London School of Economics and Political Science

Department of Economics

Cash transfers and women's labour supply: theory and evidence from Brazil

Gabriel Leite Mariante

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy. London, May 2025

Abstract

Women's participation in the labour force has been a powerful driver of economic growth, lifting millions out of poverty. Today, cash transfer programmes have come to dominate government efforts to reduce poverty. Whilst effective at providing short-term relief, these programmes may also undermine long-term poverty reduction by disincentivising labour supply.

To study the direction and magnitude of the impact of cash transfers on labour supply, as well as the underlying mechanisms, I measure the effect of an exogenous increase in Brazil's main transfer programme on the employment of men and women. I find no effect for men, whereas women increase their labour supply by 7.4% over two years. This is driven by mothers with children of pre-school age, for whom the transfer relaxes childcare constraints, enabling them to join the labour force.

Exploring regional disparities, I find that the effect is stronger in areas that are poorer, but that have better availability of local public services. I causally estimate complementarities between cash transfers and local public good provision levering discontinuities on the allocation of public funds to Brazil's 5570 municipalities. I show that this has potentially important fiscal implications for the long-term costs and benefits of cash transfer programmes.

Overall, my paper illustrates that there is no trade-off between short-term relief and long-term poverty reduction. Rather, cash transfers encourage women's labour force participation, particularly when complementary public goods, such as educational facilities, are available.

Acknowledgements

My PhD journey has only been possible with the contribution of many people, to whom I owe an enormous debt of gratitude.

First, I feel extremely lucky to have had an incredible team of advisors. I thank Oriana Bandiera for being the most inspiring researcher I have ever met; for showing me that our research matters beyond the papers we write; for opening more doors than I can count; for her unwavering enthusiasm, sense of humour, and for the countless cups of tea over meaningful discussions at LSE (or Hampstead Heath). I thank Robin Burgess for believing in my potential, for teaching me to see the bigger picture, for sharing his vast knowledge, and for his incredible institution-building effort bringing such positive externalities for the LSE community. I thank Camille Landais for taking me on board with his insightful advising, for pushing me to go beyond my comfort zone, for sharing so much of his knowledge in so many different areas, and for his early-morning open door - literally and metaphorically.

I have also benefited immensely from insightful discussions with many members of the LSE faculty. I thank Gharad Bryan, Maitreesh Ghatak, Rachael Meager, Nava Ashraf, Johannes Spinnewijn, Kate Smith, and Christiane Szerman. I am especially thankful to Xavier Jaravel and Joana Naritomi for their inspiring mentorship since even before my PhD, and to Matthias Doepke for his incredible support during the stressful job market period. Beyond the LSE, I thank Joana Silva and Henrik Kleven for taking me on board in their research agenda, and teaching me so much along the way.

I am grateful to Economics Department and STICERD staff for their incredible support during the PhD years: Sadia Ali, Marta Wasik, Mike Rose, Hitesh Patel, Nic Warner, Mini Staskus, Emma Taverner. I am especially thankful to Lubala Chibwe and Samira Mezroui, both for our chats over coffee and for their incredible role in setting up the Hub – one of the most important research initiatives to which I have had the privilege of being associated. I am thankful to all the members of the LSE community, who are responsible for making it the special place it has always been.

I thank Joana Costa and Matheus Leal at the Brazilian Ministry for Social Development, and Sergio Guimarães, Paulo Tafner and Flavio Riva at IMDS Brasil for sharing their knowledge and insights, and for offering me the opportunity to take this project beyond the London seminar rooms and into the Brazilian policy discussions.

To the social workers of the Brazilian government: thank you for your hard work, your kindness, and for doing so much to improve the lives of people who are victims of an unfair system - you are true heroes. Thank you also for hosting me in different parts of

the country, for sharing your knowledge, and for showing me first-hand how a well-designed policy can be, quite literally, life changing.

I am extremely grateful for the many friends with which the LSE has gifted me over the past years. Canishk Naik, Will Parker, Caterina Soto Vieira, Gaby Deschamps, Anton Heil, Patrick Schneider, Nilmini Herath, Michelle Rao, Roberto Maura, Chloé de Meulenaer, Sarah Winton, Amen Jalal, Veronica Salazar, Seb Ernst, Arnaud Dyevre, Jack Thiemel, Nacho Bañares – thank you for sharing this journey through all the steps of the way; thank you for the pints, the coffees, the lunches in the park, the group therapy sessions and the countless laps around the Fields. Your friendship made all the difference.

I thank my carioca friends back in Rio, who never stopped making me feel at home under the sun, and I thank my family spread across multiple cities and countries: Pati, Ciça, Tito, Rita, Dani, Ana, Aluísio, Barbara, Glorinha, Roberto. The unstranslatable *saudade* has definitely been one of the hardest part of this journey.

I thank my wonderful in-laws Nuria, José and Ana, for welcoming me *con los brazos abiertos* into their family, and for all their support over the past few years.

Nothing of what I have done would have been possible without the unconditional love and support of my parents, Mariana and Armando. *Mãe, pai*: thank you for teaching me to be curious and compassionate, and for always supporting me in pursuing my goals. To you, I owe all the opportunities I have had in life.

Finally, I am grateful beyond words to my wife Paula. Thank for your love, for your support at every single step of this journey - the good ones, and the difficult ones. Thank you for teaching me to be kind, and to take myself less seriously. Thank you for your patience, your joy, and for believing in me even when I did not. Sharing life with you has been my biggest privilege.

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it). The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party. I declare that my thesis consists of 32902 words.

Summary

Abstract i

Acknowledgements ii

Chapter 1 2

1.1 Introduction: cash transfers, labour supply and the rise of women's employment . . . 2
1.2 Summary of methodology and results . . . 4
1.3 Contribution to academic literature . . . 6
1.4 Thesis outline . . . 7

Chapter 2 9

2.1 Institutional Background . . . 9
2.2 Data sources . . . 11

Chapter 3 14

3.1 Empirical Strategy . . . 14
3.2 Main Results . . . 18
3.3 The role of gender and parenthood . . . 20
3.4 Robustness to alternative specifications . . . 22
3.5 Threats to identification and alternative explanations . . . 23

Chapter 4 27

4.1 Heterogeneity by child's age . . . 27
4.2 Survey evidence on expenditure and availability for work . . . 28
4.3 Administrative evidence on self-reported school enrolment . . . 33
4.4 Alternative Mechanisms . . . 33
4.5 Discussion . . . 34

Chapter 5 36

5.1 Overview: same programme, different contexts . . . 36
5.2 Municipality-level evidence: local supply-side complementarities . . . 37
5.3 Micro-region level evidence: exploring local covariates . . . 40

5.4 State-level evidence: individual constraints 45
5.5 Discussion 47

Chapter 6 **49**
6.1 Fiscal externalities 49

Chapter 7 **52**
7.1 Theoretical framework: Economic forces at play 52
7.2 A Model of Occupational Choice 53
7.3 Discussion 60

Chapter 8 **62**
8.1 Concluding remarks 62

Appendix A: Additional Tables **72**

Appendix B: Additional Figures **91**

Appendix C: Summary of Qualitative Work **112**

List of Tables

3.1 **Descriptive statistics by eligibility** . . . 15
3.2 **Effects on self-declared types of employment** . . . 19
3.3 **Effects on formal employment by gender and parenthood** . . . 21

4.1 **Effects on formal employment for young mothers, by age of oldest child** . . . 28
4.2 **Effect of transfer receipt on unavailability to work and expenditure on different types of goods** . . . 31
4.3 **Effects on children's school enrolment, by age group** . . . 33

5.1 **Effects on formal employment around the FPM discontinuity for different bandwidths** . . . 40
5.2 **Potential correlates of treatment effect size** . . . 42
5.3 **Quadratic relationship between state-level treatment effect and constraints to work** . . . 47

6.1 **Time for full repayment under different values of treatment effect** 51

A.1 **Proportion of beneficiaries assumed to fall under the Permanence Rule by eligibility to Basic Benefit in December/2012** . . . 73
A.2 **Effects on different types of formal employment** . . . 74
A.3 **Effects on formal employment by quartile of reform shock** . . . 75
A.4 **Average baseline monthly earnings by self-declared job type** . . . 76
A.5 **Model parameters for simulations in Chapter 7** . . . 77
A.6 **Effects on different types of employment by gender** . . . 78
A.7 **Effects on different types of employment by parenthood** . . . 79
A.8 **Effects on formal employment by region** . . . 80
A.9 **Effects on formal employment by formal market attachment between 1999 and 2013** . . . 81
A.10 **Average school enrolent by age in 2010 Census (in % of each age)** 82
A.11 **Covariates used in Propensity Score matching from survey data** . 83
A.12 **Effects on log-formal salary, conditional on being employed** . . . 84
A.13 **Components of spending categories from POF 2008-2009** . . . 85

A.14 **Most common education items purchased below poverty line, by income quintile** 86
A.15 **Effect on formal employment by municipality of employment** . . . 87
A.16 **Frequentist regressions of micro-region level treatment effect** . . . 88
A.17 **Measures of pooling and explained variance for different specifications of Bayesian Hierarchical Models** 89
A.18 **Effect of FPM discontinuity on economic outcomes** 90

List of Figures

1.1 **Reasons for work unavailability, adults age 18-45 in Brazil** 3

3.1 **First stage: share receiving the transfer around 2014 reform, by eligibility** 17
3.2 **Event-study estimates of the impact of transfer on formal employment (Equation 3.1.1)** 18
3.3 **Event study estimates on formal employment by gender** 20

4.1 **Average spending in different categories of education by income** . 32

5.1 **Population cutoffs determining FPM transfers from federal government** 38
5.2 **Local government spending per capita around FPM population cutoffs** 39
5.3 **Bayesian Meta-Regression estimates by gender** 44
5.4 **State-level estimates of treatment effect across Brazil** 45
5.5 **State-level treatment effect vs share of women unavailable to work due to care work** 46

6.1 **Benefit and employment effects over a longer time horizon** 50

7.1 **Distribution of occupational choice in a no-transfer equilibrium** . . 58
7.2 **Simulated share in wage employment as a function of transfer size** 59
7.3 **Impact of transfer on wage employment vs magnitude of upfront cost to work** 60

B.1 **Log GDP per capita across Brazilian micro-regions, equivalent countries highlighted** 92
B.2 **Share in wage employment vs transfer size, for varying magnitudes of c_{high}** 93
B.3 **Event study estimates by gender - varying first-stage** 94
B.4 **Event study estimates by gender - micro-region x quarter fixed effects** 95
B.5 **Event study estimates by gender - recursive dynamic specification** 96

B.6 **Event study estimates at the household level** 97
B.7 **Evidence on pre-reform income manipulation** 98
B.8 **Reported household income per capita of ineligible group in future years** 99
B.9 **Evolution of transfer receipt: control group vs placebo inframarginal group** 100
B.10 **Impact of the 2014 reform on size of formal market** 101
B.11 **Evolution of enrolment in education for children of different age groups** 102
B.12 **Share of women working part-time, by income decile and type of employment** 103
B.13 **First stage estimates of propensity score** 104
B.14 **Posterior estimate of state-level treatment effect for men vs share of men unavailable to work due to caring duties** 105
B.15 **Observed and law-implied FPM transfers per capita around nearest population cutoff** 106
B.16 **RD estimates of the impact of FPM "jump" on municipal spending on local public goods** 107
B.17 **Dispersion of estimates across different countries from the literature** 108
B.18 **Correlation between local education supply and spending in educaton** 109
B.19 **Correlation between local health supply and spending in health** . 110
B.20 **Correlation between share of low income households with access to basic infrastructure services and spending in infrastructure** . . 111

"Development consists of the removal of various types of unfreedoms that leave people with little choice and little opportunity of exercising their reasoned agency."

Amartya Sen - *Development as Freedom* (1999)

Chapter 1

1.1 Introduction: cash transfers, labour supply and the rise of women's employment

The rise in women's paid employment is strongly associated with economic development (Duflo, 2012; Goldin, 2014; Bandiera et al., 2022a). Paid employment does not only confer women more autonomy within the household, higher educational attainment and improved health (Qian, 2008; Heath and Mobarak, 2015; Jayachandran, 2021), but it is also associated with better outcomes for children and improved allocation of talent in the economy (Doepke and Tertilt, 2009; Hsieh et al., 2019; Ashraf et al., 2023). As such, while significant gender disparities remain (Olivetti and Petrongolo, 2016), the inclusion of women in the labour market has been a key driver of global increase in living standards and reduction in extreme poverty (Boserup, 1970; Young, 2012).

The main policy used by governments around the world to address the persistence of extreme poverty are cash transfer programmes[1], often explicitly targeted at women (Attanasio and Lechene, 2014; Heath and Jayachandran, 2018). While empirically shown to be effective at short-term poverty relief (Bastagli et al., 2016; Niehaus and Suri, 2024), cash transfers are often touted as handouts that discourage employment, and hence contribute to long-term persistence of poverty (Gerard and Gonzaga, 2021; Bergolo and Cruces, 2021; Gerard et al., 2024).

Theoretically, the effect of a cash transfer on labour supply is ambiguous (Baird et al., 2018). On one hand, the standard neoclassical economic view predicts that an unearned transfer decreases work through a pure income effect (Becker, 1965)[2]. On the other hand, a cash transfer can enable employment by funding upfront costs to work under credit constraints, such as investment, job search, or outsourcing of home production and care work (Bandiera et al., 2017; Caria et al., 2024). It can also increase productivity through better

[1] Cash transfers were first implemented in low- and middle-income countries, but have seen a recent rise in high-income settings. As of 2022, more than one billion people received some modality of cash transfer (Gentilini, 2022).

[2] In the case of a means-tested transfer, it also predicts a disincentive to increase earnings in order to maintain eligibility (Bergolo and Cruces, 2021; Bergstrom et al., 2022).

health and nutrition (Dasgupta and Ray, 1986; Banerjee et al., 2020). These mechanisms suggest that the effect on men and women can differ, depending on the relevance of each channel. Empirically, meta-analyses fail to find a consistently negative impact - ruling out a clear dominance of the income effect (Banerjee et al., 2017) - but estimates have broad support (Bastagli et al., 2016; Diaz-Pardo and Rao, 2024; Crosta et al., 2024).

This ambiguity, both theoretical and empirical, suggests important - and thus far understudied - heterogeneities in treatment effect. This study addresses the following research question: what determines the effect of cash transfers on the labour supply of men and women?

I address this question in the context of an unconditional cash transfer in Brazil, implemented as part of *Bolsa Família*[3], the world's largest cash transfer programme by number of recipients. I combine administrative data on programme recipients and the universe of formal employment contracts with survey data on household expenditures and work availability. The large sample size from administrative data sources allows me to explore heterogeneities and mechanisms driving the effect of Bolsa Família on employment.

In Brazil, the context of my study, barriers to employment are remarkably gendered. Figure 1.1 documents self-reported reasons preventing men and women from working, according to a nationally representative survey. The differences between men and women are striking: one third of women aged 18 to 45 report being unavailable to work, but only 11% of men. The gap in unavailability is entirely driven by household caring duties, which constrains 21% of women, but less than 1% of men. This pattern suggests that, if a transfer is able to relax constraints to employment, it must do so by easing care work duties for women.

Figure 1.1: Reasons for work unavailability, adults age 18-45 in Brazil

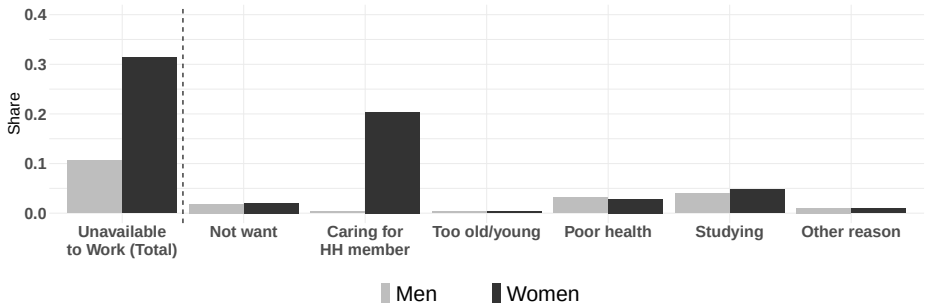

Note: This figure illustrates the share of all adults between the ages of 18 and 45 who self-declare being unavailable to work by gender, as reported in the 2014-2016 waves of PNAD-C. The first pair of bars shows the share of all adults who report being unavailable, computed as the share who respond "No" to the questions: "Are you currently working?" and "If you were offered a job today, would you be available to take it?". All respondents who report being unavailable to work are subsequently asked to choose between one of six possible reasons - the share declaring each reason is shown in the subsequent bars.

[3]Translated from Portuguese as "family stipend".

1.2 Summary of methodology and results

To identify the effect of the cash transfer on employment, I exploit plausibly exogenous variation in receipt generated be a policy reform that increased Bolsa Família coverage. I find that the transfer significantly increases women's formal employment by increasing their propensity to incur expenses related to children's education, which in turn frees up time previously dedicated to care work. The same transfer has no significant effect on men's employment. To explore regional disparities, I combine a Bayesian Hierarchical Model with a quasi-experiment generated from the discontinuity in the allocation of public funds across Brazil's municipalities. I find that the effect is stronger in areas that are poorer, but have higher public spending in local education.

In the first part of my empirical analysis, I quantify the aggregate country-wide effect of Bolsa Família's unconditional cash transfer on employment by exploiting plausibly exogenous variation in benefit receipt. This is generated by an unanticipated reform enacted by the Brazilian government in 2014, which raised the monthly income per capita threshold for eligibility[4]. I leverage administrative data from the universe of low-income individuals in Brazil observed before the policy announcement to create a treatment group, consisting of those in households with per-capita income just below the post-reform threshold, and a control group, with income just above[5]. I use these small pre-reform income differences as an instrument for actual benefit receipt in the estimation of an event study. I find that the transfer increases women's formal wage employment by around 1.13pp (7.4% of control mean) over the two year period after the reform, but the effect on men is zero.

To investigate potential mechanisms behind this result, I zoom in on the sub-population of mothers, and show that the treatment effect is driven by women with children of pre-school age[6]. Additionally, I turn to two nationally representative survey datasets (a labour force survey and a household expenditure survey), and to administrative data on beneficiaries' children. I find that benefit receipt is associated with higher spending on goods and services related to children's education, such as school material and after-school activities[7], as well as on medication. I also find that recipients are less likely to report being unavailable to work due to caring duties at home. I also find that children of pre-school age from beneficiary households are about 1.3pp (2.1%) more likely to be enrolled in school following the reform, whereas the transfer has no impact on enrolment of children in other age groups. These findings are reassuring, as the effect on mothers' employment and on children's enrolment are of similar magnitude and concentrated in the same age group. Taken together, these findings suggest that the unconditional transfer helps mothers relax

[4]The policy raised the threshold from 70 BRL (approximately 20 USD) to 77 BRL (22 USD).

[5]Crucially, income is observed four months before the policy announcement, at a moment in which 77 BRL had no policy relevance, which alleviates concerns of income manipulation for benefit eligibility.

[6]Consistent with the large cross-country variation in levels of pre-school enrolment, as opposed to very low-levels of nursery enrolment and almost universal primary and middle school enrolment.

[7]Public schools are free in Brazil, so schools fees are not an expenditure in most low-income households' budget.

childcare constraints, and free up their time for paid employment.

In the second part of my empirical analysis, I show that the aggregate country-wide effect of the transfer on formal employment masks large regional variation within Brazil, comparable to cross-country variation of estimates from the literature. I leverage data from different sub-national levels to draw insights into relevant local economic factors driving the relationship between cash transfers and labour supply. First, at the municipality level[8], I exploit a feature of the Brazilian rules for inter-governmental allocation of public funds that creates quasi-random variation in local spending in education. The magnitude of federal funds received by municipalities follows a step-wise function of estimated population, with discontinuities at arbitrary cut-offs. This feature generates substantial differences in local budget between otherwise similar municipalities - a minimum share of which is earmarked for spending in education. This allows causal estimation of the individual-level interaction between transfer receipt and the level of local public spending. Estimating the treatment effect of the transfer on employment at each side of the population discontinuities, I find large positive complementarities. The effect of Basic Benefit on formal employment in municipalities that benefit from the exogenous shock in public education spending is between 2pp and 3.2pp (more than twice the national average), whereas the effect in municipalities that receive exogenously less funding is not significantly different from zero.

Having found causal evidence of complementarities between spending in local public goods and the effect of the transfer on employment, I move to the level of the micro-region[9] to more broadly explore relevant local economic correlates through a Bayesian Meta-Regression (Meager, 2019; Angrist and Meager, 2023). I find that effects on women's employment are larger in areas that are poorer, but have higher local public spending in the provision of education and health services. These findings are a supply-side counterpart to the mechanisms of increased individual spending in education. The effect on men's employment, on the other hand, has no clear correlation with regional economic variables. Finally, I exploit survey data representative at the state level to document an inverted U-shaped relationship between treatment effect on labour supply and level of self-reported constraints to employment - consistent with a setting in which a transfer can increase employment if it has the right magnitude to release ex-ante constraints to employment, but not otherwise (i.e. if constraints are either non-binding or too high to be overcome with the benefit).

My findings suggest that the supply of local public goods can be an important complement to cash transfers by enabling the relaxation of individual constraints preventing employment. These complementarities can have fiscal implications. In the final part of my analysis, I estimate the extent to which local public spending changes the fiscal exter-

[8]The smallest unit of local government, of which there are 5570 across the country

[9]Of which there are 556 - the definition of a Brazilian micro-region is equivalent to that of a US commuting zone: it is usually considered to be a local labour market, with the vast majority of the population living and working within the same micro-region.

nalities of cash transfers by means of increased tax revenue. Following beneficiaries over a longer time horizon, I see that the difference in benefit receipt caused by the reform fades after two years, but the effects in formal employment persist for at least four years. I leverage the natural experiment on public spending to show that, under constant treatment effect on employment, a 4% increase in public spending can reduce the time for full fiscal repayment of transfer cost from a baseline of 23 years to 7 years.

To rationalise my empirical findings and lay out the mechanisms at play, I propose a theoretical framework that incorporates gender differences in barriers to work. A simple extension of standard models of occupational choice in equilibrium (Banerjee and Newman, 1993; Lloyd-Ellis and Bernhardt, 2000; Feng and Ren, 2023), in which agents are heterogeneous in a binary upfront cost to wage employment and receive a means-tested transfer, yields two related predictions. First, the relationship between share in salaried employment and the size of the transfer is non-monotonic: transfers of intermediate size increase labour supply, whereas very small or very large transfers decrease labour supply. Second, the effect of a transfer on wage employment follows an inverted U-shaped relationship with the magnitude of fixed costs faced by women, in line with empirical evidence at the state level.

Overall, my paper shows that an unconditional transfer can increase employment through the relaxation of constraints keeping beneficiary out of the labour force. I show that this effect is particularly relevant for women with children of pre-school age, and that policies that ease constraints on groups that face higher barriers to employment can complement a cash transfer in promoting employment.

1.3 Contribution to academic literature

This study contributes to several strands of the existing literature. First, I contribute to the large body of evidence documenting the role of care work as a constraints to female labour force participation (Bertrand et al., 2010; Kleven et al., 2019; Müller and Wrohlich, 2020; Berlinski et al., 2024; Kleven et al., 2024). I add to previous studies by showing evidence of the role of cash transfers in relaxing this constraint and increasing female labour supply.

Given the unearned nature of cash transfers, I also add to the large body of empirical evidence on the relationship between unearned income and labour supply in general (Gruber, 2000; Cesarini et al., 2017; Jones and Marinescu, 2022; Golosov et al., 2024), and cash transfers and labour supply in low- and middle-income countries specifically (Bastagli et al., 2016; Banerjee et al., 2017; Diaz-Pardo and Rao, 2024; Crosta et al., 2024). This literature has heterogeneous findings, to which I contribute in two meaningful ways. First, I show empirical evidence of a particular mechanism driving the effects: care work constraints keeping women out of the labour force (building on Ardington et al., 2009 and

Ervin et al., 2017). Second, I exploit cross-regional variation within Brazil to systematically study contextual drivers and policy complementarities that interact with treatment effect.

Within the cash transfers literature, several papers have documented the relationship between transfer programmes and the local economic context. One set of studies have found substantial general equilibrium effects of cash transfer programmes in the broader economy, and in particular on aggregate employment (Angelucci and De Giorgi, 2009; Egger et al., 2022; Gerard et al., 2024). A second set of studies has documented a relevant role of *ex ante* contextual factors mediating the effect of transfers on other outcomes, in particular those directly targeted by the design of cash transfers, such as consumption and children's education (Kondylis and Loeser, 2021; Molina and Vidiella-Martin, 2023; Fassarella et al., 2024). My study brings these two strands together, by investigating the role of *ex ante* contextual factors in shaping labour supply responses to cash transfers.

A number of past studies have specifically focussed on Brazil's Bolsa Família. Firpo et al. (2014) and Bergstrom et al. (2022) examine income reporting incentives created by the programme rules - relevant features of which are discussed here in chapters 2 and 3. More closely related, De Brauw et al. (2014) finds a relationship between programme receipt and increased female empowerment, consistent with the more pronounced effects on women found here. Also consistent with results found here, Cruz and Ziegelhofer (2014) and Chioda et al. (2016) find positive impacts of the programme on children's human capital and negative impact on teenager criminal behaviour. Most studies focus on the conditional component of Bolsa Família.

Of particular relevance to this study is Gerard et al. (2024), who find significant effects of Bolsa Família on many dimensions of local economies in Brazil beyond beneficiaries themselves, including increasing formal employment. I differ from their study in three meaningful ways: first, my main question can be thought of as the reverse - I study the role of *ex ante* characteristics of the local economy in shaping the impact of cash transfers; second, my identification strategy combining an income discontinuity with a policy change identifies the absolute dynamic effect of the transfer on the new beneficiaries, rather than the contemporaneous effects around the income distribution; third, I focus on gender differences of impacts on employment and specific constraint-relaxation mechanisms.

1.4 Thesis outline

The remainder of this thesis is organised as follows. Chapter 2 describes the relevant institutional settings of Bolsa Família and of public funds allocation in Brazil, and presents my data sources. Chapter 3 presents aggregate impacts at the national level. Chapter 4 shows evidence of mechanisms driving the results. Chapter 5 documents regional correlates of geographical variation on treatment effects, as well as causal evidence of complementarity

between cash transfers and local public spending. Chapter 6 discusses fiscal externalities of the programme. Chapter 7 presents a theoretical framework with a model to rationalise the empirical analysis. Chapter 8 concludes.

Chapter 2

In this chapter, I describe institutional features that are relevant to my findings, as well as the datasets used in my empirical analysis. Section 2.1 provides an overview of the Bolsa Família cash transfer progamme. Section 2.1 provides an overview of the mechanism governing the allocation of public funds across Brazilian municipalities (further explored in Chapter 5). Section 2.2 describes the different data sources used in my analysis.

2.1 Institutional Background

Bolsa Família

Brazil's Bolsa Família is the world's largest cash transfer programme (Gerard et al., 2024). It was created in 2003, when the incoming presidential administration combined a number of pre-existing local social protection programmes into a large national scheme[1]. By 2014, the year of interest in this study, Bolsa Família covered more than 14 million households across the country (Viana et al., 2018). The programme has several components, all of which consist of a direct cash transfer via debit card to beneficiary households, for which eligibility depends on a combination of household characteristics such as income, family composition or children's school attendance.

In this study, I focus on Bolsa Família's largest component, the **Basic Benefit**: an unconditional cash transfer, for which eligibility is based solely on household per-capita income. It is targeted at households considered by the government to be living in extreme poverty. Until May/2014, the definition of extreme poverty was having a monthly per-capita income of up to 70 BRL (approximately 20 USD/capita). Each eligible household was entitled to a monthly transfer worth 70 BRL, equivalent to, on average, 80% of eligible household's pre-benefit total income.[2]

[1] A unified benefit combining all pre-existing transfer schemes was first proposed in Camargo and Ferreira (2001), inspired by the experience of other countries in Latin America.

[2] Other components are targeted at households who are considered to be living in poverty which, until May/2014, meant having a per-capita income of up to 140 BRL. The main other benefits are: a variable benefit of BRL 32 (BRL 38) per child for households with children up to the age of 16 (between ages

Household income is calculated either at the time of registration or at the time of information update on Cadastro Unico, the federal government's national database of low-income households (see section 2.2), used as an overarching register to assess eligibility for all national social programmes. In order to be eligible for any benefit managed through Cadastro Unico (including Bolsa Família), households must update their information at least once every two years. Information is collected from an interview with the household head, either during a home visit by a worker from the local social assistance centre[3], or during a visit to the local social assistance offices by the household head. The household head is asked to report the monthly income from a pre-specified list of non-labour sources for each household member[4], as well as labour income both over the past year and over the past month for each one. Total household income is calculated by summing all non-labour income with the minimum between labour income over the past month and 1/12th of labour income over the past year for each household member[5]. This value is then divided by the number of residents in the household to obtain the calculated household per-capita income. The final calculated household per-capita income is used to determine eligibility to all means-tested benefits.

Even if calculated from a mathematical formula instead of declared directly, the eligibility system has the obvious drawback of, in theory, incentivising households to under-report their true income or to choose to not take paid work elsewhere in order to maintain eligibility. To address this, in 2010, the federal government amended the programme regulation with the introduction of the so-called **Permanence Rule** (MDS, 2010). Under the Permanence Rule, a household that, while receiving any Bolsa Família benefit, voluntarily discloses an increase in income that takes it beyond the eligibility threshold (and hence, based on income alone, would become ineligible) is allowed to remain a beneficiary for two years. However, if a household is caught misreporting their income or failing to report an increase in income, the benefit is cut immediately. This way, the government does not

16 and 18) conditional on school attendance, for up to 5 children under 16 and two children between 16 and 18; and a transfer of BRL 32 for households with pregnant or lactating women. All benefits can be accumulated.

[3]Workers from local social assistance centres have a regular schedule of home visits to beneficiaries or potential beneficiaries of all social programmes. These visits, which are often unannounced, have the purpose of verifying information declared by the household, flagging potential issues that might be exacerbating vulnerability of specific household members (e.g. domestic violence), and ensuring that households are registered for all benefits and services for which they are eligible.

[4]These sources are: pension, disability benefits, donations from non-household members, unemployment insurance, alimony and other non-specified sources of income which include, but are not limited to: interest from savings, rent, and support for family members of prison inmates.

[5]The formula for household i is:

$$\text{Income}_i = \sum_j \left(\sum_k \text{Non-Labour Income}_{ijk} + \min \Big(\text{Labour Income Last Month}_{ij}, \frac{1}{12} \text{Labour Income Last Year}_{ij} \Big) \right)$$

where j indexes each household member, and k indexes each different source of non-labour income for household member j

discourage Bolsa Família beneficiaries from increasing their pre-benefit income beyond the eligibility threshold.

The 2014 Reform On April 30th 2014, the federal government announced a reform that increased the income per-capita threshold below which a household is considered to be in extreme poverty from BRL 70/capita to BRL 77/capita (Casa Civil, 2004)[6]. This reform, largely unanticipated, established that the new thresholds governing eligibility would be valid from June/2014. The law increased the programme's total national coverage by 400 thousand households (equivalent to around 3%), and the total programme spending by around BRL 90 million per month.

Fundo de Participação dos Municipios

To investigate potential complementarities between local public spending and cash transfer in Chapter 5, I exploit a key institutional feature in the design of Brazil's intergovernmental transfers. Brazil is subdivided into 27 states[7], containing a total of 5570 municipalities. Whereas the bulk of tax revenue is raised federally, the municipal authorities are responsible for most local public goods and services, such as education, local infra-structure and certain aspects of health care (Brollo et al., 2013). To address the gap between tax revenues and spending requirements across different government's spheres, the Brazilian constitution provides a number of intergovernmental transfer mechanisms, the largest of which being Fundo de Participação dos Municipios (FPM), responsible for around one third of all municipal budget (Corbi et al., 2019).

According to the rules of FPM, each municipality that is not a state capital is entitled to a yearly transfer that depends on (i) the state where it is located, and (ii) its population estimate as calculated by the National Statistics Office. Crucially, within each state, the amount of funds received by a given municipality as a step-wise function of its population estimates, creating plausibly exogenous variation that allows for causal identification. Municipal population is either calculated via the national census (if it is a census year) or centrally estimated by the National Statistics Office based on data from the latest census and trends in population growth. Chapter 5 discusses more details on fund allocation, and on how this feature generates plausibly exogenous variation in local public spending.

2.2 Data sources

In my analysis, I combine data from a variety of sources including administrative records, household surveys, the national census and local public finances, which I describe next.

[6]The same policy change also increased the monthly transfer amount from BRL 70 to BRL 77
[7]Including the Federal District, which only contains the capital Brasilia.

Administrative data

Cadastro Unico My main source of administrative data are the administrative records from Cadastro Unico, the national registry of low-income households maintained by the Brazilian federal government. All households with a monthly per-capita income lower than 50% of the national minimum wage are eligible to be registered[8]. Cadastro Unico contains information on household and individual characteristics such as key demographics, employment, dwelling conditions, school enrolment and income. The household head (usually an adult woman) is interviewed on behalf of all members, and asked about employment and income from several sources of each household member separately, which are then combined on a final value of monthly household income per-capita to be used as an eligibility criterion for Bolsa Família (see section 2.1 for details on the calculation). Each individual registered in Cadastro Unico is assigned a unique identifier, the NIS (Social Identification Number). I use yearly extractions from Cadastro Unico taken every December between 2012 and 2017. The extract from December 2013 is my main sampling frame, as it the last one available before the 2014 reform (see section 3.1 for details on sample selection).

Bolsa Família Payment Sheets I combine data from Cadastro Unico with the Bolsa Família's official monthly payment sheets. For each month between October 2003[9] and December 2019, the payment sheet contains the amount paid to each household of every Bolsa Família benefit. The household is identified via the NIS of the household head, which allows the payment sheets to be merged with Cadastro Unico.

RAIS The third source of administrative data is the *Relação Anual de Informacoes Sociais* (RAIS), a national database from the Ministry of Labour that contains the universe of all formal employment contracts in the country between the years 1985 and 2019. For each contract, RAIS contains information on salary, occupational code, industry code and date of hiring and termination. The contract is identified by the pair employer-employee, where each has part has unique identifier. This allows for the creation of a monthly panel, both at the worker level and at the firm level. The worker's identifier is the same as the NIS number used in Cadastro Unico, which allows for a merge between the two datasets. I also use RAIS to identify aggregate characteristics of local labour markets, such as share of working-age population in formal employment, and average formal wage.

[8] In 2014, the national minimum wage was BRL 724 (approximately USD 200) per month.
[9] The official start of the programme.

Household Surveys

POF I use the 2008/2009 wave of the national Family Budget Survey[10] (POF), a household expenditure survey conducted every 7 to 8 years representative at the national level. POF contains information regarding household spending on a large number of finely defined products and services, as well as information on household demographics and income from various sources – including Bolsa Família. For the analysis in Chapter 4, I create broader categories of spending through aggregating specific items, as described in Appendix Table A.13.

PNAD-C I use data from the Continuous National Household Sample Survey[11] (PNAD-C) for the 2014 to 2016 period. PNAD-C is a quarterly survey that contains information on employment, income by source (including Bolsa Família), job search status, and self-reported reasons for unavailability to work. PNAD-C is, in origin, a rotating panel where each household is interviewed for five quarters. However, income from Bolsa Família is only included in a supplementary questionnaire that is administered in the first quarter of every year. As such, I treat the data as an annual repeated cross-section. The survey is representative at the state level.

Regional economic variables

Municipal Public Finances Data on local government spending comes from the publicly available database *Finanças do Brasil* (FINBRA), maintained by the National Treasury. Spending is reported on a number of disaggregated categories (e.g.: education, health, energy), and is available annually between 2013 and 2019 at the municipal level. Data on all sources of inter-government transfers, including the yearly transfers through FPM, is publicly available from the records of the National Treasury.

Local Public Goods To construct an index for local public goods provision in Chapter 5, I leverage data from a number of different sources at the municipal level. Data on educational outcomes comes from the school-level Education Census, provided by INEP, the official data and research institute of the Ministry for Education. Data on health indicators is collected by the Institute for Health Policy Studies[12] (IEPS), a think tank that collects and curates a database of municipal health indicators extracted from the Ministry of Health. Data on local infrastructure availability to low-income households, such as access to electricity, piped water and sewage network, comes from the Cadastro Unico section on dwelling conditions. I take regional averages to create local indicators.

[10] *Pesquisa de Orçamentos Famíliares* in the original Portuguese name.
[11] *Pesquisa Nacional por Amostra de Domicilios Continua* in the original Portuguese name.
[12] Instituto de Estudos para Politicas de Saude

Chapter 3

In this chapter, I describe the aggregate country-wide results of estimating the effect of Bolsa Família's Basic Benefit on labour supply. I outline the identification strategy, sample selection, and estimating equations. I show the main results for different types of employment, and explore the role gender and parenthood. I then discuss potential threats to my identification and interpretation of results.

3.1 Empirical Strategy

Sample Selection

The main sample for my analysis comes from the December/2013 snapshot of Cadastro Unico, extracted six months before the 2014 reform took effect, and four months before it was announced. I consider all individuals between the ages of 18 and 60 who, as of December/2013, lived in households with a monthly per-capita income between 75 BRL and 80 BRL, i.e. within a narrow range of the post-reform threshold of 77 BRL.

I exclude the following individuals: (i) those enrolled in full-time education, (ii) those whose household per-capita income as per Cadastro's computed variable does not equal the value obtained from my own computation using the Ministry's formula and each individual's income from various sources[1], (iii) those who have not updated their Cadastro records in more than two years, and (iv) those who receive Basic Benefit in the quarter before the reform, likely due to falling under the Permanence Rule[2].

[1] This indicates entry error, and corresponds to fewer than 0.5% of the individuals. I allow for a difference in absolute value of up to BRL 2 to accommodate rounding errors.

[2] Based on income alone, all individuals from my sample are ineligible for the benefit pre-reform, as the threshold as of December/2013 was 70 BRL. Regardless, a number of households report receiving Basic Benefit in the quarter before the reform (i.e. Q1/2014). This is primarily due to the Permanence Rule, which allows households who truthfully report an increase in income to continue receiving Basic Benefit for two years. I exclude those from the sample, so my analysis focusses on new beneficiaries post-reform. Although I cannot directly observe information on the Permanence Rule, I can check the plausibility of it being the main driver of households with income above 70 BRL receiving Basic Benefit pre-reform by looking at their income in December/2012 and checking whether families considered to be under the Permanence rule are indeed more likely to have their income below 70 BRL then (ideally, I would observe

Identification

The key institutional aspect for my identification strategy is that, in December/2013 when my sample is observed, the threshold for Basic Benefit eligibility was 70 BRL and the 2014 reform had neither taken place nor been announced. At the time, the future post-reform threshold of 77 BRL was meaningless. As such, an unanticipated increase in the eligibility threshold creates two comparable groups around the new value: a **treatment** group, with households who become eligible for the benefit by the reform due to having an income below 77 BRL; and a **control** group, with households who narrowly miss out by having an income just above 77 BRL. From my sample, I consider a window of 5 BRL around the new threshold: the treatment group consists of all individuals whose household income observed in December/2013 is between 75 BRL and 77 BRL (inclusive), and the control group consists of all individuals whose household income observed in December/2013 is between 77 BRL (exclusive) and 80 BRL.

For identification to be valid, whether an individual is part of a household that falls just above or just below the cut-off must be as good as randomly assigned. Whereas this is impossible to verify empirically, Table 3.1 shows that these two groups are similar in a number of observables. This similarity, together with the fact that households have no reason to strategically bunch on either side of the future 77 BRL threshold in December/2013, lends credibility to the identification.

Table 3.1: Descriptive statistics by eligibility

	Ineligible	Eligible
Employment		
Share formal in 2014-Q1 (admin data)	17.6%	17.3%
Av. months in formal employment throughout 2013 (if > 0)	7.75	7.65
Av. annual formal income throughout 2013 (if > 0)	6514 BRL	6399 BRL
Share informal in Dec 2013 (self-reported)	31.4%	28.3%
Demographics		
Share women	63.4%	68.5%
Share in Southeast	36.3%	35.1%
Share in Northeast	34.3%	34.1%
Av. age	34.9	33.5

Note: This table shows descriptive statistics from the main estimation sample. It includes all individuals between the ages of 18 and 60 who, in December 2013, lived in households with monthly income per-capita between 75 BRL and 80 BRL. It excludes those in full-time education, and those who received the benefit while ineligible due to Permanence Rule.

their income up to two years before, but data is only available from December/2012 onwards). Table A.1 in the Appendix confirms that households considered to fall under Permanence Rule (and excluded from the sample) are much more likely to have an income below 70 BRL in December/2012, which lends credibility to this assumption. This is also in agreement with evidence from interviews with workers from social assistance centres, who mention the Permanence Rule as the main reason why families with income above the threshold receive the benefit (see Appendix C for a brief summary of the interviews).

Estimating Equations

I leverage the quasi-random variation generated by the 2014 reform to estimate two main specifications. The first is a two-way fixed effect event study according to Equation 3.1.1, used for RAIS outcomes on formal employment, where I have a quarterly panel.

$$Y_{it} = \theta_i + \alpha_t + \sum_{s \neq -1} \delta_s \cdot I(t = s) \cdot \hat{\text{Benefit}}_i + \epsilon_{it} \tag{3.1.1}$$

Y_{it} is my outcome for interest for individual i at quarter t. θ_i and α_t are individual and time fixed-effects. Benefit_i is a dummy variable that is equal to 1 if individual i received Basic Benefit for the six months immediately following the reform (i.e. quarters 3 and 4 of 2014). My coefficients of interest are the set of $\{\delta_s\}$ capturing the interaction between treatment and time fixed-effects.

For quarterly RAIS data, Y_{it} is a dummy variable that indicates formal employment by individual i in quarter t. I consider an individual as formally employed in a quarter t if they have a formal employment contract in RAIS in all three months of quarter t. The reason for this somewhat stringent definition is so that it captures regular contractual formal employment, as opposed to temporary short-term work.

The second specification is a difference-in-differences equation, in which the periods pre- and post- reform are collapsed on a $Post_t$ dummy. This specification is used to obtain a point estimate for the overall treatment effect in quarterly RAIS data, as well as on annual self-reported Cadastro data, as per Equation 3.1.2:

$$Y_{it} = \theta_i + \alpha_t + \beta \cdot \hat{\text{Benefit}}_i \cdot Post_t + \epsilon_{it} \tag{3.1.2}$$

Where α_t is now either a year fixed-effect for yearly self-reported Cadastro Único data, or a quarter fixed-effect for quarterly RAIS data. $Post_t$ is equal to one if period t is after the second quarter of 2014, and zero otherwise – in the quarterly RAIS data, this variable turns to one from the third quarter of 2014; in the yearly Cadastro Único data, it is equal to zero for years 2012 and 2013, and equal to one from 2014 onwards[3]. The coefficient of interest capturing the treatment effect of Basic Benefit is β.

First Stage Because of imperfect take-up, which is potentially endogenous, I instrument the treatment variable Benefit_i with pre-reform eligibility, as defined by household per-

[3]Since the yearly Cadastro Único snapshots are always extracted in the month of December, the observations for 2014 are all coded as post-reform, since the reform happened in June.

capita income in December 2013 within the narrow window around the new cut-off of 77 BRL. The first stage is estimated according to Equation 3.1.3.

$$\text{Benefit}_i = \gamma_0 + \gamma_1 \cdot Z_i + \eta_i \tag{3.1.3}$$

Where Z_i is equal to one if per-capita income as observed in the December/2013 Cadastro Único is between 75 BRL and 77 BRL (inclusive), and is equal to zero if it is between 77 BRL (exclusive) and 80 BRL.

Figure 3.1 shows the evolution of benefit receipt for both eligible and ineligible groups over time around the policy change. A sudden increase in receipt is visible for the eligible group from the second quarter of 2014, when the new eligibility criterion takes effect, whereas the ineligible sees only a gradual increase of beneficiaries over time.

Figure 3.1: First stage: share receiving the transfer around 2014 reform, by eligibility

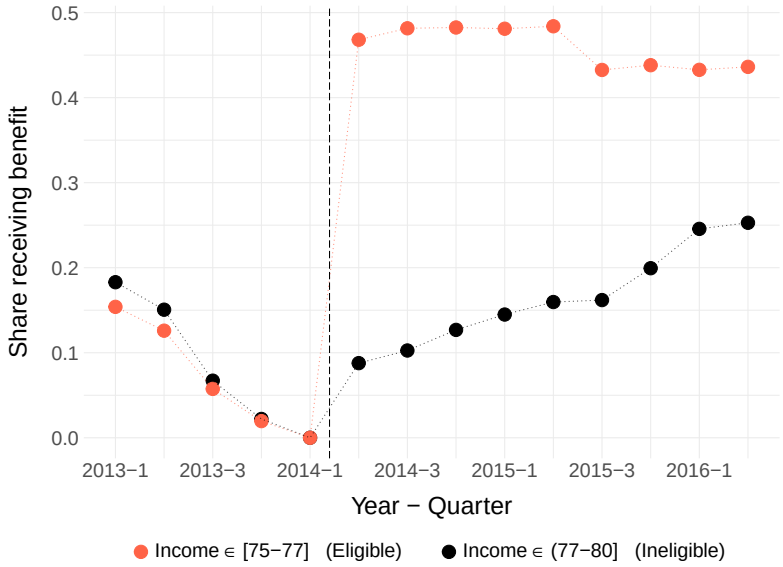

Note: This figure illustrates the first-stage of the empirical specification. It shows the evolution of the share of individuals receiving Basic Benefit at each quarter around the 2014 reform. by eligibility. Eligibility is defined by household per-capita income in December 2013. The red (black) dots are individuals living in households whose income in December 2013 were below (above) the future threshold of BRL 77. Benefit receipt in each quarter is defined as receiving the benefit in all three months of the quarter, with the exception of 2014-Q2, the quarter during which the reform was announced (April) and implemented (June), in which receipt is defined as receiving the benefit at any point in the quarter.

3.2 Main Results

Formal employment from administrative records

Figure 3.2 illustrates the event-study estimated from Equation 3.1.1. The policy reform that expanded Basic Benefit increased regular formal employment of the recipients over the subsequent two years by, on average, 1.14pp, which corresponds to a 6% increase on a baseline rate of formal employment of approximately 17%. The overall trend suggests a small immediate effect in the quarter where the reform starts to take place, followed by a gradual increase over the two subsequent years.

Figure 3.2: Event-study estimates of the impact of transfer on formal employment (Equation 3.1.1)

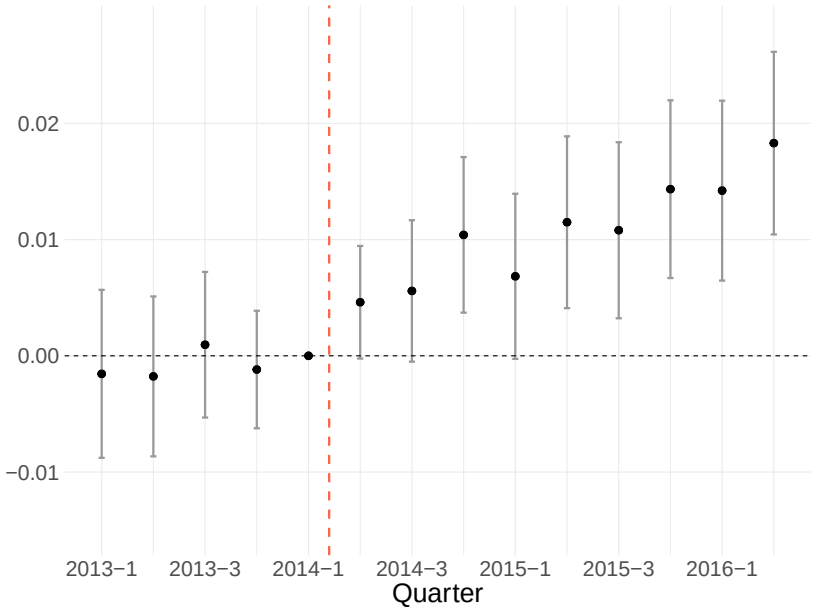

Note: This figure illustrates the set of event-study estimates $\{\hat{\delta}_s\}$ obtained from Estimating equation 3.1.1 on individual formal employment for the full sample. The grey bars represent the 95% confidence interval at each quarter. Standard errors are clustered at the individual level.

Appendix Table A.2 reports the results of estimating Equation 3.1.2 on total formal employment (a collapsed version of the event study from Figure 3.2), as well as the effect on different categories of formal jobs. Columns 2 and 3 show that the effect comes entirely from jobs in the private sector, whereas the effect on the public sector is a precisely estimated zero. Columns 4 to 7 show the effect on different occupation groups according to their ISCO 1-digit classification. The benefit has a positive effect on occupations as service workers (0.7pp) and as office clerks and technicians (0.2pp).

Self-declared employment

Administrative data from RAIS has the drawback of only covering formal employment. As in most middle-income countries, Brazil's informal sector is very large, accounting for about two thirds of businesses and more than one third of employees (Ulyssea, 2018). This is particularly relevant for this study, as the informal sector is even more prevalent in low-income populations, such as individuals in my sample (Bergolo and Cruces, 2021).

Table 3.2: Effects on self-declared types of employment

	(1) Any employment	(2) Self-employed	(3) Formal employee	(4) Informal employee
Post · Benefit	0.061***	0.042***	0.013***	0.0044***
	(0.0047)	(0.0042)	(0.0030)	(0.0017)
N	301024	301024	301024	301024
Mean Ineligible	0.44	0.29	0.076	0.04

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 where Y_{it} is an indicator for self-declared employment of each type, conditional on remaining in Cadastro Único until 2017. Column (1) shows the effect on any employment, column (2) on self-employment, column (3) on formal wage employment, and column (4) on informal wage employment.

By definition, no official administrative records exist on informal work in Brazil. To investigate the transfer effect on employment beyond the formal sector, I turn to the self-declared employment status variable from the Cadastro Único records. I consider a yearly panel from Cadastro Único snapshots every December between 2012 and 2017[4]. Table 3.2 shows the result of estimating Equation 3.1.2 on a dummy variable equal to one if individual i self-declares to be in any employment, in self-employment, in formal employment or in informal employment at year t[5]. Column 1 shows that the impact on employment of any kind is around 6pp, which suggests that only considering the formal sector underestimates of true effect of Basic Benefit on overall employment. This effect is reassuringly similar to the median impact of 5.9pp of monthly transfers found in the meta-analysis by Crosta et al. (2024). Column 2 shows that the largest impact comes from an increase in self-employment of around 4.2pp. Column 3 shows that the increase in formal employment is around 1.3pp, which is remarkably close to the point estimate obtained by using RAIS administrative records. This match between the impact as measured on self-reported formal employment and on administrative records is reassuring with respect to the

[4] An unavoidable caveat in this analysis is that all regressions on Cadastro outcomes can only be estimated conditional on individuals being present in future waves of Cadastro. An individual may leave Cadastro if their household crosses the income threshold to be registered, because of death, or because of irregularities captured by social assistance workers. From my core sample defined in 2013, 97% are present in at least one future wave of Cadastro and 66% of individuals remain in Cadastro in 2017.

[5] The other categories possible employment types covered in Cadastro Único are: apprentice, entrepreneur, and casual agricultural worker, all of which have fewer than 1% of individuals and have been omitted for brevity.

validity of using Cadastro's self-reported employment outcomes in this analysis. Column 4 shows that the impact on the probability of being an informal employee increases slightly, by 0.4pp.

3.3 The role of gender and parenthood

Figure 3.3 illustrates the event studies on formal employment for men and women separately. Panel 3.3a shows a clear positive effect for women, similar to overall pattern observed for the aggregate population. Panel 3.3b shows no significant trend for men, with relatively noisy estimates that fluctuate around zero.

Figure 3.3: Event study estimates on formal employment by gender

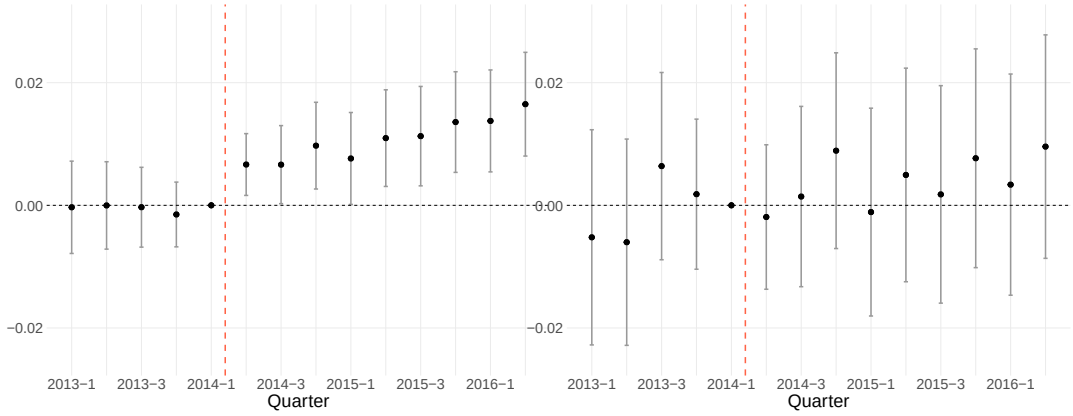

(a) Women **(b)** Men

Note: This figure illustrates the set of event-study estimates $\{\hat{\delta}_s\}$ obtained from estimating equation 3.1.1 on individual formal employment, splitting the sample by gender. The grey bar represents the 95% confidence interval at each quarter. Standard-errors are clustered at the individual level.

Table 3.3 breaks down the aggregate effects on formal employment by gender, collapsing the event-studies from Figure 3.3 into a difference-in-difference specification as per Equation 3.1.2, and by parenthood.

Column 1 shows the aggregate effects for everyone, columns 2 and 3 show the effect by gender. The coefficient for men is very small, and not significantly different from zero. Consistent with the event-studies, the effect on the overall sample is entirely effect driven by beneficiary women, who increase formal employment by 1.13pp, equivalent to 7.4% of the mean in the ineligible group. Columns 4 and 5 show the result of a similar exercise in which the sample of women between mothers and non-mothers. Here, I define a woman as being a mother if she has a son or daughter age 18 or younger living in the same household. The effect is driven by the subsample of mothers, whereas the coefficient on the subsample of non-mothers is equal to that of men's, and not significantly different from zero.

Table 3.3: Effects on formal employment by gender and parenthood

	(1)	(2)	(3)	(4)	(5)
		Full sample, by gender		*Women, by motherhood*	
	Full sample	Men	Women	Mothers	Non-mothers
Post · Benefit	0.011***	0.004	0.011***	0.012***	0.004
	(0.003)	(0.006)	(0.004)	(0.004)	(0.006)
N	310755	99583	211172	134434	76738
Mean Ineligible	0.178	0.236	0.148	0.159	0.130

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 for formal employment for different subsamples. Column (1) shows the estimate for the full sample, columns (2) and (3) show estimates for men and women separately, and columns (4) and (5) show estimates for mothers and non-mothers separately. Mothers are defined as having at least one child under the age of 18 living in the same household. Standard errors are clustered at the individual level.

Tables A.6 and A.7 in the Appendix show the same results for self-declared measures of self-employment and and total employment from yearly Cadastro snapshots. The conclusions are similar to those on formal employment: the effect is larger for women than for men (although men also see a positive effect in self-employment) and, among women, it's largest for mothers.

Overall, these results suggest that the effect of Basic Benefit on employment is strongest for the sub-population of mothers. This finding suggests that potential mechanisms through which the transfer is increasing labour supply might be more at play for mothers than for other groups of individuals.

Other dimensions of heterogeneity

Region Appendix Table A.8 explore the effect for men and women across different regions of Brazil, according to the municipality of residence at baseline. The effects are stronger in the South/Southeast (Panel A) and in the Norhteast (Panel B) - which are, respectively, the richest and the poorest regions of the country, and concentrate over 80% of my sample. In those regions, the effect mirrors the overall country-wide results: a positive significant effect on the overall sample, entirely driven by women. Panel C show that the estimates for the less populated North and Centre-West of the country are noisier, with no significant effect overall, and some evidence of a positive effect on men (although care must be taken in interpreting this, as men in the North/Centre-West represent less than 7% of my entire sample). Chapter 5 investigates geographical variation across the country in details.

Formal market attachment Appendix Table A.9 explores the role of *ex-ante* formal labour market attachment, as measured by presence in the formal market at any point in the 15 years before the reform. In levels, the effects are very similar both for attached and for non-attached individuals, and mirror the overall sample. However, the average baseline employment rate of the ineligible group for those with prior attachment to the formal market is about 34%, while the average for those with no prior attachment is approximately 2%. Such difference suggests that, as a percentage of the ineligible group, the effect is much larger for those with no prior attachment to the formal market (an increase of 37.5%) than for those with prior attachment (an increase of 3.8%).

Participation vs earnings

Finally, Appendix Table A.12 shows the treatment on formal salary conditional on employment. While the point estimates are negative for all subsamples, none of them is significantly different from zero. This result suggests that, conditional on joining the formal market, the benefit has no impact on earnings - the labour market response is purely on the extensive margin.

3.4 Robustness to alternative specifications

For robustness, I consider four alternatives to the main individual-level event-study specification defined in Equation 3.1.1.

First, I allow treatment status to vary over time. For the periods after the reform, treatment status Benefit_i becomes Benefit_{is} and is equal to 1 if individual i receives the transfer in quarter s. The motivation for this specification is the fact that, as illustrated in Figure 3.1, a share of ineligible group gradually becomes eligible over time, so my main specification is likely to be a lower bound of the full reform effect by considering that those who do not receive the benefit immediately after the reform are always untreated. The drawback of this alternative specification is that it does not take into account the persistence of the outcome variable: if an individual is formally employed in a quarter, they are likely to be formally employed in the following quarter regardless of benefit receipt. This alternative specification is hence likely to be an upper bound on total treatment as it attributes all the effect to contemporaneous benefit receipt. The resulting estimates are shown in Appendix Figure B.3 - the coefficients get larger with time, by virtue of the first stage becoming smaller as more ineligible individuals get treated.

Second, I consider my main specification with the addition of micro-region $\times$ quarter fixed effects. If results were different, one could worry that my main estimates are partly driven by different trends in overall local formal employment, potentially due to general

equilibrium effects of the reform itself. The resulting estimates are shown in Appendix Figure B.4, and look almost identical to my main specification.

Third, I consider the dynamic nature of treatment as illustrated in Figure 3.1, and estimate a set of treatment effect parameters $\{\delta_s\}$ which, instead of representing the effect in calendar quarter s, represent the effect of treatment on employment s periods after receipt, allowing for potential endogeneity of take-up at each period. I follow Giupponi and Landais (2023) and assume that the reduced form ITT estimate of treatment assignment on employment at period s[6] can be written as $ITT_s = \sum_s \delta_s \cdot \frac{d\text{Benefit}_{i,s}}{dT_s}$, so the set of coefficients $\{\delta_s\}$ can be recovered recursively from the ITT's and first stage regressions. The resulting estimates are shown in Appendix Figure B.5 for men and women. Quarterly estimates for women are significant at the 10% level in quarter 0 - i.e. immediately upon transfer receipt - and fluctuate noisily around zero for the remaining quarters, with suggests that the bulk of the effect on employment comes contemporaneously. For men, the effects are all close to zero with large standard errors, indicating no clear pattern, consistent with my main estimates.

Fourth, I consider estimates at the household level, where the outcome variable Y_{it} from Equation 3.1.1 is equal to one if any adult from household i is in formal employment at quarter t and zero otherwise. Appendix Figure B.6 show that the resulting estimates are very similar to the individual-level estimates - the difference-in-differences point estimate (standard error) is 0.011 (0.003). The similarity in results rules out concerns that the individual-level results might be coming from within-household changes in the allocation of labour supply.

3.5 Threats to identification and alternative explanations

The validity of the previous results and their interpretation rely on two key assumptions. First, the 2014 reform can be considered a quasi-experiment and treatment eligibility within the considered pre-reform income band is as good as random. Second, positive results on employment are driven by increased employment on the treated group, and not by decreased employment on the control group. Next, I show empirical evidence supporting the validity of these assumptions.

Income manipulation

The results of the empirical analysis would be invalid if, in the December 2013 snapshot of Cadastro, households were able to endogenously choose to be just to left of the 77 BRL income threshold instead of just to the right, in order to be eligible for the Basic

[6] Obtained by estimating $Y_{it} = \theta_i + \alpha_t + \sum_{s \neq -1} ITT_s \cdot I(t = s) \cdot Z_i + \eta_{it}$.

Benefit. This is unlikely given that, in December 2013, the precise value of 77 BRL had no significance since the threshold for Basic Benefit eligibility was 70 BRL, and the reform was only announced in April 2014.

The conceptual implausibility of pre-reform income manipulation at BRL 77 is further reinforced by two pieces of empirical evidence show in Appendix Figure B.7. Panel B.7a shows the density of household income per-capita around BRL 77. The distribution is clearly non-smooth, with spikes at round numbers and no specific excess mass at BRL 77. In line with the approach by Gerard et al. (2024), Panel B.7b shows the t-statistics of a McCrary test for smoothness of density at every integer between 60 BRL and 90 BRL. The McCrary test rejects smoothness at every integer tested, and the t-statistic at 77 BRL is not larger than other surrounding values. The conclusion from this analysis is that, while the distribution is not smooth, there is no evidence of manipulation around BRL 77[7].

Behavioural responses

The results presented in section 3.1 come from difference-in-differences specifications, with positive coefficients interpreted as increased employment for Basic Benefit recipients. However, an alternative interpretation of the same positive coefficients could be that non-recipients reduce labour supply, possibly in order to increase the probability of benefit receipt post-reform. Appendix Figures B.8 and B.9 show empirical evidence suggesting that this is not the case, and the results are indeed driven by increased labour supply from Basic Benefit recipients.

Figure B.8 shows the reported per-capita household income of the ineligible group (i.e. those in households with income above 77 BRL and below 80 BRL in December 2013) in the two subsequent yearly snapshots of Cadastro Único: December 2014 (panel B.8a) and December 2015 (panel B.8b). If the ineligible group was responding to the reform by reducing labour supply in order to gain eligibility, we would see an excess mass between BRL 70 and BRL 77 in their post-reform reported income: the income range that was previously ineligible and became eligible with the reform. There is no such excess mass - but rather strong persistence over time, as the majority of beneficiaries continue reporting the same income between 77 BRL and 80 BRL.

Appendix Figure B.9 illustrates the evolution of benefit receipt for the ineligible group, compared to a placebo group of "never eligible", who are not at the margin of the cut-off introduced by the reform. The trajectory in benefit receipt of the never-eligible group

[7]Other forms of income misreporting have been documented in the literature or mentioned by social assistance workers in interviews conducted for this project (see Appendix C), such as bunching at current thresholds (Bergstrom et al., 2022), or omission of income sources or household members altogether. These are not a problem for my identification strategy insofar as there are no plausible reasons to expect differential behaviour in these dimensions in either side of an arbitrary cut-off that was not the threshold at time of observation.

resembles the evolution of the control group shown in Figure 3.1. This similarity suggests that the increase in the ineligible group's benefit receipt post-reform is in line with regular churn in and out of benefits experienced by other infra-marginal ineligible individuals, rather than a behavioural response adjusting income to assure eligibility post-reform.

General equilibrium effects

Another concern to the interpretation of my findings is the existence of different general equilibrium effects on local labour markets caused by the reform itself, as described in Egger et al. (2022) and Gerard et al. (2024). An alternative interpretation of the effects would be that increased employment is driven not by the relaxation of supply side constraints, but by an increase in labour demand. Two pieces of empirical evidence allow me to rule out this explanation.

Firstly, if general equilibrium effects of the reform itself were the driving force behind increase beneficiary employment, I would find stronger employment effects in areas that saw a larger increase in benefit injection as a result of the reform itself. Appendix Table A.3 shows the result of separately estimating Equation 3.1.2 for quartiles of intensity of the reform shock according the beneficiaries' micro-region of residence[8], as defined by the difference in total regional-level Basic Benefit beneficiaries post- and pre-reform. The estimates do not follow a clear upwards pattern with benefit intensity, and are not statistically different from each other across all quartiles.

Secondly, general equilibrium effects driving the results would be possible if an increase in local labour demand as a result of the reform affects beneficiaries and non-beneficiaries differently. This could be the case if baseline employment between the two groups was different, e.g. the beneficiaries had more slack capacity to take up employment than non-beneficiaries. Table 3.1 shows that this is not the case - both groups have very similar rates of formal employment at baseline.

Job Displacement

Related to general equilibrium effects, a final key concern is the issue of job displacement: is it the case that employment increase by beneficiaries comes at the expense of non-beneficiaries which would otherwise be employed? If so, the net effect of an increase in Bolsa Família would be null - beneficiaries would simply be displacing non-beneficiaries.

To rule out job displacement as the mechanism driving the results, I follow Gerard et al. (2024) and look at the overall effects of the 2014 Bolsa Família reform on aggregate employment. I split municipalities according to the size of the 2014 reform impact in terms

[8]The are 556 micro-regions in Brazil, generally considered to be the spatial unit corresponding to a local labour market - see Chapter 5 for more details.

of total amount of Bolsa Família received. Appendix Figure B.10 shows the evolution in total number of formal jobs on municipalities who were above the median versus those who were below the median impact of the reform. If the effects were fully driven by displacement, the reform would have no net effect on total formal jobs. Instead, I find a 0.7% increase in total formal jobs in the municipalities above the median, when compared to those below the median.

Chapter 4

In this chapter, I explore potential mechanisms driving the positive impact of a large unconditional transfer on labour supply, with a focus on factors explaining the difference in effects by gender and parenthood. First, I zoom in on the subsample of mothers to investigate heterogeneity in treatment effect by child's age. Second, I use two nationally representative household surveys to establish correlational evidence on the impact of Basic Benefit on household spending and individual work availability. Third, I estimate the effect of Basic Benefit on self-reported school enrolment of beneficiaries' children.

4.1 Heterogeneity by child's age

The results from the previous chapter show that the effect of Basic Benefit on individual formal employment is mainly driven by mothers, indicating that this is the group of people who are most likely to be able to use the transfer to relax constraints to employment. This finding is consistent with descriptive evidence presented in Figure 1.1, suggests that caring for children might be an especially relevant constraining factor. To explore this hypothesis in more detail, I investigate heterogeneity in treatment effect according to children's age. I estimate the following equation:

$$Y_{it} = \theta_i + \alpha_t + \sum_{k \in G} \gamma_k \cdot \hat{\text{Benefit}}_i \cdot \text{Post}_t \cdot \mathbf{1}(\text{AgeOldestChild}_i = k) + \epsilon_{it} \tag{4.1.1}$$

Where AgeOldestChild_i is the age of mother i's oldest child at baseline grouped according to ages broadly corresponding to different schooling levels: $G = \{0\text{-}2, 3\text{-}6, 7\text{-}10, 11\text{-}18\}$. The set coefficient estimates $\{\hat{\gamma}_k\}$ represent the local average treatment effect on mothers whose oldest child is in age group k. Table 4.1 shows the resulting estimates for all ages between 0 and 10.

Table 4.1: Effects on formal employment for young mothers, by age of oldest child

	Age group of oldest child			
	(1)	(2)	(3)	(4)
	0 - 2	3 - 6	7 - 10	11 - 18
Post · Benefit	-0.03	0.012*	-0.003	0.001
	(0.008)	(0.006)	(0.008)	(0.005)
Schooling stage	Nursery	Pre-School	Primary School	Middle/High School
N	21967	35644	18638	47838
Mean Ineligible	0.11	0.16	0.16	0.16

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of transfer on formal emplyoment for a similar specification to Equation 4.1.1, except the age of the oldest child is grouped into three categories according to the stage of education that they are supposed to be enrolled according to Brazilian law. The sample contains all mothers younger than 35 with children between the ages of 0 and 10. Standard errors are clustered at the individual level.

The resulting estimates are zero for the ages 0 to 2, become positive between the ages of three and six, and then decrease to zero for the oldest ages. The effect is mainly driven by children of pre-school age.

This finding can be put into context within the broader pattern of children's enrolment in education shown in Appendix Table A.10, where I compute the average share of children of each age between 0 and 20 enrolled in each schooling stage, according to the National Census of 2010. The vast majority of children under the age of three are not enrolled in any education. On the other hand, from age of seven onwards, enrolment is almost universal, which is consistent with the age of mandatory schooling as determined by the Brazilian Law at the time[1]. This suggests that the age of children at which mothers see the largest effect is also the one in which there is more variation in school enrolment. Between ages 3 and 6 (i.e. around pre-school age), school enrolment is not universal, but possible and to some extent widespread. This setting is where the transfer's capacity to relax constraints preventing employment due to care work duties are most relevant.

4.2 Survey evidence on expenditure and availability for work

By and large, descriptive evidence suggests that the main binding constraint for women's employment is care work (see Figure 1.1). To uncover the relationship between the benefit and care work constraints, I combine data from PNAD-C with data from another nationally representative household survey: the 2008/2009 wave of POF (*Pesquisa de Orcamen-*

[1]Since then, the law has been changed to make schooling mandatory from the age of four. However, supply-side bottlenecks still remain, meaning that pre-school enrolment is still not universal across the country.

tos Famíliares), a national budget survey where households are asked to report weekly expenditure on a large number of finely detailed categories. Crucially, both PNAD-C and POF contain a module on different sources of income, including Bolsa Família[2].

Both surveys are repeated cross-sectional datasets[3], which precludes causal identification. To overcome this limitation, I employ a Propensity Score Matching (PSM) approach. In the first stage, I estimate a probit regression of a dummy indicating household receipt of Basic Benefit on a number of observable household characteristics[4] on the right-hand side (Martins and Monteiro, 2016). I use the estimated propensity score to estimate the effect of Basic Benefit on POF and PNAD-C variables of interest. Given the lack of exogenous benefit variation, results in this section should be interpreted as correlational, rather than causal. Figures B.13a and B.13b illustrate the PSM first-stage, where the density of estimated propensity score is illustrated for both recipients and non-recipients of Basic Benefit - the density for recipients lies primarily to the right of the one for non-recipients, which suggests that the set of observables used to estimate the propensity score does predict receipt with good accuracy.

I estimate three empirical specifications: (1) an OLS regression where I regress each outcome of interest on the dummy for Basic Benefit receipt controlling for the estimated propensity score; (2) a "nearest neighbour" matching strategy, where each recipient household is matched to its closets non-recipient household on the covariate space[5], and the resulting dataset is then used to estimate an OLS regression of each outcome of interest on the dummy for Basic Benefit receipt, where each observation is weighted by its inverse propensity score; and (3) a simple OLS without the propensity score, where I regress each outcome of interest on the dummy for Basic Benefit receipt and the set of observables used to estimate the propensity score.

Table 4.2 summarises the findings. Panel A shows the effects of receiving Basic Benefit on self-reported unavailability for work in PNAD-C. The estimate from column 1 suggests that individuals living in households receiving the transfer are between 3.6pp and 4.2pp less likely to declare not being available for work[6]. The effect comes mainly from a reduction

[2]Neither questionnaire breaks down income from the different benefits that make up Bolsa Família. In order to specifically flag households that receive the Basic Benefit, I look at the total monetary amount received by each household, and compare with the values of different kinds of benefits at the time of the surveys. For example, between June/2014 and June/2016 the amount paid by the basic benefit was BRL 77 and the amount paid by the variable benefit for children under 16 was BRL 32. A household receiving the Basic benefit plus two variable benefits will hence receive $77 + (2 \times 32) =$ BRL 141.

[3]PNAD-C is actually a rotating quarterly panel where each household is interviewed for 5 quarters. However, income from Bolsa Família is only part of a supplementary questionnaire that is asked once a year, so in practice it is akin to a yearly repeated cross-section.

[4]The set of variables used in the prediction is detailed in Appendix Table A.11. It is slightly different for each of the surveys, based on constraints on data availability, but it includes income, demographic composition of household members and some basic information about the dwelling conditions.

[5]Each non-recipient household can be matched to more than one recipient household, as sampling is done with replacement.

[6]All individuals who are working are coded as zero for all dummies, as being in employment necessarily implies being available to work. As such, the estimates should be interpreted in absolute terms, not conditional on employment

in the probability of being unavailable to work due to caring duties (column 4). Panels B and C show the effects of receiving Basic Benefit on a dummy for non-zero expenditure in each category of goods. The estimates from Panel B show that households who receive the transfer are around 11p more likely to incur in education-related expenses (column 1), with most of the effect coming from expenses with school material (column 6), after-school activities (column 5) and pre-school (column 3). Column 7 shows that beneficiary households are also around 3pp more likely to have non-zero pharmacy expenditure. The estimates from Panel C show that households who receive Basic Benefit are not more likely to purchase non-durable goods (columns 1, 2 and 3), the exception being bikes and motorcycles (column 4). The negative coefficient on spending on public transport (column 6) suggests a substitution towards private transport. Finally, households are between 1.7 and 3.6pp marginally more likely to have expenses with alcohol and tobacco.

Table 4.2: Effect of transfer receipt on unavailability to work and expenditure on different types of goods

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	A. Labour Force Survey (PNAD-C): Reasons for unavailability to work						
	Any reason	Studying	Too old/young	Caring duties	Health issues	Does not want	Other
Simple Prop Score	-0.039***	-0.001	-0.001**	-0.016***	-0.010***	-0.006***	-0.003**
	(0.006)	(0.002)	(0.001)	(0.005)	(0.000)	(0.002)	(0.001)
Nearest Neighbour	-0.042***	-0.001	-0.001**	-0.020***	-0.008***	-0.006***	-0.003**
	(0.006)	(0.001)	(0.001)	(0.006)	(0.000)	(0.001)	(0.001)
OLS with controls	-0.036***	-0.001	-0.005	-0.018***	-0.008***	-0.005**	-0.002*
	(0.006)	(0.001)	(0.001)	(0.005)	(0.001)	(0.002)	(0.001)
	B. Expenditure Survey (POF): Health and Education						
	Education (any)	Childcare	Pre-school	School Fees	After-school	School Material	Pharmacy
Simple Prop Score	0.111***	0.004**	0.014***	0.002	0.012	0.118***	0.029***
	(0.010)	(0.002)	(0.004)	(0.004)	(0.004)	(0.010)	(0.009)
Nearest Neighbour	0.115***	0.005**	0.012***	0.002	0.014	0.114***	0.028***
	(0.010)	(0.002)	(0.004)	(0.004)	(0.004)	(0.009)	(0.009)
OLS with controls	0.110***	0.004*	0.014***	-0.000	0.014	0.116***	0.026***
	(0.009)	(0.002)	(0.004)	(0.004)	(0.004)	(0.010)	(0.009)
	C. Expenditure Survey (POF): Durables and Other Goods						
	Durables (any)	Appliances	Electronics	Bike/Moto	Car	Public Transp	Alcohol/Tobacco
Simple Prop Score	0.047***	0.004***	-0.014***	0.040***	-0.008	-0.036***	0.017*
	(0.009)	(0.004)	(0.002)	(0.008)	(0.003)	(0.010)	(0.009)
Nearest Neighbour	0.048***	0.005***	-0.008***	0.036***	-0.004***	-0.026***	0.036***
	(0.010)	(0.004)	(0.002)	(0.008)	(0.003)	(0.009)	(0.009)
OLS with controls	0.046***	0.004***	-0.007***	0.040***	-0.004	-0.026***	0.017**
	(0.009)	(0.004)	(0.002)	(0.007)	(0.003)	(0.010)	(0.009)

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows regression estimates for the impact of receiving Basic Benefit on survey responses of interest. In panel A, the outcome is a dummy for a positive answer to the question "Would you be available to work if offered a job today?" (column 1), followed by dummies for each reason for work unavailability (columns 2-7). In panels B and C, the outcome are dummies for non-zero expenditure in different categories of goods and services. Each panel includes estimates of three specifications. "Simple Prop Score" shows the coefficient of a regression on the outcome on a dummy variable for receiving Basic Benefit, controlling for predicted probability of receipt as a function of observables described in Appendix Table A.11. "Nearest Neighbour" shows the coefficient of the same regression, but now the sample is balanced so each recipient household is matched to the nearest non-recipient household in the hyper-space of considered observables. "OLS with covariates" reports the estimate of a simple OLS regression of each outcome on a dummy for Basic Benefit receipt and a set of observables. The set of observables considered in all specifications is listed in Appendix Table A.11.

Figure 4.1 complements the budget analysis by illustrating the relevance of upfront fixed costs in low-income households' education-related budget. Conditional on having non-zero spending in education, households at the very bottom of the income distribution dedicate the largest share of their spending to school material. Appendix Table A.14 shows the 5 most common goods that make up the education budget of households below the poverty line. For households at the bottom two quintiles, all five items are non-fee costs required for school attendance education, such as stationary, textbooks, uniforms, and backpacks. Some of these items are purchased by a relatively small share of households, but represent a significant share of income - e.g.: about 9% of households in the bottom quintile report purchasing a school uniform, but those that do spend on average 12.8% of household income in uniforms.

This illustrates the relevant of upfront cost to schooling, which is also the category of education-related spending that sees the highest impact of Basic Benefit. As income increases, the monetary amount spent in school material remains roughly constant, and other spending categories such as school fees and after-school activities become more relevant. This piece of descriptive evidence suggests that, while fees and after-school activities may or may not be part of low-income households' budget, other upfront costs of schooling are a relevant expenditure.

Figure 4.1: Average spending in different categories of education by income

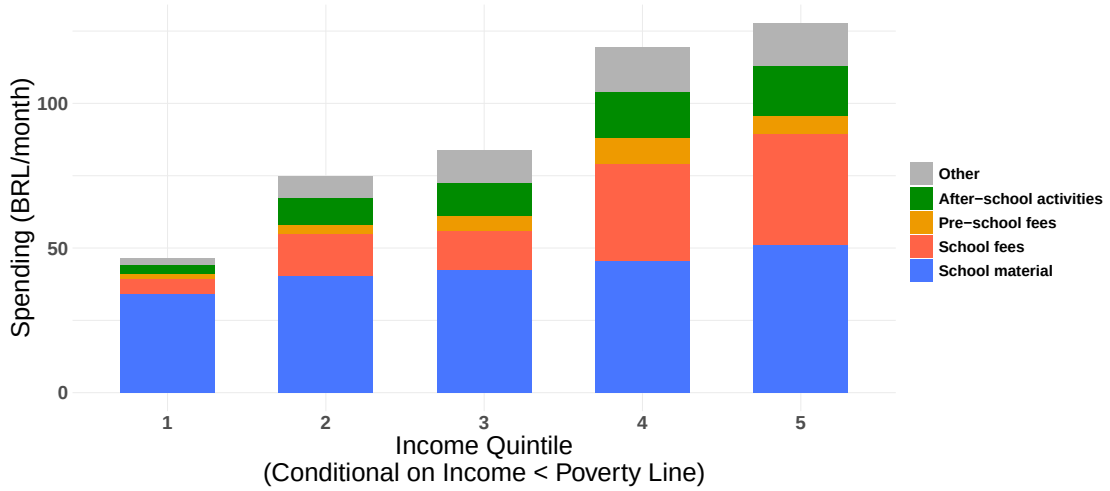

Note: This figure illustrates the breakdown of education-related spending in POF 2008-09 for different quintiles of household income. The sample includes households with per-capita income of less than half of a minimum wage (the threshold to be registered in Cadastro Único) and that have non-zero spending in education. Spending items are grouped into categories as per Appendix Table A.13.

4.3 Administrative evidence on self-reported school enrolment

Having established a correlation between receiving the Basic Benefit and increased expenditure on education and availability for work, I now turn to data on the universe of all children (up to age of 18) who live in the same household as my main sample of beneficiaries described in Section 3.1. I estimate Equation 3.1.2 with Y_{it} being a dummy for child i's self-reported school enrolment in year t in Cadastro Único.

Table 4.3: Effects on children's school enrolment, by age group

	Child's age group			
	(1)	(2)	(3)	(4)
	0 - 2	3 - 6	7 - 10	11 - 18
Post · Benefit	0.000	0.013**	0.005	0.001
	(0.006)	(0.006)	(0.004)	(0.004)
Schooling stage	Nursery	Pre-School	Primary School	Middle/High School
N	54094	75212	69510	124818
Mean Ineligible	0.22	0.62	0.95	0.88

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of transfer on self-declared school enrolment for children of different age groups, according to the approximate stage of education that they are usually enrolled. The sample contains all mothers younger than all children from households of the main sample between ages of 0 and 18. Standard errors are clustered at the individual level.

Table 4.3 reports the results by age group. Appendix Figure B.11 illustrates the raw evolution of school enrolment for eligible and ineligible groups for each group, normalised at pre-reform enrolment levels, as a check that pre-trends are not a major concern for the diff-in-diff estimation.

The result shows a positive significant effect for children between the ages of three and six, and no significant effect for the other age groups. This result is consistent with the transfer's effect on the employment of mothers by age group of oldest child: the same age group is driving both the effect on mothers' employment and the effect on children's school enrolment.

4.4 Alternative Mechanisms

Health A potentially relevant mechanism linking the transfer with increased labour supply could be via better health, which increases productivity and availability to work

(Dasgupta and Ray, 1986). Table 4.2 shows that, indeed, the unavailability to work due to health reasons is also significantly lower for women who receive the benefit compared with non-beneficiaries (panel A, column 5) - albeit to a less extent than the effect on unavailability due to caring duties. Relatedly, Panel B, column 7 shows a significant positive association with expenditure on pharmacy goods. This could be an indication of both increase spending in medications for beneficiaries themselves and on increasing spending in medication for beneficiaries' children, which is consistent with the increased enrolment in education mechanism - separating those two channels is not possible in the data[7].

Travel costs A potentially relevant mechanism behind my results could be the use of the benefit to fund commuting and travel costs, both for the job itself and for job search (Caria et al., 2024). While not possible to fully rule out this mechanism, two pieces of empirical evidence points towards its small relevance. First, Table 4.2, Panel C, column 6, shows that beneficiaries are substantially less likely to use the benefit's money to pay for public transportation – instead, there is an increase in the use of bicycles and motorcycles[8]. Second, Appendix Table A.15 shows separate estimates for Equation 3.1.2 on a binary indicator for formal employment in the same municipality (Panel A), as well as for a different municipality to where beneficiaries live[9]. Estimates in levels are relatively similar for both outcomes in levels (0.7pp vs 0.4pp), and are very similar in magnitude when considered as a percentage of the ineligible groups' mean (6.3% vs 6.6%). This similarity in results suggests that funding travel costs for employment in a different municipality is not the main driver of results[10].

4.5 Discussion

The evidence presented here sheds light on mechanisms through which the transfer has a positive impact on employment. My findings suggest that mothers of young children use the transfer to overcome liquidity constraints related to the children's education and free up time for employment outside the household. This is in line with a number of previous

[7] Qualitative evidence from focus groups and individual interviews points towards the stronger relevance of the second channel, since a large number of interviewed mothers report having children with some health condition that hampers their education, and not being able to afford appropriate medication.

[8] While it is possible that this increase expenditure in the purchase of a vehicle is associated with job search and commuting costs, it is more plausibly an investment in a durable asset, potentially useful for certain occupations with high prevalence in the informal market - e.g. delivery drivers.

[9] Municipality of residence is fixed at baseline, so a job in a different municipality could be both a result of inter-municipal commuting and of migration.

[10] It is, of course, possible that those who receive the benefit are able to travel further for job search or commute within the same municipality, but exact addresses for beneficiaries' home and work are not available in my data.

studies documenting that women are disproportionately responsible for childcare when compared to men, and as such are more likely out of paid employment in the presence of children (Ardington et al., 2009; Kleven et al., 2024). This mechanism is in also consistent with evidence from focus group discussions and interviews with beneficiaries, social assistance workers, and school principals, conducted as part of this project and briefly summarised in Appendix C.

The estimate on school enrolment for children of pre-school age is of approximately the same magnitude as the treatment effect on formal employment for mothers of pre-school age children. Taken at face value, this evidence suggests a strong relationship between increasing school enrolment of youg children, and allowing mothers to join the formal market.

Moreover, the estimate on school enrolment is about 18% of the estimate on women's self-reported employment of any kind (Appendix Table A.6), suggesting that the transfer might help overcome other constraints beyond childcare, that might be binding for self-employment. Relatedly, Appendix Figure B.12 shows the share of self-employed and formally employed women who work part-time, by income decile. Whereas the vast majority of formal jobs are full-time (i.e. 40h per week or more) across the entire income distribution, self-employed women are much more likely to work part-time, especially in the lowest income settings: in the bottom decile, this share is around 70%. Appendix Table A.4 reports average monthly earnings by self-reported job type among individuals from my main sample: self-employed individuals earn around one third of formal employees. These striking differences on hours worked and earnings suggest that: (i) time constraints that need to be overcome by women entering the labour force are much higher in formal employment than in self-employment, and (ii) self-employment is, in general, a much more precarious occupation than formal wage work. Given these striking differences in economic relevance, as well as data limitations, the remaining of the analysis focusses exclusively on formal employment.

Taken together, (i) the survey evidence on spending and availability to work, (ii) the positive effect on children's self-reported school enrolment, and (iii) the stronger effect on employment for women with pre-school age children, are all consistent with previously documented patterns, and suggests that a large unconditional cash transfer can help mothers into labour by outsourcing home production and freeing up time for work.

Chapter 5

In this chapter, I explore the extent to which the average effect estimated in Chapter 3 masks geographical heterogeneities across the different parts of Brazil, and how these heterogeneities can be connected to the mechanisms discussed in Chapter 4.

5.1 Overview: same programme, different contexts

The results from the previous chapters show a positive impact of the transfer on the employment of women with children. So far, the estimated impacts represent average effects across the entire country. However, previous literature on the impacts of cash transfers has found large dispersion of point estimates across different countries, programmes and time periods (Bastagli et al., 2016; Banerjee et al., 2017; Crosta et al., 2024; Diaz-Pardo and Rao, 2024).

The mechanisms briefly discussed in Chapter 1 and further backed by the theoretical framework outlined in Chapter 7 suggests that, under a constant transfer amount, the direction and magnitude of the treatment effect of a transfer depend on the magnitude of the upfront cost to employment faced by individuals. The heterogeneous results by gender and parenthood presented in Section 3.3, as well as the evidence of beneficiaries' increased propensity to incur in education expenses and increase children's school enrolment found in Chapter 4 suggest that the effects might be influenced by local characteristics enabling constraint relaxation, in particular those related education supply.

In this chapter, I take advantage of the Brazilian setting, in which a transfer of equal magnitude is conducted across vastly different contexts[1]. Brazil is divided into 27 states, 557 micro-regions, and 5570 municipalities, varying greatly in level of economic development (Appendix Figure B.1). I exploit of a combination of data availability, economic and

[1]Micro-regions in Brazil are remarkably diverse in their broad levels of economic development: the poorest micro-regions have the GDP per capita of Chad, whereas the richest has the GDP per capita of France (see Appendix Figure B.1 for a distribution).

institutional features of each of these three levels of geographic aggregation to investigate potential drivers of the relationship between transfer and labour supply.

I start by exploiting plausibly exogenous variation in public funds at the municipality level to causally identify complementarity between local spending in education and treatment effect. I then move to the micro-region level - the best approximation to a local labour market - to explore regional covariates of treatment effect more broadly. Finally, I move to the state level[2] in order to document the relationship between treatment-effect and individual-level constraints.

5.2 Municipality-level evidence: local supply-side complementarities

Municipalities are Brazil's smallest level of local government. Despite having limited ability to raise taxes, municipal governments are responsible for the bulk of local public goods provision, such as education, health and local physical infrastructure. To adjust for this fiscal disparity, the Brazilian Constitution has established several channels of intergovernmental transfers. The largest of these is FPM (see Section 2.1), a yearly mandated transfer of funds from the federal to municipal governments. The FPM transfers are typically very large, accounting for an average of one third of total municipal budget (Corbi et al., 2019).

The allocation of funds via the FPM follows a pre-determined rule, that have been unchanged since the scheme's inception in the Brazilian constitution of 1988. Specifically, the amount of funds received by municipality i in state k in year t is given by the following formula:

$$\text{FPM}_{i,t}^{k} = \underbrace{\text{TotalFPM}_{t}^{k}}_{\text{Budget for state } k \text{ in year } t} \cdot \frac{\overbrace{\lambda_i}^{\text{Share to munic. } i}}{\sum_{j \in k} \lambda_j} \tag{5.2.1}$$

Where TotalFPM_{t}^{k} is a the total budget pre-allocated to state k in year t, calculated according to each state's GDP per capita[3], and λ_i is a coefficient that represents the within-state share of FPM that goes to municipality i. Crucially, the municipal coefficients λ_i follow a pre-determined schedule that is a step-wise function of municipal i's

[2] Most survey data sources, including POF and PNAD-C are representative at the state level, but not at any finer.

[3] The total share destined to each state is pre-published by the federal government, and is inversely proportional to the state's estimated GDP per capita in year $t-1$.

population in year t[4]. The function is partly illustrated in Figure 5.1, and consists of sequence of discrete "jumps" at arbitrary pre-specified population cut-offs[5].

Figure 5.1: Population cutoffs determining FPM transfers from federal government

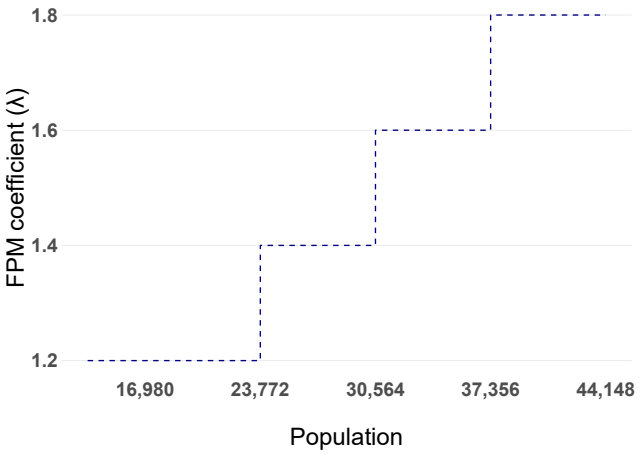

Note: This figure illustrates a part of the schedule governing the allocations of the FPM transfers from the federal government to municipalities. I shows the evolution of the coefficient λ_i as a function of municipal population for a range of population. Within each state, the value of λ_i determines the share of the total state FPM budget that is assigned to municipality i, as defined in Equation 5.2.1. Only part of the function is shown here, to illustrate the stepwise nature of the funds allocation mechanism. The full schedule is composed of 18 jumps, the first being at a population 10188, and the last at a population of 156216.

Appendix Figure B.15 shows the law-implied and the observed transfers (in BRL/capita) for municipalities within a range of 1500 inhabitants of the nearest cut-off[6]. The figure illustrates the plausibly random mechanisms governing FPM funds: municipalities that belong to the same state but happen to have its population estimated as being one inhabitant above than the nearest cut-off will receive, on average, around BRL 130 per capita when compared to a municipality with a population estimate for that year falling one inhabitant below the same threshold.

Impact on local spending

The Brazilian Constitution mandates all municipalities to spend at least 25% of total municipal budget education[7] (Brasil, 1988). Data on municipal public finances contains

[4]Every year, the Brazilian National Statistics Office (IBGE), an organism centrally controlled by the federal government, produces an estimate of each municipality's population, which are used for FPM allocation. The exception are census years, in which the actual population count from the national census is used.

[5]The stepwise schedule was introduced in 1985, and has remained the same ever since.

[6]This is equivalent to the residuals of a regression of the FPM transfers across all municipalities on state, year and nearest cutoff fixed-effects.

[7]The Constitution also mandates a 15% share of municipal budget to be spent in local health services.

information of municipal spending disaggregated across various categories. Figure 5.2 illustrates the discontinuity in total municipal spending, as well as in municipal spending on education, around the FPM cut-offs. The discontinuity in public funds is translated almost one-to-one in a discontinuity in total public spending per capita: municipalities just to the right of the nearest cut-off spend about 130 BRL/capita more than those just to left. In line with the Constitutional requirements, around a quarter of this increase is directed to spending on education.

Appendix Figure B.16 presents the effect of the FPM discontinuity in other categories of public spending related to the provision of local goods and services. For robustness, I estimate standard RD models both linear and quadratic, as well as robust RD with optimal bandwidth selection following Calonico et al. (2015). Municipalities to the right of the nearest FPM transfer schedule cut-off spend significantly more on health, education and urbanism than those to the left of the nearest cut-off, whereas spending in sanitation, housing, energy and transportation is not significantly different. As a check on the validity of the Regression Discontinuity approach, Appendix Table A.18 shows similar RD estimates for other economic outcomes. Measures of total formal employment, total female formal employment, total formal employment only considering low-paying occupations[8], and municipal GDP are all smooth around the pooled FPM discontinuity.

Figure 5.2: Local government spending per capita around FPM population cutoffs

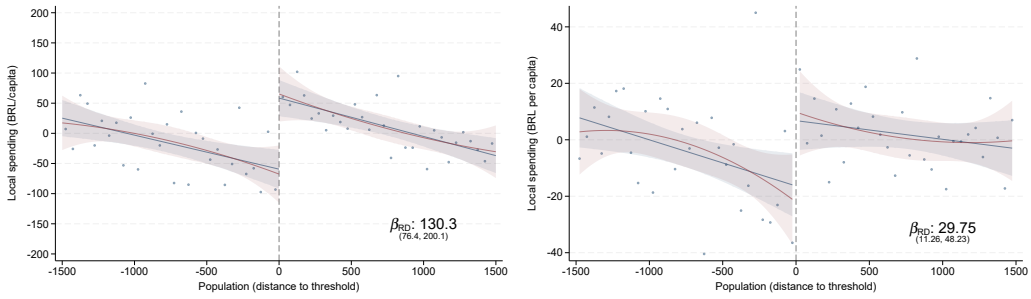

(a) Total Spending **(b)** Spending on Education

Note: This figure illustrates average municipal spending across Brazilian municipalities within 1500 inhabitants of the closest FPM cut-off for the year 2013. Dots represent the average spending of all municipalities within a bin of 50 inhabitants. Panel (A) shows total spending, and panel (B) shows spending on education. The lines on each side of the cut-off show linear and quadratic fits on each side of the cut-offs, and the shaded represent are 95% confidence intervals. The reported coefficients in each graph are obtained by a linear Regression Discontinuity design with robust bandwidth selection, following Calonico et al. (2015).

[8]Defined as an occupation that falls in bottom quintile of the distribution of average occupation-specific salaries across Brazil.

Complementarity between cash transfers and public spending

Given the plausibly exogenous variation on benefit receipt generated by the 2014 Bolsa Família reform, and the plausibly exogenous variation on local educational spending generated by the FPM transfer schedule, I can causally identify the complementarity between these two policies by estimating equation 3.1.2 separately on each side of the closest population discontinuity.

Table 5.1 shows the resulting estimates, using an increasingly large range of population bandwidths around the cutoff. The resulting estimates show a clear pattern: the point estimates for municipalities just to the left of the cut-off are small and not significantly different from zero, whereas for municipalities just to the right of the cut-off are positive, significant, and around 3pp in magnitude - i.e. around three times larger than the overall average effect of 1.1pp. This result is causal evidence of the complementarity between the cash transfer and local public spending at the municipal level.

Table 5.1: Effects on formal employment around the FPM discontinuity for different bandwidths

	(1) *Bandwidth: 500*		(2) *Bandwidth: 1000*		(3) *Bandwidth: 1500*	
	Left	Right	Left	Right	Left	Right
Post · Benefit	0.003	0.032**	-0.003	0.028***	0.004	0.020**
	(0.017)	(0.014)	(0.012)	(0.009)	(0.009)	(0.008)
Mean Ineligible	0.171	0.143	0.163	0.141	0.165	0.144
N	9604	13100	19546	29459	31335	40088

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 where Y_{it} is an indicator for individual i's formal employment in quarter t for the sample of individuals who, in December 2013, resided in municipalities close to an FPM cutoff. Closeness to the cutoff are defined by progressive large bandwidths of 500, 1000, and 1500 inhabitants according to the population estimates used to define the 2013 FPM allocation.

5.3 Micro-region level evidence: exploring local covariates

Having established the causal interaction between spending in local public goods and treatment effect at the municipality level, I turn to a larger level of spatial aggregation to study more broadly the role of potential correlates of treatment magnitude. I separately estimate equation 3.1.2 in each of Brazil's 557 micro-regions[9]. A micro-region is roughly

[9]For my analysis, I keep the 503 micro-regions for which the first-stage (Equation 3.1.3) is significant at the 95% level, dropping the remaining 54.

the equivalent of a US Commuting Zone, and has been used in the literature as the usual definition of a local market, as the vast majority of people live and work within the same micro-region (Dix-Carneiro and Kovak, 2017).

To explore relevant correlates of treatment effect while dealing with the noise contributing to the large dispersion in point-estimates, I employ a Bayesian Hierarchical model. I assume that, for each micro-region k, the estimated treatment effect $\hat{\tau}_k$ is a realised draw from a normal distribution centred around the true treatment regional effect τ_k with standard deviation $\hat{se}_k$, where $\hat{se}_k$ is estimated by the standard error of region k's point-estimate.

$$\hat{\tau}_k \sim N(\tau_k, \hat{se}_k) \tag{5.3.1}$$

The true regional estimates $\{\tau_k\}$ are in turn draws from a parent normal distribution with a hypermean τ and a hyper-standard deviation σ in a hierarchical structure. The parent distribution represents the distribution of true average treatment effect across the country, from which each true regional effect τ_k is a realisation.

$$\tau_k \sim N(\tau, \sigma) \tag{5.3.2}$$

This is akin to the classic Rubin (1981) model, and this technique has been applied in the Development Economics literature to aggregate evidence from different studies estimating the same parameter in different contexts (Meager, 2019; Bandiera et al., 2016, 2022b). Effectively, I treat each micro-region as a different study estimating the same parameter. This setting has a key advantage over cross-country comparisons: policy design, implementation and measurement are identical across regions, so these are not dimensions contributing to treatment heterogeneity.

To study the role of potential correlates of treatment magnitude, I follow Crosta et al. (2024) and Angrist and Meager (2023) and augment the Rubin model into a Bayesian Meta-Regression by including a vector of micro-region level covariates $\mathbf{X}_k$. The parent distribution from which true regional estimates $\{\tau_k\}$ are drawn can now be written conditional on values taken by $\mathbf{X}_k$.

$$\tau_k \sim N(\tau_0 + \mathbf{X}_k'\beta, \sigma) \tag{5.3.3}$$

All covariates are normalised to have mean zero and standard deviation one, such that the hyper-parameter τ represents the posterior mean of treatment effect in a micro-region

with all covariates equal to their sample means. The j^{th} element of the vector of coefficients β represents the expected increase on regional treatment effect if the j^{th} covariate increases by one standard deviation. I follow Gelman and Hill (2006) and consider weakly informative priors for all parameters[10].

Potential Covariates

A key aspect of the analysis is the choice of covariates in $\mathbf{X}$[11]. Table 5.2 summarises the considered variables, described in more detail next.

Table 5.2: Potential correlates of treatment effect size

	Type of variable	List
1	Local public spending in different categories	Education Health Physical infrastructure
2	Baseline economy	Size of formal market GDP/capita
3	Growth over study period	Growth in size of formal market Growth in average wage
4	Size of the reform	Change in Bolsa Família spending after the reform

Note: This table lists variables included in the vector $\mathbf{X}$ used in the Bayesian Meta-Regression specified in Equation 5.3.3. Local public spending is defined as average micro-region spending per-capita between 2014 and 2017. Size of the formal market is defined as the share of working-age adults in formal employment in 2013, where the number of working age adults comes from interpolating working-age population counts from the Censuses of 2010 and 2022, and number of formal employees is obtained from RAIS. GDP per capita comes from the National Statistics Office (IBGE). Growth in size of formal market and growth in average wage are computed from RAIS for the period between 2014 and 2017. Change in Bolsa Família spending after the reform is computed by log-differences in total Bolsa Família transfers between October 2014 (post-reform) and April 2014 (pre-reform).

Categories of local public spending Chapters 3 and 4 show that individual-level effects are driven by mothers via increased spending on education. This points towards a role played by the supply-side provision of public goods in enabling households to relax constraints to work. As such, one potentially important factor is the local level of public goods provision. To investigate that, I consider average local public spending during the study period in each micro-region across different categories of public goods provision: education, health and physical infrastructure.

[10] I consider a normal prior with mean zero and variance 1000 for each element j of the vector β and for the hypermean τ, and a half-normal prior with mean zero and variance 1000 for the hyper-standard deviation σ.

[11] Crosta et al. (2024), who compare different UCT programmes in different countries, include variables related to programme design, such as whether the transfer was lump-sum or monthly. Here, programme design is fixed across the country, and my question of interest is the role played by local economic variables in shaping treatment effect.

Baseline economy Given the local economic disparities in Brazil, variables capturing different level of economic activity at the micro-region level are potentially relevant as proxies for aggregate demand. To that, I include as covariates baseline GDP/capita, the size of the formal market (in percentage of the working-age population), the share of working-age adults who report being inactive (as a proxy for available labour, following Egger et al. (2022)) as well as growth in formal employment and in average formal wage over the study period.

Intensity of the 2014 reform Finally, as mentioned in Section 3.5, one potential concern could be that results are driven by general equilibrium effects of the reform itself (Gerard et al., 2024), which can act as stimulus to the local economy and generate employment. Given the nation-wide nature of the threshold increase, micro-regions were affected differently depending on their local income distribution. To address this concern, I include in my vector of potential covariates the size of the reform, represented by the increase in total Bolsa Família budget in each micro-region.

Figure 5.3 illustrates the meta-regression estimates for women and men separately. The dots represent the posterior means for the different elements of the coefficient vector β on formal employment, with their respective 95% credible intervals[12].

For women, the impact on formal employment has a clear positive correlation with local public spending on health and education, and with the availability of slack labour. This is consistent with the mechanisms outlined in Chapter 4 - individual propensity to consume education requires the availability of a supply of schools services. Reassuringly, the size of the reform is largely uncorrelated with treatment effect, which further rules out concerns that results might be driven by general equilibrium effects. Finally, treatment effect is negatively correlated with GDP per capita, suggesting that poorer regions see larger effects. For men, none of the coefficients is significantly different from zero, suggesting that none of the considered regional covariates is significantly related to the variation in treatment effect. The equivalent frequentist OLS is shown in Appendix Table A.16, yielding similar conclusions but somewhat noisier estimates.

[12]In this analysis, I keep only micro-regions where the first stage of equation 3.1.2 is significant at the 95% confidence level - this results in 483 micro-regions for women, and 423 for men

Figure 5.3: Bayesian Meta-Regression estimates by gender

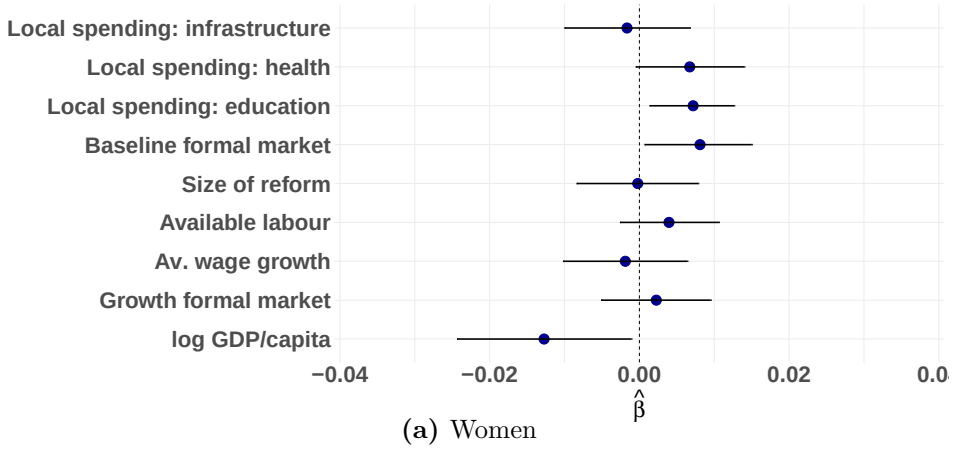

(a) Women

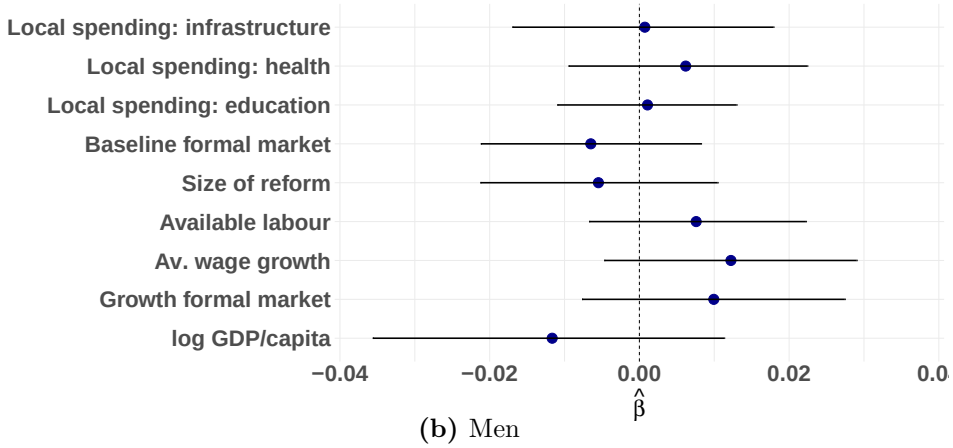

(b) Men

Note: This figure illustrates posterior distributions estimates for the set of $\{\hat{\beta}\}$ in the Bayesian Meta-Regression (Equation 5.3.3). Panel (A) shows estimates for women, and Panel (B) for men. Dots represent the posterior means, and lines the 95% credible intervals from the posterior distribution. Estimates are obtained by Markov Chain Montecarlo Simulations with 4 parallel chains of 3000 iterations, discarding the first 1000 as warm-up. Each coefficient is assumed to have a weakly-informative prior distribution $N(0, 1000)$.

In standard frequentist OLS regressions, the R^2 measures the share of outcome variance that is explained by the included covariates. I consider the Bayesian R^2 extension proposed by Gelman and Pardoe (2006)[13], which can be interpreted as the share of variance

[13]It is not straightforward to extend the notion of share of explained variance from an OLS regression to a Bayesian Meta-Regression, because the latter combines both the addition of explanatory covariates, as well as the pooling of information inherent to the Bayesian "shrinkage. The R^2 proposed by Gelman and Pardoe (2006) is defined as:

$$R^2 = 1 - \frac{E[\underset{k}{\text{Var}}(\tau_k^{post} - X_k'\hat{\beta})]}{E[\underset{k}{\text{Var}}(\tau_k^{post})]}$$

, where τ_k^{post} is micro-region k's posterior mean, the expectation operator $E[.]$ is taken across posterior simulations, and the variance operator $\underset{k}{\text{Var}}(.)$ is taken across micro-regions.

explained *within* each model specification[14]. I find a Bayesian R^2 of 50.8% for women and 40% for men. Appendix Table A.17 reports several other relevant metrics related to explained variance and degree to which information can be pooled across different micro-regions.

5.4 State-level evidence: individual constraints

Next, I turn to the largest level of national subdivision within Brazil - the 27 states - to investigate the relationship between treatment effect and individual constraints to employment. Figure 5.4 illustrates the distribution of estimates across the states, which shows large amounts of geographical variation [15].

Figure 5.4: State-level estimates of treatment effect across Brazil

Note: This figure illustrates the distribution of treatment effects of transfer on formal employment across all 27 Brazilian states. Each value corresponds to the coefficient estimate $\hat{\beta}$ from Equation 3.1.2 estimated separately for each state.

I revisit survey data from PNAD-C, which is representative at the state level, and contains information on self-reported unavailability to work (as illustrated in Figure 1.1. I examine the observed relationship between state-level treatment effects and the share of working-age adults who report being unavailable to work - a proxy for state-level magnitude of constraints to employment.[16].

Figure 5.5 illustrates the empirical relationship between state-level posterior means of treatment effect for women and state-level share of women who self-report being con-

[14] As such, it should not be used to compare different meta-regressions.

[15] Akin to cross-country variation in documented treatment effect, as illustrated by a combination of findings from the literature shown in Appendix Figure B.17.

[16] I consider state k's estimate to be the posterior mean of a Bayesian Hierarchical Model pooling information across all states.

strained to work due to care work constraints. The empirical relationship has an inverted-U shape, consistent with the channel of constraint relaxation: treatment effect is lower if women either face very high constraints, so that the transfer is not enough to relax them, or very low constraints such that the income effect dominates; and is highest if constraints are intermediate. The relationship for men, illustrated in Appendix Figure B.14, acts as a placebo: the share of men constrained by care work is minimal across all states, and has no relationship with treatment effect.

Figure 5.5: State-level treatment effect vs share of women unavailable to work due to care work

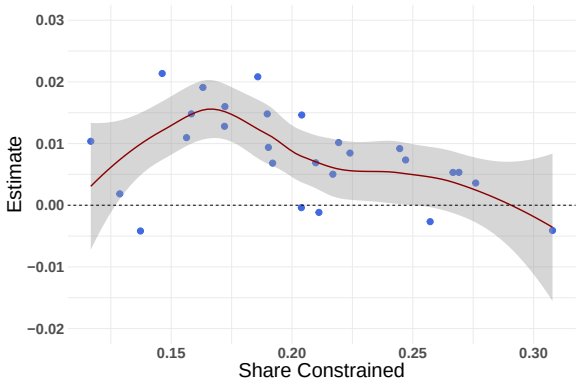

Note: This figure illustrates the relationship between state-level estimates and share of women constrained due to care work. State-level estimates are calculated by estimating Equation 3.1.2 separately for each state, and then pooling estimates from all states via a Bayesian Hierarchical Model. The values on the y-axis are the posterior means for each state. The share of women constrained due to care work is constructed from PNAD-C 2014-16, by computing the proportion of women age 18-45 who: (j) are not employed, (ii) answer no to the question "If you were offered a job today, would you be available to take it?", and (iii) report the reason for unavailability as being "Caring duties for another household member".

Table 5.3 shows the result of estimating a quadratic OLS regression of state-level treatment effect on share of adults reporting constraints to employment. I consider both any constraint (Panel A) and only constraints related to care work (Panel B), as well as both raw estimated state-level treatment effect (columns 1 and 3) and the posterior mean from estimated distributions via the Bayesian Hierarchical Model (columns 2 and 4). For women, the relationship has a negative and significant quadratic term, indicating the inverted-U shaped relationship predicted by the model and observed in Figure 5.5. For men, the treatment effect and the share unavailable are not related.

Table 5.3: Quadratic relationship between state-level treatment effect and constraints to work

	(1)	(2)	(3)	(4)
	A. Unavailability to work: any reason			
	Women		Men	
	Estimated Effect	Bayesian HM Posterior	Estimated Effect	Bayesian HM Posterior
Share Unavailable	-0.046*	-0.016**	-0.067	-0.004
	(0.026)	(0.006)	(0.043)	(0.003)
Share Unavailable^2	-0.047*	-0.010	-0.038	-0.003
	(0.026)	(0.006)	(0.043)	(0.003)
	B. Unavailability to work: caring for another household member			
	Women		Men	
	Estimated Effect	Bayesian HM Posterior	Estimated Effect	Bayesian HM Posterior
Share Unavailable	-0.047*	-0.016**	-0.035	-0.002
	(0.026)	(0.006)	(0.45)	(0.003)
Share Unavailable^2	-0.048*	-0.012*	-0.008	-0.003
	(0.026)	(0.006)	(0.045)	(0.003)

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows estimates for a regression of the form $Y_i = \beta_0 + \beta_1 \cdot \text{Share Unavailable} + \beta_2 \cdot \text{Share Unavailable}^2 + \epsilon_i$, where Y_i is treatment effect of transfer on formal employment in state i, and "Share Unavailable" is the share of respondents age 18-45 who are not employed and answer "No" to the question "If you were offered a job today, would you be available to take it?" in PNAD-C 2014-16. Panel A considers all unavailability for any reason, and panel B considers only due to "Having to care for another household member". Columns (1) and (2) show estimates for women, and columns (3) and (4) for men. Columns (1) and (3) consider Y_i to be the $\hat{\beta}$ estimate for Equation 3.1.2 estimated in state i, and columns (2) and (4) consider Y_i to be the posterior mean of a Bayesian Hierarchical Model that pools information from all 27 states, reweighing the difference-in-differences estimates.

5.5 Discussion

Brazil's unique setting allows me to take advantage of a standardised programme design across a large diversity of local economic contexts (Burgess, 2023). Within-country variation in treatment effects is informative of contextual factors shaping the relationship between an unconditional cash transfer and labour supply. Consistent with the predictions of the occupational choice framework to be outlined in Chater 7, I find that the state-level treatment effects on women have an inverted-U relationship with magnitude of care work constraints.

A finer analysis at the more granular micro-region spatial level suggests that local spending in health and education is strongly correlated to a larger treatment effect for women, but

not for men. This is a supply-side counterpart to the mechanisms enabling women to use a cash transfer to overcome barriers preventing their entry into employment.

The analysis at the municipality level brings a causal underpinning to the strong correlational analysis between treatment effect and supply of local education, pointing towards a strong complementarity between a large unconditional cash transfer and the supply-side of local public services, particularly those in education.

While the findings in Sections 5.3 and 5.2 relate to public spending, rather than the provision of public goods per se, local public spending is strongly correlated with the provision of a number of relevant local public goods and services - Appendix Figures B.18, B.19 and B.20 illustrate this relationship by showing the correlation between spending on, respectively, education, health and infrastructure and several metrics of actual provision of the corresponding local service.

The findings in this section go in line with the suggestive mechanisms discussed in Chapter 4: women with children are the largest beneficiaries of a large unconditional transfer in terms of their increased probability of formal employment, as the transfer represents the overcoming of childcare, education, and/or health constraints. However, for such constraints to be relaxed at the individual level, beneficiaries require the supply of public goods and services in their local area. As such, providing local schools, pre-schools, creches, health centres, and other services for children, are a necessary condition to unlock the positive impacts of the transfer on employment: the two policies are best seen as complements.

Taken together, these findings suggest that women who face constraints to work are more likely to increase their employment if barriers are low enough to be overcome by the transfer amount, and if there is enough supply of education to enable increased school enrolment. These relationships suggest complementarities between a cash transfer and other supply-side policies that reduce barriers to female employment.

Chapter 6

In this chapter, I explore potential fiscal implications from the results on aggregate effect estimated in Chapter 3, as well as the complementarities discussed in Chapter 5.

6.1 Fiscal externalities

A policy that increases formal employment may have fiscal implications via increased tax revenue. In particular, if a time-limited policy has impacts on tax revenue beyond its implementation period, the future flow of revenue must be considered in the present value of its net cost. This parameter is key to welfare analysis and policy decisions on allocation of public funds.

Under the framework introduced by Hendren and Sprung-Keyser (2020), different government funded policies can be compared with respect to one key metric: its Marginal Value of Public Funds (MVPF), defined as the ratio between beneficiaries' willingness to pay and net cost to the government. A policy that pays for itself over has, by definition, a MVPF of infinity.

Figure 6.1 illustrates the evolution of Basic Benefit receipt, as well as its impact on formal employment over a time horizon of approximately four years. Figure 6.1a shows that the gap in benefit receipt caused by the 2014 reform slowly fades with time, and is virtually closed after approximately two years[1]. Figure 6.1b shows that the impact on formal employment persists beyond the closing of the benefit gap, and plateaus at around 2.5pp after four years.

[1]This is due to a combination of: (i) usual churn of beneficiaries in and out of the benefit (see Figure B.9 - the difference on benefit receipt between a group of always eligible and a group of never eligible decreases by about 30pp in two years), and (ii) a subsequent policy reform enacted in 2016, which further raised the threshold of extreme poverty from BRL 77 to BRL 84, making all households with an income from the original range established in my December/2013 sample eligible.

Figure 6.1: Benefit and employment effects over a longer time horizon

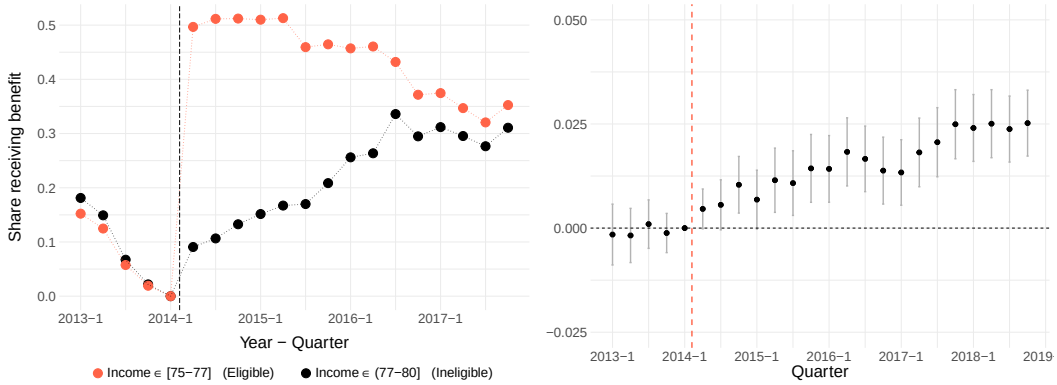

(a) Benefit receipt **(b)** Event study estimates

Note: Panel (A) illustrates the evolution of the quarterly share of individuals receiving Basic Benefit around the 2014 reform by eligibility. Eligibility is defined by household per-capita income in December 2013. It is the extension of Figure 3.1 for a longer time horizon of 3.5 years following the reform. Red (black) dots are individuals living in households with pre-reform income just below (above) the future threshold of BRL 77. Benefit receipt in each quarter is defined as receiving the benefit in all three months of the quarter, except for 2014-Q2, the quarter during which the reform was announced (April) and implemented (June), in which receipt is defined as receiving the benefit at any point in the quarter. Panel (B) illustrates event-study estimates $\{\hat{\delta}_s\}$, obtained by estimating equation 3.1.1 on formal employment for the full sample. It is an extension of Figure 3.2 over a longer time horizon of 3.5 years after the reform. Grey bars represent 95% confidence intervals at each quarter. Standard errors are clustered at the individual level.

This persistent impact of a time-limited transfer difference suggests potential fiscal externalities of Basic Benefit via increased tax revenue. In order to capture this externality, let the parameter T^* be the number of years of constant treatment effect on employment that would be required for the benefit to pay for itself - i.e. for the transfer to have an MVPF of infinity. T^* must be such that:

$$\underbrace{\text{BF} \cdot (1+\beta)}_{\text{Cost of two years of BF}} = \overbrace{\sum_{t=0}^{T^*} \beta^t \cdot \tau \cdot (\text{Av. Tax Revenue per Job})}^{\text{Expected increase in tax revenue per beneficiary}} \tag{6.1.1}$$

Where β is a discount factor, BF is the amount spent per beneficiary and τ is treatment effect on formal employment, assumed to be constant over time. The left-hand side is the cost of two years of Basic Benefit per beneficiary, and the right-hand side is the expected increase in tax revenue due to an increase in individual probability of formal employment. T^* is such that costs equal revenues, i.e. the policy pays for itself under the assumption of constant treatment effect. If T^* is small, then the policy is likely to be fiscally sustainable (or even revenue-increasing) since, as illustrated in Figure 6.1b, the impact on employment outlives the benefit by at least two years. If T^* is very large, then fiscal neutrality becomes implausible, as it would require a treatment effect persisting over a long period of time.

In order to calculate T^*, a key input is the average tax raised per job. To calculate this quantity, I make the following assumptions: (i) conditional on employment, Basic Benefit has no impact on salary[2]; (ii) beneficiaries pay no income tax, so tax revenue comes only from employer's mandatory Social Security contribution which amounts to, on average, 23.89% of formal wage[3]; (iii) real wage grows at 3.17% per year, which is the average wage growth in Brazil between 2013 and 2018. As such, the extra tax revenue raised per job in year t will be equal to the average wage of those who are employed (BRL 790 per month) multiplied by a factor $0.2389 \cdot (1.0317)^{t-1}$. Additionally, I consider a yearly discount factor $\beta = 0.95$ and an average household size of 3.4 persons[4].

Table 6.1: Time for full repayment under different values of treatment effect

	Scenario	τ	T^*
1	Baseline: diff-in-diff coefficient from Table A.2	1.13pp	23 years
2	High public spending: coefficient on the **right** side of the FPM jump from Table 5.1	3.20pp	7 years
3	Low public spending: coefficient on the **left** side of the FPM jump from Table 5.1	0.3pp	$>$ 100 years

Note: This table shows the time for full repayment of Basic Benefit under the assumption of constant treatment effect, as well as assumptions (i)-(iii) outlined in the main text. I also assume: an average salary of BRL 790/month (the observed average salary for the set of formally employed beneficiaries in 2013), a discount factor of 0.95' and an average household size of 3.4 persons. I consider different scenarios of treatment effect: baseline, immediately to right and immediately to the left of the FPM cut-off using a bandwidth of 500 inhabitants (The third panel in Figure **??**). T^* is defined as the number of years under constant treatment effect for full fiscal repayment, as per Equation 6.1.1.

The resulting values of T^* under various scenarios of treatment effect τ are described in Table 6.1. Under the baseline scenario as per the main results presented in Table A.2, Basic Benefit would pay for itself after 23 years of constant treatment effect. If, however, the treatment effect is raised to 3pp, as in the right-hand side of the FPM discontinuity shown in Table 5.1, the time for full repayment decreases drastically to 7 years. This striking difference shows that, in a scenario with better public goods and lower constraints to employment, fiscal returns of the benefit are such that its full repayment through higher tax revenues is plausible on a medium-term time horizon.

[2]See Table A.12 in the Appendix

[3]Employer's Social Security contribution is defined as 20% of monthly formal wage, but it is applied over 13 monthly wages, as Brazilian employees receive an extra month worth of pay in December, as well as over holiday pay, which amounts to 23.89% of average salary.

[4]The benefit is paid at the household level, whereas the impact on employment is calculated at the individual level.

Chapter 7

The impact on employment of a stream of unearned income, such as an unconditional cash transfer, contrasts two economic forces in opposite directions: a neoclassical income effect, discouraging labour, and a constraint relaxation effect, enabling the entry to the labour market. In the previous chapters, I show empirical evidence of the prevalence of the constraint relaxation channel effect over the income effect in the context of Brazil's flagship unconditional cash transfer, as well as the interaction of this effect with local economic variables and complementary policies. In this chapter, I propose a theoretical framework to discipline my analysis through a model of occupational choice under constraints.

7.1 Theoretical framework: Economic forces at play

Income effect On one hand, the neoclassical view in which agents maximise utility over consumption and leisure, both assumed to be normal goods, delivers the unambiguous prediction of a negative effect. This is true as long as no upfront fixed costs to work are present, and agents respond to any increase in unearned income by increasing both consumption and leisure, hence decreasing labour supply (Becker, 1965). Indeed, studies explicitly measuring the effect of large wealth shocks, such as lottery prizes (Golosov et al., 2024), guaranteed basic income (Vivalt et al., 2024), or windfall proceedings from investment funds (Jones and Marinescu, 2022), usually in high-income countries, start from this assumption and develop a theoretical framework to justify empirically-backed negative effects[1]. This consideration has been a key factor behind policy decisions regarding government-funded transfers in high-income countries, such as the widely studied Earned Income Tax Credit scheme in the USA (Kleven, 2024).

Constraint relaxation On the other hand, the impact can be positive if individuals face constraints preventing employment that are relaxed by an increase in non-labour income.

[1] See Cesarini et al. (2017) for a comprehensive example of dynamic theoretical framework.

Previous literature has identified three main mechanisms through which this effect might take place (Baird et al., 2018). First, and most relevant to this study, a transfer can increase employment if it helps recipients overcome a liquidity constraint under imperfect credit markets (Ghatak, 2015). For self-employment, prior studies have documented a causal link between a cash transfer and initial lump-sum investments (Gertler et al., 2012; Banerjee et al., 2015; Bandiera et al., 2017), potentially enabling individuals to escape a poverty trap (Balboni et al., 2022). For wage employment, this constraint relaxation has been shown to happen via funding job search (Caria et al., 2024). A less studied but potentially relevant channel is the possibility of outsourcing home production and care work. A substantial literature has shown a link between outsourcing house work and wage employment, especially for women, via childcare vouchers (Berlinski et al., 2024), supply-side policies (Baker et al., 2008; Müller and Wrohlich, 2020), or technology adoption (Dinkelman, 2011).

Second, a transfer can increase individual productivity, and consequently labour supply, if it helps individuals improve their physical or mental health (Banerjee et al., 2020).

Third, a transfer can have an insurance effect if it increases individual propensity to undertake costly job search with uncertain outcomes, such as migration, which may lead to increased employment prospects (Bryan et al., 2014).

The nature of the mechanisms at play suggests that men and women might see very different effects, depending on the relative relevance on the different channels. Indeed, descriptive evidence shown in Figure 1.1 shows a remarkably gendered pattern of barriers to employment: women are three times more likely to report being unavailable to work as men. This difference is almost exclusively due to care work responsibilities, that constraints 21% of women, but virtually no men. As such, the relaxation of care work constraints is more likely to be a key mechanism enabling employment for women than for men.

The combination of these two forces results in a theoretically ambiguous aggregate effect, depending on which is stronger. The direction of the effect may also be heterogeneous, both in individual characteristics and in contextual factors: for example, if the main constraint to employment is the cost of childcare, it may depend on the provision of local education facilities, and is empirically more likely to be binding for women than for men.

7.2 A Model of Occupational Choice

To illustrate the forces at play, I propose a model of occupational choice that incorporates both economic channels: the standard income effect and constraints to work in the form

of an upfront fixed cost to joining the labour market. I follow the long-standing literature of occupational choice models with heterogeneous agents, drawing elements from Banerjee and Newman (1993), Lloyd-Ellis and Bernhardt (2000), Ghatak and Jiang (2002), Buera et al. (2011) and Feng and Ren (2023).

In line with prior studies, I consider individual heterogeneity in endowment and productivity. Motivated by the differences in practical constraints to employment between men and women documented in Figure 1.1, I introduce a third dimension of heterogeneity not usually considered in this strand of the literature: a binary upfront fixed cost of entering the labour market. This extra dimension of heterogeneity represents the difference in constraints to employment faced by different groups. Empirically, it proxies for gender: women are *ex-ante* more constrained to work than men, mainly due to care work, as documented in Figure 1.1. Theoretically, it is more general than a model of gender differences, and in a different context it can proxy for any group that faces higher employment barriers. I also introduce a government that funds a mean-tested unconditional cash transfer with income taxes levied in the labour market.

Model Setup

A population with measure one of individuals indexed by i draw wealth endowments a_i from a cumulative distribution $F_A(a)$ and productivity z_i from a cumulative distribution $G_Z(z)$. Additionally, individuals draw an upfront fixed-cost of working from home $c_i \in \{0, c_{\text{high}}\}$, with $c_{\text{high}} > 0$. A fraction p of individuals draws $c_i = c_{\text{high}}$, whilst the remaining $1-p$ draw $c_i = 0$, i.e. they have no upfront cost. All three draws are independent of each other.

The model is static and deterministic: individuals live for one period and maximise their present utility over consumption and labour supply $u(x_i, h_i)$, assumed to be separable.

Occupations – Upon learning the realisation for their draws, individuals choose one of the four possible occupations listed below, where Y_i^j is the income generated by occupation j.

- **Idle** - the individual does not engage in labour, and gets a utility from leisure that can be translated into income terms as $\underline{w}$[2].

$$Y_i^{\text{idle}} = \underline{w}$$

[2]This is isomorphic to the subsistence production technology in Banerjee and Newman (1993), which does require labour and produces fixed earnings.

- **Wage employment** - the individual works as a wage employee hired by an entrepreneur, and earns the market equilibrium wage w. In order to enter the labour market, the individual must pay the upfront cost c_i, which is non-zero for a fraction p of individuals.

$$Y_i^{\text{wage}} = w - c_i$$

- **Self-employment** - the individual works autonomously by combining their labour with an of capital k_i hired at a fixed interest rate r. Earnings are a function of their productivity, and capital has diminishing marginal returns - i.e. $\alpha \in [0, 1)$.

$$Y_i^{\text{self}} = z_i \cdot k_i^{\alpha} - r \cdot k_i$$

- **Entrepreneurship** - the individual starts a firm by supplying their labour endowment on supervisory work, and hiring an endogenous number of workers l_i at market wage w. The firms also requires capital, hired at the exogenous fixed interest rate r. Output follows a Cobb-Douglas production function that depends on the entrepreneur's productivity, as well as on a constant scaling factor A. Entrepreneurs also pay the upfront cost of working[3].

$$Y_i^{\text{ent}} = A \cdot z_i \cdot k_i^{\alpha} \cdot l_i^{\beta} - r \cdot k_i - w \cdot l_i - c_i$$

I assume that all occupations, with the exception of choosing to remain idle, require the same amount of labour input, so disutility of labour is not explicitly written, but instead the value of leisure is considered for the idle option. Given the static nature of the model, individuals ultimately only care about income Y_i^j[4].

Credit market – Individuals face credit frictions, and can only hire capital up to a multiple of their endowment:

$$k_i \leq \phi \cdot a_i \text{ , with } \phi \in [1, \infty)$$

The parameter ϕ determines the strictness of the credit constraint. If $\phi = 1$, the economy is an autarky - i.e. individuals can only invest their endowment; if $\phi \to \infty$, credit markets are perfect and individuals can borrow unlimitedly to invest in capital (Buera et al., 2011).

[3] This is based on the assumption that wage workers and entrepreneurs work for a firm (i.e., not from home), and must pay the upfront cost. Relaxing this assumption for entrepreneurs - i.e. only wage workers pay the upfront cost - has no practical consequences for the results.

[4] This is analogous to a dynamic setting with the assumption of a constant savings rate.

Moreover, the credit constraint also binds for the fixed cost of work paid by wage workers and entrepreneurs. Due to its upfront nature, individuals can only choose these two occupations if they can borrow enough to pay c_i, i.e. if $c_i \leq \phi \cdot a_i$.

Taxes and transfers – The economy features a government who implements a means-tested unconditional transfer funded by income taxes (and has no other expenditures). The transfer is defined as a pre-determined amount θ given to all individuals whose endowment falls below an eligibility threshold $\overline{a}$. It is funded by income taxes, levied on wage workers and entrepreneurs[5] as a constant percentage t applied to: (i) wage earned by wage workers above a certain tax-free allowance $w_{\min}$, and (ii) wage paid by entrepreneurs[6].

Competitive equilibrium

Let:

- $O_i(a_i, z_i, c_i)$ by the occupation chosen by individual i whose set of draws for wealth, productivity and upfront cost of work are $\{a_i, z_i, c_i\}$. For every individual, it must be that $O_i \in$ {idle, wage, self, ent}
- $H_{A,Z,C}(a, z, c)$ be the joint distribution of endowment a, productivity z and (binary) upfront cost of work c
- T_i be the net transfer to individual i, defined as:

$$T_i = \theta \cdot \mathbf{1}(a_i \leq \overline{a}) - t \cdot (w - w_{\min}) \cdot \mathbf{1}(O_i = \text{wage}) - t \cdot (w \cdot l_i) \cdot \mathbf{1}(O_i = \text{ent})$$

A competitive equilibrium in the labour market is defined as:

- An equilibrium wage w
- A mapping between the initial draws and an occupational distribution $\{a_i, z_i, c_i\} \rightarrow O_i(a_i, z_i, c_i)$
- Capital choices for individuals choosing self-employment and entrepreneurship $\{k_i^* \mid O_i \in \{\text{self, ent}\}\}$
- Labour choices for individuals choosing entrepreneurship $\{l_i^* \mid O_i = \text{ent}\}$

[5] Which, in this setup, can be seen as a proxy for formality.

[6] This setup mirrors an economy where income tax is levied both on employees and on employers (e.g. in Brazil, my empirical context, this is the case for employer National Insurance).

Such that:

1. Given wage w, all **unconstrained** individuals (i.e. those for whom $c_i \leq \phi \cdot a_i$) choose

$$O_i(a_i, z_i, c_i) = \underset{j}{\text{argmax}} \quad Y_i^j \quad \text{, with} \quad j \in \{\text{idle, wage, self, ent}\}$$

2. Given wage w, all **constrained** individuals (i.e. those for whom $c_i > \phi \cdot a_i$) choose

$$O_i(a_i, z_i, c_i) = \underset{j}{\text{argmax}} \quad Y_i^j \quad \text{, with} \quad j \in \{\text{idle, self}\}$$

3. All individuals who pick **self-employment** choose capital to solve:

$$\{k_i^*\} = \underset{k}{\text{argmax}} \quad Y_i^{\text{self}}(k) \quad \text{s.t.} \quad k_i \leq \phi \cdot a_i$$

4. All individuals who pick **entrepreneurship** choose capital and labour to solve:

$$\{k_i^*, l_i^*\} = \underset{k,l}{\text{argmax}} \quad Y_i^{\text{ent}}(k, l) \quad \text{s.t.} \quad k_i \leq \phi \cdot a_i$$

5. The labour market clears:

$$\underbrace{\int_{(a,z,c)|O_i=\text{wage}} dH(a,z,c)}_{\text{Labour supply from wage workers}} = \underbrace{\int_{(a,z,c)|O_i=\text{ent}} l_i^*(a,z,c) \cdot dH(a,z,c)}_{\text{Labour demand from entrepreneurs}}$$

6. The government balances the budget:

$$\int T_i \cdot dH(a,z,c) = 0$$

Equilibrium distribution of occupations in the absence of transfers

Figure 7.1 illustrates the distribution of occupational choices in equilibrium across the three dimensions of heterogeneity in the absence of a transfer (i.e. $\theta = 0$).

Panel 7.1a, on the left, illustrates the occupational choices for individuals who are not constrained by the upfront fixed-cost of work. The most productive individuals choose to be entrepreneurs, whereas the least productive choose wage employment. The productivity threshold for entrepreneurship is decreasing in endowment, due to the credit constraints[7]. No unconstrained individual chooses to be idle or self-employed.

[7] See Feng and Ren (2023) for a discussion on the misallocation in entrepreneurship caused by this setting

Panel 7.1b, on the right, illustrates the occupational choices for individuals who are constrained by the upfront fixed-cost of work. Up to a certain level of wealth, individuals cannot choose neither wage employment nor entrepreneurship, so the most productive choose self-employment, whereas the least productive choose to remain idle. Beyond the threshold of wealth under which individuals are constrained, the picture looks somewhat similar to the unconstrained panel, except some individuals with an intermediate level of productivity choose self-employment: they are too productive to be wage employees, but not productive enough for entrepreneurship to be worth paying the fixed cost of work.

Figure 7.1: Distribution of occupational choice in a no-transfer equilibrium

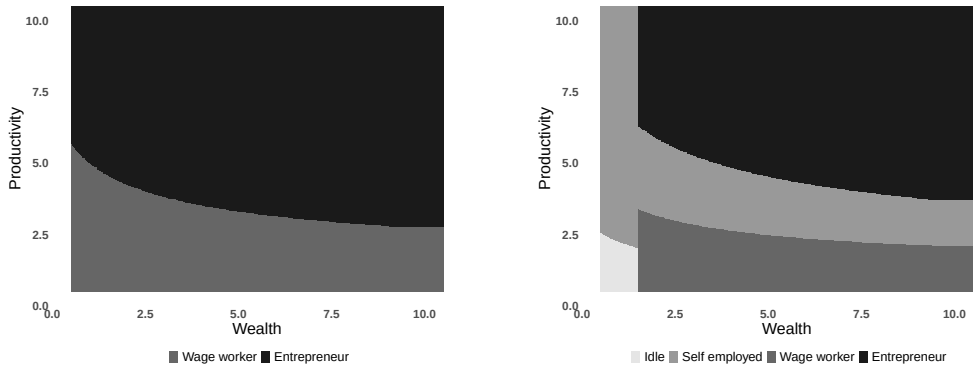

(a) For **unconstrained** individuals ($c_i = 0$) **(b)** For **constrained** individuals ($c_i = c_{\text{high}}$)

Note: This figure illustrates the mapping between wealth (x-axis) and productivity (y-axis) draws, and occupational choice in equilibrium (colours), separately for individuals who draw the upfront cost parameter $c_i = 0$ (unconstrained, in the left panel) and who draw the upfront cost parameter $c_i = c_H$ (constrained, in the right panel). This equilibrium is computed by simulating the model described in Section 7.2 without any transfer, i.e. with $\theta = 0$, and using parameters from Appendix Table A.5.

The model yields two key insights on the relationship between a cash transfer, the magnitude of costs to work faced by women, and the prevalence of wage employment.

Predictions

Wage employment as a function of transfer size – Figure 7.2 illustrates the prevalence of wage employment as a function of transfer size, keeping the magnitude of the upfront cost to work faced by women constant. This relationship is non-monotonic. For transfers which are either too small to relax constraints, or too large, the income effect dominates and the overall effect is negative. For intermediate transfer sizes the effect is positive, as the constraint relaxation channel dominates the income effect.

Figure 7.2: Simulated share in wage employment as a function of transfer size

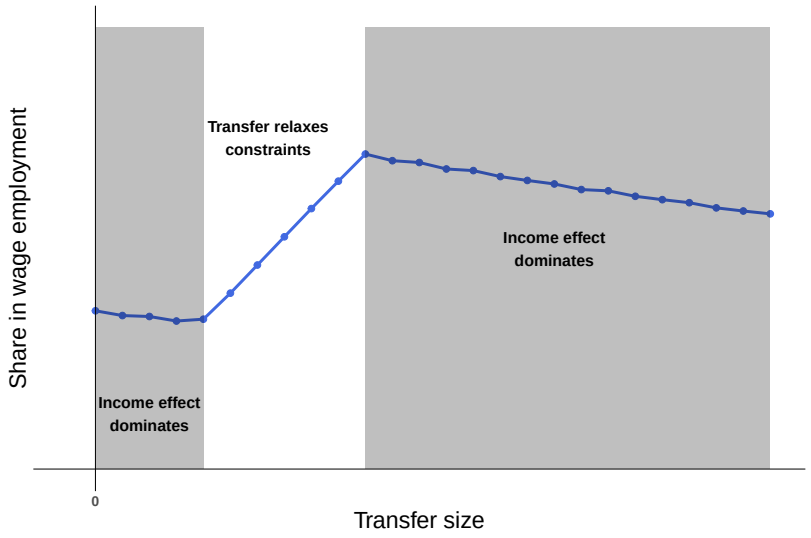

Note: This figure illustrates the equilibrium share of individuals who choose wage employment as a function of transfer size θ for a fixed magnitude of upfront cost c_{high} faced by the constrained group. The remaining parameters for simulations are outlined in Appendix Table A.5.

Appendix Figure B.2 shows an extension of this illustration for three magnitudes of upfront cost: high, medium, and low. The case with medium cost is the same as in Figure 7.2. The other cases illustrate scenarios in which either the transfer is never large enough to overcome the upfront cost (high cost), or the upfront cost is negligible and the income effect always dominates (low cost).

Effect of transfer as a function of cost – Figure 7.3 illustrates the predicted impact of the transfer on the prevalence of wage employment as a function of how costly it is for women to enter the labour market, keeping the transfer amount fixed. The y-axis shows by how much a transfer would increase the share of the population in wage employment compared to an economy without the transfer. For either very small or very high costs, the transfer has a negative impact as the income effect dominates - either the cost to work is not binding even in the absence of a transfer, or the cost is so high that the transfer is not enough to fund it. For intermediate costs, the transfer increases wage employment by relaxing constraints to employment, and this effect dominates the negative income effect. The overall relationship has the shape of an inverted U.

Figure 7.3: Impact of transfer on wage employment vs magnitude of upfront cost to work

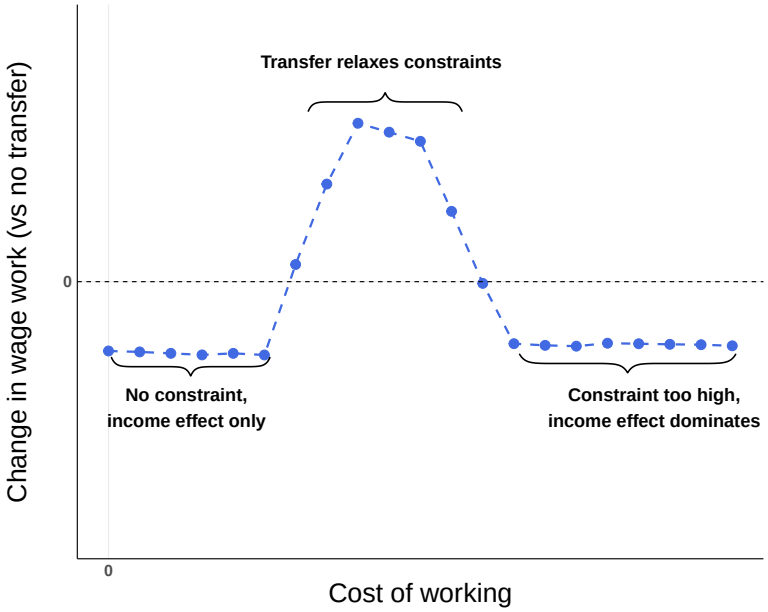

Note: This figure illustrates the effect of a means-tested transfer of fixed size on the share of the population who chooses wage employment, as a function of the value of c_{high}, i.e. of the magnitude of the upfront cost to work faced by the share of the population who draws $c_i = c_{\text{high}}$. The effect at each point is calculated as the difference between the equilibrium share in wage employment with the transfer minus the equilibrium share when the transfer is set to zero. The remaining parameters for simulation are outlined in Appendix Table A.5.

7.3 Discussion

The model presented in this chapter predicts that the direction and magnitude of the impact of a transfer on wage employment depends on the magnitude of the constraints relative to the transfer. A number of findings from the empirical analysis in the previous chapters can be interpreted in the light of this theoretical predictions.

First, the difference in treatment effects for men and women, combined with the descriptive evidence on the overall gender difference in observed constraints to employment (Figure 1.1), empirically validates the idea that the treatment effect of a transfer can be positive in a subgroup of the population that faces constraints, whilst absent for unconstrained individuals.

Second, the relationship between the state-level treatment effect and the share of individuals that report being unavailable to work (a proxy for constraint magnitude, which cannot be directly observed), as illustrated in Figure 5.5 follows an inverted U shape, in similar fashion to the predicted relationship by the model, as illustrated in Figure 7.3. This is intuitive: if individuals are not constrained, or constrained to such an extent that the

transfer is insufficient, there is no positive effect on employment. If, however, individuals face a level of constraint that can be relaxed by the transfer, the impact is positive.

Third, the model also offers an interpretation on the complementarity between a cash transfer and the supply of local public goods. Better public good provision can be thought of as a reduction in the costs of joining the labour force, e.g.: higher supply of education infra-structure can make it easier for mothers to send children to school. Within this framework, an increase in local public good provision can be interpreted as a move leftwards from the right-hand side in Figure 7.3 - i.e. from a high-cost scenario to an intermediate cost scenario.

From a policy perspective, this points towards a direction of possible complementarities. In a world in which baseline constraints are large (i.e. the right-side of Figure 7.3), policies that reduce the upfront cost of working have an impact on how a transfer impacts wage employment. If, as per Figure 1.1, the main constraining factor is care work, these policies could, for example, facilitate outsourcing through supply side policies.

Chapter 8

8.1 Concluding remarks

The rise in women's participation in paid employment is strongly connected to economic development, better living standards, and reduction in extreme poverty. To that end, most governments around the world employ cash transfers programmes as a key anti-poverty policy, and often target them at women. The impact of cash transfers on the labour supply of women is a key parameter for the optimal design of social protection, as its direction determines whether poverty alleviation and employment are at odds, or whether extra gains to employment implies that long-run benefits of cash transfers are higher. Recent meta-analyses of interventions measuring this impact have found large heterogeneities in treatment effects, and whilst concerns of "lazy welfare recipients" have been found to be largely unwarranted (Banerjee et al., 2017), no consensus has emerged over the broader nature of this parameter. The empirical dispersion, paired with a widely discussed theoretical ambiguity on the direction of the effect (Baird et al., 2018), suggests a relevant role for local context.

In this study, I take advantage of the institutional setting of Brazil's Bolsa Família, a national programme with a large unconditional component, that is designed and implemented in identical fashion across the entire country. This feature allows me to isolate the role of local context in shaping regional variation in the effect on employment of men and women. Using plausibly exogenous variation in benefit receipt generated by a policy reform that increases benefit coverage, I find that, on average across the country, the benefit increases formal employment of women by 1.13pp (7.4%), but has no effect on men. The positive effect on women in highest for mothers, particularly those with children of pre-school age, who increase their propensity to spend money on children's education, resulting in increased school enrolment and freeing up of time previously dedicated to care work.

This average result masks large amount of geographical heterogeneity. Using a Bayesian

meta-regression, I study the main geographical covariates of treatment effect. I find that impacts are largest in poorer areas, but that have higher spending in local public goods provision - education. This correlation is a supply-side counterpart to the individual level evidence on increased school enrolment freeing up time.

I take advantage of the rules of public funds allocation across Brazilian municipalities to show causal evidence of complementarity between spending on local public goods provision and the effect of transfer on employment. I find that a plausibly exogenous 4% increase to local budget, most of which is directed to spending on local education, is associated with an increase in treatment effect to around 3pp.

Taken together, my empirical findings point towards the relevance of supply-side policies improving the availability of education infra-structure in complementing the impacts of cash transfer on women's labour supply. These results can be interpreted through a theoretical framework of occupational choice with heterogeneous upfront costs to employment - proxying for observed gender differences in care work constraints. Women can use a stream of cash to increase their choice set of occupations and join the labour market; whereas men are a priori unconstrained. A supply-side policy can complement the transfer if it reduces the magnitude of constraints such that they can be overcome with the transfer.

Policy relevance

My findings have important policy implications: provision of local public goods may have previously unaccounted for labour market externalities if they increase the impact of cash transfer on labour supply. Conversely, long-term evaluations of cash transfers undervalue their true benefit if future gains are not taken into consideration. I show that, beyond welfare gains from optimal policy design, this complementarity has direct fiscal implications for the government. A cash transfer that increases formal employment will generate stream of future tax revenue - if this impact is persistent beyond the immediate payment period, the transfer can be fiscally neutral in the long run. In the MVPF framework widely adopted to evaluate government-funded policies (Hendren and Sprung-Keyser, 2020), this would be the equivalent of infinite marginal value of public funds.

Moreover the differential impact by gender - strong and positive for women, not significantly different from zero for men - suggests that cash transfers can be a strong driver of gender parity in the labour market. Equal opportunities for employment are not only fair, but also efficient, as higher barriers for women imply misallocation at a cost to aggregate output.

Future work

This study suggests a number of potentially fertile avenues for future research. First, the theoretical framework developed can be taken further into a quantitative exercise embedding the reduced form moments at the individual level to fully quantify an optimal spending mix. This expansion can also incorporate other relevant dimensions not explored here, such as migration, external economic shocks, gender norms, and sectoral composition.

Second, the supply-side complementarity angle developed here can be applied to other relevant dimensions of cash transfers, such as formation of human capital and inter-generational mobility - which are first order targets of cash transfers in their origin.

Third, the idea that the local economic factors present *ex-ante* can play a role in shaping the impact of a policy can be applied to other interventions where individuals are likely to interact with market forces, especially in the field of Development Economics (e.g.: Graduation, microfinance, in-kind transfers). Such contextual heterogeneity can play an important role in helping researchers understand how and where interventions aimed at poverty reduction can be most effective.

Bibliography

M. Angelucci and G. De Giorgi. Indirect effects of an aid program: how do cash transfers affect ineligibles' consumption? *American economic review*, 99(1):486–508, 2009. 7

N. Angrist and R. Meager. Implementation matters: Generalizing treatment effects in education. *Available at SSRN 4487496*, 2023. 5, 41

C. Ardington, A. Case, and V. Hosegood. Labor supply responses to large social transfers: Longitudinal evidence from south africa. *American economic journal: Applied economics*, 1(1):22–48, 2009. 6, 35

N. Ashraf, O. Bandiera, V. Minni, and V. Quintas-Martínez. Gender and the misallocation of labor across countries. *Work in progress, June*, 2023. 2

O. P. Attanasio and V. Lechene. Efficient responses to targeted cash transfers. *Journal of political Economy*, 122(1):178–222, 2014. 2

S. Baird, D. McKenzie, and B. Özler. The effects of cash transfers on adult labor market outcomes. *IZA Journal of Development and Migration*, 8:1–20, 2018. 2, 53, 62

M. Baker, J. Gruber, and K. Milligan. Universal child care, maternal labor supply, and family well-being. *Journal of political Economy*, 116(4):709–745, 2008. 53

C. Balboni, O. Bandiera, R. Burgess, M. Ghatak, and A. Heil. Why do people stay poor? *The Quarterly Journal of Economics*, 137(2):785–844, 2022. 53

O. Bandiera, G. Fischer, A. Prat, and E. Ytsma. Do women respond less to performance pay? building evidence from multiple experiments. 2016. 41

O. Bandiera, R. Burgess, N. Das, S. Gulesci, I. Rasul, and M. Sulaiman. Labor markets and poverty in village economies. *The Quarterly Journal of Economics*, 132(2):811–870, 2017. 2, 53

O. Bandiera, A. Elsayed, A. Heil, and A. Smurra. Presidential address 2022: Economic development and the organisation of labour: Evidence from the jobs of the world project. *Journal of the European Economic Association*, 20(6):2226–2270, 2022a. 2

O. Bandiera, N. Parekh, B. Petrongolo, and M. Rao. Men are from mars, and women too: A bayesian meta-analysis of overconfidence experiments. *Economica*, 89:S38–S70, 2022b. 41

A. Banerjee, E. Duflo, N. Goldberg, D. Karlan, R. Osei, W. Parienté, J. Shapiro, B. Thuysbaert, and C. Udry. A multifaceted program causes lasting progress for the very poor: Evidence from six countries. *Science*, 348(6236):1260799, 2015. 53

A. Banerjee, D. Karlan, H. Trachtman, and C. R. Udry. Does poverty change labor supply? evidence from multiple income effects and 115,579 bags. Technical report, National Bureau of Economic Research, 2020. 3, 53

A. V. Banerjee and A. F. Newman. Occupational choice and the process of development. *Journal of Political Economy*, 101(2):274–298, 1993. 6, 54

A. V. Banerjee, R. Hanna, G. E. Kreindler, and B. A. Olken. Debunking the stereotype of the lazy welfare recipient: Evidence from cash transfer programs. *The World Bank Research Observer*, 32(2):155–184, 2017. 3, 6, 36, 62, 108

F. Bastagli, J. Hagen-Zanker, L. Harman, V. Barca, G. Sturge, T. Schmidt, and L. Pellerano. Cash transfers: what does the evidence say? *A rigorous review of programme impact and the role of design and implementation features. London: ODI*, 1(7):1, 2016. 2, 3, 6, 36, 108

G. S. Becker. A theory of the allocation of time. *The economic journal*, 75(299):493–517, 1965. 2, 52

M. Bergolo and G. Cruces. The anatomy of behavioral responses to social assistance when informal employment is high. *Journal of Public Economics*, 193:104313, 2021. 2, 19

K. Bergstrom, W. Dodds, and J. Rios. Welfare analysis of changing notches: Evidence from bolsa famılia. 2022. 2, 7, 24

S. Berlinski, M. M. Ferreyra, L. Flabbi, and J. D. Martin. Childcare markets, parental labor supply, and child development. *Journal of Political Economy*, 132(6):000–000, 2024. 6, 53

M. Bertrand, C. Goldin, and L. F. Katz. Dynamics of the gender gap for young professionals in the financial and corporate sectors. *American economic journal: applied economics*, 2(3):228–255, 2010. 6

E. Boserup. *Women's Role in Economic Development*. Allen & Unwin, 1970. 2

Brasil. Constituicao da república federativa do brasil, artigo 212, 1988. 38

F. Brollo, T. Nannicini, R. Perotti, and G. Tabellini. The political resource curse. *American Economic Review*, 103(5):1759–1796, 2013. 11

G. Bryan, S. Chowdhury, and A. M. Mobarak. Underinvestment in a profitable technology: The case of seasonal migration in bangladesh. *Econometrica*, 82(5):1671–1748, 2014. 53

F. J. Buera, J. P. Kaboski, and Y. Shin. Finance and development: A tale of two sectors. *American Economic Review*, 101(5):1964–2002, 2011. 54, 55

R. Burgess. Igc/yale conference on firms/trade/development – keynote lecture on job variety. 2023. 47

S. Calonico, M. D. Cattaneo, and R. Titiunik. Optimal data-driven regression discontinuity plots. *Journal of the American Statistical Association*, 110(512):1753–1769, 2015. 39, 90, 107

J. M. Camargo and F. H. G. Ferreira. O beneficio social único: uma proposta de reforma da política social no brasil. 2001. 9

A. S. Caria, G. Gordon, M. Kasy, S. Quinn, S. O. Shami, and A. Teytelboym. An adaptive targeted field experiment: Job search assistance for refugees in jordan. *Journal of the European Economic Association*, 22(2):781–836, 2024. 2, 34, 53

Casa Civil. Presidential decree number 8232. Technical report, 2004. 11

D. Cesarini, E. Lindqvist, M. J. Notowidigdo, and R. Östling. The effect of wealth on individual and household labor supply: evidence from swedish lotteries. *American Economic Review*, 107(12):3917–3946, 2017. 6, 52

L. Chioda, J. M. De Mello, and R. R. Soares. Spillovers from conditional cash transfer programs: Bolsa família and crime in urban brazil. *Economics of Education Review*, 54: 306–320, 2016. 7

R. Corbi, E. Papaioannou, and P. Surico. Regional transfer multipliers. *The Review of Economic Studies*, 86(5):1901–1934, 2019. 11, 37

T. Crosta, D. Karlan, F. Ong, J. Rueschenpoehler, and C. Udry. Unconditional cash transfers: A bayesian meta-analysis of randomized evaluations in low and middle income countries. *Working Paper*, 2024. 3, 6, 19, 36, 41, 42, 108

M. Cruz and Z. Ziegelhofer. Beyond the income effect: impacts of conditional cash transfer programs on private investments in human capital. *World Bank Policy Research Working Paper*, (6867), 2014. 7

P. Dasgupta and D. Ray. Inequality as a determinant of malnutrition and unemployment: Theory. *The Economic Journal*, 96(384):1011–1034, 1986. 3, 34

A. De Brauw, D. O. Gilligan, J. Hoddinott, and S. Roy. The impact of bolsa família on women's decision-making power. *World Development*, 59:487–504, 2014. 7

G. Diaz-Pardo and M. Rao. Women and cash transfers: How program design and local conditions relate to causal estimates of impact. *G2LM—LIC Synthesis Paper No. 13*, 2024. 3, 6, 36, 108

T. Dinkelman. The effects of rural electrification on employment: New evidence from south africa. *American Economic Review*, 101(7):3078–3108, 2011. 53

R. Dix-Carneiro and B. K. Kovak. Trade liberalization and regional dynamics. *American Economic Review*, 107(10):2908–2946, 2017. 41

M. Doepke and M. Tertilt. Women's liberation: What's in it for men? *Quartely Journal of Economics*, 124(4):1541–1591, 2009. 2

E. Duflo. Women empowerment and economic development. *Journal of Economic literature*, 50(4):1051–1079, 2012. 2

D. Egger, J. Haushofer, E. Miguel, P. Niehaus, and M. Walker. General equilibrium effects of cash transfers: experimental evidence from kenya. *Econometrica*, 90(6):2603–2643, 2022. 7, 25, 43

P. Ervin, E. Elisenda, D. Silvio, D. Benjamin, V. U. Paul, M. Stanfeld, H. Sudhanshu, S. David, and T. Gelson. Learning about labour impacts of cash transfers in zambia. *Journal of African Economies*, 26(4):433–442, 2017. 7

E. Fassarella, S. Ferreira, S. Franco, V. P. Neto, G. Ribeiro, V. Schuabb, and P. Tafner. Social mobility and cct programs: The bolsa família program in brazil. *World Development Perspectives*, 35:100624, 2024. 7

Y. Feng and J. Ren. Skill bias, financial frictions, and selection into entrepreneurship. *Journal of Development Economics*, 162:103046, 2023. 6, 54, 57

S. Firpo, R. Pieri, E. Pedroso Jr, and A. P. Souza. Evidence of eligibility manipulation for conditional cash transfer programs. *EconomiA*, 15(3):243–260, 2014. 7

A. Gelman and J. Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press, 2006. 42

A. Gelman and I. Pardoe. Bayesian measures of explained variance and pooling in multilevel (hierarchical) models. *Technometrics*, 48(2):241–251, 2006. 44, 89

U. Gentilini. *Cash transfers in pandemic times: Evidence, practices, and implications from the largest scale up in history*. World Bank, 2022. 2

F. Gerard and G. Gonzaga. Informal labor and the efficiency cost of social programs: Evidence from unemployment insurance in brazil. *American Economic Journal: Economic Policy*, 13(3):167–206, 2021. 2

F. Gerard, J. Naritomi, and J. Silva. Cash transfers and formal labor markets: Evidence from brazil. 2024. 2, 7, 9, 24, 25, 43, 98, 101

P. J. Gertler, S. W. Martinez, and M. Rubio-Codina. Investing cash transfers to raise long-term living standards. *American Economic Journal: Applied Economics*, 4(1):164–192, 2012. 53

M. Ghatak. Theories of poverty traps and anti-poverty policies. *The World Bank Economic Review*, 29(suppl_1):S77–S105, 2015. 53

M. Ghatak and N. N.-H. Jiang. A simple model of inequality, occupational choice, and development. *Journal of Development Economics*, 69(1):205–226, 2002. 54

G. Giupponi and C. Landais. Subsidizing labour hoarding in recessions: the employment and welfare effects of short-time work. *The Review of economic studies*, 90(4):1963–2005, 2023. 23, 96

C. Goldin. A grand gender convergence: Its last chapter. *American economic review*, 104 (4):1091–1119, 2014. 2

M. Golosov, M. Graber, M. Mogstad, and D. Novgorodsky. How americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. *The Quarterly Journal of Economics*, 139(2):1321–1395, 2024. 6, 52

J. Gruber. Disability insurance benefits and labor supply. *Journal of political economy*, 108(6):1162–1183, 2000. 6

R. Heath and S. Jayachandran. *The causes and consequences of increased female education and labor force participation in developing countries*. Oxford University Press, 2018. 2

R. Heath and A. M. Mobarak. Manufacturing growth and the lives of bangladeshi women. *Journal of development Economics*, 115:1–15, 2015. 2

N. Hendren and B. Sprung-Keyser. A unified welfare analysis of government policies. *The Quarterly Journal of Economics*, 135(3):1209–1318, 2020. 49, 63

C.-T. Hsieh, E. Hurst, C. I. Jones, and P. J. Klenow. The allocation of talent and us economic growth. *Econometrica*, 87(5):1439–1474, 2019. 2

S. Jayachandran. Social norms as a barrier to women's employment in developing countries. imf economic review, 69 (3), 576–595, 2021. 2

D. Jones and I. Marinescu. The labor market impacts of universal and permanent cash transfers: Evidence from the alaska permanent fund. *American Economic Journal: Economic Policy*, 14(2):315–340, 2022. 6, 52

H. Kleven. The eitc and the extensive margin: A reappraisal. *Journal of Public Economics*, 236:105135, 2024. 52

H. Kleven, C. Landais, and J. E. Søgaard. Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209, 2019. 6

H. Kleven, C. Landais, and G. Leite-Mariante. The child penalty atlas. *Review of Economic Studies*, 2024. 6, 35

F. Kondylis and J. Loeser. Intervention size and persistence. 2021. 7

H. Lloyd-Ellis and D. Bernhardt. Enterprise, inequality and economic development. *The Review of Economic Studies*, 67(1):147–168, 2000. 6, 54

A. P. B. Martins and C. A. Monteiro. Impact of the bolsa família program on food availability of low-income brazilian families: a quasi experimental study. *BMC Public Health*, 16:1–11, 2016. 29

MDS. Portaria numero 617, 2010. 10

R. Meager. Understanding the average impact of microcredit expansions: A bayesian hierarchical analysis of seven randomized experiments. *American Economic Journal: Applied Economics*, 11(1):57–91, 2019. 5, 41, 89

T. Molina and J. Vidiella-Martin. Conditional cash transfers, schooling decisions, and export manufacturing. *Working Paper*, 2023. 7

K.-U. Müller and K. Wrohlich. Does subsidized care for toddlers increase maternal labor supply? evidence from a large-scale expansion of early childcare. *Labour Economics*, 62:101776, 2020. 6, 53

P. Niehaus and T. Suri. Cash transfers. In R. Hanna and B. Olken, editors, *The Handbook of Social Protection: Evidence to Inform Policy in Low- and Middle-Income Countries*. 2024. 2

C. Olivetti and B. Petrongolo. The evolution of gender gaps in industrialized countries. *Annual review of Economics*, 8(1):405–434, 2016. 2

N. Qian. Missing women and the price of tea in china: The effect of sex-specific earnings on sex imbalance. *The Quarterly journal of economics*, 123(3):1251–1285, 2008. 2

D. B. Rubin. Estimation in parallel randomized experiments. *Journal of Educational Statistics*, 6(4):377–401, 1981. 41

G. Ulyssea. Firms, informality, and development: Theory and evidence from brazil. *American Economic Review*, 108(8):2015–2047, 2018. 19

I. A. V. O. Viana, M. O. Kawauchi, and T. V. O. Barbosa. Bolsa família 15 anos (2003-2018). 2018. 9

E. Vivalt, E. Rhodes, A. W. Bartik, D. E. Broockman, and S. Miller. The employment effects of a guaranteed income: Experimental evidence from two u.s. states. *National Bureau of Economic Research WP32719*, 2024. 52

A. Young. The african growth miracle. *Journal of Political Economy*, 120(4):696–739, 2012. 2

Appendix A: Additional Tables

Table A.1: Proportion of beneficiaries assumed to fall under the Permanence Rule by eligibility to Basic Benefit in December/2012

Assumed to fall under Permanence Rule	Share with income pre capita $\leq$ BRL 70 in Dec/2012
YES	41.2%
NO	16.7%

Note: This table assesses the plausibility of my assumption that individuals who receive the benefit in the first quarter of 2014 do so because they fall under the government's Permanence Rule, which allows beneficiaries to continue receiving the benefit for two years if they truthfully report an increase in income. The first row computes the share of all individuals who receive benefit in the first quarter of 2013 (by income alone, none should be receiving - and are hence assumed to fall under Permanence Rule) that were eligible for the benefit (i.e. had an income below BRL 70) the year before. The second row computes the share of individuals who do not receive the benefit in the first quarter of 2013 (and as such do not fall under the Rule) who were eligible for the benefit the year year before. Given that the Permanence Rule lasts two years, the ideal check would be to look at the evolution of their household income over the two year prior to 2013, but data is only available at one point in time: the yearly 2012 Cadastro Único snapshot. The fact that, under this imperfect check, the share eligible in 2012 among those assumed to fall under the Permanence Rule is much larger than among those not in the Permanence Rule lends plausibility to the assumption.

Table A.2: Effects on different types of formal employment

	(1) All Formal	(2) Public Sector	(3) Private Sector	(4) Clerk, Technician (ISCO 3,4)	(5) Services (ISCO 5)	(6) Agriculture (ISCO 6)	(7) Construction, Industry (ISCO 7,8)
Post · Benefit	0.0113***	0.0000	0.0113***	0.0021***	0.007***	0.0003	0.0003
	(0.003)	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
N	310755	310755	310755	310755	310755	310755	310755
Mean Inelig.	0.178	0.017	0.161	0.045	0.073	0.017	0.067

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 for different outcome variables Y_{it} on the full sample. Column (1) shows the effect on any formal employment, columns (2) and (3) the effect on public vs private formal employment, and columns (4)-(7) show the effect on different occupation categories. Standard errors are clustered at the individual level.

Table A.3: Effects on formal employment by quartile of reform shock

	Quartile of reform shock			
	(1)	(2)	(3)	(4)
Post · Benefit	0.0125**	0.0132***	0.0074	0.0109
	(0.006)	(0.003)	(0.007)	(0.008)
N	64654	127543	63326	55206
Mean Ineligible	0.214	0.188	0.153	0.143

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 for formal employment separately for individuals living in areas of different exposure to the 2014 reform. All Brazilian municipalities are split into quartiles according to the log-change in total Bolsa Família budget as a result of the reform, and Equation 3.1.2 is estimated for each quartile separately. Standard errors are clustered at the individual level.

Table A.4: Average baseline monthly earnings by self-declared job type

Job Type	Av. Monthly Earnings (BRL)
Formal employee	749.83
Informal employee	446.07
Self-employed	290.40

Note: This table shows the average self-reported monthly earnings according to self-reported occupational category at baseline for all individuals in my main sample. The sample includes all individuals between the ages of 18 and 60 who, in December/2013, live in household with per-capita income around the future threshold of BRL 77, excluding individuals who are in full-time education and those who are assumed to fall under the Permanence Rule.

Table A.5: Model parameters for simulations in Chapter 7

Parameter	Value	Obs
c_{high}	0 - 4	Fixed at 2 in Figure 7.2
$\underline{w}$	3	
α	0.3	
β	0.2	
r	0.05	
θ	3	Varying from 0 to 2.5 in Figure 7.2
ϕ	1	
$\overline{a}$	5	
w_{min}	4	
	Distributions	
a_i	$\sim$ log-Normal (1, 30)	
z_i	$\sim$ log-Normal (1, 30)	

Note: This table lists the value for all model parameters used in the simulations presented in Chapter 7.

Table A.6: Effects on different types of employment by gender

	(1) Any job	(2) Self-employed	(3) Formal Employee
	Panel A: Everyone		
Post · Benefit	0.061***	0.042***	0.011***
	(0.005)	(0.004)	(0.003)
Mean Ineligible	0.444	0.288	0.179
N	301024	301024	310755
	Panel B: Women		
Post · Benefit	0.068***	0.047***	0.011***
	(0.005)	(0.005)	(0.004)
Mean Ineligible	0.382	0.260	0.148
N	207498	207498	211172
	Panel C: Men		
Post · Benefit	0.036***	0.024***	0.004***
	(0.008)	(0.008)	(0.006)
Mean Ineligible	0.564	0.342	0.236
N	93526	93526	99583
Source	Cadastro (Self-declared)	Cadastro (Self-declared)	RAIS (Admin)

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 where Y_{it} is an indicator for self-declared employment of each type, conditional on remaining in Cadastro Único until 2017, for different sub-samples. Panel A shows the effects on the full sample. Panel B show the effects on women, and panel C shows the effects on men. Column (1) shows the effect on any employment, column (2) on self-employment, column (3) on formal wage employment, and column (4) on informal wage employment.

Table A.7: Effects on different types of employment by parenthood

	(1) Any job	(2) Self-employed	(3) Formal Employee
	Panel A: Mothers		
Post · Benefit	0.070***	0.050***	0.012***
	(0.006)	(0.005)	(0.004)
Mean Ineligible	0.426	0.290	0.159
N	133056	133056	134434
	Panel B: Non-mothers		
Post · Benefit	0.051***	0.032***	0.004
	(0.011)	(0.009)	(0.006)
Mean Ineligible	0.314	0.212	0.130
N	74442	74442	76738
	Panel C: Fathers		
Post · Benefit	0.032***	0.020***	0.000
	(0.007)	(0.007)	(0.007)
Mean Ineligible	0.707	0.445	0.266
N	46333	46333	49185
	Panel D: Non-Fathers		
Post · Benefit	0.048**	0.035**	0.016
	(0.020)	(0.017)	(0.014)
Mean Ineligible	0.409	0.230	0.205
N	47193	47193	50398
Source	Cadastro (Self-declared)	Cadastro (Self-declared)	RAIS (Admin)

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 where Y_{it} is an indicator for self-declared employment of each type, conditional on remaining in Cadastro Único until 2017, for different sub-samples. Panel A shows the effects on mothers. Panel B show the effects on non-mothers. Panel C shows the effects on fathers Panel D shows the effects on non-fathers. Mothers and fathers are defined as those who have either a son or a daughter under the age of 18 living in the same household. Column (1) shows the effect on any employment, column (2) on self-employment, column (3) on formal wage employment, and column (4) on informal wage employment.

Table A.8: Effects on formal employment by region

	(1) Everyone	(2) Men	(3) Women
Panel A: South and Southeast			
Post · Benefit	0.0149***	0.001	0.0170***
	(0.004)	(0.008)	(0.004)
N	144065	43690	100375
Mean Ineligible	0.211	0.272	0.179
Panel B: Northeast			
Post · Benefit	0.010**	-0.003	0.0122***
	(0.005)	(0.012)	(0.005)
N	101159	36121	69038
Mean Ineligible	0.142	0.200	0.1066
Panel C: North and Centre-West			
Post · Benefit	0.007	0.024**	-0.003
	(0.005)	(0.012)	(0.006)
N	61531	19772	41759
Mean Ineligible	0.168	0.217	0.139

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 for formal employment separately for individuals living in each Brazilian region, according to their municipality of residence at baseline. Standard errors are clustered at the individual level.

Table A.9: Effects on formal employment by formal market attachment between 1999 and 2013

	(1) Everyone	(2) Men	(3) Women
Panel A: Never had a formal job			
Post · Benefit	0.009***	0.008	0.011***
	(0.002)	(0.005)	(0.002)
N	157256	36868	120388
Mean Ineligible	0.024	0.024	0.024
Panel B: Had formal job at some point			
Post · Benefit	0.013***	0.006	0.011**
	(0.004)	(0.008)	(0.005)
N	153499	62715	90784
Mean Ineligible	0.337	0.359	0.318

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 for formal employment, separating individuals by whether or not they had previously had a formal employment at any point in the 15 years prior to the reform (1999-2013). Standard errors are clustered at the individual level.

Table A.10: Average school enrolent by age in 2010 Census (in % of each age)

				Regular school, mandatory				
Age	None	Nursery	Pre-School	Primary	Middle	Secondary	Tertiary	Other
0	95.95	4.05	0	0	0	0	0	0
1	88.42	11.58	0	0	0	0	0	0
2	76.13	23.87	0	0	0	0	0	0
3	54.62	19.93	25.45	0	0	0	0	0
4	29.14	15.44	46.34	8.91	0.17	0	0	0
5	13.23	0	61.47	25.20	0.09	0	0	0
6	5.03	0	22.86	72.02	0.09	0	0	0
7	2.75	0	6.46	90.67	0.12	0	0	0
8	2.26	0	3.38	94.22	0.14	0	0	0
9	2.05	0	2.22	93.04	2.69	0	0	0
10	2.18	0	1.59	75.33	20.90	0	0	0
11	2.65	0	0.95	35.13	61.27	0	0	0
12	3.01	0	0	18.50	75.04	0	0	3.45
13	3.68	0	0	11.41	76.01	4.31	0	4.58
14	5.85	0	0	7.90	65.41	15.88	0	4.96
15	9.43	0	0	6.20	37.57	41.92	0	4.89
16	15.56	0	0	4.60	22.76	51.93	0.40	4.75
17	28.47	0	0	3.18	12.68	42.80	3.01	9.85
18	51.01	0	0	2.22	6.80	22.50	9.35	8.12
19	62.85	0	0	1.71	3.91	12.90	13.07	5.57
20	69.30	0	0	1.35	2.69	8.19	14.26	4.20

Note: This table shows the school enrolment rates of the Brazilian population younger than 21, by age and schooling stage according to the 2010 Census. All numbers are percents of each age. "Tertiary" education includes any post-secondary schooling. "Other" inlcudes special schooling regimes for older students who are not able to attend regular school, and usually takes place in the evenings.

Table A.11: Covariates used in Propensity Score matching from survey data

Survey	Variables
POF	State (set of dummies)
(Unit of observation:	Income per-capita excluding Bolsa Família
household)	Rooms per inhabitant
	Bathrooms per inhabitant
	Household has connection to piped water (dummy)
	Household is located an urban area (dummy)
	Education of household head (set of dummies)
	Household size
	Share of household members age 0-9
	Share of household members age 10-15
	Share of household members age 16-20
	Share of household members age 21-65
	Share of women
PNAD-C	State (set of dummies)
(Unit of observation:	Household income per-capita excluding all social programmes
individual)	Household is located an urban area (dummy)
	Gender
	Race (set of dummies)
	Education (set of dummies)
	Household size
	Share of household members age 0-9
	Share of household members age 10-15
	Share of household members age 16-20
	Share of household members age 21-65
	Share of women in the households

Note: This table describes the set of covariates used as predictors of Basic Benefit receipt in the Propensity Score Matching (PSM) analysis of mechanisms in Chapter 4.

Table A.12: Effects on log-formal salary, conditional on being employed

	(1)	(2)	(3)	(4)	(5)
		Full sample, by gender		*Women, by motherhood*	
	Full sample	Men	Women	Mothers	Non-mothers
Post · Benefit	-0.042*	-0.024	-0.042	-0.050	-0.025
	(0.025)	(0.047)	(0.031)	(0.033)	(0.065)
N	15034	7269	7755	4999	2756
Mean Ineligible	7.40	7.38	7.43	7.41	7.45

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 where Y_{it} is defined as log-salary in the formal market, for different sub-samples. I only include individuals in formal employment, so the effect should be interpret as the impact on salary conditional on employment. Column (1) shows the effects on the full sample. Columns (2) and (3) show the effect for men and women separately. Columns (4) and (5) show the effect for mothers and non-mothers. Mothers (fathers) fathers are defined as women (men) who have either a son or a daughter under the age of 18 living in the same household.

Table A.13: Components of spending categories from POF 2008-2009

Category	Set of items
Childcare	Nanny
Pre-school	Pre-school, creche, nursery
School Fees	Private classes, regular school fees (primary school, middle school, high school, higher education), school enrolment fees, material fees, school library fees, exam fees, special arrangements for learners with disability, fees for preparation course
After-school	School vacation camps, equipment for extra-curricular activities, after-school classes of: gymnastics, ballet, dancing, music, computer use, swimming, language, capoeira, football, volleyball, basketball, tennis, handball, martial arts, boxing, judo, karate, cooking, theatre, arts, painting
School Material	School books, dictionaries, apostilles, school uniform, school lunch box, calendar, backpack, notebook, pencil, pen, eraser, pencil sharpener, pencil case, paintbrush, scissors, paint, paper, paper clips, staples, stapler, notepad, carton, glue, tape, ruler, protractor, compass, setsquare, envelope, lead for mechanical pencils, tag, binder, sheets for binder, binding paper
Pharmacy	Analgesic, antipyretic, anti-inflammatory, cold remedy, antitussive, expectorant, syrup, medicinal honey, antiallergic pills, antiallergic cream, antiallergic soap, antacid, anti-reflux, antiemetic, vomiting and nausea pills, laxative, purgative, laxative oil, antirheumatic, corticosteroid, anti-infective, antibiotic, antibacterial, antiviral, herpes, diabetes, antihypertensive, medicine for asthma or bronchitis, nebulization, vasodilator, blood pressure, medicine for cardiovascular insufficiency, cholesterol reducer, triglyceride reducer, bronchodilator, antianaemic, dewormer, anti-helminthic, protozoa, medicine for bones and joints, anti-osteoporosis, calcium, anti-osteopenia, antidepressant, tranquilliser, anxiolytic, sleeping pill, appetite suppressant, appetite controller, slimming product, weight loss medicine, vitamins, contraception, hormone replacement, intradermal patches, phytohormones, medicine for liver, pancreas or bile ducts, medicine for hepatitis, medicine for gallbladder or kidneys, diuretic, nervous system, antiepileptic, anticonvulsant, antipsychotic, medicine for Alzheimer's, medicine for stroke, anti-Parkinsonian, neurological, memory enhancer, medicine for dizziness or vertigo, medicine for circulation, dermatological, antifungal, antimycotic, dermatological cream, body cream, cream for burns, dermatological oil, medicine for lupus, dermatological lotion, petroleum jelly, oil for diaper rash, diaper rash cream, medicine for digestive problems and gastritis, vaginal cream, gynaecological remedy, vaginal ointment, anti-abortion, eye problems, eye drops, ophthalmic solution, antidiarrhoeal, anti-gas, antispasmodic, anticolic, sweet oil, fortifiers, appetite stimulants, food supplement, amino acid, sexual stimulant, varicose veins, haemorrhoids, AIDS treatment, mouth, ear, nose, and throat problems, otologicals, throat lozenge, medicine for sinusitis, medicine for labyrinthitis, throat infections, canker sores, nasal decongestant, medicine for gum infections, medicine for bad breath, mouthwash, medicine for tonsillitis or ear infections, thyroid regulator, hypothyroidism, other
Appliances	Electric stove, dishwasher, vacuum cleaner, washing machine, drying machine, microwave oven, lawnmower
Electronics	Computer, computer accessories, printer, scanner, keyboard, mouse, cd-rom, dvd-rom, cd and dvd recorder, monitor, computer speakers, laptop, internet equipment, USB stick
Bike/Moto	Bicycle or motorcycle
Car	Car
Public Transport	Bus, train, taxi, rental car, metro, bus-metro integration, train-metro integration, funicular, train-bus integration, inter-municipal or inter-state coach, kombi, alternative transportation, vans, rides in someone else's car, moto-taxi
Alcohol/Tobacco	(includes drugs and gambling) Cachaca, fortified wine, beer, caipirinha, draught beer, whisky, vodka, wine, cognac, organic wine, light beer, cigarettes, cigarillos, lighter, cigars, tobacco, match, cigarette holder, liquid and gas for lighter, stone for lighter, rolled tobacco, snuff, paper for tobacco, marijuana, cocaine, lottery tickets, horse betting, cock-fight betting, snooker betting, football betting, domino betting, other types of gambling and lottery

Note: This table describes the spending items from POF 2008/09 who are grouped into each spending category in my analysis in Chapter 4.

Table A.14: Most common education items purchased below poverty line, by income quintile

(1) Quintile	(2) Item	(3) Share of households purchasing item	(4) Av. spending (if > 0, BRL)	(5) Av. spending (if > 0, % of HH income)
1	Notebook	38.8%	17.90	6.79%
	Pencils, pens, erasers	38.0%	8.16	3.09%
	Uniform	9.4%	33.80	12.82%
	Textbooks	2.1%	138.73	52.62%
	Backpack	2.1%	29.43	11.16%
2	Pencils, pens, erasers	35.3%	8.50	1.69%
	Notebook	35.0%	20.89	4.15%
	Uniform	10.6%	32.67	6.49%
	Textbooks	2.5%	172.11	34.18%
	Backpack	2.5%	30.17	5.99%
3	Notebook	33.3%	22.57	3.47%
	Pencils, pens, erasers	33.0%	11.22	1.72%
	Uniform	10.6%	33.80	5.19%
	Computer course	3.5%	162.69	25.00%
	Backpack	3.2%	36.99	5.68%
4	Pencils, pens, erasers	34.7%	10.58	1.36%
	Notebook	32.8%	21.37	2.75%
	Uniform	11.4%	39.89	5.13%
	Computer course	3.5%	178.65	23.00%
	Backpack	2.1%	46.07	5.93%
5	Pencils, pens, erasers	31.6%	12.21	1.38%
	Notebook	30.3%	22.96	2.60%
	Uniform	10.7%	36.29	4.10%
	Computer course	4.0%	193.94	21.93%
	Backpack	3.1%	44.14	4.99%

Note: This table describes the five most common education-related items purchased by households below the poverty line who receive a any Bolsa Família benefit, by income quintile, from the 2008/2009 Household Budget Survey. It is the same sample as the one illustrated in Figure 4.1. Within each income quintile (column 1), the items (column 2) are ranked according to the share of households who report having non-zero spending on that item (column 3). Column 4 illustrates the average amount conditional on it being larger than zero (i.e. for the households who do purchase that item, how much do they spend on average). Column 5 reports the same amount, but as a percentage of average total household income for that quintile. As an example: 9.4% of households from quintile 1 report purchasing a school uniform, and those who do on average spend BRL 33.80, which is equivalent to 12.62% of the average income across all households in quintile 1.

Table A.15: Effect on formal employment by municipality of employment

	(1) Everyone	(2) Men	(3) Women
Panel A: Same as home municipality			
Post · Benefit	0.0075***	0.004	0.006***
	(0.002)	(0.005)	(0.002)
N	310755	99583	211172
Mean Ineligible	0.114	0.138	0.100
Panel A: Different to home municipality			
Post · Benefit	0.004**	0.000	0.005***
	(0.002)	(0.004)	(0.002)
N	310755	99583	211172
Mean Ineligible	0.065	0.098	0.047

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table shows difference-in-differences estimates $\hat{\beta}$ of Equation 3.1.2 where Y_{it} is an indicator for quarterly formal employment by gender, according to municipality of hiring firm. In Panel A, outcome Y_{it} is a binary indicator equal to 1 when individual i is in formal employment in quarter t and the job is located in the same municipality where individual i lives. In Panel A, outcome Y_{it} is a binary indicator equal to 1 when individual i is in formal employment in quarter t and the job is located in a different municipality to where individual i lives. Municipality of residence is fixed at baseline (i.e. as of December/2013). Standard errors are clustered at the individual level.

Table A.16: Frequentist regressions of micro-region level treatment effect

	(1) Estimates - Women	(2) Estimates - Men
Local spending: infrastructure	−0.002	0.001
	(0.004)	(0.009)
Local spending: health	0.007**	0.005
	(0.003)	(0.008)
Local spending: education	0.008***	0.002
	(0.003)	(0.006)
Baseline formal market	0.009***	−0.007
	(0.003)	(0.007)
Size of reform	0.0003	−0.005
	(0.004)	(0.008)
Available labour	0.004	0.007
	(0.003)	(0.007)
Av. wage growth	−0.003	0.013
	(0.004)	(0.008)
Growth formal market	0.002	0.009
	(0.004)	(0.009)
log GDP/capita	−0.012**	−0.011
	(0.006)	(0.012)
N	483	423
R^2	0.038	0.018

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Note: This table reports coefficients of OLS regressions of the form $\hat{\tau}_k = \mathbf{X}_k'\beta + \epsilon_k$, where $\hat{\tau}_k$ is the treatment effect of transfer on employment on micro-region k and $\mathbf{X}_k$ is the set of micro-region characteristics included in my analysis of local correlates of treatment effect in Chapter 5. Each row represents one of the attributes included in $\mathbf{X}_k$. Regressions are weighted by the inverse variance of the regional estimate $\hat{\tau}_k$. Column (1) shows the estimates only for treatment effect on women, and column (2) for men. These estimates are the frequentist counterparts for the posterior coefficient estimates of the Bayesian Meta-Regression illustrated in Figure 5.3.

Table A.17: Measures of pooling and explained variance for different specifications of Bayesian Hierarchical Models

	Metric	Interpretation	Value	
			Women	Men
	Meta-Regression - Bayesian HM model with covariates			
1	R^2 (Gelman and Pardoe, 2006)	Within-model share of variance explained by covariates	50.3%	40.8%
2	λ (Gelman and Pardoe, 2006)	Share of information pooling across micro-regions	96.3%	95.7%
3	R^2 from WLS of estimated effects on posterior	Share of variance in "raw" estimates explained by posterior estimates	14.9%	17.5%
	Standard Bayesian HM - no covariates			
4	Pooling factor (Meager, 2019)	Share of information pooling across micro-regions	91.8%	95.1%
5	R^2 from WLS of estimated effects on posterior	Share of variance in "raw" estimates explained by posterior estimates	36.3%	42.8%

Note: This table shows several metrics of pooling and explained variance for the Bayesian Hierarchical Model at the micro-region level, as specified in Chapter 5 for men and women. Rows 1 to 3 refer to a Bayesian Meta-Regression which includes a pre-specified set of explanatory covariates, as detailed in Equation 5.3.3, and consider both information pooling across regions and conditional distributions governed by regional factors. Rows 4 and 5 refer to a standard Bayesian Hierarchical Model, with no covariates, so only information pooling is considered. Rows 1, 2 and 4 compute standard metrics from the Bayesian literature. Rows 3 and 5 report the resulting R^2 of running a weighted least-squares regression of the "raw" (i.e. pre-Bayesian) estimated region-level treatment effect $\hat{\tau}_k$ on the posterior estimate $\hat{\tau}^{p}ost_k$ obtained from the Bayesian model, where the weights are defined as the inverse variance (i.e. precision) of the "raw" estimate.

Table A.18: Effect of FPM discontinuity on economic outcomes

	(1)	(2)	(3)	(4)
Outcome:	Total formal employment (log)			log GDP/capita
	Everyone	*Women's*	*Low-paying jobs*	
FPM Jump	0.021	0.031	0.031	0.010
	(0.035)	(0.035)	(0.047)	(0.034)
Bias-corrected BW	1629	1350	1650	1280
N (within BW)	11756	11756	11756	11756

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table shows Regression Discontinuity estimates using optimal bandwidth selection and bias-correction as in Calonico et al. (2015) for the impact of the FPM discontinuity on a number of economic and labour market indicators across Brazilian municipalities. Column 1 shows the effect of the discontinuity on total formal employment, column 2 on total formal employment by women, column 3 on total formal employment only considering occupations which, on average, fall within the bottom 20% of the salary distribution, and column 4 on log-GDP per capita. The running variable is normalised population at the nearest FPM cut-off. All specifications consider a year, cut-off and state fixed effects.

Appendix B: Additional Figures

Figure B.1: Log GDP per capita across Brazilian micro-regions, equivalent countries highlighted

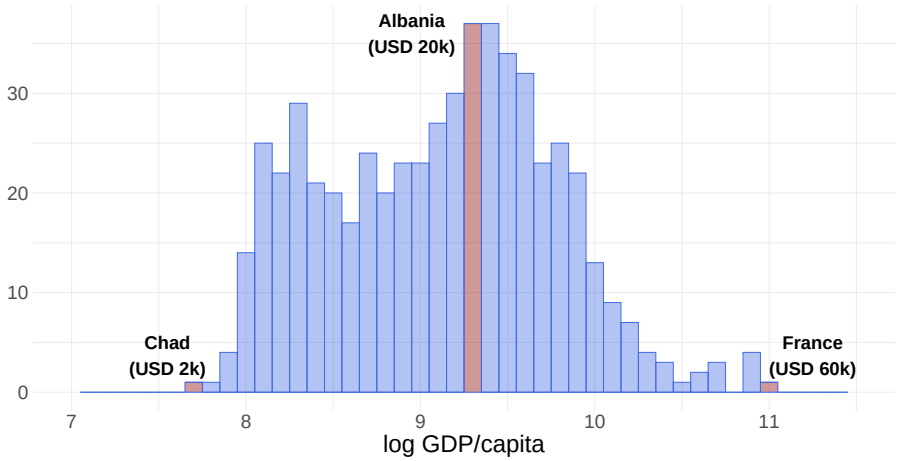

Note: This figure illustrates the distribution of log GDP per capita in 2013 across Brazil's 557 micro-regions, converted to 2022 USD in PPP, with countries of equivalent GDP per capita highlighted. Data for Brazilian micro-regions come from the National Statistics Office (IBGE), and for individual countries comes from the World Bank Development Indicators.

Figure B.2: Share in wage employment vs transfer size, for varying magnitudes of c_{high}

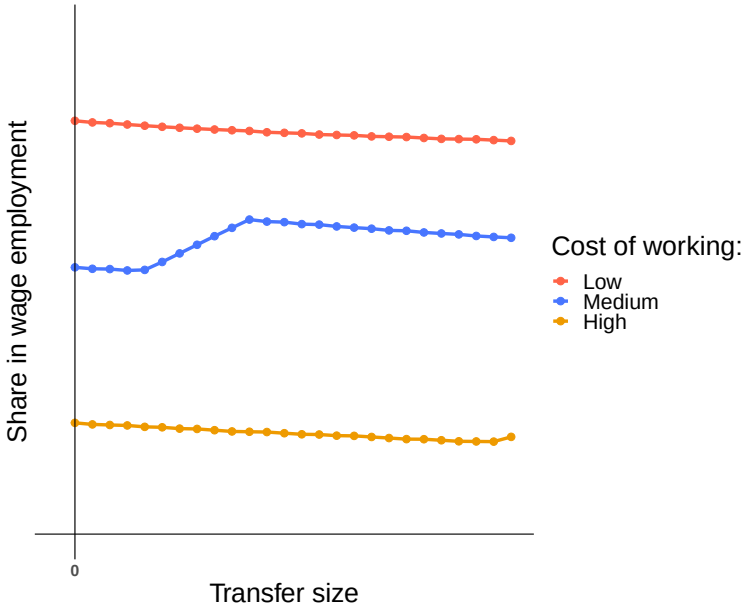

Note: This figure illustrates the equilibrium share of individuals who choose wage employment as a function of transfer size θ for three values of upfront cost c_{high}: high, medium and low. The medium line is identical to the one illustrated in Figure 7.2

Figure B.3: Event study estimates by gender - varying first-stage

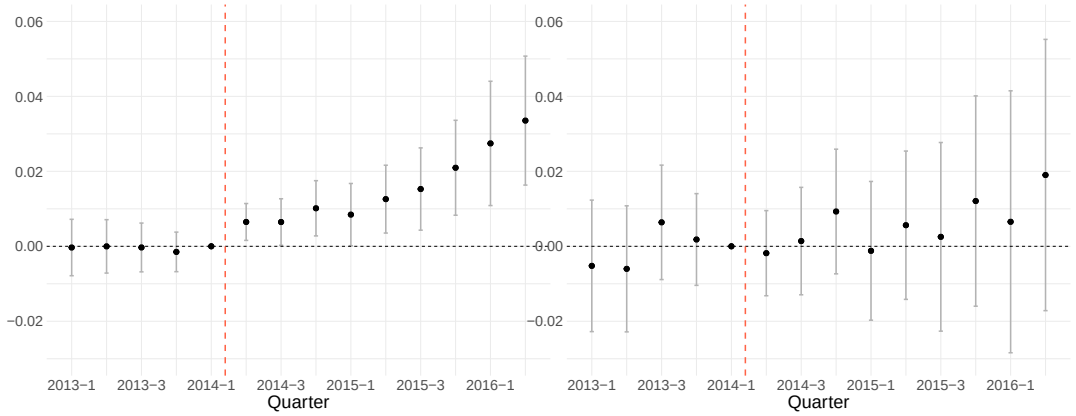

(a) Women **(b)** Men

Note: This figure illustrates the set of event-study estimates $\{\hat{\delta}_s\}$ obtained from estimating Equation 3.1.1 on individual formal employment by gender, under the first alternative specification described in Section 3.4. Treatment status is allowed to vary at every quarter, such that treatment variable is now defined as Benefit_{it} (as opposed to Benefit_i in the main specification) is a dummy variable equal to one if individual i received the benefit at quarter t. The grey bars represent the 95% confidence interval at each quarter. Standard errors are clustered at the individual level.

Figure B.4: Event study estimates by gender - micro-region x quarter fixed effects

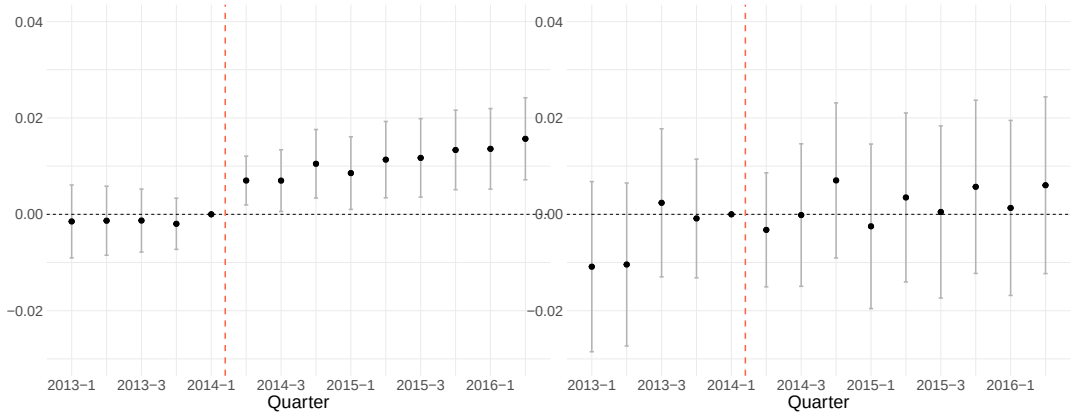

(a) Women **(b)** Men

Note: This figure illustrates the set of event-study estimates $\{\hat{\delta}_s\}$ obtained from estimating Equation 3.1.1 on individual formal employment by gender, under the second alternative specification described in Section 3.4. This specification is identical to the main specification in Equation 3.1.1 with the addition of a micro-region × year fixed-effect to capture potential differential variations in aggregate formal employment across regions. The grey bars represent the 95% confidence interval at each quarter. Standard errors are clustered at the individual level.

Figure B.5: Event study estimates by gender - recursive dynamic specification

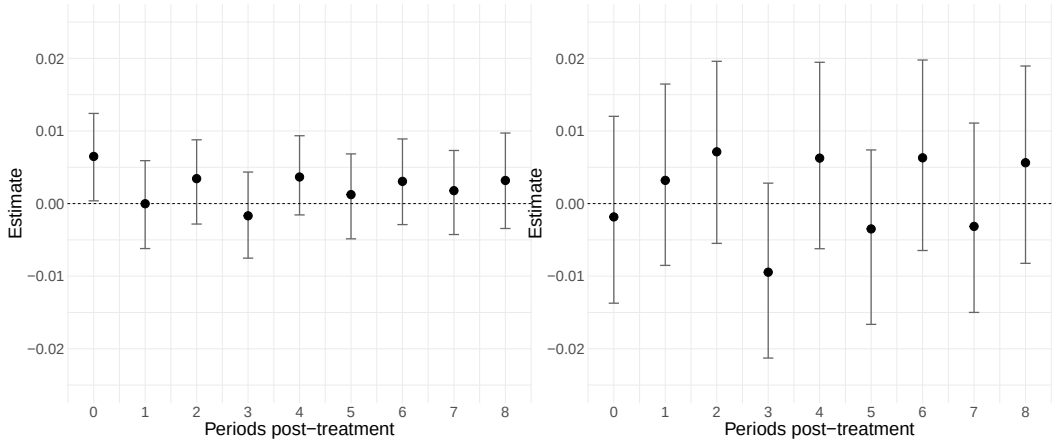

(a) Women **(b)** Men

Note: This figure illustrates the set of event-study estimates $\{\hat{\delta}_s\}$ obtained from estimating Equation 3.1.1 on individual formal employment by gender, under the third alternative specification described in Section 3.4. This specification defines the set of coefficients $\{\delta_s\}$ at period s as the impact on formal employment s periods after benefit receipt, as opposed to in calendar quarter s. I followed the methodology outlines in Giupponi and Landais (2023). Let the reduced form ITT estimate of treatment assignment on employment at period s be the set of event-study coefficients of the regression of formal employment on the interaction between benefit eligibility (fixed by December/2013 income) and quarter dummies: $Y_{it} = \theta_i + \alpha_t + \sum_{s \neq -1} ITT_s \cdot I(t = s) \cdot Z_i + \eta_{it}$. For each period s, the treatment-on-the-treated coefficients can be written as a linear function of the ITT in period s and the first-stage $\frac{d\text{Benefit}_{i,s}}{dT_i}$: $ITT_s = \sum_s \delta_s \cdot \frac{d\text{Benefit}_{i,s}}{dT_i}$. The set of coefficients $\{\delta_s\}$ can be obtained by matrix inversion. Standard errors are bootstrapped over 500 iterations. Grey bars represent the 90% confidence intervals.

Figure B.6: Event study estimates at the household level

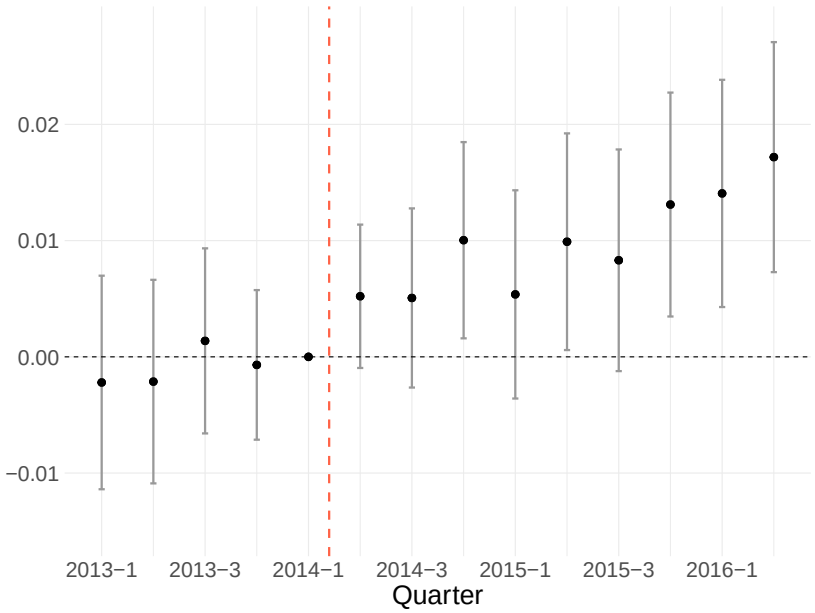

Note: This figure illustrates the set of event-study estimates $\{\hat{\delta}_s\}$ obtained from estimating Equation 3.1.1 on formal employment for the full sample at the household level. The outcome Y_{it} is equal to 1 if any adult from household i is in formal employment at quarter t, and to 0 otherwise. The grey bars represent the 95% confidence interval at each quarter. Standard errors are clustered at the individual level.

Figure B.7: Evidence on pre-reform income manipulation

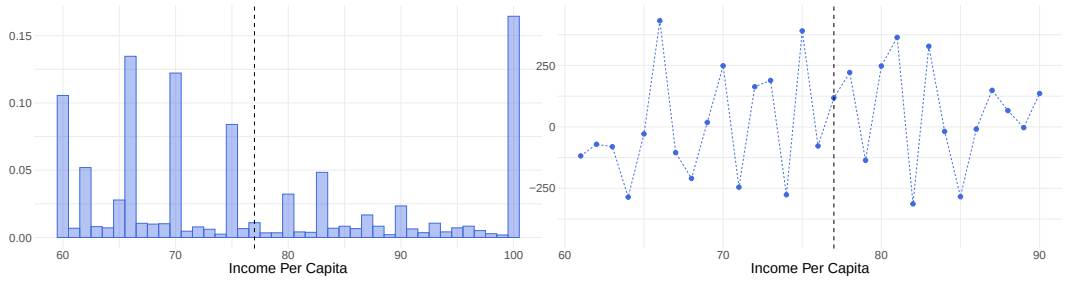

(a) Distribution of Income per Capita **(b)** T-Statistic for discontinuity

Note: This figure illustrates pieces of empirical evidence against the hypotheses of *ex-ante* income manipulation in December 2013 around the future benefit threshold of BRL 77/month. Panel (A) illustrates the distribution of household income per-capita around BRL 77 for all households present in December 2013 Cadastro Único. There is no specific bunching at BRL 77 (represented by the dotted line), which be a reason for concern. Panel (B) reproduces the method employed in Gerard et al. (2024), and shows the McCrary's test statistic for income per-capita discontinuity at every between BRL 60 and BRL 90. The test statistic is large and significant (in absolute value), hence rejecting the hypothesis of smoothness of the distribution at all integers. The value at BRL 77 (highlighted with the dotted line) is not larger or smaller than other values in the vicinity, which suggests no specific bunching at BRL 77.

Figure B.8: Reported household income per capita of ineligible group in future years

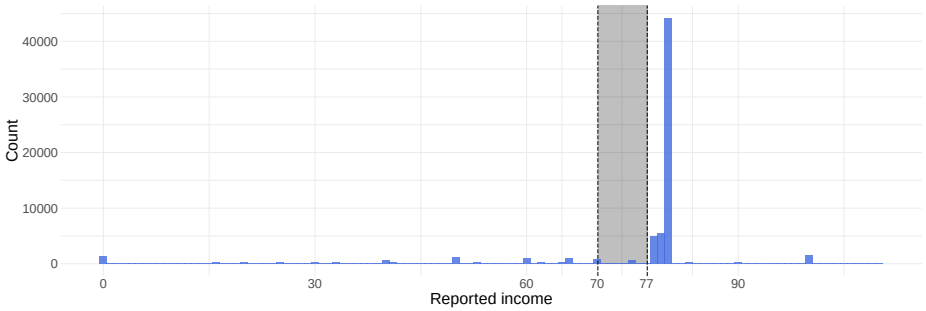

(a) 2014

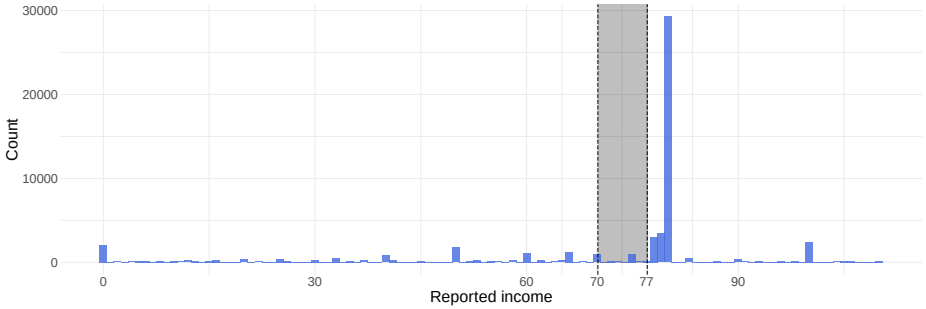

(b) 2015

Note: This figure illustrates the distribution of self-reported household income per-capita of the control group in my main sample in the two subsequent waves of Cadastro Único: 2014 (panel A) and 2015 (panel B). The control group is defined as all individuals who live in a household that, in December 2013, had a self-reported income per capita between BRL 77 (exclusive) and BRL 80 (inclusive). The area shaded in grey represents the new range of income for benefit eligibility post-reform, but not pre-reform: an income between BRL 70 (exclusive) and BRL 77 (inclusive). If the estimates in Chapter 3 were driven by the control group having less incentives to work due to the possibility of bunching in the new area of eligibility, the grey area would contain an excess mass in the years following the reform. The lack of a clear excess mass in this area suggests that the results are not driven by a behavioural response by the control group that chooses to be eligible for the benefit by reducing earnings.

Figure B.9: Evolution of transfer receipt: control group vs placebo infra-marginal group

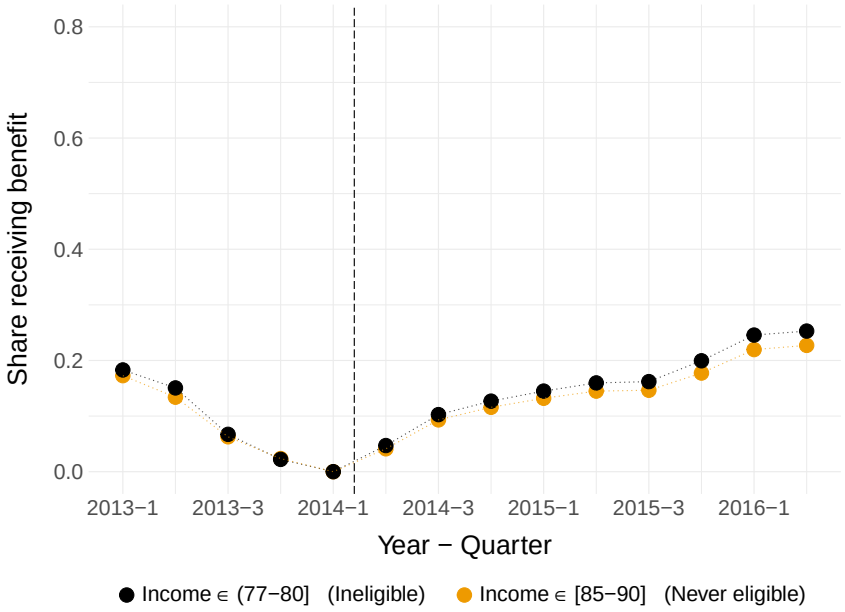

Note: This figure illustrates the evolution of the share of individuals receiving Basic Benefit at each quarter around the time of the reform for the ineligible group used as a control in my main specification (i.e. those in households with income between BRL 77 and BRL 80) and a placebo control group that is both unaffected by the reform, and not at the margin of the new threshold: those in households with income in December 2013 between BRL 85 and BRL 90. It serves as a placebo comparison to Figure 3.1, and it shows that the behavioural of the ineligible group from my main sample is not substantially different from other infra-marginal groups. This similarity helps rule out the hypothesis that the main findings are driven by changes in behaviour of the ineligible group.

Figure B.10: Impact of the 2014 reform on size of formal market

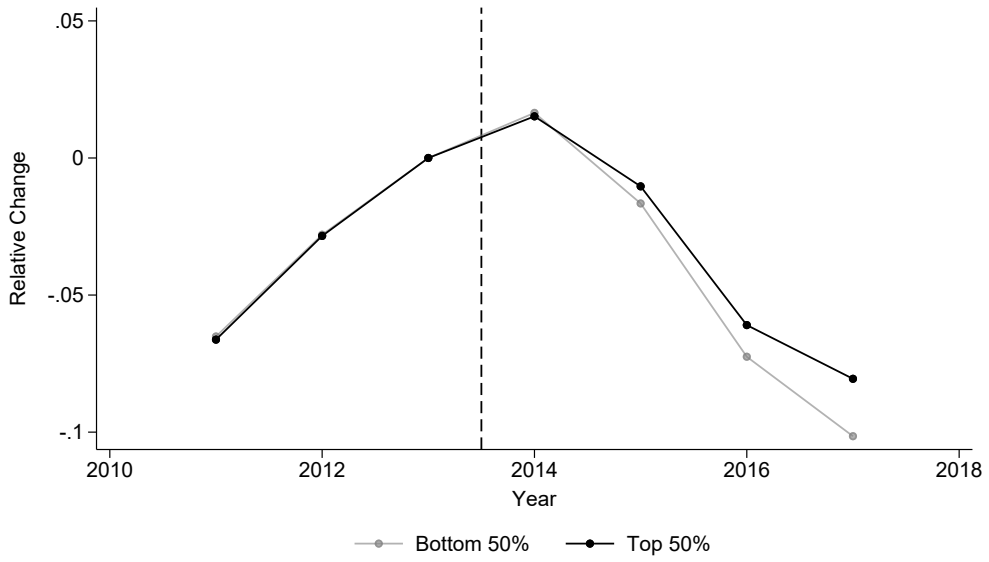

Note: This figure replicates the empirical strategy of Gerard et al. (2024) to quantify the impact of the 2014 Bolsa Família reform on total formal employment at the micro-region level. First, I calculate the change in total monthly Bolsa Família budget generated by the 2014 reform for all Brazilian micro-regions as the difference between the total budget in October 2014 (post-reform) and April 2014 (pre-reform). I then group micro-regions into whether they are above or below the median change in total monthly Bolsa Família spending. This figure illustrates the evolution of total number of private-sector formal employment contracts for both groups of micro-regions around the time of the reform, relative the total number of formal jobs in 2013 (normalised to zero). It aims to show that the net-effect of the 2014 reform in total formal employment is an increase in total jobs, so my findings on the increase in employment are not purely driven by displacement of non-recipients.

Figure B.11: Evolution of enrolment in education for children of different age groups

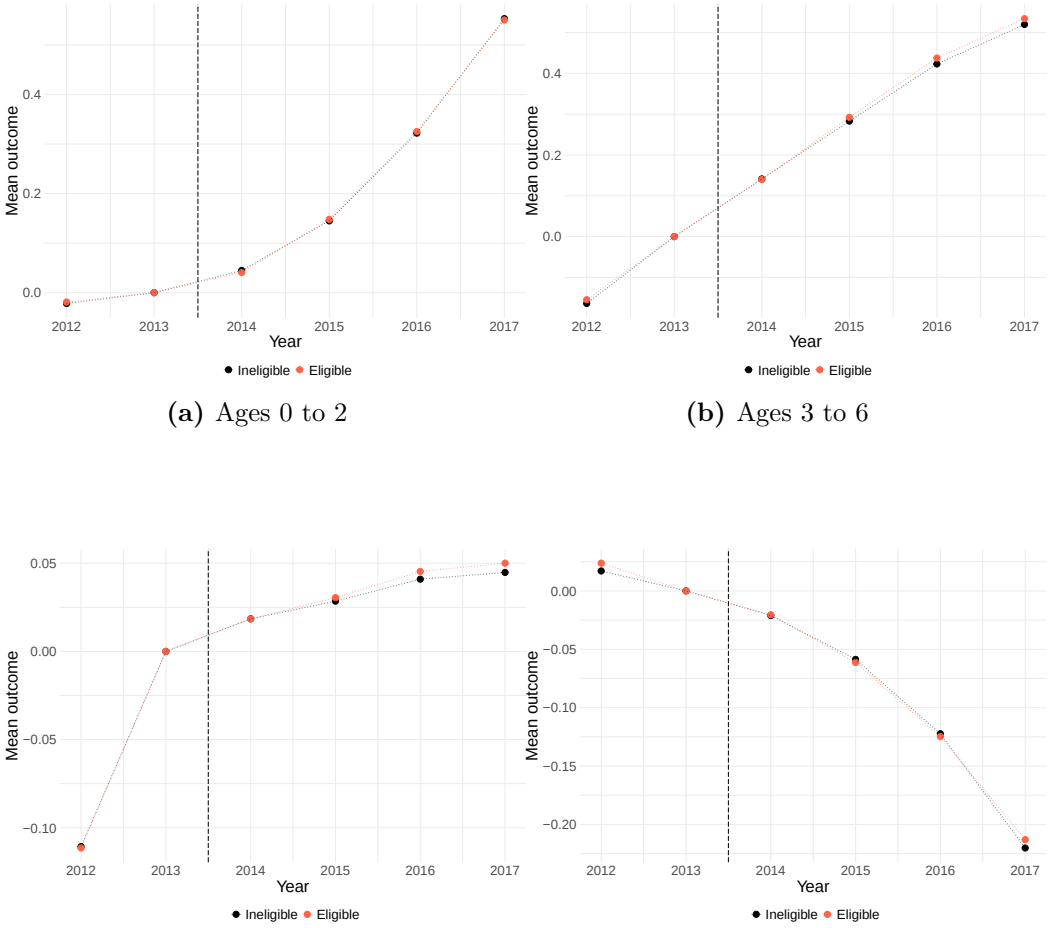

(a) Ages 0 to 2

(b) Ages 3 to 6

(c) Ages 7 to 10

(d) Ages 11 to 18

Note: This figure illustrates the evolution of self-reported enrolment rate in education for all children who, in 2013, were youbger than 18. Children are group by ages in which they are normally enrolled in each schooling stage: nursery (0-2), pre-school (3-6), primary school (7-10), middle/high school (11-18). Enrolment in education is measured as enrolment in any stage of formal education at each yearly snapshot of Cadastro Único. Enrolment rate is computed separately for the eligible and the ineligible group, and normalised such that enrolment rate in 2013 - just before the reform - is set to zero. Enrolment rates are calculated conditional on children remaining in Cadastro Único in each year.

Figure B.12: Share of women working part-time, by income decile and type of employment

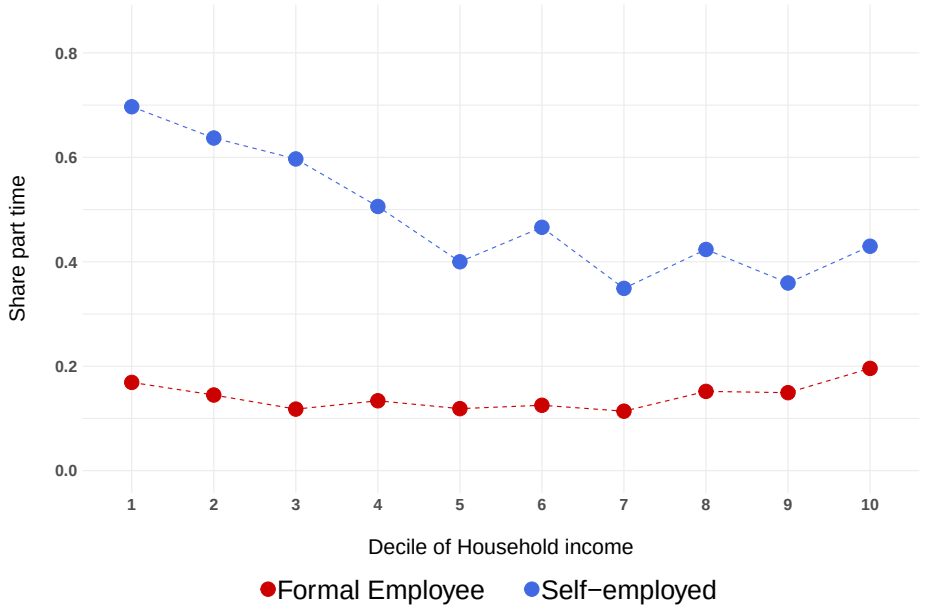

Note: This figure illustrates the share of women in part time employment as a function of deciles of household income per capita, by type of employment (formal employee vs self-employed). Data comes from the National Household Survey between 2014-2016, and the sample includes all women between the ages of 18 and 45 who are either in formal employment or in self-employment. Household income is calculated excluding any social benefits. Part time employment is defined as working less than 40 hours a week.

Figure B.13: First stage estimates of propensity score

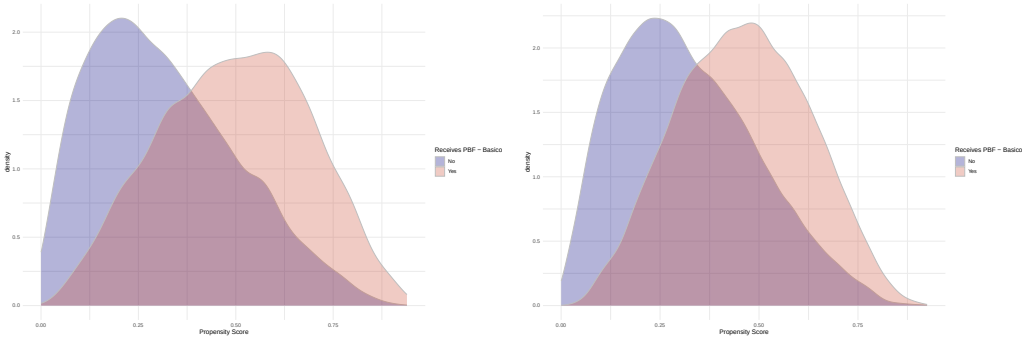

(a) POF 2008 **(b)** PNAD 2014-16

Note: This figure illustrates the first stage of the Propensity Score Matching (PSM) algorithm employment on the analysis of mechanisms from survey data in Chapter 4. Panel A illustrates the results for POF 2008/09 (Household Budget Survey), and Panel B illustrates the result for PNAD-C 2014-16 (National Household Survey). I fit a probit regression of a dummy for Basic Benefic receipt on a rich set of observables, described in Appendix Table A.11. I then compute the density of predicted receipt probability, and compare with observed receipt status. Each panel shows overlapping densities of predicted probability of receipt for those who actually receied (red) and those who did not receive (blue). The red density lies mostly to the right of the blue, which indicates good prediction accuracy.

Figure B.14: Posterior estimate of state-level treatment effect for men vs share of men unavailable to work due to caring duties

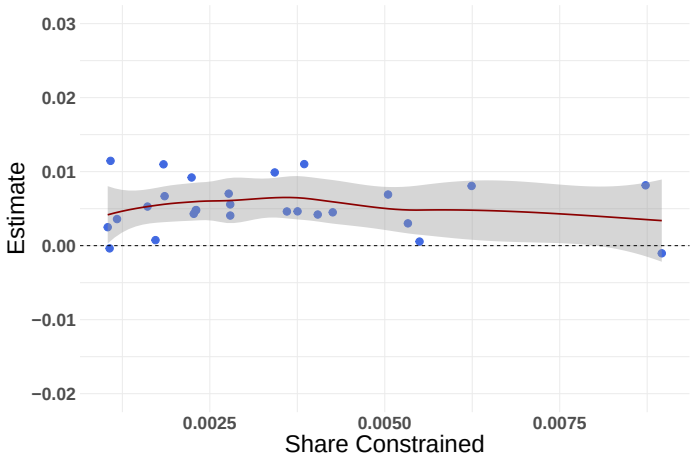

Note: This figure illustrates the relationship between state-level estimates and share of men constrained due to care work. State-level estimates are calculated by estimating Equation 3.1.2 separately for each state, and then pooling estimates from all states via a Bayesian Hierarchical Model. The values on the y-axis are the posterior means for each state. The share of men constrained due to care work is constructed from PNAD-C 2014-16, by computing the proportion of women age 18-45 who: (j) are not employment, (ii) answer no to the question "If you were offered a job today, would you be available to take it?", and (iii) report the reason for unavailability as being "Caring duties for another household member". This figure is analogous to Figure 5.5 in the main body of the paper, which shows the same comparison for women.

Figure B.15: Observed and law-implied FPM transfers per capita around nearest population cutoff

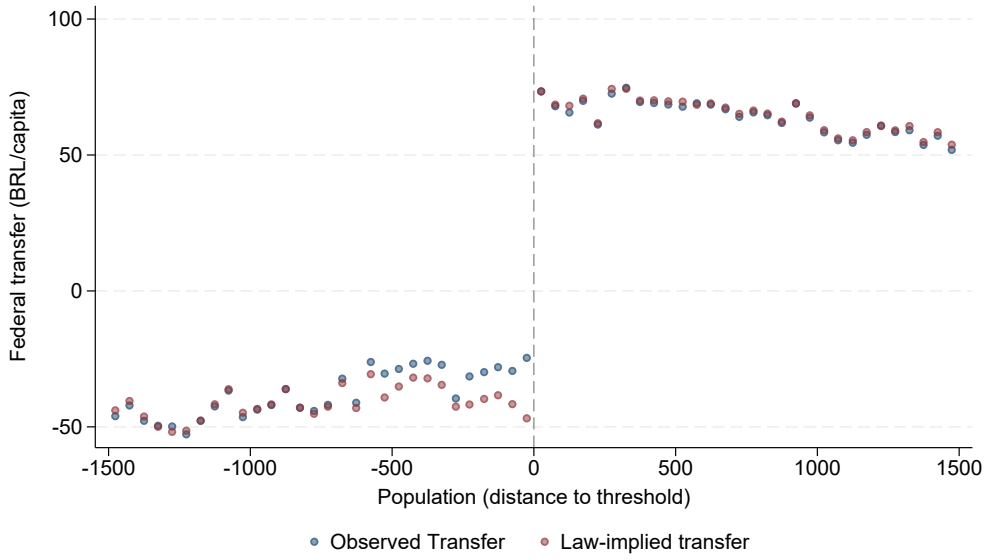

Note: This figure illustrates the comparison between the the amount of yearly FPM funds per capita that municipalities closed to FPM cutoffs are entitle to receive (red dots) and dots reported in data from Brazilian inter-governmental sources. The figure pools all municipalities within 1500 inhabitants of all FPM cutoffs. It illustrates the "jump" in assigned transfers to municipalities around the population cutoffs.

Figure B.16: RD estimates of the impact of FPM "jump" on municipal spending on local public goods

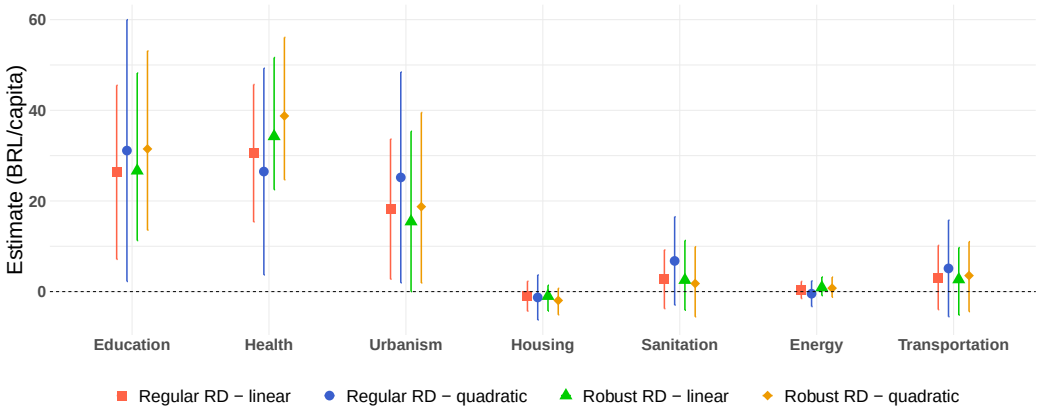

Note: This figure illustrates coefficient estimates of the effect of the FPM jump on yearly municipal spending a number of categories related to the provision of public goods. Reported coefficients are estimated from Regression Discontinuity Design, specified as $Y_{it} = \alpha_i + \gamma_t + f(Pop_{it}) + \beta \cdot I(Pop_{it} > 0) + \epsilon_i t$ where Y_{it} is spending in each category for municipality i in year t, and Pop_{it} is population centered around the nearest FPM cutoff. The dots represents point estimates, and the bars are 95% confidence intervales. Red estimates report $hat\beta$ coefficients from a linear RDD (i.e. $f(Pop_{it}) = a \cdot Pop_{it}$) with a bandwidth fo 1500 people. Blue estimates report estimates from a quadratic RD (i.e. $f(Pop_{it}) = a \cdot Pop_{it} + b \cdot Pop_{it}^2$ with a bandwidth of 1500 people. Green and yellow estimates come, respectivelly, from linear and quadratic specifications using optimal bandwidth selection algorithm following Calonico et al. (2015).

Figure B.17: Dispersion of estimates across different countries from the literature

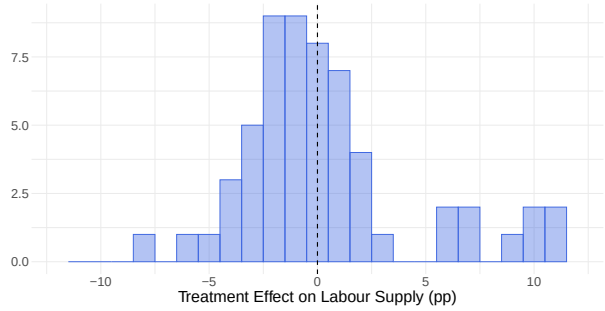

Note: Panel A illustrates the distribution of estimates of the effect of unearned income on extensive margin labour supply obtained from different studies in the literature. Data comes from a combination of four meta-analyses on cash transfers (Bastagli et al., 2016; Banerjee et al., 2017; Diaz-Pardo and Rao, 2024; Crosta et al., 2024), combined with 17 individual studies of the impacts of shocks to unearned income on extensive margin labour supply.

Figure B.18: Correlation between local education supply and spending in educaton

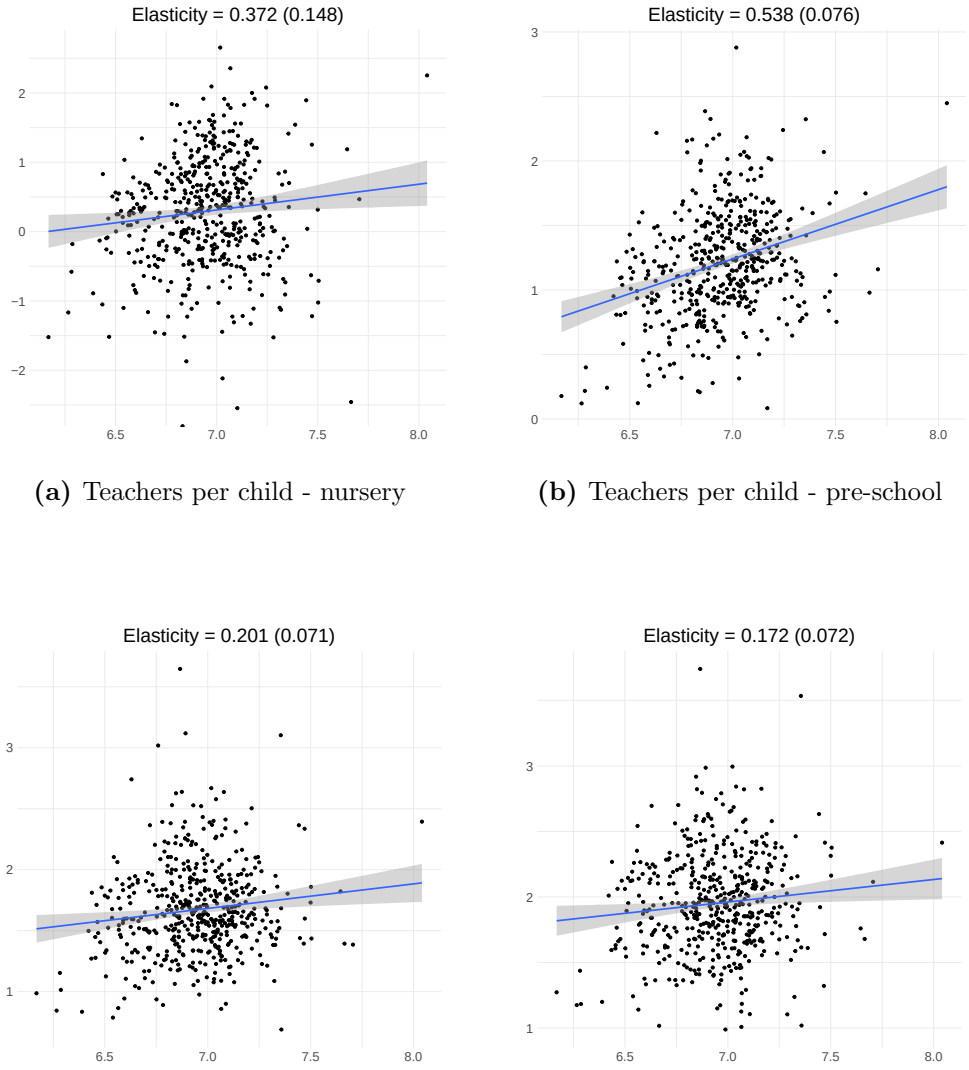

(a) Teachers per child - nursery

(b) Teachers per child - pre-school

(c) Teachers per child - primary school

(d) Teachers per child - middle school

Note: This figure illustrates the relationship between measures of local supply of education and yearly local spending in education, at the micro-region level in 2013. In each panel, I show the relationship between the logarithm of an indicator and the logarithm of the total municipal spending per capita in education, aggregated at the micro-region level. The education indicators are measured as number of teachers in the public education system per 100 resident children, for different levels of schooling: nursery (panel A), pre-school (panel B), primary school (panel C), and middle school (panel D). Data on education supply comes from the School Census. Data on resident children for different age groups is obtained from the National Census of 2010 and 2022, linearly interpolating to obtain estimates for 2013. Data on public spending in education comes from the FINBRA database on municipal public finances.

Figure B.19: Correlation between local health supply and spending in health

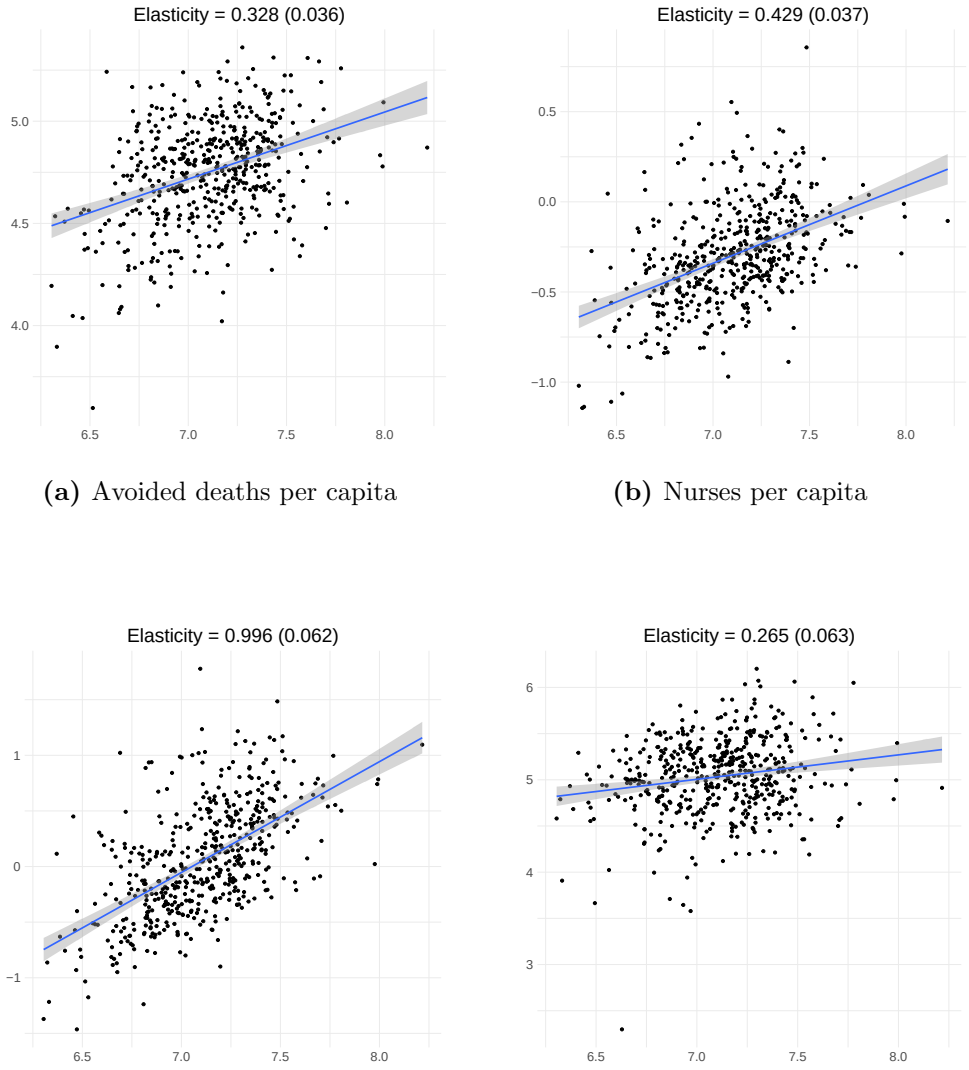

(a) Avoided deaths per capita

(b) Nurses per capita

(c) Doctors per capita

(d) Hospital beds per capita

Note: This figure illustrates the relationship between measures of local supply of public health and yearly local spending in health, at the micro-region level in 2013. In each panel, I show the relationship between the logarithm of an indicator and the logarithm of the total municipal spending per capita in health, aggregated at the micro-region level. The education indicators are measured as: estimate of number of avoided deaths (panel A), number of nurses in the public health systen (panel B), number of doctors in the public health system (panel C) and number of hospital beds in the public health system (panel D). All indicators are calcualted per 100,000 inhabitants. Data on health indicators comes from the Health Policy Research Institute (IEPS). Data on public spending in health comes from the FINBRA database on municipal public finances.

Figure B.20: Correlation between share of low income households with access to basic infrastructure services and spending in infrastructure

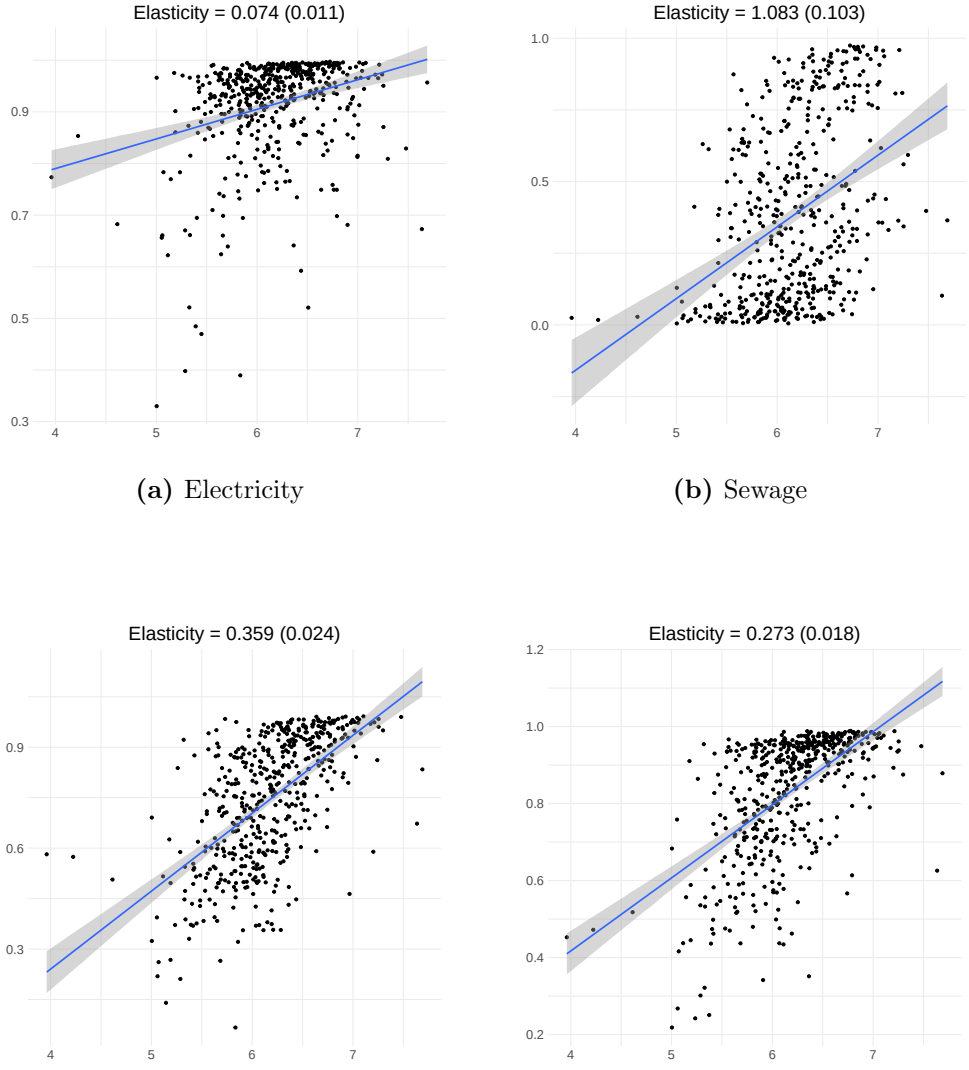

(a) Electricity

(b) Sewage

(c) Trash collection

(d) Running water

Note: This figure illustrates the relationship between measures of local access to infrastructure and yearly local spending in infrastructure, at the micro-region level in 2013. In each panel, I show the relationship between an indicator and the logarithm of the total municipal spending per capita in infrastructure, aggregated at the micro-region level. The education indicators are measured as the share of households registered in Cadastro Único 2013 with access to a basic category of local service: grid electricity (panel A), sewage network (panel B), regular trash collection (panel C) and running water (panel D). Data on infrastructure comes from the Cadastro Único. Data on public spending in infrastructure comes from the FINBRA database on municipal public finances, and aggregates the following categories of spending: urbanism, housing, sanitation, communications, and energy.

Appendix C: Summary of Qualitative Work

As part of this project, I conducted the following qualitative work activities in the summer of 2023 and 2024:

- 5 interviews with workers in the local Social Assistance Centre (CRAS) in the municipalities of Água Branca (AL), Tarrafas (CE), Lavras da Mangabeira (CE), Itaúba (MT) and Sinop (MT)
- 8 home visits in the municipalities of Água Branca (AL) and Tarrafas (CE)
- 4 focus group discussions with beneficiaries in the municipalities of Lavras da Mangabeira (CE), Itaúba (MT) and Rio de Janeiro (RJ)
- 2 phone interviews with principals of municipal public schools in the municipalities of Manaus (AM) and Itaperuna (RJ)

Next, I briefly summarise the main points of relevance to this analysis raised in each of these activities.

Interviews with social assistance workers

Interviewed social assistance workers in all locations spontaneously mentioned the "Permanence Rule" as a reason for continuous receipt of Bolsa Família by families with income above the threshold. Beneficiaries frequently self-report increases in income that would normally disqualify them, yet they are able to retain benefits.

Although income misreporting is a concern, it generally does not involve precise income bunching; instead, social workers report two common forms of misreporting: omission of employment altogether, and omission of household members as to decrease registered family size. The government has been particularly concerned about the latter, and has

increased enforcement on households misreport being one-person households in order to receive extra benefit payments.

Additionally, social assistants highlight a lack of discretion in benefit allocation at the local level. Information entered into an electronic system produces eligibility decisions through a process described by workers as a "black box," with limited transparency. The process for manually adjusting eligibility and benefit receipt involves going up many levels in the administrative hierarchy of the programme and, in general, has not been thought of as a concern.

In towns visited in the Northeastern states of CE and AL, data from regional centres show that Bolsa Família reaches over 60% of the population, reflecting the program's prominence in these areas. In those municipalities, the Social Assistance Centres play an important role in the local community life, serving as a point of contact for both beneficiary families as well as low income non-beneficiary families, and organising a number of activities related to local culture, health, and education. In these smaller towns, employment in either local public-sector roles as well as public-sector adjacent positions is often mentioned as of the main reason individuals leave Bolsa Família. On the other hand, in Itaúba (MT) and Sinop (MT), in the affluent Center-West, the beneficiary rate is below 10%, and the role played by local Social Assistant Centres in community life is somewhat limited.

Individual home visits

I visited seven households and conducted a video interview with one in the urban, peri-urban and rural districts of the municipalities of Água Branca (AL) and Lavras da Mangabeira (CE). In seven out of the eight households, the head was a woman with either children or grandchildren of school age or younger, whereas one household had a single male member.

Subjects report differing degrees of economic hardship. Bolsa Família makes up very large share of household income, often 100%. The main other reported sources of income were:

- Selling of own fruits and vegetables in local market
- Day wages from casual agricultural labour by men in rural and peri-urban areas,
- Informal domestic work in other houses

Several women report not being able to afford any childcare. Most children of schooling age are enrolled in local public school, but are often unable to attend classes for different

reasons, for example, due to lack of transport or material. Children younger than school age are often at home. There have been no reports of enrolment in nursery, and some of enrolment in pre-school. Untreated health conditions of beneficiaries themselves, or of their children, have been also reported as a barrier to work.

All beneficiaries report buying groceries as the main expenditure using Bolsa Família benefit. Other relevant expenditures were: supplies for children, medication, rent and bills, and everyday household items. One beneficiary reported saving a significant share of the benefit and used it to improve the house.

Focus groups

I conducted 4 focus group discussions, two organised via the local social assistance authorities, and two via a partner NGO in Rio de Janeiro. Without pre-specifying a gender breakdown, only 3 out of 40 participants of my focus groups were men, all the others were women. Almost all women who participated in the discussions have either children or grand-children of school or pre-school age.

The majority of discussants have no occupation, some report working on an irregular basis as:

- Domestic workers
- Manicure/hairdresser
- Street cleaner for the local authority
- Salesperson in the streets or local markets

A share of beneficiaries reported having had a formal job in the past, more so in Rio de Janeiro than in the other municipalities. The diversity of occupations was much higher in Rio de Janeiro (RJ) than in the other locations.

Children of nursery and pre-school age were often brought to the focus group discussion, as participants reported being unable to afford any form of childcare. Many mentioned health conditions, of themselves or of their children, as additional hardships and barriers to either working or studying.

All beneficiaries report buying groceries as the main expenditure using Bolsa Família benefit. Other relevant expenditures were: supplies for children, medication, rent and bills, transport.

In Lavras da Mangabeira (CE), a few participants reported often using small savings to improve their house, one participant reported being able to save a substantial part of benefit for a few years, and using it open a waxing salon together with other associates. In Itauba (MT) and Rio de Janeiro (RJ), most participants reported being unable to save any share of the benefit. In Rio de Janeiro, many participants reported using the money to buy goods for re-sale in the streets or local markets.

Interviews with school principals

Principals in both Manaus (AM) and Itaperuna (RJ) mentioned a serious problem of evasion for older children, towards the transition between middle school and high school. Both mention school material as a relevant expense for children's education. In Manaus (AM): books are provided but stationary and uniform must be purchased privately, transport is not typically an issue as it is free for those who live far, but the majority lives within walking distance. In Itaperuna (RJ), uniform is in theory provided for free but the supply is often insufficient, while school material is given for free but usually does not last the whole year. Both mention recent improvements in supplies of education goods, and report having had more frequent problems in the past. In Manaus (AM), the principal reported having had to gift children's uniforms and supplies out of their own pocket in more than one occasion.