



UNIVERSIDADE FEDERAL DE PERNAMBUCO
CENTRO DE CIÊNCIAS SOCIAIS APLICADAS
PROGRAMA DE PÓS-GRADUAÇÃO EM ECONOMIA

Alexandre de Andrade Fonseca

**ESSAYS IN INTERGENERATIONAL
MOBILITY AND PUBLIC ECONOMICS**

RECIFE
2024

Alexandre de Andrade Fonseca

ESSAYS IN INTERGENERATIONAL MOBILITY AND PUBLIC ECONOMICS

*Tese apresentada ao PROGRAMA DE PÓS-GRADUAÇÃO
EM ECONOMIA do DEPARTAMENTO DE ECONOMIA
da UNIVERSIDADE FEDERAL DE PERNAMBUCO como
requisito parcial para obtenção do grau de Doutor em
Economia.*

Orientador: *Prof. Dr. Breno Ramos Sampaio*
Co-orientador: *Prof. Dr. Diogo Gerhard Castro de Britto*

RECIFE
2024

.Catalogação de Publicação na Fonte. UFPE - Biblioteca Central

Fonseca, Alexandre de Andrade.

Essays in Intergenerational Mobility and Public Economics /
Alexandre de Andrade Fonseca. - Recife, 2024.
135f.: il.

Tese (Doutorado) - Universidade Federal de Pernambuco, Centro
de Ciências Sociais Aplicadas, Programa de Pós-Graduação em
Economia, 2024.

Orientação: Breno Ramos Sampaio.

Coorientação: Diogo Gerhard Castro de Britto.

Inclui referências e apêndices.

1. Mobilidade intergeracional; 2. Transferência condicional
de renda; 3. Taxação propriedade; 4. Economia aplicada. I.
Sampaio, Breno Ramos. II. Britto, Diogo Gerhard Castro de. III.
Título.

UFPE-Biblioteca Central

**UNIVERSIDADE FEDERAL DE PERNAMBUCO
CENTRO DE CIÊNCIAS SOCIAIS APLICADAS
DEPARTAMENTO DE ECONOMIA
PIMES/PROGRAMA DE PÓS-GRADUAÇÃO EM ECONOMIA**

ESSAYS IN INTERGENERATIONAL MOBILITY AND PUBLIC ECONOMICS

PARECER DA COMISSÃO EXAMINADORA DE DEFESA DE TESE DO
DOUTORADO EM ECONOMIA DE:

ALEXANDRE DE ANDRADE FONSECA

A Comissão Examinadora composta pelos professores abaixo, sob a presidência do primeiro, considera o Candidato Alexandre de Andrade Fonseca **APROVADO**.

Recife, 08/11/2024

Prof. Dr. Breno Ramos Sampaio
Orientador

Prof. Dr. Giuseppe Trevisan Cruz
Examinador interno

Prof. Dr. Paulo Henrique Pereira de Meneses Vaz
Examinador interno

Prof. Dr. Bladimir Carrillo Bermúdez
Examinador externo/EESP-FGV

Prof. Dr. Diogo Gerhard Castro de Britto
Examinador externo/University of Milan-Bicocca

À minha esposa, Tatiana e aos meus filhos, Mariana e Felipe.

Acknowledgements

Essa tese é o resultado de um trabalho conjunto. Literalmente, dezenas de pessoas contribuíram, de alguma forma, para a sua conclusão. Citarei poucos nomes, para cometer menos injustiças, mas a todas elas dedico essas linhas.

Em primeiro lugar, à minha esposa, amiga e companheira, Tati, que sempre não só me incentivou a buscar novos desafios como, paciente, compreendeu as renúncias e assumiu tarefas e obrigações. Obrigado por ser tanto! Aos meus filhos, Mari e Lipe por entenderem as noites e os finais de semana em que tive de estar ausente e dedicado ao doutorado. À minha família estendida, em especial à minha sogra, Socorro e à minha cunhada Manuella que deram apoio à minha esposa e aos meus filhos e me permitiram um tranquilo e produtivo período de pesquisa no exterior.

Aos muitos amigos da Receita Federal, em especial Assis, que acreditaram e apoiaram esse projeto de doutorado e as pesquisas propostas.

Aos colegas, professores, coordenadores e administrativos do PIMES pelo apoio durante todo o doutorado. Aos muitos amigos que fiz no PIMES, em especial, àqueles que não conseguiram completar o programa, em meio a questões pessoais e a uma inacreditável pandemia global, mas que foram fundamentais em meio a listas de exercícios, provas, monitorias e pesquisas.

Ao meu orientador e coorientador, Breno e Diogo, que se tornaram grandes amigos e com os quais aprendi muito mais do que economia. Aos meus coautores, em especial Lucas pelas inúmeras horas de discussões e trocas de ideias.

Aos amigos que fiz no CLEAN/Bocconi, em especial Paolo e Zach, não só pela oportunidade e suporte mas sobretudo pela calorosa e gentil recepção.

À Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES), pelo financiamento do período de pós-doutorado no exterior.

Por fim, a todos aqueles que, de alguma forma deram suporte para mim e para minha família nessa longa jornada de doutorado.

O autor

"Este paper é leitura obrigatória, um triste retrato do que o Brasil construiu em 200 anos de independência. Não surpreende, mas consegue chocar"

— FILIPE CAMPANTE (Bloomberg Distinguished Professor na Johns Hopkins University - comentário não solicitado, feito em rede social, sobre o primeiro capítulo desse tese)

Abstract

This thesis addresses issues of intergenerational mobility and public economy, studying intergenerational mobility in Brazil, the effects of conditional cash transfer policies and rural property taxation. In the first chapter, I provide the first estimates of intergenerational income mobility using population wide tax data for a large developing country, namely Brazil. I measure formal income from tax and payroll data, and train machine learning models on census and survey data to predict informal income. I quantify the estimation bias resulting from income imputation and other sources of measurement error, and show that such bias remains negligible in our context. A 10 percentile increase in parental income rank is associated on average with a 5.5 percentile increase in child income rank, with considerable variation across sociodemographic groups and geographical areas. The second essay examines the intergenerational effects of the world's largest conditional cash transfer (CCT) program, Brazil's Bolsa Família Program (PBF). Employing a differences-in-differences design and a comprehensive dataset covering cohorts born between 1970–1994. I reveal that PBF significantly promotes human capital accumulation, leading to reduced dependency on the social safety net, higher earnings, and intergenerational income mobility among the next generation. These effects are more pronounced for younger children and females, highlighting the importance of early exposure to the program and its role as a driver of greater equality. My findings underscore the effectiveness of CCTs in breaking the intergenerational cycle of poverty and fostering social mobility. The third chapter uses a difference-in-differences research design to evaluate a program that partially decentralized the administration of rural land taxes to local authorities in exchange of increases in their share of tax revenues. Using microdata from tax returns, I find that the program led to an expansion of tax revenues by 20% after five years. Decentralization expanded tax revenues mainly by increasing reported land values. Using satellite data, I find that partial decentralization did not influence farmer behavior significantly. A cost-benefit exercise indicates that partial decentralization had large returns. Overall, the findings indicate that cooperation between local and central authorities can increase property taxation in contexts with incomplete information and weak enforcement capacity.

Keywords: Intergenerational Mobility, Conditional Cash Transfers, Property Taxation, Applied Economics.

Resumo

Esta tese aborda questões de mobilidade intergeracional e economia pública, analisando a mobilidade intergeracional no Brasil, os efeitos de políticas de transferência condicional de renda e a tributação de propriedades rurais. No primeiro capítulo, apresento as primeiras estimativas de mobilidade intergeracional de renda utilizando dados fiscais de abrangência nacional para um grande país em desenvolvimento. A renda formal é medida a partir de dados fiscais e administrativos e utilizo modelos de aprendizado de máquina, treinados em dados censitários e de pesquisas para prever a renda informal. Quantifico o viés de estimação resultante da imputação de renda e outras fontes de erro de medição, demonstrando que tal viés permanece insignificante em nosso contexto. Um aumento de 10 percentis no ranking de renda dos pais está associado, em média, a um aumento de 5,5 percentis no ranking de renda dos filhos, com considerável variação entre grupos sociodemográficos e áreas geográficas. O segundo ensaio examina os efeitos intergeracionais do maior programa de transferência condicional de renda (CCT) do mundo, o Programa Bolsa Família (PBF) do Brasil. Empregando um modelo de diferenças em diferenças e um conjunto de dados abrangente, cobrindo coortes nascidas entre 1970 e 1994, revelo que o PBF promove significativamente a acumulação de capital humano, levando à redução na dependência da rede de proteção social, a maiores rendimentos e à mobilidade intergeracional de renda na geração seguinte. Esses efeitos são mais pronunciados para crianças mais jovens e mulheres, destacando a importância da exposição precoce ao programa e seu papel como motor de maior igualdade. Meus achados ressaltam a eficácia das CCTs em romper o ciclo intergeracional da pobreza e promover mobilidade social. O terceiro capítulo utiliza um modelo de diferenças em diferenças para avaliar um programa que descentralizou parcialmente a administração do Imposto Territorial Rural para os municípios em troca da elevação na participação dessas receitas. Usando microdados de declarações fiscais, constato que o programa levou a uma expansão das receitas tributárias em 20% após cinco anos. A descentralização expandiu as receitas principalmente por meio do aumento nos valores declarados das terras. Utilizando dados de satélite, verifico que a descentralização parcial não influenciou significativamente o comportamento dos produtores rurais. Um exercício de custo-benefício indica que a descentralização parcial gerou retornos elevados. No geral, os resultados indicam que a cooperação entre autoridades locais e centrais pode aumentar a tributação de propriedades em contextos de informação incompleta e capacidade de fiscalização limitada.

Palavras-chave: Mobilidade Intergeracional, Transferência Condicional de Renda, Taxação Propriedade, Economia Aplicada.

List of Figures

1.1	Baseline Mobility Curve in Brazil	13
1.2	Transition Probability Matrix by Quintile	14
1.3	Mobility Curve: Baseline vs. Administrative Income Data Only	15
1.4	Alternative Measures	17
1.5	The Great Gatsby Curve	18
1.6	Intergenerational Income Elasticity (IGE)	19
1.7	Mobility Curves by Gender and Race	20
1.8	Racial transition matrix	21
1.9	Long-Term Outcomes	22
1.10	Absolute Mobility Map: Predicted Rank for a Below-Median Income Child	24
1.11	The Great Gatsby Curve across Brazilian regions	25
1.12	Educational mobility across Brazilian regions	26
1.13	Exposure Effect Estimates for Children's Income Rank in Adulthood	28
2.1	PBF effects in the Long Run: ITT event studies	38
3.1	Evolution of the Decentralization Agreements	49
3.2	Effects of decentralization on ITR on revenues	54
3.3	Placebo Tests	55
3.4	Decomposing the effects of decentralization on ITR revenues	57
3.5	Self-reported versus minimum land values	58
3.6	The Effects of Decentralization on Land Use	59
A.1	Number of parent-child links relative to the population by cohort	77
A.2	Average share of imputed income by income decile	79
A.3	Income Distribution	81
A.4	Sensitivity of child and parental income to timing	87
A.5	Evolution of children's position at the income distribution	88
A.6	Mobility curve using household income for children	90
A.7	Neighborhood-based measure: children changing address	90
A.8	Individual mobility curves in Fortaleza and Belo Horizonte	92
A.9	Price-Adjusted Absolute Mobility Map	93
A.13	Placebo tests: Cohort-specific convergence	97
A.10	Mobility Correlates	100
A.11	Mobility Correlates: Principal Components	101
A.12	Mobility Correlates: LASSO Regularization	101

C.1	The Effects of Decentralization Agreements – No Controls	110
C.2	The Effects of Decentralization Agreements – Weights	111
C.3	The Effects of Decentralization Agreements – (CALLAWAY; SANT’ANNA, 2021)	112
C.4	IHS(ITR Due) by firm size	113

List of Tables

1.1	Income Distribution Statistics	10
1.2	Quantifying IGM biases due to formal income imputation	16
2.1	Mother's characteristics by deciles of predicted PBF score.	36
2.2	PBF effects in the Long Run: Age-specific ITT estimates by gender	42
2.3	PBF effects in the Long Run: Heterogeneity by Family and Place Inputs	43
3.1	Tax Rates of the ITR	48
3.2	Descriptive Statistics	50
3.3	Treatment and Control Group Characteristics	52
3.4	Cost Benefit Analysis (2021)	60
A.1	Descriptive statistics of main sample	78
A.2	IGM estimates fully based on years when tax data are available	82
A.3	Quantifying IGM estimation biases due to measurement error: the impacts of informal and formal non-labor income imputation	83
A.4	Quantifying IGM estimation biases due to measurement error: the impacts of formal income imputation, additional IGM measures	83
A.5	Quantifying IGM estimation biases due to measurement error: the impacts of informal income and formal non-labor income imputation, additional IGM measures	84
A.6	Robustness to larger samples	85
A.7	Reweighting Procedure	86
A.8	Alternative Income Measures	89
A.9	Labor market differences by gender and race	91
A.10	Siblings comparisons by parental income quintile	91
A.11	Summary of mobility estimates for the 50 largest metropolitan areas	98
A.12	List of municipal socioeconomic indicators and data sources	99
A.13	Placebo test: Gender- and race-specific convergence	102
A.14	Placebo test: Distributional convergence	102
B.1	Description of Outcome Variables	103
B.2	First-stage relationships and outcomes counterfactuals	104
B.3	PBF effects in the Long Run: Age-specific ITT estimates by race	105
B.4	PBF effects in the Long Run: Age-specific ITT estimates by region	106
B.5	PBF effects in the Long Run: Robustness to Alternative Treatment Definitions	107

C.1	Average effect of ITR by region	114
-----	---------------------------------	-----

Contents

1	Intergenerational Mobility in the Land of Inequality	1
1.1	Introduction	1
1.2	Institutional Background	5
1.3	Mobility measures	6
1.4	Data and income measurement	6
1.4.1	Family links	7
1.4.2	Income	8
1.5	Measurement error	10
1.5.1	Bias decomposition	10
1.5.2	Alternative approaches to rank incomes	11
1.6	Income mobility at the national level	12
1.6.1	IGM estimates	12
1.6.2	Measurement error	13
1.6.3	Additional robustness exercises	16
1.6.4	Alternative mobility measures	17
1.6.5	Cross-country comparisons	18
1.6.6	Mobility by gender and race	19
1.6.7	Parental income and children's long-term outcomes	20
1.7	Geographic variation in mobility	23
1.7.1	Geographical units and IGM measures	23
1.7.2	Regional mobility patterns	23
1.7.3	Income mobility and educational mobility	24
1.8	Causal place effects	25
1.8.1	Data and research design	26
1.8.2	Results	27
1.9	Conclusion	29
2	Do CCTs Create Conditions to Thrive?	
	<i>Bolsa Família</i> and Social Mobility in Brazil	31
2.1	Introduction	31
2.2	Institutional Background	34
2.2.1	The <i>Programa Bolsa Família</i> (PBF)	34
2.2.2	PBF in the Short-run and Expected Long-term Effects	35
2.3	Data and Methodology	35
2.3.1	Data	35

2.3.2	Empirical Strategy	36
2.4	Results	37
2.4.1	PBF Long-term Effects	37
2.4.1.1	Human Capital	39
2.4.1.2	Labor Market	40
2.4.1.3	Intergenerational Mobility	40
2.4.1.4	Migration	40
2.4.1.5	Health and Behavior	41
2.4.2	Racial and Regional Heterogeneity	41
2.4.3	Heterogeneity in PBF Effectiveness	41
2.5	Conclusion	43
3	Decentralization, Tax Administration, and Taxation: Evidence from Brazil's Rural Land Tax	45
3.1	Introduction	45
3.2	Institutional background	47
3.2.1	Rural land taxation in Brazil	47
3.2.2	The Decentralization of the Rural Property Tax	48
3.3	Data	50
3.3.1	Land Tax Returns	50
3.3.2	Additional Datasets	51
3.4	Empirical Strategy	51
3.5	Results	52
3.5.1	ITR revenues	52
3.5.2	Mechanisms	56
3.5.3	Efficiency Costs	58
3.5.4	Cost Benefit	60
3.6	Conclusion	60
	References	63
		63
A	Intergenerational Mobility in the Land of Inequality	75
A.1	Appendix to Section 4	75
A.1.1	Description of data sources	75
A.1.2	Family links	76
A.1.3	Sample selection	77
A.1.4	Imputation method based on random forests	78
A.1.5	Imputation method accuracy	79
A.1.6	Income distribution in our main sample and survey data sources	81
A.2	Appendix to Section 5	81
A.2.1	Decomposition of biases due to measurement error	81
A.3	Appendix to Section 6	82
A.3.1	Estimates fully based on years when tax data are available	82

A.3.2	Quantifying measurement error biases, additional IGM measures	83
A.3.3	Additional robustness exercises	84
A.3.3.1	Sample selection	84
A.3.3.2	Timing of income measurement	86
A.3.3.3	Alternative income and occupation definitions	88
A.3.4	Individual vs. household income for child ranks	89
A.3.5	Neighborhood-based measure for movers	90
A.3.6	Additional results	91
A.4	Appendix to Section 7	92
A.4.1	Individual mobility curves	92
A.4.2	Robustness of subnational estimates, price differences	92
A.4.3	Mobility estimates for the 50 largest metropolitan areas	93
A.4.4	Mobility correlates	93
A.5	Appendix to Section 8	94
A.5.1	Sample construction	94
A.5.2	Defining the predicted outcomes of permanent residents	95
A.5.3	Parametric specification and family fixed effects	95
A.5.4	Overidentification tests	95
B	Do CCTs Create Conditions to Thrive?	
	<i>Bolsa Família</i> and Social Mobility in Brazil	103
B.1	Appendix to Section 3	103
B.1.1	Definition of Outcomes	103
B.2	Appendix to Section 4	104
B.2.1	First-stage relationships and outcomes counterfactuals	104
B.2.2	Racial and Regional Heterogeneity	105
B.2.3	Robustness to Alternative Treatment Definitions	107
C	Decentralization, Tax Administration, and Taxation: Evidence from Brazil's Rural Land Tax	109
C.1	Additional Results	109

Intergenerational Mobility in the Land of Inequality

1.1 Introduction

Intergenerational mobility (IGM) is a long-standing interest in social sciences and the public debate. The extent to which children’s opportunities are determined by parental income is a relevant question from both equity and efficiency perspectives. Moreover, evidence on the actual degree of IGM can shift preferences for redistributive policies, as recently shown by (ALESINA; STANTCHEVA; TESO, 2018).

A large literature has used survey data to document that parental income is a relevant predictor of child income in adulthood, and that the strength of this association substantially varies across countries (e.g., see (NARAYAN et al., 2018; WEIDE et al., 2021)). More recently, a new wave of studies starting with (CHETTY et al., 2014) have relied on nation-wide tax data to study these questions. The use of large-scale data allowed for innovative analyses and new stylized facts – e.g., showing that mobility sharply varies for children growing up in nearby neighborhoods. Because similar data are typically not available for medium- and low-income economies, such studies have been largely restricted to the context of high-income countries.¹

In this paper, we study income mobility in Brazil, a large developing country characterized by extreme inequality in socioeconomic conditions. In 2019, the Gini index was as high as 0.53 – the 9th highest worldwide – and the top 10% of the population holds 43% of the country’s income (IBGE, 2019). To conduct our analysis, we combine rich individual-level data from multiple population-wide administrative registries and large-scale household surveys. In addition to estimating IGM at the national level, these data allow us to document in detail how income mobility varies by groups and fine geographical units, and to study the role of causal place effects for upward mobility.

Importantly, we address a major challenge in the estimation of IGM in Brazil that is common to other developing countries. Almost a third of the Brazilian economy is informal and, as such, is not reported in administrative registries. We predict informal income by training machine learning (ML) models on rich survey data reporting income from all sources.² We use the same method to impute formal non-labor income – notably, dividends and capital gains – for

¹For example, see (ABBAS; SICSIC, 2022; ACCIARI; POLO; VIOLANTE, 2021; BRATBERG et al., 2017; CHETTY et al., 2014; CONNOLLY; CORAK; HAECK, 2019; DEUTSCHER; MAZUMDER, 2020; HEIDRICH, 2017; HELSØ, 2021).

²We show that our ML algorithm yields improvements in accuracy relative to saturated OLS regressions. Specifically, we show that saturated OLS predictions perform poorly out-of-sample due to overfitting issues, which are not present for our ML algorithm.

earlier time periods when tax data are not available or for individuals not required to file taxes.³

This approach allows us to measure both formal and informal income at the individual level for a large, representative sample of 1.3 million children born between 1988-1990 and their parents. We use these data to estimate several IGM measures. Our main measures are based on the relationship between the percentile income rank of children and their parents', in line with (CHETTY et al., 2014). The estimated slope coefficient of the rank-rank regression equals 0.55, meaning that a 10 percentile increase in parental income is associated with an average 5.5 percentile increase in child income during adulthood.

In terms of absolute mobility, children born to below-median income parents reach on average the 36th income percentile in adulthood. A transition matrix between parental and child income quintiles shows that only 2.5% of children born to parents in the bottom quintile surge to the top quintile, and only 4% of those born to parents in the top quintile fall to the bottom quintile. In turn, almost one in two children born in the bottom and top quintiles remain at the same quintiles when adults. We show that our main results are unaffected by several robustness checks that address measurement and estimation issues, notably sample selection, attenuation, and life-cycle bias.

The approach that we develop for predicting income is crucial for obtaining these results, as relying exclusively on payroll and tax data on formal income would result in a much flatter rank-rank regression (slope equal to .35). This large attenuation bias is due to administrative data neglecting informal income for a large share of individuals at the bottom of the distribution and, also, dividends and other types of capital income at the top of the income distribution when tax data are not available.

Importantly, we develop simulation exercises to quantify biases arising from the fact that a significant portion of income in our analysis is imputed. They address different types of potential biases due to income imputation, discussed in earlier literature (e.g., see (CROSSLEY; LEVELL; POUPAKIS, 2022; INOUE; SOLON, 2010; JERRIM; CHOI; SIMANCAS, 2016; ZIMMERMAN, 1992)). They also address potential biases related to measurement error in survey-based income measures (ABOWD; STINSON, 2013; BOUND et al., 1994; GOTTSCHALK; HUYNH, 2010; KIM; SOLON, 2005). The central idea is that we can learn about such biases by replacing accurately measured income components (from administrative data sources) with predicted income and studying the impacts of these changes on IGM estimates. These exercises suggest that such biases are quantitatively small in our context and unlikely to significantly affect our estimates. We also develop a formal decomposition of these biases, which is close in spirit to (GOTTSCHALK; HUYNH, 2010) who study intragenerational mobility. We show that different bias components are small in magnitude and partially offset each other, explaining the small overall bias.⁴ These exercises make no assumption on the distribution of the measurement error components or on the model determining income, and hold for any model and covariates that one might use to predict child and parental income.

³Formal labor income is available, in all these cases, from employee payroll data.

⁴The key assumption for these exercises is that measurement error on informal income predictions follows a similar structure as those on formal income predictions (based on the same model and survey data). We also replicate the same exercise on a (selected) subset of the survey data where we can link the total income of parents cohabiting with adult children. We show that replacing informal income with predicted informal income has little impact on estimated mobility measures.

Our main findings are also robust to using two alternative, novel approaches that rank parents and children on socioeconomic status without the need to impute income components unobserved in administrative data. The first method exploits that (i) Brazilian workers move very frequently between formal and informal jobs (ULYSSEA, 2018, 2020), and (ii) more than 80% of individuals in our sample hold at least one formal job in our analysis period. We can thus rank them on the average income earned during periods of formal employment, which is precisely recorded in administrative employment data, as a measure of individual-specific “productivity”. The second approach ranks parents and children on a “neighborhood-based” income measure, defined by the average (formal) income across 300 thousand census tracts, leveraging data on more than 500 million residential addresses. The rationale for this second approach is that residential choices strongly correlate with socioeconomic status, particularly in highly unequal contexts such as the Brazilian one.

Each method has advantages and disadvantages, but the rank-rank curves estimated using both these two approaches are largely consistent with the one obtained using our baseline method. Importantly, these alternative approaches may be viable in other contexts characterized by a paucity of data on informal income.

Our large-scale data allow us to explore how upward mobility varies with individual characteristics and across geographical areas. A girl born to below-median income parents ranks on average 14 percentiles below boys born with the same parental income, and this gap is unaffected when we restrict the comparison to siblings. In turn, whites rank on average 7 percentiles above non-whites with the same parental income, and the gap is larger for below-median income families. While these results are broadly in line with previous evidence on differences in IGM by race in the US (DAVIS; MAZUMDER, 2018; CHETTY et al., 2020), they are all the more remarkable in the context of Brazil, where non-whites are not a minority group but instead represent about half of the population. We also document that higher parental income is associated with an improvement in several long-term outcomes – e.g., related to education, mortality, teenage pregnancy, welfare dependency and victimization.

Turning to heterogeneity across local areas, we uncover a mobility divide between the wealthier Center-South regions and the poorer Northern regions. A second key finding is that poor children born in the largest economic centers such as Sao Paulo and Rio de Janeiro do not achieve the best outcomes. Instead, southern regions colonized by European immigrants in the late XIX century and Center-Western regions that recently experienced a “soy boom-driven” economic growth exhibit the highest degrees of upward mobility.

Motivated by these stark regional divides, we estimate causal place effects on absolute mobility leveraging (within siblings) variation in age at move among the children of migrating families (CHETTY; HENDREN, 2018a). Movers converge linearly to the income of permanent residents in the destination area at a rate of 2.4% per year of childhood exposure, meaning that children moving at birth to a place where they are expected to rank 10 percentiles higher will increase their rank by 5.76 percentiles on average due to causal place effects.⁵ Hence, these effects explain more than half of the regional differences in absolute mobility across Brazil.

Our paper contributes to a recent body of literature estimating IGM using large-scale tax

⁵In this analysis, we measure income at the age of 24. Hence, exposure from birth to the age of 24 implies a $24 \times 2.4\% = 57.6\%$ convergence.

data. Starting with the seminal paper by (CHETTY et al., 2014) in the US, this literature has focused exclusively on rich countries, mainly due to data constraints.⁶ Our paper is the first one that studies income mobility in a large developing country using population-wide administrative registries, while previous evidence on developing countries largely relied either on survey data (DUNN, 2007; FERREIRA; VELOSO, 2003; LEONE, 2018; MAHLMEISTER et al., 2017; NARAYAN et al., 2018) or on educational mobility as a proxy for income mobility (e.g. (ALESINA et al., 2021; ASHER; NOVOSAD; RAFKIN, 2021; SAAVEDRA; ANDRES, 2022b)). Finally, (LEITES et al., 2022) and (MENESES, 2020) have access to administrative data on income but do not attempt to estimate informal income.

Our second main contribution is methodological, as we devise new approaches to measure income mobility in contexts of high labor informality.⁷ Specifically, we show how survey and admin data can be combined for improving income measurement and we develop new methods to quantify and decompose any bias from errors in income imputation, which has been a major concern in the IGM literature since the seminal work by (SOLON, 1992). Additionally, we develop two novel methods for ranking parents and children on economic status without the need to impute informal income. These tools can be adapted to estimate IGM in other countries – including many developed economies – that are also characterized by a large informal sector.⁸ More generally, they may find application in investigations tackling similar measurement challenges. For instance, in research using administrative data to study income dynamics and inequality in contexts where the underground economy is relevant (ENGBOM et al., 2022; GUVENEN; PISTAFERRI; VIOLANTE, 2022), and in studies relying on different forms of income imputation which naturally lead to measurement error (e.g. see (JÁCOME; KUZIEMKO; NAIDU, 2021)).

Finally, we contribute to the literature studying the impact of places on social mobility (e.g. (CHETTY; HENDREN, 2018a, 2018b; DEUTSCHER, 2020)) and other long-term outcomes (e.g. (CHETTY; HENDREN; KATZ, 2016; CHYN, 2018; DAMM; DUSTMANN, 2014)). In line with previous evidence from the US and other rich countries (CHYN; KATZ, 2021), we find that causal place effects explain a large share of the total variation in intergenerational mobility. These results add to recent evidence showing that places matter for educational mobility in Africa and Latin America (ALESINA et al., 2021; SAAVEDRA; ANDRES, 2022a).

The remainder of the paper proceeds as follows. Section 1.2 briefly introduces the Brazilian context, followed by Section 1.3 describing our mobility measures and Section 1.4 describing our data, family linkage and income measurement methods. Section 1.5 tackles measurement error issues, while Section 1.6 presents our main IGM estimates at the national level and by subgroups. We explore geographic variation in mobility in Section 1.7 and estimate causal place effects in Section 1.8. Finally, Section 1.9 concludes.

⁶(SOLON, 1999) and (BLACK; DEVEREUX et al., 2011) review previous studies relying mainly on household surveys, while (BLANDEN, 2013) and (BJÖRKLUND; JÄNTTI, 2020) consider alternative approaches.

⁷On the relationship between informality and economic development, see, e.g., (PORTA; SHLEIFER, 2014) and (ULYSSEA, 2020).

⁸(MEDINA; SCHNEIDER, 2018) estimate that, during the period 1991-2015, one fourth of Italian GDP is produced in the informal economy, and the size of the informal sector accounts for as much as 15% of GDP in countries like Canada, Denmark, Norway, and Sweden.

1.2 Institutional Background

Brazil is the fifth largest country in the world by area and the sixth by population size, hosting nearly one-third of the population in Latin America, and it has historically been characterized by extreme socioeconomic inequality. In 1990 – roughly the period when our cohorts of children were born – the Gini index was as high as 0.60, placing Brazil as the fifth most unequal country in the world, and the first one outside Africa. Although inequality has subsequently followed a mildly decreasing trend, the Gini index remained as high as 0.53 in 2019. According to official estimates, the top 10% of the population holds 43% of the country's income (IBGE, 2019), compared to 31% in the US, 29% in China, and around 25% in European countries.⁹

The country's colonial past, characterized by short-spanned extractive economic cycles and over 350 years of slavery, bestowed strong social disparities. The gap in income per capita between white and non-white households is over 35%. Non-whites represent nearly half of the population but account for 64% of the unemployed, 67% of the incarcerated population, and 75% of the beneficiaries of *Bolsa Família* cash transfers. Socio-economic conditions also vary widely across geographical areas. The country comprises 27 states (and 5,570 municipalities), and GDP per capita is about 40% lower in Northern states relative to the more developed Center-South. The homicide rate ranges from above 50 per 100k inhabitants in poorer states such as Roraima and Ceará to below 12 in the richest states such as São Paulo and Santa Catarina. These facts further motivate an analysis of mobility across subgroups and geographical areas.

Like in most low- and middle-income countries, the labor market is characterized by a large degree of informality. Labor turnover is also very high, with 70% of formal jobs lasting less than a year, and it is common for workers to turnover between the formal and informal sector (ULYSSEA, 2018, 2020). In our data, 82.8% of men have held at least one formal job over their lifetime, but about 40% of workers are employed in the informal sector in a given year. Hence, it is crucial to properly measure informal income in our analysis.

The bulk of income taxes in Brazil is collected on formal labor income, although around half of formal workers are fully exempted from filing yearly income taxes because they earn below the first tax bracket (BRL 22,847 in 2019).^{10,11} For the same reason, most informal workers would not pay taxes even if they had an official contract, since the majority of them earn below the first tax bracket. Dividends are fully exempt from income taxes.¹²

Individual income taxes are exclusively levied by the federal government and marginal tax rates range from 7.5% to 27.5%. Tax filings are mandatory for individuals with earnings above the first tax bracket, for all firm owners and for all individuals with any capital gains, any stock market operations, or property wealth above BRL 300,000.¹³ Individuals filing taxes must report all (formal) income sources, including tax-exempted ones.

⁹Estimates based on the World Bank's Poverty and Inequality Platform (World Bank, 2021).

¹⁰Throughout the paper, we refer to BRL at 2019 prices. In 2019, the purchasing power parity rate was 2.28 relative to the US dollar.

¹¹For instance, in 2015 only slightly more than 27 million tax forms were filled in a universe of over 60 million formal workers.

¹²For simplicity, throughout the paper we refer to all types of withdrawals by firm owners as dividends.

¹³Starting in 2010, a small share of firm owners receiving dividends below 40,000 BRL were no longer required to file taxes.

1.3 Mobility measures

Following the recent literature (e.g. (CHETTY et al., 2014; ACCIARI; POLO; VIOLANTE, 2021)), we focus on the relationship between children and parents' income ranks, as originally proposed by (DAHL; DELEIRE, 2008). Since this relationship tends to be linear, it can be summarized by a few statistical parameters that can be compared across areas and groups. We estimate the linear regression:

$$y_i = \alpha + \beta p_i + \varepsilon_i \quad (1.1)$$

where y_i and p_i are, respectively, the income percentile rank of child i and her parents' at the national level, ordered from 1 to 100. Child ranks are measured relative to their own cohorts, and parents' ranks are measured relative to other parents with children from the same cohorts.

The estimated parameters in equation (1.1) provide us with two IGM measures. The slope coefficient β measures the (inverse) *relative mobility* of children born to parents who are 1 percentile apart in the parental income distribution. A higher β means a wider gap between the two, thus implying lower IGM. In a perfectly mobile society, the rank-rank slope would equal zero as children's long-term outcomes would be unrelated to parental income.

The intercept α equals the expected rank for children at the bottom of the parental income distribution. Combining α and β , one can recover the expected rank for children born at any point of the income distribution. Following previous literature (see (CHETTY et al., 2014)), we focus on the expected rank of children born in below-median income families as our main measure of *absolute mobility*, which we also refer to as *upward mobility* throughout the paper. In turn, the latter equals the expected rank for children whose parents are in the 25th percentile of the income distribution (i.e. $\alpha + 25 \times \beta$). This measure is particularly useful to characterize geographical variation in mobility patterns, as it compares the outcomes of children born in different regions of the country while holding constant parental income.

In addition, we construct transition matrices from parental income quintiles to child income quintiles. In particular, we focus on the chances of *escaping poverty* – defined as the probability that children born to parents in the bottom quintile do not belong to the same quintile when adults –, and on the probability that children move from the bottom to the top quintile of the income distribution within one generation (CORAK; HEISZ, 1999). We also estimate intergenerational income elasticities (IGE), defined by the correlation of children's and parents' log incomes. IGE allows for a comparison with earlier survey-based studies in Brazil and other countries (see, e.g. (DUNN, 2007; LEE; SOLON, 2009; BLACK; DEVEREUX et al., 2011)).

Finally, we also document the association between parental income and several children's long-term outcomes beyond income – namely, education, access to prestigious occupations, victimization, mortality, and teenage pregnancy.

1.4 Data and income measurement

Estimating the mobility measures described in the previous section requires (i) linking one or more cohorts of children to parents at the individual level, and (ii) measuring their individual

income. Constructing such data for Brazil faces two main challenges, which are common in the context of developing countries. First, comprehensive registries of family links (of the type available, e.g., for Scandinavian countries) are not readily available. Second, a large portion of income is earned in the informal economy and, as such, it is not reported in administrative registries. We next describe how we overcome these challenges by combining several sources of individual-level data to recover family relationships, and training supervised ML models on large-scale survey data to impute informal income. In Appendix A.1.1, we describe all data sources used in the paper and how we link them to our main sample.

1.4.1 Family links

We aim to link each child's unique person code (*CPF*) to their parents'. Our starting point is dependent claims in individual tax returns data for the 2006-2020 period, provided by the Brazilian tax authority (*Receita Federal do Brasil*). Parents report children aged 0-24 for the purpose of tax deductions, in which case we can directly link them to each other through the unique person codes available in these data.¹⁴ However, only one-third of Brazilians – mainly in the upper part of the income distribution – file taxes every year (unlike in the context of rich countries, where much larger shares of the population file taxes). Therefore, we rely on additional data sources to link children who are not claimed by their parents in the tax data.

We link unclaimed children to their mothers using the Brazilian person registry (*Cadastro de Pessoas Físicas*), which covers the entire population and is provided by the Brazilian tax authority. All individuals are identified by their person code, full name, and mother's full name. If the mother can be uniquely identified by her name – as is the case for 52% of Brazilians – we link the child's person code to her mother's based on the mother's name.¹⁵ Since fathers' names are not available in the person registry, we rely on a welfare registry (*Cadastro Único*) to link children to their fathers. The registry covers around two-thirds of the Brazilian population and contains the father's name for all individuals, along with person codes.¹⁶ Since it provides the informative basis for administering social programs such as *Bolsa Família*, the registry mainly covers the low and middle parts of the income distribution. We implement the same procedure as before, linking children to their fathers conditional on the father having a unique name in the country, so that we can precisely identify his person code.

Overall, 49% and 25% of the children of the 1988-1990 cohorts can be linked to their mother and father, respectively.¹⁷ Our main sample is defined by 1.34 million children who can be linked to both parents, accounting for around 15% of the entire 1988-1990 cohorts. In Appendix A.1.3, we show that our main sample is fairly representative of the population in terms of

¹⁴Children aged 22-24 can only be reported if they are enrolled in technical school or higher education.

¹⁵The share of individuals with a unique name in the country is large because Brazilians typically carry one or more surnames from both their parents. These individuals are easily identified in the person registry as the latter includes both names and person codes. (BRITTO; PINOTTI; SAMPAIO, 2022) show that individuals with unique names do not strongly differ from the overall population along several characteristics.

¹⁶We combine yearly snapshots of this registry for the 2011-2020 period, with 135.6 million individuals in total.

¹⁷Using younger cohorts reduces the period in which we can measure their income as adults, whereas using earlier cohorts reduces our sample because tax data on dependent claims starts in 2006. Nevertheless, we show in Appendix A.3.3.1 that our main findings remain similar when using additional cohorts.

several individual characteristics, given that our procedure relies on complementary data sources covering different parts of the income distribution.¹⁸

Importantly, we also show that our main findings are robust to: (i) using the less conservative linkage procedure, which increases sample coverage from 15% to up to 45% of the population (see Appendix A.1.2); (ii) extending the sample to additional cohorts; and (iii) re-weighting the sample to eliminate any remaining differences in characteristics between our working sample and the general population.

1.4.2 Income

Our mobility measures are based on individuals' total income, defined as the sum of formal and informal income. Accounting for informal income is crucial given the size of the informal labor market (about 40% of all jobs). For this purpose, we develop a novel approach leveraging rich survey data and ML methods to estimate income that is unobserved in administrative data.¹⁹

FORMAL INCOME. Tax records cover all sources of formal income earned by an individual in a given year, including both labor and non-labor components. Tax-exempt income (e.g., dividends) must also be reported in tax filings. However, tax data are not always available for two reasons: first, only a third of Brazilians file taxes each year; and, second, tax data are available from 2006 onwards, limiting our ability to measure parental income until children in the main sample – born in 1988-1990 – are aged 16-18.²⁰

Whenever tax records are unavailable for a given individual in a given year, we measure formal income as the sum of a labor and non-labor component. The first component – formal labor income – is directly available from administrative employment data covering the population of formal jobs for the 1985-2019 period (*Relação Anual de Informações Sociais*, RAIS).²¹ The second component – formal non-labor income – includes dividends, rents, interests, and capital gains, which are not available in administrative registries other than tax data. We thus follow an imputation procedure to predict formal non-labor income leveraging survey data sources. The procedure is the same one used to input informal income, which we describe next.

INFORMAL INCOME. While the Brazilian administrative registries allow us to accurately measure formal income, they do not contain – by their very nature – information on informal income. We measure the latter using individual-level data from two large-scale surveys: *Pesquisa Nacional por Amostra de Domicílios* (PNAD), a cross-sectional household survey covering about 400,000 individuals per year for the 1992-1999 and 2001-2019 periods; and Population Census surveys covering 10% of the population in 1991, 2000, and 2010. Both surveys are collected

¹⁸Specifically, the tax registry covers the upper part of the income distribution; the person registry covers the entire distribution (for mothers); and the welfare registry covers the lower and middle part of the distribution (for fathers).

¹⁹Our main measure of parental income is the sum of father and mother's incomes. We proceed by estimating income for all individuals in our data to later aggregate the income of fathers and mothers.

²⁰We show that our results remain similar when measuring parental income only for years when tax data are available (Section A.3.3).

²¹RAIS has been extensively used in previous research on the Brazilian labor market, see e.g. (FERRAZ; FINAN; SZERMAN, 2015) and (GERARD; GONZAGA, 2021).

by the Brazilian Institute of Geography and Statistics (IBGE), which has a long tradition of measuring informal income for estimating GDP and other national aggregates.

We impute informal income based on a rich array of individual characteristics available in both administrative registries and survey data. This is a typical prediction problem, which we address using random forests (RF), a supervised ML algorithm that endogenously splits the space of covariates to generate predictions for a given outcome – see Appendix A.1.4 for details. The key advantage relative to a fully-saturated OLS is that it avoids excessively splitting the sample.

We grow a separate RF to predict informal income in each year from 1991 to 2019 by training the algorithm developed by (ATHEY; TIBSHIRANI; WAGER, 2019) on our survey data. The vector of predictors includes a wide array of individual characteristics: state of residence (27), a dummy identifying metropolitan regions, gender, age, race (white vs. non-white), education dummies (4), and occupation category (dummies for formal worker, formally self-employed, and firm owner, with informal workers being the residual category).

After training the model, we predict informal income for all individuals in our main sample, including formal workers and owners who may earn part of their total income in the informal sector. We repeat the same process for estimating formal non-labor income, which is necessary for measuring total formal income when tax data are not available.

In Appendix A.1.5, we estimate that our procedure based on the RF model predicts income ranks based on a single year with a fairly high R-squared of .57, which helps mitigating measurement error issues. As a comparison, using a fully-saturated OLS leads significantly smaller R-squared of .29 due to overfitting issues.²² We also provide evidence that averaging out income over multiples years further increases the precision of our predictions.

MAIN SAMPLE. In the main analysis, we measure the average income of children born in 1988-1990 over the period 2015-2019, when they are 25-31 years old, and relate it to the average income of their parents (father plus mother) at the time when children were 3-18 years old.²³ The median parental and child annual income is BRL 47,068 and BRL 19,730, respectively, while the share of total income held by the top decile is around 40% for both populations. Table 1.1 displays descriptive statistics for the full sample and separately by gender and race. Appendix Figure A.3 shows that the distribution of total income in our sample matches the distribution of total income in the PNAD survey, apart from some (small) differences at the bottom of the father's income distribution.

²²R-squared statistics are based on out-of-sample predictions using a random subsample of the survey data – not used for training the prediction models.

²³Three years old is the earliest age at which we can measure parental income for our oldest cohort born in 1988, since survey data (in the format that we use) is available since 1991. In turn, we measure child income setting a five-year window as late as possible. Income data from the tax authority cover children in our main cohorts over the period 2015-2019, and their parents over the period 2006-2010. In addition, RAIS data are available until 2019. Appendix A.1.4 shows that averaging income over several years significantly increases the precision of our predictions, while Appendix A.3.3 shows that our main results are not affected by life-cycle bias and alternative income definitions.

Table 1.1: Income Distribution Statistics

	Parents			Children		
	5%	50%	95%	5%	50%	95%
All	16,044	47,068	249,468	8,515	19,730	102,068
Males	9,509	31,997	193,521	10,604	22,226	117,414
Females	5,202	13,046	64,874	7,736	16,005	87,762
White	19,906	53,931	282,546	10,419	22,274	111,560
Non-white	13,797	32,917	187,921	7,388	16,267	81,658

Notes: The table reports the average yearly income at the 5th, 50th, and 95th percentiles of both parents and children (in 2019 BRL). The first row refers to the entire sample, while the other rows present separate statistics by gender and race. In the columns for “Parents”, the entries for “Males” and “Females” report individual incomes of fathers and mothers, respectively, while all other entries report household income. The columns for “Children” always report individual income.

1.5 Measurement error

Even though imputation is crucial for properly measuring total income, it carries with it some degree of measurement error that may bias our mobility estimates. Such error is unlikely to be classical for at least three reasons. First, errors are likely correlated across generations. Specifically, our algorithm underestimates the income of individuals with high unobserved ability, who earn above the average of their group. Such error will be positively correlated for parents and children if ability is transmitted across generations. Second, measurement error may also be non-classical because of mean-reversion: income errors tend to be negatively correlated with income levels in survey data (BOUND et al., 1994; GOTTSCHALK; HUYNH, 2010). Third, our imputation procedure is equivalent to an instrumental variable approach discussed in earlier literature (e.g., see (CROSSLEY; LEVELL; POUPAKIS, 2022; INOUE; SOLON, 2010; JERRIM; CHOI; SIMANCAS, 2016; ZIMMERMAN, 1992)). As such, violations of the exclusion restriction could drive a correlation between measurement error in parental income and unexplained child income (ε , equation (1.1)) and generate bias to our IGM estimates.

1.5.1 Bias decomposition

We provide a formal decomposition for the bias that non-classical measurement error may generate to our estimates. When estimating equation (1.1), we effectively estimate $y = y^* + \eta$ on $p = p^* + \mu$, where $\{y^*, p^*\}$ are actual child and parental income ranks and $\{\eta, \mu\}$ are the respective measurement error terms. We make no assumption on the distribution of such errors or on the income generation process. In fact, error components may come from any prediction models using any characteristics which are relevant for income determination. These characteristics could be fix, such as race, or correlated across generations such as ability and education. In Appendix A.2.1, we show that they lead to the following estimation bias:²⁴

²⁴The same appendix provides an intermediate decomposition for the case where estimates are not based on income ranks.

$$\hat{\beta} - \beta = -\frac{1}{2}\beta \frac{v(\mu)}{v(p)} + \beta_{\varepsilon\mu} \frac{v(\mu)}{v(p)} + \beta_{\eta p^*} + \beta_{\eta\mu} \frac{v(\mu)}{v(p)}, \quad (1.2)$$

where β is our coefficient of interest (i.e., the regression of y^* on p^*) and β_{ab} denotes the coefficient of a hypothetical OLS regression of a on b ; $v(\cdot)$ denotes the variance operator; and ε is the error-term in Eq. (1.1) (i.e., child income that is unexplained by parental income). The decomposition is close in spirit to (GOTTSCALK; HUYNH, 2010) who study the impacts of measurement error on intragenerational mobility.

The first term in the decomposition is a downward bias, which grows larger in magnitude as our estimates for parental income becomes more imprecise (i.e., larger $v(\mu)$), working in a similar way to attenuation bias caused by classical measurement error. The second term shows that a positive correlation between unexplained child income (ε) and measurement error on parental income will lead to an upward bias in the rank-rank slope. This captures biases due to the violation of the exclusion restriction documented in earlier literature.²⁵ The third term shows that a correlation between measurement error for child (v) and parental income (p^*) leads to an upward bias in the rank-rank slope. Finally, a positive correlation in measurement error across generations, $\beta_{\eta\mu}$, will bias our estimates of the rank-rank slope upward. The biases in the second, third and fourth components are explained by fact that inflating the left- and right-hand-side of the equation (1.1) at the same time drives a spurious correlation between child and parental income, leading to an upward bias in the estimation of the rank-rank slope.

In Section 1.6.2, we develop a simulation exercise to learn about the impact of measurement error on our estimates. The key insight is that we can learn about the impact of measurement error by replacing income components which are precisely measured in administrative data with predicted counterparts. This exercise suggests that measurement error leads only to relatively small biases in our context.

1.5.2 Alternative approaches to rank incomes

In addition to our main analysis, we rank parents and children on two novel measures of their overall economic conditions that do not require imputating informal income.

PRODUCTIVITY-BASED MEASURE. A large share of individuals in Brazil frequently turnover between the formal and informal sector, and about 80% of individuals in our main sample hold at least one formal job throughout their career. We can thus rank individuals based on the average monthly income earned during employment spells in the formal labor market as a measure of their individual-specific productivity. The underlying assumption is that the average productivity during employment spells in the formal sector – as measured by formal earnings – is a reasonable proxy for individual productivity when employed in the informal sector. Even though this method is unable to cover individuals who have never held formal jobs, it has the key advantage of exclusively relying on high-quality data on formal labor income.

²⁵Parental income measurement error directly depends on parental characteristics which may have a direct effect on child income. If that is the case, such characteristics will be part of the unexplained child income component (ε), driving this bias.

NEIGHBORHOOD-BASED MEASURE. Our second approach ranks parents and children based on the average income in the census tract in which they reside. Census tracts are small geographical areas designed to cover homogeneous groups of about 400 families throughout the country. The rationale for this measure is that residential choices are strongly correlated with income, particularly in poorer countries characterized by high inequality and socioeconomic spatial segregation. In addition, neighborhoods have a major impact on living standards and access to opportunities, as determined by access to public goods, job opportunities, and exposure to violence (e.g. see (BILAL; ROSSI-HANSBERG, 2021; CARD; ROTHSTEIN; YI, 2021)). In our data, variation between census tracts explain 29% of total variation in formal labor income. Another important advantage of this measure is that it may better capture the high living standards of individuals benefiting from inherited wealth or living on in-kind and informal family donations. While we acknowledge that this measure may differ in nature from a pure income measure, it tracks a relevant dimension of socioeconomic status and may be useful for validating our main results based on individual income.

To implement this strategy, we geocode unique data from the Brazilian tax authority tracking 500+ million residential addresses for the entire population in the 2000-2020 period, and assign them to a census tract using shape files provided by IBGE.²⁶ We then measure the average income in each location as the average labor income of residents holding formal jobs.

1.6 Income mobility at the national level

1.6.1 IGM estimates

Figure 1.1 plots the average and median income rank in adulthood for children born to parents in each income percentile, along with the inter-quartile range (i.e., the range between the 25th and 75th percentiles of the children's income distribution). The ranks are based on our main measure of total income, described in Section 1.4.2. The rank-rank relationship is approximately linear, with the exception of the very top percentiles of the distribution, which exhibit a steeper slope. Although similar patterns have been documented for Canada, Denmark, Italy, Norway, Sweden, and the United States (BRATBERG et al., 2017; CHETTY et al., 2014; CORAK, 2020), for the case of Brazil the change in slope is more concentrated at the very top of the distribution.


The rank-rank slope coefficient in Equation (1.1) equals 0.546, meaning that a 10 percentile increase in parental income is associated on average with a 5.46 percentile increase in children's income in adulthood. Based on this estimate, it would require seven generations for a family starting in the 25th percentile to reach the same rank of a family in the 75th percentile.²⁷

Regarding absolute mobility, a child born to parents in the 25th percentile is expected to

²⁶Specifically, we consider the place of residence for children in 2000, when they are aged 10-12, and the place where they live when adults in 2019.

²⁷Assuming that permanent income over generations is an $AR(1)$ process, the number of generations N required for families that are Δ percentiles apart to converge to the same percentile solves the equation $\beta^N \Delta = 1$, where β is the rank-rank slope coefficient (ACCIARI; POLO; VIOLANTE, 2021). This back-of-the-envelope calculation might be a lower bound given that recent empirical estimates find a stronger correlation between the grandparents' and grandchildren's incomes than an $AR(1)$ process would suggest (LINDAHL et al., 2015; BRAUN; STUHLER, 2018).

Figure 1.1: Baseline Mobility Curve in Brazil



Notes: The figure shows the relationship between parental and child income ranks at the national level, for our main sample (1988-1990 cohorts). For each parental income percentile, it plots the mean (blue dots), median (red triangles) and inter-quartile range (shaded area) of child income rank during 2015-2019, i.e. at the age of 25-31. Parental income is the sum of the father's and mother's average income when children are aged 3-18 years old. The figure also displays our absolute ($\alpha + \beta * 25$) and relative mobility (β) measures based on Equation (1.1).

reach the 36th percentile in adulthood. Figure 1.1 also shows that – even conditional on parental income – there is considerable variation in children's outcomes. For instance, the inter-quartile range of child ranks for parents at the 25th percentile is [17, 53].

Figure 1.2 shows the transition matrix between quintiles of the parental and child income distributions. The probability of raising from the bottom to the top quintile within one generation is only 2.5%, and the probability of falling from the top to the bottom is only 4%. Indeed, roughly half of the children born to parents in the bottom quintile fail to escape poverty, remaining in the bottom quintile when adult; similarly, half of the children born to parents in the top quintile remain at the top of the income distribution when adult.

1.6.2 Measurement error

In Figure 1.3, we show that the relationship between parent and child rank is severely attenuated when relying exclusively on administrative data, the slope coefficient decreasing from .546 to .357. The difference is particularly marked for informal workers with zero formal income, resulting in a flat relationship over the bottom 10% of the parental income distribution. In addition, the inter-quartile range is substantially larger at the upper side of the parental income distribution, which is likely due to neglecting dividends and other sources of non-labor income for years when tax data are not available. Therefore, imputing all these sources of income unreported in administrative data is crucial for correctly estimating IGM.

Nevertheless, as discussed in Section 1.5, our imputation process leads to some degree of measurement error that may bias our mobility estimates. To gauge the magnitude of such bias, we replace income components that are precisely measured in administrative data with predicted values based on the same ML models and survey data used to impute income in the

Figure 1.2: Transition Probability Matrix by Quintile

5	2.5%	5.1%	15%	29%	48.5%
4	10.2%	16%	23.6%	26.9%	23.4%
3	17.3%	24.5%	23.7%	20.3%	14.3%
2	24%	27.1%	22.8%	16.1%	9.9%
1	46.1%	27.4%	14.9%	7.6%	4%
	1	2	3	4	5

Parent Household Income Quintile

Notes: The figure shows the probability that children born to parents in a given quintile of the parental income distribution (horizontal axis) move to a given income quintile in adulthood (vertical axis). Darker red tones indicate higher probabilities.

main analysis. Then, we study how different IGM measures vary with the imputation process and decompose the implied bias in estimated rank-rank slope using equation (1.2).


First, we replace formal labor income for parents and children who lie below given income percentiles in our benchmark sample. The goal is emulating the imputation of informal income in our main analysis, which is mainly based on labor income and disproportionately falls upon low-income individuals – see Appendix Figure A.2. Table 1.2, columns 2-4, presents the results. In all cases, the rank-rank slope remains in the range .549-.570, close to our benchmark estimates (.546). Indeed, all the different bias components tend to be small – below .028 – and the attenuation bias in the first component is more than offset by the other components, resulting in a (small) upward bias. Other mobility measures, reported in the bottom panel of the table, also remain close to our benchmark estimates: absolute mobility lies in the range 36.0-36.5 (vs. 36.6 in the main analysis) and the transition probability from the first to the top income quartile is in the range 1.4%-2.4% (vs. 2.5%).

Second, we address the fact that we impute formal non-labor income for several years when measuring parental income, due to the lack of tax data before 2006. We extend the initial simulation by replacing formal non-labor income with predicted counterparts for measuring parental income in the period 2006-2010 onwards (when tax data is available). The results in Table 1.2, columns 5-7, reveal similar patterns to the initial exercise. The rank-rank slope remains in the range .539-.551, close to our benchmark, and the same holds true for other mobility estimates. Bias components continue to be small in magnitude and do not change direction.²⁸

Appendix Table A.2 provides yet another robustness tackling the issue that we need to

²⁸Bias components are only somewhat larger when replacing formal non-labor income for all parents, including those in the top quartile. This can be explained by the fact that survey data is more inaccurate at the top of the income distribution, highlighting the importance of using tax data in the analysis.

Figure 1.3: Mobility Curve: Baseline vs. Administrative Income Data Only



Notes: The figure shows our baseline mobility curve displayed in Figure 1.1 (blue dots) and the mobility curve obtained when solely relying on administrative data sources to measure income (red triangles). For each parental income percentile, we plot the mean child income rank during 2015-2019, i.e. when the cohorts of children in our main sample (1988-1990) were aged 25-31, along with the interquartile range. Parental income is the sum of the father's and mother's average income when the child is aged 3-18 years old. For each curve, the figure also displays the estimated β coefficient in Equation (1.1).

impute formal non-labor income for several years when measuring parental income. We show how our results vary when we measure parental income using only years for which tax data is available: we take the 5-year average during the period 2006-2010 for measuring parental income, while child income is measured as in our baseline. We find a rank-rank slope (.537 vs. .545) and absolute mobility close to our main estimates (36.8 vs. 36.58).

Next, we repeat our simulation exercise on PNAD data. The key advantage of this exercise is that we can replace exactly the same income components which we impute in our main analysis: informal income and formal non-labor income. We focus on a sample of adult children aged 25-34 who live with their parents, so that we can observe their incomes and estimate a rank-rank regression which we use as the benchmark for this simulation.²⁹ Given that this is a small and an extremely selected subsample of individuals and that income is observed for a single year only, we are not directly interested in the mobility estimates.³⁰ However, studying how estimates change as we impute informal income and formal non-labor income is informative on the size of measurement error biases that income imputation may generate in our main analysis. The results in Appendix Table A.3 show that key mobility estimates remain close to the benchmark after the imputation of these income components. For instance, the rank-rank slope is in the

²⁹We focus on the period 2006-2014, and also restrict the sample to fathers aged 45-64. To avoid overfitting concerns, we train the ML model on the survey data after dropping those households in our estimation sample (children cohabiting with their parents).

³⁰Incidentally, estimates based on this selected sample are similar to our main estimates. However, it would be difficult to draw strong conclusion based on the survey data only given these important limitations. In addition, the small sample size would not allow for analyses across subgroups or small geographical areas as done in the remainder of the paper.

Table 1.2: Quantifying IGM biases due to formal income imputation

		Replacing income components with predicted counterparts for individuals in different income quartiles					
		Formal labor income (all)			Formal labor income (all) and formal non-labor income (parents only)		
Benchmark		Q1	Q1-Q3	All	Q1	Q1-Q3	All
(1)		(2)	(3)	(4)	(5)	(6)	(7)
PANEL A. RELATIVE MOBILITY							
Rank-rank slope	0.546	0.549	0.561	0.570	0.549	0.551	0.539
SE	0.001	0.001	0.001	0.001	0.001	0.001	0.001
ME Bias decomposition							
Term 1: $-\frac{1}{2}\beta\frac{v(\mu)}{v(p)}$		-0.002	-0.013	-0.028	-0.002	-0.017	-0.052
Term 2: $\beta_{\varepsilon\mu}\frac{v(\mu)}{v(p)}$		0.005	0.013	0.027	0.004	0.007	0.019
Term 3: β_{vp^*}		0.001	0.014	0.021	0.001	0.014	0.021
Term 4: $\beta_{v\mu}\frac{v(\mu)}{v(p)}$		0.000	0.002	0.005	0.000	0.002	0.006
Total bias		0.004	0.016	0.024	0.003	0.006	-0.006
PANEL B. OTHER IGM MEASURES							
Exp. rank p=25	36.6	36.5	36.2	36.0	36.5	36.4	36.7
Q1Q5	2.5%	2.4%	1.8%	1.4%	2.4%	2.0%	2.1%
Q5Q5	48.5%	48.8%	51.3%	52.6%	48.8%	50.9%	50.4%
IGE	0.500	0.503	0.516	0.534	0.503	0.510	0.584

Notes: This table shows how benchmark relative mobility (Panel A) and several IGM measures based on our main sample change after replacing formal income with predicted counterparts for different groups (Panel B). Column 1 reports the benchmark estimates, while columns 2-6 reports IGM estimates after replacing formal income for specific groups of parents and children based on their income quartiles. Q1Q5 (Q5Q1) defines the probability that children born in income quintile 1(5) reach income quintile 5(1) in adulthood. Panel A also provides a decomposition of the total bias resulting from income imputation following the decomposition presented in Section ??

range .514-.52 (vs. .52) and absolute mobility is in the range 35.24-35.37 (vs. 35.23).³¹

These results offer a transparent assessment of the potential consequences of measurement error to our mobility estimates and clarify how different sources of biases interact. Overall, they suggest that the magnitudes of measurement error biases are reasonably small for different mobility measures and unlikely to overturn our key results showing strong persistence in income across generations in Brazil. In Appendix A.3.2, we show that similar results emerge for group and area-level mobility measures presented in the remainder of the paper

1.6.3 Additional robustness exercises

In Appendix A.3.3, we show that our main IGM estimates are not significantly affected by other sources of bias in IGM measurement, namely selection, life-cycle, and attenuation bias. In Appendix A.3.4, we show that focusing on household income to measure child ranks also has


³¹The first and second bias components are larger in this simulation because we measure income based on a single survey year since it is not possible to follow individuals across years in PNAD, so income imputation is less precise than in our main analysis averaging income over many years (see Table ??). In particular, $v(\mu)$ is larger, increasing the magnitudes of the first two bias components. Thanks to the fact that they compensate each other, the overall bias remains small in magnitude. In turn, IGE estimates are much more unstable, which is an additional reason for focusing on rank based measures.

little impact on our results.

1.6.4 Alternative mobility measures

Figures 1.4a-1.4b show how our main results change when ranking parents and children based on the “productivity” and “neighborhood-based” measures described in Section 1.5.2. The productivity-based curve, which relies only on precisely measured formal labor income, yields a rank-rank slope of .44, somewhat flatter relative to our baseline curve. This is the case because the minimum wage is binding in the bottom part of the distribution and because omitting capital income contributes to flattening the curve in the upper half of the distribution. Although this measure excludes individuals who have never held formal jobs in our sample, it offers some additional support to our main finding of high intergenerational income persistence in Brazil.

Figure 1.4: Alternative Measures



Notes: The figure plots mobility curves based on the productivity-based ranking (a) and neighborhood-based ranking (b), along with our baseline mobility curve. The productivity-based measure is based on the average formal labor income for parents and children in periods when they hold formal jobs. The neighborhood-based measure is based on the average formal income in the census tract in which children grew up (parental rank) and where they live as adults (child rank). Section 1.6.4 provides a detailed description of these measures. For each curve, the figure also displays the estimated β coefficient in Equation (1.1).


In turn, the neighborhood-based curve has a slope of .57, and it is steeper in the top quartile of the parental income distribution. This is consistent with the intuition that ranking individuals on income may underestimate the high living standards of children raised in affluent families, who may enjoy amenities and transfers beyond the income that they produce.³² Although it may differ to some extent in nature to our main income measure, it also supports the notion that there is high intergenerational persistence in socioeconomic status in Brazil.

³²To the extent that individuals born in a given place may develop preferences for that area, the neighborhood-based measure could underestimate mobility since such preferences will mechanically create persistence in our analysis. Although such selection is endogenous and should be interpreted with some caution, Appendix Figure A.7 shows that dropping children who did not change area flattens the neighborhood-based measure, but the rank-rank slope remains as high as 0.48 and continues to show strong persistence in the top 20% of the distribution.

1.6.5 Cross-country comparisons

Although comparisons of IGM estimates across countries must be interpreted with caution (BRATBERG et al., 2017; HECKMAN; LANDERSØ, 2021), social mobility in Brazil appears to be much lower than in any other country for which similar estimates are available. In particular, the rank-rank slope is estimated at 0.34 in the US (CHETTY et al., 2014), and ranges from 0.19 to 0.24 in Australia, Canada, France, Italy, and Scandinavian countries (ABBAS; SICSIC, 2022; ACCIARI; POLO; VIOLANTE, 2021; BRATBERG et al., 2017; CONNOLLY; CORAK; HAECK, 2019; DEUTSCHER; MAZUMDER, 2020; HEIDRICH, 2017; HELSØ, 2021). Figure 1.5 plots these estimates against the Gini index of income inequality. Interestingly, both intergenerational persistence of income and inequality are much higher in Brazil than the other (richer) countries, and Brazil lies perfectly on the Great Gatsby Curve depicted by other countries.³³

Figure 1.5: The Great Gatsby Curve



Notes:


Notes: The figure plots the relationship between the Gini index (horizontal axis) and relative mobility (vertical axis) using this paper's estimates for Brazil and available rank-based mobility estimates for developed countries. The latter are obtained from (DEUTSCHER; MAZUMDER, 2020) (Australia), (CONNOLLY; CORAK; HAECK, 2019) (Canada), (HELSØ, 2021) (Denmark), (ACCIARI; POLO; VIOLANTE, 2021) (Italy), (BRATBERG et al., 2017) (Norway), (HEIDRICH, 2017) (Sweden), and (CHETTY et al., 2014) (US).

Absolute upward mobility is also lower, as below-median income children reach an income rank around 6 percentiles lower in Brazil than in the US. We reach similar conclusions when comparing the full mobility matrix across income quintiles. For instance, children born in the bottom income quintile in Brazil have only a 2.5% chance of reaching the top quintile, while the same figure is three times larger in the US (7.5%) and 4-6 times larger in Italy (11.2%) and Sweden (15.7%).

³³See (CORAK, 2013) for a discussion on the factors driving the relationship between inequality and income mobility. In line with previous evidence for the US and Italy (see, respectively, (CHETTY et al., 2014; ACCIARI; POLO; VIOLANTE, 2021)), we also document a within-country Great Gatsby Curve, as mobility is inversely correlated with income inequality across Brazilian areas.

The stark contrast between Brazil and developed countries is also evident when we turn to the intergenerational income elasticity (i.e., the log-log relationship between parent and child income) as an alternative measure of income persistence (Figure 1.6). We estimate an IGE coefficient of .50, significantly larger than the estimates available for high-income countries, e.g., (CHETTY et al., 2014) finds .34 for the US and (ACCIARI; POLO; VIOLANTE, 2021) .23 for Italy. (DUNN, 2007) finds an even larger IGE of .69 in Brazil using survey data and instrumenting parental income by education. Such a higher estimate may reflect – among other things – the fact that parental education increases child income through other mechanisms beyond parental income.

Figure 1.6: Intergenerational Income Elasticity (IGE)




Notes: The figure plots the relationship between child and parental log income for our main sample (1988-1990 cohorts). For each level of log parental income (100 bins), it plots the mean log child income during 2015-2019, at the age of 25-31. It also reports the estimated IGE slope across all individuals and when restricting to parents between the 10th and 90th income percentiles. The vertical dashed line marks the median income in the parental income distribution.

1.6.6 Mobility by gender and race

Opportunities depend not only on parental income but also race and gender, especially in a country characterized by strong segregation such as Brazil. Figures 1.7a and 1.7b show the gender- and race-specific mobility curves, respectively. Importantly, the ranks on both axes indicate the positions relative to all individuals within the same cohort (rather than separately by gender and race), so the graphs show between-group differences in child ranks keeping constant parental income.

Female children's income is on average 14 percentiles below males with the same parental income. This gap largely reflects gender differences in labor market participation and wages (Appendix Table A.9). The mobility gap is virtually identical when restricting the same comparison to siblings, whereas the gap between siblings unconditional on gender is near zero (Appendix Table A.10). Interestingly, the rank-rank slope is steeper for females than for males (.61 vs. .47).

Figure 1.7: Mobility Curves by Gender and Race



Notes: This figure plots separate mobility curves by gender (a) and race (b), for our main sample (1988-1990 cohorts). For each parental income percentile, it plots the mean child income rank in ages 25-31. Parental income is the sum of the father's and mother's average income when the child is aged 3-18 years old. The ranks on both axes indicate the income positions relative to all individuals within the same cohort (rather than separately by gender and race). For each curve, the figure also displays our relative mobility measure based on Equation (1.1), the between-group gap conditional on having parents at the 25th and 75th income percentiles, and the average between-group gap across parental income percentiles.


Consequently, the gender gap declines from 17 to 10 percentiles when moving from the 25th to the 75th percentile of the parental income distribution.

Turning to differences by race, Figure 1.7b shows that non-white children rank on average 7 percentiles lower than white children with the same parental income. Race-specific transition matrices show that non-whites born in the first income quintile are much more likely to remain at the bottom (52.8% vs. 33.7%) and less likely to climb to the top (2% vs. 3.4%) compared to white children (Figure 1.8). Although differences are strongly reduced at the top of the distribution, non-whites born in the top quintile are twice as likely to fall to the bottom relative to white children (5.7% vs. 2.8%). The large mobility gap is remarkable given that non-whites – mainly comprising black and mixed-race individuals – are far from a minority in Brazil, representing about half of the population. The gap by race in Brazil is similar to the black-white gap in the US, where the former group is a minority.

1.6.7 Parental income and children's long-term outcomes

Next, we show that parental income is associated with improvement in several other long-term outcomes. Figure 1.9 shows how parental income relates to a wide array of children's outcomes other than income. Figure 1.9a plots college attainment over parental income ventiles, showing that it is convex over income: while children in the bottom ventile have almost no chances of completing college, roughly 80% of children in the upper ventile do so. Girls exhibit higher educational attainment than boys over the entire parental income distribution, yet they experience lower income later in life (Figure 1.7a). Children in higher-income families are

Figure 1.8: Racial transition matrix



Notes: The figure shows the transition probability matrix by quintiles of the income distribution for the 1988-1990 cohorts separately for each race group. Each cell displays the share of children born in that parental income quintile (horizontal axis) who end up in a given income quintile in adulthood (vertical axis). Income quintiles in both axes indicate the income positions relative to all individuals in their own cohorts (rather than each group). Cells are colored according to the quintile-quintile transition probability, with darker red tones indicating higher likelihoods.


disproportionately more likely to hold prestigious occupations – such as doctors and lawyers – and this relationship is highly convex at the top (Figure 1.9b).

Figures 2.1b, 1.9d and 1.9e show that low parental income is also strongly associated with markers of socioeconomic struggle. Children born to below-median income families are four times more likely to receive conditional cash transfers (*Bolsa Família*), five times more likely to become teenage mothers, and twice as likely to be the victim of a crime leading to hospitalization compared to richer children.³⁴ Finally, low parental income is associated with early mortality (Figure 1.9f): children in low-income families are up to three times more likely to die before they turn 30.

These results suggest that income persistence may be explained (or amplified) by gaps in educational achievement and other factors that emerge early in life such as teenage fertility. The fact that all children outcomes are correlated in the expected direction with parental income and that most of these relationships are smooth bolsters our estimates of the rank-rank curve.

³⁴We measure victimization as the probability of hospitalization due to an assault.

Figure 1.9: Long-Term Outcomes



Notes: This figure plots the relationship between parental income, measured when children are aged 3-18, and several children long-term outcomes in adulthood: college degree attainment (a), the probability of working as a doctor or lawyer (b), the likelihood of receiving *Bolsa Família* transfers when adult (c), teenage pregnancy rates (d), the probability of being hospitalized due to violent assault (e), and mortality rates (f).

1.7 Geographic variation in mobility

1.7.1 Geographical units and IGM measures

Brazil exhibits extreme variability in local socioeconomic conditions. We investigate social mobility across the 510 “immediate geographic regions” (IGRs), which are aggregations of neighboring municipalities sharing the same urban network and a common local hub (similar to the US commuting zones).³⁵ We assign children to the area where they grew up, which we proxy by their father’s place of residence (or, when the latter is missing, the mother’s) in 2000, i.e. when children in our sample were aged 10-12.³⁶ Like in the main analysis, we rank parents and children relative to the national income distribution.

1.7.2 Regional mobility patterns

The rank-rank relationship between parental and child income remains linear within regions – see, e.g., the plots for Belo Horizonte and Fortaleza, two of the largest metropolitan areas in the country, in Appendix Figure A.8. Therefore, we can compare mobility between regions using the measures of relative and absolute mobility introduced in Section 1.3, which rely on such linearity.

Figure 3.1a visualizes spatial variation in absolute mobility across IGRs. The map highlights three striking patterns. The first pattern is that absolute mobility strongly varies across regions, with the expected rank of below-median income children ranging between the 10th and the 51st percentile. More developed areas in the Center-South display higher upward mobility relative to the less affluent North and Northeast regions. A natural concern is that this map reflects different costs of living across regions. In Appendix Figure A.4.2, we show that adjusting for prices does not alter the main patterns in the map.³⁷

The second striking pattern is that several regions in the countryside display higher absolute mobility than large and rich metropolitan areas such as São Paulo and Rio de Janeiro. By contrast, children in high-income families in these two areas achieve excellent outcomes – see Table A.11 reporting mobility estimates for the 50 largest metropolitan areas of the country.

The third pattern is that the top 5% areas in terms of absolute mobility are all concentrated in a large mobility hotspot crossing three southern states: Paraná, Santa Catarina, and Rio Grande do Sul. This region has historically been characterized by the presence of agricultural communities established by European settlers maintaining a strong cultural heritage. In such regions, below-median income children reach on average the 47th percentile in adulthood and about 80% of children born in the bottom quintile escape poverty, transiting to higher income quintiles (see Appendix Table A.11).


Figure 1.11 documents a Great Gatsby curve within Brazil, even after controlling for variation in GDP per capita. Appendix Section A.4.4 presents an analysis of the factors that better explain

³⁵IGRs replaced the *microrregião* used in earlier studies on Brazil.

³⁶In our main sample, the father and mother live in the same IGR in 83% of cases.

³⁷The correlation between baseline and price-adjusted estimates of both absolute and relative mobility across regions is above .9. This high correlation is explained by the fact that, although prices significantly vary across regions, most children live in the same area where they grew up (or in areas with similar price levels).

Figure 1.10: Absolute Mobility Map: Predicted Rank for a Below-Median Income Child




Notes: The figure visualizes spatial variation absolute mobility (in deciles) across Brazil's 510 immediate geographical regions (IGRs) for our main sample (1988-1990). Parent and child incomes are ranked in the national income distribution and measured when children are aged 3-18 and 25-31, respectively. Absolute mobility indicates the expected rank for children in below-median income families, based on Equation (1.1). Darker green tones indicate higher absolute mobility. Children are assigned to IGRs according to the location of their fathers in 2000.

the substantial regional variation in mobility. Although entirely correlational, this analysis may inform future work aimed at understanding the causal determinants of upward mobility. Interestingly, we find that factors related to the quality of education provision yield by far the highest explanatory power on absolute mobility across IGRs, followed by indicators related to family structure, demographics (including the racial composition), household characteristics, and the local infrastructure. Although there is some overlap with the main mobility predictors found by (CHETTY et al., 2014) and (ACCIARI; POLO; VIOLANTE, 2021) for the US and Italy, in Brazil the quality of education stands out as the strongest factor.

1.7.3 Income mobility and educational mobility

Figure 1.12 shows the relationship between income mobility and educational mobility across regions. Following (ALESINA et al., 2021), we compute upward mobility as the likelihood that a child born to parents who did not complete primary school manages to do so; similarly, we compute downward mobility as the likelihood that a child born to parents who completed primary school fails to achieve the same level of education. These measures of mobility have

Figure 1.11: The Great Gatsby Curve across Brazilian regions



Notes: The figure plots the relationship between income inequality measured by the Gini index in 2000 (horizontal axis) and absolute mobility (vertical axis) across Brazilian regions. Both variables are residualized with respect to GDP per capita in 2002. The series in blue (dots) displays regions in the Center-South of the country and the series in red (triangles) displays regions in the North-Northeast. The figure also reports the correlation coefficient between the two (residualized) variables.

the advantage of being available for a much larger number of countries than income-based measures, including many developing economies. (ALESINA et al., 2021) estimate educational mobility across 2,800 regions in 27 African countries; the two graphs in Figure 1.12 also plot their estimates for some of these countries.


Three striking patterns emerge. First, the stark contrast between the North-Northeast and the Center-South of Brazil emerges for both upward and downward mobility. Second, educational mobility varies widely across Brazilian regions and, overall, it is comparable to that observed in some of the most mobile African countries – Nigeria, Botswana, and South Africa.³⁸ Finally, although income and educational mobility are strongly correlated with each other, there is a large amount of variation in income mobility for given levels of educational mobility, which further motivates the use of income-based measures.

1.8 Causal place effects

Motivated by the stark regional disparities in IGM documented in the previous section, we next estimate the causal effect of the place where children grew up on their perspectives of upward mobility. To disentangle such effect from sorting, we compare migrant children (or siblings) who moved to new areas at different ages (CHETTY; HENDREN, 2018a).

³⁸We achieve a similar conclusion when comparing (SAAVEDRA; ANDRES, 2022b) estimates for Latin American countries.

Figure 1.12: Educational mobility across Brazilian regions



Notes: The figure plots estimates of upward (a) and downward (b) educational mobility across Brazilian regions (horizontal axis) versus baseline regional absolute mobility measures estimated in Section 1.7. Blue dots (red triangles) indicate regions in the Center-South (North-Northeast) region of Brazil. Vertical lines mark estimates of educational mobility for selected African countries from (ALESINA et al., 2021). Upward (downward) mobility is the likelihood of a child born to parents who did not (did) complete primary school succeeding (failing) to do so. The figure also reports the average upward (downward) educational mobility in Brazil and the cross-regional correlation between educational mobility and income mobility.

1.8.1 Data and research design

For this analysis, we use a sample that covers all children born during the 1983-1992 period that can be linked to their fathers. We distinguish between permanent residents and movers based on parents' residency in the 1992-2019 period. Like in Section 1.7, the geographical unit of analysis is the IGR. We track moves using formal employment data, because address coverage prior to 2000 is low in the person registry (see Appendix A.5.1 for details).

Our empirical strategy and specifications closely follow (CHETTY; HENDREN, 2018a) (see also (DEUTSCHER, 2020) for an application to Australian data). We first characterize the predicted outcomes of permanent residents using rank-rank regressions for each cohort and region (see Appendix A.5.2 for additional details). We then use these estimates to compute the predicted rank difference for each mover based on the origin and destination region, the child's cohort, and parental income rank. Finally, we estimate causal place effects by relating movers' income rank at the age of 24 to their predicted difference in ranks across children moving at different ages.³⁹ Intuitively, to the extent that location exert causal effects, movers' outcomes should display greater convergence to that of permanent residents the earlier they move (and the longer they are exposed) to the destination place. Specifically, our main analysis is based on the following equation:

$$y_i = \alpha_{ocpa} + \sum_{a=1}^{33} b_a I_a(a_i = a) \Delta_{odpc} + \sum_{c=1983}^{1991} \kappa_c I_c(c_i = c) \Delta_{odpc} + \varepsilon_i, \quad (1.3)$$

³⁹Like (CHETTY; HENDREN, 2018a), we focus on income at an earlier age relative to our main analysis (Section 1.6), so that we can measure income for (older) cohorts who move at older ages.

where y_i is the child's income rank at the age of 24; α_{ocpa} is a fixed effect by origin o , cohort c , parental income decile p , and age at move a ; I_a and I_c are indicators for each age at move a and cohort c ; and Δ_{odpc} is the difference in permanent residents' predicted outcomes between origin o and destination d for parental income decile p and cohort c . The coefficients of interest b_a for $a \leq 24$ give the expected increase in rank associated with moving at age a to a destination with a 1 percentile higher predicted rank. Since we measure income at 24, moving at an older age cannot possibly have a causal effect on income, so b_a for $a > 24$ captures solely selection effects. The coefficients κ_c control for our varying ability to track moves across cohorts, ensuring that we only use within-cohort variation in the age at move.⁴⁰ One advantage relative to (CHETTY; HENDREN, 2018a) is that we track moves from age 1 (instead of 11). This implies that we can flexibly study how convergence varies by age from early life, without the need to rely on linear extrapolations.

Our econometric specification effectively compares the extent of convergence to destination outcomes by children who move at different ages. The key identifying assumption is that selection effects driving some children to move to better (or worse) areas are orthogonal to the child's age when they move. Importantly, fixed effects α_{ocpa} ensure that the estimation of convergence coefficients for moves at each age exclusively relies on variation between children who have the same parental income background, and belong to the same cohort and place of origin. We provide several pieces of evidence that strongly support our main identification assumption in the next Section 1.8.2.

1.8.2 Results


Causal place effects are summarized in Figure 1.13. The series in blue displays the estimated coefficients b_a on the predicted difference in outcomes for each child given her origin, destination, cohort, and age at move. Supporting the intuition behind our identification strategy, the extent of convergence decreases with age, following a roughly linear pattern: children who move earlier to better places benefit more. The positive coefficients from the age of 25 purely reflect positive selection into migration, as child income measured at the age of 24 – by construction – cannot be affected by future moves. The flat pattern from the age of 24 supports our identifying assumption that parental selection into migration to a given destination does not vary with age.

Since exposure effects decline linearly with age at move, we substitute the coefficients b_a with a linear counterpart to estimate an average convergence by year of exposure.⁴¹ Each additional year of exposure to the destination area increases convergence in children's outcomes by .024 (baseline slope), meaning that children moving at birth to a place where they are expected to rank 10 percentiles higher will increase their rank in the national income distribution

⁴⁰For instance, since we track parents' locations starting from 1993, we observe children born in 1983 moving from 10 years old onwards, while children born in 1992 since the age of 1. To avoid collinearity, we omit the indicator for the 1992 cohort.

⁴¹Specifically, we substitute the non-parametric term $\sum_{a=1}^{33} b_a I_a(a_i = a) \Delta_{odpc}$ in eq. 1.3 with a linear counterpart $I(a_i \leq 24)(b_0 + (24 - a_i)\gamma) \Delta_{odpc} + I(a_i > 24)(\delta + (24 - a_i)\delta') \Delta_{odpc}$. The linear term is split at the age of 24, since moves above this age capture selection effects. γ is the main coefficient of interest identifying convergence by year of exposure. We repeat the same procedure to summarize convergence for an alternative specification used in the robustness analyses.

Figure 1.13: Exposure Effect Estimates for Children's Income Rank in Adulthood



Notes: This figure plots the estimated b_a coefficients in equation (1.3) (blue dots) and in an alternative specification including family fixed effects in Appendix equation (A.3) (red triangles). The sample includes all father-linked children from the 1983-1992 cohorts whose father moved once between 1993-2019, and the dependent variable – child income – is measured at the age of 24 (dashed vertical line). Vertical bars report 95% confidence intervals. Each coefficient b_a indicates the degree of convergence to the outcomes of permanent residents in the destination relative to those at the origin. Coefficients b_a for moves until the age of 24 estimate causal place effects, while b_a coefficients for older ages estimate selection effects as moves when aged 25 or older cannot explain income at age 24. The slope of the blue and red lines, as estimated by linear approximations for b_a in equations (1.3) and (A.3), summarize convergence per year of exposure.

by $.024 \times 24 \times 10 = 5.76$ percentiles on average due to causal place effects.⁴²

These estimates suggest that about 57% of the substantial mobility gap across Brazilian regions (Section 1.7) is due to causal place effects. Hence, some areas in Brazil offer significantly better opportunities for low-income children, notwithstanding the high levels of income persistence at the national level.

Importantly, Figure 1.13 shows that our results are virtually unaffected by the addition of family fixed effects, thus relying exclusively on within-family variation (Appendix A.5.3 provides the details on this specification). The latter rules out the possibility that our findings are driven by family selection over child age at move. Our results are also robust to overidentification tests – see Appendix A.5.4. They follow from the intuition that children's outcomes should converge to the average outcomes of their own group, whereas predicted outcomes of other groups are used as placebos.

⁴²Focusing on educational mobility and moves between 1-11 years old, (ALESINA et al., 2021) find yearly exposure effects of 0.03 in Africa while (SAAVEDRA; ANDRES, 2022a) report 0.035 for Latin America.

1.9 Conclusion

In this paper, we provide the first estimates of income intergenerational mobility using large-scale tax data for a developing country, while addressing in details income measurement issues related to the informal economy. We find that income mobility in Brazil is much lower than comparable estimates available only for developed countries. Moreover, we uncover wide disparities across areas, genders, and racial groups, depicting a “*land of inequality*” in which children’s opportunities are deeply dependent on their parents’ socioeconomic status.

Importantly, we develop new methods for imputing unobserved income and studying the consequences of measurement error for IGM measures, in addition to providing alternative measures for ranking individuals on socioeconomic status. The same methods could be applied to estimate income and social mobility in other contexts characterized by a large unofficial sector. This is typically the case in low- and middle-income countries, but it is also relevant in several high income contexts [e.g., see]medina2018shadow. More generally, these methods may find application in any study where the underground economy is a challenge for income measurement.

This work is also relevant in public and policy debates. Even though Brazil has long been perceived as a place of high inequality and low mobility, hard evidence on IGM may contribute to shifting people’s perceptions and potentially their preferences for distributive policies (ALESINA; STANTCHEVA; TESO, 2018). Moreover, revealing dramatic penalties for long-neglected groups and places – in particular, non-whites and the North-Northeast of the country – can encourage public policies targeted at increasing access to opportunities. In particular, our results on causal place effects and drivers of mobility across regions can motivate place-based policies aimed at improving the quality of public education provision in the poorest areas of Brazil.

Do CCTs Create Conditions to Thrive?

Bolsa Família and Social Mobility in Brazil

2.1 Introduction

Conditional Cash Transfers (CCTs) are the backbone of poverty-reduction policies in many countries, particularly in the developing world (World Bank, 2018). By design, CCTs have two main goals. In the short-term, cash transfers to poor households aim to ease financial constraints and alleviate poverty immediately. By conditioning such transfers to children’s education and health checks, CCTs’ long-term purpose is to enable sustained social mobility, i.e., break the intergenerational cycle of poverty. While several studies document that CCTs are effective in the short-term (FISZBEIN; SCHADY; FERREIRA, 2009; BASTAGLI et al., 2016), whether short- and medium-term effects on poor children translate into better living conditions in adulthood remains an open question (MILLÁN et al., 2019; GARCIA; SAAVEDRA, 2023).

In this paper, we estimate the long-term effects of the largest CCT program in the world, Brazil’s *Programa Bolsa Família* (PBF). Created in 2004, PBF is the main welfare program in the country, assisting over 14 million households as of 2019. Many studies show that PBF had positive short-term impacts on poor children’s education and health outcomes (VIANA; KAWAUCHI; BARBOSA, 2018). Still, the main purpose of the program was to offer conditions for these kids to move away from economic hardship in adulthood. Thus, we combine different administrative datasets tracking adult outcomes for over 20 million people born between 1970–1994 to ask whether individuals exposed to PBF during childhood display improved socio-economic conditions later on. Such large-scale assessment of long-term impacts is still rare in the literature and key for policy-making in both developing and developed countries (BLATTMAN et al., 2017; GENTILINI et al., 2020).

Similar to other analyses of intergenerational effects of the social safety net (GOODMAN-BACON, 2021b; BAILEY et al., 2023), we employ a differences-in-differences design in which temporal variation comes from cohorts born earlier versus later relative to PBF implementation. We flexibly assess PBF’s effects by age at implementation with event studies in which control cohorts have adult outcomes effectively determined when the program began (ages 25–34 in 2004). Using mother fixed effects, we compare children young enough to benefit from PBF with their older siblings in families with low versus high exposure to the program. In particular, we measure exposure using a Machine Learning (ML) model that identifies households’ likelihood to participate in the program given pre-determined characteristics.¹ Finally, to account for

¹This approach is akin to (CENGIZ et al., 2022), in which the authors train an ML model to identify workers more likely to be exposed to a minimum wage increase.

convergence between regions and the expansion of Brazil's welfare state in the period, all baseline regressions include city-by-cohort fixed effects and directly control for children's exposure to the roll-out of social programs contemporary to PBF.

Our main result is that PBF promoted human capital accumulation for its beneficiaries, leading to intergenerational income mobility among the next generation. We first document that boys and girls exposed to the program's conditionalities between 10–15 years old attain 1.3 and 1.9 more years of education, respectively – implying gains of 14% and 19% relative to their baseline counterfactuals.² This higher educational attainment is reflected in better chances of completing high school – around 9 p.p. (15%) for boys and 12 p.p. (21%) for girls – but remarkably is a substantial increase in the probability of obtaining a college degree, which more than doubles for both genders. Notably, we also find small gains in attainment – around 0.5 years, or less than 5% – for young adults aged 16–21 at the start of the program. However, such gains are fully concentrated on higher high school completion rates, with no improvements in the obtainment of a college degree.

These heterogeneous effects across cohorts are likely related to PBF's design, as only children who are 15 years or younger are directly impacted by the program's conditionalities.³ Indirect effects on older children (16–21) may arise from the relatively large household-level income shock promoted by the cash transfer – which potentially reduces the opportunity cost of young adults remaining in school but is unlikely to foster meaningful changes in skill development (HECKMAN; MOSSO, 2014).⁴ In turn, the more significant effects on young children are probably linked to a lengthened exposure to PBF's conditionalities from an earlier age, which can improve conditions related to human capital acquisition.

In line with this notion, we next show that PBF is also related to higher chances of working in the formal labor market – but increased labor supply is associated with higher earnings only for children impacted by the program's conditionalities. In particular, boys and girls aged 10–15 at the start of the program are, respectively, 8 p.p. (14%) and 14 p.p. (40%) more likely to hold at least one formal job between 23–25 years old, while the effect for their older siblings is around 5 p.p. (8–14%). In turn, effects on earnings only emerge for younger children and amount to around 2,300 BRL/year (30%) for males and 1,450 BRL/year (40%) for females. Consistent with higher earnings, children who grew up in families which received PBF transfers are 4 p.p. less likely to participate in the program as adults, or a 30% drop relative to their baseline counterfactual. This result is an important indication that PBF promoted social mobility and contradicts the common-held argument against CCTs that its beneficiaries become “dependent” or “lazy” in the long-term.

Nevertheless, even with higher earnings and lower dependency on social programs, bene-

²In line with the differences-in-differences design, we compute counterfactuals for exposed children's outcomes by adjusting the baseline of non-exposed cohorts in the treated group by the relative change across cohorts in the control group.

³This is particularly true in our setting, which analyses effects on children impacted in the first years of the program. In 2008, a slight reform introduced conditionalities for children aged 16–17 years old, but children in our sample potentially affected by it were already younger than 15 in 2004.

⁴Despite high school completion in Brazil is expected at 18 years old, significant levels of grade distortion – especially among the poor – can help explain the effects on young adults in their early 20s. Moreover, the possibility of seeking supplementary education later in life also explains part of the observed effects.

ficiary children could still grow up to be at the bottom of the income distribution. To assess this, we replicate the methodology in (BRITTO et al., 2022) to account for informal income and construct cohort- and gender-specific income distributions. We then show that PBF is associated with intergenerational income mobility, but again only for children directly impacted by its conditionalities. Boys and girls exposed to the program rank, respectively, 6.3 and 4.3 percentiles higher in their respective income distributions – corresponding to relative gains of nearly 15%. Besides impacting mean income ranks, PBF also affects probabilities of being in the tails of the income distribution. For instance, boys and girls are, respectively, 10.4 p.p. (15%) and 14 p.p. (23%) more likely to be outside the poorest 20% and 5.8 p.p. (69%) and 5.2 p.p. (87%) more likely to be among the richest 20%. This set of results underscores that PBF promoted both the exit from poverty and extreme upward mobility, succeeding in its long-term goal of enabling sustained social mobility.

Furthermore, we show that PBF is related to significant reductions in migration rates across cohorts, which relates to anecdotal and empirical evidence of welfare programs lessening the need for “desperate moves” among the extremely poor in Brazil. We also find that girls exposed to the program early on observe significant decreases in teenage pregnancy rates – around 2 p.p., or 20% – which may point to an important mechanism behind social mobility, as teenage pregnancy is pervasive among vulnerable households in Brazil and has lasting consequences for girls. Finally, PBF marginally improves survival rates for children and young adults exposed to the program, which is potentially linked to better health outcomes overall and reduced engagement in risky behavior.

Across all outcomes, we find that relative effects are consistently larger for girls. In additional heterogeneity analysis, we also document that relative effects are larger for non-whites and children growing up in the North/Northeast regions of the country. Since these groups are relatively more disadvantaged, this goes in hand with the focused approach of PBF, which aims at targeting the most vulnerable groups in society. Likewise, PBF can be (cautiously) interpreted as reducing racial and regional disparities in income mobility, which speaks to broad evidence that the program contributed to the reduction in income inequality in Brazil between 2000–2010 (SOUZA et al., 2019).

Moreover, we show that PBF effects are higher in places with above-median quality of the public educational system and for children whose mothers are relatively more educated – suggesting that both place and family inputs are relevant to the effectiveness of the program. This suggests that the availability of opportunities in a given place mediates CCTs’ capability to promote children’s social mobility and that it does not replace important family inputs in the process of child development.

In sum, we document that PBF impacted a wide range of long-term outcomes. While the household-level income shock promoted by cash transfers was able to increase school completion rates and formal labor supply of young adults in exposed families, younger children directly impacted by the program’s conditionalities also observed lasting improvements in earnings and income mobility. Our results are an important empirical support for CCTs’ design, showing that conditionalities can be successful in achieving the long-term goal of breaking the intergenerational cycle of poverty, while immediate cash transfers can have positive indirect effects on other members of the household. To the best of our knowledge, this paper conducts

the first robust in-depth empirical analysis of the intergenerational impact of a CCT program – and does so for the largest CCT in the world, PBF.

Our work makes contributions to three strains of the literature. First, to the wide set of studies on the impacts of CCTs reviewed by (FISZBEIN; SCHADY; FERREIRA, 2009; BASTAGLI et al., 2016; MILLÁN et al., 2019; GARCIA; SAAVEDRA, 2023). In particular, we add robust evidence to a growing literature showing that CCTs promote sustained education gains for poor children (BARHAM; MACOURS; MALUCCIO, 2018; MILLÁN et al., 2020; ARAUJO; MACOURS, 2021; BARRERA-OSORIO; LINDEN; SAAVEDRA, 2019; CAHYADI et al., 2020; ATTANASIO et al., 2021), deepening the analysis on a longer time horizon and across more outcomes. In this regard, the most similar paper to ours is (PARKER; VOGL, 2021) long-term follow-up of Mexico's *Progreso*, which has a more limited scope.⁵ Importantly, we are the first to combine an analysis of CCTs and carefully computed intergenerational mobility measures, which is crucial to evaluate the long-term performance of such programs relative to their goals.

Furthermore, our results relate to the rich literature on child development documenting that skill formation is an intricate process that impacts several life outcomes (HECKMAN; MOSSO, 2014). In particular, our results dovetail nicely with the finding that investing in skill formation in disadvantaged children turns them more productive in the future. In the context of PBF, we show that this leads to upward social mobility. Finally, we also speak to studies analyzing the intergenerational effects of the social safety net, which have focused on the U.S. context (BAILEY et al., 2023; BARR; EGGLESTON; SMITH, 2022; GOODMAN-BACON, 2021b).

The paper is structured as follows: Section 2.2 gives details on PBF's institutional background and the expected long-term effects of the program. In Section 2.3, we introduce the data and describe our empirical strategy. Our main results of PBF effects on adult outcomes are discussed in Section 2.4. Finally, Section 2.5 concludes.

2.2 Institutional Background

2.2.1 The *Programa Bolsa Família* (PBF)

Created in 2004, PBF is the largest CCT program in the world. As of 2019, it reached 15 million families – roughly 30% of the population – with an average stipend of 180 BRL per family, costing 0.5% of the GDP annually. It is broadly considered a successful program among policy-makers, and contributed to the spread of CCTs around the developing world.

Under PBF, every family with per capita income below a national poverty threshold is eligible for a variable benefit that depends on the number and age of the children. As of 2019, this threshold was 140 BRL and families could accumulate up to five children (0–15 years old) and two youth (16–17 years old) monthly benefits of 32 and 38 BRL each, respectively.⁶

⁵In particular, the use of a municipal-level variation on Census data with high attrition rates complicates (PARKER; VOGL, 2021) identification, which also relies on smaller samples and fewer outcomes.

⁶At implementation, families could only receive up to two child benefits. The youth benefit and the higher child limit were implemented in 2008. Monetary values of transfers and the poverty threshold are regularly adjusted for inflation.

Transfers are conditional on children's school attendance and vaccination. In addition to the variable benefit, families with per capita income less than half of the poverty threshold receive an unconditional monthly benefit of 70 BRL. Eligible households can self-enroll at a social assistance center or be prospected to the program by a social assistance team. Transfers are preferentially entitled to the mother of the household and conditionalities are enforced jointly by municipalities and the federal government.

2.2.2 PBF in the Short-run and Expected Long-term Effects

The design of CCTs relies on some key assumptions (GARCIA; SAAVEDRA, 2023). Foremost is that enhancing poor children's human capital will lead to improved socioeconomic well-being in adulthood. This is related to a rich literature showing that cognitive and non-cognitive skills developed during childhood are highly associated with a wide range of adult outcomes (HECKMAN; MOSSO, 2014). Moreover, CCTs consider that human capital accumulation is a function of time spent in school and, importantly, these programs believe that they improve human capital acquisition by helping households overcome demand-side constraints such as educational externalities, informational constraints, and opportunity costs of children's time.

A large literature seems to indicate that PBF is successful in lessening such constraints, improving children's school enrollment and attendance rates in the short-run (VIANA; KAWAUCHI; BARBOSA, 2018). Furthermore, these studies underscore that PBF also improves children's health outcomes and overall household conditions, which could also benefit the process of human capital acquisition. Therefore, we aim to test whether the alleviation of constraints indeed leads to higher human capital and improved adult outcomes for exposed children, especially on earnings and social mobility – which would largely support the program's design.

2.3 Data and Methodology

2.3.1 Data

We combine several administrative records to perform our analysis. Our sample of children is the Brazilian person registry (*Cadastro de Pessoa Física*, CPF), which covers the entire population and is provided by the Brazilian tax authority. All individuals are identified by their person code, full name, and mother's full name. We restrict the sample to children born between 1970 and 1994 (aged 10 to 34 years old when PBF started) and whose mothers are uniquely identified by their names.⁷ In this way, we can link them to their mothers' and retrieve their characteristics before the launch of PBF. Our final sample comprises over 23 million children, around half of the cohorts studied.

We track individuals' participation in PBF using administrative records on the program's payments from 2004 to 2020. We observe children's educational and labor market outcomes from administrative employment data covering the population of formal jobs for the 2002-2019 period (*Relação Anual de Informações Sociais*, RAIS) and entrepreneurship activity via the

⁷Due to accumulation of surnames, around 52% of Brazilians have a unique name. In (BRITTO; PINOTTI; SAMPAIO, 2022), it is shown that the uniquely-named population is representative of the overall population.

Brazilian firm registry (CNPJ). We observe fertility, migration, and location decisions from updates in CPF. Finally, we follow the same methodology in (BRITTO et al., 2022) to measure individual-level informal income and compute social mobility measures.⁸

2.3.2 Empirical Strategy

We expect PBF's long-term effects to be higher for children impacted at younger ages, while young adults impacted by the program after adult outcomes are determined could not be affected by it, by definition. Accordingly, we use a differences-in-differences setup in which temporal variation comes from cohorts born earlier versus later relative to PBF implementation. We measure family-level exposure to PBF and compare children young enough to benefit from the program with their older siblings (first difference) in families with low versus high probability of being targeted by the program (second difference).

We identify families' exposure to PBF by predicting their likelihood to participate in the program in 2004 using a Machine Learning (ML) algorithm, akin to (CENGIZ et al., 2022). The baseline model uses mothers' education, age, number and age of their children, formal employment background, and city of residence as features.⁹ Table 2.1 shows mothers' characteristics by deciles of their predicted "PBF score", showing that highly-exposed mothers are concentrated in the North and Northeastern regions of the country, have more children, and lower educational levels. In our baseline specification, the treatment group comprises children whose mothers are at the top 25% of the predicted-probabilities distribution – in line with participation rates as shown in Table reftab:mothers.¹⁰

Table 2.1: Mother's characteristics by deciles of predicted PBF score.

Decile	Age (2004)	High School	Formal Income	Lives in N/NE	Age 1 st childbirth	No. of kids	PBF transfers 2004-08 (BRL/y)	Children's PBF eleg.	PBF score
1	50.81	0.98	35754.06	0.07	26.52	1.93	0.31	0.00	0.00
2	50.02	0.85	7894.13	0.07	24.82	2.13	1.31	0.00	0.00
3	47.27	0.76	4163.67	0.13	24.48	2.20	3.09	0.01	0.01
4	45.29	0.64	3619.99	0.15	23.99	2.33	6.88	0.04	0.01
5	44.77	0.50	2897.17	0.19	23.57	2.43	15.08	0.09	0.02
6	44.32	0.42	2352.48	0.22	23.05	2.57	31.58	0.19	0.03
7	41.96	0.39	1636.16	0.27	22.61	2.55	63.21	0.48	0.06
8	39.73	0.28	1071.93	0.32	22.10	2.67	127.67	1.12	0.12
9	38.08	0.12	535.00	0.35	21.74	2.84	241.26	2.21	0.21
10	37.11	0.05	232.94	0.61	20.80	3.69	478.79	3.63	0.38

Notes. The table reports mothers' characteristics according to deciles of their predicted probability of participating in PBF, as computed by the ML model. In our baseline specification, the treatment group includes children whose mothers are in the top 25% of the predicted-probabilities distribution.

We estimate reduced-form event-study specifications to validate the research design and flexibly assess PBF's dynamic effects relative to age at the start of the program:

⁸A detailed description of the variables used as outcomes in the following sections can be found in Appendix Table B.1.

⁹All covariates are measured before the start of PBF. We use data only for the first year of the program to avoid endogeneity concerns about selection into the program and its roll-out.

¹⁰In Appendix Table B.5, we show that our results are robust to alternative treatment definitions.

$$y_{iamt} = \sum_{t=10}^{34} (I_t \times Treat_a) \mu_t + Treat_a \vec{X}_{mt}' \Omega + \alpha_a + \gamma_{m,t} + \varepsilon_{iamt} \quad (2.1)$$

where y_{iamt} is an adult outcome of individual i with mother a born in municipality m who was t years old when PBF began. The key independent variable is an interaction between the mother-level binary treatment defined above, $Treat_a$, and a set of age t indicators I_t ranging from 10 to 34 years old. We normalize $\mu_{28} = 0$ since adult outcomes are mostly determined “children” aged 28 when the program began. With mother fixed effects α_a , our design leverages only within-family variation, absorbing other forms of selection.

To rule out the possibility of capturing the effects of phenomena other than PBF, we add two more terms in Equation 2.1. First, city-by-cohort fixed effects ($\gamma_{m,t}$) account for convergence in outcomes across Brazil’s regions during the period. Still, even within a city-cohort group, there could be differences (other than PBF) between rich and poor households confounding adult outcomes, such as contemporaneous welfare policies. Hence, we interact our treatment indicator $Treat_a$ with a vector of city-by-cohort measures of exposure to other policies of the period, $\vec{X}_{mt}'$. Specifically, we control for the supply of secondary schools, the expansion of higher education establishments, and the roll-out of a major health program targeted at poor families. Finally, we cluster the error term ε_{imt} at the family level.

Hence, our identifying assumption is that, conditional on exposure to other policies, absent the creation of PBF trends in siblings differences between poorer versus richer households within the same municipality would remain the same. We show that this assumption is strongly supported by parallel trends for cohorts too old to benefit from PBF. If this design captures the effect of PBF, the intention-to-treat (ITT) estimates μ_t should be decreasing in age and indistinguishable from zero for cohorts too old to benefit from the program.

2.4 Results

2.4.1 PBF Long-term Effects


Figure 2.1 plots μ_t coefficients from Equation 2.1 for four key long-term outcomes. The absence of pre-trends for individuals impacted by the program after 24 years old and the larger effects for children who spent more time under PBF reinforce the credibility of our research design. Figures 2.1a–2.1d summarize our main findings: PBF promoted human capital accumulation for its beneficiaries, leading to reduced dependency on the social safety net, higher earnings, and intergenerational income mobility among the next generation.

Figure 2.1a shows ITT estimates for total years of education. High exposure to PBF is associated with gains in educational attainment ranging from less than 0.5 years for young adults aged 18–22 at the start of the program to up to 2 years for children aged 10 when PBF began. Notably, there is no discernible effect for individuals older than 24 years old, as they had likely completed their educational trajectory before the program’s implementation. Similarly, Figure 2.1b reveals that individuals exposed to PBF are less likely to be PBF beneficiaries in adulthood, indicating improved social mobility. Furthermore, Figure 2.1c shows that exposed children have higher earnings in the formal labor market, and Figure 2.1d document that they also achieve

higher positions at the income distribution – a clear sign of intergenerational mobility. Notably, the effects on earnings and income mobility depicted in Figures 2.1c and 2.1d only arise for children aged 15 or younger when the program started.

In summary, Figures 2.1a–2.1d underscore that PBF is associated with long-term improvements in human capital and socio-economic well-being, suggesting its effectiveness in promoting sustained social mobility. Nevertheless, the different age patterns across outcomes might indicate the presence of both income and substitution effects (GARCIA; SAAVEDRA, 2023). To assess this, we analyze adult outcomes across five dimensions: Human Capital, Labor Market, Intergenerational Mobility, Migration, and Health and Behavior.

Figure 2.1: PBF effects in the Long Run: ITT event studies



Notes. The figure shows ITT coefficients μ_t and 99% confidence intervals obtained from Equation 2.1. A detailed description of outcomes used as dependent variables is given in Appendix Table B.1.

Table 2.2 displays ITT estimates obtained from gender-specific regressions. We report coefficients from a modified version of Equation 2.1, replacing age indicators with indicators for being 10–15 years old in 2004, 16–21 years old, or older. This approach allows us to compare cohorts that experienced both the income shock and PBF’s conditionalities (10–15) to those only impacted by the cash transfer (16–21). From now on, we refer to these two groups as “fully-treated” and “partially-treated”, respectively. The young adults older than 21 serve as control cohorts, based on the graphical evidence from event studies in Figure 2.1.

2.4.1.1 Human Capital

Human Capital accumulation is key for CCTs, as the design of such programs relies on the assumption that allowing poor children to spend more time in school will foster the development of skills that enable sustained social mobility in the long-term. Panel A of Table 2.2 displays ITT estimates for years of education and the likelihoods of completing school (high school degree) and obtaining a college degree. Columns 1 and 3 of Table 2.2 reveal that fully-treated boys and girls attain 0.7 and 0.9 more years of education, respectively. In Appendix Table B.2 we compute counterfactuals for children by adjusting the baseline outcome in the treatment group by the relative change across cohorts in the control group. As children in the treatment group are around 50 percentage points (p.p.) more likely to participate in PBF, these estimates translate into educational gains of 1.3 years (14%) for boys and 1.9 years (19%) for girls after effects are inflated and normalized relative to the counterfactual.¹¹

These improvements in educational attainment are evident in higher school completion rates: 9 p.p. (15%) for boys and 12 p.p. (21%) for girls. Nevertheless, the most substantial effects are observed in college degree attainment, where estimates range from 9 to 11 p.p., implying relative increases of more than 100% across all groups. On the other hand, estimates for years of education are substantially smaller for partially-treated young adults (columns 2 and 4 in Table 2.2), implying effects smaller than 5%. For this age group, the greater educational attainment arises primarily from increased school completion rates, with effects around 4 p.p. (5–6%).¹²

The economically significant effects suggest that PBF successfully increased the human capital of young generations in vulnerable households. The observed heterogeneity across cohorts, in turn, hints to different types of shocks experienced by younger versus older children. Partially-treated cohorts experience a household-level income shock, potentially reducing the opportunity cost of staying in school or seeking additional education. However, this late income shock is unlikely to promote the development of new skills in these young adults, who grew up under unfavorable conditions.¹³ In contrast, fully-treated cohorts were exposed earlier and for longer to conditionalities that positively influenced educational and health outcomes related to the development of cognitive and non-cognitive skills, thus enhancing their ability to acquire human capital (VIANA; KAWAUCHI; BARBOSA, 2018; HECKMAN; MOSSO, 2014). This difference is reflected, for instance, in the steeper pattern for fully-treated kids in Figure 2.1a and the effects on college degree attainment in Table 2.2. Importantly, these heterogeneous effects on human capital accumulation might manifest in other dimensions, particularly in the labor market.

¹¹Henceforth in the paper, all absolute estimates refer to treatment on the treated (TOT) estimates obtained by inflating ITT estimates with participation rates. Also, estimates in relative terms refer to TOT estimates relative to counterfactuals listed in Appendix Table B.2.

¹²Our estimates of additional educational attainment are close to the 0.6–1.6 years reported by (PARKER; VOGL, 2021) in the context of Mexico's *Progreso* for similarly-treated cohorts. The authors also find increases in school completion rates and larger effects for girls, but no effect on tertiary enrollment. However, given the different contexts and identification strategies, we avoid drawing strong conclusions from these comparisons.

¹³In particular, there is a large literature documenting that interventions in late childhood are rarely effective in promoting skill development for disadvantaged young adults (HECKMAN; MOSSO, 2014).

2.4.1.2 Labor Market

Panel B, columns 1 and 3, of Table 2.2 reveals that fully-treated children are more likely to work in the formal labor market. The increase in the probability of holding at least one formal job between 23–25 years old is 8 p.p. (14%) for boys and 14 p.p. (40%) for girls. The effects for the partially-treated (columns 2 and 4) are also positive but smaller, at around 5 p.p. (8–14%). Importantly, this increased labor supply is associated with higher earnings only for children impacted by the program's conditionalities (see Figure 2.1c). The gains are economically significant, amounting to around 2,300 BRL/year (30%) and 1,450 BRL/year (40%) for males and females, respectively.

Additionally, PBF is associated with substantial increases in entrepreneurial activity. The effects on the probability of opening a firm range from 0.5 to 2 p.p. for partially-treated and up to 6 p.p. for fully-treated cohorts. Given the low baseline, these imply large relative effects of 50% to 100%.

2.4.1.3 Intergenerational Mobility

The ultimate goal of CCTs is sustained social mobility, i.e., breaking the intergenerational cycle of poverty. In the specific case of PBF, the program intended to “*allow the sustained emancipation of its participants*”. A straightforward manner to look into this is to check whether children in beneficiary households grow up to be adults that participate in PBF as heads of households themselves. As in Figure 2.1b, panel C of Table 2.2 documents significant effects for fully-treated children. They are around 4 p.p. less likely to be part of the program in adulthood, or a 30% drop.

The smaller likelihood of depending on social welfare is in line with the higher formal earnings documented before. Despite these gains, given the substantial income inequality in Brazil, it could still be the case that children who grew up under PBF experience zero income mobility, remaining at the bottom of the income distribution. To investigate this further, we use the methodology in (BRITTO et al., 2022) to account for informal income and create cohort- and gender-specific income distributions. As shown in Figure 2.1d, columns 1 and 3 of panel C in Table 2.2 demonstrate that fully-treated cohorts attain higher positions in their respective income distributions. The effects are around 6.3 (16%) and 4.3 (12%) percentiles for boys and girls, respectively. Besides impacting mean income ranks, PBF also affects probabilities of being in the tails of the income distribution. For instance, boys and girls are, respectively, 10.4 p.p. (15%) and 14 p.p. (23%) more likely to be outside the poorest 20% and 5.8 p.p. (69%) and 5.2 p.p. (87%) more likely to be among the richest 20%. Therefore, PBF promoted both the exit from poverty and extreme upward mobility, succeeding in its long-term goal of enabling sustained social mobility.

2.4.1.4 Migration

In Panel D of Table 2.2, we explore PBF's effects on the migration and location decisions of children and young adults. Remarkably, we find substantial reductions in the probability of migration in both cohorts. Fully-treated children are 7 p.p. (17%) less likely to move away from their hometown, while the drop for partially-treated cohorts is approximately 4.5 p.p.

(10%). These reductions are also pronounced when considering cross-state (2–3 p.p., 12.5%) and cross-regional (1–2 p.p., 10%) migrations. These findings align with existing empirical evidence, suggesting that welfare programs in Brazil can diminish the occurrence of “desperate” moves prompted by vulnerable socio-economic conditions.

Interestingly, for fully-treated children, reduced migration does not lead to a lower likelihood of living in a large urban area. On the contrary, this age group is on average around 4 p.p. (9%) more likely to reside in a city with more than 100k inhabitants. Conversely, partially-treated young adults exhibit a slightly lower propensity to live in large cities (1 p.p., 2%). This could be attributed to higher relocation costs for individuals with stronger connections to their hometowns. Additionally, this pattern might help explain why partially-treated cohorts observe a slight decrease in earnings and income rank, as migration to major urban centers often provides access to higher-paying job opportunities.

2.4.1.5 Health and Behavior

Finally, panel E of Table 2.2 inspects outcomes related to health, broadly related to risky behavior. Fully-treated children are 0.5–1.0 p.p. (0.75%) more likely to survive up to 2019, suggesting overall better health conditions and reduced engagement in risky activities, particularly among boys. Notably, fully-treated girls exhibit a striking 2 p.p. (20%) decrease in the likelihood of childbearing before 18 years old. This finding is crucial as teenage pregnancy is pervasive among poor households in Brazil and has lasting implications throughout the life cycle and could be a key driver behind the human capital and income mobility gains we observe for girls.

2.4.2 Racial and Regional Heterogeneity

Appendix Tables B.3 and B.4 report ITT estimates by race (whites versus non-whites) and region (Center-South versus North and Northeast).¹⁴ Across all adult outcomes we consider, we find relative effects to be larger among the more disadvantaged group, that is, children who are non-white and born in the North and Northeast regions of the country. This goes in hand with the focused approach of PBF, which aims at targeting the most vulnerable groups in society. Moreover, by moving up poor children but moving relative more to those in more underprivileged groups, PBF can be (cautiously) interpreted as force reducing racial and regional disparities in income mobility – such as those documented in (BRITTO et al., 2022).

2.4.3 Heterogeneity in PBF Effectiveness

As noted in (GARCIA; SAAVEDRA, 2023), a key assumption behind CCTs’ design is that relaxing demand-side constraints such as children’s opportunity cost of time will result in increased human capital acquisition. However, in many contexts the supply of schooling – schools, teachers, textbooks – is insufficient. Analogously, increased human capital is a necessary but not sufficient condition for income mobility, as it is usually related to the local availability of opportunities. (BRITTO et al., 2022; CHETTY et al., 2014) In this context, PBF

¹⁴We track race by the modal race reported in employment and welfare registries, and as non-whites blacks (*pretos*), mixed (*pardos*), and indigenous people.

Table 2.2: PBF effects in the Long Run: Age-specific ITT estimates by gender

	Males		Females	
	10–15 (1)	16–21 (2)	10–15 (3)	16–21 (4)
% treated, Treated vs. Control	49.5	47.6	49.3	47.1
<i>A. Human Capital</i>				
Years of Education	0.663*** (0.023)	0.119*** (0.016)	0.918*** (0.024)	0.187*** (0.017)
School completion	0.047*** (0.003)	0.021*** (0.002)	0.060*** (0.003)	0.019*** (0.002)
College degree	0.046*** (0.002)	-0.003*** (0.001)	0.053*** (0.002)	-0.004*** (0.001)
<i>B. Labor Market</i>				
Formal Job	0.039*** (0.003)	0.027*** (0.002)	0.072*** (0.003)	0.025*** (0.002)
Months Worked	0.412*** (0.027)	0.286*** (0.019)	0.631*** (0.026)	0.137*** (0.018)
Formal Income (BRL/year)	1,298.3*** (77.8)	55.3 (47.4)	732.5*** (58.4)	-489.0*** (35.0)
Entrepreneurship	0.028*** (0.001)	0.006*** (0.001)	0.027*** (0.001)	0.008*** (0.001)
<i>C. Intergenerational Mobility</i>				
Adult PBF enrollment	-0.007*** (0.001)	-0.003*** (0.0008)	-0.030*** (0.003)	-0.001 (0.002)
Income pct. rank	3.29*** (0.151)	-0.465*** (0.103)	2.15*** (0.143)	-0.538*** (0.096)
Income rank > 20	0.054*** (0.002)	-0.023*** (0.002)	0.070*** (0.002)	-0.010*** (0.002)
Income rank > 80	0.029*** (0.002)	0.004** (0.001)	0.026*** (0.002)	-0.003* (0.001)
<i>D. Migration</i>				
City Migration	-0.035*** (0.003)	-0.024*** (0.002)	-0.035*** (0.003)	-0.019*** (0.002)
State Migration	-0.021*** (0.002)	-0.013*** (0.001)	-0.018*** (0.002)	-0.010*** (0.001)
Region Migration	-0.015*** (0.002)	-0.010*** (0.001)	-0.011*** (0.001)	-0.006*** (0.001)
Live in Metro Area	0.014*** (0.002)	-0.010*** (0.002)	0.021*** (0.002)	-0.006*** (0.002)
<i>E. Health and Behavior</i>				
Survive up to 2019	0.004*** (0.001)	0.001* (0.001)	0.002*** (0.001)	0.001** (0.000)
Teenage parenthood	-0.002 (0.001)	0.001* (0.001)	-0.010*** (0.002)	0.007*** (0.002)
Observations	11,386,335	11,386,335	11,381,264	11,381,264

Notes. The table reports ITT coefficients on interactions of cohort indicators and our treatment variable. They are obtained from a modified version of Equation 2.1 in which age indicators I_t are replaced by indicators for being 10–15, 16–21, or older (omitted category) when PBF began. We run separate regressions for males and females. A detailed description of outcomes used as dependent variables is given in Appendix Table B.1. (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

Table 2.3: PBF effects in the Long Run: Heterogeneity by Family and Place Inputs

	Mothers' Education		City's School Quality	
	Low (1)	High (2)	Low (3)	High (4)
<i>A. Human Capital</i>				
Years of Education	0.329*** (0.019)	1.08*** (0.050)	0.971*** (0.026)	0.763*** (0.016)
School completion	0.027*** (0.002)	0.076*** (0.005)	0.064*** (0.003)	0.055*** (0.002)
College degree	0.019*** (0.002)	0.061*** (0.005)	0.054*** (0.002)	0.050*** (0.001)
<i>B. Labor Market</i>				
Formal Income (BRL/year)	299.7*** (61.7)	1,206.2*** (148.6)	808.9*** (69.4)	1,239.0*** (50.0)
Entrepreneurship	0.016*** (0.001)	0.024*** (0.003)	0.020*** (0.001)	0.032*** (0.001)
<i>C. Intergenerational Mobility</i>				
Adult PBF enrollment	-0.013*** (0.002)	-0.016*** (0.004)	-0.021*** (0.002)	-0.018*** (0.001)
Income pct. rank	0.389*** (0.135)	3.43*** (0.324)	1.02*** (0.155)	3.66*** (0.104)
<i>D. Migration</i>				
City Migration	-0.025*** (0.002)	-0.049*** (0.006)	-0.039*** (0.003)	-0.037*** (0.002)
Live in Metro Area	0.002 (0.002)	0.024*** (0.005)	0.014*** (0.001)	0.019*** (0.002)
<i>E. Health and Behavior</i>				
Teenage parenthood	-0.006*** (0.002)	-0.005 (0.003)	-0.007*** (0.002)	-0.009*** (0.001)

Notes. The table reports ITT coefficients on interactions of cohort indicators and our treatment variable. They are obtained from a modified version of Equation 2.1 in which age indicators I_t are replaced by indicators for being 10–15, 16–21, or older (omitted category) when PBF began. We run separate regressions for each group and only report coefficients for children aged 10–15. A detailed description of outcomes used as dependent variables is given in Appendix Table B.1. (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

could display heterogeneous effects according to family characteristics related to the relaxing of demand-side constraints and according to place supply-side inputs, such as educational infrastructure.

To gauge this, we look for heterogeneity in long-term effects according to family and place inputs. Table 2.3 reports ITT estimates for fully-treated children, i.e., those aged 10–15 years old when exposed to the program. Columns 1 and 2 contrast the effects for poor children whose mothers were below or above the median educational attainment, respectively. Remarkably, human capital and intergenerational mobility effects are substantially larger for children of mothers with relatively higher education. In turn, columns 3 and 4 compare outcomes for children growing up in cities below and above the median level of (pre-PBF) public school quality, respectively. Despite similar effects for human capital attainment in both groups, labor market, and intergenerational mobility effects are significantly higher in cities with better infrastructure.

2.5 Conclusion

In conclusion, this paper provides robust evidence on the long-term impacts of Brazil's *Programa Bolsa Família* (PBF), the largest Conditional Cash Transfer (CCT) program in the world. Remarkably, our findings demonstrate that PBF successfully achieves its goal of promoting sustained social mobility among the next generation.

Children exposed to PBF's conditionalities between ages 10 and 15 experienced significant

gains in education, with notable increases in high school completion and college degree attainment. Additionally, PBF was associated with improved income mobility, enabling beneficiaries to rise in income distribution and escape poverty. The program also displayed positive impacts on labor market participation, earnings, and reduced dependence on social assistance, debunking the notion of long-term dependency on CCTs. Moreover, we find indirect effects on educational attainment and labor supply for young adults aged 16–21 when PBF began – suggesting spillovers to other household members impacted by the cash transfer. However, young adults do not observe improvements in earnings or social mobility, highlighting the role of conditionalities in breaking the intergenerational cycle of poverty.

The paper’s comprehensive analysis fills a gap in the literature, providing much-needed insights into the intergenerational impact of CCT programs (MILLÁN et al., 2019; PARKER; VOGL, 2021; GARCIA; SAAVEDRA, 2023). These findings have significant implications for policymakers in both developing and developed countries, offering valuable lessons for designing effective poverty reduction strategies in a world of growing inequality (BLATTMAN et al., 2017; GENTILINI et al., 2020).

Decentralization, Tax Administration, and Taxation: Evidence from Brazil's Rural Land Tax

3.1 Introduction

Property taxes represent a much lower proportion of the overall tax revenues in developing countries than in developed ones (BROCKMEYER et al., 2021). Increasing property taxation could help those countries to meet the goal of simultaneously increasing tax revenues and making the tax system more progressive (OECD, 2023). However, limited administrative resources and weak enforcement capacity often constrain developing countries ability to change its tax structure (GORDON; LI, 2009; BESLEY; PERSSON, 2014; SLEMROD, 2019). While there is increasing evidence on strategies to overcome these challenges for other types of taxes (e.g, (CARRILLO; POMERANZ; SINGHAL, 2017), (BROCKMEYER et al., 2019), (PEREZ-TRUGLIA; TROIANO, 2018), (NARITOMI, 2019), (BASRI et al., 2021)), there is much less evidence on the effectiveness of strategies to increase property taxation.¹

This paper studies the effects of a reform that partially decentralized the administration of Brazil's rural land tax (*Imposto Territorial Rural*, ITR). ITR is a progressive tax on landholdings created in the 1960s whose revenues have been historically low due to informational constraints and low enforcement capacity (ASSUNÇÃO; MOREIRA, 2001; FENDRICH et al., 2022).² To tackle these issues, a reform enacted in 2008 authorized the Federal Revenue of Brazil (*Receita Federal do Brasil*, RFB) to sign decentralization agreements with municipalities. Under these agreements, municipalities agreeing to collect information on land values and audit properties for RFB increase their share of ITR revenues from 50 to 100 percent. The reform maintained the RFB's authority in selecting taxpayers to be audited audits and to impose fines to prevent local officials from being captured.³

We investigate the effects of this partial decentralization on tax revenues using microdata on 120 million tax returns from the period 2002-2021. This data enables us to estimate the effects of the decentralization on different margins and across different types of properties. To identify causal effects, we explore the fact that municipalities often sign decentralization agreements with Brazil's Federal Revenue Service, but fail to fully implement the program. Using this

¹One exception is the work of (BROCKMEYER et al., 2021) who studies the effects of increasing tax rates and tax enforcement on the collection of property taxes in Mexico City.

²ITR due is the product of the value of the land, the share of the property not covered by forests, and the tax rate. The tax rate is increasing on property size and decreasing on the intensity of land use. All parameters are self-reported. See section 3.2.1 for details.

³The reform also made RFB responsible for providing IT infrastructure to the municipalities that sign decentralization agreements.

feature, we build a staggered differences-in-differences design comparing municipalities that signed agreements in a given year and implemented the program and municipalities that signed agreements in the same year, but did not implement the program. This design ensures that treatment and controls municipalities are relatively similar, mitigates the concern that time-varying shocks correlated with tax revenues and the timing of entry in the program drive the results, and deals with the negative weighting problem identified in the recent literature on staggered differences-in-differences designs (e.g., (GOODMAN-BACON, 2021a))

We obtain four main results. First, we document that the partial decentralization reform increased ITR revenues by more than 20% after five years. The effects increase over time and are relatively homogeneous across treatment cohorts. Extrapolating the effects found for the earlier treatment cohorts to the full sample suggest that partial decentralization generates more than 40% growth in ITR revenues after ten years. These effects are quite similar across properties from different sizes and slightly larger in regions more intensive in crop cultivation.

Second, we find that increases in reported land values of properties that already paid land taxes explain most of the increases in revenues caused by the program. We do find other margins that are influenced by partial decentralization. At the intensive margin, taxable area (the share of the properties not covered by forests) increases and the effective tax rate (based on reported size and land use) decreases. At the extensive margin, the number of tax paying properties goes up. However, albeit statistically significant, the latter effects are economically small and do not explain the large increase of tax revenues observed as a consequence of the partial decentralization. We further provide evidence that the increase in reported land values is not caused by taxpayers bunching at the minimum land values reported by municipalities to RFB, but rather by a shift in the distribution of reported land values. This is suggestive that the increase in reported land values is not driven simply by improvements in the quality of information about tax parameters, but by a combination of better information and tougher enforcement.

Third, we use satellite-based information to document that the program did not lead to changes in observed land use. Neither the area covered by pastures neither the area cultivated with different crops is affected by the partial decentralization. Because differences in land use are a key source of distinctions in agricultural productivity, this is *prima facie* evidence that the increase in taxation caused by the program did not have substantial efficiency costs.

Fourth, we find that the program is highly cost-effective. The estimates reported above indicate that the typical municipality participating in the program increased tax revenues by BRL 940,000 annually. The overall increase in ITR collection explains 57% of this growth in municipal revenues, whereas the increase in the share of the ITR revenues transferred to the municipalities explains 43%. In contrast, information on implementation costs obtained from RFB indicates that the reform increased costs related to tax administration by BRL 90,000-130,000 annually. Given the small efficiency costs documented previously, these numbers imply that the program has a high benefit-cost ratio either from the perspective of individual municipalities or from the perspective of the public sector as a whole.

Our work contributes to the growing body of empirical research on taxation on developing countries.⁴ Within this literature, there is limited work on property taxes. One exception

⁴See (POMERANZ, 2015), (CARRILLO; POMERANZ; SINGHAL, 2017), (PEREZ-TRUGLIA; TROIANO, 2018), (NARITOMI, 2019), (BROCKMEYER et al., 2019), (BROCKMEYER et al., 2021), (LONDOÑO-VÉLEZ;

is (BROCKMEYER et al., 2021) who document that higher tax rates and more stringent enforcement both lead to increased collection of property taxes in Mexico City, but that increases in rates generate much smaller efficiency costs than increases in enforcement.⁵ Similar to them, we document that higher enforcement increase the collection of property taxes, but, different from them, finds that the efficiency costs of increasing enforcement are quite low. This difference is likely explained by the fact that under-reporting (as opposed to tax delinquencies in their work) is the main challenge for increasing property taxation in our context.

We contribute to this literature by showing that cooperation between local and central authorities can increase tax revenues in contexts with incomplete information and weak enforcement capacity. This finding adds to long standing debates on fiscal centralization (see (OATES et al., 1972) and (BARDHAN; MOOKHERJEE, 2000)) for opposing views on the topic and (BALAN et al., 2022) for recent evidence on this issue). They also add to a growing body of literature documenting high returns of better tax administration (see (BASRI et al., 2021) for recent evidence).

3.2 Institutional background

This section provides an overview of the main facts about the rural property taxation in Brazil and discuss some features of its decentralization that are relevant to our study.

3.2.1 Rural land taxation in Brazil

Land taxation in Brazil begins with the first republican constitution of 1891. However, the ITR only arises in the 1930's. The ITR administration switched over the years between all three tiers of government in Brazil, which reflects the administration challenges associated with this tax. Eventually, in more recent years, the Federal Government was assigned the duty to administer and enforce ITR. However, until the Constitution of 1988, tax revenues were entirely allocated to the municipalities where the land was located. Hence, the tax design created a clear disincentive for the creation and maintenance of an adequate administration of ITR which involves informational challenges and expensive inspection activities. It is important to acknowledge that Brazil is a very unequal country, where 1% of the largest rural properties concentrates 48% of the total rural land. (IBGE, 2017)

Since its creation, there has been a widespread dissatisfaction with the ITR's results not only in terms of tax revenue but also in terms of extrafiscal outcomes.⁶ In fact, the revenue collected

ÁVILA-MAHECHA, 2021), (BASRI et al., 2021), (BALAN et al., 2022), (CARRILLO et al., 2023), among others for examples of work in this literature.)

⁵(CALDEIRA; EHRL; MOREIRA, 2023) studies the same partial decentralization program we study in this paper using aggregate data. They find that the program more than doubled revenues from rural land taxes in participating municipalities. We advance their work by using an identification strategy that better deals with time-varying shocks that simultaneously affect the entry in the program and ITR revenues and by using more detailed data that enables us to understand different margins of adjustment to the partial decentralization.

⁶The ITR is sometimes pejoratively referred to as “the ten reais tax” in allusion to the minimum value of the tax, which is the amount collected by a significant portion of taxpayers. The tax has also failed in achieving results in terms of reducing non used land and in landscapes preservation

Table 3.1: Tax Rates of the ITR

Property Area (ha)	Degree of Utilization - GU (%)				
	Over 80	65 - 80	50 - 65	30 - 65	Up to 30
Up to 50	0.03	0.20	0.40	0.70	1.00
50 - 200	0.07	0.40	0.80	1.40	2.00
200 - 500	0.10	0.60	1.30	2.30	3.30
500 - 1000	0.15	0.85	1.90	3.30	4.70
1000 - 5000	0.30	1.60	3.40	6.00	8.60
Over 5000	0.45	3.00	6.40	12.00	20.00

Notes. This table presents the tax rates of the ITR according to the Federal Law n° 9,393/1996. The tax rate is defined by the total area of the property and the degree of land use.

from the ITR remained around 0.15% of GDP in the period 1970 to 1985, dropping to 0.11% of the GDP in the period 1985 to 1990, according to (BLANCO; REIS, 1996).

Since 1996, the tax has been levied by self declaration. The ITR due is calculated as shown in equation 1.

$$ITR_{due} = (Value\ of\ the\ land) \cdot \left(\frac{taxable\ area}{total\ area} \right) \cdot (tax\ rate) \quad (3.1)$$

where ITR_{due} is the amount of ITR to be paid. The value of the land (*Valor da Terra Nua - VTN*) is the value of the bare land. The taxable area is the total area of the property less the areas occupied by forests and of ecological interest.⁷

The tax rate of the ITR goes from 0,03% up to 20% according to the size of the property and to the degree of utilization (*Grau de Utilização - GU*) as shown in Table 3.1. The GU is calculated by the percentage ratio of the area actually used for rural production⁸ in relation to the total usable area of the property.⁹ Small properties are exempted from the tax, when the owner do not have another rural or urban property.

3.2.2 The Decentralization of the Rural Property Tax

A Constitutional Amendment opening the possibility of municipalities to carry out inspections and collect the ITR was approved in 2008.¹⁰ Municipalities could keep 100% of ITR taxation


⁷The non taxable areas are the parts of permanent preservation, legal reserve, Private Natural Heritage Reserve (RPPN), of ecological interest, of environmental easement, covered by native, primary or secondary forests in medium or advanced stage of regeneration and flooded for purposes of constituting a reservoir for hydroelectric plants authorized by the Public Authority

⁸Area effectively used by rural activity is the area planted, of pasture, of extractive exploration or used for farming or aquaculture activities

⁹The total usable area is considered to be the part of the property suitable for agricultural, livestock, poultry, aquaculture or forestry. It is the total area of the property, excluding the non-taxable areas and the areas occupied with useful and necessary improvements intended for rural activity.

¹⁰The decentralization of the ITR arises with the Constitutional Amendment N° 42/2003. The amendment was regulated by Law N° 11,250/2005 but, only in 2008, with decree n° 6.433/2008 the ITR Management Committee was established, in order to regulate the requirements and conditions necessary for the conclusion of the agreements, municipalities were able to sign the agreement.

Figure 3.1: Evolution of the Decentralization Agreements



Notes. The figure shows the evolution of the decentralization agreements. In panel (a) all municipalities are included whereas in (b) only municipalities that were not excluded are shown.

but they had to attend some requirements. They had to assess and publish the minimal value of its rural land¹¹ and to provide infrastructure and personnel to the task. The municipality employees also had to undergo training carried out by the RFB, before being able to use the system.

The decentralization of the ITR does not represent a full decentralization of tax power and the design of the agreement also presents provisions to prevent the capture of the tax by local elites. In fact, only the tax collection and auditing are transferred to municipalities that sign the agreement in exchange for receiving 100% of the tax revenue. The legislation remains with the Federal Government and the Federal Tax Administration - RFB is still the one to choose who is going to be audited and the order that taxpayers are going to be audited, based in the expected tax notice.

In 2016, a new regulation of the decentralization agreements was published and municipalities had one year to adjust under the penalty of exclusion. In the following year, RFB started revising the agreements and excluded 1,149 municipalities, in total, that were not complying with the obligations of the agreement. In most of the cases, these municipalities did not even had trained personnel with access to the system used to manage the tax. In Figure 3.1, it can be seen the evolution of the decentralization agreements according to the year the municipality signed the agreement. In 2021, out of the 5,570 municipalities of Brazil, the ITR was decentralized in 1,260 (22,6%). The total ITR revenue of that year was R\$ 1,950 million (USD 360 million) and

¹¹The minimal value has to be assessed by engineers according to technical criterias and the value of the land has to be classified in 6 different categories

the 1,260 decentralized municipalities administered 86.7% of that tax. The total ITR revenue corresponds only to 0.02% of the GDP and (FENDRICH et al., 2022) estimates that the total tax collection should be almost four times the current value.

3.3 Data

In this section, we present the main administrative databases used in this paper and discuss the procedures employed to merge and to unidentified the confidential information.

3.3.1 Land Tax Returns

We use the universe of individual land tax returns administered by the Federal Revenue of Brazil - RFB (*Imposto Territorial Rural - ITR* from 2002 to 2021, totaling over 120 million returns. The returns identify the property by a unique tax id number (*Número do Imóvel na Receita Federal - NIRF* and the owner by the unique tax id of natural people (*Cadastro de Pessoa Física - CPF* and of legal entity (*Cadastro Nacional da Pessoa Jurídica - CNPJ*). The NIRF, CNPJ and CPF were unidentified to protect the confidential taxpayer information.

The ITR return is filled on a yearly basis and brings information of the municipality the property is located, the size of the property, the areas of construction, native forest, crops, pasture, the number of animals and the ITR tax rate and the amount due.

In table 3.2, we have the descriptive statistics of the data. It calls attention the reduction of the number of properties paying 10 BRL, the minimum tax due. Although the reduction began prior to the municipalization, it accelerates after that.

Table 3.2: Descriptive Statistics

	2000	2005	2010	2015	2020
<i>Tax descriptives</i>					
Total ITR due (R\$)	30,403.29	42,032.57	80,619.71	183,962.09	292,989.08
Mean ITR due per property (R\$)	65.25	81.03	138.58	295.28	481.99
Number of declarations	849.28	957.27	1083.88	1168.43	1110.76
Share of exempt properties	0.45	0.46	0.46	0.47	0.45
Share paying 10 R\$	0.33	0.31	0.27	0.21	0.16
Average land value (R\$)	38,715.89	53,899.51	91,192.53	190,356.78	306,185.7
Average tax rate (%)	0.19	0.21	0.25	0.26	0.26
Share of taxable land	0.75	0.78	0.78	0.75	0.72
<i>Land use</i>					
Share Agriculture	0.15	0.16	0.17	0.17	0.19
Share Pasture	0.30	0.27	0.244	0.21	0.2

Notes. This table presents the mean values by municipality for each variable for the years shown. *Share of paying 10 R\$* is the share of the properties that pay R\$ on ITR, the minimum value. *Average land value* is the average assessed value accepted by the tax authority in R\$, this is the value of the land without any improvements.

3.3.2 Additional Datasets

For this study, the FAO-GAEZ v3 (Food and Agriculture Organization - Global Agro-Ecological Zones) dataset was used to assess the suitability of soy and maize pastures, to control for increases in land value caused by soy production. Additionally, we use the Mapbiomas dataset, which provides us with information on land use and land cover changes yearly, for robustness checks on the pasture and crop data. For robustness tests on other taxes, IPTU and ISS, we use data from the National Treasury (*Tesouro Nacional*), FINBRA (Municipalities Finance), which provides the yearly tax levied by each municipality. For analysis of land usage, we use the MapBiomas dataset, which extracts, with satellite images, how land is used yearly in each of the Brazilian municipalities.

3.4 Empirical Strategy

As discussed in Section 3.2, the legislation does not force municipalities to become responsible for the collection of land taxes. The legislation merely authorizes the RFB to sign “decentralization” agreements with municipalities, thereby implying the decision to sign these agreements is endogenous.

We deal with this issue by comparing municipalities that signed decentralization agreements in year t and implemented the minimum infrastructure to collect these taxes (treatment) with municipalities that signed decentralization agreements in year t , but were not able to implement the minimum infrastructure to collect land taxes and therefore had their agreements denounced by RFB (control group). Because treatment and control municipalities showed interest in signing agreements in the same period, this comparison reduces concerns that municipalities anticipated their revenues from land taxes would change differentially and entered the program because of this.

Formally, we pool municipalities that signed agreements on different periods and estimate the following dynamic differences-in-differences equation:

$$\log(y_{ict}) = \sum_{k=-6}^{-1} \beta_k T_{ic} + \sum_{k=0}^5 \beta_k T_{ic} + \gamma' \mathbf{X}_{ict} + \lambda_i + \lambda_{tc} + \varepsilon_{ict}, \quad (3.2)$$

in which y_{ict} represents ITR revenues of municipality i , of treatment cohort c in the period t , T_{ic} is an indicator that municipality i made an agreement with RFB in period c and was never denounced. X_{it} is a vector of controls, λ_i is a municipality fixed effect, λ_{tc} is a cohort $\times$ year fixed effect, and ε_{ict} is an idiosyncratic error term.

The parameters of interest in equation (3.2) are the β_k . The inclusion of cohort $\times$ year fixed effects implies that these parameters are a weighted average of within-cohort treatment effects (ATT) obtained for each k . These weights are strictly positive, depending solely on the relative size of the different cohorts. Therefore, despite treatment occurring in different time periods, our empirical design does not suffer from the negative weighting problem identified in the recent literature on differences-in-differences designs with staggered entry (e.g., (GOODMAN-BACON, 2021a)).

Table 3.3: Treatment and Control Group Characteristics

	Treatment	Control	Diff. (1-2)	p-value
	(1)	(2)	(3)	(4)
Municipality GDP	809531.86	682388.61	127143.24	0.65
Public workers per capita	0.04	0.05	-0.00	0.00
Public workers per capita (“estatutários”)	0.03	0.03	-0.00	0.87
Previous Year ISS (per Capita)	22.13	18.25	3.87	0.17
Previous Year IPTU (per Capita)	7.88	6.73	1.15	0.06
Previous Year ITR (per Capita)	5.38	2.14	3.23	0.00
Maize Productivity	8084.93	8197.25	-112.32	0.27
Soybean productivity	4244.20	4109.69	134.51	0.00

Notes. This table reports descriptive statistics of the estimating sample. Column 1 reports means for treatment municipalities, column 2 reports means for control municipalities, column 3 reports the differences in these means, and column 4 reports the p-value associated with this difference.

The β_k estimated for $k < 0$ test whether ITR revenues were evolving similarly in treatment/control municipalities before decentralization agreements were signed. The β_k estimated for $k \geq 0$, on their turn, test whether ITR revenues changed differentially in treatment/control municipalities after these agreements were signed, that is, it tests whether decentralization affected ITR revenues.

One concern with our empirical design is that treatment/control municipalities might face different shocks after agreements are signed. These shocks might in turn influence the dynamics of ITR revenues. To mitigate this concern, we include as controls a series of municipality characteristics (the log of revenues from other municipality taxes in 2002, the log of ITR revenues in 1997, and soy and maize suitability) interacted with time dummies. This ensures that shocks related to these characteristics do not influence our empirical design.

Our treatment group is composed by 1,074 municipalities that signed agreement and were not excluded. In the control group, we assigned the 1,149 municipalities that signed the agreement and were later excluded. Table 3.3 reports descriptive characteristics of the treatment/control groups. The treatment group has higher initial ITR revenues and agricultural suitability, slightly higher collection of local taxes, similar GDP, and slightly lower number of public workers.

3.5 Results

3.5.1 ITR revenues


Figure 3.2, Panel A reports the coefficients obtained from estimating (3.2). Before the decentralization agreements were signed, there was no noticeable difference in the evolution of ITR revenues between treatment and control municipalities. However, immediately after the agreements are signed, ITR revenues begin to increase faster in treatment municipalities relative to control municipalities. We observe a relative increase in ITR revenues in control municipalities of 5% in the year the agreements were signed. The effects increase over time to over 10% one to

three years and to roughly 20% four to five years after the agreements after the agreements were signed. The appendix provides evidence that these results are robust to different specifications. We find that the results are robust to the choice of controls (Figure C.1), to different weighting procedures (Figure C.2), and to using the differences-in-differences estimators proposed by (CALLAWAY; SANT'ANNA, 2021) and (SUN; ABRAHAM, 2021) (Figure C.3).


Figure 3.2, Panel B reports the coefficients obtained from estimating (3.2) separately for the different treatment cohorts. The effects of decentralization on ITR revenues are very similar across cohorts. Data from cohorts that signed the agreement earlier indicates that treatment effects continue to increase up to one decade after the decentralization agreements were signed – for these cohorts, ITR revenues increased 40% more in treatment municipalities than in control municipalities after one decade. The appendix explores other dimensions of heterogeneity of the results. We find no heterogeneity of treatment effects across farm size (Figure C.4) – the effects are quite similar for farms across four size bins (< 50 hectares, 50-100 hectares, 100-1000 hectares, > 1000 hectares). We do find some weak evidence of heterogeneity across regions (Table C.1) – the effects of decentralization are stronger in the center-west and weaker in the northeast although the differences from the mean effects are not statistically significant.

The effects reported in Figure 3.2 might reflect the fact that municipalities that sign decentralization agreements and implement them are municipalities in which tax revenues are increasing across the board due to economic growth. To test this hypothesis, we conduct placebo tests using local taxes as the outcome of interest. Figure 3.3 the results. Panel A examines revenues from the services tax (ISS) and panel B from the urban property tax (IPTU) revenues. Panel A reports small and statistically insignificant increases in ISS revenues in treatment relative to control municipalities. Panel B reports no differences in the evolution of IPTU revenues in treatment and control municipalities. These findings reduces concerns that increases in ITR revenues are reflecting increases in overall taxation.

Figure 3.2: Effects of decentralization on ITR on revenues




(a) Effect on log(ITR revenues)




(b) Effects on log(ITR revenues) by cohort

Notes. These figures report estimates of the effects of the decentralization agreements on log(ITR revenues) obtained using equation (3.2). The x-axis shows the relative time since decentralization while the y-axis plot the coefficients and confidence intervals obtained. Panel A reports the average effect across all cohorts. Panel B reports the effects for different cohorts.

Figure 3.3: Placebo Tests



(a) Effect on log(ISS revenues)



(b) Effects on log(IPTU revenues)

Notes. These figures report estimates of the effects of the decentralization agreements on log(ISS revenues) (Panel A) and log(IPTU revenues) (Panel B). Effects are obtained using equation (3.2). The x-axis shows the relative time since decentralization while the y-axis plot the coefficients and confidence intervals obtained.

3.5.2 Mechanisms

Figure 3.4, Panels A and B decompose the effects of decentralization on ITR revenues between changes in the average tax rate and the (self-reported) land values.¹² Panel A documents that, after decentralization agreements are signed, tax rates immediately decline by 0.01 p.p. (5%) in treatment municipalities. Panel B documents that, after decentralization agreements are signed, self-reported land values increase strongly in treatment municipalities (25% after five years). The results indicate that the overall 20% increase in ITR revenues is the result of a combination of a 25% increase in self-reported land values and a 5% decrease in tax rates. We interpret these results as an indication that decentralization increased tax revenues operated mainly by inducing taxpayers to report land values closer to the real values. Taxpayers responded to this by changing the reporting of other parameters to reduce rates (possibly because these parameters are more difficult to monitor than land values), but this only mitigates part of the increase in self-reported land values.

Figure 3.4, Panels C and D decompose the effects of decentralization on ITR revenues between the intensive (average ITR paid) and extensive (number of taxpayers) margin.¹³ The figures document that, after decentralization agreements are signed, average ITR paid and the number of taxpayers increase in treatment municipalities. However, increases in average ITR paid are much larger than increases in the number of taxpayers (20% vs. 1%) and explain most of the increases in overall taxation.

Together, the results from Figure 3.4 indicate that increases in ITR revenues are mostly led by increases in self-reported land values by existing taxpayers. Municipalities might be inducing taxpayers to report land values differently either by providing better information to RFB about land values or by increasing the likelihood of audits (or other enforcement measures in general). The two explanations have different implications for the distribution of self-reported land values – the former would generate bunching at the minimum values reported by RFB, whereas the latter would shift the distribution to the right.

Figure 3.5 tests across these explanations. It plots the distribution of self-reported land values divided by the minimum land value for the treatment and the control groups.¹⁴ There are two noticeable patterns. First, there is bunching on minimum values for both groups, but substantially less in the treatment group than in the control group. Second, the entire distribution of self-reported land values for the treatment group is to the right of the distribution control for the control group. These results provide suggestive evidence that the reform affected taxpayer behavior not only by changing the quality of information that the RFB receives about the taxpayers, but also by changing the overall enforcement. However, because the data on


¹²As discussed in section 3.2, ITR due is a product of self-reported land values and the tax rate. Let T_{it} denote ITR due and V_{it} denote the sum of self-reported land values in municipality i and period t . Define $r_{it} = T_{it}/V_{it}$ as the average tax rate. The effect of the decentralization agreements on ITR revenues ($\log(T_{it})$) can be thus written as the sum of the effects on the tax rate ($\log(r_{it})$) and self-reported land values ($\log(V_{it})$).

¹³Let T_{it} denote ITR due, N_{it} denote the number of taxpayers, and t_{it} the average ITR paid. By definition, $T_{it} = t_{it} \times N_{it}$. The effect of the decentralization agreements on ITR revenues ($\log(T_{it})$) can be thus written as the sum of the effects on the average ITR paid ($\log(t_{it})$) and the number of taxpayers ($\log(N_{it})$).

¹⁴The minimum land value is the value of one hectare of land covered with native vegetation. Municipalities that sign decentralization agreements report these values to RFB. RFB relies on other sources of data to construct these values for the other municipalities.


minimum land values is just available for the post-treatment period, these results should be interpreted with caution.

Figure 3.4: Decomposing the effects of decentralization on ITR revenues



Notes. These figures report estimates of the effects of the decentralization agreements on different outcomes obtained using equation (3.2). The x-axis shows the relative time since decentralization while the y-axis plot the coefficients and their respective 95% confidence intervals. Panel A reports the effect of the decentralization agreements on the (log) of self-reported land values. Panel B reports the effects of the decentralization agreements on the average tax rate. Panel C reports the effects of the decentralization agreements on the number of paying properties. Panel D reports the effects of the decentralization agreements on the average ITR paid.

Figure 3.5: Self-reported versus minimum land values




Notes. The figure plots the distribution of self-reported land values divided by the minimum land value for the treatment and the control group.

3.5.3 Efficiency Costs


Figure 3.6 uses remote sensing data to investigate whether decentralization influences taxpayers' land use choices. Whether the land is used as pastureland or as cropland is an important dimension of intensification in Brazil's agriculture (BRAGANÇA, 2018; BUSTOS; CAPRETTINI; PONTICELLI, 2016). Examining decentralization's effects on land use is therefore important to determine whether it influenced taxpayers behavior and therefore whether its efficiency costs were significant. We find that decentralization neither influenced the share of pastureland (Panel A) nor the share of cropland (Panel B).¹⁵ This provides suggestive evidence that the efficiency costs of decentralization are negligible. This is possibly linked to the fact that, even after the increase caused by decentralization, the effective ITR rates continue quite low and therefore unlikely to affect behavior considerably.

¹⁵We consider three different mutually exclusive land uses: cropland, pastureland and forests. We omit forests from the analysis as it is just the residual of cropland and pastureland.

Figure 3.6: The Effects of Decentralization on Land Use



(a) Fraction of land used for pasture



(b) Fraction of land used for agriculture

Notes. These figures report estimates of the effects of the decentralization agreements on the shares of pastureland (Panel A) and cropland (Panel B). Effects are obtained using equation (3.2). The x-axis shows the relative time since decentralization while the y-axis plot the coefficients and their respective 95% confidence intervals.

3.5.4 Cost Benefit

We monetize the costs and benefits of decentralization to compute its cost-benefit ratio. We focus on year 2021 – the last period in our data. Table 3.4 reports the results.

Table 3.4: Cost Benefit Analysis (2021)

Description	Amount (BRL)
Total Revenue Gain by Municipality (1260)	940,000
Average increase in revenue collection by municipality	537,000
Tax transference from federal to municipality	403,000
Cost of Assessing Land	50,000
Cost of Land Tax Administration by Municipality	40,000 - 80,000
Total Cost by Municipality	90,000 - 130,000
Total Net Gain in Tax Revenue by Municipality	407,000 - 850,000
Total Per Capita Gain in Taxes (Including transference)	55
Net Per Capita Gain in Taxes	28

Notes. This table reports the cost benefit analysis for the year 2021.

Combining the mean ITR revenues in 1,260 municipalities with active decentralization agreements with our estimates, we estimate that the reform increased municipal tax revenues by BRL 940,000 (USD 188,000). The overall increase in ITR collection explains 57% of this growth (BRL 537,000), whereas the increase in the share of the ITR revenues transferred to the municipalities explains 43% (403,000).

We compare these benefits with the costs of implementing the decentralization agreements.¹⁶ These implementation costs were obtained from interviews with the officials from RFB responsible for ITR administration. We consider two types of costs – the costs of assessing land values and the costs of the personnel involved in ITR collection at the local level (usually 1/2 of the time of a public employee). While there is some heterogeneity in the costs depending on local characteristics, RFB officials estimate these costs to be roughly between BRL 90,000-130,000. These numbers imply that ITR decentralization had a cost/benefit ratio of 7-10 from the perspective of the municipalities and 4-6 from the perspective of the public sector.

3.6 Conclusion

This paper evaluated the effects of a program that partially decentralized the administration of rural land taxes to local authorities in Brazil. Using microdata from tax returns, we find that the program led to an expansion of tax revenues by 20% after five years. While taxpayers responded to the decentralization in different margins, we find that it expanded tax revenues mainly by

¹⁶We ignore efficiency costs as our findings indicate these costs are negligible.

increasing self-reported land values. This increase in self-reported land values seems to be at least partially connected to increased enforcement. Using satellite data, we further found that partial decentralization did not influence farmer behavior significantly. A cost-benefit exercise indicates that the reform had large returns.

Our findings have important implications for policy design. First, the robust response of tax revenues (almost 50% in one decade) indicates that partial decentralization schemes in which central governments use local officials for information collection and enforcement purposes, while keeping oversight responsibilities might be an effective way for increasing taxation in developing countries (BALAN et al., 2022). Second, the negligible response of taxpayer behavior suggests that the current ITR rate is below the optimal rate. Third, the findings that taxpayers respond to the tax mainly by reporting land values that are closer to real land values indicates that moving from self-reported land values to market assessments (as done for the collection of urban property taxes) is a low hanging fruit to increase ITR revenues.

References

- ABBAS, H.; SICSIC, M. Who Moves Up the Income Ladder Relative to their Parents? An Analysis of Intergenerational Income Mobility in France. 2022. Disponível em: <<https://www.insee.fr/en/statistiques/6445472>>. Citado 2 vezes nas páginas page.11 e page.1818.
- ABOWD, J. M.; STINSON, M. H. Estimating measurement error in annual job earnings: A comparison of survey and administrative data. *Review of Economics and Statistics*, The MIT Press, v. 95, n. 5, p. 1451–1467, 2013. Citado na página page.22.
- ACCIARI, P.; POLO, A.; VIOLANTE, G. 'and yet, it moves': Intergenerational mobility in Italy. *American Economic Journal: Applied Economics*, 2021. Citado 7 vezes nas páginas page.11, page.66, page.1212, page.1818, page.1919, page.2424 e page.9494.
- ALESINA, A. et al. Intergenerational Mobility in Africa. *Econometrica*, v. 89, n. 1, p. 1–35, 2021. ISSN 1468-0262. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA17018>. Disponível em: <<https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA17018>>. Citado 5 vezes nas páginas page.44, page.2424, page.2525, page.2626 e page.2828.
- ALESINA, A.; STANTCHEVA, S.; TESO, E. Intergenerational Mobility and Preferences for Redistribution. *American Economic Review*, v. 108, n. 2, p. 521–554, fev. 2018. ISSN 0002-8282. Disponível em: <<https://www.aeaweb.org/articles?id=10.1257/aer.20162015>>. Citado 2 vezes nas páginas page.11 e page.2929.
- ARAUJO, M. C.; MACOURS, K. Education, Income and Mobility: Experimental Impacts of Childhood Exposure to Progresa after 20 Years. *Working Papers*, out. 2021. Number: halshs-03364972 Publisher: HAL. Disponível em: <<https://ideas.repec.org/p/hal/wpaper/halshs-03364972.html>>. Citado na página page.3434.
- ASHER, S.; NOVOSAD, P.; RAFKIN, C. Intergenerational Mobility in India: Estimates From New Methods and Administrative Data. 2021. Citado na página page.44.
- ASSUNÇÃO, J. J.; MOREIRA, H. Towards a truthful land taxation mechanism in Brazil. *Brazilian Review of Econometrics*, Sociedade Brasileira de Econometria, v. 21, n. 1, p. 49–99, 2001. Citado na página page.4545.

ATHEY, S.; TIBSHIRANI, J.; WAGER, S. Generalized random forests. *The Annals of Statistics*, Institute of Mathematical Statistics, v. 47, n. 2, p. 1148 – 1178, 2019. Disponível em: <<https://doi.org/10.1214/18-AOS1709>>. Citado 3 vezes nas páginas page.99, page.7878 e page.7979.

ATTANASIO, O. et al. Working Paper, *Long Term Effects of Cash Transfer Programs in Colombia*. National Bureau of Economic Research, 2021. (Working Paper Series). Disponível em: <<https://www.nber.org/papers/w29056>>. Citado na página page.3434.

BAILEY, M. J. et al. Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program. *The Review of Economic Studies*, p. rdad063, jun. 2023. ISSN 0034-6527. Disponível em: <<https://doi.org/10.1093/restud/rdad063>>. Citado 2 vezes nas páginas page.3131 e page.3434.

BALAN, P. et al. Local elites as state capacity: How city chiefs use local information to increase tax compliance in the democratic republic of the congo. *American Economic Review*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, v. 112, n. 3, p. 762–797, 2022. Citado 2 vezes nas páginas page.4747 e page.6161.

BARDHAN, P.; MOOKHERJEE, D. Capture and governance at local and national levels. *American economic review*, American Economic Association, v. 90, n. 2, p. 135–139, 2000. Citado na página page.4747.

BARHAM, T.; MACOURS, K.; MALUCCIO, J. A. SSRN Scholarly Paper, *Experimental Evidence of Exposure to a Conditional Cash Transfer During Early Teenage Years: Young Women's Fertility and Labor Market Outcomes*. Rochester, NY: [s.n.], 2018. Disponível em: <<https://papers.ssrn.com/abstract=3247237>>. Citado na página page.3434.

BARR, A.; EGGLESTON, J.; SMITH, A. A. Investing in Infants: the Lasting Effects of Cash Transfers to New Families*. *The Quarterly Journal of Economics*, v. 137, n. 4, p. 2539–2583, nov. 2022. ISSN 0033-5533. Disponível em: <<https://doi.org/10.1093/qje/qjac023>>. Citado na página page.3434.

BARRERA-OSORIO, F.; LINDEN, L. L.; SAAVEDRA, J. E. Medium- and Long-Term Educational Consequences of Alternative Conditional Cash Transfer Designs: Experimental Evidence from Colombia. *American Economic Journal: Applied Economics*, v. 11, n. 3, p. 54–91, jul. 2019. ISSN 1945-7782. Disponível em: <<https://www.aeaweb.org/articles?id=10.1257/app.20170008>>. Citado na página page.3434.

BASRI, M. C. et al. Tax administration versus tax rates: Evidence from corporate taxation in indonesia. *American Economic Review*, v. 111, n. 12, p. 3827–71, 2021. Citado 2 vezes nas páginas page.4545 e page.4747.

BASTAGLI, F. et al. Cash transfers: what does the evidence say? 2016. Citado 2 vezes nas páginas page.3131 e page.3434.

BESLEY, T.; PERSSON, T. Why do developing countries tax so little? *Journal of economic perspectives*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203-2418, v. 28, n. 4, p. 99–120, 2014. Citado na página page.4545.

BILAL, A.; ROSSI-HANSBERG, E. Location as an asset. *Econometrica*, v. 89, n. 5, p. 2459–2495, 2021. Citado na página page.1212.

BJÖRKLUND, A.; JÄNTTI, M. Intergenerational mobility, intergenerational effects, sibling correlations, and equality of opportunity: a comparison of four approaches. *Research in Social Stratification and Mobility*, Elsevier, v. 70, p. 100455, 2020. Citado na página page.44.

BLACK, S. E.; DEVEREUX, P. J. et al. Recent developments in intergenerational mobility. *Handbook of Labor Economics*, v. 4, p. 1487–1541, 2011. Citado 2 vezes nas páginas page.44 e page.66.

BLANCO, F.; REIS, E. A capacidade tributária dos estados brasileiros (1970-90). *Rio de Janeiro: IPEA*, 1996. Citado na página page.4848.

BLANDEN, J. Cross-country rankings in intergenerational mobility: a comparison of approaches from economics and sociology. *Journal of Economic Surveys*, Wiley Online Library, v. 27, n. 1, p. 38–73, 2013. Citado na página page.44.

BLATTMAN, C. et al. Cash as Capital. 2017. Citado 2 vezes nas páginas page.3131 e page.4444.

BOUND, J. et al. Evidence on the validity of cross-sectional and longitudinal labor market data. *Journal of Labor Economics*, University of Chicago Press, v. 12, n. 3, p. 345–368, 1994. Citado 2 vezes nas páginas page.22 e page.1010.

BRAGANÇA, A. The effects of crop-to-beef relative prices on deforestation: evidence from the tapajós basin. *Environment and development economics*, Cambridge University Press, v. 23, n. 4, p. 391–412, 2018. Citado na página page.5858.

BRATBERG, E. et al. A Comparison of Intergenerational Mobility Curves in Germany, Norway, Sweden, and the US. *The Scandinavian Journal of Economics*, v. 119, n. 1, p. 72–101, 2017. ISSN 1467-9442. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/sjoe.12197>. Disponível em: <<https://onlinelibrary.wiley.com/doi/abs/10.1111/sjoe.12197>>. Citado 3 vezes nas páginas page.11, page.1212 e page.1818.

BRAUN, S. T.; STUHLER, J. The transmission of inequality across multiple generations: testing recent theories with evidence from germany. *The Economic Journal*, v. 128, n. 609, p. 576–611, 2018. Citado na página page.1212.

BRITTO, D. G. et al. Intergenerational Mobility in the Land of Inequality. 2022. Disponível em: <https://www.dropbox.com/s/u1juunq4syd983q/2022\09\30\IGM_BFPSW.pdf?dl=0>. Citado 5 vezes nas páginas page.3333, page.3636, page.4040, page.4141 e page.103103.

BRITTO, D. G.; PINOTTI, P.; SAMPAIO, B. The effect of job loss and unemployment insurance on crime in brazil. *Econometrica*, Wiley Online Library, v. 90, n. 4, p. 1393–1423, 2022. Citado 2 vezes nas páginas page.77 e page.3535.

BROCKMEYER, A. et al. *Taxing property in developing countries: theory and evidence from Mexico*. [S.l.], 2021. Citado 3 vezes nas páginas page.4545, page.4646 e page.4747.

BROCKMEYER, A. et al. Casting a wider tax net: Experimental evidence from costa rica. *American Economic Journal: Economic Policy*, American Economic Association, v. 11, n. 3, p. 55–87, 2019. Citado 2 vezes nas páginas page.4545 e page.4646.

BUSTOS, P.; CAPRETTINI, B.; PONTICELLI, J. Agricultural productivity and structural transformation: Evidence from brazil. *American Economic Review*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, v. 106, n. 6, p. 1320–1365, 2016. Citado na página page.5858.

CAHYADI, N. et al. Cumulative Impacts of Conditional Cash Transfer Programs: Experimental Evidence from Indonesia. *American Economic Journal: Economic Policy*, v. 12, n. 4, p. 88–110, nov. 2020. ISSN 1945-7731. Disponível em: <<https://www.aeaweb.org/articles?id=10.1257/pol.20190245>>. Citado na página page.3434.

CALDEIRA, T. C. M.; EHRL, P.; MOREIRA, T. B. S. Fiscal decentralization and tax collection: Evidence from the rural property tax in brazil. *European Journal of Political Economy*, Elsevier, p. 102396, 2023. Citado na página page.4747.

CALLAWAY, B.; SANT’ANNA, P. H. Difference-in-differences with multiple time periods. *Journal of econometrics*, Elsevier, v. 225, n. 2, p. 200–230, 2021. Citado 3 vezes nas páginas page.xivxiv, page.5353 e page.112112.

CARD, D.; ROTHSTEIN, J.; YI, M. *Location, Location, Location*. [S.l.], 2021. Disponível em: <<https://EconPapers.repec.org/RePEc:cen:wpaper:21-32>>. Citado na página page.1212.

CARRILLO, P. et al. Ghosting the tax authority: Fake firms and tax fraud in ecuador. *American Economic Review: Insights*, 2023. Citado na página page.4747.

CARRILLO, P.; POMERANZ, D.; SINGHAL, M. Dodging the taxman: Firm misreporting and limits to tax enforcement. *American Economic Journal: Applied Economics*, v. 9, n. 2, p. 144–64, 2017. Citado 2 vezes nas páginas page.4545 e page.4646.

CENGIZ, D. et al. Seeing beyond the Trees: Using Machine Learning to Estimate the Impact of Minimum Wages on Labor Market Outcomes. *Journal of Labor Economics*, v. 40, n. S1, p. S203–S247, abr. 2022. ISSN 0734-306X. Publisher: The University of Chicago Press. Disponível em: <<https://www.journals.uchicago.edu/doi/10.1086/718497>>. Citado 2 vezes nas páginas page.3131 e page.3636.

CHETTY, R.; HENDREN, N. The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The Quarterly Journal of Economics*, v. 133, n. 3, p. 1107–1162,

2018. ISSN 0033-5533. Disponível em: <<https://doi.org/10.1093/qje/qjy007>>. Citado 6 vezes nas páginas page.33, page.44, page.2525, page.2626, page.2727 e page.9595.

CHETTY, R.; HENDREN, N. The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates*. *The Quarterly Journal of Economics*, v. 133, n. 3, p. 1163–1228, ago. 2018. ISSN 0033-5533. Disponível em: <<https://doi.org/10.1093/qje/qjy006>>. Citado 2 vezes nas páginas page.44 e page.9595.

CHETTY, R. et al. Race and Economic Opportunity in the United States: an Intergenerational Perspective*. *The Quarterly Journal of Economics*, v. 135, n. 2, p. 711–783, maio 2020. ISSN 0033-5533. Disponível em: <<https://doi.org/10.1093/qje/qjz042>>. Citado na página page.33.

CHETTY, R.; HENDREN, N.; KATZ, L. F. The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, v. 106, n. 4, p. 855–902, 2016. Citado na página page.44.

CHETTY, R. et al. Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, v. 129, n. 4, p. 1553–1623, 2014. ISSN 0033-5533. Disponível em: <<https://doi.org/10.1093/qje/qju022>>. Citado 9 vezes nas páginas page.11, page.22, page.44, page.66, page.1212, page.1818, page.1919, page.2424 e page.4141.

CHYN, E. Moved to opportunity: The long-run effects of public housing demolition on children. *American Economic Review*, v. 108, n. 10, p. 3028–56, 2018. Citado na página page.44.

CHYN, E.; KATZ, L. F. Neighborhoods matter: Assessing the evidence for place effects. *Journal of Economic Perspectives*, v. 35, n. 4, p. 197–222, 2021. Citado na página page.44.

COHEN, J. *Statistical power analysis for the behavioral sciences*. [S.l.]: Routledge, 2013. Citado na página page.7777.

CONNOLLY, M.; CORAK, M.; HAECK, C. Intergenerational mobility between and within canada and the united states. *Journal of Labor Economics*, v. 37, p. S595–S641, 2019. ISSN 0734-306X. Disponível em: <<https://www.journals.uchicago.edu/doi/abs/10.1086/703465>>. Citado 2 vezes nas páginas page.11 e page.1818.

CORAK, M. Income Inequality, Equality of Opportunity, and Intergenerational Mobility. *Journal of Economic Perspectives*, v. 27, n. 3, p. 79–102, set. 2013. ISSN 0895-3309. Disponível em: <<https://www.aeaweb.org/articles?id=10.1257/jep.27.3.79>>. Citado na página page.1818.

CORAK, M. The canadian geography of intergenerational income mobility. *The Economic Journal*, Oxford University Press, v. 130, n. 631, p. 2134–2174, 2020. Citado na página page.1212.

CORAK, M.; HEISZ, A. The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data. *The Journal of Human Resources*, v. 34, n. 3, p. 504, 1999. ISSN 0022166X. Disponível em: <<https://www.jstor.org/stable/146378?origin=crossref>>. Citado na página page.66.

CROSSLEY, T. F.; LEVELL, P.; POUPAKIS, S. Regression with an imputed dependent variable. *Journal of Applied Econometrics*, Wiley Online Library, 2022. Citado 2 vezes nas páginas page.22 e page.1010.

DAHL, M.; DELEIRE, T. The association between children's earnings and fathers' lifetime earnings: Estimates using administrative data. *Institute for Research on Poverty, University of Wisconsin*, v. 1342, 2008. Citado na página page.66.

DAMM, A. P.; DUSTMANN, C. Does growing up in a high crime neighborhood affect youth criminal behavior? *American Economic Review*, v. 104, n. 6, p. 1806–32, 2014. Citado na página page.44.

DAVIS, J. M.; MAZUMDER, B. Racial and ethnic differences in the geography of intergenerational mobility. *Unpublished Working Paper*, 2018. Citado na página page.33.

DEUTSCHER, N. Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia. *American Economic Journal: Applied Economics*, v. 12, n. 2, p. 220–249, abr. 2020. ISSN 1945-7782. Disponível em: <<https://www.aeaweb.org/articles?id=10.1257/app.20180329>>. Citado 3 vezes nas páginas page.44, page.2626 e page.9595.

DEUTSCHER, N.; MAZUMDER, B. Intergenerational mobility across australia and the stability of regional estimates. *Labour Economics*, v. 66, 2020. Disponível em: <<https://ideas.repec.org/a/eee/labeco/v66y2020ics0927537120300658.html>>. Citado 2 vezes nas páginas page.11 e page.1818.

DUNN, c. The Intergenerational Transmission of Lifetime Earnings: Evidence from Brazil. *The B.E. Journal of Economic Analysis & Policy*, v. 7, n. 2, p. 1–42, 2007. Publisher: De Gruyter. Disponível em: <<https://ideas.repec.org/a/bpj/bejeap/v7y2007i2n2.html>>. Citado 3 vezes nas páginas page.44, page.66 e page.1919.

ENGBOM, N. et al. Earnings inequality and dynamics in the presence of informality: The case of brazil. *Quantitative Economics*, Wiley Online Library, v. 13, n. 4, p. 1405–1446, 2022. Citado na página page.44.

FENDRICH, A. N. et al. Taxation aiming environmental protection: The case of brazilian rural land tax. *Land Use Policy*, Elsevier, v. 119, p. 106164, 2022. Citado 2 vezes nas páginas page.4545 e page.5050.

FERRAZ, C.; FINAN, F.; SZERMAN, D. *Procuring firm growth: the effects of government purchases on firm dynamics*. [S.l.], 2015. Citado na página page.88.

FERREIRA, S.; VELOSO, F. Mobilidade intergeracional de educação no brasil. *Pesquisa e Planejamento Econômico*, v. 33, n. 3, 2003. Citado na página page.44.

FISZBEIN, A.; SCHADY, N. R.; FERREIRA, F. H. G. *Conditional cash transfers: reducing present and future poverty*. Washington D.C: World Bank, 2009. (A World Bank policy research report). OCLC: 271104749. ISBN 978-0-8213-7352-1. Citado 2 vezes nas páginas page.3131 e page.3434.

GARCIA, S.; SAAVEDRA, J. E. Conditional cash transfers for education. In: HANUSHEK, E. A.; MACHIN, S.; WOESSMANN, L. (Ed.). Elsevier, 2023, (Handbook of the Economics of Education, v. 6). p. 499–590. Disponível em: <<https://www.sciencedirect.com/science/article/pii/S1574069222000046>>. Citado 6 vezes nas páginas page.3131, page.3434, page.3535, page.3838, page.4141 e page.4444.

GENTILINI, U. et al. Social Protection and Jobs Responses to COVID-19: A Real-Time Review of Country Measures. abr. 2020. Publisher: World Bank, Washington, DC. Disponível em: <<http://hdl.handle.net/10986/33635>>. Citado 2 vezes nas páginas page.3131 e page.4444.

GERARD, F.; GONZAGA, G. Informal labor and the efficiency cost of social programs: Evidence from unemployment insurance in brazil. *American Economic Journal: Economic Policy*, 2021. Citado na página page.88.

GOODMAN-BACON, A. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, Elsevier, v. 225, n. 2, p. 254–277, 2021. Citado 2 vezes nas páginas page.4646 e page.5151.

GOODMAN-BACON, A. The Long-Run Effects of Childhood Insurance Coverage: Medicaid Implementation, Adult Health, and Labor Market Outcomes. *American Economic Review*, v. 111, n. 8, p. 2550–2593, ago. 2021. ISSN 0002-8282. Disponível em: <<https://www.aeaweb.org/articles?id=10.1257/aer.20171671>>. Citado 2 vezes nas páginas page.3131 e page.3434.

GORDON, R.; LI, W. Tax structures in developing countries: Many puzzles and a possible explanation. *Journal of public Economics*, Elsevier, v. 93, n. 7-8, p. 855–866, 2009. Citado na página page.4545.

GOTTSCHALK, P.; HUYNH, M. Are earnings inequality and mobility overstated? the impact of nonclassical measurement error. *The Review of Economics and Statistics*, The MIT Press, v. 92, n. 2, p. 302–315, 2010. Citado 3 vezes nas páginas page.22, page.1010 e page.1111.

GUVENEN, F.; PISTAFERRI, L.; VIOLANTE, G. L. Global trends in income inequality and income dynamics: New insights from grid. *Quantitative Economics*, Wiley Online Library, v. 13, n. 4, p. 1321–1360, 2022. Citado 2 vezes nas páginas page.44 e page.8888.

HAIDER, S.; SOLON, G. Life-Cycle Variation in the Association between Current and Lifetime Earnings. *American Economic Review*, v. 96, n. 4, p. 1308–1320, 2006. Publisher: American Economic Association. Disponível em: <<https://ideas.repec.org/a/aea/aecrev/v96y2006i4p1308-1320.html>>. Citado na página page.8686.

HAINMUELLER, J. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, v. 20, n. 1, p. 25–46, 2012. Citado 4 vezes nas páginas page.7777, page.7878, page.8585 e page.8686.

HECKMAN, J. J.; LANDERSØ, R. *Lessons from Denmark about Inequality and Social Mobility*. [S.l.], 2021. Disponível em: <<https://www.nber.org/papers/w28543>>. Citado na página page.1818.

- HECKMAN, J. J.; MOSSO, S. The Economics of Human Development and Social Mobility. *Annual Review of Economics*, v. 6, n. 1, p. 689–733, 2014. _eprint: <https://doi.org/10.1146/annurev-economics-080213-040753>. Disponível em: <<https://doi.org/10.1146/annurev-economics-080213-040753>>. Citado 4 vezes nas páginas page.3232, page.3434, page.3535 e page.3939.
- HEIDRICH, S. Intergenerational mobility in sweden: a regional perspective. *Journal of Population Economics*, v. 30, n. 4, p. 1241–1280, 2017. ISSN 1432-1475. Disponível em: <<https://doi.org/10.1007/s00148-017-0648-x>>. Citado 2 vezes nas páginas page.11 e page.1818.
- HELSØ, A.-L. Intergenerational income mobility in denmark and the united states. *The Scandinavian Journal of Economics*, Wiley Online Library, v. 123, n. 2, p. 508–531, 2021. Citado 2 vezes nas páginas page.11 e page.1818.
- IBGE. *Censos of Agriculture*. [S.l.], 2017. Citado na página page.4747.
- IBGE. *Pesquisa Nacional por Amostra de Domicílios (PNAD)*. 2019. Citado 2 vezes nas páginas page.11 e page.55.
- INOUE, A.; SOLON, G. Two-sample instrumental variables estimators. *The Review of Economics and Statistics*, The MIT Press, v. 92, n. 3, p. 557–561, 2010. Citado 3 vezes nas páginas page.22, page.1010 e page.8080.
- JÁCOME, E.; KUZIEMKO, I.; NAIDU, S. *Mobility for all: Representative intergenerational mobility estimates over the 20th century*. [S.l.], 2021. Citado na página page.44.
- JERRIM, J.; CHOI, Á.; SIMANCAS, R. Two-sample two-stage least squares (tstsls) estimates of earnings mobility: how consistent are they? In: *Survey Research Methods*. [S.l.: s.n.], 2016. v. 10, n. 2, p. 85–101. Citado 3 vezes nas páginas page.22, page.1010 e page.8080.
- JOHANNEMANN, J. et al. Sufficient representations for categorical variables. *Unpublished Working Paper*, 2019. Citado na página page.7979.
- KIM, B.; SOLON, G. Implications of mean-reverting measurement error for longitudinal studies of wages and employment. *Review of Economics and statistics*, MIT Press 238 Main St., Suite 500, Cambridge, MA 02142-1046, USA journals . . . , v. 87, n. 1, p. 193–196, 2005. Citado na página page.22.
- LEE, C.-I.; SOLON, G. Trends in intergenerational income mobility. *The Review of Economics and Statistics*, v. 91, n. 4, p. 766–772, 2009. Citado na página page.66.
- LEITES, M. et al. Intergenerational mobility and top income persistence for a developing country: estimates using administrative data from Uruguay. 2022. Citado na página page.44.
- LEONE, T. The geography of intergenerational mobility: Evidence of educational persistence and the “great gatsby curve” in brazil. *Review of Development Economics*, Wiley Online Library, 2018. Citado na página page.44.

LINDAHL, M. et al. Long-term intergenerational persistence of human capital an empirical analysis of four generations. *Journal of Human Resources*, v. 50, n. 1, p. 1–33, 2015. Citado na página page.1212.

LONDOÑO-VÉLEZ, J.; ÁVILA-MAHECHA, J. Enforcing wealth taxes in the developing world: Quasi-experimental evidence from colombia. *American Economic Review: Insights*, v. 3, n. 2, p. 131–48, 2021. Citado na página page.4747.

MAHLMEISTER, R. et al. Revisitando a mobilidade intergeracional de educação no brasil. *Inspere Policy Paper*, n. 26, 2017. Citado na página page.44.

MEDINA, L.; SCHNEIDER, M. F. *Shadow economies around the world: what did we learn over the last 20 years?* [S.l.]: International Monetary Fund, 2018. Citado na página page.44.

MELLO, U.; NYBOM, M.; STUHLER, J. A lifecycle estimator of intergenerational earnings mobility. *Unpublished Working Paper*, 2021. Citado 2 vezes nas páginas page.8686 e page.8787.

MENESES, F. Intergenerational mobility in chile: A year-to-year analysis of a national cohort of students (rr). *Unpublished Working Paper*, 2020. Citado na página page.44.

MILLÁN, T. M. et al. Long-Term Impacts of Conditional Cash Transfers: Review of the Evidence. *The World Bank Research Observer*, v. 34, n. 1, p. 119–159, fev. 2019. ISSN 0257-3032. Disponível em: <<https://doi.org/10.1093/wbro/lky005>>. Citado 3 vezes nas páginas page.3131, page.3434 e page.4444.

MILLÁN, T. M. et al. Experimental long-term effects of early-childhood and school-age exposure to a conditional cash transfer program. *Journal of Development Economics*, v. 143, p. 102385, mar. 2020. ISSN 0304-3878. Disponível em: <<https://www.sciencedirect.com/science/article/pii/S0304387818312768>>. Citado na página page.3434.

NARAYAN, A. et al. *Fair progress?: Economic mobility across generations around the world*. [S.l.]: World Bank Publications, 2018. Citado 2 vezes nas páginas page.11 e page.44.

NARITOMI, J. Consumers as tax auditors. *American Economic Review*, v. 109, n. 9, p. 3031–72, 2019. Citado 2 vezes nas páginas page.4545 e page.4646.

OATES, W. E. et al. Fiscal federalism. *Books*, Edward Elgar Publishing, 1972. Citado na página page.4747.

OECD. Tax on property (indicator). 2023. Citado na página page.4545.

PARKER, S. W.; VOGL, T. S. Do conditional cash transfers improve economic outcomes in the next generation? evidence from mexico. *Working Paper*, September 2021. Citado 3 vezes nas páginas page.3434, page.3939 e page.4444.

PEREZ-TRUGLIA, R.; TROIANO, U. Shaming tax delinquents. *Journal of Public Economics*, Elsevier, v. 167, p. 120–137, 2018. Citado 2 vezes nas páginas page.4545 e page.4646.

POMERANZ, D. No taxation without information: Deterrence and self-enforcement in the value added tax. *American Economic Review*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, v. 105, n. 8, p. 2539–2569, 2015. Citado na página page.4646.

PORTA, R. L.; SHLEIFER, A. Informality and development. *Journal of Economic Perspectives*, v. 28, n. 3, p. 109–26, 2014. Citado na página page.44.

SAAVEDRA, M.; ANDRES, E. *Does It Matter Where You Grow Up? Childhood Exposure Effects in Latin America and the Caribbean*. Washington, DC, 2022. Accepted: 2022-05-13T15:39:38Z. Disponível em: <<https://openknowledge.worldbank.org/handle/10986/37415>>. Citado 2 vezes nas páginas page.44 e page.2828.

SAAVEDRA, M.; ANDRES, E. *The Geography of Intergenerational Mobility in Latin America and the Caribbean*. Washington, DC, 2022. Accepted: 2022-05-13T15:36:20Z. Disponível em: <<https://openknowledge.worldbank.org/handle/10986/37414>>. Citado 2 vezes nas páginas page.44 e page.2525.

SLEMROD, J. Tax compliance and enforcement. *Journal of Economic Literature*, v. 57, n. 4, p. 904–54, 2019. Citado na página page.4545.

SOLON, G. Intergenerational Income Mobility in the United States. *The American Economic Review*, v. 82, n. 3, p. 393–408, 1992. ISSN 0002-8282. Publisher: American Economic Association. Disponível em: <<https://www.jstor.org/stable/2117312>>. Citado 2 vezes nas páginas page.44 e page.8080.

SOLON, G. Intergenerational mobility in the labor market. In: *Handbook of labor economics*. [S.l.]: Elsevier, 1999. v. 3, p. 1761–1800. Citado na página page.44.

SOUZA, P. H. G. F. d. et al. Os Efeitos do Programa Bolsa Família sobre a pobreza e a desigualdade : um balanço dos primeiros quinze anos. <http://www.ipea.gov.br>, ago. 2019. Accepted: 2019-10-08T15:09:46Z Publisher: Instituto de Pesquisa Econômica Aplicada (Ipea). Disponível em: <<https://repositorio.ipea.gov.br/handle/11058/9356>>. Citado na página page.3333.

SUN, L.; ABRAHAM, S. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, Elsevier, v. 225, n. 2, p. 175–199, 2021. Citado na página page.5353.

ULYSSEA, G. Firms, informality, and development: Theory and evidence from brazil. *American Economic Review*, v. 108, n. 8, p. 2015–47, August 2018. Disponível em: <<https://www.aeaweb.org/articles?id=10.1257/aer.20141745>>. Citado 2 vezes nas páginas page.33 e page.55.

ULYSSEA, G. Informality: Causes and consequences for development. *Annual Review of Economics*, v. 12, n. 1, p. 525–546, 2020. Disponível em: <<https://doi.org/10.1146/annurev-economics-082119-121914>>. Citado 3 vezes nas páginas page.33, page.44 e page.55.

VIANA, I. A. V.; KAWAUCHI, M.; BARBOSA, T. V. *Bolsa Família 15 Anos (2003-2018)*. [S.l.: s.n.], 2018. Citado 3 vezes nas páginas page.3131, page.3535 e page.3939.

WEIDE, R. Van der et al. Intergenerational mobility around the world. *Available at SSRN 3981372*, 2021. Citado na página page.11.

World Bank. *The State of Social Safety Nets 2018*. Washington, DC: World Bank, 2018. ISBN 978-1-4648-1254-5. Disponível em: <<http://hdl.handle.net/10986/29115>>. Citado na página page.3131.

World Bank. *Poverty and Inequality Platform*. 2021. Citado na página page.55.

ZIMMERMAN, D. J. Regression toward mediocrity in economic stature. *The American Economic Review*, JSTOR, p. 409–429, 1992. Citado 2 vezes nas páginas page.22 e page.1010.

Intergenerational Mobility in the Land of Inequality

A.1 Appendix to Section 4

A.1.1 Description of data sources

- Person Registry: The *Cadastro de Pessoa Física* is the administrative population registry maintained by *Receita Federal*, the Brazilian tax authority. It contains all individuals who have ever held a Brazilian person code (CPF) – 255 million people in total. The CPF is similar to the social security number in the United States. Every individual in the country is identified by this unique and non-exchangeable code. Besides the person code, each observation has the person's full name, date of birth, gender, and the full name of the mother. If the person is dead, it contains the death year, which we use to create mortality outcomes.
- Address Registry: The tax authority provided us with a dataset containing the history of individuals' place of residence. The tax authority updates these addresses from several administrative sources, such as electoral registries and tax declarations, and when individuals autonomously update their information in the person registry. Each observation is identified by the individual's person code, the year when the address was updated and the full residential address (street name, number, apartment/house/unit, neighborhood and postal code). Overall, there are more than 500 million addresses, which we geocoded to longitude and latitude coordinates.
- Tax Returns: The tax authority also provided us with all personal income tax returns filed during the period 2006-2020. Each observation is identified by the returnee person code, all dependents' tax codes, and reported income divided into three categories: taxable income (mainly labor earnings and rents), tax-exempted income (mainly dividends, donations, and bequests), and income subjected to withheld or definitive taxation (mainly investment earnings and capital gains from real estate transactions). These data cover the period 2015-2019 for children in our main sample (cohorts 1988-1990) and the period 2006-2010 for their parents.
- Firm Ownership: The *Cadastro Nacional de Pessoa Jurídica* (CNPJ) is maintained by the tax authority and contains the universe of (formal) firms in Brazil, which are identified by a unique code (*CNPJ*), dates of opening and (eventual) closing, tax regime, city of registry, and a list of all shareholders identified by their person codes.
- Formal Employment: The *Relação Anual de Informações Sociais* (RAIS) is a linked

employer-employee administrative dataset covering the universe of firms and workers in the formal labor market, provided by the Ministry of Labor. We use all years of RAIS available, from 1985 to 2019. Employment spells are identified by the worker's person code and the firm's unique identifier (*CNPJ*),¹ workers' full name, gender, race, date of birth, and education; and complete information on the work contract such as dates of start and (eventual) termination, hours, wages, occupation.

- **Welfare Registry:** The *Cadastro Único* (CadÚnico) is an administrative registry maintained and constantly updated by the Ministry of Social Development to track the socioeconomic conditions of families with per capita income below half minimum wage or with total income below three minimum wages. It also includes all individuals of every family that has ever been a beneficiary of a federal social welfare program. We construct a yearly panel of CadÚnico from 2011 to 2020 with the individual's full name, gender, birth year, race, education, and mother's and father's full names for more than 135 million individuals identified by their person codes. Each household is also identified by a unique identifier, allowing for the recovery of family structures.
- **Hospitalization Records:** Individual-level data on admissions to public hospitals SIH-SUS (*Sistema de Internações Hospitalares*) for the period 2002-2019. It includes information on individual characteristics such as age, sex, municipality and zip code of residence, and descriptive information on the hospital admission, including the ICD-10 diagnostic, and date of admission. We use ICD-10 codes on hospitalization due to assaults to generate a measure of crime victimization. To merge these records to other datasets, we focus on individuals who can be uniquely identified by their postal code, gender and birth date – all of which can be observed for the entire population in the person registry, maintained by the Brazilian Tax Authority.

A.1.2 Family links


Our main analysis is based on a conservative family linkage procedure focused on avoiding any erroneous links. For robustness purposes, we assemble an expanded sample based on additional, less conservative linkages between parents and children. Namely, we proceed by rounds and expand our conservative family linkages by additionally matching children to mothers using the following information: (i) mother's name in the person registry, conditional that mother has a unique name within the postal code where the child lives; (ii) mother's name and state (uf) in the person registry, conditional that mother has a unique name in the state of residence; (iii) mother's name in the welfare registry, conditional that mother has a unique name in the country; (iv) mother's name in the welfare registry, conditional that mother has a unique name in the state of residence; (v) mother's name in the 2014 School Census, covering all enrolled students in Brazilian schools, conditional that the mother has a unique name in the country; (v) mother's name in the 2014 School Census, conditional that the mother has a unique name in the the state of residence; (vi) household composition in welfare registry. We follow the same procedure

¹From 1985 to 2001, workers are identified by a different (unique) code, the PIS. We retrieve PIS-CPF pairs for all workers matching individuals across RAIS waves by their full name and date of birth.

for fathers, with the exception that rounds i-iii are not available because fathers' names are not available in the person registry. Although these linkages are somewhat less conservative relative to our baseline, conservative linkage, they remain highly accurate as they are based on high quality data on names and addresses.

Figure A.1 plots the share of children from each cohort linked to their parents when using either procedure. Many more children can be linked to their parents – particularly mothers, since mothers' names are available for the entire population whereas fathers' names are available for roughly two-thirds of the population in the welfare registry. In addition, the share of successful links is increasing over time because younger cohorts can be claimed throughout more childhood years in the tax data, which start in 2006.

Figure A.1: Number of parent-child links relative to the population by cohort



Note: The figure plots the share of the population that can be linked to their parents by cohort following our baseline, more conservative method (left graph) and the alternative, less conservative method (right graph). The first method links children to parents using unique person codes in dependent claims tax data and using names for uniquely named parents in population and welfare registries. The second method allows for additional links using individual names and addresses.

A.1.3 Sample selection

Columns 1-2 in Table A.1 provide descriptive statistics for the population and our main sample. The standardized differences in column 3 are below the critical value .2 for all but three variables that slightly exceed the cutoff (race, college education and living in the North-East), indicating only small differences in the underlying distributions (COHEN, 2013). Nevertheless, in light of these small differences, we will show as a robustness test that our main findings are unaffected: (i) when substantially enlarging the sample by using the less conservative procedure to link families (see Section A.1.2 above), and (ii) re-weighting the sample to perfectly match the first and second moments of several characteristics in the population using the entropy algorithm by (HAINMUELLER, 2012) (Table A.1, columns 4-5) – see Section A.3.3.1.

Table A.1: Descriptive statistics of main sample

	Population	Sample	Std. Diff.	Weighted	Std. Diff., W
Female	0.494	0.512	0.037	0.494	0.000
Non-White	0.532	0.417	0.232	0.532	0.000
Primary	0.060	0.033	0.130	0.060	0.000
Elementary	0.350	0.270	0.172	0.350	0.000
High School	0.501	0.503	0.004	0.501	0.000
College	0.089	0.194	0.305	0.089	0.000
Welfare	0.624	0.581	0.086	0.622	0.003
Formal Job	0.858	0.899	0.124	0.856	0.006
Cohort 1988	0.344	0.319	0.054	0.344	0.000
Cohort 1989	0.337	0.336	0.001	0.337	0.000
Cohort 1990	0.319	0.345	0.055	0.319	0.000
North	0.089	0.077	0.044	0.089	0.000
Northeast	0.277	0.192	0.201	0.277	0.000
Southeast	0.414	0.453	0.080	0.414	0.000
South	0.143	0.195	0.140	0.144	0.005
Center-West	0.077	0.082	0.018	0.076	0.006
State capital	0.251	0.290	0.088	0.251	0.000

Note: The table compares the average characteristics of our main sample of children born in 1988-1990 (omitting missing values) with the average characteristics of the same cohorts in the general population. The means for each variable are presented (columns 1-2), along with the standardized difference (column 3), the mean in the main sample after re-weighting observations to match the first and second moments of population characteristics (HAINMUELLER, 2012) (column 4), and the standardized difference between the samples in columns 1 and 4. All variables are recorded as dummy indicators.

A.1.4 Imputation method based on random forests

We use PNAD and population censuses to assemble a repeated cross-section from 1991 to 2019 of all adults aged 18-65 in any occupation – formal, informal, firm owner, or self-employed. We leverage the high-quality information contained in these surveys to train a generalized RF model (ATHEY; TIBSHIRANI; WAGER, 2019) to predict informal income each year. We repeat the same process for estimating formal non-labor income (which is used when tax data are not available).

An RF is a collection of trees, each one endogenously splitting the covariate space to predict our outcomes of interest. To generate each tree, the algorithm starts by sampling without replacement from the survey dataset defining a root node. The root node is split over the space of covariates into child nodes as follows. A random subset of covariates are selected as candidates to split on, and the algorithm selects the split that maximizes heterogeneity in the prediction outcome. The sample splits take place recursively until a stop criterion met, so that overfitting is avoided.


For each outcome (informal income and formal non-labor income) and each year of data, we grow a RF with 5,000 trees, feeding the algorithm with a random subset of the original survey data. We grow *honest* forests, meaning that separate sets of data, selected at random, are used for splitting the node and estimating prediction improvements (ATHEY; TIBSHIRANI;

WAGER, 2019). We tune all parameters of the model via cross-validation – namely, the number of variables selected at each split, the minimum node size, the penalization for imbalanced splits, and the subsample fraction for honest splits. The model covariates are state of residence (27), state capital dummy, gender, age, four education dummies, a white/non-white dummy, and worker category. Following the literature on the representation of categorical variables in ML models (JOHANNEMANN et al., 2019), we encode state of residence dummies as a set of real-valued covariates related to demographic and socioeconomic indicators.² We build the RF model to perform out-of-sample predictions of unobserved income components: namely, informal and formal non-labor income.

A.1.5 Imputation method accuracy

We use our RF prediction model to impute informal income and, when not available in administrative records, formal non-labor income. Figure A.2 plots the share of imputed income by income deciles for parents and children in our main sample. As expected, this share is highest at the bottom and decreases over the income distribution due to informal income being more prevalent among low-income individuals.

Figure A.2: Average share of imputed income by income decile



Note: The figure plots the average share of income imputed by income decile for children (blue dots) and parents (red triangles) in our main sample. Imputed income comprises informal income and, when not available in administrative data sources, formal non-labor income.

To assess the accuracy of our measure for income ranks, we use a random test sample from our survey for the period 1991-2019 – which has not been used to train the model so to avoid

²We define formally self-employed as owners of formal firms with zero employees, and firm owners as those owning firms with at least one formal spell in the year. We consider all working-age adults who are not formal workers or firm owners as informal workers (i.e., the residual occupation category). For census years, we use more detailed information on the municipality of residence instead of the state of residence, since the former is available in the data.


overfitting issues.³ For each year in this data, we rank individuals into income percentiles based on their total income. Next we rank them again after replacing their informal and formal non-labor income with the model predictions – emulating the procedure for estimating income ranks in our main analysis – and then compute accuracy statistics. These estimates yield a R-squared of .57 and a root mean squared error (RMSE) of 19 ranks. For comparison, we repeat the procedure and compute the same statistics using a fully-saturated OLS prediction model – i.e., interacting fixed-effects for all covariates –, which yields a .29 R-squared and RMSE of 24 ranks. The OLS model is less accurate because of overfitting issues. Repeating this exercise on the sample used to train the model results in a much higher R-square for the OLS model (.75 in-sample vs. .29 out-of-sample), while it makes little difference for our RF model (.59 in-sample vs. .57 out-of-sample). The fairly high R-squared of our model helps reducing potential biases in mobility estimates, as pointed by earlier literature studying the consequences of income imputation for mobility estimates (e.g., see (SOLON, 1992; INOUE; SOLON, 2010; JERRIM; CHOI; SIMANCAS, 2016).

Moreover, this exercise estimates accuracy statistics on income rank predictions for a single year, while our main analysis averages out income from several years before ranking individuals. We argue that such procedure further increases the accuracy of our income measures, as it reduces the influence of transitory measurement error components. We cannot demonstrate this point based on the (cross-sectional) survey data because it does not follow the same individuals over time. Hence, we provide evidence on this by predicting formal labor income for individuals in administrative employment data (RAIS) using our RF model. Then, we compare the R-square and RMSE for income rank predictions for parents and children in our main sample averaged over the periods 1995-1999 and 2015-2019, respectively. This exercise shows that averaging income over these periods for parents and children increases the R-squared by 27.5% and 12.2% and reduces the RMSE by 18.6% and 6.7%, respectively. Hence, averaging predictions over multiple years leads to significantly lower measurement error.

³Since we use all survey observations to train the model for our main analysis, we re-estimate again the prediction models on a 50% random sample of the survey data, and use the remaining observations to generate out-of-sample accuracy statistics.

A.1.6 Income distribution in our main sample and survey data sources

Figure A.3: Income Distribution



Note: The graphs compare the income distribution in our main sample (blue dots) and in the PNAD survey (red triangles), separately for children, fathers, and mothers. The PNAD sample includes “children” born in 1988-1990 that were interviewed in the wave 2019, and “parents” of children born in 1988-1990 that were interviewed in the wave 2005. PNAD sample weights are used to compute the income distribution in PNAD data.

A.2 Appendix to Section 5

A.2.1 Decomposition of biases due to measurement error

We formally study how measurement error in child and parental income may bias our estimates of the rank-rank slope based on equation (1.1). Since $\alpha = 50(1 - \beta)$, it is straightforward to extend the decomposition to other IGM measures based on the rank-rank regression. Our measures of child (y) and parental (p) income ranks in equation (1.1) are measured with error as $y = y^* + \eta$ and $p = p^* + \mu$. Actual child and parental income are defined by y^* and p^* , and measurement errors are defined by η and μ , respectively. We make no assumption on the distribution of the errors. Since we observe income with measurement error, in practice we regress $y^* + \eta$ on $p^* + \mu$. This yields a OLS estimate for β with the following bias:

$$\begin{aligned}
 \hat{\beta} - \beta &= \frac{c(y^* + \eta, p^* + \mu)}{v(p^* + \mu)} - \beta \\
 &= \frac{c(y^*, p^*) + c(y^*, \mu) + c(\eta, p^*) + c(\eta, \mu)}{v(p)} - \beta \\
 &= \beta \left(\frac{v(p^*)}{v(p)} - 1 \right) + \beta_{y^*\mu} \frac{v(\mu)}{v(p)} + \beta_{\eta p^*} \frac{v(p^*)}{v(p)} + \beta_{\eta\mu} \frac{v(\mu)}{v(p)}
 \end{aligned} \tag{A.1}$$

where β is our coefficient of interest (i.e., the regression of y^* on p^*) and β_{ab} denotes the coefficient of a hypothetical OLS regression of a on b , and $v(\cdot)$ and $c(\cdot)$ denote the variance and covariance operators, respectively.

This decomposition can be applied for different mobility measures based on income measured in any form: ranks, levels or logs. For the case of ranks, the variance of income ranks measured with error equals, by construction, the variance of actual income ranks. This is because the rank

measure always takes the same values: they range from 1 to 100, grouping the population into their income percentile. Hence, $v(p) = v(p^*)$ which implies that:

$$\begin{aligned} v(p) &= v(p^* + \mu) = v(p^*) + v(\mu^*) + 2c(p, \mu) \iff \\ -v(\mu) &= +2c(p^*, \mu) \iff \\ \beta_{p^*\mu} &= -1/2 \end{aligned}$$

Using this result and $v(p^*) = v(p)$ in eq. (A.1), we immediately have our final bias-decomposition formula:

$$\hat{\beta} - \beta = -\frac{1}{2}\beta \frac{v(\mu)}{v(p)} + \beta_{\varepsilon\mu} \frac{v(\mu)}{v(p)} + \beta_{\eta p^*} + \beta_{\eta\mu} \frac{v(\mu)}{v(p)}$$

A.3 Appendix to Section 6

A.3.1 Estimates fully based on years when tax data are available

Table A.2: IGM estimates fully based on years when tax data are available

	Rank-rank slope (1)	Exp. rank p=25 (2)	Q1Q5 (3)	Q5Q5 (4)	IGE (5)
Estimate	0.537	36.8	2.0%	49.5%	0.557

Note: The table shows mobility estimates obtained using only years when tax data are available, using our main sample (1988-1990 cohorts). Parental income is measured in the period 2006-2010, when tax data are available and children are in age ranges 16-20 (1990 cohort), 17-21 (1989) and 18-22 (1988). Children's income is measured as in our baseline at ages 25-29. Q1Q5 (Q5Q1) defines the probability that children born in income quintile 1(5) reach income quintile 5(1) in adulthood.

Table A.3: Quantifying IGM estimation biases due to measurement error: the impacts of informal and formal non-labor income imputation

		Replacing income components with predicted counterparts		
	Benchmark (based on survey data)	Informal and formal non-labor	Informal	Formal non-labor
	(1)	(2)	(3)	(4)
PANEL A. RELATIVE MOBILITY				
Rank-rank slope (RRS)	0.520	0.516	0.514	0.520
SE	0.004	0.004	0.004	0.004
ME bias decomposition				
Term 1: $-\frac{1}{2}\beta \frac{v(\mu)}{v(p)}$		-0.150	-0.136	-0.013
Term 2: $\beta_{\varepsilon\mu} \frac{v(\mu)}{v(p)}$		0.134	0.123	0.011
Term 3: β_{vp^*}		-0.016	-0.017	0.001
Term 4: $\beta_{v\mu} \frac{v(\mu)}{v(p)}$		0.028	0.024	0.001
Total bias		-0.005	-0.007	-0.001
PANEL B. OTHER IGM MEASURES				
Exp. rank p=25	35.23	35.32	35.37	35.24
Q1Q5	5.5%	3.5%	3.9%	5.0%
Q5Q5	47.1%	48.1%	48.1%	46.2%
IGE	0.100	0.393	0.300	0.102
Observations	45,718	45,718	45,718	45,718

Note: This table shows how benchmark relative mobility estimates (Panel A) and several IGM measures (Panel B) change after replacing income components with predicted counterparts for different groups in a sample of cohabiting parents and working children aged 25-34 in PNAD survey data. Column 1 reports the benchmark estimates, while columns 2-5 reports IGM estimates after replacing income components. Q1Q5 (Q5Q1) defines the probability that children born in income quintile 1(5) reach income quintile 5(1) in adulthood. Panel A also provides a decomposition of the total bias resulting from income imputation following the decomposition presented in Section 1.5.1

A.3.2 Quantifying measurement error biases, additional IGM measures

Table A.4: Quantifying IGM estimation biases due to measurement error: the impacts of formal income imputation, additional IGM measures

		Replacing income components with predicted counterparts for individuals in different income quartiles					
		Formal labor income (all)			Formal labor income (all) and formal non-labor income (parents only)		
Benchmark		Q1	Q1-Q3	All	Q1	Q1-Q3	All
(1)		(2)	(3)	(4)	(5)	(6)	(7)
ADDITIONAL IGM MEASURES							
RRS - Female	0.611	0.615	0.628	0.632	0.614	0.616	0.597
RRS - Male	0.469	0.472	0.484	0.496	0.472	0.476	0.472
Gender gap, p=25	-17.5	-17.9	-18.8	-19.0	-17.9	-18.7	-18.8
Gender gap, p=75	-10.4	-10.8	-11.6	-12.2	-10.8	-11.7	-12.6
RRS - Non-White	0.571	0.580	0.589	0.585	0.580	0.577	0.550
RRS - White	0.520	0.528	0.544	0.545	0.528	0.533	0.509
Race gap, p=25	-8.48	-8.84	-9.24	-9.89	-8.85	-9.50	-10.56
Race gap, p=75	-5.92	-6.25	-6.96	-7.88	-6.25	-7.29	-8.48
Cor. Regional Ab. p=25	1.000	0.998	0.994	0.989	0.998	0.993	0.986

Note: This table shows how additional IGM measures change after replacing formal income with predicted counterparts for different groups in our main sample (Panel B). Column 1 reports our main estimates, while columns 2-6 reports IGM estimates after replacing formal income for specific groups of parents and children based on their income quartiles. The last row shows the correlation of relative and absolute mobility across Brazil's 510 IGRs in each simulation (columns 2-6) with our benchmark estimates (column 1).

Table A.5: Quantifying IGM estimation biases due to measurement error: the impacts of informal income and formal non-labor income imputation, additional IGM measures

	Benchmark (based on survey data) (1)	Predicting (replacing) income components		
		Informal and formal non-labor (2)	Informal (3)	Formal non-labor (4)
		PANEL A. ADDITIONAL MEASURES		
RRS - Female	0.511	0.536	0.533	0.511
RRS - Male	0.534	0.513	0.511	0.533
Gender gap, p=25	-3.58	-5.57	-5.38	-3.74
Gender gap, p=75	-4.69	-4.39	-4.28	-4.83
RRS - Non-White	0.476	0.430	0.439	0.469
RRS - White	0.481	0.501	0.490	0.485
Race gap, p=25	-8.86	-7.08	-7.56	-8.61
Race gap, p=75	-9.12	-10.63	-10.12	-9.38

Note: This table shows how benchmark relative mobility estimates (Panel A) and several IGM measures (Panel B) change after replacing income components with predicted counterparts for different groups in a sample of cohabiting parents and working children aged 25-34 in PNAD survey data. Column 1 reports the benchmark estimates, while columns 2-5 reports IGM estimates after replacing income components. Panel A also provides a decomposition of the total bias resulting from income imputation following the decomposition presented in Appendix A.2.1.

A.3.3 Additional robustness exercises

We now present a series of additional robustness exercises that further support our main results. In some of these analyses, we use additional birth cohorts born in the 1983-1990 period, additional parent-child links, and also vary the period when income is measured. Since we were granted access to tax data on our cohort of parents only for the period 2006-2010 and on our cohort of children only for the period 2015-2019, we rely on the procedure laid out in Section 1.4.2 to measure formal income when running these robustness. The procedure sums up formal labor income based on formal employment data and predicted formal non-labor income based on our RF prediction model. Therefore, we first show our benchmark mobility estimates for running these robustness tests in column 1 of Table A.6, Panel A. The rank-rank slope is .453, somewhat smaller relative to our main estimates (.546).

A.3.3.1 Sample selection

LARGER SAMPLES. Our baseline sample comprises 1.3 million children born during the period 1988-1990 that we can link to both parents, comprising 15% of all children in such cohorts. We show that our main results are robust to expanding the sample along three dimensions. First, we include all children that can be linked to their father (regardless of whether they are linked to their mother), which increases the sample size by 1 million, and run the analysis solely based on the father's income. The results in columns 2-3 of Table A.6, Panel A, show that the father-child rank correlations in the baseline and enlarged samples are nearly identical (.44 and .45), and they are also identical to the baseline rank-rank slope estimated without tax data, reported in column 1.

Second, we expand the sample to include all cohorts born in 1983-1992, for a total of 6.9 million children. In this case as well, the estimated rank-rank coefficient remains identical.

Finally, Panel B of Table A.6 replicates the analysis using the less conservative linking

procedure described in Appendix Section A.1.2. This increases the data coverage for our main cohorts (1988-1990) from 15% to 45%. Once again, all estimated coefficients are virtually unaffected (0.44 - 0.47).

Table A.6: Robustness to larger samples

	1988-1990 cohorts			1983-1992 cohorts		
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. Baseline links</i>						
Parent Rank	0.459***			0.480***		
Father Rank		0.445***	0.455***		0.467***	0.472***
Observations	1,304,586	1,304,586	2,361,010	3,797,639	3,797,639	6,949,075
Only father links			Yes			Yes
<i>Panel B. Complete links</i>						
Parent Rank	0.459***			0.468***		
Father Rank		0.440***	0.448***		0.448***	0.456***
Observations	3,416,131	3,416,131	3,901,433	9,976,431	9,976,431	11,478,370
Only father links			Yes			Yes

Note: The table reports the estimated slope of the rank-rank regressions in equation (1.1), i.e. our (inverse) measure of income mobility, in different samples. In Panel A, we link parents to children using our baseline, conservative procedure; in Panel B, we expand the sample using the less conservative procedure described in Section A.1.2. Columns 1-3 cover the 1988-1990 cohorts – as in our main sample – while columns 4-6 cover the 1983-1990 cohorts. Columns 1-2 and 4-5 are based on children who can be linked to both parents, while columns 3 and 6 are based on children who can be linked at least to their fathers. The dependent and explanatory variables are always the child and parental income percentile rank. For consistency, in all specifications we measure income without using tax data, which are only available for the 1988-1990 cohorts. Income for the 1988-1990 cohorts is measured between 2015-2019, at ages 25-31 (as of baseline), while income for other cohorts is measured when they are 25-29 years old (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

REWEIGHTING. To address any residual concern about the representativeness of our sample, we re-weight the data to match a rich set of characteristics (up to their 2nd moment) in the general population, using the algorithm proposed by (HAINMUELLER, 2012). Specifically, we balance our baseline sample with respect to gender, race, month and year of birth, state of residence (27), state capital dummy, education (4), and indicators for being in social welfare registries, formal labor market participation, and having a unique name in the country.⁴ In Table A.7, we report estimates of the rank-rank coefficient on the raw and re-weighted data (row 1 and 2, respectively) for the full sample (column 1), males (2), females (3), whites (4), and non-whites (5). Reweighting leads only to small changes in the estimated rank-rank slope in our full sample (from .546 to .566) or subgroups of the population.

⁴Table A.1 in Section A.1.3 displays descriptive statistics of the sample before and after reweighting together with standardized differences with respect to the population.

Table A.7: Reweighting Procedure

	Full Sample	Males	Females	Whites	Non-whites
	(1)	(2)	(3)	(4)	(5)
Baseline	0.546***	0.469***	0.613***	0.521***	0.573***
Weighted	0.566***	0.504***	0.624***	0.546***	0.574***
Observations	1,304,586	633,489	671,097	672,015	546,773

Note: The table reports the estimated slope of the rank-rank regressions in equation (1.1), i.e. our (inverse) measure of income mobility, in the raw data (first row) and in the re-weighted data (second row). The re-weighted data match the first and second moments of the population of children in the same birth cohorts along several characteristics, using the entropy algorithm by (HAINMUELLER, 2012) (see Table A.1). All samples cover the 1988-1990 cohorts and the dependent and explanatory variables are the child and parental income percentile rank (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).


A.3.3.2 Timing of income measurement

The timing when we measure income may lead to two main types of estimation bias, which we describe next. We provide several tests showing that our main results are not substantially affected by the timing of income measurement. We estimate income mobility without relying on tax data, as the latter are not available for some of the cohorts and years required for such tests.

ATTENUATION BIAS. Measuring income for short time spans may attenuate estimates of income mobility due to temporary income shocks. This is not a main concern when we measure parental income since we virtually cover children's entire childhood (from age 3 to 18). However, this could be a more relevant issue for child income in our main analysis, which uses a five-year window (age ranges 25-29, 26-30, 27-31 for the cohorts born in 1988, 1989 and 1990, respectively). In light of this, we show in Figure A.4a how the rank-rank slope changes as we vary the number of years used to measure parental income. The estimates are remarkably stable regardless of how many years are used in the analysis. Moreover, in Figure A.4b we show that estimates also remain largely stable when using 1 to 5 years to measure children's income. The estimates vary by less than 5 percentage points in both exercises, relative to the 0.453 rank-rank slope benchmark (without tax data). Overall, these results support the idea that the five-year window used to measure child income is sufficient to prevent meaningful attenuation bias in our main analysis.

LIFE-CYCLE BIAS. Measuring income too early may not adequately capture permanent income, possibly leading to life-cycle bias (HAIDER; SOLON, 2006; MELLO; NYBOM; STUHLER, 2021). Again, this could be relevant when measuring children's income with a five-year window at the age range 25-31 in our main analysis. We use the fact that we can track parental income for a long period of time to study how our estimates change when measuring parental income at different ages. In particular, we focus on parents in our main sample born in 1960-1965. Figure A.4c shows that our estimates do not vary much when using a 3-year window to measure father's income centered from age 31 to 45. Next, in Figure A.4d, we show how the rank-rank slope changes as we center a three-year window around different ages for measuring children's income. We focus on the 1988 cohort, for which we can track income up to age 31. Again, the estimated rank-rank slope remains fairly constant within cohorts when income is measured at varying ages.

Figure A.4: Sensitivity of child and parental income to timing




Note: This figure plots robustness exercises for attenuation bias (a, b) and life-cycle bias (c, d). In Panels (a) and (c) child income is held constant and measured as in our baseline estimates, and we vary how parental income is measured. In Panels B and D, we measure parental income as in our baseline and vary how child income is measured. Panel A displays estimates of the rank-rank slope from separated rank-rank regressions in which we vary the number of years used to compute parental income, from 1 to 17 years, and centered at the age of 11. Panel (b) displays an analogous exercise in which we measure children's income using from 1 to 5 years, centered at age 27. In Panel (c), we run rank-rank estimates using father's income (rather than parental income) and vary the age when father's income is measured using a three-year window from ages 31 to 45. Finally, in Panel (d) we vary the age when we center the three-year window to measure children's income, from 20 to 30 years old. In Panels (a) and (b) we use our full baseline sample of the 1988-1990 cohorts. In Panel (c), we restrict the sample to children whose fathers are born between 1960-65 and focus on fathers' rather than parental income – to precisely gauge the sensitivity concerning different age windows. In Panel (d), the working sample is the 1988 cohort since income data is only available until 2019.

This result can be explained by the fact that there are little positional changes in annual income in Brazil from early ages, especially from the age of 24 (Figure A.5a); this is due, in turn, to the fact that most Brazilians enter the labor market relatively early, given that college enrollment is low. Positional changes below the age of 24 are concentrated at the very top of the income distribution, and they are driven by a large share of high-income children who attend college and delay entry into the labor market (see also Figure 1.9a). Figure A.5b provides additional evidence on these aspects by showing a near parallel evolution of income by quartiles of completed education and parental income. These patterns are in contrast with the case of developed countries – as documented by (MELLO; NYBOM; STUHLER, 2021) for the US

and Sweden – where a much larger share of individual attend college. In fact, (GUVENEN; PISTAFERRI; VIOLANTE, 2022) document that Brazil displays the highest intragenerational persistence in income in a group of 13 middle- and high-income countries – for instance, 16% and 30% larger than the U.S. and Sweden respectively.

Figure A.5: Evolution of children’s position at the income distribution



Note: The figure plots the evolution of children’s income distribution over time. Panel (a) shows the mean income percentile rank (on the vertical axis) when aged 18 to 36 (horizontal axis) for individuals in each decile of total lifetime income distribution. In turn, panel (b) shows the evolution of log incomes of the 1983 cohort when 22 to 36 years old, by quartiles of children’s educational level (left) and parental income (right).

A.3.3.3 Alternative income and occupation definitions

We now show that two potentially relevant choices that we take to define total annual income have virtually no impact on our IGM estimates. Table A.8 presents our baseline mobility estimate (column 1), along with alternative estimates that we describe next (columns 2-4). First, we rely on survey questions on “normal” monthly income to predict annual informal income and formal non-labor income (used when tax data are not available). In our main analysis, we extrapolated such income to the entire year, multiplying it by 12. We show that results do not change if we adopt a similar procedure for measuring formal labor income (derived from administrative employment data). Namely, we take the average monthly formal income while formally employed and multiple by 12 each year, instead of considering the sum over the year (column 2). Alternatively, we move back to our baseline but multiply predicted monthly informal income and formal non-labor income by the number of months that individuals spend out of formal employment in the year (rather than by 12) (column 3). Second, when predicting unobserved income in the main analysis, we label as informal workers in the administrative data those who do not hold any formal job in the entire year and who are not firm owners. We vary this assumption by defining informal workers as those who work formally for less than three consecutive months in the year and who are not firm owners (column 4).

Table A.8: Alternative Income Measures

	(1)	(2)	(3)	(4)
Baseline	0.546*** (0.001)	0.545*** (0.001)	0.547*** (0.001)	0.547*** (0.001)
Income definition	Baseline	Alternative 1	Alternative 2	Baseline
Informal workers definition	Baseline	Baseline	Baseline	Alternative
Observations	1,304,586	1,304,586	1,304,586	1,304,586


Note: The table reports relative mobility estimates based on the slope of rank-rank regressions – as in eq. (1.1) – using different income and occupation definitions. It presents estimates when using our baseline income and occupation definitions (column 1); when measuring formal labor income by multiplying its monthly average in each year by 12 (column 2), when measuring predicted informal and formal non-labor income by multiplying the predicted monthly quantities by the number of months out of formal employment in the year (column 3); when defining informal workers as those who are formally employed for less than three months in the year and who are not firm owners. All samples cover the 1988-1990 cohorts and the dependent and independent variables are the child and parental income percentile rank (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

A.3.4 Individual vs. household income for child ranks

We assess whether using household instead of individual income for children affects our results. We focus on the 38% of children who can be linked to their spouses in tax declarations and welfare registries (*CadÚnico*). We compute our baseline income measures for spouses in the same calendar years that their partner's income is measured starting from the first year when we observe both together. Household income is defined as the sum of both partners' individual income. Figure A.6 plots the mobility curves using individual (baseline) and household income for children in the married sample.

Both curves perfectly overlap and are similar to the baseline mobility curve based on our main sample (Figure 1.1). Overall, the exercise indicates that taking household or individual income has little impacts on our IGM estimates, which only become slightly larger, increasing from .529 to .58. They are also similar to our main estimates based on children's individual income, using the full sample rather than the married sample (.546).


Figure A.6: Mobility curve using household income for children



Note: The figure plots the mobility curves using individual (blue dots) and household (red triangles) income for children, while parental income is the sum of father and mother's income, as in our baseline. The sample comprises the individuals in the baseline sample for which we are able to recover partners from tax declarations or *CadÚnico*. Household income is the sum of both partners' individual annual income starting from the year we observe them as a couple. For each curve, the figure also reports our relative mobility measure based on Equation 1.1.

A.3.5 Neighborhood-based measure for movers

Figure A.7: Neighborhood-based measure: children changing address



Note: This figure plots mobility curves based on the neighborhood-based measure, for our main sample (red triangles) and for a sample of movers who live in a different zip code than the place where they grew up (blue dots), both covering the 1988-1990 cohorts. For each parental rank in each curve, it plots the mean child rank. The neighborhood-based measure is given by the average formal income in the census tract where children grew up (parental rank) and where they live as adults (child rank). See Section 1.6.4 for a detailed description of these measures. For each curve, the figure also displays our relative mobility measure based on Equation (1.1).

A.3.6 Additional results**Table A.9: Labor market differences by gender and race**

Parent Quintile	Gender			Race		
	Rank Gap	LFP Gap (pp.)	Wage Ratio	Rank Gap	LFP Gap (pp.)	Wage Ratio
1	17.7	17.6	0.84	10.4	10.0	0.95
2	18.8	13.2	0.84	9.2	8.4	0.95
3	15.2	8.2	0.84	8.6	7.2	0.96
4	10.8	4.6	0.87	6.4	6.3	1.01
5	7.1	1.9	0.89	3.4	6.5	1.12

Note: The table reports average gaps in child income ranks and labor market outcomes over gender and race, for each parental income quintile. Income rank gaps are calculated as the difference between average adult ranks for males (whites) and females (non-whites). The labor force participation (LFP) gap is the difference in average participation rate in the formal labor market between the two groups, in percentage points. Finally, the wage ratio is the ratio of the formal average wages of females (non-whites) to males (whites).

Table A.10: Siblings comparisons by parental income quintile


Parental Quintile	Siblings Gap	Brother-Sister Gap
1	0.11	16.23
2	0.22	19.00
3	0.19	16.67
4	1.39	11.79
5	2.36	7.29

Note: The table reports average gaps in income ranks between siblings for each parental income quintile. Siblings gaps are calculated as the difference between adult income rank of the older and younger siblings, regardless of gender. Brother-sister gaps are calculated as the difference between the male and female siblings, regardless of birth order. Both are calculated for individuals in our baseline sample of the 1988-1990 cohorts.

A.4 Appendix to Section 7

A.4.1 Individual mobility curves

Figure A.8: Individual mobility curves in Fortaleza and Belo Horizonte




Note: The figure plots separate mobility curves for Fortaleza (red triangles) and Belo Horizonte (blue dots). Both curves are non-parametric binscatters constructed by plotting mean child income rank for children born in each parental percentile income rank in both regions. The figure is based on our main 1988-1990 cohorts sample. Income is defined as our baseline measure and both children and parents continue to be ranked according to the respective national income distribution. Children are assigned to regions according to the location of their parents in 2000, regardless of where they live in adulthood.

A.4.2 Robustness of subnational estimates, price differences

We construct a regional price index to generate mobility maps that account for price differences across Brazil. To create the index, we use the POF (*Pesquisa de Orçamentos Familiares*), a household budget survey conducted by IBGE. POF gathers rich demographic and expenditure data at a fine geographic level. We use the 2003 and 2019 editions to calculate – across Brazilian areas – the average price of the reference basket used to compute the main Brazilian consumer inflation index (IPCA). Specifically, we compute prices at the state level, distinguishing between the (state) capital, the metropolitan areas around the capital, and the countryside. Next, we rescale parents' and children's income by the index computed for 2003 and 2019, respectively, according to their location. Finally, we re-estimate our regional mobility measures based on the price-adjusted income and correlate them with our main estimates. Figure A.9 shows that such adjustment has little impact on regional mobility patterns. It shows that absolute mobility remains similar across space, and that both absolute and relative mobility are strongly correlated with our original, region-specific, mobility measures (see Figure 3.1a). The correlations for our relative and absolute mobility measures are as high as .92 (.906,.933 - 95% C.I.) and .93 (.917,.941 - 95% C.I.), respectively.

Figure A.9: Price-Adjusted Absolute Mobility Map



Note: The figure displays price-adjusted absolute mobility – scaled by deciles – in Brazil’s 510 immediate geographical regions (IGRs) for our main sample (1988-1990). Parent and child incomes are deflated by regional price indexes constructed with POF survey data and ranked in the national income distribution (measured when children are aged 3-18 and 25-31, respectively). Absolute mobility indicates the expected rank for children in below-median income families, based on Equation (1.1). Darker green tones indicate higher absolute mobility. Children are assigned to IGRs according to the location of their fathers in 2000.

A.4.3 Mobility estimates for the 50 largest metropolitan areas

A.4.4 Mobility correlates

We explore the correlates of social mobility by estimating univariate regressions of absolute mobility on a wide range of local indicators covering thirteen broad categories: demographics, economic structure, education, family structure, health, household, income, inequality, local infrastructure, labor market, municipal budget, public safety, and social capital. Table A.12 provides a detailed description and data sources for the variables in each category. Figure A.10 plots the results of these regressions when normalizing both the dependent and explanatory variables so that coefficients can be interpreted as correlations and more easily compared with each other.⁵ Overall, coefficients have the expected sign, and nearly all of them are statistically significant. Several variables related to education quality show up among the top mobility predictors – in particular, literacy rates and students’ performance in standardized test scores. In line with the analysis by race in Section 1.6.6, the racial composition is also a strong mobility predictor: the share of white population displays the second highest correlation, while the share of black and mixed-race individuals yield negative coefficients. Other variables related to the number of formal firms per capita, the number of bank agencies, and labor market participation by men are also among the strongest mobility predictors. In turn, markers of socioeconomic struggle such as large or high density households, and the share of individuals without earnings are among the top predictors of low mobility, followed by the GDP share of the public sector.

⁵Specifically, we recenter them around the mean and rescale them so that their standard deviation is equal to one.

One difficulty when interpreting these results is the strong correlation between the indicators considered. Thus, we reduce the dimensionality of the problem in two steps. First, we create a single index for each category based on the principal components of the initial variables, similarly to (ACCIARI; POLO; VIOLANTE, 2021).⁶ Figure A.11 report the results of multivariate regressions of absolute mobility on such indexes. Education quality yields the largest correlation with absolute mobility by far, with a positive sign (blue coefficients). Other categories showing strong correlation with mobility are the indexes related to the family structure, demographics (including the racial composition), household characteristics and local infrastructure. Once we control for region fixed effects (5 categories), the education index continues to stand out relative to other factors, being by far the strongest mobility predictor within regions (light blue coefficients).

Nonetheless, the regressors in Figure A.11 still have a high degree of multi-collinearity, which could harm the interpretation of the results. Hence, we employ a standard LASSO regularization procedure to select robust predictors of mobility. We validate the choice of parameters for the LASSO regression via ten-fold cross-validation. Figure A.12 summarizes the results of such an exercise, displaying the regressors' coefficients (y-axis) against increasing values of the regularization parameter (x-axis). As we increase the penalization for the number of regressors, coefficients are shrunk toward zero and variables leave the model. Again, the education index dominates the other variables.

A.5 Appendix to Section 8

A.5.1 Sample construction

For the analysis on causal place effects, we focus on children born in the 1983-1992 period who we can link to their fathers (following the baseline, conservative method described in Section 1.4). For every father in this sample, we retrieve all regions in which they worked during the 1992-2019 period using (formal) employment data (RAIS). We focus on the latter to track moves since address updates in the Brazilian person registry are largely incomplete before 2000. Although this choice implies that our sample is more representative of formal workers, doing so allows us to track migration more precisely, which is crucial for this analysis.

We define a mover as someone leaving a job in region a and taking a job in a new region b for at least two years. To increase precision, we focus on fathers showing up in employment data for at least five years in the 1993-2019 period.⁷ Fathers who never move are defined as permanent residents of their regions. In turn, movers are those who move at least once. To simplify the analysis, we focus on families moving only once. Our final sample comprises 3,172,145 children and 2,260,645 fathers, with around 18% of them being movers.

⁶Specifically, for each group of variables, we keep the number of principal components needed to explain 90% of the variation in them. Subsequently, we compute the index as an average weighted by the amount of variation each component absorbs.

⁷Our results remain similar when varying this threshold to ten or fifteen years. We use the five-year cut-off to enlarge the final sample and enhance precision.

A.5.2 Defining the predicted outcomes of permanent residents

We closely follow the research design and specifications in (CHETTY; HENDREN, 2018a) and (DEUTSCHER, 2020). First, we characterize outcomes of permanent residents of each region m and cohort c by running several rank-rank regressions of the type, for each region and cohort:

$$y_{imc} = \alpha_{mc} + \beta_{mc}p_{imc} + \varepsilon_{imc} \quad (\text{A.2})$$

where y_{imc} denotes the income percentile rank at the age of 24 of a child from cohort c who spent her entire childhood in region m . We focus on income at the age of 24 that we can measure for all cohorts (1983-1992) in our sample. To ensure precision, we keep only region-cohort pairs for which we have at least 400 observations. We then calculate the predicted income rank of residents for every parental income rank p , region m , and cohort c : $\hat{y}_{pmc} = \hat{\alpha}_{mc} + \hat{\beta}_{mc} \times p$.

A.5.3 Parametric specification and family fixed effects

Our baseline specification (1.3) includes nearly 180 thousands of fixed effects α_{ocpa} . While they ensure that we exclusively compare very similar children (with the same origin, cohort, parental income decile and age at move) to estimate place effects, they also strongly restrict the variation used in the analysis. Consequently, they leave little space for adding family fixed effects, which also strongly restricts the variation in the analysis. Hence, we follow (CHETTY; HENDREN, 2018b) and rely on a less restrictive, parametric specification to assess the robustness of our findings to family fixed effects:

$$y_i = \sum_{a=1}^{33} b_a I_a(a_i = a) \Delta_{odpc} + \sum_{c=1983}^{1991} \kappa_c I_c(c_i = c) \Delta_{odpc} + \sum_{a=1983}^{1992} I_c(c_i = c) (\eta_c^1 + \eta_c^2 \hat{y}_{poc}) + \sum_{a=1}^{33} I_a(a_i = a) (\zeta_a^1 + \zeta_a^2 p_i) + \lambda_f + \varepsilon_i \quad (\text{A.3})$$

Rather than controlling for fixed effects α_{ocpa} , this specification linearly controls for the quality of origin – which is allowed to vary by parental income and cohort – and age at move by parental income, accounting for the disruption effects of moving at different ages. Specifically, the first term in the second line is defined by cohort fixed effects η_c^1 and an interaction between the cohort dummies η_c^2 and the quality of origin $\hat{y}_{poc}$, modeled as the predicted income of permanent residents at origin o . In turn, the second term is defined by age at move fixed effects ζ_a^1 and age at move dummies ζ_a^2 interacted with parental percentile rank, p_i . Finally, the specification controls for family fixed effects λ_f , ensuring that causal place effects solely rely on variation across siblings.

A.5.4 Overidentification tests

The next results confirm that children's outcomes converge to those of permanent residents with precisely the same age, gender, and race, while coefficients on other groups are generally an order of magnitude smaller, close to zero, and statistically insignificant. In addition, since children's

outcomes in different areas not only differ at the mean but over the entire distribution, we show that movers' outcomes track different moments of the distribution of permanent residents' outcomes beyond the mean. For instance, two areas may have the same mean child rank for low-income children but different probabilities that children end up in the top decile of the income distribution. These tests address additional concerns such as the possibility that moves to better places are driven by different shocks producing positive effects on children that decrease with age, e.g., positive income or wealth shocks. Specifically, they indicate that following these shocks, parents would need highly accurate knowledge to select better places for our results to be driven by selection. Accordingly, for them to drive our main findings, parents would need to select places that offer better opportunities for children from the same cohort, gender and race. Finally, the potential shocks driving such selection process would need to replicate not only the mean outcomes but also the distribution of outcomes for children in the destination.


We start by showing that place effects are cohort-specific. The blue line in Figure A.13 displays the estimates for the exposure rate obtained from eleven separate regressions. In each of them, we replace the main independent variable – the predicted difference in outcomes for children of the *same cohort* – by the predicted difference for *other* cohorts, born from five years before to five years after. The coefficients obtained using adjacent cohorts are quite similar to the baseline, as regions with better opportunities for a given cohort usually are also good for other cohorts. In turn, the red line in Figure A.13 plots estimates when all cohort-specific predictions are simultaneously included in a single regression. Conditional on the predicted outcomes of their *own* cohort, all other cohorts' predictions are statistically insignificant, while the true cohort coefficient approaches the baseline estimate. Hence, children's outcomes converge to the outcomes of permanent residents of their *own* cohort and other cohorts' outcomes have little explanatory power. Thus, any omitted variable possibly driving our results would have to precisely emulate cohort-specific place effects.

In Table A.13, we conduct an analogous exercise for gender and race. For this purpose, we construct gender-specific predictions and estimate exposure rates in three different ways: using the predicted outcomes of the child's own gender (column 1), the opposite gender's prediction (column 2), and both together (column 3). Like the blue line in the previous exercise, both regressors yield statistically significant estimates, since there is a considerable correlation between outcomes for boys and girls within regions. Nonetheless, the one based on the child's own gender has higher explanatory power. Column 3 replicates the red line in Figure A.13: controlling for the own gender's prediction, the coefficient on the opposite gender (placebo) is negligible. Columns 4-6 conduct the same exercise for race, showing similar results and supporting our main analysis.

All estimates up to now have been based on the predicted differences in *mean* outcomes across locations. Now we show that place effects also replicate permanent residents' outcomes along the income distribution. We construct permanent residents' predictions for the probability of being in the top and bottom deciles of the national income distribution in adulthood. Subsequently, we estimate exposure rates contrasting the distributional predictions with the mean prediction (placebo). Columns 1-2 in Table A.14 show that the top ten probability is better explained by the distributional prediction than by the mean prediction when running separate regressions. In column 3, a simultaneous regression yields a significant coefficient for

the distributional prediction, while the coefficient on the placebo is zero. Columns 4-6 replicate 1-3 but for the bottom ten probability, with equivalent results. Thus, the distribution of children's incomes converges to the distribution of incomes in the destination in proportion to exposure time.

Figure A.13: Placebo tests: Cohort-specific convergence



Note: This figure presents estimates of the annual childhood exposure effect on children's income ranks in adulthood using permanent resident predictions for the child's own birth cohort and surrounding "placebo" birth cohorts. The series in blue plots estimates of the exposure effect γ_t from nine separate regressions, using permanent resident predictions from cohort $c+t$ (where t ranges between -5 and 5) as the key independent variables and the outcomes of children in birth cohort c as the dependent variable. The series in red plots estimates from a single multivariate regression that simultaneously includes all nine permanent resident predictions $t = -5, \dots, 5$.

Note:

Table A.11: Summary of mobility estimates for the 50 largest metropolitan areas

Region	2021 pop. (thousands)	Slope	$E[y p = 25]$	$E[y p = 75]$	Q1Q1	Q1Q5
São Paulo, SP	22,049	0.66	34.3	67.4	44.4	1.7
Rio de Janeiro, RJ	12,901	0.55	35.6	63.0	30.1	1.6
Belo Horizonte, MG	5,348	0.50	39.3	64.4	23.2	2.2
Fortaleza, CE	4,179	0.49	33.1	57.8	45.8	2.0
Recife, PE	4,108	0.54	31.8	58.6	48.5	1.7
Salvador, BA	4,065	0.54	29.1	56.3	52.6	1.5
Curitiba, PR	3,732	0.53	41.7	68.3	17.9	2.4
Porto Alegre, RS	3,267	0.50	39.5	64.5	21.8	2.4
Campinas, SP	3,201	0.53	40.4	66.9	33.1	2.3
Distrito Federal, DF	3,094	0.54	36.7	63.6	41.5	2.4
Belém, PA	2,773	0.58	25.1	54.3	63.0	0.8
Goiânia, GO	2,628	0.49	42.4	66.8	24.4	3.8
Manaus, AM	2,605	0.57	24.5	53.0	58.1	1.4
Vitória, ES	2,100	0.48	40.1	64.2	25.4	3.1
Santos, SP	1,927	0.58	34.1	62.9	44.3	2.9
Sorocaba, SP	1,840	0.53	39.3	65.6	32.3	2.8
Natal, RN	1,734	0.49	33.9	58.6	44.9	1.7
São Luís, MA	1,657	0.46	32.9	55.8	49.6	2.1
Ribeirão Preto, SP	1,534	0.60	36.3	66.3	42.8	2.7
João Pessoa, PB	1,430	0.48	33.9	57.9	46.8	2.6
Maceió, AL	1,316	0.52	30.9	56.8	52.7	0.7
Feira de Santana, BA	1,242	0.44	31.1	53.0	57.9	2.4
Aracaju, SE	1,233	0.51	31.9	57.3	49.1	2.1
Florianópolis, SC	1,181	0.42	46.5	67.5	11.2	4.0
Campo Grande, MS	1,131	0.49	41.1	65.8	31.5	3.4
São José dos Campos, SP	1,125	0.54	36.3	63.2	37.3	2.2
Teresina, PI	1,116	0.48	37.1	61.1	43.8	2.8
Londrina, PR	1,114	0.44	43.9	65.7	25.7	2.8
Cuiabá, MT	1,105	0.48	40.3	64.2	24.5	2.3
Joinville, SC	1,044	0.46	46.6	69.7	15.7	2.0
Jundiaí, SP	973	0.53	41.2	67.7	37.5	9.1
Uberlândia, MG	959	0.43	43.3	64.7	30.5	4.4
São José do Rio Preto, SP	934	0.53	41.5	67.8	30.9	4.0
Novo Hamburgo - São Leopoldo, RS	908	0.48	42.0	65.9	17.8	2.3
Pelotas, RS	845	0.39	42.1	61.8	28.1	3.3
Caxias do Sul, RS	841	0.42	46.6	67.7	23.3	3.4
Maringá, PR	801	0.43	45.3	66.8	25.7	1.4
Montes Claros, MG	770	0.44	40.1	62.0	33.4	2.7
Juiz de Fora, MG	753	0.44	38.7	60.6	31.7	2.3
Macapá, AP	675	0.50	23.9	49.0	59.1	0.5
Bauru, SP	668	0.54	38.4	65.4	34.8	0.6
Volta Redonda - Barra Mansa, RJ	668	0.51	36.5	61.8	36.4	1.6
Porto Velho, RO	667	0.49	30.0	54.6	42.9	2.4
Campos dos Goytacazes, RJ	661	0.48	36.5	60.5	32.0	2.9
Ipatinga, MG	651	0.41	39.1	59.5	34.0	3.5
Ponta Grossa, PR	648	0.51	41.6	67.1	31.1	2.2
Taubaté - Pindamonhangaba, SP	637	0.48	37.1	61.3	26.1	0.8
Araraquara, SP	631	0.55	38.6	66.0	28.7	2.3
Piracicaba, SP	617	0.56	39.1	66.9	31.4	3.9
Santa Maria, RS	485	0.41	45.2	65.7	28.0	4.5


Note: The table summarizes mobility estimates in the 50 largest metropolitan areas (IGRs) of Brazil, according to IBGE's population count in 2021. Mobility estimates are the rank-rank slope (relative), the expected income rank of below- and above-median income children (absolute), the bottom-bottom persistence probability (Q1Q1), and the bottom-to-top quintile transition probability (Q1Q5). Mobility measures are based on our baseline sample of the 1988-1990 cohorts. Children are assigned to IGRs based on the location of their parents in 2000.

Table A.12: List of municipal socioeconomic indicators and data sources

Group	Indicator	Year	Source
Demographics	% evangelical, % catholic, % w/o religion, % black, % mixed-race, % white, % HH in rural area	2000	IBGE
Economic Structure	% GDP agriculture, % GDP pub. sector, % GDP industry, (value-added), # companies pc, % income from wages, % family agriculture	2000	IBGE
Education	Test scores (portuguese and math), % teachers w/ college degree, school quality, drop-out rate, grade distortion, literacy rate (5-9yo, 10-14yo, and 40-49yo), % of adults w/ high school	2000/2005	IBGE/Inep
Family Structure	% families > 5 people, % HH single mother, % HH married couple, % HH dependency rate > .75, % of divorcees	2000	IBGE
Health	Dengue cases pc, teen pregnancy (10-14yo and 15-17yo), child mortality, life expectancy, % of vaccinated, SUS coverage	2000	DataSUS/IBGE
Household	% HH in favelas, % HH single family, % HH w/ land property, % HH w/ trash service, % HH w/ paved street, % HH w/ bathroom, % HH w/ piped water, % HH w/ electricity, % HH w/ street light, % HH w/ adequate housing, % HH w/ sewage, % HH w/ high people/room density	2000	IBGE
Income	Average family income pc, % families above the poverty line, % families earning more than 1 MW	2000	IBGE
Inequality	Gini Index, top 10/bottom 40 income ratio, top 10 income share, land inequality	2000	IBGE
Infrastructure	Broadband access, distance from state capital, # of bank agencies	2000/2007	ANATEL/IBGE
Labor Market	% of workers < 1 MW, % of workers w/o earnings, % of workers > 10 MW, female/male participation rate, female/male employment rate, % public servants, % firm owners, % self-employed, % formal workers, % employees, informality rate, # of labor unions	2000	IBGE
Municipal Budget	Government spending (health, welfare, education, housing), % of federal government transfers, IPTU revenue (property-tax)	2000	FINBRA
Public Safety	Homicide rate (total, males, young), % of deaths by homicide, % of deaths by car accident	2007	Ipeadata
Social Capital	Religious associations pc, cultural organizations pc, political organizations pc, civil associations pc, % of deaths by suicide	2000	IBGE


Note: The table list all indicators used in the mobility correlates analysis, along with their source, year, and category group (used in the principal components analysis). All of them are obtained at the municipal level and then aggregated to immediate region level by population-weighted averages.

Figure A.10: Mobility Correlates




Note: The figure summarizes a series of cross-regional univariate regressions of absolute mobility on a series of demographic, political, and socioeconomic indicators. The horizontal axis marks coefficients and .95 confidence intervals for each indicator, which are labelled in the vertical axis. Both dependent and independent variables are normalized so coefficients can be interpreted as straightforward correlations. Indicators are colored according to broad categories.

Figure A.11: Mobility Correlates: Principal Components



Note: The figure reports the results of two multivariate regressions between absolute mobility and principal components of regional characteristics. The horizontal axis marks coefficients and .95 confidence intervals for each indicator, which are labelled in the vertical axis. The blue dots includes only the principal components and the red triangles include region fixed-effects (North, Northeast, Center-West, Southeast, and South).

Figure A.12: Mobility Correlates: LASSO Regularization



Note: The figure plots the results of LASSO regularization for the correlation between absolute mobility and indexes constructed from socioeconomic indicators. The horizontal axis plots the regularization parameter lambda, while coefficients of each index are represented in the vertical axis. As lambda grows, the penalization for the number of regressors grows and coefficients are shrunk towards zero.

Table A.13: Placebo test: Gender- and race-specific convergence

	Exposure effect γ					
	Gender			Race		
	(1)	(2)	(3)	(4)	(5)	(6)
Own Group	0.025*** (0.001)		0.025*** (0.002)	0.022*** (0.001)		0.022*** (0.002)
Opposite Group		0.020*** (0.001)	-0.000 (0.001)		0.017*** (0.001)	0.000 (0.002)
Observations	285,912	285,912	285,912	267,614	267,614	267,614

Note: The table reports estimates of annual childhood exposure effects γ using gender- (Columns 1-3) and race-specific (Columns 4-6) permanent resident predictions. In all columns, the dependent variable is the child's family income rank at the age of 24. In both panels, column 1 (4) replaces the predicted outcomes based on all permanent residents in the origin and destination with predictions based on the outcomes of children who have the same gender (race) as the child who moves. Column 2 (5) replicates column 1, replacing the own-gender (race) predicted outcomes with the predicted outcomes of the opposite gender (race). Column 3 (6) combines the variables in columns 1 and 2, including both the own-gender (race) and placebo other-gender (race) predictions (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

Table A.14: Placebo test: Distributional convergence

	Upper Tail			Lower Tail		
	(1)	(2)	(3)	(4)	(5)	(6)
Distributional prediction	0.031*** (0.002)		0.034*** (0.002)	0.021*** (0.003)		0.022*** (0.003)
Mean Rank prediction		0.000*** (0.000)	0.000 (0.000)		0.000*** (0.000)	0.000** (0.000)
Observations	285,912	285,912	285,912	285,912	285,912	285,912

Note: This table reports estimates of annual childhood exposure effects γ for upper- and lower-tail outcomes: being in the top or bottom 10% of the cohort-specific income distribution at the age of 24. Column 1 reports estimates from a regression of an indicator for being in the top 10% on the difference between permanent residents' predicted probabilities of being in the upper tail in the destination vs. the origin. Column 2 replicates column 1 but uses the difference between permanent residents' predicted *mean* ranks on the right-hand side of the regression. Column 3 includes both the (distributional) and the mean rank prediction. Columns 4-6 replicate columns 1-3 using an indicator for being at the bottom 10% at the age of 24 as the outcome. In all columns, the sample comprises all children in the primary analysis sample of one-time movers (* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$).

APPENDIX B

Do CCTs Create Conditions to Thrive? *Bolsa Família* and Social Mobility in Brazil

B.1 Appendix to Section 3

B.1.1 Definition of Outcomes

Table B.1: Description of Outcome Variables

Outcome	Sources and Description
<i>A. Human Capital</i>	
Years of Education	Calculated from educational level reported in administrative employment registries (RAIS) or welfare registries (CadÚnico), in this order of priority. We go from education level to years of education using the standard grade progression in Brazil.
School completion	Dummy equal to 1 if years of education is greater or equal than 12.
College degree	Dummy equal to 1 if years of education is greater or equal than 16.
<i>B. Labor Market</i>	
Formal Job	Dummy equal to 1 if held at least one formal job (RAIS) btw 23–25 years old.
Months Worked	Average months worked per year in a formal job (RAIS) btw 23–25 years old.
Formal Income (BRL/year)	Average annual earnings in formal jobs (RAIS) btw 23–25 years old.
Entrepreneurship	Dummy equal to 1 if has firm registered on the Firm Registry (CNPJ)
<i>C. Intergenerational Mobility</i>	
Adult PBF enrollment	Dummy equal to 1 if head (or her partner) in HH receiving PBF in 2015–19.
Income pct. rank	Income percentile in the gender and cohort-specific income distribution as computed in (BRITTO et al., 2022)
Income rank > 20	Dummy equal to 1 if income percentile greater than 20
Income rank > 80	Dummy equal to 1 if income percentile greater than 80
<i>D. Migration</i>	
City Migration	Dummy equal to 1 if city of residence in 2019 is different from city of birth (defined as mother's city in 2003). Residence is tracked by updates in the Person Registry (CPF).
State Migration	Dummy equal to 1 if state of residence is different from state of birth
Region Migration	Dummy equal to 1 if region of residence is different from region of birth
Live in Metro Area	Dummy equal to 1 if city of residence has population of 100k or more
<i>E. Health and Behavior</i>	
Survive up to 2019	Dummy equal to 1 alive in 2019, tracked by updated in CPF
Teenage parenthood	Dummy equal to 1 if registers a child in CPF before 18 years old

B.2 Appendix to Section 4

B.2.1 First-stage relationships and outcomes counterfactuals

Table B.2: First-stage relationships and outcomes counterfactuals

	Males		Females		Whites		Non-Whites	
	10–15 (1)	16–21 (2)	10–15 (3)	16–21 (4)	10–15 (5)	16–21 (6)	10–15 (7)	16–21 (8)
% treated	49.5	47.6	49.3	47.1	46.1	44.0	50.3	47.8
# years treated	3.81	1.27	3.81	1.26	3.25	1.14	4.07	1.31
<i>A. Human Capital</i>								
Years of Education	9.6	10.3	9.7	10.4	11.1	11.2	9.1	10.4
School completion	0.60	0.63	0.58	0.61	0.71	0.70	0.54	0.61
College degree	0.06	0.07	0.10	0.13	0.10	0.12	0.07	0.10
<i>B. Labor Market</i>								
Formal Job	0.58	0.61	0.35	0.36	0.62	0.61	0.43	0.46
Months Worked	4.5	4.9	2.7	2.7	5.0	5.0	3.2	3.5
Formal Income (BRL/year)	7,503	7,652	3,564	3,308	8,022	7,564	5,162	5,028
Entrepreneurship	0.03	0.04	0.02	0.03	0.03	0.05	0.03	0.03
<i>C. Intergenerational Mobility</i>								
Adult PBF enrollment	0.03	0.04	0.31	0.48	0.14	0.20	0.18	0.31
Income pct. rank	38.3	38.1	35.3	34.7	49.8	46.6	31.1	32.5
Income rank > 20	0.68	0.70	0.61	0.62	0.85	0.80	0.53	0.61
Income rank > 80	0.09	0.09	0.06	0.06	0.13	0.11	0.06	0.05
<i>D. Migration</i>								
City Migration	0.43	0.46	0.42	0.44	0.46	0.49	0.41	0.42
State Migration	0.20	0.21	0.18	0.18	0.21	0.23	0.18	0.17
Region Migration								
Live in Metro Area	0.43	0.45	0.42	0.45	0.50	0.50	0.40	0.44
<i>E. Health and Behavior</i>								
Survive up to 2019	0.98	0.97	0.99	0.99	0.99	0.99	0.99	0.98
Teenage parenthood	0.01	0.03	0.11	0.17	0.06	0.09	0.06	0.11

Notes. The table reports first-stage relationships and outcomes counterfactuals by gender and race for fully- and partially-treated cohorts. First-stage relationships are the probability of being in a household treated by PBF in the first 3 years of the program (2004–06) and the actual number of years between 0–21 years old spent under the program. For each cohort, they are calculated as the difference between raw averages in the treatment and control groups in our research design. Counterfactual outcomes are computed as average outcomes for the non-exposed cohorts (older than 21 at the start of the program) in the treatment group (baseline) adjusted by the relative change across cohorts in the control group (trend).

B.2.2 Racial and Regional Heterogeneity

Table B.3: PBF effects in the Long Run: Age-specific ITT estimates by race

	Whites		Non-Whites	
	10–15 (1)	16–21 (2)	10–15 (3)	16–21 (4)
<i>A. Human Capital</i>				
Years of Education	0.359*** (0.014)	0.197*** (0.010)	0.943*** (0.024)	0.167*** (0.016)
School completion	0.031*** (0.002)	0.027*** (0.002)	0.062*** (0.002)	0.018*** (0.002)
College degree	0.037*** (0.002)	-0.001 (0.001)	0.050*** (0.002)	-0.001 (0.001)
<i>B. Labor Market</i>				
Formal Income (BRL/year)	807.1*** (70.1)	-25.0 (44.8)	925.1*** (55.4)	-186.1*** (32.8)
Entrepreneurship	0.031*** (0.001)	0.009*** (0.001)	0.019*** (0.001)	0.003*** (0.001)
<i>C. Intergenerational Mobility</i>				
Adult PBF enrollment	-0.018*** (0.002)	-0.005*** (0.001)	-0.018*** (0.002)	-0.002 (0.001)
Income pct. rank	1.47*** (0.135)	-0.261*** (0.094)	2.91*** (0.130)	-0.703*** (0.087)
<i>D. Migration</i>				
City Migration	-0.035*** (0.002)	-0.026*** (0.002)	-0.033*** (0.002)	-0.018*** (0.002)
Live in Metro Area	-0.009*** (0.002)	-0.009*** (0.001)	0.032*** (0.002)	-0.003** (0.001)
<i>E. Health and Behavior</i>				
Teenage parenthood	-0.010*** (0.002)	0.001 (0.001)	-0.006*** (0.002)	0.004*** (0.001)
Observations	11,806,187	11,806,187	10,961,412	10,961,412

Notes. The table reports ITT coefficients on interactions of cohort indicators and our treatment variable. They are obtained from a modified version of Equation 2.1 in which age indicators I_t are replaced by indicators for being 10–15, 16–21, or older (omitted category) when PBF began. We run separate regressions for whites and non-whites. A detailed description of outcomes used as dependent variables is given in Appendix Table B.1. (*p<0.1; **p<0.05; ***p<0.01).

Table B.4: PBF effects in the Long Run: Age-specific ITT estimates by region

	Center-South		North and Northeast	
	10–15 (1)	16–21 (2)	10–15 (3)	16–21 (4)
<i>A. Human Capital</i>				
Years of Education	0.700*** (0.016)	0.116*** (0.012)	1.10*** (0.026)	0.211*** (0.017)
School completion	0.052*** (0.002)	0.021*** (0.001)	0.069*** (0.002)	0.020*** (0.002)
College degree	0.047*** (0.001)	-0.006*** (0.001)	0.062*** (0.002)	0.001 (0.001)
<i>B. Labor Market</i>				
Formal Income (BRL/year)	1,229.9*** (49.9)	-135.0 (32.2)	857.7*** (69.0)	-318.5*** (38.8)
Entrepreneurship	0.031*** (0.001)	0.010*** (0.001)	0.023*** (0.001)	0.005*** (0.001)
<i>C. Intergenerational Mobility</i>				
Adult PBF enrollment	-0.018*** (0.001)	-0.003*** (0.001)	-0.020*** (0.002)	-0.002 (0.002)
Income pct. rank	3.66*** (0.103)	-0.593*** (0.073)	1.20*** (0.156)	-0.422*** (0.102)
<i>D. Migration</i>				
City Migration	-0.038*** (0.002)	-0.025*** (0.001)	-0.037*** (0.003)	-0.021*** (0.002)
Live in Metro Area	0.014*** (0.002)	-0.009*** (0.001)	0.027*** (0.002)	-0.009** (0.002)
<i>E. Health and Behavior</i>				
Teenage parenthood	-0.010*** (0.002)	0.001 (0.001)	-0.005*** (0.002)	0.006*** (0.001)
Observations	16,466,417	16,466,417	6,301,182	6,301,182

Notes. The table reports ITT coefficients on interactions of cohort indicators and our treatment variable. They are obtained from a modified version of Equation 2.1 in which age indicators I_t are replaced by indicators for being 10–15, 16–21, or older (omitted category) when PBF began. We run separate regressions for each region. A detailed description of outcomes used as dependent variables is given in Appendix Table B.1. (*p<0.1; **p<0.05; ***p<0.01).

B.2.3 Robustness to Alternative Treatment Definitions

Table B.5: PBF effects in the Long Run: Robustness to Alternative Treatment Definitions

	Top 25% (baseline)		Top 20%		Top 30%	
	10–15 (1)	16–21 (2)	10–15 (3)	16–21 (4)	10–15 (5)	16–21 (6)
<i>First stage</i>						
% treated	49.4	47.4	50.1	50.1	47.9	44.1
# years treated	3.81	1.27	4.00	1.37	3.65	1.16
<i>A. Human Capital</i>						
Years of Education	0.839*** (0.014)	0.165*** (0.010)	0.808*** (0.014)	0.166*** (0.010)	0.860*** (0.013)	0.160*** (0.009)
School completion	0.058*** (0.001)	0.022*** (0.001)	0.054*** (0.002)	0.020*** (0.001)	0.059*** (0.001)	0.022*** (0.001)
College degree	0.052*** (0.001)	-0.003*** (0.001)	0.049*** (0.001)	-0.003*** (0.001)	0.054*** (0.001)	-0.003*** (0.001)
<i>B. Labor Market</i>						
Formal Income (BRL/year)	1,114.5*** (40.2)	-208.0*** (24.9)	973.0*** (40.0)	-231.56*** (25.0)	1,191.7*** (41.1)	-201.9*** (25.1)
Entrepreneurship	0.029*** (0.001)	0.008*** (0.001)	0.027*** (0.001)	0.007*** (0.001)	0.030*** (0.001)	0.009*** (0.001)
<i>C. Intergenerational Mobility</i>						
Adult PBF enrollment	-0.019*** (0.001)	-0.003*** (0.001)	-0.020*** (0.001)	-0.003*** (0.001)	-0.019*** (0.001)	-0.003*** (0.001)
Income pct. rank	2.94*** (0.086)	-0.504*** (0.059)	2.49*** (0.090)	-0.507*** (0.063)	3.24*** (0.083)	-0.521*** (0.057)
<i>D. Migration</i>						
City Migration	-0.037*** (0.002)	-0.024*** (0.001)	-0.037*** (0.002)	-0.024*** (0.001)	-0.037*** (0.002)	-0.024*** (0.001)
Live in Metro Area	0.019*** (0.001)	-0.008*** (0.001)	0.015*** (0.001)	-0.009*** (0.001)	0.020*** (0.001)	-0.008*** (0.001)
<i>E. Health and Behavior</i>						
Teenage parenthood	-0.008*** (0.001)	0.003*** (0.001)	-0.009*** (0.001)	0.002*** (0.001)	-0.009*** (0.001)	0.002*** (0.003)
Observations	22,767,599	22,767,599	22,767,599	22,767,599	22,767,599	22,767,599


Notes. The table reports ITT coefficients on interactions of cohort indicators and our treatment variable. They are obtained from a modified version of Equation 2.1 in which age indicators I_t are replaced by indicators for being 10–15, 16–21, or older (omitted category) when PBF began. Columns 1 and 2 report coefficients using our baseline treatment definition (mothers in the top 25% of the predicted-probabilities distribution), while columns 3–6 report coefficients for alternative definitions for treatment. A detailed description of outcomes used as dependent variables is given in Appendix Table B.1. (*p<0.1; **p<0.05; ***p<0.01).

APPENDIX C

Decentralization, Tax Administration, and Taxation: Evidence from Brazil's Rural Land Tax


C.1 Additional Results

Figure C.1: The Effects of Decentralization Agreements – No Controls




Notes. The figure plot estimates of the effects of the decentralization agreements on different outcomes obtained using equation (3.2) without controls. The x-axis shows the relative time since decentralization while the y-axis plot the coefficients and their respective 95% confidence intervals.

Figure C.2: The Effects of Decentralization Agreements – Weights




Notes. The figure plot estimates of the effects of the decentralization agreements on different outcomes obtained using equation (3.2) weighting observations by the municipal area. The x-axis shows the relative time since decentralization while the y-axis plot the coefficients and their respective 95% confidence intervals.

Figure C.3: The Effects of Decentralization Agreements – (CALLAWAY; SANT'ANNA, 2021)



Notes. The figure plot estimates of the effects of the decentralization agreements on different outcomes obtained using the (CALLAWAY; SANT'ANNA, 2021) estimator.

Figure C.4: IHS(ITR Due) by firm size



Notes. The figure plot estimates of the effects of the decentralization agreements on different outcomes obtained using equation (3.2). The x-axis shows the relative time since decentralization while the y-axis plot the coefficients and their respective 95% confidence intervals. Each panel reports results for farms in different size intervals.

Table C.1: Average effect of ITR by region

Sub-samples:	North	Northeast	Southeast	Midwest	South
Post $\times$ Non-denounced	0.1300** (0.0524)	0.0724 (0.0996)	0.1549*** (0.0261)	0.2473*** (0.0241)	0.0735*** (0.0186)
<i>Controls</i>					
Municipality	Yes	Yes	Yes	Yes	Yes
Year-Cohort	Yes	Yes	Yes	Yes	Yes
Baseline Characteristics	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
Observations	14,975	14,677	17,493	17,774	18,469
# Municipalities	1,207	1,178	1,410	1,433	1,483
Mean Dep. Var	10.63	10.64	10.81	10.91	10.79

Notes. This table presents the differences-in-differences for the $\text{Log}(ITR_{due})$ for municipalities in each region. We use the same control group for each regression. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1. Errors are clustered on the municipality level. Baseline characteristics are first-year IPTU, ISS and ITR, also maize and soy productivity interacted with year dummies.

