ESSAYS ON PUBLIC POLICY AND POVERTY

AN ABSTRACT

SUBMITTED ON THE EIGHTH DAY OF APRIL 2016

TO THE DEPARTMENT OF ECONOMICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

OF THE SCHOOL OF LIBERAL ARTS

OF TULANE UNIVERSITY

FOR THE DEGREE

OF

DOCTOR OF PHILOSOPHY

BY

Sean Higgins

APPROVED:
Nora Lustig, Ph.D.
Director

Alan Barreca, Ph.D.

John H. Y. Edwards, Ph.D.

Abstract

Public policy has important implications for the lives of the poor. This dissertation analyzes how three types of policy impact the poor in developing countries. First, tax and transfer systems can benefit many poor while still making some poor worse off, and this phenomenon is overlooked by measures currently used to assess transfers in tandem with the taxes used to pay for them. I show that comparisons of poverty before and after taxes and transfers, as well as measures of horizontal equity and progressivity—which are often used to analyze anti-poverty policies in tandem with the taxes used to pay for them—can fail to capture an important aspect: that a substantial proportion of the poor are made poorer (or non-poor made poor) by the tax and transfer system. I call this fiscal impoverishment, and axiomatically derive a measure of its extent. Second, the government's choice of how to measure poverty—specifically, the choice between a unidimensional (usually income or consumption-based) measure and a multidimensional measure that incorporates other dimensions such as health and education—can affect the strategic interactions between government ministers, leading to changes in the amount of resources spent to alleviate poverty. In a game-theoretic framework, I show that despite introducing free riding, a multidimensional measure usually leads to an increase in total antipoverty spending; antipoverty expenditures can be further increased by publishing partial dimensional indices alongside the scalar multidimensional one. Third, efforts to digitize government transfer programs through savings accounts and debit cards can enable the poor to build trust in financial institutions and save more. I study a natural experiment in which debit cards were rolled out to beneficiaries of a Mexican conditional cash transfer program, who were al-

ready receiving their transfers in savings accounts through a government bank. Using a rich combination of administrative and survey data, I find beneficiaries initially used their cards to check their balances and build trust in the bank, after which they used the account to save. Formal and overall savings increased, and this effect was higher for women with low baseline bargaining power who may have the most difficulty saving at home.

ESSAYS ON PUBLIC POLICY AND POVERTY

A DISSERTATION

SUBMITTED ON THE EIGHTH DAY OF APRIL 2016

TO THE DEPARTMENT OF ECONOMICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

OF THE SCHOOL OF LIBERAL ARTS

OF TULANE UNIVERSITY

FOR THE DEGREE

OF

DOCTOR OF PHILOSOPHY

BY

____

Sean Higgins

APPROVED:________________________

Nora Lustig, Ph.D.
Director

____

Alan Barreca, Ph.D.

____

John H. Y. Edwards, Ph.D.

Acknowledgement

I am greatly indebted and very grateful to my dissertation committee, Nora Lustig (chair), Alan Barreca, and John Edwards, for their guidance and detailed comments and suggestions, which helped to greatly improve this dissertation. I am also very grateful to my faculty sponsors during my time as a visiting researcher at the Center for Economic Studies at El Colegio de México and as a visiting scholar at Haas School of Business at University of California, Berkeley: Antonio Yúnez-Naude and Paul Gertler. I gratefully acknowledge research funding from the Bill & Melinda Gates Foundation, Fulbright–García Robles Public Policy Initiative, National Science Foundation, and Tulane School of Liberal Arts; without this funding, the research in this dissertation would not have been possible.

The first chapter is co-authored with Nora Lustig, and is forthcoming in the *Journal of Development Economics*. For detailed comments on earlier versions of the paper, we are grateful to Francesco Andreoli, Alan Barreca, Jean-Yves Duclos, John Edwards, Charles Kenny, Peter Lambert, Darryl McLeod, Mauricio Reis, Kathleen Short, Rafael Salas, Jay Shimshack, Harry Tsang, Paolo Verme, two anonymous referees, and Maitreesh Ghatak, the editor. We are also grateful to Karim Araar, Abhijit Banerjee, François Bourguignon, Satya Chakravarty, Nachiketa Chattopadhyay, Chico Ferreira, Gary Fields, James Foster, Ravi Kanbur, Doug Nelson, Jukka Pirttilä, John Roemer, Jon Rothbaum, and Shlomo Yitzhaki for useful feedback. For applying our measures to microdata in a number of developing countries and sharing the results with us for this paper, we are grateful to the many country study authors we cite in the paper. We thank Claudiney Pereira for collaboration on the analysis of Brazilian data and Ruoxi Li, Sandra Martínez-Aguilar, Adam Ratzlaff, Mel Reitcheck, and William Smith for research assistance. Work on this project was partially completed when S. Higgins was visiting Haas School of Business at UC Berkeley and the Center for Economic Studies at El Colegio de México with funding from the Fulbright–García Robles Public Policy Initiative. This paper is part of a larger project that received funding from the Bill & Melinda Gates Foundation (grant numbers OPP1097490 and OPP1335502). These institutions and funders are gratefully acknowledged.

The second chapter is co-authored with Stefano Barbieri. Work on this project was partially completed when S. Barbieri was visiting IIEP at the Elliott School of George Washington University and S. Higgins was visiting Haas School of Business at UC Berkeley and CEE at El Colegio de México with funding from the Fulbright–García Robles Public Policy Initiative. The hospitality of these institutions is gratefully acknowledged. The authors thank Sabina Alkire, Rodrigo Aranda, Enrique de la Rosa, Gerardo Escaroz, Francisco Ferreira, Esfandiar Maasoumi, Antonio Martín del Campo, Iñaki Permanyer, Suman Seth, Stephen Younger, Claudio Zoli, and especially Nora Lustig for their generous comments and discussions, which helped to greatly improve the paper.

The third chapter is co-authored with Pierre Bachas, Paul Gertler, and Enrique Seira. This study would not have been possible without the generous time, access to data, and tireless answers to questions provided by officials in the offices of Mexico's conditional cash transfer program Prospera (formerly Oportunidades), and the government bank Bansefi. At Prospera, we are indebted to Martha Cuevas, Armando Gerónimo, Rodolfo Sánchez, José Solis, Carla Vázquez, and especially Rogelio Grados and Raúl Pérez. At Bansefi, we are indebted to Miguel Ángel Lara, Oscar Moreno, Ramón Sanchez, and especially Benjamín Chacón and Ana Lilia Urquieta. For comments that helped to greatly improve the paper, we thank Bibek Adhikari, Rodrigo Aranda, Stefano Barbieri, Alan Barreca, Carlos Chiapa, Laura Chioda, Alain de Janvry, John Edwards, Gerardo Escaroz, Fred Finan, Keith Finlay, Louise Fox, Jens Hainmueller, Anne Karing, Dean Karlan, John Loeser, Paty López, Jeremy Magruder, Aprajit Mahajan, Justin McCrary, Ted Miguel, John Morgan, Doug Nelson, Christine Parlour, Gérard Roland, Betty Sadoulet, Steven Sheffrin, Ben Sperisen, Stephen Younger, and numerous seminar participants. We are also grateful to Ignacio Camacho, Austin Farmer, Joel Ferguson, Isaac Meza, and Mel Reitcheck for research assistance. S. Higgins gratefully acknowledges funding from the Fulbright–García Robles Public Policy Initiative and National Science Foundation (Doctoral Dissertation Research Improvement Grant 1530800). E. Seira gratefully acknowledges funding from the Consortium on Financial Systems and Poverty and the Institute for Money, Technology & Financial Inclusion.

List of Tables

List of Figures

Contents

Chapter 1

Can a Poverty-Reducing and Progressive Tax and Transfer System Hurt the Poor?

Abstract

To analyze anti-poverty policies in tandem with the taxes used to pay for them, comparisons of poverty before and after taxes and transfers are often used. We show that these comparisons, as well as measures of horizontal equity and progressivity, can fail to capture an important aspect: that a substantial proportion of the poor are made poorer (or non-poor made poor) by the tax and transfer system. We illustrate with data from seventeen developing countries: in fifteen, the fiscal system is poverty-reducing and progressive, but in ten of these at least one-quarter of the poor pay more in taxes than they receive in transfers. We call this fiscal impoverishment, and axiomatically derive a measure of its extent. An analogous measure of fiscal gains of the poor is also derived, and we show that changes in the poverty gap can be decomposed into our axiomatic measures of fiscal impoverishment and gains.

1.1 Introduction

Anti-poverty policies are often evaluated in isolation from the taxes used to pay for them.[1] If, however, taxes cancel out the benefits of transfers for many poor households, so that some poor pay more in taxes than they receive in transfers, the objective of these policies might be compromised. This is especially important when poverty traps exist at the individual level (e.g., Ghatak, 2015; Ravallion, 2015): a tax and transfer system in which many poor pay more in taxes than they receive in transfers risks pushing the transiently poor into chronic poverty by shifting their after tax and transfer incomes below their individual-specific poverty trap thresholds.

Recently, the connection between anti-poverty policies and the taxes used to pay for them has come into the spotlight in the debates over the United Nations' Post-2015 Sustainable Development Goals. In recognition of the resources necessary to achieve these ambitious development goals, and partly as a consequence of austerity in advanced countries (and thus lower anticipated flows of international aid to developing countries), much of the discussion has focused on how developing countries should collect the revenue necessary to achieve the goals.[2] Influential organizations such as the International Monetary Fund and World Bank emphasize the importance of efficient taxes with minimal exemptions (International Monetary Fund, 2013; World Bank, 2013). When concerns are raised about these taxes—such as a no-exemption value added tax—falling disproportionately on the poor, many argue that higher tax burdens on the poor are acceptable if they are accompanied by sufficiently large targeted transfers: "spending instruments are available that are better targeted to the pursuit of equity concerns" (Keen and Lockwood, 2010, p. 141). Similarly, Engel et al. (1999, p. 186) assert that "it is quite obvious that the disadvantages of a proportional tax are moderated by adequate targeting" of transfers, since "what the poor

[1] We focus on anti-poverty policies that are redistributive in nature, one of the three categories of anti-poverty policies described in Ghatak (2015).

[2] See, for example, the focus on domestic resource mobilization in United Nations (2015).

individual pays in taxes is returned to her." These taxes "might conceivably be the best way to finance pro-poor expenditures, with the net effect being to relieve poverty" (Ebrill et al., 2001, p. 105).

How can we be sure that what the poor individual pays in taxes is returned to her? Even if the net effect of taxes and transfers is to relieve poverty, are some poor made worse off? When taxes and transfers are analyzed in tandem to determine how they affect the poor, it is common to compare poverty before taxes and transfers ("pre-fisc") to poverty after taxes and transfers ("post-fisc"). As we show in this paper, however, a fiscal system can be unambiguously poverty-reducing for a range of poverty lines and any poverty measure, yet still make a substantial proportion of the poor worse off. This phenomenon does not only occur with regressive taxes: we show that taxes and transfers can be globally progressive, unambiguously equalizing, and unambiguously poverty-reducing and *still* make many poor worse off. In other words, conventional tools used to measure how the poor are affected by the tax and transfer system are inadequate to measure whether some of the poor pay more in taxes than they receive in transfers, a phenomenon we call fiscal impoverishment (FI).

We also show that in practice, there are a number of countries with poverty-reducing and progressive tax and transfer systems that nevertheless make a substantial proportion of the poor poorer (or non-poor poor), illustrating with data from seventeen developing countries.[3] In fifteen of these countries, post-fisc poverty is unambiguously lower than pre-fisc poverty (measured with any poverty line up to $1.25 per person per day in low and lower-middle income countries and $2.50 per day in upper-middle income countries)[4]

[3] Our illustration uses results provided to us by the authors of country studies conducted as part of the Commitment to Equity (CEQ) Institute, located at Tulane University (`www.commitmenttoequity.org`). The countries included are Armenia (Younger and Khachatryan, forthcoming), Bolivia (Paz Arauco et al., 2014), Brazil (authors' calculations), Chile (Martínez-Aguilar and Ortiz-Juarez, 2015), the Dominican Republic (Aristy-Escuder et al., forthcoming), Ecuador (Llerena Pinto et al., 2015), El Salvador (Beneke et al., 2015), Ethiopia (Hill et al., forthcoming), Ghana (Younger et al., 2015), Guatemala (Cabrera et al., 2015), Indonesia (Afkar et al., forthcoming), Mexico (Aranda and Scott, 2015), Peru (Jaramillo et al., 2015), Russia (Lopez-Calva et al., forthcoming), South Africa (Inchauste et al., forthcoming), Sri Lanka (Arunatilake et al., forthcoming), and Tunisia (Shimeles et al., forthcoming). For an overview of the impact of taxes and social spending on inequality and poverty in many of these countries, see Lustig (2015).

[4] The $1.25 per person per day poverty line (in 2005 US dollars adjusted for purchasing power parity) is approximately equal to the median poverty line of the fifteen poorest countries for which poverty line data

and the tax and transfer system is globally progressive and unambiguously equalizing, i.e., we would conclude that the tax and transfer system unambiguously benefits the poor using conventional measures, potentially overlooking impoverishment. In all of these countries, some degree of FI occurs, and in ten of them we find that at least one-quarter of the poor pay more in taxes than they receive in transfers.

In light of the debate about financing anti-poverty policies and the Sustainable Development Goals, it is necessary to fill this gap in the measurement arsenal and develop a measure of this phenomenon that adheres to certain properties. We axiomatically derive a measure of FI, as well as an analogous measure for fiscal gains of the poor (FGP), which captures the extent to which some poor receive more in transfers than they pay in taxes.[5] We then show how a commonly used measure of poverty that overlooks the extent of FI, the poverty gap, can be decomposed into FI and FGP components using our axiomatic measures, again illustrating with data from seventeen developing countries. Because the extent of FI and FGP depend on the particular poverty line used, we also propose dominance criteria that can be used to determine whether one fiscal system (such as the one that would occur after a proposed reform) causes unambiguously less FI or more FGP than another (such as the current system) over a range of poverty lines. We analyze FI and FGP over a range of poverty lines in Brazil, which is a pertinent example due to the coexistence of high tax burdens on the poor (Baer and Galvão, 2008; Goñi et al., 2011) and lauded poverty-reducing cash transfer programs: a large-scale conditional cash transfer program that reaches over one-fourth of all Brazilian households and a non-contributory pension program for the elderly poor that reaches one-third of all elderly (Levy and Schady, 2013, Table 1).

Section 1.2 uses hypothetical and empirical examples to show that common tools to

are available, and the $2.50 line to the median of the world's low and middle income countries excluding the fifteen poorest (Chen and Ravallion, 2010).

[5] Our axioms are adapted from the axiomatic poverty and mobility measurement literatures (see Foster, 2006, and Zheng, 1997, for surveys of axiomatic poverty measurement and Fields, 2001, for a survey of axiomatic mobility measurement). Our resulting measure can be viewed as a censored directional version of the mobility measure derived by Fields and Ok (1996).

assess how the tax and transfer system affects the poor can fail to capture FI. Section 1.3 axiomatically derives a measure that does capture FI; it then proposes a partial FI ordering that can be used to compare the level of FI induced by two fiscal systems for any poverty line. Section 1.4 derives an analogous measure and partial ordering for FGP and shows that the poverty gap can be decomposed into our axiomatic measures of FI and FGP. Section 2.4 uses data from seventeen developing countries to illustrate the axiomatic measures and poverty gap decomposition. Section 2.5 concludes, and the formal axioms and proofs are collected in the Appendix.

1.2 The Problems with Conventional Measures

Through a number of examples, we illustrate and explain the problems with conventional measures of poverty, horizontal equity, and progressivity. Of course, these measures are still quite important for assessing a tax and transfer system; we merely aim to show that they do not capture everything we are interested in. First, in Section 1.2.1 we show the problem with poverty measures when they are used to compare poverty before and after taxes and transfers. Although comparisons of pre-fisc and post-fisc poverty are common in empirical studies (e.g., DeFina and Thanawala, 2004; Hoynes et al., 2006), poverty measures can overlook fiscal impoverishment because they obey the anonymity axiom (which is usually taken as an innocuous and desirable axiom): the tax and transfer system can reduce poverty while simultaneously making a substantial portion of the poor poorer, or making some non-poor poor. The anonymity axiom is not the only culprit for the shortcomings of existing measures, however: in Section 1.2.2 we show that measures designed to incorporate information about individuals' pre-fisc positions, such as measures of horizontal equity and progressivity, can also fail to capture FI.[6] To show that these shortcomings of

[6] Other measures that are sometimes used, such as the percent of income gained or lost by each pre-fisc income decile, overlook FI for a distinct reason: they average over individuals, so for example the poorest decile could gain income on average while a substantial number of poor within the first decile lose income. We do not include these measures in this paper since the reason they overlook FI is obvious.

conventional measures are not confined to contrived hypothetical examples, but rather occur frequently in practice, in Section 1.2.3 we present examples from seventeen developing countries: in ten, the tax and transfer system is poverty-reducing and progressive, but hurts a substantial portion of the poor by pushing them deeper into poverty.

1.2.1 Poverty Measures

Suppose the change in poverty caused by the fiscal system will be evaluated over a range of poverty lines, including lines greater than 6 and less than or equal to 10. Suppose there are three individuals in society with pre-fisc incomes of 5, 8, and 20, and (retaining the order of the individuals) post-fisc incomes 9, 6, and 18. For any poverty line in the range we are considering, and for any poverty measure in a broad class of measures, poverty has either not changed or decreased. This is because the poorest individual in the pre-fisc income distribution has an income of 5 and the second-poorest 8, while in the post-fisc distribution, the poorest has an income of 6 and the second-poorest 9. Poverty comparisons do not take into account that the poorest individual in the post-fisc distribution, with an income of 6, is not the poorest individual in the pre-fisc distribution who has an income of 5, but instead had an income of 8 in the pre-fisc distirbution and paid 2 more in taxes than she received in transfers. Depending on the exact poverty line chosen within the range we are considering, this individual was either pre-fisc poor and lost income to the fiscal system, or pre-fisc non-poor and pushed into poverty by the fiscal system.

It is clear, then, that poverty measures are inadequate to measure whether some of the poor pay more in taxes than they receive in transfers. Stochastic dominance tests, which are used to determine whether poverty is unambiguously lower in one income distribution than another for any poverty line and a broad class of poverty measures (Atkinson, 1987; Foster and Shorrocks, 1988), are also inadequate. This is because poverty measures and stochastic dominance tests are anonymous with respect to pre-fisc income: they compare the pre- and post-fisc income distributions without paying attention to the specific pre-fisc to post-fisc trajectory of particular individuals' incomes. The anonymity axiom, normally considered an innocuous and desirable property, becomes problematic when we are concerned with how the fiscal system affects the poor: in the words of Amiel and Cowell (1994, p. 448–9), "anonymity itself may be questionable as a welfare criterion when the social-welfare function is to take into account something more than the end-state distribution of incomes." Anonymity implies that poverty measures fail to take into account individuals' initial positions, and thus whether some are being made poorer by the tax and transfer system.[7]

Figure 1.1: Stylistic Illustration of Fiscal Impoverishment and Gains to the Poor

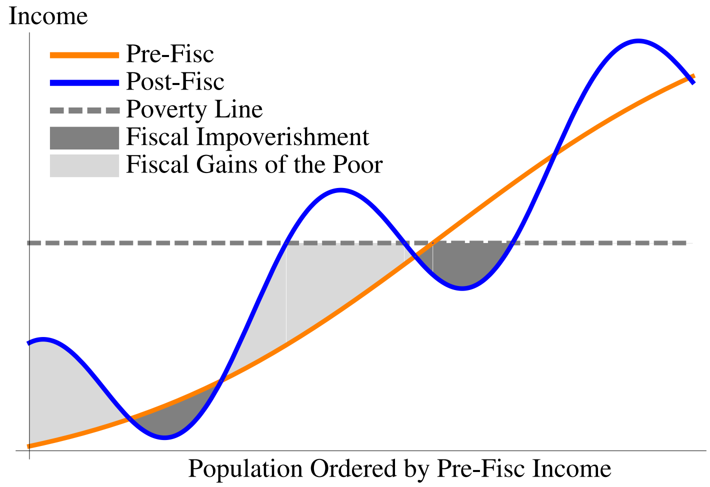

To illustrate visually, Figure 1.1 shows a stylistic representation of the pre- and post-fisc incomes of a population ordered by pre-fisc income. The orange curve represents pre-fisc income, blue post-fisc income, and dashed gray the poverty line; because some individuals receive more in transfers than they pay in taxes, while others pay more in taxes than they receive in transfers, the blue post-fisc income curve is sometimes above and sometimes

[7] Amiel and Cowell (1994) also point out that the respect for income dominance axiom is only equivalent to the monotonicity axiom when anonymity is imposed. In the example from the previous paragraph, the post-fisc income distribution first order stochastically dominates the pre-fisc distribution on the domain from 0 to the maximum poverty line, so it would be evaluated as superior by any measure satisfying poverty focus and respect for income dominance (or, equivalently, poverty focus and both monotonicity and anonymity). It would not necessarily be evaluated as superior by a measure satisfying poverty focus and monotonicity but not anonymity, however. Other concerns with the anonymity axiom have also been pointed out: for example, it can clash with the Pigou–Dalton transfer axiom when there are households of different types (Ebert, 1997) and with the subgroup sensitivity axiom, an extension of the Pigou–Dalton transfer axiom to subgroups (Subramanian, 2006).

below the orange pre-fisc income curve. Although post-fisc poverty is lower than pre-fisc poverty because the losses of some poor are more than compensated by the gains of other poor, there is FI. The extent of FI is shown by the dark-shaded areas, while the light-shaded areas represent the extent of FGP (using the measures we axiomatically derive in Sections 1.3 and 1.4).

1.2.2 Horizontal Equity and Progressivity

Anonymity is not the only reason conventional measures overlook fiscal impoverishment: non-anonymous measures such as horizontal equity and progressivity, which are designed to incorporate information about an individual's pre-fisc position, can fail to capture FI because they are not concerned with whether her net tax burden (taxes paid minus transfers received) is positive or negative. Denote income before taxes and transfers by $y_i^0 \in \mathbb{R}_+$ and income after taxes and transfers by $y_i^1 \in \mathbb{R}_+$ for each $i \in S$, where S is the set of individuals in society. Consider a range of potential poverty lines $\mathcal{Z} \subset \mathbb{R}_+$. Each individual's income before or after taxes and transfers is arranged in the vector y^0 or y^1, both ordered in ascending order of pre-fisc income y_i^0—even if reranking occurs, the order of the y^1 vector reflects the pre-fisc income ranking.

Horizontal equity can be defined in two ways: the reranking definition, which requires that no pair of individuals switch ranks, and the classical definition, which requires that pre-fisc equals are treated equally by the tax and transfer system. Under either definition, the existence or absence of horizontal equity among the poor does not tell us whether FI has occurred. Even if some are impoverished by the tax and transfer system, the ranking among the poor may not change (so there is horizontal equity by the reranking definition) and pre-fisc equals may be impoverished to the same degree (so there is classical horizontal equity): e.g., $\mathcal{Z} = (6, 10], y^0 = (1, 1, 7, 7, 13), y^1 = (3, 3, 6, 6, 11)$. Nor does horizontal inequity among the poor necessarily imply FI, because there could be reranking among the poor or unequal treatment among pre-fisc equals when the tax and transfer system lifts incomes of some of the poor without decreasing incomes of any poor: e.g., $\mathcal{Z} = (6, 10], y^0 = (5, 5, 6, 20), y^1 = (5, 7, 6, 18)$.

A tax and transfer system is everywhere progressive when net taxes (i.e., taxes minus benefits), relative to pre-fisc income, increase with income (Duclos, 1997; Lambert, 1988). The tax and transfer system can be progressive (and unambiguously equalizing) but cause fiscal impoverishment: e.g., $\mathcal{Z} = (6, 10]$, $y^0 = (1, 3, 7, 13)$, $y^1 = (3, 4, 6, 11)$; net taxes relative to pre-fisc income increase with income, but the third individual whose income falls from 7 to 6 is fiscally impoverished; thus, progressivity is not a sufficient condition to ensure that FI does not occur. Nor is progressivity a necessary condition for the absence of FI: e.g., $\mathcal{Z} = (6, 10]$, $y^0 = (1, 3, 7, 14)$, $y^1 = (1, 5, 8, 11)$, which involves no FI but is not everywhere progressive because net taxes first decrease with income when moving from the poorest to the second-poorest, then increase with income thereafter.

Table 1.1: Summary of the Problems with Conventional Measures

Measure	Issue	Example with $\mathcal{Z} = (6, 10]$
Poverty (and stochastic dominance)	$\downarrow$ poverty $\not\Rightarrow$ no FI (anonymity)	$y^0 = (5, 8, 20)$, $y^1 = (9, 6, 18)$
Horizontal equity	Horizontally equitable $\not\Rightarrow$ no FI	$y^0 = (1, 1, 7, 7, 13)$, $y^1 = (3, 3, 6, 6, 11)$
	No FI $\not\Rightarrow$ horizontally equitable	$y^0 = (5, 5, 6, 20)$, $y^1 = (5, 7, 6, 18)$
Progressivity	Progressive $\not\Rightarrow$ no FI	$y^0 = (1, 3, 7, 13)$, $y^1 = (3, 4, 6, 11)$
	No FI $\not\Rightarrow$ progressive	$y^0 = (1, 3, 7, 14)$, $y^1 = (1, 5, 8, 11)$

Table 1.1 summarizes the examples presented in Sections 1.2.1 and 1.2.2 to show that conventional tools—specifically, poverty measures (and stochastic dominance tests) and measures of or tests for horizontal equity and progressivity—can overlook FI.

1.2.3 Real-World Examples

The problems with conventional measures are not limited to contrived hypothetical examples. In a number of countries, we observe an unambiguous reduction in poverty and a globally progressive tax and transfer system, while a significant proportion of the poor are fiscally impoverished. Using the income concepts from Higgins et al. (2015), we compare market income (before taxes and transfers) to post-fiscal income (after direct and indirect taxes, direct cash and food transfers, and indirect subsidies) in seventeen developing countries. We use post-fiscal income as the after taxes and transfers income concept even though taxes are used to fund more than just direct cash and food transfers and indirect subsidies from the government (e.g., they are used to fund public goods and services, many of which also reach the poor) because this is the income concept relevant for measuring poverty: it is "disposable money and near-money income" that should be compared to the poverty line when the latter is based on "a poverty budget for food, clothing, shelter, and similar items" (Citro and Michael, 1995, p. 212, 237). For low and lower-middle income countries, we use a poverty line of $1.25 per person per day; for upper middle income countries, $2.50 per day. Table 1.2 column 1 shows the pre-fisc (market income) poverty headcount and column 2 shows the change in poverty from the pre-fisc to the post-fisc income distribution; countries in which poverty increased due to the fiscal system are excluded.[8]

Moving to the progressivity of the tax and transfer system and change in inequality in each country, column 3 shows the pre-fisc Gini coefficient and column 4 shows the Reynolds and Smolensky (1977) index, which is a summary indicator corresponding to tests of global progressivity; the Reynolds-Smolensky equals the pre-fisc Gini minus the concentration coefficient of post-fisc income with respect to pre-fisc income, and thus globally progressive systems have a positive Reynolds-Smolensky index. Column 5 shows the change in inequality, with negative numbers indicating that inequality fell as a result of the tax and transfer system.[9]

Since we do not derive an axiomatic measure of FI until Section 1.3, here we use two intuitively appealing measures likely to have policy traction. Column 6 shows the percent *of the population* that are fiscally impoverished and column 7 the percent *of the post-fisc*

[8] Although the table only shows poverty for a particular poverty line and poverty measure, it is also true that the post-fisc distribution first order stochastically dominates the pre-fisc distribution from 0 to the poverty line used for each country, meaning that poverty unambiguously fell for all poverty lines up to $1.25 or $2.50 and all poverty measures in a broad class.

[9] We test global progressivity by dominance of the concentration curve of post-fisc with respect to pre-fisc income over the pre-fisc Lorenz curve, and test unambiguously equalizing by comparing the post-fisc and pre-fisc Lorenz curves.

Table 1.2: Poverty, Inequality, and Fiscal Impoverishment in Developing Countries

Country (survey year)	(1) Pre-Fisc Poverty Headcount (%)	(2) Change in Poverty Headcount (p.p.)	(3) Pre-Fisc Inequality (Gini)	(4) Reynolds-Smolensky (Post-Fisc w.r.t. Pre-Fisc)	(5) Change in Inequality (ΔGini)	(6) Fiscally Impoverished as % of Population	(7) Fiscally Impoverished as % of Post-Fisc Poor
Panel A: Upper-middle income countries, using a poverty line of $2.50 per day							
Brazil (2008–2009)	16.8	−0.8	57.5	4.6	−3.5	5.6	34.9
Chile (2013)	2.8	−1.4	49.4	3.2	−3.0	0.3	19.2
Ecuador (2011–2012)	10.8	−3.8	47.8	3.5	−3.3	0.2	3.2
Mexico (2012)	13.3	−1.2	54.4	3.8	−2.5	4.0	32.7
Peru (2011)	13.8	−0.2	45.9	0.9	−0.8	3.2	23.8
Russia (2010)	4.3	−1.3	39.7	3.9	−2.6	1.1	34.4
South Africa (2010–2011)	49.3	−5.2	77.1	8.3	−7.7	5.9	13.3
Tunisia (2010)	7.8	−0.1	44.7	8.0	−6.9	3.0	38.5
Panel B: Lower-middle income countries, using a poverty line of $1.25 per day							
Armenia (2011)	21.4	−8.4	47.4	12.9	−9.2	6.2	52.3
Bolivia (2009)	10.9	−0.5	50.3	0.6	−0.3	6.6	63.2
Dominican Republic (2007)	6.8	−0.9	50.2	2.2	−2.2	1.0	16.3
El Salvador (2011)	4.3	−0.7	44.0	2.2	−2.1	1.0	27.0
Guatemala (2010)	12.0	−0.8	49.0	1.4	−1.2	7.0	62.2
Indonesia (2012)	12.0	−1.5	39.8	1.1	−0.8	4.1	39.2
Sri Lanka (2009–2010)	5.0	−0.7	37.1	1.3	−1.1	1.6	36.4

Sources: For Brazil, authors' calculations. For other countries, provided to us by the authors of the studies cited in footnote 3.
Notes: p.p. = percentage points. w.r.t. = with respect to. Ethiopia and Ghana are not included in the table because poverty with a $1.25 per day poverty line increased from pre-fisc to post-fisc income (and hence they do not illustrate shortcomings of conventional measures). Country classifications are from the World Bank for the year of the survey.

poor that are fiscally impoverished. Although all of the countries in Table 1.2 experienced a reduction in poverty and inequality due to the tax and transfer system, the amount of FI varies greatly between countries. In ten countries—Armenia, Bolivia, Brazil, El Salvador, Guatemala, Indonesia, Mexico, Russia, Sri Lanka, and Tunisia—between one-quarter and two-thirds of the post-fisc poor lost income to the fiscal system.[10] In other countries, this figure is much lower, at 13.3% of the post-fisc poor in South Africa (but, due to the high proportion of the total population that is poor, still 5.9% of the total population) and 3.2% of the post-fisc poor in Ecuador.

Even when poverty increases from pre-fisc to post-fisc income and hence we know

[10] If we instead scale down taxes so that they equal the transfers included in our analysis, which we avoid in the main analysis for the reasons mentioned above in defense of post-fiscal income as the after taxes and transfers income concept, FI is lower: for example, in Brazil 10.8% of the post-fisc poor are fiscally impoverished using this method.

that FI has occurred (as in Ghana and Ethiopia), it is impossible to tell its extent without explicit measures like the ones we propose in Section 1.3. A stark example of this comes from Ethiopia, where looking at poverty and progressivity numbers alone greatly masks the extent of FI: the headcount ratio at $1.25 per day increases from 31.9% to 33.2% of the population, while squared poverty gap and Gini coefficient fall as a result of taxes and transfers (World Bank, 2015). Nevertheless, applying our measures to the same data, Hill et al. (forthcoming) find that 28.5% of Ethiopians and over 80% of the post-fisc poor experience FI.

Even if we add the value of public spending on education and health (imputed at their government cost to families who report a child attending public school or who report using public health facilities), fiscal impoverishment is still high in several countries: in Armenia, Ethiopia, Indonesia, Tunisia, and Russia, between 25 and 50% of those who are fiscally impoverished before adding in benefits from public spending on health and education are still fiscally impoverished when these benefits are included as transfers.

1.3 Measures of Fiscal Impoverishment

To assess anti-poverty policies in tandem with the taxes used to finance them, it is important to have measures of the extent of fiscal impoverishment. In the last section, we provided a glimpse of FI in several developing countries using two simple, straight-forward, and intuitive measures that—given these features—can be useful for policy discussions. These two measures also have drawbacks, however. To illustrate their limitations, we begin by providing more detail about the two measures. For a particular poverty line $z \in \mathcal{Z}$, there is *fiscal impoverishment* if $y_i^1 < y_i^0$ and $y_i^1 < z$ for some individual $i \in S$. In other words, the individual could be poor before taxes and transfers and made poorer by the fiscal system, or non-poor before taxes and transfers but poor after. Both straight-forward measures count the number of individuals who meet this condition (and are thus fiscally impoverished) in the numerator. The proportion *of the population* who are fiscally impoverished (column

6 of Table 1.2) divides this numerator by the number of individuals in society, while the proportion *of the post-fisc poor* who are fiscally impoverished (column 7) divides it by the number who are post-fisc poor (with $y_i^1 < z$).

In the context of poverty measurement, Sen (1976, p. 219) proposes a monotonicity axiom requiring that, all else equal, "a reduction in income of a person below the poverty line must increase the poverty measure." We propose a similar axiom for FI measures requiring that a larger decrease in post-fisc income for an impoverished person, all else equal, must increase the FI measure. Monotonicity is violated by the straight-forward measures, which do not increase when an impoverished person becomes more impoverished because she counts as one impoverished individual in the measure's numerator regardless of how much income she loses to the fiscal system.[11]

1.3.1 Axioms

We propose eight properties desirable for a robust measure of FI; we describe these properties here and formally define them in the Appendix. Throughout, we assume that income is measured in real terms and has been converted to a common currency such as US dollars adjusted for purchasing power parity, thereby simplifying away concerns about inflation or currency conversions if comparing FI over time or across countries.

Our **FI monotonicity** axiom described above implies not only that the FI measure must be strictly increasing in the extent to which an impoverished individual is impoverished (ceteris paribus), but also that the measure must be strictly increasing in the number of individuals that are impoverished, holding fixed the amount of FI experienced by others. The **focus** axiom, analogous to Sen's (1981) focus axiom for poverty measurement, says that different income changes to the non-impoverished—provided that they remain non-impoverished—leave the FI measure unchanged. Given the focus axiom, it is natural to

[11] Another simple tool to examine FI is the $q \times q$ transition matrix P, whose typical element p_{kl} represents the probability of being in post-fisc income group $l \in \{1, \ldots, q\}$ for an individual in pre-fisc income group $k \in \{1, \ldots, q\}$. Measures based on P also fail to satisfy FI monotonicity and have the large drawback of not capturing FI among the poorest pre-fisc group ($k = 1$).

impose a **normalization** that if no one is impoverished, the FI measure equals zero. Note that this normalization axiom is not instrumental to our result: if we did not impose it, our result would be that our axioms uniquely determine a measure of FI up to a linear (rather than proportional) transformation.[12]

Similar to Chakravarty's (1983) **continuity** axiom for poverty measures, we require the FI measure to be continuous in pre-fisc income, post-fisc income, and the poverty line (since we may want to assess FI for a range of possible poverty lines). This is stronger than Foster and Shorrocks's (1991) restricted continuity axiom which only requires the measure to be continuous in incomes *below* the poverty line and left-continuous *at* the poverty line, thus allowing the measure to jump discontinuously at the poverty line; see Zheng (1997) and Permanyer (2014) for arguments in favor of using the stronger continuity axiom in the contexts of unidimensional and multidimensional poverty measures.

Because "the names of income recipients do not matter" (Zheng, 1997, p. 131), we impose a **permutability** axiom requiring that if we take each individual's pre- and post-fisc income pair and (keeping each pre- and post-fisc income pair as a bundle) shuffle these around the population, FI is unchanged. We use the term "permutability" rather than symmetry or anonymity because—although both have been used in the same way we use permutability above (e.g., Cowell, 1985; Fields and Fei, 1978; Plotnick, 1982)—symmetry and anonymity have also taken on different definitions. Symmetry can instead mean, for two income distributions X and Y and a distance measure d, that $d(X, Y) = d(Y, X)$; the two income distributions are treated symmetrically: losses are not distinguishable from gains (Ebert, 1984; Fields and Ok, 1999). Anonymity can instead mean that the measure compares the cumulative distribution of pre-fisc income, F_0, to that of post-fisc income, F_1, without regard to where a particular individual at position j in F_0 ended in F_1 (e.g., Bourguignon, 2011a,b). In other words, an anonymous measure would compare the pre-

[12] It is also possible to normalize by the measure's upper bound so that it always lies on the interval $[0, 1]$ by specifying an axiom that if everyone loses all of their income to the fiscal system (the maximum possible FI), the measure of FI equals 1. We prefer to avoid normalizing in this way so that the class of axiomatic FI measures is more general.

fisc income of the jth poorest individual in F_0 to the post-fisc income of the jth poorest individual in F_1, even though "they are not necessarily the same individuals" because of reranking (Bourguignon, 2011a, p. 607).

Next, we must decide whether our measure of FI should be absolute or relative (recalling that we assume income to be in real terms of a constant currency, so arguments about inflation or currency exchange should not affect the decision). Suppose each poor individual's pre-fisc income increases by \$1, taxes and transfers are held fixed, and the price of one essential good in the basic goods basket, normalized to have one unit in the basket, also increases by \$1 per unit.[13] Each poor individual remains the same distance below the poverty line; that distance represents the amount of additional income she needs to afford adequate nutrition and other basic necessities. For those who experience FI, it is the absolute increase in the distance between that individual's income and the poverty line that matters in terms of the quantity of basic goods she can buy. Hence, we assume that if all pre- and post-fisc incomes increase by \$1 and the poverty line also increases by \$1, FI should remain unchanged. We thus impose **translation invariance**.

Given our above argument for absolute measures, we also impose **linear homogeneity**: if all incomes and the poverty line are multiplied by the same factor, the measure of FI changes by that factor. Instead, specifying homogeneity of degree zero (scale invariance) would be incompatible with translation invariance for the reasons explored in Zheng (1994). Since we assume income is expressed in real terms and a common currency, our measure is nevertheless insensitive to inflation or currency changes. The translation invariance and linear homogeneity axioms have been used together in axiomatic derivations of measures of inequality (Kolm, 1976), poverty (Blackorby and Donaldson, 1980), economic distance (Chakravarty and Dutta, 1987; Ebert, 1984), and mobility (Fields and Ok, 1996; Mitra and Ok, 1998).[14]

[13] To avoid inflation in this thought experiment, assume that there is an offsetting fall in the price of a good *not* in the basic good basket and not consumed by the poor.

[14] By requiring translation invariance and linear homogeneity, we are deriving a measure of *absolute* FI; from there, the measure can nevertheless be modified to obtain other types of desired measures such as a scale

Our final axiom is based on a concept introduced to the poverty literature by Foster et al. (1984, p. 761), who argue that "at the very least, one would expect that a decrease in the poverty level of one subgroup ceteris paribus should lead to less poverty for the population as a whole." Similarly, it would be desirable for a measure of FI if a decrease in the measured FI for one subgroup of the population and no change in the measured FI for all other subgroups results in a decrease in the measured FI of the entire population. Hence, we impose a **subgroup consistency** axiom analogous to the one used for poverty measurement by Foster and Shorrocks (1991). In his survey of axiomatic poverty measurement, Zheng (1997, p. 137) notes that subgroup consistency "has gained wide recognition in the literature."

1.3.2 An Axiomatic Measure of Fiscal Impoverishment

Proposition 1. *A measure satisfying FI monotonicity, focus, normalization, continuity, permutability, translation invariance, linear homogeneity, and subgroup consistency is uniquely determined up to a proportional transformation, and given by*

$$f(y^0, y^1; z) = \kappa \sum_{i \in S} \big(\min\{y_i^0, z\} - \min\{y_i^0, y_i^1, z\} \big). \tag{1.1}$$

The summand for individual i behaves as follows. For an individual who was poor before taxes and transfers and is impoverished ($y_i^1 < y_i^0 < z$), it is equal to her fall in income, $y_i^0 - y_i^1$. For an individual who was non-poor before taxes and transfers and is impoverished ($y_i^1 < z \le y_i^0$), it equals her post-fisc poverty gap, or the amount that would need to be transferred to her to move her back to the poverty line (equivalently, to prevent her from becoming impoverished), $z - y_i^1$. For a non-impoverished pre-fisc non-poor individual ($y_i^0 \ge z$ and $y_i^1 \ge z$) it equals $z - z = 0$. For a non-impoverished pre-fisc poor individual ($y_i^0 < z$ and $y_i^1 \ge y_i^0$) it equals $y_i^0 - y_i^0 = 0$. Hence, f sums the total amount

invariant measure. This is similar to the approach taken by Fields and Ok (1996), who axiomatically derive a measure of absolute mobility from which other desired measures such as mobility proportional to income can be obtained.

of FI, multiplied by a factor of proportionality. This constant can be chosen based on the preferences of the practitioner: for example, $\kappa = 1$ gives total FI (the dark-shaded area in Figure 1.1), while $\kappa = |S|^{-1}$ gives per capita FI.[15]

1.3.3 Fiscal Impoverishment Dominance Criteria

Having identified the existence of FI in a country, a useful implementation of our FI measure would be to compare the degree of FI in two situations, e.g. by comparing the current fiscal system to a proposed reform. The choice of poverty line might, however, influence our conclusion about which situation entails higher FI. We thus present a partial FI ordering that can be used to determine if FI is unambiguously lower in one situation than another for any poverty line and any measure that satisfies FI monotonicity, focus, normalization, continuity, permutability, translation invariance, linear homogeneity, and subgroup consistency. Since we have already shown that a FI measure satisfies these axioms if and only if it takes the form in (1.1), a simple way to test for FI dominance for any measure satisfying those axioms and any poverty line in the domain of poverty lines $\mathcal{Z}$ is to simply compare the curves $f(y^0, y^1; z)$ and $f(x^0, x^1; z)$ across $\mathcal{Z}$. Interestingly, if the minimum poverty line being considered is 0 (so $\mathcal{Z} = [0, z^+]$, where z^+ is the maximum poverty line), there is an alternative (equivalent) way to test whether FI is unambiguously lower in one situation than another that uses a dominance test already developed in the mobility literature: Foster and Rothbaum's (2014) second order downward mobility dominance.

Proposition 2. *The following are equivalent.*

a) FI is unambiguously lower in (y^0, y^1) than (x^0, x^1) for any poverty line in $[0, z^+]$ and any measure satisfying FI monotonicity, focus, normalization, continuity, permutability, translation invariance, linear homogeneity, and subgroup consistency.

[15] We do not impose a population invariance axiom; this axiom is commonly imposed but is criticized by Hassoun and Subramanian (2012). A subset of measures of form (1.1) are population invariant: choosing $\kappa = |S|^{-1}$ gives a measure that satisfies population invariance, while $\kappa = 1$ gives a measure that does not.

b) $f(y^0, y^1; z) < f(x^0, x^1; z) \; \forall \, z \in [0, z^+]$.

c) (y^0, y^1) *second order downward mobility dominates* (x^0, x^1) *on* $[0, z^+]$.

1.4 Fiscal Gains of the Poor

Most likely, we will be interested in more than just the extent to which some poor are not compensated for their tax burden with transfers: we will also want to know about the gains of other poor families, and the way in which a comparison of poverty before and after taxes and transfers can be decomposed into the losses and gains of different poor households. In this section, we formally define fiscal gains of the poor, briefly present the axioms for a measure of FGP analogous to those in Section 1.3.2 for a measure of FI, and present an axiomatic measure and partial ordering of FGP. We then show that a commonly used measure of poverty, the poverty gap, can be decomposed into our axiomatic measures of FI and FGP.

1.4.1 An Axiomatic Measure of Fiscal Gains of the Poor

There are *fiscal gains of the poor* if $y_i^0 < y_i^1$ and $y_i^0 < z$ for some individual $i \in S$. The individual may or may not receive enough in net transfers to be post-fisc non-poor (i.e., it is possible that $z \leq y_i^1$ or $y_i^1 < z$). Consider a pre-fisc poor individual who receives more in transfers than she pays in taxes. If she is given even more transfer income, while the pre- and post-fisc incomes of all others experiencing FGP do not change, FGP should not decrease; if she would have remained in poverty post-fisc without the additional transfer income, FGP should increase with the additional transfer. We impose these conditions in the **FGP monotonicity** axiom; we also impose FGP analogues of the other axioms from Section 1.3.2.

Proposition 3. *A measure satisfying FGP monotonicity, focus, normalization, continuity, permutability, translation invariance, linear homogeneity, and subgroup consistency is* *uniquely determined up to a proportional transformation, and given by*

$$g(y^0, y^1; z) = \kappa \sum_{i \in S} \big(\min\{y_i^1, z\} - \min\{y_i^0, y_i^1, z\}\big). \tag{1.2}$$

An individual who is pre-fisc poor and gains income from the tax and transfer system, but remains post-fisc poor ($y_i^0 < y_i^1 < z$), contributes the amount of her income gain, $y_i^1 - y_i^0$, to the measure of FGP. A pre-fisc poor individual that gains income and as a result has post-fisc income above the poverty line ($y_i^0 < z \leq y_i^1$) contributes the amount of net transfers that pulled her pre-fisc income to the poverty line, $z - y_i^0$. Someone who is pre-fisc poor and does not gain income ($y_i^1 \leq y_i^0 < z$) contributes $y_i^1 - y_i^1 = 0$. Someone who is pre-fisc non-poor ($z < y_i^0$) also contributes 0 (for her, the summand equals $z - z$ if she remains non-poor or $y_i^1 - y_i^1$ if she loses income and becomes poor). For $\kappa = 1$, g equals the light-shaded area in Figure 1.1.

As with fiscal impoverishment orderings, a fiscal gain partial ordering can be used to make unambiguous FGP comparisons for any poverty line and any measure satisfying our axioms. The ordering compares $g(y^0, y^1; z)$ to $g(x^0, x^1; z)$ for all $z \in \mathcal{Z}$, and for $\mathcal{Z} = [0, z^+]$ coincides with Foster and Rothbaum's (2014) second order upward mobility dominance (the proof proceeds similarly to the proof of Proposition 2 for FI).

1.4.2 Decomposition of the Difference between Pre-Fisc and Post-Fisc Poverty

The most common measures of poverty used in both policy circles and scholarly papers (e.g., Chen and Ravallion, 2010; Ravallion, 2012) are the poverty headcount ratio, which enumerates the proportion of the population that is poor, and the poverty gap, which takes into account how far the poor fall below the poverty line. The latter might be expressed in absolute terms, summing the gap between each poor person's income and the poverty line, in which case it can be thought of as the total amount that would need to be given to the poor to eliminate poverty (if targeting were perfect). Or it can be normalized, dividing the absolute poverty gap by the poverty line and population size, for example, to create a scale- and population-invariant measure. We use a general definition of the poverty gap that encompasses its absolute and normalized forms:

$$p(y;z) = \nu(S,z)\sum_{i\in S}(z-y_i)\mathbb{I}(y_i<z), \tag{1.3}$$

where $\nu(S,z)$ is a normalization factor. Two special cases are the *absolute poverty gap*, where $\nu(S,z)=1$, and the *poverty gap ratio*, where $\nu(S,z)=(z|S|)^{-1}$. For simplicity and because a comparison of pre- and post-fisc poverty usually occurs for a fixed population and given poverty line, we assume that S and z are fixed in what follows.

Proposition 4. *A change in the poverty gap before and after taxes and transfers is equal to the difference between the axiomatic measures of FI and FGP from* (1.1) *and* (1.2)*, multiplied by a constant.*

Given the assumption that the population and poverty line are fixed, $\nu(S,z)$ is a constant that we denote $\bar{\nu}$. The poverty gap in (1.3) can be rewritten as $p(y;z)=\bar{\nu}\sum_{i\in S}(z-y_i)\mathbb{I}(y_i<z)=\bar{\nu}\sum_{i\in S}(z-\min\{y_i,z\})$, so we have $p(y^1;z)-p(y^0;z)=\bar{\nu}\sum_{i\in S}\left(z-\min\{y_i^1,z\}\right)-\bar{\nu}\sum_{i\in S}\left(z-\min\{y_i^0,z\}\right)$, or

$$\begin{aligned}
p(y^1;z)-p(y^0;z) &= \bar{\nu}\Bigg[\sum_{i\in S}\left(\min\{y_i^0,z\}-\min\{y_i^0,y_i^1,z\}\right) \\
&\quad -\sum_{i\in S}\left(\min\{y_i^1,z\}-\min\{y_i^0,y_i^1,z\}\right)\Bigg] \\
&= \frac{\bar{\nu}}{\kappa}\left[f(y^1,y^0;z)-g(y^1,y^0;z)\right].
\end{aligned}$$

Comparisons of pre- and post-fisc poverty are often used to assess whether the tax and transfer system helps or hurts the poor. This decomposition can be used to dig deeper into that net effect and observe the extent to which a net reduction in poverty masks the offsetting gains of some poor and impoverishment of others at the hands of the (possibly progressive) tax and transfer system.

1.4.3 Results for Seventeen Developing Countries

We saw in Section 1.2 that in fifteen of seventeen developing countries for which we have data, the tax and transfer system is poverty-reducing and progressive but, in many cases, fiscally impoverishes a significant proportion of the poor. In Table 1.3, we present FI and FGP results for these countries using the axiomatic measures derived in Sections 1.3 and 1.4. Column 1 gives total FI (i.e., the axiomatic measure from (1.1) with $\kappa = 1$) and column 2 total FGP, both expressed in millions of 2005 US dollars per year using purchasing power parity adjusted exchange rates. Because the axiomatic measure with $\kappa = 1$ is population variant, FI and FGP tend to be higher in more populous countries; these absolute amounts of FI and FGP can be useful, for example, in comparisons to the size of a country's main cash transfer program, as we show for Brazil below. To ease interpretation and comparison across countries, column 3 shows FI expressed as a percent of FGP, while columns 4 and 5 show FI and FGP per capita (where per capita refers to dividing by the entire population), normalized by the poverty line; each of these is population invariant.

There is large heterogeneity in the extent to which some poor are hurt by the tax and transfer system relative to the extent to which other poor gain, despite that the same range of policies, including direct taxes, direct cash and near-cash transfers, indirect consumption taxes, and indirect subsidies were considered in each country study. Among the upper-middle income countries, FI as a percent of FGP (using a poverty line of $2.50 per day) ranges from less than 1% in Ecuador to 40% in Tunisia. In low and lower-middle income countries, FI as a percent of FGP (using a poverty line of $1.25 per day) is even higher in some countries, reaching 55% in Guatemala and 81% in Bolivia; in Ethiopia and Ghana—the two countries in which post-fisc poverty is higher than pre-fisc poverty—FI exceeds FGP.

Column 6 shows the change in the poverty gap ratio from pre-fisc to post-fisc income,

Table 1.3: Fiscal Impoverishment and Gains of the Poor in Developing Countries

Country (survey year)	(1) Total FI ($ millions per year)	(2) Total FGP ($ millions per year)	(3) FI as % of FGP	(4) Per capita FI as % of z	(5) Per capita FGP as % of z	(6) Change in Poverty Gap Ratio (p.p.)
Panel A: Upper-middle income countries, using a poverty line of $2.50 per day						
Brazil (2008–2009)	676.0	3503.6	19.3	0.39	2.02	−1.63
Chile (2013)	2.0	93.3	2.1	0.01	0.59	−0.58
Ecuador (2011–2012)	1.1	277.8	0.4	0.01	2.00	−1.99
Mexico (2012)	227.7	1446.5	15.7	0.21	1.35	−1.14
Peru (2011)	53.7	177.0	30.3	0.20	0.65	−0.45
Russia (2010)	84.9	1561.4	5.4	0.07	1.24	−1.17
South Africa (2010-2011)	186.6	5964.0	3.1	0.41	12.96	−12.56
Tunisia (2010)	20.8	52.0	40.0	0.23	0.59	−0.35
Panel B: Low and lower-middle income countries, using a poverty line of $1.25 per day						
Armenia (2011)	6.3	117.9	5.3	0.44	8.17	−7.74
Bolivia (2009)	25.9	32.2	80.6	0.55	0.68	−0.13
Dominican Republic (2007)	4.4	105.1	4.2	0.02	0.53	−0.51
El Salvador (2011)	1.2	11.1	11.1	0.04	0.39	−0.35
Ethiopia (2010–2011)	408.9	392.8	104.1	1.18	1.13	0.05
Ghana (2013)	25.9	9.9	262.1	0.22	0.08	0.13
Guatemala (2010)	20.7	37.8	54.9	0.33	0.61	−0.27
Indonesia (2012)	150.2	531.5	28.3	0.13	0.47	−0.34
Sri Lanka (2009–2010)	4.4	25.5	17.1	0.05	0.27	−0.23

Sources: For Brazil, authors' calculations. For other countries, provided to us by the authors of the studies cited in footnote 3.
Notes: p.p. = percentage points. z denotes the poverty line. "$ millions" denotes millions of 2005 US dollars, at purchasing power parity adjusted exchange rates. Country classifications are from the World Bank for the year of the survey.

which by Propostion 4 can be decomposed into FI per capita minus FGP per capita, both normalized by the poverty line like the poverty gap ratio. This decomposition reveals some interesting traits of each country's tax and transfer system. For example, Ecuador achieves the same FGP per capita as Brazil but with nearly no FI, compared to substantial FI in Brazil; as a result, the poverty gap is reduced by more in Ecuador. The difference in FI might be attributable to the the multiple consumption taxes levied at the state and federal levels in Brazil: these are high and often cascading, and consumption tax exemptions for basic goods are almost non-existent (Corbacho et al., 2013), compared to a system that exempts food, basic necessities, and medicine in Ecuador (Llerena Pinto et al., 2015). Interestingly, most of those experiencing FI are *not* excluded from the safety net; they *do* receive government transfers or subsidies: 65% of the impoverished in Brazil receive cash transfers from Bolsa Família, for example. It is also noteworthy that Peru, one of the countries in which *less than* a quarter of the post-fisc poor experience FI, nevertheless redistributes low amounts to the poor, and thus has a low reduction in the poverty gap; this is consistent with Jaramillo's (2014, p. 391) finding that Peru's low poverty reduction induced by fiscal policy is "associated with low social spending rather than with inefficient spending." Among three lower-income countries that each reduce the poverty gap ratio by about 0.3 percentage points (El Salvador, Guatemala, and Indonesia), Guatemala has high FI but also higher FGP, while El Salvador has lower FGP but very low FI, and Indonesia falls in the middle. We do not attempt to answer whether a lower-FI, lower-FGP or higher-FI, higher-FGP system is preferable from a welfare perspective, but note that this decomposition enables a substantially richer analysis than the typical comparison of poverty before and after taxes and transfers.

1.4.4 Results for a Range of Poverty Lines in Brazil

So far, the FI and FGP results we have presented use a fixed poverty line ($1.25 in low and lower-middle income countries and $2.50 in upper-middle income countries). We now extend the analysis to a range of poverty lines, focusing the illustration on data from Brazil, using the Pesquisa de Orçamentos Familiares (Family Expenditure Survey) 2008–2009. The precise direct and indirect taxes, direct cash and food transfers, and indirect subsidies included in our analysis are described in detail in Higgins and Pereira (2014).

As we stated in Section 1.2.3, the tax and transfer system in Brazil is unambiguously poverty-reducing for any poverty line up to $2.50 per person per day, globally progressive, and unambiguously equalizing.[16] This is shown in Figure 1.2, where cumulative distri-

[16] Nevertheless, the tax and transfer system reduces poverty by less than its potential under the type of optimal redistribution considered by Fellman et al. (1999), which follows a lexicographic maximin principle. Replacing the actual tax system with optimal taxes of this type (which, in total, equal the size of actual taxes), and replacing the actual distribution of Bolsa Família benefits with the optimal one (redistributing all transfers this way would completely eliminate poverty, so we only optimally redistribute Bolsa Família for illustration), the lowest income in the population would be $1.92 per day, the post-fisc poverty gap ratio

bution functions reveal that the post-fisc distribution first order stochastically dominates the pre-fisc distribution on the domain $[0, 2.5]$, which implies an unambiguous reduction in poverty for any poverty line in this domain and any measure in a broad class (Atkinson, 1987; Foster and Shorrocks, 1988);[17] the post-fisc concentration curve with respect to pre-fisc income dominates the pre-fisc Lorenz curve, which implies global progressivity (in the income redistribution sense; see Duclos, 2008); and the post-fisc Lorenz curve dominates the pre-fisc Lorenz curve, which implies that the fiscal system is unambiguously equalizing (Atkinson, 1970). If, however, we extend the maximum poverty line to, say, \$4 per person per day—a poverty line frequently used by the World Bank when studying middle-income Latin American countries (e.g., Ferreira et al., 2013)—poverty is no longer unambiguously lowered by the fiscal system: for poverty lines above about \$3 per day, the poverty headcount is higher after taxes and transfers than before. We would thus know that FI occurred using conventional measures and a poverty line above \$3 per day, but would still be unaware of its extent without FI measures.[18]

Using the \$2.50 line, we know that 5.6% of Brazil's population and over one-third of its post-fisc poor experience FI (Table 1.2); these impoverished individuals pay a total of \$676 million more in taxes than they receive in transfers annually (Table 1.3), which is equivalent to 10% of the 2009 budget of Bolsa Família, Brazil's flagship anti-poverty program that reaches over one-fourth of the country's population. While substantial in size, this FI is dwarfed by FGP from Brazil's transfer programs, which totals over \$3.5 billion. The absolute poverty gap, or the minimum amount that would need to be transferred to the poor to eliminate poverty if transfers were perfectly targeted, falls from \$12.4 billion before taxes and transfers to \$9.6 billion after. The change in the absolute poverty gap, \$2.8 billion, looks impressive, but masks differential trends in two groups of the poor: those who

would be 2.7% of the poverty line rather than 5.5%, and the post-fisc Gini would be 45.3 rather than 53.9.

[17] We verify that this first order dominance is statistically significant at the 5% level using the asymptotic sampling distribution derived by Davidson and Duclos (2000) with a null hypothesis of non-dominance; the result is also robust to the type of data contamination considered in Cowell and Victoria-Feser (2002).

[18] It is easy to show that if the post-fisc distribution does not first order stochastically dominate the pre-fisc distribution on the domain from 0 to the maximum poverty line, then FI has occurred.

Figure 1.2: Conventional Tools to Assess the Tax and Transfer System in Brazil

(a) First Order Stochastic Dominance (Cumulative Distribution Functions) | (b) Global Progressivity (Lorenz and Concentration Curves)

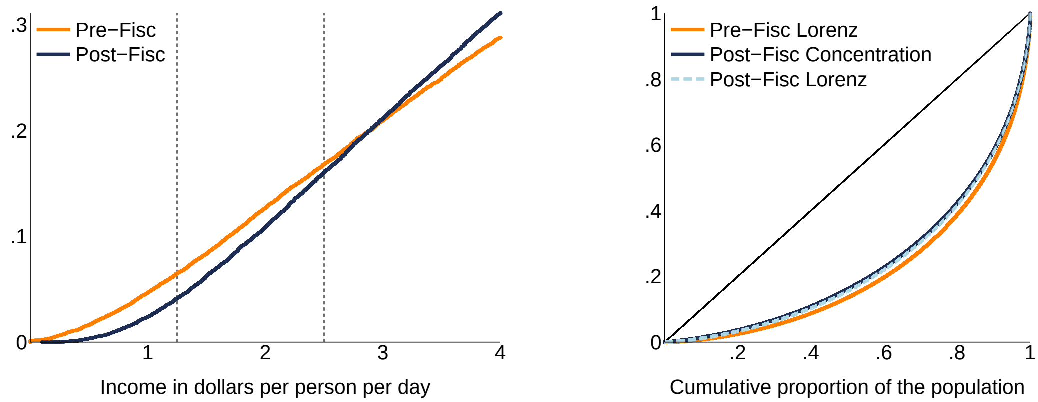

Note: Dashed vertical lines included at common "international" poverty lines of $1.25 and $2.50 per person per day.

gain (a total of $3.5 billion) and those who lose (a total of $676 million), as revealed by the decomposition of the change in the poverty gap derived in Section 1.4.

Figure 1.3: FI, FGP, and Poverty Gaps in Brazil for Various Poverty Lines

(a) Total FI and FGP (Billions of Dollars per Year) | (b) Total Poverty Gaps (Billions of Dollars per Year)

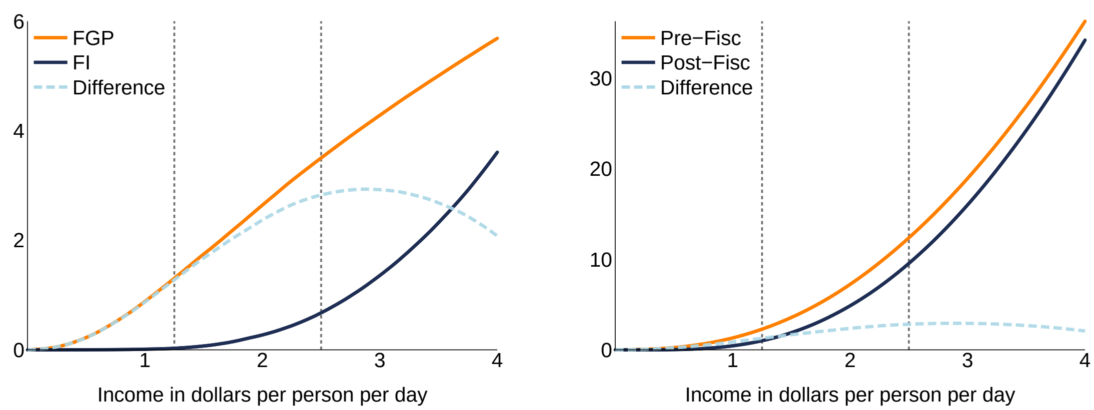

Note: Dashed vertical lines included at common "international" poverty lines of $1.25 and $2.50 per person per day.

Figure 1.3 shows how this decomposition and our axiomatic measures of total FI and FGP in Brazil vary with the poverty line. For low poverty lines, FI is essentially non-existent: at $1.25 per day, for example, total FI is $28 million per year, or 0.4% of the 2009 budget of Bolsa Família (Figure 1.3a). This is not surprising in light of the unconditional component of the government cash transfer program Bolsa Família, available to households with income below 70 reais per person per month ($1.22 per day), regardless of whether the household has children or elderly members, and without conditions. At higher poverty lines, FI begins to increase more rapidly, and at a poverty line of $2.88 the rate of increase of FI exceeds the rate of increase of FGP: this can be seen by comparing the slopes of the solid curves in Figure 1.3a, or by looking at the point where the difference between the two curves (plotted as the dashed curve in Figure 1.3a) is at its maximum. By Proposition 4, this is also the point at which the absolute poverty gap reduction acheived by the fiscal system reaches its maximum, as seen by the dashed curve in Figure 1.3b.

At this poverty line of $2.88 per day, where maximum poverty reduction is achieved, the difference between the pre-fisc and post-fisc poverty gaps is $2.9 billion. The eligibility cut-off for the conditional component of Bolsa Família, available to families with children who comply with certain education and health requirements, is $2.45 per person per day. Just above this line, a number of families still receive benefits due to program leakages, variable and mismeasured income, or components of income we are measuring that are not taken into account in the estimation of eligible income; not far above the line, however, families become much less likely to receive the program and we see a simultaneous deceleration of fiscal gains and acceleration of impoverishment.

1.5 Conclusions

Anti-poverty policies are increasingly being discussed in the same breath as the taxes used to pay for them. One example is the focus on mobilizing domestic resources to finance the policies necessary to achieve the United Nations' Post-2015 Sustainable Development Goals. To analyze transfers, subsidies, and taxes together, poverty comparisons and progressivity measures are often used. These measures, however, can lead us to conclude that the tax and transfer system unambiguously benefits the poor, when in fact a substantial number of poor are not compensated with transfers for their tax burdens. Indeed, we observe this in a number of developing countries: out of seventeen developing countries for which we have data, fifteen have tax and transfer systems that unambiguously reduce poverty and are globally progressive, but in ten of these at least one-quarter of the poor pay more in taxes than they receive in transfers and subsidies. In Brazil, for example, over one-third of the post-fisc poor experience fiscal impoverishment, paying a total of \$676 million more in taxes than they receive in transfers and subsidies.

Given this shortcoming of conventional criteria and the debate about anti-poverty policies and the taxes used to pay for them, we propose a set of axioms that should be met by a measure of FI, and show that these uniquely determine the measure up to a proportional transformation. We also propose a partial ordering to determine when one fiscal system, such as that under a proposed reform, induces unambiguously less FI than another, such as the current system, over a range of possible poverty lines. To obtain a complete picture of the fiscal system's effect on the poor, we propose an analogous measure of fiscal gains of the poor, and show that the difference between the pre-fisc and post-fisc poverty gaps can be decomposed into our axiomatic measures of FI and FGP.

Our results can be extended to comparisons between two points in time or before and after a policy reform, rather than pre- and post-fisc. In comparison to the tools used to assess whether the tax and transfer system hurts the poor, tools from the literatures on pro-poor growth and policy reforms (tax and subsidy reforms, trade liberalization, etc.) suffer from similar limitations. For pro-poor growth,[19] poverty measures and stochastic dominance tests are often used to assess whether poverty is unambiguously reduced over time; it directly follows from the first row of Table 1.1 that these will not necessarily capture that some of the poor become poorer over time. Hence, growth can appear unambiguously

[19] Here, we are using the poverty-reducing or weak absolute definition of pro-poor (in the respective taxonomies of Kakwani and Son (2008) and Klasen (2008)), by which "growth is pro-poor if the poverty measure of interest falls" (Kraay, 2006, p. 198). We could instead adopt a relative definition of pro-poor growth (Kakwani and Pernia, 2000); growth-adjusted stochastic dominance tests can be used to determine when growth is unambiguously relatively pro-poor (Duclos, 2009), and it can be shown that this type of dominance can also occur despite a significant portion of the poor becoming poorer.

pro-poor even if a significant proportion of the poor are immiserized. Growth incidence curves (Ravallion and Chen, 2003) and related pro-poor partial orderings (Duclos, 2009) can fail to capture impoverishment for the same reason that stochastic dominance tests do: they are anonymous with respect to initial income. Although their non-anonymous counterparts (Bourguignon, 2011a; Grimm, 2007; Van Kerm, 2009) resolve this issue in theory, in practice—to become graphically tractable—they average within percentiles, and hence impoverishment can still be overlooked if within some percentiles, some poor are "hurting behind the averages" (Ravallion, 2001, p. 1811).

For consumption tax and subsidy reform, Besley and Kanbur (1988) derive poverty-reducing conditions for reallocating food subsidies; these results are extended to commodity taxes and a broader class of poverty measures by Makdissi and Wodon (2002) and Duclos et al. (2008). Again, by the first row of Table 1.1, unambiguous poverty reduction does not guarantee that a substantial portion of the poor are not hurt by the reform. Studies that evaluate indirect tax reform with measures that take pre-fisc positions into account but average within groups, such as the percent gain or loss caused by the reform for each income or expenditure decile (Mirrlees et al., 2011, Chapter 9), can again overlook FI that occurs within each group.

In the literature on trade liberalization, Harrison et al. (2003, p. 97) note that "even the most attractive reforms will typically result in some households losing," and recent efforts to measure welfare impact at the household level have been made following Porto (2006). Nevertheless, because results are presented at some aggregate level (e.g., by state or percentile), impoverishment due to trade reform could still be overlooked. For example, Nicita's (2009, p. 26) finding that "on average all income groups benefited from [Mexico's] trade liberalization, but to a varying extent" does not tell us the extent to which some households within each group were made worse off by the reform.

In each of these cases, our axiomatically derived FI measure could be used to quantify the impoverishment of those becoming poorer over time or the extent to which losers are hurt by policy reforms. Our decomposition could be used to examine the extent to which a decrease in poverty over time or due to a reform balances out the gains and losses of different households. Doing so, we will cease to overlook cases where growth, policy reform, or the tax and transfer system is poverty-reducing and progressive, yet hurts a substantial proportion of the poor.

Chapter 2

Towards an Understanding of the Political Economy of Multidimensional Poverty Measurement

Abstract

Does adopting a multidimensional poverty index lead to higher government spending on the poor? If so, why? And how does it affect the allocation of government budgets across ministers? We answer these questions in a game theoretic framework that accounts for the strategic interactions between government agents. Government ministers—such as the education, health, and housing ministers—share a common interest in reducing measured poverty; adopting a multidimensional measure may induce them to spend more on the poor since they can now directly impact measured poverty, whereas they have little to no short-run impact on, say, a unidimensional income poverty measure. Because an improvement in the scalar multidimensional poverty index is a public good for ministers, however, they can also free ride on each other's antipoverty spending. Despite introducing free riding, a multidimensional measure usually leads to an increase in total antipoverty spending. In addition to incentivizing ministers to spend resources to reduce measured poverty and thereby increase prestige for all government members, the multidimensional measure creates a new set of policy tools that serve as levers affecting total spending on the poor: dimension

weights in the index and resource allocations across ministers. In the use of these tools, a conflict arises between maximizing reductions in measured poverty and maximizing equilibrium antipoverty spending; its resolution depends on whether the authority deciding weights and the authority allocating budgets have the same or opposing incentives. We illustrate using data from Mexico, the first country to adopt an official multidimensional poverty measure.

2.1 Introduction

Poverty is a multidimensional phenomenon. Although "inadequate income is a strong predisposing condition for an impoverished life," Sen (1999, p. 87) asserts that "poverty must be seen as the deprivation of basic capabilities rather than merely as lowness of incomes." The recent debate about whether multiple dimensions of poverty can be credibly aggregated into a single scalar measure has revived scholarly interest in multidimensional poverty measures.[1] Meanwhile, scalar indices that capture multiple dimensions of poverty in a single axiomatically consistent measure are gaining significant policy traction. In particular, the measure proposed by Alkire and Foster (2011a, henceforth AF), which combines information about joint deprivations in multiple dimensions and is decomposable by subgroup and dimension, is becoming part of many countries' official poverty measurement and antipoverty program targeting methods.[2]

While the implications of using the AF measure to target antipoverty programs have been studied (Alkire and Seth, 2013; Azevedo and Robles, 2013; Duclos et al., 2014), the

[1] On one hand, Ravallion (2011, p. 235) argues that "we should aim for a credible set of multiple indices rather than a single multidimensional index"; on the other, Deaton (2011, p. 14) argues that multidimensional measures are "required," and that they "need to be calculated from surveys that collect multiple measures for each respondent" due to the correlation of deprivations across dimensions. Others have argued for intermediate approaches that analyze deprivations in particular dimensions while still capturing interdependence across dimensions (e.g., Ferreira et al., 2013).

[2] In 2009, Mexico replaced its official (unidimensional) poverty measure with a multidimensional one, based on AF. Colombia adopted the AF class of multidimensional measures in its official poverty measurement in 2011, and is also using it to design and target its social programs. Bhutan, Ecuador, El Salvador, Costa Rica, Chile, Malaysia, and the Philippines have also implemented an official multidimensional poverty measure based on AF.

political economy implications of using it to measure multi- rather than unidimensional poverty have not. Does using a multidimensional poverty index such as AF lead to higher government spending on the poor than using a unidimensional poverty index? If so, why? How does the allocation of government budgets across ministries change when a multidimensional measure is adopted? And how should the parameters of the multidimensional measure be chosen? We answer these questions in a game theoretic framework that accounts for the strategic interactions between government agents.

In the public choice tradition, we consider government members that do not simply maximize the welfare of the poor, but divide their resources between two objectives. On the one hand, each minister fights poverty in order to obtain "prestige," which consists of two components: first, prestige is bestowed upon all government members as a function of the reduction in measured poverty, and second, each minister is independently evaluated on reductions in deprivations in that minister's dimension of competence. On the other hand, each minister derives utility from private consumption, which can be thought of as corruption, patronage, spending on the non-poor, or any other form of spending that does not benefit the other ministers.[3] Further, we assume that "line ministers"—such as the education, health, housing, and social development ministers—can spend resources to reduce deprivations in a particular dimension of the multidimensional poverty measure (e.g., the health minister can only act to reduce deprivations in health). Therefore, since a fall in measured multidimensional poverty reflects positively on all line ministers, poverty reduction is a public good among them and free riding can arise, as in Bergstrom et al. (1986). Even when ministers are evaluated not only on the multidimensional index but on this index and its breakdown by dimension, the partial evaluation of all government ministers using the scalar multidimensional index leads to free riding.[4]

[3] Several papers document instances of corruption reducing the resources that reach the poor from, for example, food programs in Indonesia (Olken, 2012), education services in Brazil (Ferraz et al., 2012), and "cash for work" employment guarantee schemes in India (Niehaus and Sukhtankar, 2013).

[4] Even if ministers are evaluated solely on their dimension's contribution to the AF measure, free riding can occur since the dimensional breakdown uses the censored deprivation matrix, each column of which depends on the contributions of other ministers for the identification of the poor (see Alkire and Foster,

Our game begins with a technical committee deciding whether to adopt a multidimensional poverty measure and selecting the weights assigned to each dimension of the measure, as well as the deprivation cut-offs that determine whether an individual is deprived in a particular dimension.[5] Next, taking the measure's weights and deprivation cut-offs as given, a finance minister allocates a fixed amount of resources among line ministers, within certain constitutional or political constraints. Both of these agents may be interested in making the government *look* as good as possible by maximizing measured poverty reduction (i.e., prestige), or they may be benevolent, a term that we use as shorthand for "maximizing total antipoverty spending by the line ministers." Then, taking as given budgets and dimension weights, two line ministers[6] independently maximize a utility function that is increasing in private consumption, deprivation reduction in the minister's specific dimension, and overall poverty reduction—as measured by the official poverty measure, which may be uni- or multidimensional. To summarize our main result, the use of a multidimensional index increases resources devoted to the poor compared to a unidimensional index (assuming optimal weight choices and under some technical conditions); antipoverty expenditures can be further increased by publishing partial dimensional indices alongside the scalar multidimensional one.

We make three main contributions to the literature. First, as described above, we begin the exploration of the political economy consequences of measuring poverty using a multidimensional scalar index. Doing so, we extend the literature on optimal antipoverty spending by government agents who seek to maximize *measured* poverty reduction (Bour-

2011a, p. 482). It appears that this occurred in practice in India (Alkire and Seth, 2015). Furthermore, even if scalar measures are avoided altogether and multidimensional stochastic dominance tests are used instead (Duclos et al., 2006, 2011; Maasoumi and Racine, forthcoming), complementarities across dimensions, and thus externalities across ministers and the opportunity for free riding, would still exist.

[5] The decision of whether to adopt a multidimensional measure can be thought of theoretically as an elementary aspect of the choice of weights, since a unidimensional measure is equivalent to assigning all of the weight to one dimension (Sen, 1987). In the case of Mexico, weights and deprivation cut-offs were chosen by the National Council for the Evaluation of Social Development Policy (CONEVAL) in consultation with a team of experts (e.g., Foster, 2007), while in Chile the government created a Poverty Measurement Commission in 2012.

[6] We present results of extending the model to three dimensions and three line ministers in the Appendix; further generalizations to n dimensions and line ministers are immediate.

guignon and Fields, 1997) to poverty measures with more than one dimension, multiple decision makers, and competing objectives. Second, our findings contribute to the ongoing debate about which axioms are desirable for a multidimensional poverty measure. Alkire and Foster (2016) show that no measure can simultaneously satisfy two popular axioms from the literature: *poverty non-decreasing rearrangement* (Tsui, 2002)—which Duclos and Tiberti (forthcoming) proclaim as "arguably the most important property" for multidimensional poverty measurement—and *factor decomposability* (Chakravarty et al., 1998)—which Alkire and Foster (2016) claim is "largely responsible" for the growing use of multidimensional measures.[7] By decomposing the multidimensional measure into partial indices as well as the overall measure, transparency is enhanced and free riding is reduced; we show that this leads to an unambiguous increase in antipoverty spending and prestige. Thus, our study illustrates an underexplored benefit of the factor decomposability axiom and lends theoretical support for Alkire et al.'s (2011) recommendation to not rely solely on the scalar multidimensional index, but also publish partial indices. Third, we propose a novel framework for determining the optimal weights of a multidimensional poverty measure, seeking to maximize spending on the poor by government ministers who free ride on each other's efforts to reduce poverty; this contrasts with the arbitrary selection of weights often employed in practice and complements other advocated approaches, such as expert opinion (Chowdhury and Squire, 2006), participatory processes (Bossert et al., 2013; Decancq et al., 2013), and hybrid data-driven/normative approaches (Cavapozzi et al., 2015; Decancq et al., 2015; Maasoumi and Xu, 2015).

While we make several simplifying assumptions (described in the next section), many interesting results emerge from the model. Although using a multidimensional index cre-

[7] More precisely, Alkire and Foster (2016) show that no measure can satisfy related axioms they call *dimensional transfer* and *dimensional breakdown*, modified to allow for ordinal variables and the Alkire and Foster (2011a) dual cut-off approach to identify the poor. Variants of the *poverty non-decreasing rearrangement* axiom, first discussed in the context of multidimensional welfare by Atkinson and Bourguignon (1982), have been called *weak rearrangement* (Alkire et al., 2015, Ch. 2), *non-decreasing poverty under a correlation-increasing switch* (Bourguignon and Chakravarty, 2003), and *sensitivity to multiple deprivation* (Duclos and Tiberti, forthcoming), while variants of factor decomposability have been called *additive decomposability in attributes* (Bossert et al., 2013) and *attribute additivity* (Duclos and Tiberti, forthcoming).

ates a free-rider problem, its net effect on overall antipoverty spending is often positive. The reason is twofold: first, compared to a unidimensional measure, more line ministers are incentivized to spend on poverty reduction (in excess of the amount they spend when their dimensions are not included in the unidimensional measure) because their actions now have a direct impact on measured poverty. Second, additional policy tools (e.g., the dimension weights in the AF measure) become available to the government to encourage line ministers to increase their antipoverty spending. The main idea is to use these tools to make the pursuit of prestige *hard* so that, if there is at least some degree of complementarity between prestige and private consumption, more antipoverty expenditures result in equilibrium.

The results from the model are usefully classified according to whether only one line minister exerts antipoverty efforts beyond the amount that minister would spend if her dimension were not included in the official poverty measure, or both line ministers do. When only one line minister exerts antipoverty efforts beyond this amount, as when the poverty measure is unidimensional, the budget allocations chosen by a prestige-maximizing and a benevolent finance minister coincide. In contrast, when both line ministers exert antipoverty efforts beyond this amount, the outcome depends on whether the finance minister is prestige-maximizing or benevolent. A prestige-maximizing finance minister assigns a larger budget to the minister whose dimension is responsible for a larger fraction of initial poverty and for whom improvements require smaller investments, thereby reaping "low-hanging fruit." A benevolent finance minister does the opposite. Intuitively, the benevolent finance minister makes it *hard* for the line ministers to obtain prestige (in the form of higher *measured* poverty reduction) by giving more resources to the minister whose dimension accounts for a smaller fraction of initial poverty or involves a higher marginal cost of reducing deprivations; because measured poverty reduction is harder to achieve and prestige and private consumption are not perfect substitutes, line ministers end up contributing more resources to the task.

For the committee, the distinction between prestige-maximization and benevolence is again important: a self-interested agent will always prefer to assign zero weight to one of the two dimensions. In other words, governments interested in “looking good” may have an incentive to continue measuring poverty as unidimensional, even though a multidimensional measure could increase the resources spent on the poor. Vietnam, for example, still measures poverty unidimensionally; if a multidimensional measure were used instead, the government would appear to be reducing poverty by less (Tran et al., 2015). A benevolent committee selecting weights not only chooses a multidimensional measure by assigning a large enough positive weight to each dimension that both line ministers contribute, but, if the finance minister is also benevolent, assigns weights to make achieving prestige hard, for instance by assigning higher weight to the dimension of competence of a minister with fewer resources or a higher cost of reducing deprivations. If instead a benevolent committee is facing a prestige-maximizing finance minister, then the committee chooses weights so that the finance minister cannot “game the system” by shifting resources around; in our model, if improvements in each dimension are equally costly, this is implemented by equal relative weight in initial poverty for each dimension. Further, if deprivations are equal across dimensions, this translates to assigning equal absolute weights to each dimension, as is often done in practice. If deprivations differ significantly across dimensions, however, assigning equal absolute weights is not optimal, and higher antipoverty spending can be achieved by adjusting the dimensions’ absolute weights to equalize their relative importance.

Our model’s predictions played out in Mexico after it implemented a multidimensional measure: between 2008 and 2010 (the official poverty measure changed to multidimensional in 2009), the total proportion of transfers and subsidies reaching the poor—including in-kind transfers in the form of public education and health services—increased from 23.2% in 2008 to 26.2% in 2010, despite there being less poor people in 2010 (Scott, 2014).[8] Furthermore, our model predicts that moving from a unidimensional measure to a multidimensional one would cause a reduction in spending allocated to the poor by the minister in charge of the dimension included in the unidimensional measure (now that she can free ride) and an increase from the ministers in charge of the dimensions that have been newly added to the measure. Indeed, Scott (2014) finds that the proportion of cash transfers reaching the poor from Mexico's two main antipoverty programs—the Oportunidades conditional cash transfer program and a non-contributory pension for the elderly poor—decreased, while the proportion of education and health spending reaching the poor increased.

The rest of the paper is organized as follows. The next section presents the model and discusses our assumptions. Section 2.3 derives the game's equilibrium and our theoretical results. Section 2.4 provides an illustration through a detailed example using Mexican data. Section 2.5 concludes.

2.2 The Model

2.2.1 Preferences and Strategic Interaction

We consider a simple setup with two line ministers, each in charge of one of the two dimensions relevant for the measurement of poverty.[9] Minister j's utility function ($j = 1, 2$) has the Cobb-Douglas formulation $U_j\left(P_j^T, s_j\right) = \left(P_j^T\right)^{\gamma_j} \cdot \left(s_j\right)^{\delta_j}$, where P_j^T is a prestige term that depends on poverty reduction, s_j is private consumption, and $\gamma_j > 0$ and $\delta_j > 0$ are parameters affecting the marginal rate of substitution between private consumption and prestige.[10] Each minister j divides her total financial resources r_j between private

[8] Of course, the agreement between our model's predictions and changes in Mexico after implementing a multidimensional measure are merely suggestive and cannot be interpreted causally.

[9] The restriction to two dimensions is for ease of exposition only; the basic economic forces of the model are unchanged by including more than two dimensions. To show this, we derive the interior solution for three dimensions in the Appendix; generalizations to n dimensions are immediate.

[10] No qualitative changes arise in our results by considering a more general constant elasticity of substitution (CES) utility function, as long as prestige and private consumption are not perfect substitutes.

consumption—which can be thought of as corruption, patronage, or any type of spending that does not reach the poor—and investment in a poverty reduction program that reduces deprivations in dimension j. We denote with x_j the financial resources that minister j devotes to poverty reduction and, for simplicity, we conduct our analysis with multiplicative policies: spending of x_j reduces the proportion of people deprived in dimension j by $100x_j/c_j$ percent, where $1/c_j$ is the effectiveness of one dollar spent in dimension j.

We assume P_j^T is derived from two sources. First, line minister j benefits from a decrease in the proportion of people deprived in dimension j. We assume this is a function of x_j only and denote it $Q_j(x_j)$. Second, both ministers benefit from a decrease in the multidimensional poverty index. This depends on antipoverty spending by each minister, so we denote it $P(x_1, x_2)$. The strategic interaction between line ministers is modeled as a simultaneous public good provision game in which $P(x_1, x_2)$ is a public good enjoyed by the ministers, since both line ministers benefit when one spends resources on poverty reduction, thereby increasing $P(x_1, x_2)$.[11] Therefore, taking as given x_2, minister 1 maximizes for $x_1 \in [0, \min\{r_1, c_1\}]$ the function

$$U_1(x_1, x_2) = \left(P\left(x_1, x_2\right) + \theta_1 Q_1(x_1)\right)^{\gamma_1} \cdot \left(r_1 - x_1\right)^{\delta_1}, \tag{2.1}$$

where θ_1 is the non-negative weight minister 1 assigns to the decrease in dimension-1 deprivations. Minister 2 solves a similar problem taking x_1 as given.

The other two decision makers in our game are a finance minister and a technical committee. We consider two possible objectives for these agents to maximize: prestige, i.e. they are concerned with making the government look as good as possible, or "benevolence." Truly benevolent agents would like to maximize the welfare of the poor. However, in the absence of a specific, universally-accepted social welfare function and an accurate

[11] This description favors simplicity over realism but it is sufficient to highlight the main forces we are interested in. Clearly, refining the description of how government policy is formed is important; however, as long as multiple agents are responsible for policy and these agents cannot write complete and enforceable binding contracts with one another, then a public good problem among policymakers can arise.

measure of utility for each poor individual, it is reasonable to assume that in practice they seek to maximize total government spending aimed at reducing poverty.[12] Indeed, policymakers directly cite the amount spent on the poor as a policy objective, indicating its importance as a yardstick;[13] furthermore, spending on the poor has been recognized as an important variable in empirical studies (e.g., Alderman, 2002; Inchauste et al., forthcoming; Paz Arauco et al., 2014) and theoretical work in public economics (e.g., Bardhan and Mookherjee, 2005; Galasso and Ravallion, 2005; Ravallion, 1999).

2.2.2 Poverty Measurement and Production Technology

The way poverty is measured enters into the production function for prestige. To make our main point in the simplest possible framework, we assume poverty is measured with a simplified version of an AF index over two dimensions. Consider what Atkinson (2003) terms the union approach, that is, if an agent is deprived in any one dimension, then that agent is poor. Following the notation and terminology in Alkire and Foster (2011a), let y_{ij} indicate individual i's achievement in dimension j (for $i = 1, \ldots, n$ and $j = 1, 2$), and let z_j be a deprivation cut-off for dimension j such that individuals with $y_{ij} < z_j$ are considered deprived in dimension j. The normalized deprivation or gap of individual i in dimension j is $g^1_{ij} = (z_j - y_{ij})/z_j$ if $y_{ij} < z_j$ and 0 otherwise. Consider the matrix of normalized gaps g^1 with typical element g^1_{ij}, in which columns represent dimensions relevant for poverty and rows represent individuals.[14]

Now, denote with g^1_{*j} the jth column of g^1 and let μ be the mean operator. Also, for $\alpha \geq 1$, let g^α_{*j} be the column vector derived from g^1_{*j} by raising to the power α every element

[12] The proxy for welfare of the poor cannot, of course, be the multidimensional poverty index, since this index is a function of the dimension weights and deprivation cut-offs, which are choice variables for the committee.

[13] For example, the Mexican government's National Development Plan argues that the amount of government spending allocated to the poor is too low (Gobierno de la República, 2013, p. 50).

[14] To extend this framework to the intersection approach where an individual is poor only if she is deprived in both dimensions, or Alkire and Foster's (2011a) intermediate approach for more than two dimensions where an individual is poor if the weighted sum of the dimensions in which she is deprived is greater than some poverty cut-off k, let g^1 be the *censored* matrix of deprivations, with $g^1_{ij} = 0$ for all j if i is non-poor.

of g^1_{*j}; for $\alpha = 0$, g^0_{*j} is the jth column of g^0 whose typical element $g^0_{ij} = \mathbb{I}\big(g^1_{ij} > 0\big)$, where $\mathbb{I}$ is the indicator function taking a value of 1 if its argument is true and 0 otherwise. Using the dimension-specific decomposition, and assigning weights w_1 and $w_2 = 1 - w_1$ to each dimension, the Alkire–Foster multidimensional poverty index M_α is calculated as

$$M_\alpha = \left(\frac{w_1}{n}\sum_{i=1}^{n} g^\alpha_{i1} + \frac{w_2}{n}\sum_{i=1}^{n} g^\alpha_{i2}\right) = w_1\mu(g^\alpha_{*1}) + w_2\mu(g^\alpha_{*2})\,. \tag{2.2}$$

In this paper, we focus on the $\alpha = 0$ case because M_0 is fully compatible with categorical variables. This is an important feature: “with the exception of income, most available deprivation indicators are categorical” (Dewilde, 2004, p. 339); as a result, the governments that have implemented official multidimensional measures have all opted for M_0, with the exception of Colombia where results are published for M_0, M_1, and M_2.[15] Recalling that g^0 is a matrix of zeros and ones, where $g^0_{ij} = 1$ if individual i is deprived in dimension j and 0 otherwise, $\mu\big(g^0_{*j}\big)$ gives the proportion of the population deprived in dimension j. For simplicity, we assume that the poverty reduction technology is linear, and that the constant marginal cost of reducing the proportion of the population deprived in dimension j is c_j.[16] The cost is meant to capture various differences across dimensions that influence how antipoverty spending is mapped into a reduction in the number of individuals deprived in that dimension. Achievement shortfalls might differ across dimensions, such that *ceteris paribus* it would be more costly to reduce deprivations in the dimension in which people are more deeply deprived (i.e., have a larger shortfall between their achievements and the deprivation cut-off). Or, conditional on the distribution of deprivations, in some dimensions it might be costlier to reduce the proportion deprived: for example, it might be cheaper to

[15] In a companion paper, we analyze the $\alpha \geq 1$ case, focusing on $\alpha = 1$ and $\alpha = 2$.

[16] Because the measure is sensitive to the depth of deprivations when $\alpha \geq 1$, in our companion paper we reinterpret the productivity of antipoverty expenditures: instead of reducing the proportion of the population that is deprived in dimension j by x_j, spending of x_j proportionally reduces each poor individual’s deprivation in j (i.e., g^1_{ij}) by x_j. For $\alpha = 1$, results are identical to those presented here; for $\alpha = 2$, the kind of explicit solutions we derive in the next section are usually impossible to achieve and one must make additional simplifications to obtain interesting analytical results; however, these results are broadly consistent with what we derive in the next section.

provide a solid floor than to ensure that children do not drop out of school before grade 8—both of which are dimensions of the Global Multidimensional Poverty Index (Alkire and Santos, 2014).

After minister j spends x_j on poverty reduction, the proportion of the population deprived in dimension j will be $\mu\big(g^0_{*j}\big)\,(1 - x_j/c_j)$, so

$$M_0' = w_1\mu\big(g^0_{*1}\big)\left(1 - \frac{x_1}{c_1}\right) + w_2\mu\big(g^0_{*2}\big)\left(1 - \frac{x_2}{c_2}\right). \tag{2.3}$$

Government agents earn prestige when the level of multidimensional poverty, as measured by the scalar index (2.2), falls. Specifically, prestige is awarded based on the percentage fall in multidimensional poverty:

$$P\left(x_1, x_2\right) \equiv \frac{M_0 - M_0'}{M_0} = \frac{1}{M_0}\left[w_1\mu\big(g^0_{*1}\big)\,\frac{x_1}{c_1} + w_2\mu\big(g^0_{*2}\big)\,\frac{x_2}{c_2}\right] = \pi_1\frac{x_1}{c_1} + \pi_2\frac{x_2}{c_2}, \tag{2.4}$$

where $\pi_j \equiv w_j\mu\big(g^0_{*j}\big)/M_0$ is the fraction of initial multidimensional poverty explained by dimension j. Nothing would change in our analysis if prestige were based on the absolute fall in multidimensional poverty, $M_0 - M_0'$, as the marginal rate of substitution between prestige and private consumption would remain the same. We refer to π_j as a *relative weight* throughout the paper; it is a function of the *absolute weight* w_j and initial deprivations, which are in turn a function of individuals' achievements and the chosen deprivation cut-offs in each dimension. In practice, technical committees often choose both the absolute weights and deprivation cut-offs.[17]

Similarly, we assume Q_j reflects the percentage fall in dimension-j poverty, so $Q_j(x_j) = x_j/c_j$, for $j = 1, 2$.

[17] The choice of deprivation cut-offs will also affect who is considered poor (the *identification* step of poverty measurement), which will likely effect the effectiveness $1/c_j$ of each dollar spent on reducing deprivations in dimension j; we abstract from the latter detail for simplicity, treating c_j as given.

2.2.3 Timing of the Game

Our game begins with a technical committee selecting the weights assigned to each dimension of the multidimensional measure. Then, taking as given the weights, a finance minister allocates resources among the two line ministers.[18] Both the committee and the finance minister may be interested either in maximizing measured poverty reduction or maximizing total antipoverty spending in equilibrium;[19] the game is solved under all four possible combinations of objectives. We always maintain the assumption that weights are chosen before resources are allocated, since, in practice, the choice of weights is a long-term decision, while government budgets are frequently altered. We also consider the case in which budget allocations are fixed, as the finance minister may not have political leeway to reallocate resources among line ministers; in this case, the technical committee knows the fixed values of r_1 and r_2 and incorporates them into its choice of weights. Finally, taking as given dimension weights and budgets, two line ministers choose their antipoverty spending independently and simultaneously.

When the finance minister makes her decision, the equilibrium consequences on antipoverty expenditures chosen by the line ministers are fully anticipated. Similarly, when the committee decides weights, equilibrium consequences on resource allocations and line ministers' antipoverty expenditures are fully taken into account.

2.3 Equilibrium

We now characterize equilibrium; since at each step agents know the outcome of previous play, the appropriate notion is subgame perfect equilibrium. Using the logic of backwards

[18] Alternatively, one may think of budget allocations as the outcome of bargaining among ministries, affected by the bargaining protocol, outside options, and bargaining strengths. While this would be a different setup, free riding still arises as long as ministers cannot commit to expenditures while bargaining; indeed, some of our results would go through unchanged, such as the direct conflict between maximizing prestige and maximizing resources spent on the poor.

[19] In Section 2.4, we also explore the outcome of the game between line ministers when weights and/or budgets are chosen "non-optimally," in the sense that they do not maximize prestige or total antipoverty spending in equilibrium.

induction, we begin with the simultaneous poverty reduction game between line ministers. Then, we move to the choice of the finance minister, and finally to that of the committee. Standard results show that equilibrium is generally unique (unless the finance minister is indifferent across resource allocations).

2.3.1 The Game between Line Ministers

Because ministers are unlikely, in practice, to be able to completely eliminate deprivations, for simplicity only we consider $r_1 < c_1$ and $r_2 < c_2$, avoiding corner solutions where a line minister completely eliminates deprivations in her dimension of competence. The other possible corner solution, where one line minister contributes nothing to reducing poverty, is however fully discussed in what follows. Minister 1 maximizes the utility in (2.1) by choosing x_1 for given x_2, and minister 2 behaves similarly. Standard first-order conditions (FOCs) imply that

$$x_j = \max\left\{ \frac{\frac{(\pi_j+\theta_j)\gamma_j}{c_j\delta_j} r_j - P}{\frac{(\pi_j+\theta_j)\gamma_j}{c_j\delta_j} + \frac{\theta_j}{c_j}}\ ,\ 0 \right\},\ \text{for } j = 1, 2. \tag{2.5}$$

We now define $\beta_j \equiv \delta_j/\gamma_j$ as the parameter that governs the marginal rate of substitution between private consumption and prestige. In what follows up to Section 3.5, we maintain θ_j fixed for $j = 1, 2$; for ease of exposition then, we now set $\theta_j = 0$ and equation (2.5) simplifies to $x_j = \max\{r_j - P\beta_j c_j/\pi_j, 0\}$.[20] This simplification does not change the qualitative results of the model unless θ_j is large or not at all similar between the two line ministers. By normalizing $\theta_j = 0$, we are in essence normalizing to 0 the amount a minister contributes to reducing deprivations when her dimension is not included in the unidimensional poverty measure—which would still be positive for $\theta_j > 0$ since her utility function is increasing in her dimension-specific deprivation reduction $Q_j(x_j)$. Thus, when we discuss corner solutions where one minister contributes 0 to antipoverty spending, we

[20] We will return to the full formulation in (2.5) when discussing the relation to the dashboard approach in Section 3.5.

can think of 0 as a normalization of the amount the minister would spend to reduce deprivations in her own dimension when that dimension is not included in the country's official poverty measure; we refer to this spending level as "the minimum" throughout the paper.

Simple algebra then yields the following proposition, stated without proof.

Proposition 5 (Equilibrium of the game between line ministers)**.** *Equilibrium antipoverty expenditures and prestige are as follows:*

1. *If*
$$\frac{\pi_2 r_2}{c_2} < \frac{\beta_2}{1+\beta_1}\frac{\pi_1 r_1}{c_1}, \tag{2.6}$$
then the equilibrium outcome has $x_2^* = 0$,
$$x_1^* = \frac{P_{k_1}^* c_1}{\pi_1} = \frac{r_1}{1+\beta_1}, \tag{2.7}$$
and prestige equals
$$P_{k_1}^* = \frac{\pi_1 r_1}{c_1(1+\beta_1)}, \tag{2.8}$$
where the k_1 *subscript indicates a corner solution in which only minister 1 contributes beyond the minimum.*

2. *If*
$$\frac{\pi_1 r_1}{c_1} < \frac{\beta_1}{1+\beta_2}\frac{\pi_2 r_2}{c_2}, \tag{2.9}$$
then the equilibrium outcome has $x_1^* = 0, x_2^* = \frac{r_2}{1+\beta_2}$, *and* $P_{k_2}^* = \frac{\pi_2 r_2}{c_2(1+\beta_2)}$,

3. *If both* (2.6) *and* (2.9) *are violated, then both line ministers contribute beyond the minimum in equilibrium, with*
$$x_1^* = r_1 - \frac{\beta_1 c_1}{\pi_1}\frac{\frac{\pi_1 r_1}{c_1} + \frac{\pi_2 r_2}{c_2}}{1+\beta_1+\beta_2}, \qquad x_2^* = r_2 - \frac{\beta_2 c_2}{\pi_2}\frac{\frac{\pi_1 r_1}{c_1} + \frac{\pi_2 r_2}{c_2}}{1+\beta_1+\beta_2}, \tag{2.10}$$

and equilibrium prestige is

$$P^* = \frac{\frac{\pi_1 r_1}{c_1} + \frac{\pi_2 r_2}{c_2}}{1 + \beta_1 + \beta_2}. \tag{2.11}$$

If both (2.6) and (2.9) are violated, using (2.10) and substituting the value of equilibrium prestige in (2.11), we see that total antipoverty expenditures equal

$$x_1^* + x_2^* = r_1 + r_2 - \left(\frac{\beta_1 c_1}{\pi_1} + \frac{\beta_2 c_2}{\pi_2}\right) P^*. \tag{2.12}$$

2.3.2 Resource Decision

Consider now the problem of a decision maker (e.g., the finance minister) who is allocating a total amount of resources R to r_1 and r_2, under the constraint that some minimum funding level $\underline{r}_j$ be provided to line minister j, $j = 1, 2$. For instance, these lower bounds may be due to constitutional mandates or they may be the result of bargaining between government coalition partners. Let $\bar{r}_j$ be the maximum amount a line minister can receive, determined by the total resources available to the finance minister and the minimum the other line minister can receive, i.e., $\bar{r}_1 = R - \underline{r}_2$ and $\bar{r}_2 = R - \underline{r}_1$,. To avoid corner solutions where some antipoverty contribution fully eliminates deprivations in that minister's dimension, assume $\bar{r}_1 < c_1$ and $\bar{r}_2 < c_2$. We consider two possible objectives for the finance minister: maximize measured poverty reduction or maximize total antipoverty expenditures. In either case, the piecewise linearity and continuity of the objectives implies that at the optimal solution, one of the line ministers is allocated only the minimum amount of resources, while the remaining resources are allocated to the other minister.

Before proceeding further with the determination of how the parameters of the antipoverty measure affect resource distributions between line ministers, it is opportune to restate both boundary conditions (2.6) and (2.9) in terms of r_1 using the fixed quantity $R = r_1 + r_2$: if

$$r_1 > r_1^h \equiv \frac{\frac{\pi_2}{c_2} R\,(1 + \beta_1)}{\frac{\pi_1}{c_1}\beta_2 + \frac{\pi_2}{c_2}(1 + \beta_1)},$$

then line minister 2 does not contribute beyond the minimum in equilibrium, while if

$$r_1 < r_1^l \equiv \frac{\beta_1 \frac{\pi_2}{c_2} R}{(1+\beta_2)\frac{\pi_1}{c_1} + \beta_1 \frac{\pi_2}{c_2}},$$

then line minister 1 does not contribute beyond the minimum in equilibrium. (Simple difference shows that $r_1^h > r_1^l$.)

First, note than for equilibria in which a line minister does not provide antipoverty efforts, the incentives of a finance minister that maximizes measured poverty reduction are the same as those of a finance minister that maximizes total antipoverty expenditures, as described in the following remark, immediate from equations (2.7) and (2.8).

Remark 2.3.1 (No conflict for corner solutions). Consider the range $r_1 > r_1^h$ in which line minister 2's equilibrium expenditure is zero. A redistribution of resources towards line minister 1 improves both prestige and antipoverty expenditures. Therefore, if $\bar{r}_1 > r_1^h$ so that $[r_1^h, \bar{r}_1]$ is nonempty, then a prestige-maximizing or benevolent finance minister's best choice among all $r_1 \in [r_1^h, \bar{r}_1]$ is $\bar{r}_1$.

When one line minister does not contribute to poverty reduction beyond the minimum, the antipoverty measure is "behaviorally unidimensional," in the sense that any changes in poverty are completely driven by changes in the behavior of the only line minister that adjusts its expenditure in response to marginal changes in resources. The previous remark (which of course holds *mutatis mutandis* when line minister 1 does not contribute) illustrates the common-sense implication that the availability of more resources to the only ministers who make more than minimal efforts to reduce poverty tends to increase both measured poverty reduction and actual antipoverty expenditures. Things are not as straightforward when parameter values are such that both line ministers' choices are strictly interior, i.e., above the minimum. As the next proposition shows, there arises a conflict between the maximization of antipoverty expenditures and of measured poverty reduction.

Proposition 6 (Maximization of antipoverty expenditures vs. that of measured poverty reduction). *Suppose parameter values are such that line ministers' expenditure choices are strictly positive in equilibrium*: $\bar{r}_1 \leq r_1^h$ *and* $\underline{r}_1 \geq r_1^l$. *If* $\frac{\pi_1}{c_1} > \frac{\pi_2}{c_2}$, *then* $r_1 = \bar{r}_1$, $r_2 = \underline{r}_2$ *maximizes measured poverty reduction and minimizes total antipoverty expenditures. In contrast,* $r_1 = \underline{r}_1$, $r_2 = \bar{r}_2$ *maximizes total antipoverty expenditures and minimizes measured poverty reduction. The implication is reversed for* $\frac{\pi_1}{c_1} < \frac{\pi_2}{c_2}$. *And if* $\frac{\pi_1}{c_1} = \frac{\pi_2}{c_2}$, *i.e., for* $\pi_1 = \frac{c_1}{c_1+c_2}$, *then any possible distribution of resources generates the same measured poverty reduction and antipoverty expenditures.*

Proof. Follows from the linearity of (2.11) and (2.12), and the fact that the two objectives are directly in contrast, as (2.12) clearly shows. □

First, it is worth pointing out a contrast with Bergstrom et al. (1986) in which marginal redistribution has no effects. Here it does, because the expenditure of line ministers 1 and 2 have different marginal productivity on P^*, as in Cornes and Hartley (2007), for instance. Second, note that for interior solutions $\partial P^*/\partial x_j = \pi_j/c_j$; it is common sense that achieving prestige becomes easier after transferring resources to the line minister whose expenditures have the largest marginal impact on measured poverty reduction (e.g., to minister 1 when $\pi_1/c_1 > \pi_2/c_2$, and hence $\partial P^*/\partial x_1 > \partial P^*/\partial x_2$). Third, note that redistributing resources to the line minister with less impact on the poverty measure (e.g., to minister 2 when $\pi_1/c_1 > \pi_2/c_2$) makes achieving prestige harder, but, because of the complementarity between private consumption and prestige, this induces higher total antipoverty spending.

If total antipoverty spending is the policy objective, then the policy prescription is to redistribute resources to the minister whose dimension is weighted less (lower π_j) and for which improvements are harder to achieve (higher c_j). More generally, Proposition 6 illustrates that a larger set of policy tools becomes available when poverty is measured multidimensionally.

When the conditions of Proposition 6 are not met, it is not possible to pin down whether the solution is $(\bar{r}_1, \underline{r}_2)$ or $(\underline{r}_1, \bar{r}_2)$ *a priori*. However, the logic behind Remark 2.3.1 and Proposition 6 can be used to draw a picture of the objective function of the finance minister.

We do so in Figure 2.1, where the objective of the finance minister is prestige. We see that, if $\underline{r}_1 = r_1^A$, then the solution to a prestige-maximizing finance minister's decision problem is $(\underline{r}_1, \bar{r}_2)$ and only one line minister contributes beyond the minimum in equilibrium. But if $\underline{r}_1 = r_1^B$, then the solution is $(\bar{r}_1, \underline{r}_2)$ because prestige is higher at $\bar{r}_1$ than at $\underline{r}_1 = r_1^B$, and hence both line ministers exert antipoverty efforts beyond what they would spend to achieve dimension-specific deprivation reductions.

Figure 2.1: Prestige as a Function of r_1, $\pi_1/c_1 > \pi_2/c_2$ Case

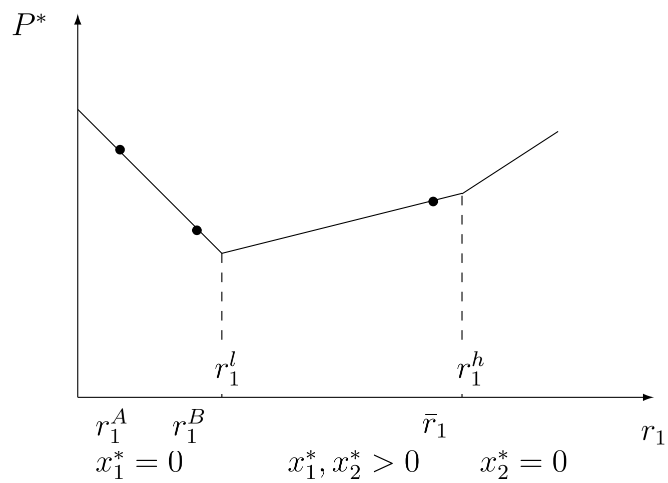

2.3.3 The Choice of Weights

The following proposition describes the choice of weights of a committee that maximizes prestige, and shows that prestige is larger for a unidimensional poverty measure than for a multidimensional one.

Proposition 7 (Prestige is maximized by the choice of a unidimensional index)**.** *The value of prestige obtained when* $\pi_1 \in (0, 1)$ *is strictly less than that obtained for* $\pi_1 = 0$ *or* $\pi_1 = 1$.

Proof. We first consider the case where the finance minister has no control over budgets because budgets may already fixed for a number of years, giving the finance minister little to no leeway to reallocate resources across line ministers. Then we analyze the full model in which the finance minister optimally responds to the committee's choice of weights. Assuming that resources are unaffected by the choice of the committee (e.g., because the finance minister does not have political leeway to change the budgets assigned to each ministry), note that $P^*_{k_1}$ in (2.8) is strictly increasing in r_1 and π_1. Therefore, for all equilibria where line minister 2 does not provide antipoverty expenditures, the largest possible prestige obtains for $\pi_1 = 1$, at a level equal to $r_1/(c_1(1+\beta_1))$. This is exactly what occurs with a unidimensional poverty measure that focuses all weight on dimension 1. With a similar reasoning, for all equilibria with $x_1^* = 0$, the committee achieves the largest prestige, at level $r_2/(c_2(1+\beta_2))$, by setting $\pi_1 = 0$. Now suppose π_1 induces an equilibrium outcome of the line ministers' game with both ministers contributing and first suppose $r_1/c_1 \geq r_2/c_2$. Then we have equilibrium prestige

$$P^* = \frac{\frac{\pi_1}{c_1}r_1 + \frac{\pi_2}{c_2}r_2}{1+\beta_1+\beta_2} \leq \frac{\frac{r_1}{c_1}}{1+\beta_1+\beta_2} < \frac{r_1}{c_1(1+\beta_1)} = P^*_{k_1},$$

where $P^*_{k_1}$ is the prestige obtained if the committee instead chooses $\pi_1 = 1$. The case $r_1/c_1 < r_2/c_2$ similarly gives $P^* < P^*_{k_2}$, where $P^*_{k_2}$ is the prestige obtained by setting $\pi_1 = 0$.

We now consider the full model in which resources do adjust to the choice of weights of the committee, either because the finance minister is benevolent or prestige-maximizing. Suppose the equilibrium of the line ministers' game has both of them contributing. Recall that, by Proposition 6, the optimal choice of the finance minister is either $(\bar{r}_1, \underline{r}_2)$, or $(\underline{r}_1, \bar{r}_2)$. Suppose it is $(\bar{r}_1, \underline{r}_2)$ and first consider the case $\bar{r}_1/c_1 \geq \underline{r}_2/c_2$. Then we have equilibrium prestige

$$P^* = \frac{\pi_1\frac{\bar{r}_1}{c_1} + \pi_2\frac{\underline{r}_2}{c_2}}{1+\beta_1+\beta_2} \leq \frac{\frac{\bar{r}_1}{c_1}}{1+\beta_1+\beta_2} < \frac{\bar{r}_1}{c_1(1+\beta_1)} = P^*_{k_1},$$

where $P^*_{k_1}$ is the prestige obtained if the committee chooses $\pi_1 = 1$. To see this, note that with $\pi_1 = 1$, (2.6) must hold, which gives $x_2^* = 0$ in equilibrium; prestige then equals $r_1/(c_1(1+\beta_1))$ and antipoverty expenditures equal $r_1/(1+\beta_1)$. The finance minister maximizes both by setting $r_1 = \bar{r}_1$. Therefore, for either objective of the finance minister, the final level of prestige is $\bar{r}_1/(c_1(1+\beta_1))$. Now consider the case $\bar{r}_1/c_1 < \underline{r}_2/c_2$. Then we have that the equilibrium prestige equals

$$P^* = \frac{\pi_1 \frac{\bar{r}_1}{c_1} + \pi_2 \frac{\underline{r}_2}{c_2}}{1+\beta_1+\beta_2} \leq \frac{\frac{\underline{r}_2}{c_2}}{1+\beta_1+\beta_2} < \frac{\underline{r}_2}{c_2(1+\beta_2)} < \frac{\bar{r}_2}{c_2(1+\beta_2)} = P^*_{k_2},$$

where $P^*_{k_2}$ is the prestige obtained if the committee chooses $\pi_1 = 0$, following a reasoning similar to the one above. All other possibilities are analyzed similarly and lead to the same outcome. □

It is worth pointing out that Proposition 7 treats the two dimensions interchangeably. However, in some situations one of the dimensions—income, say—may be so salient that an index that completely disregards it is not feasible. In this case, it is possible for prestige to be maximized for a multidimensional index, rather than a unidimensional one. For instance, if we identify dimension 1 with income and impose that $\pi_1 \geq \pi_1^l > 0$, then, using Proposition 5, we see that a multidimensional index that induces both line ministers to contribute generates the largest feasible prestige if

$$\frac{\frac{\underline{r}_2}{c_2} - \frac{\underline{r}_1}{c_1}}{\frac{\underline{r}_1}{c_1}} > \frac{\beta_2}{(1-\pi_1^l)(1+\beta_1)}.$$

Proposition 7 demonstrates that the incentive to free-ride reduces the quantity of the public good "prestige" that line ministers are able to provide in equilibrium. By choosing $\pi_1 = 1$, the responsibility to provide prestige falls fully with line minister 1, and at the same time the effectiveness of her expenditures in providing prestige is maximal. It is then not surprising that a unidimensional index would record a larger percentage improvement than a multidimensional one.[21]

We now consider a benevolent committee. The following proposition analyzes the

[21] A partial counterpoint arises when the marginal effectiveness of expenditures in reducing deprivations is not constant as we consider, but is decreasing. In that case, one may benefit by splitting resources and responsibilities across two ministries to reduce the effect of decreasing marginal returns. Further, while our analysis treats both dimensions identically, it may be in practice impossible to assign a very low weight to an especially salient dimension such as income, as we previously discussed.

choice of a committee that maximizes antipoverty expenditures if the finance minister cannot change the resource allocation.

Proposition 8 (Antipoverty spending maximized by multidimensional index for any fixed r_1, r_2)**.** *If*

$$\frac{\sqrt{\beta_1\beta_2}}{(1+\beta_2)} \leq \sqrt{\frac{r_1}{r_2}} \leq \frac{(1+\beta_1)}{\sqrt{\beta_1\beta_2}}, \tag{2.13}$$

then the largest antipoverty expenditures obtain with a multidimensional poverty measure in which $\pi_1 \in (0,1)$*, and in particular*

$$\pi_1^*(r_1, r_2) = \left(1 + \sqrt{\frac{r_1}{r_2}\frac{c_2^2}{c_1^2}\frac{\beta_2}{\beta_1}}\right)^{-1}. \tag{2.14}$$

Proof. By Proposition 5, total antipoverty expenditures are

$$x_1^* + x_2^* = \begin{cases} \frac{r_1}{1+\beta_1} & \text{if } \frac{\pi_1}{\pi_2} > \frac{1+\beta_1}{\beta_2}\frac{\frac{r_2}{c_2}}{\frac{r_1}{c_1}}, \\ r_1 + r_2 - \frac{\pi_1\frac{r_1}{c_1}+\pi_2\frac{r_2}{c_2}}{1+\beta_1+\beta_2}\left(\frac{\beta_1 c_1}{\pi_1} + \frac{\beta_2 c_2}{\pi_2}\right) & \text{if } \frac{\beta_1}{1+\beta_2}\frac{\frac{r_2}{c_2}}{\frac{r_1}{c_1}} \leq \frac{\pi_1}{\pi_2} \leq \frac{1+\beta_1}{\beta_2}\frac{\frac{r_2}{c_2}}{\frac{r_1}{c_1}}, \\ \frac{r_2}{1+\beta_2} & \text{if } \frac{\pi_1}{\pi_2} < \frac{\beta_1}{1+\beta_2}\frac{\frac{r_2}{c_2}}{\frac{r_1}{c_1}}. \end{cases}$$

Simple algebra shows that the above is continuous in π_1. Further, we see that

$$\frac{d}{d\pi_1}\left(r_1 + r_2 - \frac{\pi_1\frac{r_1}{c_1}+\pi_2\frac{r_2}{c_2}}{1+\beta_1+\beta_2}\left(\frac{\beta_1 c_1}{\pi_1} + \frac{\beta_2 c_2}{\pi_2}\right)\right) = \frac{(1-\pi_1)^2\beta_1 r_2 c_1^2 - (\pi_1)^2\beta_2 r_1 c_2^2}{(1+\beta_1+\beta_2)(\pi_1)^2(1-\pi_1)^2 c_1 c_2}, \tag{2.15}$$

which starts positive for π_1 small and ends negative for π_1 large, meaning that the critical value of $\pi_1 = \pi_1^*(r_1, r_2)$, as defined in the statement of the proposition, is a maximum for $x_1^* + x_2^*$, if the appropriate condition on $\pi_1^*(r_1, r_2)/\pi_2^*(r_1, r_2)$ holds. Indeed, condition (2.13) ensures that

$$\frac{\beta_1}{1+\beta_2}\frac{\frac{r_2}{c_2}}{\frac{r_1}{c_1}} \leq \frac{\pi_1^*(r_1, r_2)}{\pi_2^*(r_1, r_2)} \leq \frac{1+\beta_1}{\beta_2}\frac{\frac{r_2}{c_2}}{\frac{r_1}{c_1}}, \tag{2.16}$$

therefore the largest possible antipoverty expenditures obtain at the strict global maximum

$\pi_1 = \pi_1^* (r_1, r_2) \in (0, 1)$. □

Note first that (2.13) is satisfied strictly for symmetric setups, so Proposition 8 applies if asymmetries across ministers are not too great. We also see that $\pi_1^*(r_1, r_2)$ decreases in r_1; moreover, assuming a symmetric propensity to divert resources away from the poor (i.e., $\beta_1 = \beta_2$), it turns out that $\pi_1^* > (1 - \pi_1^*) \equiv \pi_2^*$ if and only if $r_1/c_1 < r_2/c_2$. In other words, to maximize total antipoverty spending, the minister with the lower budget and responsible for the dimension with the largest cost of effecting improvements should be assigned a higher relative weight in the multidimensional poverty measure. This counterintuitive result obtains because ministers do not maximize the total sum spent on the poor; rather, they are interested in prestige. And, as equation (2.12) makes clear, these two objectives are in conflict, in equilibrium. Indeed, achieving prestige is harder—that is, it requires larger antipoverty expenditures—if the richer minister's contribution, which tends to be larger, is weighed less, and if the dimension with the largest cost of effecting improvements is weighed more.

We now assume that finance minister adapts the resource allocation to best suit her objectives in response to a change in weights. This introduces complications that make the determination of the dimension weight that is an overall maximizer of antipoverty expenditures very casuistic. However, the main message of Proposition 8 remains true: under certain conditions, total antipoverty expenditures are maximized for a truly multidimensional index, rather than for a unidimensional one. The following proposition provides a simple sufficient condition, limiting the maximal budget asymmetry between line ministers, and strategy to demonstrate this fact: choose π_1 such that $\pi_1/c_1 = \pi_2/c_2$, i.e. $\pi_1 = \frac{c_1}{c_1+c_2}$, and thus eliminate any incentive of the finance minister to manipulate resource allocations.

Proposition 9 (Antipoverty spending maximized by multidimensional index)**.** *If*

$$\underline{r}_1 > \frac{\beta_1}{1+\beta_1+\beta_2} R \text{ and } \underline{r}_2 > \frac{\beta_2}{1+\beta_1+\beta_2} R \tag{2.17}$$

then antipoverty expenditures are larger with $\pi_1 = \frac{c_1}{c_1+c_2}$ *than with* $\pi_1 = 0$ *or* $\pi_1 = 1$. *Therefore, the largest antipoverty expenditures obtain with a truly multidimensional poverty measure in which* $\pi_1 \in (0, 1)$.

Proof. It is sufficient to show that the choice of $\pi_1 = \frac{c_1}{c_1+c_2}$ beats setting $\pi_1 = 0$ or $\pi_1 = 1$. First note that, if $\pi_1 = \frac{c_1}{c_1+c_2}$ and (2.17) holds, then any possible distribution of resources violates (2.6) and (2.9), so it results in an equilibrium outcome in which both line ministers contribute. To see this, consider (2.6). For $\pi_1 = \frac{c_1}{c_1+c_2}$, it is violated if $\beta_2 r_1 < r_2(1+\beta_1)$, and a sufficient condition for this to be always true is that it is true for $r_1 = \bar{r}_1$ and $r_2 = \underline{r}_2$, implying

$$\beta_2 R < \underline{r}_2(1+\beta_1+\beta_2),$$

but this is implied by (2.17). (The argument for (2.9) is similar.) Therefore, by setting $\pi_1 = \frac{c_1}{c_1+c_2}$, the committee forces an equilibrium outcome in which both line ministers expend resources to fight poverty.

Further, by setting $\pi_1 = \frac{c_1}{c_1+c_2}$, the committee ends up neutering the finance minister. Indeed, by Proposition 6, any feasible distribution of resources is indifferent to the finance minister, regardless of whether she maximizes prestige or antipoverty expenditures. Using (2.12), the total antipoverty expenditures generated by $\pi_1 = \frac{c_1}{c_1+c_2}$ is

$$\frac{R}{1+\beta_1+\beta_2}.$$

If instead the committee sets $\pi_1 = 0$ or $\pi_1 = 1$, total antipoverty spending cannot exceed $\frac{R-\underline{r}_1}{1+\beta_2}$ or $\frac{R-\underline{r}_2}{1+\beta_1}$, respectively. But (2.17) implies

$$\frac{R}{1+\beta_1+\beta_2} > \frac{R-\underline{r}_1}{1+\beta_2}, \text{ and } \frac{R}{1+\beta_1+\beta_2} > \frac{R-\underline{r}_2}{1+\beta_1};$$

so the proof is complete. □

The following proposition calculates the dimension weights that maximize total antipoverty spending, under (2.17) and a "reasonableness" condition that implies that the maximum budget for one line minister is larger than the minimum budget for the other minister, after normalizing by the propensity to divert resources to private consumption.

Proposition 10 (Optimal dimension weights)**.** *Assume* (2.17) *and*

$$\frac{\bar{r}_1}{\beta_1} > \frac{\underline{r}_2}{\beta_2} \text{ and } \frac{\bar{r}_2}{\beta_2} > \frac{\underline{r}_1}{\beta_1}. \tag{2.18}$$

1. *If the finance minister maximizes measured poverty reduction, then* $\pi_1 = \frac{c_1}{c_1+c_2}$ *maximizes total antipoverty spending.*

2. *If the finance minister maximizes antipoverty spending, then* $\pi_1 = \pi_1^*(\bar{r}_1, \underline{r}_2)$ *or* $\pi_1 = \pi_1^*(\underline{r}_1, \bar{r}_2)$*—where* $\pi_1^*(r_1, r_2)$ *is defined by* (2.14)*—maximizes total antipoverty spending.*

Proof. Part 1. By (2.17) and Proposition 5, we know that $\pi_1 = \frac{c_1}{c_1+c_2}$ is superior to any choice of π_1 that induces zero expenditure from one of the line ministers. Therefore, we only need to consider choices of π_1 and resource allocations that induce both line ministers to devote resources to fight poverty. In this case, the derivative of total antipoverty expenditure is given by (2.15). Note how this equation implies that, if the derivative (from the left) is positive at some $\pi_1 = \tilde{\pi}_1$, and if resources do not change for all $\pi_1 < \tilde{\pi}_1$, then the derivative remains positive for all $\pi_1 < \tilde{\pi}_1$. And if the derivative (from the right) is negative at some $\pi_1 = \tilde{\pi}_1$, and if resources do not change for all $\pi_1 > \tilde{\pi}_1$, then the derivative remains negative for all $\pi_1 > \tilde{\pi}_1$. These two results in turn imply that total antipoverty expenditures have a kink and a maximum for $\pi_1 = \tilde{\pi}_1$. But this is exactly what happens for $\tilde{\pi}_1 = \frac{c_1}{c_1+c_2}$. Indeed, as we know from Proposition 6, for $\pi_1 < \frac{c_1}{c_1+c_2}$ the prestige-maximizing finance minister chooses $(\underline{r}_1, \bar{r}_2)$, and condition (2.18) shows that, coming from the left, the derivative in (2.15) is positive at $\pi_1 = \frac{c_1}{c_1+c_2}$. And if $\pi_1 > \frac{c_1}{c_1+c_2}$, then the prestige-maximizing finance minister chooses $(\bar{r}_1, \underline{r}_2)$, and condition (2.18) shows that, coming from the right, the derivative in (2.15) is negative at $\pi_1 = \frac{c_1}{c_1+c_2}$.

Part 2. Instead, if the finance minister also maximizes total antipoverty expenditures, her choices as a function of π_1 are reversed: $\bar{r}_1, \underline{r}_2$ for $\pi_1 < \frac{c_1}{c_1+c_2}$ and $\underline{r}_1, \bar{r}_2$ for $\pi_1 > \frac{c_1}{c_1+c_2}$. The previous set of steps now implies that $\pi_1 = \frac{c_1}{c_1+c_2}$ is a local minimum. Since (2.17) implies $\pi_1 = 0$ and $\pi_1 = 1$ are not maximizers either, by Proposition 9, continuity and differentiability of the total expenditure function imply that the only possibilities for a maximum are where (2.15) equals zero, taking into account the choice of resources by the finance minister. For $r_1 = \bar{r}_1$, $r_2 = \underline{r}_2$, the derivative in (2.15) is zero at $\pi_1^* \left(\bar{r}_1, \underline{r}_2\right)$. Further, under (2.18), $\pi_1^* \left(\bar{r}_1, \underline{r}_2\right)$ is smaller than $\frac{c_1}{c_1+c_2}$, so the choice of the finance minister is indeed $r_1 = \bar{r}_1$, $r_2 = \underline{r}_2$. And for $r_1 = \underline{r}_1$, $r_2 = \bar{r}_2$, the derivative in (2.15) is zero at $\pi_1^* \left(\underline{r}_1, \bar{r}_2\right)$. Under (2.18), $\pi_1^* \left(\underline{r}_1, \bar{r}_2\right)$ is larger than $\frac{c_1}{c_1+c_2}$, so the choice of the finance minister is indeed $r_1 = \underline{r}_1$, $r_2 = \bar{r}_2$. Therefore, both candidates are valid and the overall maximum is determined by a direct comparison of the objective function at these two values of π_1. □

When the the finance minister's objective aligns with that of a benevolent committee, the second part of Proposition 10 confirms the main forces explored in Proposition 8, this time taking into account the endogenous allocation of resources. Interestingly, if the finance minister's objective is prestige and so it is in conflict with that of the committee, then the committee finds it optimal to neuter the finance minister. And, in this last case, if the cost of reducing deprivations across the two dimensions in the same, then it turns out that equal dimension weights are best.

The basic combined message of Propositions 8, 9, and 10 is that having a multidimensional index rather than a unidimensional one adds an important tool to the arsenal of a policymaker that wants to increase antipoverty expenditures. This not a mere mathematical statement, but also one about which choices are politically feasible. For instance, suppose line ministers are symmetric (in resources, propensities for private consumption, and costs of reducing deprivations) and suppose that an antipoverty-maximizing commission desires to set $\pi_1 = \frac{1}{2}$, either because the finance minister is prestige-motivated or because resources do not adjust to weights. Since $\pi_j \equiv w_j \mu(g^0_{*j})/M_0$, this implies that the dimension with initially lower average normalized shortfalls should receive a *higher* weight. This corresponds with what Decancq and Lugo (2013) call frequency-based weights, where a dimension's weight is an inverse function of the average deprivation in that dimension. In the literature, the normative logic behind frequency-based weights "is that if owning a refrigerator is much more common than owning a dryer, a greater weight should be given to the former indicator so that if an individual does not own a refrigerator, this rare occurrence will be taken much more into account in computing the overall degree of poverty than if some individual does not own a dryer" (Deutsch and Silber, 2005, p. 150). Therefore, the committee would have a valid, normative justification for implementing, or at least approaching, its most desired weighting scheme.

2.3.4 Dimensional Decomposability

If the poverty measure satisfies dimension decomposability, then, in addition to the overall scalar index, each dimension's contribution to poverty may receive publicity in a coherent fashion. In our model, this is likely to reduce the incentive to free-ride and we implement this as a reduction in β_1 or β_2. We assign the choice of whether to reduce β_1 or β_2 to the committee and now we ask whether the committee would do so. *Prima facie*, a reduction in free-riding appears positive. However, we have seen earlier that it is possible for a benevolent decision-maker to choose to make the line ministers' pursuit of prestige harder. The next proposition shows that the first effect always dominates.

Proposition 11. *Assume* (2.17) *and* (2.18). *The committee always prefers to implement a marginal decrease in* β_1 *or* β_2.

Proof. We conduct the proof for β_1 only. The result is obvious for equilibria in which only one line minister contributes: using Proposition 5 and, without loss of generality, considering the case $x_2^* = 0$, prestige equals

$$\frac{\pi_1 r_1}{c_1(1+\beta_1)},$$

and expended antipoverty resources equal

$$\frac{r_1}{1+\beta_1}.$$

The choice of the finance minister is unaffected by changes in π_1 and β_1. Therefore, using our previous results, a prestige-motivated committee would always reduce β_1, and so would a committe that maximizes antipoverty expenditures and finds $\pi_1 = 0$ or $\pi_1 = 1$ optimal. The only case left to consider is that of a benevolent committee's choice that induces an equilibrium of the line ministers' game in which both of them contribute. By the reasoning in Propositions 8 and 10, this may have $\pi_1 = \frac{c_1}{c_1+c_2}$ or $\pi_1 = \pi_1^*(r_1, r_2)$ and equation (2.16) must hold. In either case, using (2.12) we have

$$\begin{aligned} \frac{d\left(x_1^* + x_2^*\right)}{d\beta_1} &= \frac{\partial\left(x_1^* + x_2^*\right)}{\partial\pi_1}\frac{\partial\pi_1}{\partial\beta_1} + \frac{\partial\left(x_1^* + x_2^*\right)}{\partial\beta_1} \\ &= 0 + \frac{\partial\left(x_1^* + x_2^*\right)}{\partial\beta_1} \\ &= -\frac{\frac{r_1}{c_1}\pi_1 + \frac{r_2}{c_2}\pi_2}{\left(1+\beta_1+\beta_2\right)^2\pi_1\pi_2}(c_1\pi_2(1+\beta_2) - c_2\pi_1\beta_2), \end{aligned}$$

where the second line follows because either $\pi_1 = \frac{c_1}{c_1+c_2}$, regardless of β_1, or because $\frac{\partial\left(x_1^*+x_2^*\right)}{\partial\pi_1} = 0$ when $\pi_1 = \pi_1^*(r_1, r_2)$, as we saw earlier using (2.15).[22] When $\pi_1 = \frac{c_1}{c_1+c_2}$, simple algebra shows the above is negative. When $\pi_1 = \pi_1^*(r_1, r_2)$, the sign of the above displayed derivative is negative iff

$$c_1(1 - \pi_1^*(r_1, r_2))(1+\beta_2) > c_2\pi_1^*(r_1, r_2)\beta_2$$

[22] Note that, since we are considering only a marginal change in β_1, the change in the optimal π_1 is not sufficient to change the choice of the finance minister, see Proposition 6, because, under (2.17) and (2.18), π_1^* is strictly bounded away from $c_1/(c_1+c_2)$.

$$\iff r_1(1+\beta_2)^2 > \beta_1\beta_2 r_2$$

$$\iff r_1((1+\beta_2)^2 + \beta_1\beta_2) > \beta_1\beta_2 R,$$

and the above is implied by (2.17). Therefore, increasing β_1 always leads to a decrease in total antipoverty expenditures. □

2.3.5 Relation with the Dashboard Approach

What happens if, in addition to a scalar multidimensional poverty measure, a dashboard of dimension-specific measures receives publicity and enters into the line ministers' prestige objective? While a complete investigation of dashboards is deserving of its own analysis, we can provide some insights into this question by returning to our full framework, for which line ministers' contributions are described by equation (2.5), initializing θ_1 and θ_2 to zero to match our previous results, and then marginally increase the parameters θ_1 and θ_2 to reflect the increased, separate relevance of each dimension. As one might expect, since now ministers cannot hide their lack of contributions as easily, measured poverty reduction and antipoverty expenditures increase, if both θ_1 and θ_2 marginally increase.

Proposition 12. *Antipoverty expenditures and measured poverty reductions increase in θ_1 and θ_2, with respect to what determined in Sections 3.1-3.3.*

Proof. We conduct the proof only for the case in which both line ministers contribute to poverty reduction in equilibrium, the finance minister is prestige-maximizing, and the committee is benevolent. All other situations are dealt with similarly. In such an equilibrium, equation (2.5) leads to equilibrium prestige

$$P^* = \frac{\left(\pi_1\left(\pi_1+\theta_1\right)\frac{r_1}{c_1}\right)\left(\left(\pi_2+\theta_2\right)+\theta_2\beta_2\right)+\left(\left(\pi_1+\theta_1\right)+\theta_1\beta_1\right)\left(\pi_2\left(\pi_2+\theta_2\right)\frac{r_2}{c_2}\right)}{\left(\left(\pi_1+\theta_1\right)+\theta_1\beta_1\right)\left(\left(\pi_2+\theta_2\right)+\theta_2\beta_2\right)+\left(\left(\pi_2+\theta_2\right)+\theta_2\beta_2\right)\beta_1\pi_1+\pi_2\beta_2\left(\left(\pi_1+\theta_1\right)+\theta_1\beta_1\right)},$$

which agrees with (2.11) when $\theta_1 = \theta_2 = 0$. In this case, and by continuity in a neighborhood of $\theta_1 = \theta_2 = 0$ as well, all results derived in Sections 3.1-3.3 hold. In particular, the analogue of Part 1) of Proposition 10 now requires π_1 and π_2 to be chosen to satisfy

$$\frac{\pi_1(\pi_1+\theta_1)}{c_1}\left((\pi_2+\theta_2)+\theta_2\beta_2\right)=\left((\pi_1+\theta_1)+\theta_1\beta_1\right)\frac{\pi_2(\pi_2+\theta_2)}{c_2}, \tag{2.19}$$

so that the finance minister choice of r_1 and r_2 has no bearing on P^* displayed above. Because of this fact, we proceed with a resource allocation such that $r_1/c_1 = r_2/c_2$, to make calculations simpler. Applying the implicit function theorem to the above displayed equation for P^* and evaluating the derivatives at $\theta_1 = \theta_2 = 0$, we have

$$\begin{aligned}\Delta P^* = &\frac{\frac{r_1}{c_1}}{1+\beta_1+\beta_2}\Delta\pi_1 + \frac{\frac{r_2}{c_2}}{1+\beta_1+\beta_2}\Delta\pi_2 \\ &+ \frac{\beta_1((1+\beta_1)\pi_2\frac{r_2}{c_2}-\beta_2\pi_1\frac{r_1}{c_1})}{\pi_1(1+\beta_1+\beta_2)^2}\Delta\theta_1 + \frac{\beta_2((1+\beta_2)\pi_1\frac{r_1}{c_1}-\beta_1\pi_2\frac{r_2}{c_2})}{\pi_2(1+\beta_1+\beta_2)^2}\Delta\theta_2.\end{aligned}$$

The above displayed equation is positive, if both ministers contribute in equilibrium. Indeed, $\pi_1+\pi_2 = 1$ implies $\Delta\pi_1 = -\Delta\pi_2$, which together with $r_1/c_1 = r_2/c_2$ give us

$$\Delta P^* = \frac{\beta_1((1+\beta_1)\pi_2\frac{r_2}{c_2}-\beta_2\pi_1\frac{r_1}{c_1})}{\pi_1(1+\beta_1+\beta_2)^2}\Delta\theta_1 + \frac{\beta_2((1+\beta_2)\pi_1\frac{r_1}{c_1}-\beta_1\pi_2\frac{r_2}{c_2})}{\pi_2(1+\beta_1+\beta_2)^2}\Delta\theta_2.$$

And since both ministers are assumed to contribute, equations (2.6) and (2.9) are violated, ensuring that the right-hand side of the above displayed equation is positive.

We now show that x_1 increases in θ_1. Using (2.5) and proceeding similarly as above, we obtain

$$\Delta x_1 = \frac{\beta_1}{\pi_1^2}\left(P^*+\beta_1P^*-\pi_1\frac{r_1}{c_1}\right)\Delta\theta_1 + \frac{\beta_1P^*}{\pi_1^2}\frac{\Delta\pi_1}{\Delta\theta_1}\Delta\theta_1 - \frac{\beta_1}{\pi_1}\frac{\Delta P^*}{\Delta\theta_1}\Delta\theta_1,$$

and substituting the value of ΔP^* above calculated for $\Delta\theta_1 > 0$ and $\Delta\theta_2 = 0$, we obtain

$$\Delta x_1 = \frac{\beta_1(1+\beta_2)((1+\beta_1)\pi_2\frac{r_2}{c_2}-\beta_2\pi_1\frac{r_1}{c_1})}{\pi_1^2(1+\beta_1+\beta_2)^2}\Delta\theta_1 + \frac{\beta_1P^*}{\pi_1^2}\frac{\Delta\pi_1}{\Delta\theta_1}\Delta\theta_1.$$

Both terms in the right-hand side of the above-displayed equation are positive. The first term is positive because equations (2.6) and (2.9) are violated if ministers contribute positive amounts. The second is positive because $\frac{\Delta\pi_1}{\Delta\theta_1} > 0$. To see this, note that the equation the committee uses to choose weights, that is (2.19) can be rewritten as

$$\frac{\pi_1 (\pi_1 + \theta_1)}{\pi_1 + \theta_1 + \theta_1 \beta_1} = \frac{\pi_2 (\pi_2 + \theta_2)}{((\pi_2 + \theta_2) + \theta_2 \beta_2)} \frac{c_1}{c_2}.$$

Note that the left-hand side is increasing in π_1 and decreasing in θ_1. Similarly, the right-hand side is increasing in π_2. Therefore, if θ_1 increases, the left-hand side decreases and π_1 must increase to restore balance. If π_1 were to decrease, then the left-hand side would decrease further, while the right-hand side would increase because $\pi_2 = 1 - \pi_1$. Thus, we see that θ_1 and π_1 move in the same direction, as we wanted to show. The proof that x_2 increases in θ_2 is similar and here omitted. Note as well that the above discussion implies that, if θ_1 increases, then x_2 *falls*. □

Proposition 12 suggests that evaluating ministers on a dashboard as well as an MPI can be beneficial, keeping all other parameters of the model fixed. Whether this is a realistic assumption or not depends on a number of behavioral responses. For instance, one of the advantages of an MPI is the allure of a single number, which makes the measure especially salient in the public eye. If the inclusion of the dashboard raises θ_1, but at the same time makes messages about the evolution of inequality harder to grasp, with respect to providing a single number, then a countervailing effect arises, which can be captured in our framework as a decrease of γ_1 in equation (2.1), which implies that β_1 increases. And now Proposition 11 shows that this effects a decrease of antipovery spending.

2.4 Illustration

We illustrate the results in Section 2.3 using data from Mexico, the first country to adopt an official multidimensional poverty measure based on AF. Figure 2.2 illustrates the four possible combinations of preferences for the technical committee and the finance minister—each can be benevolent (B) or prestige-maximizing (P). For each possible combination, represented by a square, the top triangle refers to cases where both line ministers contribute, while the bottom triangle refers to cases where only one line minister contributes. In each triangle, we report triples (π_1, r_1, r_2) that are equilibrium candidate outcomes for that case, taking into account that a prestige-maximizing committee will never choose weights such that both line ministers contribute. For each square, the ultimate outcome is obtained by the committee comparing the value of its objective function at the candidates in the top and bottom triangle.

Figure 2.2: Summary of Theoretical Results

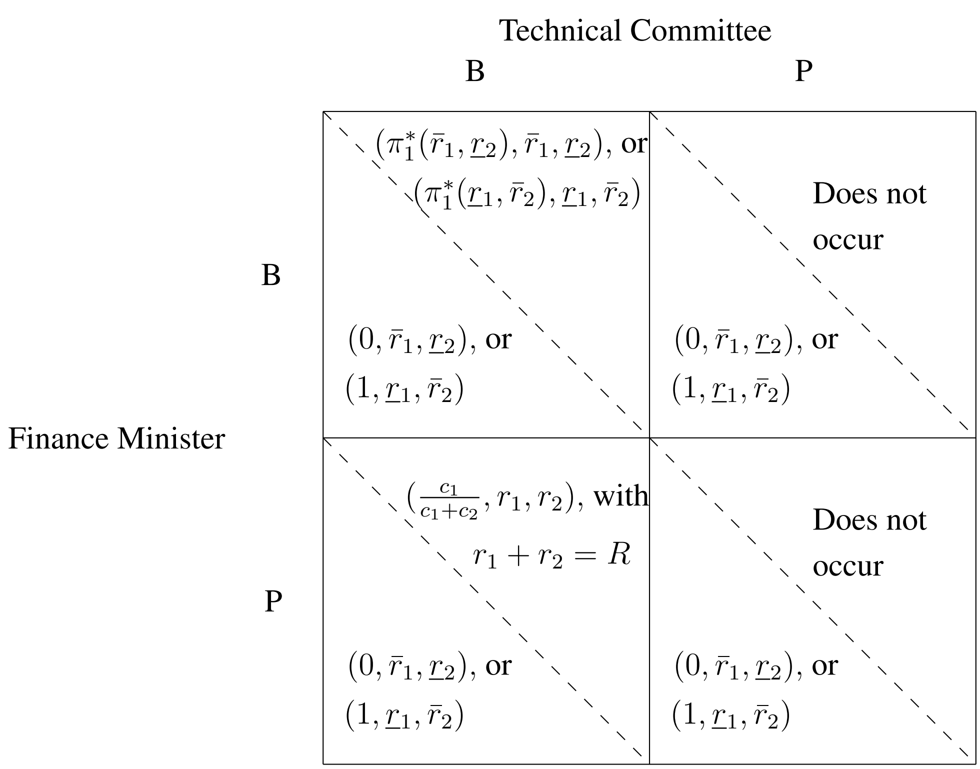

We use indicators from Mexico's official measure and the same household survey data that is used by the Mexican government. Since our theoretical model has two dimensions, we restrict ourselves to two indicators; because all of the dimensions except income are dichotomous in Mexico's measure, we focus on $\alpha = 0$, the same case that we focused on in the theoretical part of this paper.

Our illustration accomplishes four objectives. First, it illustrates the usefulness of our framework even when one or more of the dimensions of poverty is measured ordinally. Since many of the governments implementing multidimensional poverty use ordinal variables in their measure, they almost exclusively use M_0; hence, it is important to accommodate this possibility in our framework. Second, it illustrates many of the theoretical results from Section 2.3, in particular the strength of the conflict between self interest and benevolence. A self-interested government that seeks to maximize prestige will reallocate resources to the minister with the higher ratio of relative weight in the index to marginal cost of reducing deprivations. In other words, the budget allocated to ministry j by a prestige-maximizing finance minister is increasing in dimension j's weight in the multidimensional index and decreasing in the cost of reducing deprivations in dimension j; budgets are allocated to reap "low-hanging fruit." A benevolent government that seeks to maximize total antipoverty spending will do just the opposite. Similar considerations apply to the committee: in particular, when marginal costs of reducing deprivations differ across dimensions, a prestige-maximizing committee assigns as much weight as possible to the dimension with the lower cost of reducing deprivations, *ceteris paribus*; a benevolent committee instead chooses an optimal weight for dimension j that is *increasing* in the cost of reducing deprivations in dimension j.

Third, by illustrating that ministers with a low relative budget, a high propensity for diverting funds, a low effectiveness of antipoverty spending, a low relative weight in the index (due to a low absolute weight or low deprivations relative to other dimensions), have great incentives to free ride when multidimensional poverty is measured using a scalar measure, it points to the importance of decomposing multidimensional poverty by dimension and prominently publishing its partial indices as well as the overall measure. This, in turn, highlights an underexplored benefit of the factor decomposability or dimensional breakdown axiom for multidimensional poverty measures, as discussed in the Introduction; such decompositions and the publishing of partial indices have been advocated by proponents of the AF measure (e.g., Alkire et al., 2011).

Fourth, it illustrates what it means in practice for a benevolent technical committee facing a prestige-maximizing finance minister (a likely real-world scenario) to set $\pi_1 = \frac{c_1}{c_1+c_2}$. In the case that the costs of reducing deprivations in each dimension are similar, this will approximate to $\pi_1 = \pi_2 = \frac{1}{2}$. This does *not* imply setting equal absolute weights w_1 and w_2 in the index: although the absolute weights are the ones discussed by policymakers and scholars, the relative weights π_1 and π_2 also depend on initial deprivations in each dimension, which in turn depend on each dimension's deprivation cut-off. If deprivations are greater in dimension 1, setting $\pi_1 = \pi_2 = \frac{1}{2}$ requires giving *less* absolute weight to dimension 1, i.e. $w_1 < w_2$. In the alternative scenario that absolute weights are fixed, we show how the committee could adjust deprivation cut-offs (e.g., the income poverty line) to achieve optimal relative weights.

2.4.1 Data and Mexico's Multidimensional Poverty Measure

We use the 2010 Encuesta Nacional de Ingresos y Gastos de Hogares (National Income and Expenditure Survey; ENIGH). The survey was modified in 2008 to include the Módulo de Condiciones Socioeconómicas (Socioeconomic Conditions Module; MCS), which was specifically designed to enable the measurement of mutlidimensional poverty using the indicators decided on by the Mexican government's Consejo Nacional de Evaluación de la Política de Desarrollo Social (National Council for the Evaluation of Social Development Policy; CONEVAL). The 2010 survey includes 30,169 households, and is representative at the national, urban, rural, and state levels. The government is required by law to produce multidimensional poverty measures for the country as a whole and for each state every two years, as well as for all municipalities every five years (CONEVAL, 2010). Mexico's official poverty measurement uses the ENIGH MCS data and identifies seven dimensions: income, education, access to health care, access to social security, housing quality, basic housing services, and food security.

Income in our illustration is defined identically to the definition used in Mexico's official poverty measurement.[23] Two deprivation cut-offs are used in the income dimension, resulting in two reported levels of multidimensional poverty in Mexico. The first cut-off, called the minimum wellbeing line, corresponds to extreme poverty and is calculated as the per-adult-equivalent cost of a minimum basket of food, which in turn is based on caloric intake requirements and observed consumption patterns of households whose members approximately meet the caloric intake minimums. The second cut-off, called the wellbeing line, also incorporates the cost of basic non-food necessities based on consumption patterns of the same families (those who approximately meet the caloric intake minimums).[24]

Of the seven dimensions included in Mexico's multidimensional poverty measure, we choose to include income as one of the two dimensions of our illustration because it is given a much larger weight than other dimensions: the weight on income deprivations is $1/2$, while the weight on each of the remaining six dimensions is $1/12$. This accords with Foster's (2007, p. 9) point that "a weight on income that is higher than the equal-weight case, but lower than the full-weight case, represents a reasonable compromise between a traditional 'economic' view of poverty and a more inclusive multidimensional view." In addition, it is important to include income as a dimension due to its fungibility and usefulness in reducing deprivations in other dimensions and its salience in discussions about poverty (Foster, 2007), and because the unidimensional measure was income-based, so that discussions in our model of whether total antipoverty spending is higher when a unidimensional or multidimensional measure is adopted rely on income being one of the dimensions

[23] It includes income from labor, self-employment, capital, public and private transfers, the imputed value of own production, regular but not extraordinary in-kind payments, and regular but not extraordinary gifts received. The imputed value of owner-occupied housing is *not* included in the official income concept, with the argument that households cannot use this "income" component to meet their basic needs.

[24] The cut-offs are calculated separately for urban and rural areas. The minimum wellbeing line for 2010 equals 1125.42 pesos per adult equivalent per month in urban areas and 800.26 pesos in rural areas. The wellbeing line equals 2328.82 pesos in urban areas and 1489.78 in rural areas. Adult equivalence and economies of scale are taken into account as follows: the household head is assigned a weight of 1, additional adults ages 19 and above are assigned a weight of 0.9945, adolescents ages 13–18 are assigned a weight of 0.7057, children ages 6–12 are assigned a weight of 0.7382, and children ages 0–5 are assigned a weight of 0.7031.

of the multidimensional measure.

With respect to the non-income dimensions, the indicators and their deprivation cut-offs were chosen based on rights guaranteed in the Mexican Constitution and other laws; this method of selecting dimensions and cut-offs is praised by Alkire and Foster (2011b). In our illustration, we focus on access to healthcare as our second dimension. In Mexico, healthcare is not universally provided by the government; a household is defined as deprived in access to healthcare if they do not have any type of medical insurance (including the government-provided insurance for the uninsured *Seguro Popular*, insurance through the Mexican Institute of Social Security which is provided to formal sector workers, insurance for state employees, or private insurance). There are a number of reasons we chose health as the second dimension for our illustration. First, Birdsall (2011) argues that multidimensional poverty measures have independent value compared to Ravallion's (1996, p. 1332) "multiple-indicator approach" if they reveal additional information about poverty dynamics; lack of health insurance in Mexico has indeed been shown to increase the vulnerability of income-poor households to becoming poorer (López-Calva and Ortiz-Juárez, 2009). Second, there is a ministry that is clearly tied to this dimension (the Health Ministry). Although there is also a ministry clearly tied to the dimension of education, our assumption that spending results in proportional reductions in the number of deprived individuals seems more adequate for access to healthcare than for the education dimension since educational deprivation in Mexico's measure is a stock variable for those older than 15.[25]

[25] Specifically, deprivations in education are defined as follows. For children ages 3–15, those who neither attend a formal education institution nor have completed the minimum basic secondary education are considered deprived. For those older than 15 born after 1982, those who did not complete the required basic secondary education are considered deprived; for those born before 1982, those who did not complete the former requirement of primary education are considered deprived.

2.4.2 Illustration with Equal Costs

We denote income as dimension $j = 1$ and health as $j = 2$. Using the minimum wellbeing line (MWL), $\mu(g^0_{*1}) = 0.194$, while using the wellbeing line (WL), $\mu(g^0_{*1}) = 0.520$. In other words, 19% of the Mexican population had income below the MWL in 2010 and was thus unable to afford a basic food basket based on a minimum caloric intake for each of its members. Just over half of the population had income below the cost of a basket of basic needs including both food and nonfood needs. In access to health insurance, $\mu(g^0_{*2}) = 0.292$; 29% of Mexicans did not have access to some form of health care despite the fact that access is, in theory, guaranteed by law.

For illustrative purposes, we assign the following parameters. Suppose the ministry of finance has an available budget of $R = 1$. As long as there is a positive minimum budget that must be assigned to each minister, our assumption from Section 2.2 that $\bar{r}_j < 1$ for $j = 1, 2$ will hold: neither minister can be allocated enough resources to completely eliminate deprivations in her dimension. Since the health ministry has a higher budget than the social development ministry, we set the minimum that can be allocated to the development minister $0.3 = \underline{r}_1 < \underline{r}_2 = 0.4$. The health minister also has more significant other demands on her resources, however: public health spending is also intended to reach the non-poor. Hence, we set $0.1 = \beta_1 < \beta_2 = 0.3$. In this section we consider equal effectiveness of antipoverty spending: $c_1 = c_2 = 1$, and in the next we illustrate how changes to these costs affect equilibrium weights, budget allocations, and antipoverty spending.

Given these parameters and the existing deprivations in Mexico using the minimum well-being line (panel A) and well-being line (panel B) as the deprivation cut-off in the income dimension, Table 2.1 shows the relative and absolute weights assigned to each dimension by the technical committee, the budgets assigned by the finance minister, and the resulting antipoverty spending from each line minister, as well as the resulting levels of measured poverty and government prestige. As shown in Figure 2.2, a prestige-maximizing technical committee will always prefer a unidimensional measure. In this illustration, the minister assigns all the weight to the income dimension (maintaining a traditional unidimensional poverty measure based only on income) because $\beta_1 < \beta_2$ while $c_1 = c_2$, so the development minister will spend more of her resources on the poor than the health minister. This can be seen in columns 1 and 2 of Table 2.1, where $\pi_1^* = 1$ and hence $w_1^* = 1$, in both panels A and B.

Table 2.1: Illustration with Mexican Data and Equal Costs

Technical committee maximizes	P^*		$x_1^* + x_2^*$	
Finance minister maximizes	P^*	$x_1^* + x_2^*$	P^*	$x_1^* + x_2^*$
Panel A: Minimum well-being line				
$\mu(g_1^0)$	0.194	0.194	0.194	0.194
$\mu(g_2^0)$	0.292	0.292	0.292	0.292
π_1^*	1.000	1.000	0.500	0.320
π_2^*	0.000	0.000	0.500	0.680
w_1^*	1.000	1.000	0.601	0.415
w_2^*	0.000	0.000	0.399	0.585
r_1^*	0.600	0.600	any	0.600
r_2^*	0.400	0.400	any	0.400
x_1^*	0.545	0.545	0.529	0.497
x_2^*	0.000	0.000	0.186	0.254
$x_1^* + x_2^*$	0.545	0.545	0.714	0.750
M_0^*	0.194	0.194	0.233	0.252
$M_0'^*$	0.088	0.088	0.150	0.168
P^*	0.545	0.545	0.357	0.331
Panel B: Well-being line				
$\mu(g_1^0)$	0.520	0.520	0.520	0.520
$\mu(g_2^0)$	0.292	0.292	0.292	0.292
π_1^*	1.000	1.000	0.500	0.320
π_2^*	0.000	0.000	0.500	0.680
w_1^*	1.000	1.000	0.360	0.209
w_2^*	0.000	0.000	0.640	0.791
r_1^*	0.600	0.600	any	0.600
r_2^*	0.400	0.400	any	0.400
x_1^*	0.545	0.545	0.529	0.497
x_2^*	0.000	0.000	0.186	0.254
$x_1^* + x_2^*$	0.545	0.545	0.714	0.750
M_0^*	0.520	0.520	0.374	0.340
$M_0'^*$	0.236	0.236	0.241	0.227
P^*	0.545	0.545	0.357	0.331

Because the health dimension carries no weight in the measurement of poverty, the finance minister assigns as little budget as possible to the health minister, which in this case is $r_2^* = \underline{r}_2 = 0.4$, and assigns the remaining budget to the development minister, $r_1^* = \bar{r}_1 = R - \underline{r}_2 = 0.6$. This result does not depend on whether the finance minister is prestige-maximizing or benevolent: the minister fully anticipates the next stage of the game, in which line ministers simultaneously decide how much to spend on poverty, and in which—because the health minister's actions get no weight in measured poverty—she will contribute the minimum possible to poverty reduction, which we have here normalized to 0. Hence, when the technical committee is prestige-maximizing, the actions of a prestige-maximizing and benevolent finance minister coincide: as much of the total budget as possible is allocated to the minister that can reduce poverty (i.e., the minister corresponding to the unidimensional poverty measure selected by the technical committee). In this example, the development minister spends nearly all of her budget on the poor ($x_1^*/r_1 = 0.545/0.6 = 91\%$), since she has a low propensity for private spending (e.g., spending on the non-poor, clientelism, corruption) with $\beta_1 = 0.1$, while the health minister spends the minimum amount possible, which we have normalized to 0.

Turning now to the benevolent technical committee that wishes to maximize antipoverty spending by the line ministers, the committee must look ahead to how the resource minister will assign budgets to determine the optimal weights. If the finance minister is prestige-maximizing (in other words, the preferences of the technical committee and finance minister conflict; column 3 of Table 2.1), the techical committee neuters the finance minister by setting $\pi_1 = \frac{c_1}{c_1+c_2}$, which in the $c_1 = c_2$ case we are analyzing here gives $\pi_1 = \pi_2 = \frac{1}{2}$. We say that this choice neuters the finance minister because it makes any allocation by the finance minister result in the same level of prestige. If the committee did not do so, the prestige-maximizing finance minister would shuffle resources to increase measured poverty reduction, but decrease total antipoverty spending.

It is worth noting that, when initial deprivations differ, i.e. $\mu(g_1^0) \neq \mu(g_2^0)$, as will inevitably be the case in practice, setting $\pi_1 = \pi_2$ is not the same as setting equal absolute weights w_1 and w_2, which are typically the salient version of the weights in public discussions and policy debates. Instead, setting $\pi_1 = \pi_2$ requires using the absolute weights (or deprivation cut-offs) to counterbalance initial deprivations and make each dimension equally relevant in the initial level of measured poverty. This can be seen by comparing Panels A and B of Table 2.1, column 3. When initial deprivations are relatively low in the income dimension, as they are when the MWL is used as the deprivation cut-off in the income dimension, $0.601 = w_1 > w_2$ (Panel A); when deprivations are increased by increasing the deprivation cut-off, so that $\mu(g_1^0)$ rises, the committee shifts absolute weight to column 2 so that $0.360 = w_1 < w_2$. This use of low weights for dimensions with high deprivations coincides with the normative frequency-based weighting method (Decancq and Lugo, 2013). Alternatively, if w_1 and w_2 are fixed (say, at $\frac{1}{2}$), the committee can instead manipulate deprivation cut-offs to acheive $\pi_1 = \pi_2$; in the Mexican case, since the health cut-off is fixed (it is a categorical dimension), the optimal deprivation cut-off for income z_1^* would be somewhere between the currently used MWL (which results in $\pi_1 > \pi_2$ for $w_1 = w_2 = \frac{1}{2}$) and WL (which results in $\pi_1 < \pi_2$ for $w_1 = w_2 = \frac{1}{2}$).

If, on the other hand, the finance minister also maximizes total antipoverty spending, so the preferences of the technical committee and finance minister are aligned, optimal weights (with equal costs) are given by $\pi_1^* (r_1, r_2) = \left(1 + \sqrt{\frac{r_1}{r_2}\frac{\beta_2}{\beta_1}}\right)^{-1}$, from (2.14) with $c_1 = c_2$. As seen in column 4 of Table 2.1, since $\beta_1 < \beta_2$ and the committee anticipates that the benevolent finance minister will "cooperate" by shifting resources away from the health minister if the committee sets $\pi_1 < \pi_2$ (thus making achieving prestige harder), health gets a higher weight. In other words, the dimension in which the minister has a higher propensity for private spending is given a *higher* weight and *lower* resources, which makes life harder for the ministers by making achieving prestige harder. In equilibrium, the higher cost of achieving prestige leads to higher overall spending.

2.4.3 Illustration with Unequal Costs

We now illustrate how the results from Section 2.4.2 change when the two ministers have differing effectiveness of antipoverty spending $1/c_j$ (or, equivalently, differing marginal cost c_j of reducing deprivations). If $c_1 < c_2 = 1$, a prestige-maximizing committee will still assign all of the weight to income, maintaining a unidimensional measure. A benevolent committee anticipating a prestige-maximizing finance minister will still neuter the finance minister by setting $\pi_1 = \frac{c_1}{c_1+c_2}$, but now that $c_1 < c_2$ this implies lowering the weight assigned to dimension 1. A benevolent committee anticipating a benevolent finance minister will also decrease π_1, by (2.14). For example, with $c_1 = 0.6$, $c_2 = 1$, and using the MWL, $\pi_1^* = 0.220$, compared to 0.320 in Table 2.1, column 4. The intuition is again that the committee makes it *hard* to acheive prestige, lowering the weight assigned to the minister who has the easier job of reducing deprivations in her dimension.

If instead $c_1 > c_2$, the changes relative to the equal weights case are more interesting, since the higher marginal cost in dimension 1 works in the opposite direction of minister 1's lower propensity for private consumption ($\beta_1 < \beta_2$). Table 2.2 shows the results for the same parameters as in Section 2.4.2 but with $c_1 = 1$, $c_2 = 0.6$. A prestige-maximizing technical committee still prefers a unidimensional measure (as guaranteed by Proposition 7), but now that minister 2 reduces deprivations by more for each dollar spent—which could be due to lower initial deprivations in that dimension or to it being less expensive to provide access to health insurance than to transfer sufficient income to the poor for their incomes to surpass the income poverty line—the technical committee adopts a measure that is unidimensional in health (as seen by $\pi_2^* = 1$ and hence $w_2^* = 1$ in columns 1 and 2). Although, in practice, it is unlikely that a government would adopt a non-income unidimensional poverty measure, it is worth noting that some of the officially adopted multidimensional measures do assign zero weight to income,[26] possibly motivated by the proven difficulty of reducing

[26] Examples include the official multidimensional poverty measures in Chile (Herrera, 2014) and Colombia (Angulo, 2011).

poverty in this dimension by raising incomes above the income poverty line, which would be consistent with this result from our model when c_1 is relatively high.

Table 2.2: Illustration with Mexican Data and Unequal Costs

Technical committee maximizes	P^*		$x_1^* + x_2^*$	
Finance minister maximizes	P^*	$x_1^* + x_2^*$	P^*	$x_1^* + x_2^*$
Panel A: Minimum well-being line				
$\mu(g_1^0)$	0.194	0.194	0.194	0.194
$\mu(g_2^0)$	0.292	0.292	0.292	0.292
π_1^*	0.000	0.000	0.625	0.440
π_2^*	1.000	1.000	0.375	0.560
w_1^*	0.000	0.000	0.715	0.542
w_2^*	1.000	1.000	0.285	0.458
r_1^*	0.300	0.300	any	0.600
r_2^*	0.700	0.700	any	0.400
x_1^*	0.000	0.000	0.529	0.497
x_2^*	0.538	0.538	0.186	0.254
$x_1^* + x_2^*$	0.538	0.538	0.714	0.750
M_0^*	0.292	0.292	0.222	0.239
$M_0'^*$	0.135	0.135	0.133	0.153
P^*	0.897	0.897	0.446	0.455
Panel B: Well-being line				
$\mu(g_1^0)$	0.520	0.520	0.520	0.520
$\mu(g_2^0)$	0.292	0.292	0.292	0.292
π_1^*	0.000	0.000	0.625	0.440
π_2^*	1.000	1.000	0.375	0.560
w_1^*	0.000	0.000	0.484	0.306
w_2^*	1.000	1.000	0.516	0.694
r_1^*	0.300	0.300	any	0.600
r_2^*	0.700	0.700	any	0.400
x_1^*	0.000	0.000	0.529	0.497
x_2^*	0.538	0.538	0.186	0.254
$x_1^* + x_2^*$	0.538	0.538	0.714	0.750
M_0^*	0.292	0.292	0.402	0.362
$M_0'^*$	0.135	0.135	0.241	0.232
P^*	0.897	0.897	0.446	0.455

We now consider the case of a benevolent committee. For any weights other than those given by (2.14), a prestige-maximizing finance minister would still seek to assign as many resources to the dimension whose spending has a higher impact per dollar on measured multidimensional poverty. A benevolent finance minister, on the other hand, would distribute resources to the minister with a *higher* cost of reducing deprivations, the social development minister in this illustration.

2.4.4 Non-optimal Weights and Budget Allocations

How do our results change if weights are set non-optimally (e.g., by arbitrarily assigning equal absolute weights to each dimension)? And how do they change if the finance minister allocates resources across ministers in a way that maximizes neither prestige nor antipoverty spending? We explore these questions in this section using the illustration from Mexico. The main take-away is that even when weights and budget allocations are chosen in ways that differ from those in our model, total antipoverty spending is usually higher when a multidimensional index is used.

As in Section 2.4.3, we set $\beta_1 = 0.1$, $\beta_2 = 0.3$, $R = 1.0$, $\underline{r}_1 = 0.3$, $\underline{r}_2 = 0.4$, $c_1 = 1.0$, and $c_2 = 0.6$. Figure 2.3 shows total antipoverty spending as a function of π_1, for three possible budget allocations: that of a benevolent finance minister, that of a prestige-maximizing finance minister, and the actual distribution of resources between the Ministry of Social Development and Ministry of Health in 2010.[27] It is worth noting that, since the vertical axis measures total antipoverty spending (which is a benevolent finance minister's maximand), the solid black curve for total antipoverty spending under a benevolent finance minister is the upper envelope of the set of curves for each possible resource allocation chosen by the finance minister. The kink in the solid black and dashed gray curves occurs because at that value of π_1, a benevolent finance minister switches from assigning $r_1 = \bar{r}_1 = 0.6$ (for values of π_1 such that $\frac{\pi_1}{c_1} < \frac{\pi_2}{c_2}$) to $r_1 = \underline{r}_1 = 0.3$ (for values of π_1 such that $\frac{\pi_1}{c_1} > \frac{\pi_2}{c_2}$), while a prestige-maximizing finance minister makes the opposite switch (Proposition 6).

[27] The actual budget of Mexico's Ministry of Social Development in 2010 was 80,176,891,338 pesos, while the budget of the Ministry of Health was 89,892,930,927 pesos (Secretaría de Hacienda y Crédito Público, 2010). Normalizing the sum of these to $R = 1$ for this illustration, this gives $r_1 \approx 0.471$, i.e. the Ministry of Social Development's budget equaled 47.1% of the total budget allocated to the two ministries.

A multidimensional poverty measure usually results in higher antipoverty spending than a unidimensional measure, even if the technical committee does not behave optimally. Specifically, we use Figure 2.3 to compare a unidimensional measure in income, i.e. $\pi_1 = 1$, to a multidimensional measure with income and health, i.e. $\pi_1 \in (0, 1)$. For a benevolent or prestige-maximizing finance minister, total antipoverty spending with a unidimensional measure is 0.545, or 54.5% of the total budget $R = 1$ allocated between the two ministers. (The curves for the benevolent and prestige-maximizing finance ministers overlap at $\pi_1 = 1$ because large values of π_1 cause the equilibrium of the game between line ministers to be a corner solution, and recall from Remark 2.3.1 that for corner solutions, the decisions of a prestige-maximizing and benevolent finance minister coincide.) As long as $\pi_1 \in [0.235, 0.803]$, a multidimensional measure results in higher antipoverty spending than a unidimensional measure in income, *regardless of the finance minister's choice*.

Figure 2.3: Total Antipoverty Spending as a Function of Weights

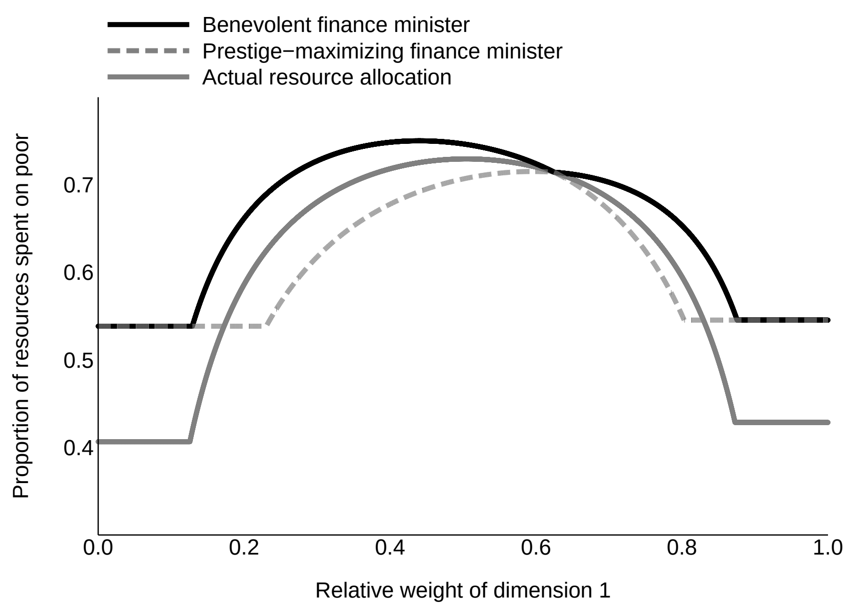

If the finance minister is prestige-maximizing, the range of values of π_1 for which a multidimensional measure results in higher antipoverty spending is the range given above. For $\pi_1 < 0.235$ a multidimensional measure results in lower spending than $\pi_1 = 1$, while for $\pi_1 > 0.803$ it results in the same level of antipoverty spending as $\pi_1 = 1$, since high values of π_1 lead to a corner solution in the game between line ministers. If the finance minister is benevolent, the multidimensional measure leads to higher antipoverty spending than a unidimensional income measure for a wider range of weights (specifically, $\pi_1 \in [0.132, 0.876]$). Again, a very low relative weight on the income dimension ($\pi_1 < 0.132$) result in less antipoverty spending than a unidimensional income measure, while a low relative weight on the health dimension ($\pi_1 > 0.876$) results in the same antipoverty spending as a unidimensional income measure.

The solid gray curve in Figure 2.3, showing total antipoverty spending using Mexico's actual allocation of budgets, allows us to consider one example of non-optimal weights combined with a resource allocation that differs from that of a benevolent or prestige-maximizing finance minister (or equals that of a benevolent or prestige-maximizing finance minister if the interval $[\underline{r}_1, \bar{r}_1]$ is smaller than we have allowed for in this illustration). For this allocation of resources ($r_1 \approx 0.471$), total antipoverty spending with a measure of unidimensional income poverty is 0.429, which is again lower than total antipoverty spending under most multidimensional measures. Specifically, antipoverty spending is larger with a multidimensional measure as long as $\pi_1 \in [0.132, 0.873]$.

Figure 2.3 also illustrates a number of points from Section 2.3. First, with highly unequal weights, we encounter a corner solution where one minister spends minimally on the poor, only spending to approach her dimension-specific target, but not spending additional funds to contribute to a reduction in the multidimensional poverty index (recall that we have normalized this minimal amount of spending to 0). Specifically, the Social Development Ministry spends minimally if $\pi_1 \leq 0.129$ and the finance minister is benevolent, or $\pi_1 \leq 0.230$ and the finance minister is prestige-maximizing; similarly, the Health Ministry spends minimally if $\pi_1 \geq 0.876$ and the finance minister is benevolent, or $\pi_1 \geq 0.803$ and the finance minister is prestige-maximizing. These corner solutions are represented in Figure 2.3 by the portions of the curve that are horizontal.

Second, maximum antipoverty spending is generally acheived with fairly equal relative weights, which again only translate to fairly equal absolute weights if initial deprivations are similar across dimensions. Specifically, maximum antipoverty spending is acheived at $\pi_1 \approx 0.440$ (and thus $\pi_2 \approx 0.560$) if the finance minister is benevolent and $\pi_1 \approx 0.505$ (and thus $\pi_2 \approx 0.495$) using the actual allocation of resources in Mexico in 2010. Both of these values are precisely $\pi_1^*(r_1, r_2)$ from (2.14) for the particular parameters used in this illustration, and for $r_1 = \bar{r}_1 = 0.6$ (the allocation chosen by a benevolent finance minister) and $r_1 \approx 0.471$ (the actual resource allocation). If the finance minister is prestige-maximizing, antipoverty spending is maximized at $\pi_1 = \frac{c_1}{c_1+c_2} = 0.595$ (and thus $\pi_2 = 0.405$).

Third, the three curves intersect at $\pi_1 = \frac{c_1}{c_1+c_2} = 0.625$ (and the curves for any other allocation of resources by the finance minister would also intersect at this point). Hence, at this value of π_1 the finance minister is "neutered" and cannot affect total antipoverty spending by shifting resources around. If the technical committee is seeking to maximize antipoverty spending but it faces a finance minister with a conflicting objective (to maximize prestige), the committee should thus choose $\pi_1 = \frac{c_1}{c_1+c_2}$.

2.5 Conclusion

We use a simple game theoretic framework to analyze the political economy issues that can come into play when poverty is measured multidimensionally. When different deprivations are combined into a scalar measure of poverty, and prestige is jointly bestowed on ministers for improvements in this measure, a contribution to a public good game arises between government ministries in charge of reducing deprivations in different dimensions. Each minister has preferences over antipoverty spending and private spending; the latter can be thought of as corruption or simply other spending that doesn't reach the poor. Each minister's antipoverty spending benefits all ministries because it reduces multidimensional poverty and increases government prestige. The weights assigned to each dimension, the ministers' propensity for private spending, their budgets, and the costs of reducing deprivations in each dimension all affect the ministers' ability to reduce multidimensional poverty and their incentives to free ride on the contributions of other ministers.

A number of interesting results emerge from our model. The allocation of resources across ministries depends critically on the policy objective of the agent assigning budgets. For $\alpha = 0$ (the most common choice by governments in practice since it allows dimensional achievements to be measured ordinally), if this agent is concerned with the government "looking good" in terms of *measured* poverty reduction, he assigns as much budget as possible to the minister whose dimension makes a larger contribution to initial poverty. If, on the other hand, he is "benevolent" and wants to maximize total antipoverty spending by the ministers, he does the opposite, assigning as much budget as possible to the minister whose dimension has lower importance.

A technical committee assigning weights for the measure will give as much weight as possible to one dimension if its objective is to maximize government prestige. In contrast, if the committee's objective is to maximize poverty-reducing effort by the ministers, weights should be chosen so that each dimension's contribution to initial poverty is much more similar. Exact values of the weights depend on the ministers' budgets, propensities for private spending, and the preferences of the finance minister.

Our study has a number of limitations, and is only a first attempt at understanding the complicated political economy issues at play when poverty is measured multidimensionally. Most obviously, our model of the political process is extremely simplified. Agents may be able to cooperate. Nevertheless, as long as multiple agents are responsible for the dimensions of a scalar measure of multidimensional poverty and these agents cannot write complete and enforceable binding contracts with one another, then a public good problem among policymakers arises to some extent. We posit that coalition governments in which coalitions are not very stable might be especially susceptible to free riding concerns.

We make a number of simplifying assumptions regarding the measurement of multidimensional poverty as well. First, we restrict the analysis to the Alkire and Foster (2011a) class of multidimensional poverty measures, which do not allow for varying complementarity or substitutability between deprivations in different dimensions. It has been argued that for this reason Bourguignon and Chakravarty's (2003) class of measures are "probably better" if the number of dimensions under consideration is small (Silber, 2011, p. 480).[28] Nevertheless, we view this restriction as justified on positive grounds, since the AF specification is the one being adopted by governments around the world, as well as by the United Nations Development Program for over 100 countries in their 2010 Human Development Report (UNDP, 2010; see also Alkire and Santos, 2014). Second, we focus on the simplified case in which the multidimensional poverty index has two dimensions. While this is unlikely to be the case in practice, it is a common simplification[29] that allows us to gain intuition about the relevant interactions between ministries. Furthermore, we show in the Appendix that the main forces identified in our framework with two dimensions and line ministers carry through to a model with three line dimensions and line ministers; further generalizations to n dimensions and line ministers are immediate. Third, our two-dimensional AF measure uses the union approach to identify the poor: an individual is poor if she is deprived in at least one dimension. Using the intersection approach, where an individual is poor only if she is deprived in both dimensions, should only increase the incentives for ministers to free ride because the other minister's actions can completely eliminate an individuals' poverty status. When the model is extended beyond two dimen-

[28] Other characteristics of Alkire and Foster's (2011a) measure have been criticized as well. One is that under the strong focus axiom, an increase in the attainment of a poor individual in one of her non-deprived dimensions does not reduce poverty, while it may be the case that attainments in non-deprived dimensions can, to some degree, be substituted for deprivations in other dimensions (Permanyer, 2014). The AF class is also not continuous in its arguments at each deprivation cut-off (except when the union approach is used to identify the poor) due to the dual cut-off approach. Furthermore, a regressive transfer in a certain dimension can decrease poverty even when $\alpha \geq 1$ if it causes the less poor individual to no longer be deprived in that dimension *and as a result* to no longer be poor (Permanyer, 2014, footnote 9).

[29] For example, Kanbur (1987) assumes two sectors when studying the political economy of unidimensional poverty reduction, and Duclos et al. (2014) assume two dimensions for multidimensional poverty targeting. In the Appendix, we show how to extend our model to more than two dimensions and demonstrate that, for equilibria in which all line minister contribute to poverty reduction, the main issues we identify—in particular the conflict between benevolence and self-interest—remain valid and are resolved in a similar manner as in our two-dimensional setup.

sions and Alkire and Foster's (2011a) dual cut-off approach is used, where a cut-off k is chosen such that an individual is identified as poor if $\sum_j w_j g^0_{ij} \geq k$, there should still be more free riding than under the union approach for the same reason. Fourth, we restrict the analysis to $\alpha = 0$, which is by far the most commonly implemented version of the AF measure in practice; nevertheless, the choice of α can have nuanced implications for antipoverty spending.[30]

We also make a number of simplifying assumptions regarding the preferences of the ministers and the poverty-reducing implications of their spending. First, we restrict the ministers' preference functions to be Cobb-Douglas. The dynamics are similar and no qualitative changes arise if utility is instead CES, but it is worth noting that if instead they are quasilinear in multidimensional poverty reduction, the amount allocated to poverty reduction by minister 2 does not enter into the FOC of minister 1's maximization problem and vice versa, so the political economy dynamics of using a multidimensional poverty measure become less interesting. Second, we assume that antipoverty spending in each dimension causes a proportional reduction in the number of individuals deprived in that dimension. If, instead of assuming this constant marginal product technology for poverty reduction, we considered a more realistic concave production function, corner solutions would be less common. Third, we assume away "spillovers" across dimensions: spending by minister j reduces deprivations in dimension j only.[31] Including such spillover effects would affect the details of our analysis, but would only exacerbate the free riding incentives we analyze. Relaxations of our simplifying assumptions are the subject of ongoing research.

[30] We address these implications in a companion paper. Increasing α increases complementarities across dimensions which reduces the potential to free ride and increases antipoverty spending, but also exacerbates initial asymmetries in relative dimension weights, reducing antipoverty spending. Which effect dominates depends on the breadth and depth of deprivations across dimensions: if the distribution of deprivations is similar across dimensions, the latter effect will be minimal and the former effect will dominate, leading to higher antipoverty spending.

[31] This ignores, for example, the potential of increased income to reduce health deprivations (e.g., Pritchett and Summers, 1996) or food insecurity (Sen, 1981) and the causal impact of schooling on health (Conti et al., 2010).

A number of policy implications emerge. To reduce free riding and increase overall antipoverty spending, governments can use a measure that satisfies the dimensional decomposability axiom and prominently publish the partial indices for each dimension in conjunction with the overall index. Our line of reasoning reinforces what has been advocated by proponents of the AF measure (e.g., Alkire et al., 2011) and what is indeed done by some governments (e.g., Colombia and Mexico). This also stresses the importance of the dimensional decomposability axiom, which is not met by some measures proposed in the literature. Nevertheless, using the dimensional decomposition of AF does not completely eliminate the opportunity to free ride; this result is not an artefact of our model: AF is based on the matrix of *censored* deprivations, where $g_{ij} = 0$ for all i if j is non-poor, i.e., i can potentially be deprived in some dimensions but not in enough dimensions to be considered poor. As a result, spending by minister j can change individual i's poverty status from poor to non-poor, which will then reduce $\mu(g_k^0)$ for $k \neq j$ and $y_{ik} < z_k$, even if minister k spends nothing on reducing deprivations.

Since relative weights depend on both the absolute weights and the vectors of deprivations, the sometimes arbitrary selection of weights and cut-offs can have important implications for antipoverty spending. If the objective is to maximize antipoverty spending and ministers have similar budgets, propensities for private spending, and costs of reducing deprivations, these should be chosen so that each dimension initially contributes to an approximately equal proportion of total multidimensional poverty. Finally, the policy objectives of a "benevolent" government maximizing total antipoverty spending and a "self-interested" government maximizing prestige (i.e., *measured* poverty reduction) often conflict, which is an important result to keep in mind when evaluating the process used to select the multidimensional poverty measure's parameters or the allocation of budgets across ministries.

Chapter 3

Banking on Trust: How Debit Cards Enable the Poor to Save More

Virtually every commercial transaction has within itself an element of trust.... It can be plausibly argued that much of the economic backwardness in the world can be explained by the lack of mutual confidence.

—Kenneth Arrow (1972)

Abstract

Trust is an essential element of economic transactions, but trust in financial institutions is especially low among the poor, which may explain in part why the poor do not save formally. Debit cards provide not only easier access to savings (at any bank's ATM as opposed to the nearest bank branch), but also a mechanism to monitor bank account balances and thereby build trust in financial institutions. We study a natural experiment in which debit cards were rolled out to beneficiaries of a Mexican conditional cash transfer program, who were already receiving their transfers in savings accounts through a government bank. Using administrative data on transactions and balances in over 300,000 bank accounts over four years, we find that after receiving a debit card, the transfer recipients do not increase

their savings for the first 6 months, but after this initial period, they begin saving and their marginal propensity to save increases over time. During this initial period, however, they use the card to check their balances frequently; the number of times they check their balances decreases over time as their reported trust in the bank increases. Using household survey panel data, we find the observed effect represents an increase in overall savings, rather than shifting savings; we also find that consumption of temptation goods (alcohol, tobacco, and sugar) falls, providing evidence that saving informally is difficult and the use of financial institutions to save helps solve self-control problems.

3.1 Introduction

Trust is an essential element of economic transactions and an important driver of economic development (Banfield, 1958; Knack and Keefer, 1997; Porta et al., 1997; Narayan and Pritchett, 1999; Algan and Cahuc, 2010). Trust is the "subjective probability with which an agent assesses that another . . . will perform a particular action" (Gambetta, 1988, p. 217). It is particularly important in financial transactions where people pay money in exchange for promises, and essential where the legal institutions that enforce contracts are weak (McMillan and Woodruff, 1999; Karlan et al., 2009). Given the nature of financial decisions, it is not surprising that trust has been shown to be key to stock market participation (Guiso et al., 2008), use of checks instead of cash (Guiso et al., 2004), and decisions to not withdraw deposits from financial institutions in times of financial crisis (Iyer and Puri, 2012; Sapienza and Zingales, 2012).

Trust in financial institutions, however, is low as evidenced by the fact that majorities in 40 percent of countries included in the World Values Survey report lack of confidence in banks (Figure 3.1). Trust is especially low among the poor. In Mexico, for example, 71% of those with less than primary school report low trust in banks, compared to 55% of those who completed primary school and 46% of those who completed university (Figure 3.2). Along with fees and minimum balance requirements, trust is frequently listed as a primary reason for not saving in formal bank accounts (e.g., Dupas et al., 2016). At the country level, low trust in financial institutions is strongly correlated with the proportion of the population without bank accounts (Figure 3.3). Despite its importance, trust as a potential barrier to the poor saving in financial institutions has not been extensively studied (Karlan et al., 2014).[1]

Figure 3.1: Low Trust in Banks Around the World

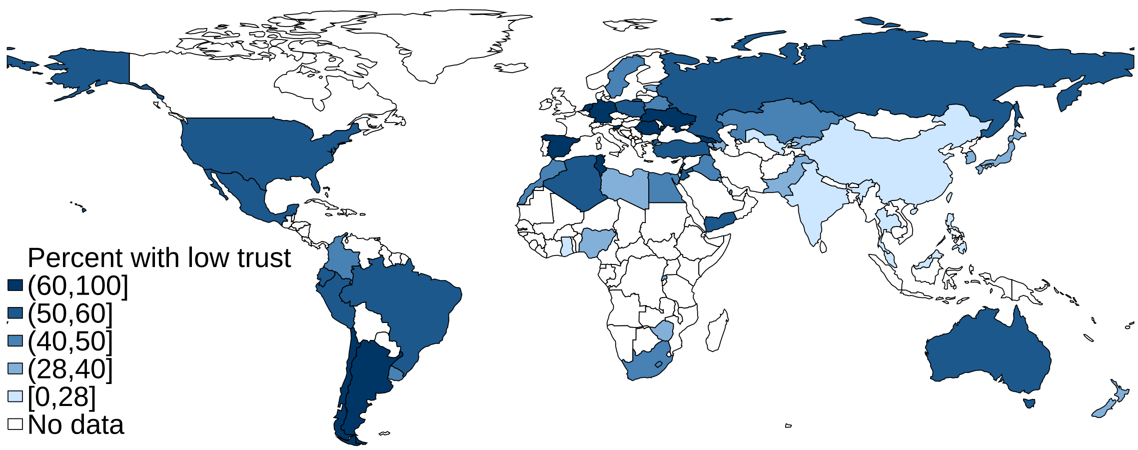

Source: World Values Survey, Wave 6 (2010–2014).
Notes: $N = 82{,}587$ individuals in 60 countries. Low trust in banks is defined as "not very much confidence" or "none at all" for the item "banks" in response to the following question: "I am going to name a number of organizations. For each one, could you tell me how much confidence you have in them: is it a great deal of confidence, quite a lot of confidence, not very much confidence or none at all?" Countries are divided into quintiles, with quintile cut-offs rounded to the nearest percentage point in the legend. Darker shades indicate countries with a higher percent of the population reporting low trust in banks.

Lack of trust in financial institutions may not be unfounded. Cohn et al. (2014) provide evidence that the banking industry fosters a culture of dishonesty relative to other industries. Bankers in Mexico have been found to loot money by directing a large portion of bank lending to "related parties," i.e. shareholders of the bank and their firms (La Porta et al., 2003). Mexican newspapers report many instances of outright bank fraud where depositors have lost their savings. For example, an extensively covered scandal involved Ficrea whose majority shareholder reportedly stole USD 200 million from savers (CNBV, 2014).[2] It is also telling that articles with financial advice in Mexican newspapers have

[1] Increased trust is proposed—but not explored further—as one channel through which no-fee savings accounts led to saving in Prina (2015).

[2] This type of fraud is not uncommon: we scraped the online news archives of all electronic newspapers

Figure 3.2: Low Trust in Banks by Education Level in Mexico

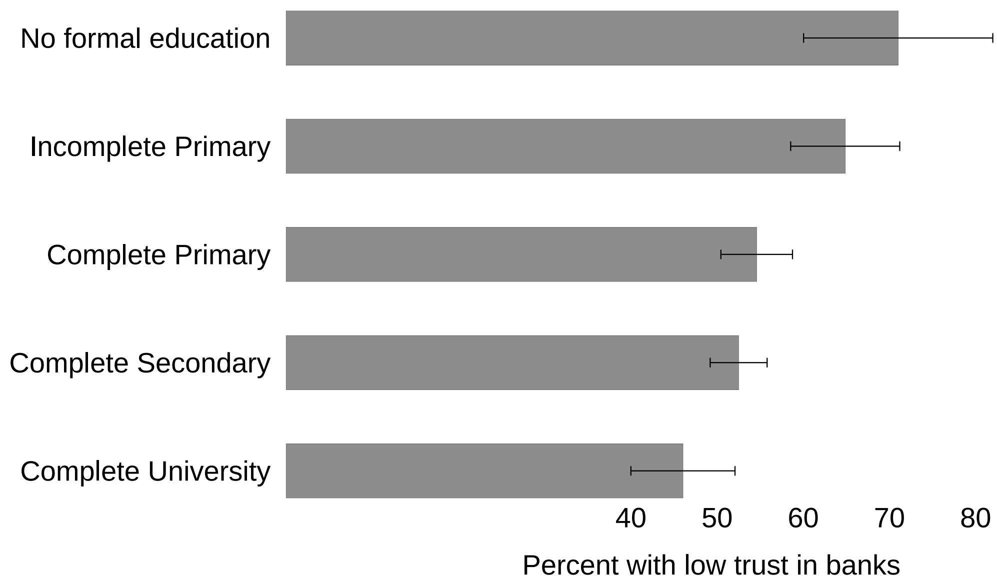

Source: World Values Survey, Mexico, Wave 6 (2012).
Notes: $N = 1993$ individuals. Low trust in banks is defined as "not very much confidence" or "none at all" for the item "banks" in response to the following question: "I am going to name a number of organizations. For each one, could you tell me how much confidence you have in them: is it a great deal of confidence, quite a lot of confidence, not very much confidence or none at all?"

titles like "How to Save for Your Graduation and Avoid Frauds" and "Retirement Savings Accounts, with Minimal Risk of Fraud." When contract enforcement is poor and fraud is rampant, trust becomes even more important (Guiso et al., 2004; Karlan et al., 2009) and people are understandably even more reluctant to use untrustworthy financial institutions (Bohnet et al., 2010).

While trust is important, it is not an innate characteristic but rather can be influenced by experience and information (Hirschman, 1984; Williamson, 1993; Attanasio et al., 2009). Debit cards (and mobile money) provide a low cost technology to monitor account balances and thereby build trust that a bank will neither explicitly steal deposits nor charge unexpectedly large hidden fees. Previous studies on debit cards and mobile money have focused on the effect of the lower transaction costs facilitated by these technologies to make purchases (Zinman, 2009), access savings and remittances (Suri et al., 2012; Schaner, 2015a), and

---

and news websites we could find in Mexico (129 total) using several keywords, then filtered the results by hand to keep only relevant stories. We found 1338 news stories associated with savings fraud in 2014 and 2015 alone.

Figure 3.3: Cross-Country Comparison of Trust in Banks and Saving in Financial Institutions

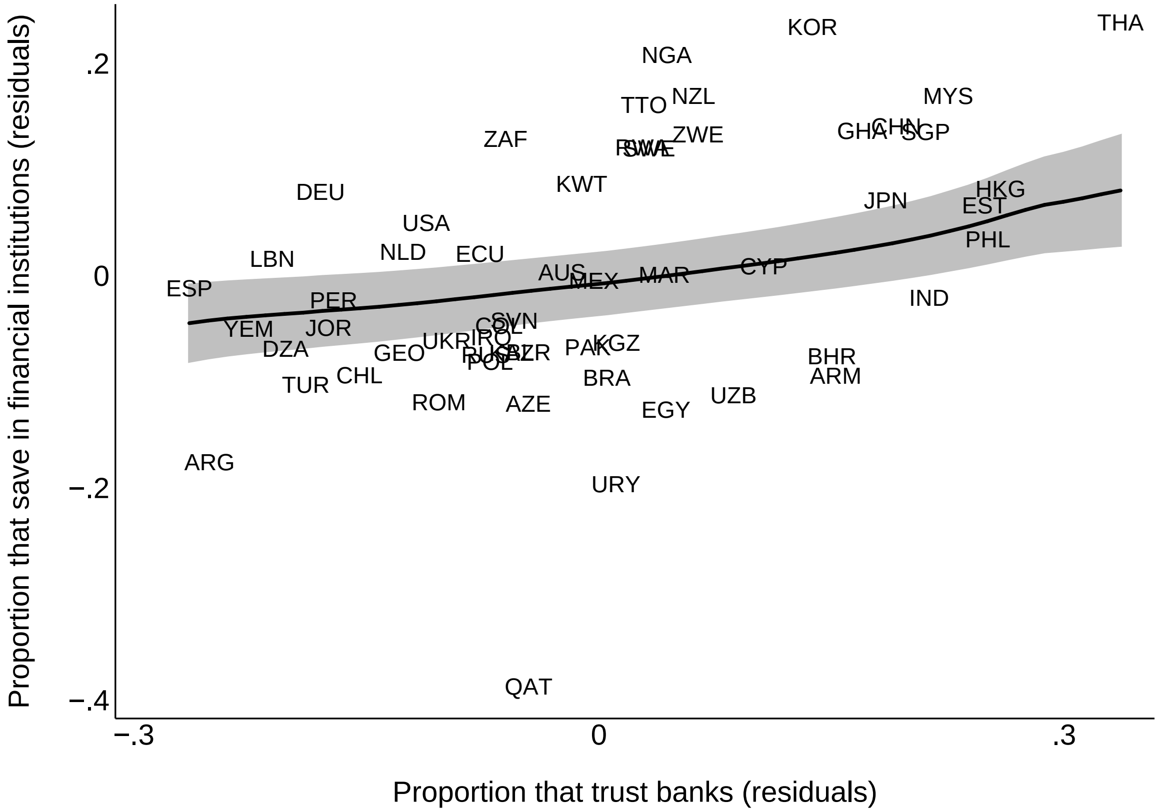

Sources: World Values Survey (WVS), Wave 6 (2010–2014); Global Findex; World Development Indicators (WDI).
Notes: $N = 56$ countries. The y-axis plots residuals from a regression of the proportion that save in financial institutions (from Global Findex) against controls (average age, education, and perceived income decile from WVS, GDP per capita and growth of GDP per capita from WDI). The x-axis plots residuals from a regression against the same controls of the proportion that respond "a great deal of confidence" or "quite a lot of confidence" in response to the WVS question "I am going to name a number of organizations. For each one, could you tell me how much confidence you have in them: is it a great deal of confidence, quite a lot of confidence, not very much confidence or none at all?" The solid line shows a kernel-weighted local polynomial regression, while the gray area shows its 95% confidence interval.

transfer money (Jack et al., 2013; Jack and Suri, 2014), but not their capacity to monitor and build trust in financial institutions. We hypothesize that new debit card clients first use the cards to check balances and thereby establish trust, after which they take advantage of the cards' lower transaction costs to use the services of formal financial institutions. In this sense, we argue that building trust in a financial institution is a necessary condition for the use of formal financial services; i.e., financial inclusion requires trust. Indeed, a lack of trust could explain why a number of randomized field experiments have found that even when take-up of accessible and affordable formal savings products is high, use is low in that most opened accounts have few transactions after the first 6 to 12 months (Ashraf et al., 2006; Dupas and Robinson, 2013a; Karlan and Zinman, 2014; Schaner, 2015b).

We examine this hypothesis in the context of a natural experiment in which debit cards were rolled out geographically over time to beneficiaries of the Mexican conditional cash transfer program Oportunidades. The beneficiaries had been receiving their transfers into savings accounts for five years on average before debit cards were attached to their accounts, but typically did not use the accounts to save as they immediately withdraw most if not all of the transfer.[3] The phased geographic rollout provides plausibility exogenous variation in assignment of debit cards to beneficiaries in a difference in difference context. For the analysis, we use high frequency administrative data on bank transactions for over 340,000 beneficiary accounts in 370 bank branches over 4 years as well as several household surveys of a sample of the same beneficiaries.

Using the high frequency administrative data, we find that beneficiaries initially used debit cards to check account balances without any increase in savings, but over time the frequency of account balance checks fell and savings rates rose. We estimate that after one year, the share of total income saved each payment period increased by 5 percentage points and that after nearly two years those with cards saved 8 percentage points more per period.

The delayed initiation of savings suggests some kind of learning. We explore three kinds of learning that may be occurring: (i) learning to trust the bank, (ii) learning to use the debit cards and ATMs, and (iii) learning that the program will not drop beneficiaries who accumulate savings. Using household survey data, we find support for the "learning to trust" hypothesis but not for the other two types of learning. Specifically, we find that 27 percent of beneficiaries who have had the debit card for less than 6 months report that they do not trust the bank, compared to just 17 percent of those who have had the card for

[3] This is consistent with findings from other countries such as Brazil, Colombia, and South Africa, in which cash transfers are paid through bank accounts, but recipients withdraw the entire transfer amount each pay period and do not save in the account (Bold et al., 2012).

more than 6 months. We find very few beneficiaries who report not knowing how to use the technology or fear the program will drop them if they accumulate savings, and no change over time comparing those that have had the debit card less than and more than 6 months. We also find that those who have had the card more than 6 months report checking their balances significantly less frequently than those who have had the card less than 6 months, consistent with our finding from administrative data that when beneficiaries first get the debit card, they check their balances often, but the frequency of checking falls over time.

We then test whether the increase in the bank account balances is an increase in total savings or a substitution from other forms of saving, both formal and informal. Using panel household survey data, we find that after one year the treatment group increases total savings by about 5 percent of income relative to the control group, which is close in magnitude to the effect we see in the administrative account data. We find no differential change in income or assets in the treatment group compared to the control. These results suggest that the increase in saving is not driven by higher income but by (voluntarily) lowering current consumption and that that the increase in bank savings does not crowd out other forms of saving (consistent with Ashraf et al., 2015; Dupas and Robinson, 2013a; Kast et al., 2012).

Finally, the increase in savings in achieved through a decrease in the consumption of alcohol, tobacco, and sugar—the most frequently mentioned temptation goods in Banerjee and Mullainathan (2010)—and transportation. Although the poor do save via cash at home (Collins et al., 2009), saving informally is harder difficult as "[cash] money is hot" and susceptible to temptation spending, either by the beneficiary herself or by her husband if she lacks control over his access to her savings (Ashraf, 2009). Indeed, we also find that among beneficiaries living with a spouse or partner, those with lower baseline bargaining power relative to their spouse have a higher increase in savings after receiving the debit card. Our results suggest that saving in formal financial institutions may help solve some of the intra-household bargaining and self-control problems associated with trying to save informally.

These results are important for public policy as building savings in formal financial institutions has positive welfare effects for the poor and nearly half of the world's adults do not use financial institutions (Demirgüç-Kunt et al., 2015). The poor have used savings products to decrease income volatility (Chamon et al., 2013), accumulate money for microenterprise investments (Dupas and Robinson, 2013a), invest in preventative health products and pay for unexpected health emergencies (Dupas and Robinson, 2013b), and invest in children's education (Prina, 2015). Various randomized experiments have found that providing affordable and accessible savings accounts to the poor increases their future agricultural/ business output and household consumption (Brune et al., 2016; Dupas and Robinson, 2013a), decreases debt (Kast et al., 2012; Atkinson et al., 2013), and improves their ability to cope with shocks (Prina, 2015). For these reasons, Mullainathan and Shafir (2009) conclude that access to formal savings services "may provide an important pathway out of poverty."

Given our results, government cash transfer programs could be a promising channel to increase financial inclusion and enable the poor to save, not only because of the sheer number of the poor that are served by cash transfers, but also because many governments are already embarking on digitizing their cash transfer payments through banks and mobile money. Furthermore, the technologies of debit cards and ATMs or point of sale (POS) terminals—which can be used to check balances and access savings—are simple, prevalent, and potentially scalable to millions of government cash transfer recipients worldwide.

3.2 Institutional Context

We examine the the roll-out of debit cards to urban beneficiaries of Mexico's conditional cash transfer program Oportunidades whose benefits were already being deposited directly into savings accounts without debit cards. Oportunidades is one of the largest and most well-known conditional cash transfer programs worldwide with a history of rigorous impact evaluation (e.g., Gertler, 2004; Parker and Teruel, 2005). The program provides bimonthly cash transfers to poor families in Mexico, seeking to alleviate poverty in the short term and break the intergenerational poverty cycle in the long term by requiring families to invest in the human capital of childen by sending their children to school and having health check-ups. It began in rural Mexico in 1997 under the name Progresa, and later expanded to urban areas starting in 2002. Today, nearly one-fourth of Mexican households receive benefits from Oportunidades (Levy and Schady, 2013).

Oportunidades opened savings accounts in banks for a portion of beneficiaries in urban localities and began depositing the transfers directly into those accounts. The original motives for paying through bank accounts were to (1) decrease corruption as automatic payments through banks lowers both the ability of corrupt local officials to skim off benefits and of local politicians associating themselves with the program through face-to-face contact with recipients when they received their transfers, (2) decrease long wait times for recipients who previously had to show up to a "payment table" on a particular day to receive their benefits, (3) decrease robberies and assaults of program officers and recipients transporting cash on known days, and (4) increase the financial inclusion of poor households. By the end of 2004, over one million families received their benefits directly deposited into savings accounts in Bansefi, a government bank created to increase savings and financial inclusion of underserved populations (Figure 3.4).[4]

The Bansefi savings accounts have no minimum balance requirement or monthly fees and pay essentially no interest.[5] Before the introduction of debit cards, beneficiaries could only access their money at Bansefi bank branches. Because there are only about 500 Bansefi branches nationwide, many beneficiaries live far from their nearest branch meaning that accessing their accounts involved large transaction costs for many beneficiaries. Over-

[4] Originally Oportunidades partnered with two banks: Bansefi and Bancomer, a commercial bank. However, working with a commercial bank proved to be difficult, and Oportunidades phased out the Bancomer accounts and transferred them to Bansefi by mid-2006.

[5] Nominal Interest rates were between 0.09 and 0.16 percent per year compared to an inflation of around 5 percent per year during our sample period.

Figure 3.4: Timing of Roll-out and Data

(a) Administrative Bank Account Data

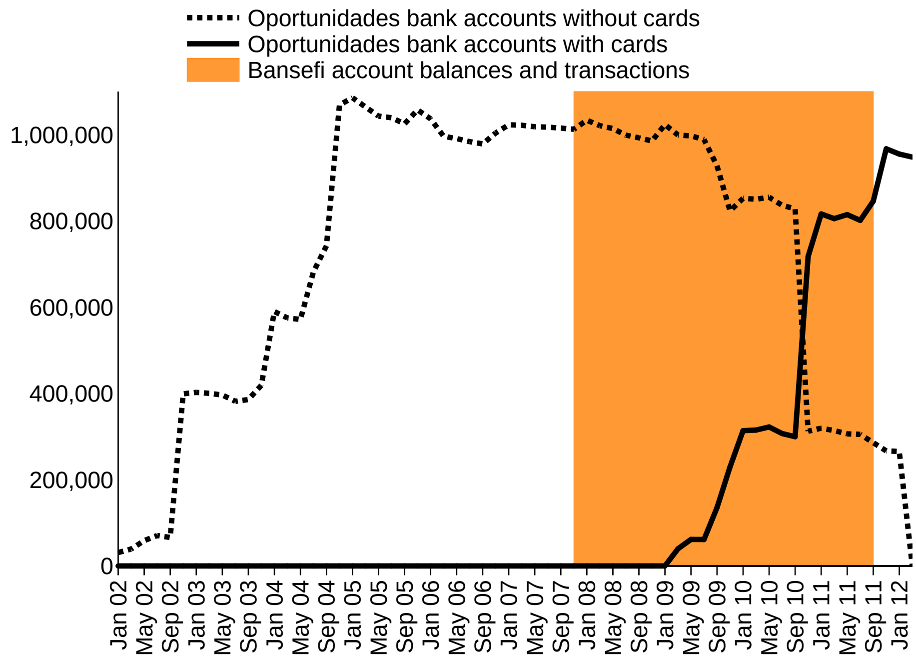

(b) Household Survey Data

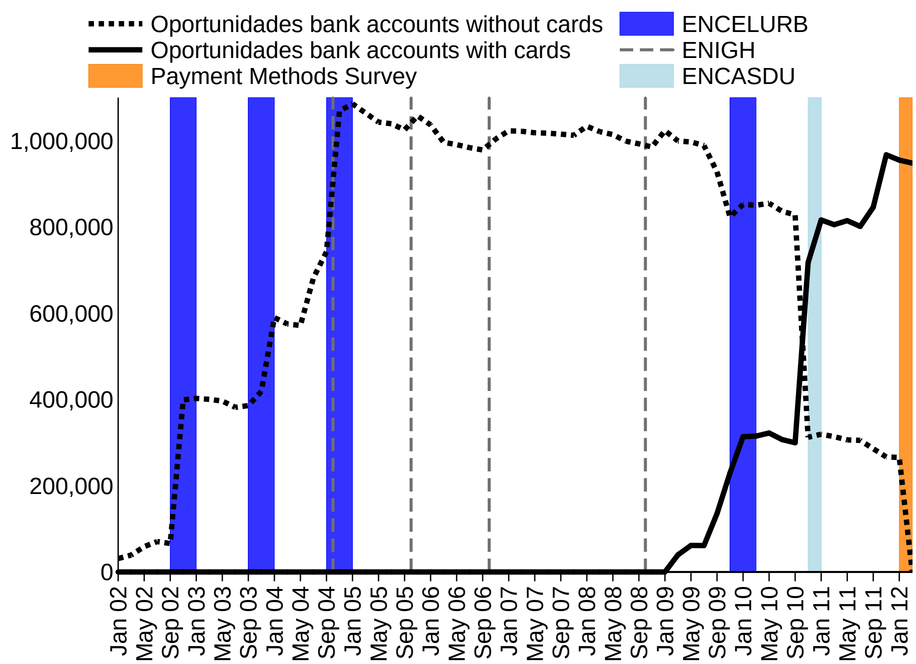

Source: Number of Oportunidades bank accounts with cards and without cards by bimester is from administrative data provided by Oportunidades.

all, the savings accounts were barely used prior to the introduction of debit cards. In 2008, the year before the rollout of debit cards, the average number of deposits per bimester[6] was 1.05 including the deposit from Oportunidades, the average number of withdrawals was 1.02, and 98.9 percent of the transfer was taken during the first withdrawal following payment.

In 2009, the government announced that they would issue Visa debit cards to beneficiaries that were receiving their benefits directly deposited into Bansefi savings accounts. The cards enabled account holders to withdraw cash from, make deposits into, and check balances of their account at any bank's ATM as well as make electronic payments at any store accepting Visa. The cards included two free ATM withdrawals and every bimester at any bank's ATM, after which ATM withdrawal fees averaged 13 pesos (about $1 using 2009 exchange rates) but varied by bank.

Opportunities used direct deposit into savings accounts for its beneficiaries in 275 out of Mexico's 550 urban localities. Of these, debit cards were rolled out to approximately 100,000 beneficiaries in 143 localities in 2009 (wave 1) and to an additional 75,000 beneficiaries in XXX localities in late 2010 (wave 2). Another 170,000 beneficiaries in the remaining localities were scheduled to receive cards between November 2011 and February 2012 (control group) after the end date of our data period. The map in Figure 3.5 shows that the treatment and control waves had substantial geographical breadth and that some treatment and control localities were physically close.

The sequence with which localities switched was determined as a function of the proportion of households in the locality that were eligible for the program but were not yet receiving benefits. This is because the introduction of debit cards to existing recipients was coupled with an effort to incorporate more beneficiaries. Table 3.1 compares the means of locality-level variables and account-level variables from the control, wave 1, and wave 2 localities using data from the population census from 2005, poverty estimates from Opor-

[6] The program is paid in two-month intervals, which we refer to throughout the paper as bimesters. (The Spanish word *bimestre* is more common than its English cognate, and is used by Bansefi and Oportunidades.)

Figure 3.5: Geographic Coverage and Expansion of Debit Cards

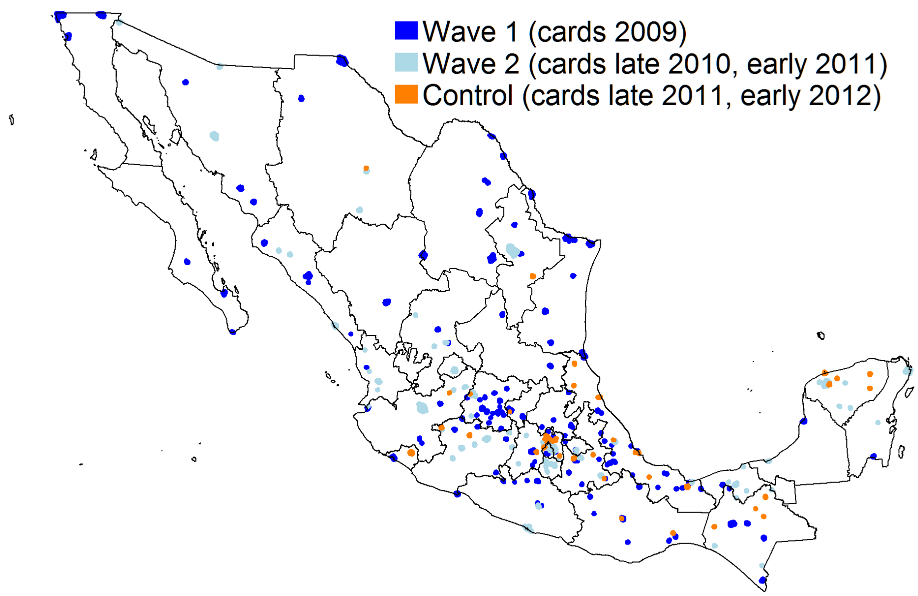

Sources: Administrative data from Oportunidades on timing of debit card receipt by locality and shape files from INEGI.
Notes: $N = 275$ localities (44 in control, 143 in wave 1, 88 in wave 2). The area of each urban locality included in the study is shaded according to its wave of treatment. Urban localities that were not included in the Oportunidades program at baseline or were included in the program but did not pay beneficiaries through Bansefi savings accounts are not included in the figure or in our study.

tunidades from 2005, Bansefi branch locations from 2008, and the administrative account data on average balances and transactions from Bansefi in 2008. Column 6 shows the p-value of an F-test of equality of means. Because the roll-out was not random, it is not surprising that there are some differences across treatment and control localities: treatment localities are slightly larger and beneficiaries in these localities receive higher transfer amounts. The percent of the transfer withdrawn also differs (it is lower in wave 1 than the control and insignificantly different but with a higher point estimate in wave 2), but is high in all cases (ranging from 97.5 percent to 99.6 percent of the transfer), indicating very low savings in the account prior to receiving the card. In Sections 3.4 and 3.9.1, we will test and show that trends of saving, income and consumption were parallel across waves.

Table 3.1: Comparison of Baseline Means

Variable	Control	Wave 1	Wave 2	Diff. W1–C	Diff. W2–C	F-test p-value
Panel A: Locality-level data						
Log population	10.57	11.18	11.48	0.60***	0.91***	0.000***
	(0.11)	(0.10)	(0.16)	(0.14)	(0.19)	
Bansefi branches per 100,000	1.27	1.23	1.58	–0.03	0.32	0.411
	(0.28)	(0.13)	(0.23)	(0.30)	(0.36)	
% HHs in poverty	15.93	13.20	12.23	–2.73	–3.71*	0.177
	(1.67)	(0.75)	(1.09)	(1.82)	(1.99)	
Occupants per room	1.18	1.11	1.12	–0.07	–0.06	0.260
	(0.04)	(0.01)	(0.02)	(0.04)	(0.04)	
Number of localities	44	143	88			
Panel B: Administrative bank account data						
Average balance	581.25	670.32	614.29	89.07	33.05	0.112
	(12.46)	(56.24)	(21.26)	(55.33)	(23.95)	
Number of deposits	1.06	1.05	1.06	–0.02	–0.01	0.907
	(0.01)	(0.04)	(0.03)	(0.04)	(0.03)	
Size of transfer	1506.55	1809.50	1761.26	302.96***	254.71***	0.000***
	(12.73)	(20.16)	(17.47)	(23.67)	(21.15)	
Number of withdrawals	1.03	1.01	1.02	–0.01	–0.01	0.757
	(0.01)	(0.03)	(0.02)	(0.03)	(0.02)	
Percent withdrawn	98.56	97.50	99.64	–1.06**	1.08	0.021**
	(0.18)	(0.45)	(0.71)	(0.46)	(0.72)	
Years with account	5.31	5.49	5.21	0.17	–0.10	0.510
	(0.08)	(0.15)	(0.25)	(0.17)	(0.26)	
Number of accounts	97,922	73,070	171,717			

Sources: Census (2005), Bansefi branch locations (2008), poverty estimates from Oportunidades (based on 2005 Census), timing of card receipt by locality from Oportunidades, and administrative data from Bansefi. Notes: W1 = wave 1, W2 = wave 2, C = control, Diff. = difference. For the administrative data from Bansefi, baseline is defined as January 2009 to October 2009 (prior to any accounts receiving cards in the data from Bansefi).

3.3 Data

We use a rich combination of administrative and survey data sources. To examine the effect of rollout of the debit cards on savings we use administrative data from Bansefi at the account level for 342,709 accounts at 380 Bansefi branches for a four-year period, from November 2007 to October 2011. These data include the bimonthly transfer amount the timing and amount of transactions made in the account, bimonthly average savings balances, the date the savings account was opened, and the month the card was awarded to the account holder. The average account had been opened 5.3 years before getting the card.

To test whether the delayed savings effect and increasing propensity to save over time can be explained by learning to use the technology, learning the program rules, or building trust in the bank, we use the Survey of Urban Households' Sociodemographic Characteristics (ENCASDU), conducted by Oportunidades at the end of 2010. We also use the Payment Method Survey, a household survey conducted by Oportunidades in 2012 aimed at eliciting satisfaction with and use of the debit cards.

To explore whether the increased savings in the Bansefi accounts is an increase in overall savings or a substitution from other forms of saving, we use Survey of Urban Household Characteristics (ENCELURB), a panel survey with three pre-treatment waves in 2002, 2003, and 2004, and one post-treatment wave conducted from late 2009 to early 2010. This survey has comprehensive modules on consumption, income, and assets. We merge these data with administrative data from Oportunidades on the transfer histories for this sample—which we use to add transfer income into total income and to identify which households are Oportunidades recipients, given the common misreporting of transfer receipt in surveys (Meyer et al., 2015)—and on the dates that debit cards were distributed in each locality.

Because the final pre-treatment wave of ENCELURB in 2004 is five years prior to wave 1 of the debit card roll-out, we supplement our parallel trends test in ENCELURB with data for the intervening period (2004-2008) from the Encuesta Nacional de Ingresos y Gastos de Hogares (National Household Income and Expenditure Survey; ENIGH), a repeated cross-section; we merge the publicly available ENIGH with restricted-access locality identifiers provided by the Instituto Nacional de Estadística y Geografía (National Institute of Statistics and Geography; INEGI) to determine which surveyed households were in treatment and control localities, and restrict the analysis to the poorest 20 percent of surveyed households to proxy for Oportunidades recipients.

Figure 3.4a shows the timing of the administrative Bansefi account balance and transaction data, while Figure 3.4b shows the timing of the household survey data (merged with additional administrative data) we use, both relative to the roll-out of debit cards.

3.4 Effect of Debit Cards on Stock of Savings

Figure 3.6 presents average balances over time; even the raw data are very telling. Panel (a) compares the first wave of debit card recipients to the control group, with a dashed vertical line indicating the time when wave 1 localities received debit cards, while Panel (b) compares the second wave to the control, with a dashed vertical line indicating the time when wave 2 localities received debit cards. Strikingly, average balances increase sharply for the first wave after receiving the card, but the effect is not immediate: it begins three to four bimesters after receiving the card and the larger increase happens after a year with the card. By October of 2011, wave 1 has average balances of around 2000 pesos, over three times that of the control group. Average balances also increase over time with the card in wave 2, although we have information for less bimesters after wave 2's later switch to debit cards.

Figure 3.6: Evolution of Average Balances

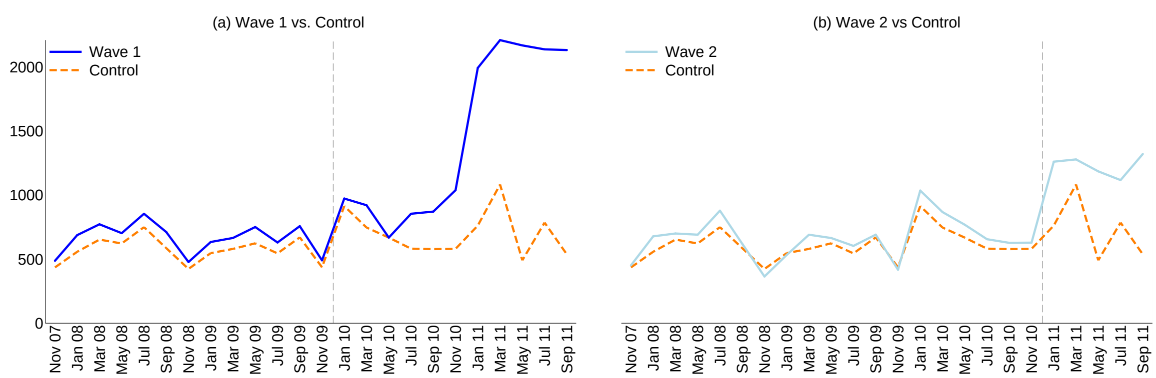

Sources: Administrative data from Bansefi on average account balances by bimester and timing of card receipt.
Notes: N = 5,834,468 account-bimester observations from 343,204 accounts. Average balances are winsorized at the 95th percentile.

Although our data on average balances is by bimester, some payments get shifted to the end of the prior bimester, so we group adjacent bimesters into four-month periods for the remainder of the analysis. Because we have four years of data, this leaves us with 12 four-month periods. To compare the stock of savings in the treatment and control groups while controlling for individual observables and unobservables, as well as any common time shocks, we use a period-by-period difference-in-differences (DID) strategy and estimate:

$$Balance_{it} = \lambda_i + \delta_t + \sum_{k=1}^{12} \phi_k T_{j(i)} \times \mathbb{I}(t = k) + \varepsilon_{it} \tag{3.1}$$

where $Balance_{it}$ is the average balance in account i over period t (specifically, end of day balances were averaged over the number of days in the bimester by Bansefi, and we average the average balances over the two adjacent bimesters that make up the four-month period), λ_i are account level fixed effects which control for observable and unobservable time-invariant characteristics of the beneficiaries, δ_t are time-period dummies that control for general macro trends such as bimester-specific shocks that affect both treatment and control groups, $T_{j(i)} = 1$ if locality j in which account holder i lives is a treatment locality, and $\mathbb{I}(t = k)$ are time period dummies. Thus, $T_{j(i)} \times \mathbb{I}(t = k)$ pick up the difference in balances between treatment and control localities in each period. We estimate cluster-robust standard errors, ε_{it}, clustering by Bansefi branch. Since one time period dummy must be omitted from (3.1), we follow the standard procedure of omitting the four-month period immediately preceding the change to cards. We estimate (3.1) separately for wave 1 and wave 2.

The coefficients of interests are the ϕ_ks, which measure the average difference in balances between the control and treatment group in bimester k. The raw data clearly suggest that pre-treatment trends of savings were parallel across control and treatment groups before getting the card; we test this statistically by testing $\phi_1 = \cdots = \phi_{\ell-1} = 0$ where ℓ is the period of switch. (In wave 1, ℓ is the period November 2009-February 2010, and in wave 2 it is the period November 2010-February 2011.) Figure 3.7 plots the ϕ_ks and shows that pretreatment coefficients are, in most periods, not individually different from zero, and we cannot reject that pre-trends are equal between treatment and control: the p-value for the F-test of $\phi_1 = \cdots = \phi_{\ell-1} = 0$ is 0.823 for wave 1 and 0.110 for wave 2.

Figure 3.7: Difference between Treatment and Control in Average Balances

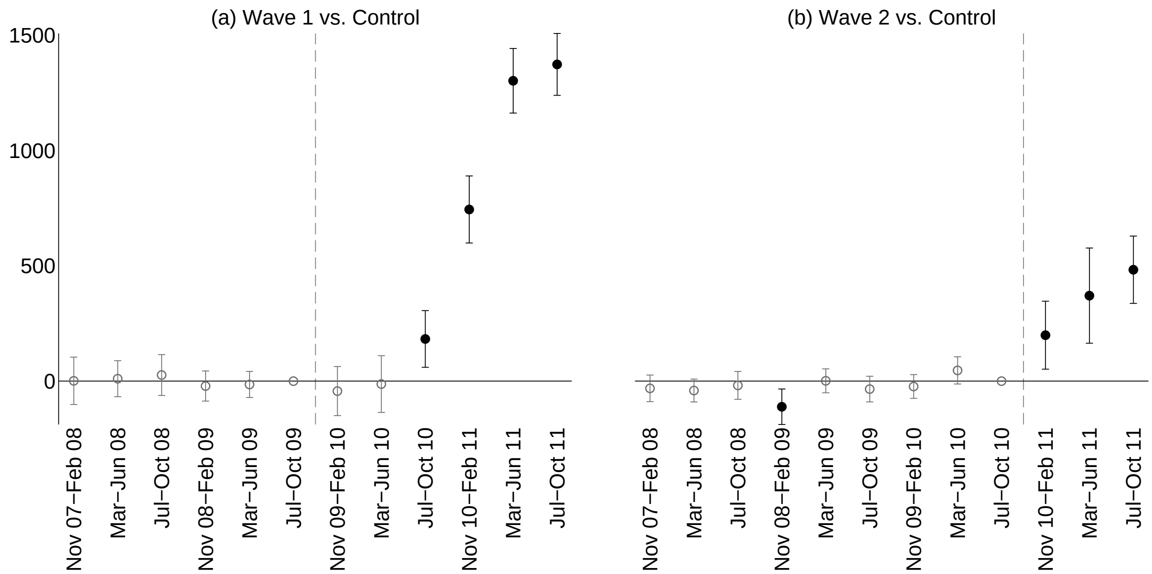

Sources: Administrative data from Bansefi on average account balances by bimester and timing of card receipt.
Notes: (a) N = 2,023,862 from 171,441 accounts. (b) N = 3,086,749 from 270,046 accounts. The figure plots ϕ_k from (3.1). Average balance over each four-month period is the dependent variable, and is winsorized at the 95th percentile. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval.

The cards also led to an increase in use of the accounts, as shown in Figure 3.8, which plots the number of withdrawals per bimester. Prior to receiving the debit card, both the treatment and control groups made about one withdrawal on average. After receiving the card, this increases to about 1.4 withdrawals per bimester. More precisely, Figure 3.9 shows that after receiving the card, 72% of beneficiaries continue to make just one withdrawal, while 22% make 2 withdrawals and 6% make 3 or more withdrawals. This immediate increase in use of the account after a decrease in the transaction costs of accessing money agrees with the prediction of the Baumol (1952) and Tobin (1956) model of money demand in the face of transaction costs, and with empirical evidence that ATMs and debit cards lead to reduced transaction costs and an increased number of withdrawals (Attanasio et al., 2002;

Alvarez and Lippi, 2009).

Figure 3.8: Withdrawals per Bimester

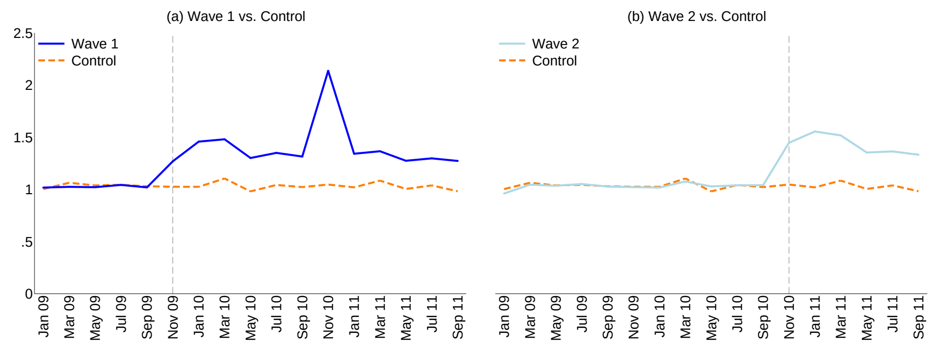

Sources: Administrative data from Bansefi on transactions by quarter and timing of card receipt.
Notes: $N = $ 2,917,234 account-quarter observations from 343,204 accounts.

Figure 3.9: Distribution of Withdrawals

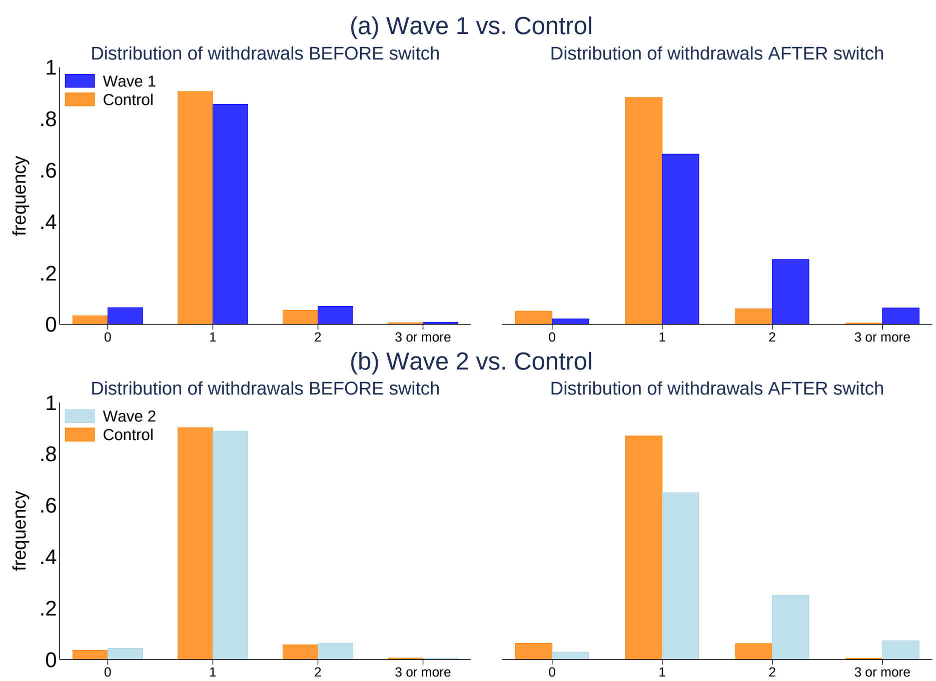

Sources: Administrative data from Bansefi on transactions by bimester and timing of card receipt.
Notes: $N = $ 5,834,468 account-bimester observations from 343,204 accounts.

This increased account use will also lead to a "mechanical" increase in our dependent variable, average balance, because beneficiaries will be leaving a portion of their transfer in the account for a longer period of time. For example, the 22% who make two withdrawals with the card withdraw during the first withdrawal 71% of the total amount withdrawn over the bimester (which might be less than the total deposited if they intend to save some), then return on average 9 days later to make a second withdrawal of the remaining 29% of the total they withdraw over the period. For these 9 days, 29% of the amount they withdrew over the bimester (and hence did not save) is nevertheless captured in the balance; we call the effect of this on the average balance over the period the "mechanical effect." Furthermore, even for those who make one withdrawal of the entire transfer, the average balance will be positive if they wait some number of days after receiving the deposit before withdrawing it. We compute the mechanical effect for each account in each bimester using data on the amounts of and timing between deposits and withdrawals during each bimester, as described in detail in the Appendix, and subtract this from the average balance to create a variable we call "net balance."

Figure 3.10 shows the ϕ_k coefficients from (3.1), using net balance (i.e., average balance minus mechanical effect) as the dependent variable. The debit cards lead to an increase in the stock of savings, with net balances in the account tending to increase over time with the debit card. In Wave 1, there is a marked delay of about one year before beneficiaries start using the account to save. As expected, after subtracting out the mechanical effect from average balances, the treatment effect is smaller in magnitude, reaching about 900 pesos after two years with the card, compared to 1400 pesos in the average balances specification.

3.5 Effect of Debit Cards on Marginal Propensity to Save

To measure the propensity to save, we control for the amount received in transfers each period. This is important since there is a large amount of variation in transfers received within accounts over time, as well as between accounts.The variation within an account

Figure 3.10: Difference between Treatment and Control in Net Balances

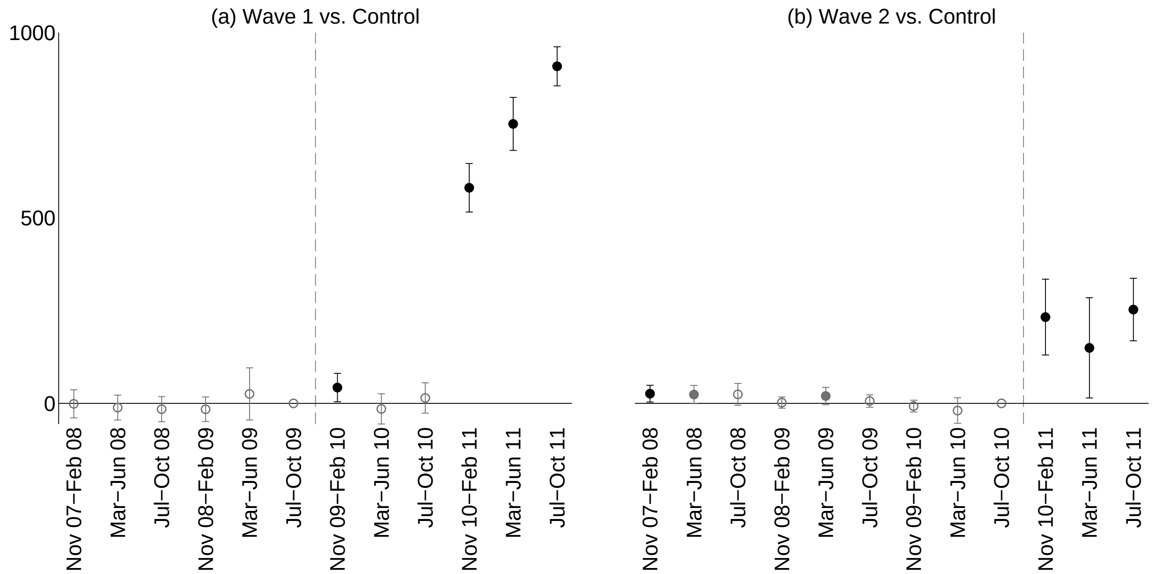

Sources: Administrative data from Bansefi on average account balances by bimester, timing and amount of transfer payments, timing and amount of withdrawals, and timing of card receipt.
Notes: (a) $N = 2{,}023{,}862$ from 171,441 accounts. (b) $N = 3{,}086{,}749$ from 270,046 accounts. Net balances refer to average balances minus the mechanical effect on average balance of leaving a portion of the deposit in the account for a certain number of days before withdrawing it. The figure plots ϕ_k from (3.1). Average balance over each four-month period is the dependent variable, and is winsorized at the 95th percentile. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval.

over time can be explained by local elections in certain localities,[7] compliance with program conditions,[8] payment amounts varying depending on the time of year,[9] and family structure.[10]

In the spirit of asset accumulation models, we assume that savings in period t is a function of assets in period $t-1$, income in period t, (time-invariant) individual preferences,

[7] When there is an election, Oportunidades has to give the transfer in advance, so that there is no payment close to the election month. In practice, this means that beneficiaries receive no payment in the bimester of the election and an additional payment toward the end of the preceding bimester.

[8] If a family does not comply with program conditions such as school attendance and health check-ups, the payment is suspended, but if the family returns to complying with the conditions, the missed payment is added into a future payment.

[9] For example, the program includes a school component that is not paid during the summer, and a school supplies component that is only paid during one bimester out of the year.

[10] One child might age into or out of the program, for example.

and period-specific shocks such as changes to prices:

$$Savings_{it} = f(Assets_{i,t-1}, Income_{it}, \lambda_i, \delta_t). \tag{3.2}$$

Linearizing f, separating assets into the savings balance in the Bansefi account and other assets, and separating income into Oportunidades transfer income and other income gives

$$\begin{aligned} Savings_{it} = \lambda_i + \delta_t + \beta Net\ Balance_{i,t-1} + \kappa Other\ Assets_{i,t-1} \\ + \gamma Transfers_{it} + \xi Other\ Income_{it} + \varepsilon_{it}, \end{aligned} \tag{3.3}$$

where ε_{it} captures period-specific idiosyncratic shocks. Our adminitrative data from Bansefi only include transfers and balances, but not other income and other assets; after removing these terms from (3.3), each household's average other income and average assets over time are captured by the fixed effect λ_i, while idiosyncratic changes in these variables over time add noise in the error term. Our measure of savings at time t is the difference in net balance between time t and time $t-1$; we thus have

$$\begin{aligned} Net\ Balance_{it} - Net\ Balance_{i,t-1} = \lambda_i + \delta_t + \beta Net\ Balance_{i,t-1} \\ + \gamma Transfers_{it} + \varepsilon_{it}, \end{aligned} \tag{3.4}$$

where γ gives the marginal propensity to save out of transfer income. Since transfers are on average about 20% of total income in our sample, dividing our estimates by five gives a rough approximation of the marginal propensity to save out of total income. Grouping terms in (3.4) gives

$$Net\ Balance_{it} = \lambda_i + \delta_t + \theta Net\ Balance_{i,t-1} + \gamma Transfers_{it} + \varepsilon_{it}, \tag{3.5}$$

where $\theta = 1 + \beta$; then to estimate the effect of receiving a debit card on the marginal propensity to save out of transfers and allow this effect to change over time, we estimate

$$\begin{aligned} Net\ Balance_{it} = \lambda_i + \delta_t + \theta Net\ Balance_{i,t-1} + \sum_{k=2}^{12} \alpha_k T_{j(i)} \times \mathbb{I}(k=t) \\ + \sum_{k=2}^{12} \gamma_k Transfers_{it} \times \mathbb{I}(k=t) \\ + \sum_{k=2}^{12} \psi_k Transfers_{it} \times T_{j(i)} \times \mathbb{I}(k=t) + \varepsilon_{it}. \end{aligned} \tag{3.6}$$

As is well-known, however, fixed effects panel data models with a lagged dependent variable (also known as dynamic panel data models) are biased and inconsistent (Nickell, 1981). We thus use the system generalized method of moments (GMM) estimator proposed by Arellano and Bover (1995) and Blundell and Bond (1998), which is consistent for fixed T, large N (as we have here) and performs well in Monte Carlo simulations, especially for large N (Blundell et al., 2001). The two-step system GMM estimator also appears to perform better than Kiviet's (1995, 1999) and Bruno's (2005) least squares dependent variable correction methods when N is large (Bun and Kiviet, 2006). The effect of the debit card on the marginal propensity to save out of transfer income in bimester k is $\alpha_k/\mu_k + \psi_k$, where μ_k is average transfers in bimester k; Figure 3.11 plots the $\alpha_k/\mu_k + \psi_k$ estimates along with their confidence intervals. Standard errors of the parameters in (3.6) are clustered at the bank branch level and corrected for finite sample bias following Windmeijer (2005); the formula for the variance of $\alpha_k/\mu_k + \psi_k$ is then approximated using the delta method. As before, we estimate (3.6) separately for wave 1 and wave 2.[11]

[11] Following the best reporting practices outlined in Roodman (2009a), the details of our two-step system GMM estimation are as follows. Lagged balance is used as an endogenous GMM-style instrument; because bias can increase in finite samples as T increases (since this leads to more lags and, hence, more instruments: see Ziliak, 1997; Roodman, 2009b), to reduce the number of instruments we only use one lag of $Balance_{i,t-1}$ as an instrument. Because $Transfers_{it}$ is predetermined but not strictly exogenous, variables on the right hand side of (3.6) interacted with $Transfers_{it}$ are valid instruments in the system's equation in levels, but not the equation in differences; as a result, we include time dummies and all interaction terms on the right-hand side of (3.6) as IV-style instruments in the system's equation in levels, and time dummies and interaction terms excluding those interacted with $Transfers_{it}$ in the equation in differences. These specification choices result in a total instrument count of 70. Because our panel does not include gaps, we use first differencing—as in Blundell and Bond (1998)—rather than the sample-maximizing forward orthogonal deviations—as in Arellano and Bover (1995)—to eliminate fixed effects in the transformed equation to be estimated.

Figure 3.11: Difference between Treatment and Control in Marginal Propensity to Save Out of Transfer

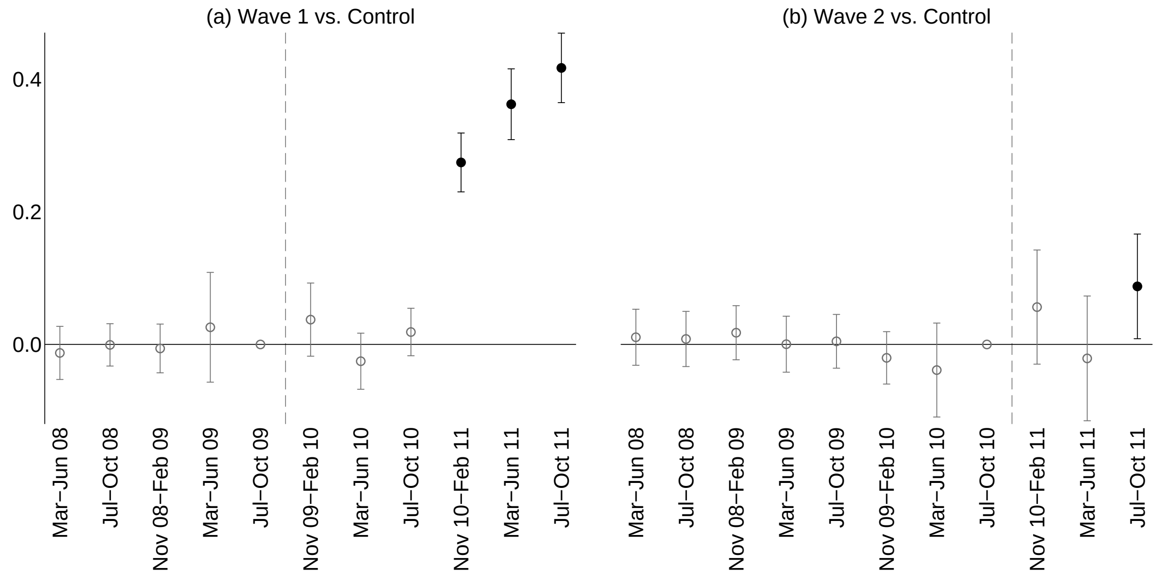

Sources: Administrative data from Bansefi on average account balances by bimester, timing and amount of transfer payments, timing and amount of withdrawals, and timing of card receipt.
Notes: (a) $N =$ 1,852,416 from 171,441 accounts. (b) $N =$ 2,816,671 from 270,046 accounts. (Total number of observations does not include the $t = 1$ observations, which are not included in the regressions but are used to generate $y_{ij,t-1}$ for $t = 2$ observations.) The figure plots $\alpha_k/\mu_k + \psi_k$ from (3.6) estimated by Blundell and Bond (1998) two-step system GMM, where μ_k is average transfers in period k. Average balances and transfer amounts are winsorized at the 95th percentile within the treatment and control groups and within each time period. The variance of $\frac{\alpha_k}{\mu_k} + \psi_k$ is estimated using the delta method. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent level, and hollow circles indicate results that are statistically insignificant from 0. The period prior to receiving the card is the omitted period, which is why its point estimate is 0 with no confidence interval.

In Figure 3.11, the marginal propensity to save out of the transfer is not significantly different between the treatment and control prior to receiving the card, and we observe a delayed effect after receiving the card: in wave 1, the effect remains statistically insignificant from 0 for the first three 4-month periods after receiving the card, while in wave 2 it is insignificant from 0 for the first two periods after they receive the card. The MPS then increases over time and, in wave 1 where we have more post-treatment data, increases substantially over the two years with the card. After one year with the card (in the November 2010–February 2011 period), account-holders save 26.8% of their transfer, which—using household survey data merged with administrative data from Oportunidades on bimonthly transfers to determine the proportion of total income coming from transfers—equals about

5.4% of total income; after close to two years (in the July–October 2011 period), it equals 39.1% of the transfer, or 7.9% of total income. In wave 2, the MPS increases sooner, reaching 10.8% of the transfer or 2.2% of total income after between 6 months and one year with the card.

3.6 Mechanisms

Why do we see a delayed savings effect after receiving the debit card, and why does the marginal propensity to save out of the transfer gradually increase with time? We conjecture that learning is at play and explore three kinds of learning: operational learning (i.e., learning how to use the technology), learning the program rules (specifically, that the program will not drop beneficiaries who accumulate savings), and learning to trust (that the bank is a safe place to save). The first involves knowledge of how to use the debit card and ATM, memorizing the card's PIN, etc. The second involves learning that the program will not use accumulated savings as a signal that the family is actually not poor enough to be receiving Oportunidades benefits. These first two explanations were conjectured by Oportunidades program officials when we shared our initial results from the administrative Bansefi data. The third involves learning that the risk of getting the money "stolen" in the form of hidden fees, operational errors, or nefarious behavior by the bank is lower than initially believed. We find evidence that beneficiaries use the card to check their account balances, and that it thus provides them with a technology to monitor bank behavior, ensure that their money is not disappearing, and subsequently build trust in the bank.

We first use data from the ENCASDU, a survey that directly asks beneficiaries "Do you leave part of the monetary support from Oportunidades in your bank account?" and, if the response is no: "Why don't you keep part of the monetary support from Oportunidades in your Bansefi bank account?" The second question includes pre-written responses and an open-ended response. An example of an answer coded as lack of knowledge is "They didn't explain the process for saving." An example of an answer coded as fear of being dropped from the program is "Because if I save in that account, they can drop me from the program." An example of an answer coded as lack of trust is "Because if I don't take out all the money I can lose what remains in the bank." The ENCASDU surveyed 8788 Oportunidades beneficiary households across rural, semi-urban, and urban areas; of these, the 1674 that received Oportunidades benefits in savings accounts tied to debit cards at the time of the survey make up our sample.

We estimate

$$y_i = \alpha + \gamma \mathbb{I}(\text{Card} \leq 6 \text{ months})_i + u_i, \quad (3.7)$$

where three regressions are run in which the dependent variable $y_i = 1$ if the beneficiary reports not saving due to (i) a lack of knowledge, (ii) fear they will be dropped from the program, or (iii) lack of trust. We estimate the unconditional probability, i.e. beneficiaries who report saving are included in the regression with $y_i = 0$. The unconditional probability is the more relevant measure; instead using the conditional probability (only including those who save in the regression) would mean that the delayed effect we have observed of debit cards on savings could drive the result. Standard errors are clustered at the locality level. We test the null hypothesis $\gamma = 0$, where a rejection of the null would imply that the dependent variable we are testing—which is related to either learning to use the technology, learning program rules, learning to trust the bank—changes over time with the card. Although this survey is cross-sectional, we exploit the variation in time with the debit card, exogenously determined by the staggered locality-level roll-out of the cards.

Figure 3.12a and Table 3.2a show the results. The first thing to note is that lack of knowledge and fear of being dropped from the program after saving are rarely cited as reasons for not saving (combined, less than 4 percent of the sample who have had the card for less than 6 months do not save for these reasons), while lack of trust is cited by 27 percent of those who have had the card for less than 6 months. Second, the proportion who report not saving due to a lack of knowledge does not change over time; in contrast, trust increases gradually with experience: beneficiaries with more than 6 months with the card are 36 percent less likely to report not saving due to low trust than those with less than 6 months with the card.

Figure 3.12: Trust and Knowledge Over Time with the ATM Card

(a) ENCASDU (2010)

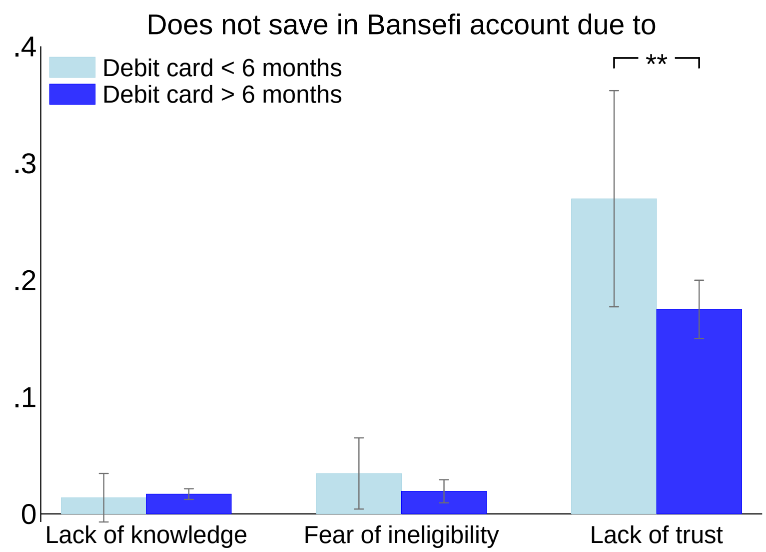

(b) Payment Methods Survey (2012)

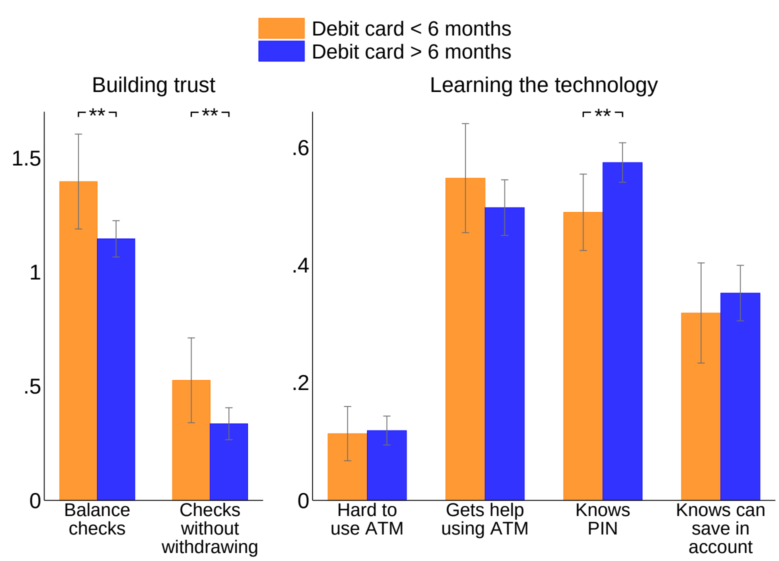

Sources: ENCASDU 2010 and Payment Methods Survey 2012.
Notes: (a) $N = 1674$. (b) $N = 1617$, or less in some regressions if there were respondents who reported "don't know" or refused to respond (see Table 3.2 for number of observations in each regression). Balance checks are measured over the past bimester. Whiskers denote 95 percent confidence intervals. Bars for "debit card $<$ 6 months" are colored light blue in (a) because at the time of ENCASDU 2010, those with the card 6 months or less were in wave 2 localities; bars for "debit card $<$ 6 months" are colored orange in (b) because at the time of Payment Methods Survey 2012, those with the card 6 months or less were in control localities.

Table 3.2: Trust and Knowledge Over Time with the ATM Card

	Mean	Has card $\leq$ 6 months	N
Panel A: ENCASDU Survey (2010): Doesn't save in Bansefi due to …			
Lack of knowledge	0.017***	–0.003	1,674
	(0.002)	(0.010)	
Fear of ineligibility	0.019***	0.015	1,674
	(0.004)	(0.015)	
Lack of trust	0.175***	0.095**	1,674
	(0.012)	(0.044)	
Panel B: Payment Methods Survey (2012)			
Lack of trust			
Times checked balance	1.146***	0.251**	1,493
	(0.039)	(0.105)	
Times checked balance without withdrawing	0.336***	0.190**	1,490
	(0.035)	(0.093)	
Lack of knowledge			
Hard to use ATM	0.106***	0.002	1,617
	(0.013)	(0.025)	
Gets help using ATM	0.498***	0.050	1,612
	(0.023)	(0.048)	
Knows PIN	0.575***	–0.085**	1,609
	(0.017)	(0.034)	
Knows can save in account	0.353***	–0.034	1,617
	(0.023)	(0.046)	

Notes: * indicates statistical significance at $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$. Standard errors are clustered at the locality level. The "Mean" column shows the mean for those who have had the card for more than six months; the "Has card $\leq$ 6 months" column shows the regression coefficient on a dummy for those who have had the debit card for six months or fewer (i.e., the difference relative to the mean column).

Next, we explore mechanisms behind operational learning and learning to trust the bank using the 2012 Payment Methods Survey. The survey includes a number of questions related to operational learning: "What have been the main problems you have had with the ATM?"; "In general, does someone help you use the ATM?"; "Do you know your PIN by heart?"; "Did they tell you that with the card you have a Bansefi savings account?" It also includes a question on balance checking ("In the last bimester, how many times did you check your balance?"), which is a mechanism that beneficiaries could use to build trust in the bank once they have a debit card. The Payment Methods Survey included 5381 households, drawn by stratified (by payment method and locality) random sampling from all Oportunidades beneficiaries; of these, our sample is made up of the 1641 who received their benefits on debit cards tied to savings accounts.

We again use specification (3.7), with y_i equal to: (i) the self-reported number of balance checks over the past bimester; (ii) the self-reported number of balance checks over the past bimester without withdrawing any money, constructed as the total number of balance checks minus the number of withdrawals; and dummies if the respondent reports (iii) it is hard to use the ATM; (iv) she gets help using the ATM; (v) she knows her PIN; (vi) she knows she can save in the account. Because this survey was conducted in 2012, those with the card for at least 6 months now include both wave 1 and wave 2, while beneficiaries in the localities we treat as control localities throughout this paper make up the group with cards for less than 6 months.

Figure 3.12b and Table 3.2b show the results. Both the number of balance checks and number of balance checks without withdrawing decrease over time with the card. Making trips to the ATM specifically to check the account balance (i.e., making a balance check without withdrawing any money) decreases by 36 percent after six months compared to the first 6 months (from an average of 0.53 balance checks without withdrawing to an average of 0.34), while most measures that indicate knowledge of how to use the technology do not change over time: the proportions who report it is hard to use the ATM (around 10 percent), that they get help using the ATM (55 percent), and that they know they can save in the account (32 percent) do not change, although there is a statistically significant increase in the proportion who know their PINs (from 49 to 58 percent).

Finally, we use the administrative transactions data from 342,709 Bansefi accounts, which include the date, time, and fee charged for each balance check at an ATM for each account, to investigate whether the mechanism that appears to be driving the increase in trust—balance checks which clients use to monitor and, over time, build trust in the bank—

holds true in the administrative data; the increased power we have from a large number of observations in the administrative data allows us to take a more granular look at balance checks over time. Note that balance checks at a Bansefi branch are possible both before and after receiving the debit card. Nevertheless, if the distance to the nearest bank branch is high, the debit cards provide a technology that greatly reduces the cost of balance checking (by enabling clients to check their balances at the closest ATM of any bank). Since balance checks at a Bansefi branch are not charged a fee—unlike balance checks at ATMs in Mexico—we do not observe them in our data, which is why the average number of balance checks (at ATMs) in the graph begin after debit card receipt.

Learning to use the technology makes a different prediction regarding the evolution of balance checks over time than learning to trust the bank. The former means that it is easier—less costly—for a beneficiary to check her balance as she learns to use the technology (e.g., by memorizing her PIN or learning how to use the ATM). Therefore, if anything, we might expect her to check her balance more over time. On the contrary, learning to trust predicts that although at the start an individual would check her balance often to monitor her savings, she learns that her money is still in the account and updates downward her belief about the risk of losing money. With simple Bayesian learning, balance checking has decreasing marginal benefit and therefore she checks her balance less over time.

Figure 3.13 plots the number of times people check their balance per bimester, with vertical lines indicating the timing of card receipt. Again due to the shifting of some payments to the end of the previous bimester (which might affect the bimester timing of balance checks), we continue grouping adjacent bimesters into four-month periods. We observe that the number of balance checks per bimester is initially high (about 2.5 checks on average in wave 1 and 1.5 checks on average in wave 2), but in both waves decreases during each four-month period after beneficiaries receive the card, consistent with the trust-building hypothesis.

Figure 3.13: Balance Checks (Administrative Data)

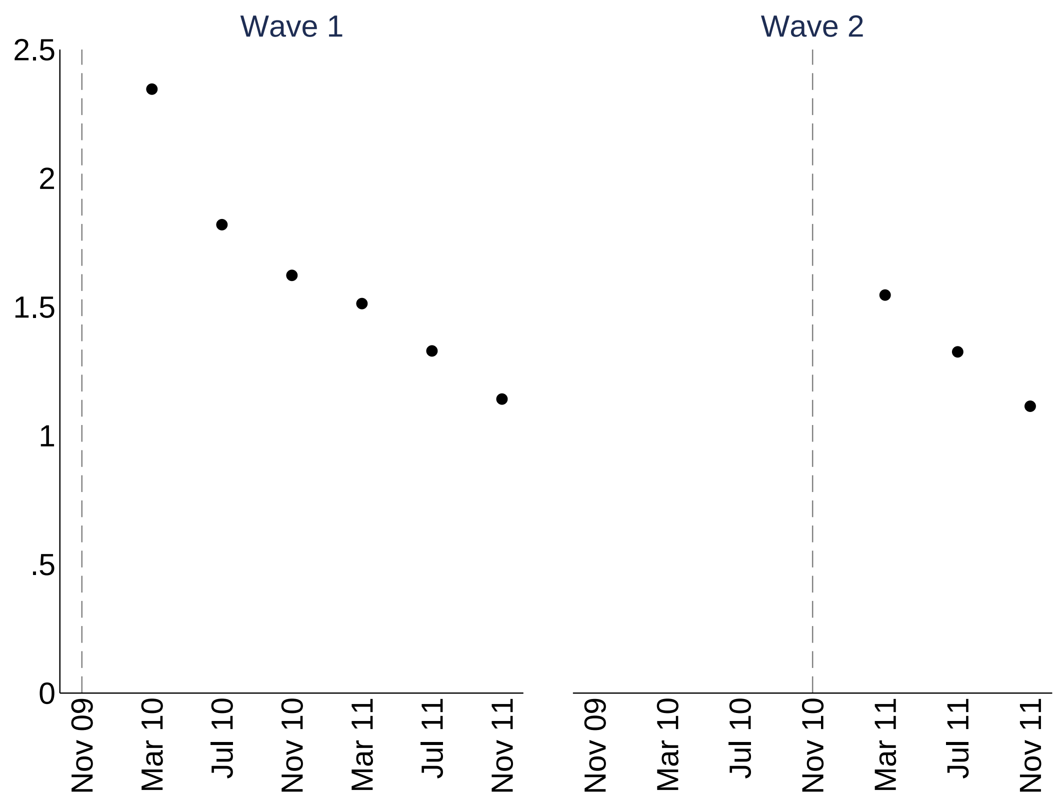

Source: Aministrative transactions data from Bansefi.
Notes: Number of balance checks per account tied to a debit card. Prior to receiving the card it was possible to check balances at Bansefi branches only, and balance checks at Bansefi branches are not recorded in our transactions data because they are free of charge.

3.7 Increase in Overall Savings vs. Substitution

The increase in formal Bansefi account savings might come at the expense of other types of savings that the household is already conducting, in such a way that total savings is not affected. The question of whether the observed increase in Bansefi savings crowds out other saving is relevant not only if one is concerned with total household savings, but also to understand the mechanics through which the effect on formal savings is operating and as a first step towards thinking about the broader welfare implications of providing a formal savings account with a debit card.

Does the provision of the debit card and the resulting increase in formal savings represent an increase in *overall* savings, or is it merely a substitution from other forms of saving? To address this question, we use Oportunidades' ENCELURB panel survey, conducted in four waves during the years 2002, 2003, 2004 and November 2009 to February 2010. This survey is conducted by Oportunidades and has comprehensive modules on consumption, income, and assets for 6272 households in urban and semi-urban areas.[12] Of the 6272 households in the post-treatment wave of ENCELURB, 2951 live in urban areas and, according to administrative data provided by Oportunidades and merged with the survey, are Oportunidades beneficiaries when interviewed in the post-treatment wave and receive their benefits in a savings account (with or without a debit card); this is the sample used in our analysis, except in the placebo tests described in Section 3.9.

As before, we use a differences-in-differences strategy where we examine changes in consumption, savings, and income across beneficiaries, exploiting the differential timing of debit card receipt. Because the ENCELURB was conducted after wave 1 localities had received cards but before wave 2 or control localities had received cards, we compare those with cards (wave 1) to those who have not yet received cards (waves 2 and control), respectively referring to them as "treatment" and "control" in this section of the paper. The identification assumption is that in the absence of the debit card, treatment and control groups would have experienced similar changes in consumption, income, and assets. We formally test for parallel trends in Section 3.9, and since we indeed find that trends were parallel prior to treatment, we now test whether there was an increase in savings, which we construct as income minus consumption from the income and consumption modules of

[12] The 2002, 2003, and 2004 waves had around 17,000 households, but due to budget constraints the number of localities was cut for the 2009-2010 wave. The consumption, income, and assets modules of Oportunidades' analogous survey for rural areas have been used by Angelucci and De Giorgi (2009), Attanasio et al. (2013), de Janvry et al. (2015), Gertler et al. (2012), and Hoddinott and Skoufias (2004), while these modules from the ENCELURB have been used by Angelucci and Attanasio (2013) and Behrman et al. (2012).

ENCELURB. We estimate

$$y_{it} = \lambda_i + \delta_t + \gamma D_{j(i)t} + \nu_{it} \tag{3.8}$$

separately for five dependent variables: consumption, income, savings (constructed as income minus consumption), purchase of durables, and an asset index. All variables except assets are measured in pesos per month, i indexes households, and t indexes survey years. The asset index dependent variable is constructed as the first principal component of dummy variables indicating ownership of the assets that are included in all rounds of the survey questionnaire: car, truck, motorcycle, TV, video or DVD player, radio, washer, gas stove and refrigerator. Time-invariant differences in household observables and unobservables are captured by the household fixed effect λ_i, common time shocks are captured by the time fixed effects δ_t, and $D_{j(i)t} = 1$ if locality j in which household i lived prior to treatment has received debit cards by time t (i.e., in the notation used in specifications (3.1) and (3.6), $D_{j(i)t} \equiv T_{j(i)} \times \mathbb{I}(t = \text{2009–10})$). We use the locality of residence prior to treatment to avoid capturing migration effects in our estimation and estimate cluster-robust standard errors at the locality level.

If the increase in formal savings constitutes an increase in total savings then we expect $\gamma > 0$ for total savings (defined as income minus consumption), and if we observe $\gamma = 0$ for income we expect $\gamma < 0$ for consumption. If there is no substitution of savings from assets (and if they are not using the formal savings accounts to save up for assets, at least in the short run), we expect $\gamma = 0$ for the purchase of durables (which measures a flow) and the asset index (which measures a stock). This is indeed what we find. Figure 3.14 shows that consumption decreased almost 200 pesos on average (statistically significant at the 5% level). Meanwhile, there is no effect on income; we also test the difference in the coefficients of consumption and income using a stacked regression (which is equivalent to seemingly unrelated regression when the same regressors are used in each equation, as is the case here); although both are noisily measured, the difference in the coefficients is significant at the 10% level. Purchase of durables and the stock of assets do not change, ruling out a crowding out of these forms of saving. The increase in savings, measured as income minus consumption, is estimated at slightly more than 200 pesos, and is significant at the 5 percent level.

The results in Figure 3.14 are from our preferred specification where we winsorize the dependent variable at 5 percent (specifically, at the 95th percentile, as well as the 5th percentile if the variable does not have a lower bound of 0) to avoid letting our results be driven by outliers. Table 3.3, columns 1–3 show that the effects are robust to using the raw data without winsorizing, winsorizing at 1 percent, or—as in our preferred specification—winsorizing at 5 percent (we follow Kast and Pomeranz (2014) who show the robustness of results to these three possibilities for their savings measures). They are also robust to including baseline characteristics interacted with time fixed effects, as well as municipality-specific time effects, both to control for specific time trends more flexibly (Table 3.3 columns 4 and 5).[13]

These results mean that total savings—not just formal savings—increase, and that this increase in being funded by lower consumption. A back of the envelope calculation reveals that the magnitude of the increase in monthly savings from this household survey is in line with the average increase of savings in the account from the administrative data: from the propensity to save specification, after 1 year, beneficiaries who received cards in wave 1 save 26.8 percent of their transfer more than the control group. Using ENCELURB, transfers are, on average, 20.2 percent of income for the treatment group, implying that the savings effect in the Bansefi administrative data is about 5.4% of income. The effect for

[13] The household characteristics interacted with time fixed effects in this robustness check are measured at baseline and include characteristics of the household head (whether the household head worked, a quadratic polynomial in years of schooling, and a quadratic polynomial in age), whether the household has a bank account, variables used to measure poverty by Oportunidades (the proportion of household members with health insurance, the proportion aged 15 or older that are illiterate, the proportion aged 6 to 14 that do not atted school, the proportion aged 15 or older with incomplete primary education, and the proportion aged 15 to 29 with less than 9 years of schooling), and dwelling characteristics (dirt floor, no bathroom, no water, no sewage, number of occupants per room).

Figure 3.14: Effect of the Debit Card on Consumption, Income, Total Savings, Purchase of Durables, and Assets

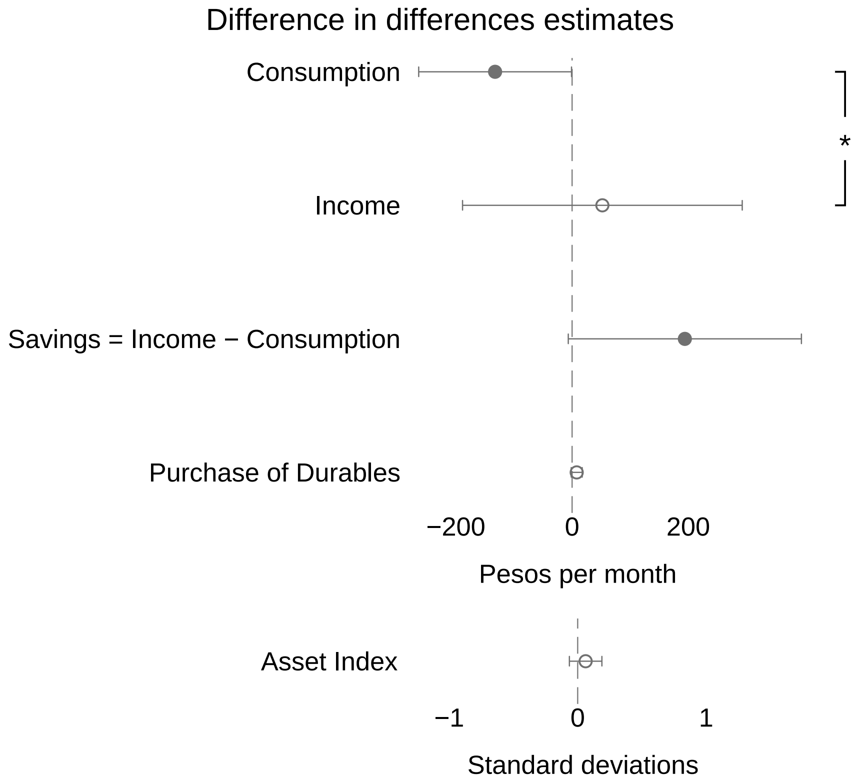

Sources: ENCELURB panel survey combined with administrative data on timing of card receipt and transfer payment histories for each surveyed beneficiary household.
Notes: $N = 11,275$ (number of households $= 2951$). Dependant variables are measured in pesos per month, with the exception of the asset index. Asset index is the first principal component of assets that are included in both the early (2002, 2003, 2004) and post-treatment (2009–2010) versions of the survey: car, truck, motorcycle, television, video or DVD player, radio or stereo, washer, gas stove, and refrigerator. Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent level, and hollow circles indicate results that are statistically insignificant from 0. The * linking consumption and income denotes that a test of equal coefficients from the consumption and income regressions is rejected at the 10 percent level using a stacked regression. Results are from the preferred specification of winsorizing variables at the 95th percentile (and 5th percentile for variables that do not have a lower bound of 0). Raw results, winsorized at 1 percent, winsorized at 5 percent, winsorized at 5 percent with baseline household characteristics interacted with time fixed effects, and winsorized at 5 percent with municipality $\times$ time fixed effects are available in Table 3.3. All regressions include household and time fixed effects, and standard errors are clustered at the locality level, using pre-treatment (2004) locality.

savings (income minus consumption) in the ENCELURB household survey data shown in Figure 13 equates to 4.8 percent of income. Taken at face value, this suggests that most of the increase in savings in the account is new saving. This result is consistent with other

Table 3.3: Effect of the Debit Card on Consumption, Income, Total Savings, Purchase of Durables, and Assets

	(1)	(2)	(3)	(4)	(5)
Consumption	–178.27**	–143.61*	–132.54*	–250.59**	–149.33**
	(85.36)	(74.56)	(67.11)	(115.68)	(68.39)
Income	73.43	78.98	52.20	40.34	50.47
	(154.24)	(137.97)	(122.83)	(132.90)	(120.47)
P-value Consumption vs. Income	[0.057]*	[0.054]*	[0.098]*	[0.010]***	[0.072]*
Savings = Income – Consumption	251.70*	214.11*	194.17*	283.47**	213.75**
	(128.76)	(113.56)	(102.38)	(121.86)	(103.11)
Purchase of durables	5.94	6.22	7.99	6.55	6.91
	(12.55)	(8.52)	(4.82)	(6.78)	(4.55)
Asset index	0.04	0.04	0.06	–0.07	0.05
	(0.07)	(0.07)	(0.06)	(0.06)	(0.06)
Number of households	2,951	2,951	2,951	2,951	2,938
Number of observations	11,275	11,275	11,275	11,275	11,243
Time fixed effects	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes
Municipality $\times$ time fixed effects	No	No	No	Yes	No
Household characteristics $\times$ time	No	No	No	No	Yes
Winsorized	No	1%	5%	5%	5%

Notes: * indicates statistical significance at $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$. Standard errors are clustered at the locality level, using pre-treatment (2004) locality. Dependant variables are measured in pesos per month, with the exception of the asset index. Asset index is the first principal component of assets that are included in both the early (2002, 2003, 2004) and post-treatment (2009–2010) versions of the survey: car, truck, motorcycle, television, video or DVD player, radio or stereo, washer, gas stove, and refrigerator. Household characteristics are measured at baseline (2004, or for households that were not included in the 2004 wave, 2003). They include characteristics of the household head (working status, a quadratic polynomial in years of schooling, and a quadratic polynomial in age), whether anyone in the household has a bank account, a number of characteristics used by the Mexican government to target social programs (the proportion of household members with access to health insurance, the proportion age 15 and older that are illiterate, the proportion ages 6-14 that do not attend school, the proportion 15 and older with incomplete primary education, the proportion ages 15-29 with less than 9 years of schooling), and dwelling characteristics (dirt floors, no bathroom, no piped water, no sewage, and number of occupants per room). The number of households in column (5) is slightly lower because 13 households have missing values for one of the household characteristics included (interacted with time fixed effects) in that specification.

studies where formal savings products were offered, which found that the increased savings in these products did not crowd out other forms of saving (Ashraf et al., 2015; Dupas and Robinson, 2013a; Kast et al., 2012).

Taken together, these results suggest that Oportunidades beneficiaries, who were in general not using their Bansefi bank accounts to save prior to receiving a debit card tied to the savings account, use the debit cards to check their balances and thereby build trust in the bank; after 4-8 months with the card, they begin saving and their marginal propensity to save increases over time. The observed increase in savings in their Bansefi accounts represents an increase in overall savings.

3.8 Does Money Burn a Hole in Your Pocket?

Because the accounts pay no interest, but there was clearly an unmet demand for savings among program beneficiaries, we explore why they were not able to save before (for example, under the mattress). Since the results in Figure 3.14 show that the debit card induces higher total savings through decreased consumption, we might expect that it influences different components of consumption differentially. We thus examine the proportion of income spent on several categories of consumption goods: temptation goods (where we group the three most frequently cited temptation goods in Banerjee and Mullainathan (2010): alcohol, tobacco, and sugar); fats and sweets (junk food, fats, soda); meat, dairy, and produce; tortillas and cereals; entertainment; transportation; and health and education. We use the proportion of total income spent on each consumption category, rather then the level of consumption in that category, because individual $\times$ time-specific shocks to income, which we expect to be passed through as shocks to various consumption categories, would otherwise add noise to the estimation through the error term; we use total income rather than total consumption in the denominator because, from the results in Figure 3.14, total income does not change differentially between the treatment and control groups.

We estimate a DID specification with household and year fixed effects and standard errors clustered at the locality level; specifically, for each consumption category g,

$$\frac{Consumption_{git}}{Income_{it}} = \lambda_{gi} + \delta_{gt} + \gamma_g D_{j(i)t} + \nu_{git}, \tag{3.9}$$

where $Consumption_{git}$ is monthly consumption of good g by household i at time t (in pesos) and $Income_{it}$ is total monthly income of household i at time t (in pesos). We find that the proportions of income spent on temptation goods and transportation decrease in the treatment group relative to the control (Figure 3.15), and that these are the only two categories where the decrease in consumption is statistically significant (at the 5% and 10% significance levels, respectively). Although our grouping of temptation goods is based on the goods most frequently mentioned by Banerjee and Mullainathan (2010), it could be viewed as arbitrary (and, indeed, we do not find a decrease in the grouping of fats and sweets—junk food, fats, and soda—which could also be classified as temptation goods); we thus look separately at each item in the temptation good category, and find a statistically significant decrease in consumption of alcohol and sugar, but not of tobacco.

Figure 3.15: Effect of the Debit Card on Different Categories of Consumption

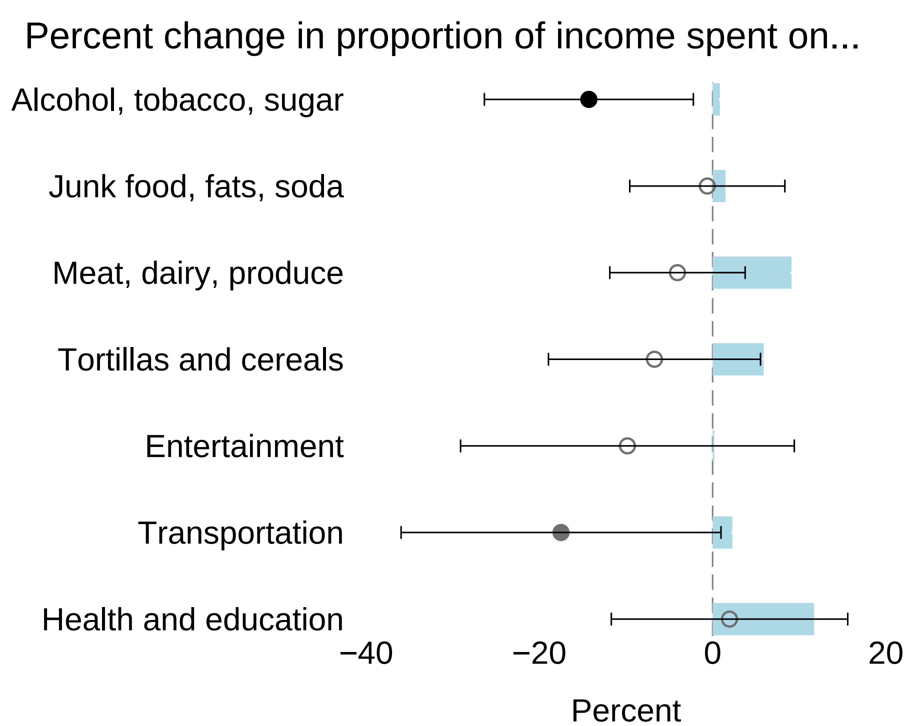

We interpret this result as evidence that it is difficult to save informally due to self-control problems, and that these problems can be partially solved by access to a formal savings account (but that low indirect transaction costs and trust in the bank are necessary conditions for these formal savings accounts to be used). This finding is consistent with the demand for commitment savings devices (e.g., Ashraf et al., 2006; Bryan et al., 2010) if the savings accounts without debit cards, which could be used as an even stronger commitment due to the high indirect cost of accessing savings, would have been too strong of a commitment (since strong commitment devices have low take-up and use relative to weak commitment devices: see Karlan and Linden, 2014; Laibson, 2015), or if the bank accounts were merely not trusted prior to being able to cheaply monitor them with debit cards. Under either explanation, trust appears to have been a necessary condition for formal saving, given the delayed savings effect and self-reported reasons for not saving initially.

The self-control problems that prevent the poor from saving prior to having access to a trusted formal savings account could result directly from an asset-based poverty trap, as in Bernheim et al. (2015), a model that is consistent with the finding that microcredit decreases temptation good consumption (Angelucci et al., 2015; Augsburg et al., 2015; Banerjee et al., 2015). Alternatively, it is possible that the self-control problems stem from the timing of access to the money: Carvalho et al. (2016), using exogenous variation in the timing of an experiment relative to payday, find that those who are more financially constrained behave in a more present-biased way. If the beneficiary withdraws her money and attempts to save at home, she has easy access to it throughout the two month period (including access to the portion she intended to save rather than spend that period); toward the end of the period she is likely to be more financially constrained. On the contrary, if she trusts the bank and decides to save in her Bansefi account, she makes her saving decision when initially withdrawing benefits, when she is less financially constrained due to having recently received the Oportunidades payment.

It is also possible that saving money informally is difficult because the beneficiary lacks control over her husband or partner's access to money saved at home, and the husband has different (perhaps more present-biased) time preferences. Consistent with these potentially differing preferences, (Rubalcava et al., 2009) find that Oportunidades income (viewed as the wife's income) tends to be spent more on investments in the future than other income does. When the spouses have differing time preferences (even if neither is present-biased), the collective decision making of the household becomes present-biased (Jackson and Yariv, 2014), making soft commitment devices such as bank accounts more attractive.

We thus test whether this constraint—that if the cash is saved at home, the husband could take or request the money, whereas it is difficult for him to access it in the account—is binding.[14] Since single beneficiaries (i.e., beneficiaries who are not living with a spouse or partner) would not be affected by this barrier, a first pass to exploring whether beneficiaries' lack of control over their husbands' access to the money is a barrier to saving informally is to test whether a single woman responds differently to the card than a woman who is living with her spouse or partner.[15] We thus estimate

$$Savings_{it} = \lambda_i + \delta_t + \gamma D_{j(i)t} + \xi D_{j(i)t} \times H_i + \sum_{k \in K} \zeta_k H_i \times \mathbb{I}(t = k) + \nu_{it} \qquad (3.10)$$

where H_i is a time-invariant measure of heterogeneity, and $K = \{2003, 2004, 2009\}$ (dropping 2002 to avoid collinearity with the household fixed effects). The $H_i \times \mathbb{I}(t = k)$ terms thus allow the evolution of savings over time to vary with H_i even in the absence of treatment. In this case, $H_i = 1$ if the beneficiary is single in the post-treatment survey wave (since marriage should not be endogenously affected by receiving the debit card and we do

[14] On the other hand, debit cards might lead the husband to have higher access to the money, especially in households where the woman has low bargaining power. The high indirect transaction costs of a bank account without a card, and the requirement that the card holder herself appear at the bank branch to withdraw money, could make control over the husband's access to the money easier without a debit card. Indeed, in Schaner (2015a), women with low bargaining power are hurt by receiving debit cards because they lose control over the money. As we have already seen, however, the Bansefi bank accounts were not being used to save prior to receiving the debit cards. Even so, the beneficiary might be able to hide the money at home but unable to prevent her husband from taking and using her card to withdraw money. In our survey data, however, only 4% of beneficiaries report that their spouse sometimes withdraws money from the account.

[15] Although it is easy to identify whether the Oportunidades beneficiary is married and living with her spouse in ENCELURB, it is difficult to determine with certainty whether unmarried beneficiaries nevertheless live with a partner. We thus include beneficiaries living in the same household as another adult who is not the household head's child or grandchild in our "non-single" group. Using this definition, the non-single group of beneficiaries is made up of 95% married women and 5% who are not married but living in the same household as another adult.

not want the effect to be driven by beneficiaries whose marital status changes between pre- and post-treatment).[16]

If the husband or partner's access to money is a barrier to saving for women living with a husband (or other adult) but not for single women, we expect $\gamma > 0$, $\gamma + \xi = 0$, and $\xi < 0$. Table 3.4 column 1 shows that we do find $\gamma > 0$ and cannot reject $\gamma + \xi = 0$, but—although the point estimate on ξ is 188 pesos (close to the average treatment effect from Table 3.3, column 3)—it is not statistically significant from 0, so we cannot reject $\xi = 0$.

If a lack of control over the husband's access to money is indeed a barrier to saving, we would also expect treatment effect heterogeneity among women who do live with a husband or partner based on their bargaining power in the household. To test for this, we proxy for baseline female bargaining power using four questions asked only in the first wave of the survey on who makes the primary decisions in the household: whether to take their children to the doctor if they are sick, whether the children have to attend school, whether to buy them new clothes when needed, and "important decisions that affect the household members (transport, moving, changing jobs)." We code these questions as $+1$ if a woman makes the decision, 0 if spouses make them jointly, and -1 if a man makes the decision, then following Kling et al. (2007), standardize the variables to each have a mean of 0 and standard deviation of 1 and average them to create a summary measure of female bargaining power. We estimate (3.10) on the subset of women living with a spouse (or other adult), with H_i as this summary measure of baseline female bargaining power. Our hypothesis that women with high bargaining power could already exercise control over money saved in the home, and thus should not have as large of a treatment effect as women with low bargaining power prior to receiving the card, would mean that $\xi < 0$.

The results of this test are shown in Table 3.4, column 2.[17] Indeed, we find $\xi < 0$,

[16] The sign and statistical significance of the point estimates on ξ and γ are the same, and magnitudes similar, if we instead define H_i at baseline, where baseline refers to using the most recent pre-treatment wave in which household i was included.

[17] Of the 2951 households in our sample, the Oportunidades beneficiary lives with a spouse or partner

Table 3.4: Other Barriers to Saving Informally

Dependent variable: savings	(1)	(2)	(3)	(4)	(5)
Has card at t	241.08*	204.81	176.54	120.80	92.38
	(125.51)	(134.48)	(107.91)	(112.81)	(221.02)
Has card at $t \times$ single	–188.28				
	(173.75)				
Has card at $t \times$ baseline female bargaining power		–196.81*			
		(114.24)			
Has card at $t \times$ household gave money to others at baseline			298.55		
			(408.20)		
Has card at $t \times$ municipal homicides per 100,000				–2.12	
				(4.11)	
Has card at $t \times$ high-crime municipality					131.85
					(247.23)
Number of households	2,951	1,625	2,951	2,951	2,951
Number of observations	11,275	6,300	11,275	11,275	11,275
Subsample	All	Not single[a]	All	All	All
Time fixed effects	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes
Winsorized	5%	5%	5%	5%	5%

Notes: [a]Not single refers to beneficiaries who live with a spouse (95% of the group) or at least one other adult (5% of the group). * indicates statistical significance at $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$. Standard errors are clustered at the locality level, using pre-treatment (2004) locality. Dependant variable is savings, constructed as income minus consumption and measured in pesos per month. "Baseline female bargaining power" uses questions only included in the 2002 wave of the survey on who decides (i) whether to whether to take their children to the doctor if they are sick, (ii) whether the children have to attend school, (iii) whether to buy them new clothes when needed, and (iv) "important decisions that affect the household members (transport, moving, changing jobs)." The measure is constructed by coding the responses to these four questions as 1 if a woman makes the decision, 0 if they make the decision jointly, and –1 if a man makes the decision, then the responses from the multiple questions are standardized and averaged following Kling et al. (2007). "Household gave money to others at baseline" is a dummy variable equal to 1 if the household reported making transfers to others in any of the pre-treatment waves of the survey. "High crime municipality" refers to municipalities with above the median homicide rate, where the median is calculated for household included in our sample.

significant at the 10% level. A one standard deviation *decrease* in baseline female bargaining power translates to an *increase* of about 196 pesos in the savings effect of the debit card, roughly equal to the average treatment effect in the full sample. This suggests that a woman with low bargaining power at baseline (and hence less control over money saved informally) receives a larger benefit from the card because it enables her to build trust in the bank and subsequently save in the account, which is out of reach of her husband. A woman

in 2098 (71%). Of these 2098, only 1625 are included in the regression for column 2; the difference of 473 households is because 93 were not included in the 2002 wave of the survey (the only wave to ask these bargaining power questions), while 380 were included but refused to answer one or more of the bargaining power questions.

with high bargaining power at baseline, on the other hand, was already able to prevent her husband from spending informal savings prior to receiving the card, and thus receives a lower benefit from the card. As a result, among women who are married or living with a partner, the savings effect caused by the debit card is higher for those with low baseline bargaining power.

Another potential barrier to saving informally is that money saved at home could be in demand from friends and relatives. It is obvious from the interviews in Baland et al. (2011) that the desire to conceal money in a savings account to avoid demands from others extends beyond one's spouse to friends and relatives. Ideally, we would test whether transfers from the household to other households decreased after receiving the card; the question on transfers out of the household was not included in the post-wave survey, however. We thus estimate (3.10) with H_i as a dummy variable equal to 1 if the household reported transfers out at baseline (specifically, in any of the three pre-treatment waves). Because those with higher demands for money from friends and relatives are more likely to have $H_i = 1$, if this is a barrier to saving informally we expect $\xi > 0$. The results of this test, shown in Table 3.4, column 3, are inconclusive: although the point estimate on the interaction term is large, at 299 pesos, the standard error is very large, and the effect is statistically insignificant from 0. It is worth noting that only 7% of the sample has $H_i = 1$. This suggests that demands for money from relatives and friends might be a barrier to saving informally, but—if so—that this barrier only affects a small fraction of Oportunidades recipients.

A final potential barrier to saving informally is that the money risks being stolen if saved at home. An anticipated reduction in crime was one of the primary motivations for the change to debit cards; in the U.S., changing the payment method of cash welfare payments to debit cards caused a significant decline in burglary, assault, and larceny (Wright et al., 2014). In developing countries, risk of theft has been anecdotally reported as a reason for not saving at home by cash transfer recipients in the Dominican Republic (Center for Effective Global Action, 2015), and is pointed out as a potential mechanism in Malawi by

Brune et al. (2016).

To test this hypothesis, we test whether high-crime municipalities—where saving informally would be more difficult due to risk of theft—have a higher treatment effect. Specifically, we use municipal homicide rates since these are the best available proxy of crime at the municipal level in Mexico.[18] We estimate (3.10) where H_i (or, more precisely, $H_{m(i)}$ is either the homicide rate per 100,000 in the municipality m in which household i lives (Table 3.4, column 4) or a dummy variable equal to 1 if the homicide rate is greater than the median, with the median calculated based on the municipal homicide rates faced by each beneficiary household in our sample (column 5). If risk of theft is a barrier to saving, we expect $\xi > 0$, i.e., there was a higher savings effect from the debit cards in localities with higher crime and thus greater risk that informal savings are stolen. In both specifications, the point estimate on the interaction term is statistically insignificant from 0, and in the first the point estimate is very close to 0 and the standard error is small: the 95% confidence interval rules out heterogeneous treatment effects outside of $[-10 \text{ pesos}, +6 \text{ pesos}]$ for an increase in the homicide rate of 1 per 100,000, or about an 11% increase in crime relative to the median homicide rate of 8.7 per 100,000. This suggests that risk of theft is not a barrier to saving informally in our context.

3.9 Robustness

3.9.1 Internal Validity Checks

The identifying assumption for (3.1) and (3.6) is that the beneficiaries that received the debit card in waves 1 and 2 would have had the same average balances and marginal propensity to save as the control group in the absence of treatment. While the assumption is inherently untestable, its plausibility was confirmed by two sets of results presented sections 3.2 and 3.4. First, although the roll-out was not random, most means between treatment and

[18] Homicide rates are not available at the locality level, which is the unit at which the roll-out was determined; a locality is a sub-unit of a municipality.

control do not have statistically significant differences; there is a difference, however, in population, transfer amount, and percent of the transfer withdrawn. (For percent of the transfer withdrawn, the F-test of equality between the three means is rejected, and a test of equality of wave 1 and the control is rejected, but the test of equality between wave 2 and the control is not rejected.) More important, average balances follow parallel pre-treatment trends in wave 1 and the control prior to wave 1 receiving debit cards, and in wave 2 and the control prior to wave 2 receiving debit cards: this can be seen visually in Figure 3.6 and is formally tested in Section 3.4. The similarity of savings in the treatment and control groups before treatment contrasts sharply with the diverging trends after debit cards are received. The fact that results comparing the control to two waves receiving debit cards in different years are similar suggests this is not an artefact of a shock in a particular month or year.

Similarly, the identifying assumption for the household survey panel data results on savings, income, consumption, purchase of durables, and ownership of assets in (3.8) is that these variables would have followed parallel trends in the absence of treatment. Figure 3.16 shows these parallel trends graphically for the pre-treatment rounds of the survey. In addition, because there are many years between the last pre-treatment ENCELURB survey year (2004) and the year of treatment (2009), we supplement the ENCELURB parallel trends tests with tests using data from the 2004–2008 rounds of the ENIGH, a national income and expenditure survey used for Mexico's official poverty measurement. This is a repeated cross-section survey conducted in even years (but additionally conducted in 2005) that sampled between 20,000 and 30,000 households during each year in this time frame.

Although the publicly available version of the survey does not include each household's locality code, which determines whether the household lives in a treatment or control locality, we obtained the locality codes for sampled households from Mexico's National Institute of Statistics and Geography. Although Oportunidades receipt is reported in the survey, there is a large discrepancy between the number of beneficiaries according to the survey (after expansion factors are applied) and the number in national accounts (Scott, 2014), a

Figure 3.16: Parallel Pre-Treatment Trends in Household Survey Data

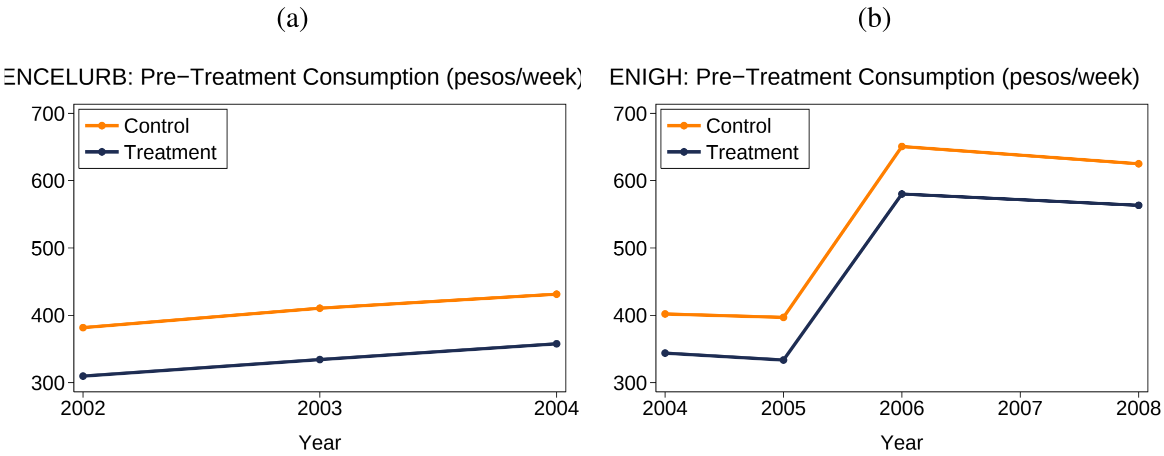

problem also common in developing countries (Meyer et al., 2015), so to have sufficient power for our test we restrict the analysis to the poorest 20 percent of surveyed households to proxy for Oportunidades recipients, rather than use self-reported Oportunidades receipt. Again, the parallel trends can be clearly seen visually in Figure 3.16.

3.9.2 Alternative Explanations and Placebo Tests

We have argued that the card allows beneficiaries to build trust in the bank by monitoring the bank's activity through balance checks. We now explore alternative explanations for the observed delayed savings effect and increasing marginal propensity to save over time. First, it could be that accumulating time *with the savings account*, rather than with the card, drives the increase over time. Second, while the hypothesis that debit cards increased trust through bank monitoring is demand-driven, the effect could be supply-driven if banks optimally responded to the increased debit card concentration by opening up more ATMs or bank branches in those localities; if such an expansion were gradual, it could explain the delayed savings effect and increasing marginal propensity to save over time. Third, the effect might be driven by locality-specific shocks unrelated to the debit cards. Fourth, the debit cards could merely make savings more salient, as in Akbas et al. (2015), by giving beneficiaries a reminder (in the form of an object carried with them) of their savings intentions.

There are a number of reasons that it is unlikely that the effects are driven by experience with the savings account leading beneficiaries to learn the benefits of saving, rather than time with the debit card itself. First, both treatment and control accounts are accumulating time with their savings accounts simultaneously. Second, because the savings accounts were mainly rolled out between 2002 and 2004 (Figure 3.4), most beneficiaries had already accumulated several years with the account by 2009, when our study begins. Indeed, the median date of account opening in our 342,709 accounts is October 18, 2004, and less than 5 percent of accounts had existed for less than two years when they received debit cards. Third, our results from Section 3.5 include account fixed effects, so any time-invariant effect of having the account for a longer period of time would be absorbed. Fourth, to test for a time-varying effect of having the account for a longer period of time, we test whether results vary when we run the analysis separately for two groups: those who have had the account for more vs. less time. We use the median date of account opening to split the accounts into these two groups, and find that results are very similar. The graphs splitting the sample into those that opened before and after the median date of October 18, 2004 both look very similar to Figure 3.11.

A second possible explanation for the increase in savings over time is that banks gradually expanded complementary infrastructure in localities where treated beneficiaries live. Depending on the costs of each branch and ATM machine, this could be a profit-maximizing response to the increase in the number of debit card holders in treated localities. The increasing marginal propensity to save over time could be the result of the staggered expansion of this infrastructure, not increased trust. If this is so, then the increase in savings would have to be reinterpreted not only as the effect of debit cards but of the expansion of the whole enabling technology. Using quarterly data for each municipality on the number of bank branches and ATMs for Bansefi and all other banks, we test if there was indeed a contemporaneous expansion of infrastructure and if this was correlated geographically with Oportunidades debit card expansion or with savings in our accounts.

We first test for a relationship between the roll-out of ATM cards and a supply-side expansion of banking infrastructure (ATMs and bank branches)[19] by estimating:

$$y_{jt} = \lambda_j + \delta_t + \sum_{k=-4}^{4} \beta_k D_{j,t+k} + \varepsilon_{jt},$$

where y_{jt} is the number of total ATMs, total bank branches, Bansefi ATMs, or Bansefi branches in municipality j in quarter t and D_{jt} equals one if at least one locality in municipality j has Oportunidades debit cards in quarter t. We include one year (four quarters) of lags and one year of leads to test for a relationship between bank the debit card roll-out and bank infrastructure. For this test, we use data on the number of ATMs and bank branches by bank by municipality by quarter from the Comisión Nacional Bancaria y de Valores (CNBV), from the last quarter of 2008 through the last quarter of 2013 (since the roll-out was from late 2009 to early 2012, when what we refer to as control group localities received debit cards). We separately test whether lags of credit card receipt predict banking infrastructure (i.e., whether there is a supply-side response to the roll-out of debit cards) by testing $\beta_{-4} = \cdots = \beta_{-1} = 0$ and whether leads of credit card receipt predict banking infrastructure (i.e., whether debit cards were first rolled out in municipalities with a recent expansion of banking infrastructure) by testing $\beta_1 = \cdots = \beta_4 = 0$. We find evidence of neither relationship, failing to reject the null hypothesis of each test for each of the four dependent variables (Table 3.5).

To rule out locality-specific shocks that could be driving the savings effect, as opposed to the effect being driven by the debit cards, we perform a placebo test using poor non-Oportunidades households in the treated vs. control localities in the ENCELURB data. The ENCELURB initially included households deemed potentially eligible for the Oportunidades program as it was expanded to urban areas; some households did not become

[19] We do not test an expansion of point of service (POS) payment terminals because the data on POS terminals by municipality does not begin until 2011, toward the end of our study period.

Table 3.5: Supply-Side Response

	Total		Bansefi	
	ATMs	Branches	ATMs	Branches
Current quarter	–1.52	0.03	–0.01	–0.01
	(4.14)	(0.30)	(0.01)	(0.02)
1 quarter lag	0.01	0.02	–0.02	0.02
	(4.11)	(0.34)	(0.01)	(0.02)
2 quarter lag	–10.83	0.08	0.01	0.01
	(5.64)	(0.36)	(0.03)	(0.01)
3 quarter lag	–5.42	0.08	–0.03	0.02
	(2.98)	(0.26)	(0.02)	(0.02)
4 quarter lag	–0.74	0.42	0.00	–0.03
	(5.97)	(0.50)	(0.01)	(0.03)
1 quarter lead	–1.10	–0.12	–0.01	0.00
	(3.66)	(0.36)	(0.00)	(0.02)
2 quarter lead	–6.09	0.25	0.00	0.01
	(4.90)	(0.34)	(0.02)	(0.01)
3 quarter lead	–7.84	0.25	–0.03	–0.01
	(8.00)	(0.65)	(0.01)	(0.03)
4 quarter lead	7.58	0.59	–0.01	–0.06
	(10.32)	(0.94)	(0.03)	(0.05)
Mean control group	198.29	36.87	0.49	1.41
F-test of lags	1.26	0.20	0.68	0.96
[p-value]	[0.29]	[0.94]	[0.61]	[0.43]
F-test of leads	0.69	0.44	0.79	0.67
[p-value]	[0.60]	[0.78]	[0.53]	[0.62]
Municipality fixed effects	Yes	Yes	Yes	Yes
Quarter fixed effects	Yes	Yes	Yes	Yes

Notes: * indicates statistical significance at $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$. The table shows β_k from

$$y_{jt} = \lambda_j + \delta_t + \sum_{k=-4}^{4} \beta_k D_{j,t+k} + \varepsilon_{jt}$$

where y_{jt} is the number of ATMs or bank branches of any bank or of Bansefi in municipality j during quarter t, $D_{jt} = 1$ if municipality j has at least one locality with Oportunidades debit cards in quarter t. The F-test of lags tests $\beta_{-4} = \cdots = \beta_{-1} = 0$; the F-test of leads tests $\beta_1 = \cdots = \beta_4 = 0$.

beneficiaries (either they were deemed ineligible or did not take up the program). Because these non-beneficiaries were "potentially eligible" for the program to be included in the survey, they are similarly (though not quite as) poor compared to the Oportunidades beneficiaries who make up our main sample. Because they did not receive debit cards during the roll-out, due to not being Oportunidades beneficiaries, these individuals in treatment and control localities serve as a good placebo test for locality-level shocks. The results are presented in Figure 3.17a. The difference-in-differences estimates on consumption, income, and savings are all insignificant from 0, although due to the low number of non-Oportunidades beneficiaries in ENCELURB (382 households), the estimates are very noisy. Nevertheless, it is comforting that the point estimates are substantially close to 0 relative to the coefficients from our main sample, and the coefficients for consumption and savings actually have the *opposite* sign as the coefficients from the main regression (shown again in panel (c) for comparison). This suggests that, although the noisy placebo estimates' 95% confidence intervals do include the point estimates from our main sample, locality-level shocks do not explain the observed results.

Finally, we test for a salience effect of the cards themselves, where the card—which a beneficiary might carry with her in a wallet or purse—serves as a salient reminder of her savings goals. In some localities, beneficiaries received their benefits through Bansefi but did not have access to a Bansefi savings account (and thus had to withdraw all of their money each pay period at a Bansefi branch); in these localities, the government decree requiring all beneficiaries to receive benefits through a plastic card led to receiving benefits on a pre-paid card, still without access to a savings account. Again using ENCELURB, we find that in localities without savings accounts that switched to a pre-paid card prior to the last round of the survey compared to localities without savings accounts that did not switch prior, there was no differential effect on consumption, income, or savings. These estimates are again noisier than the results from the main sample (here we have 2300 households), but the DID coefficient from the placebo consumption regression is statistically significant at the 10% level from the coefficient from the corresponding consumption regression in the full sample.

Figure 3.17: Placebo Tests

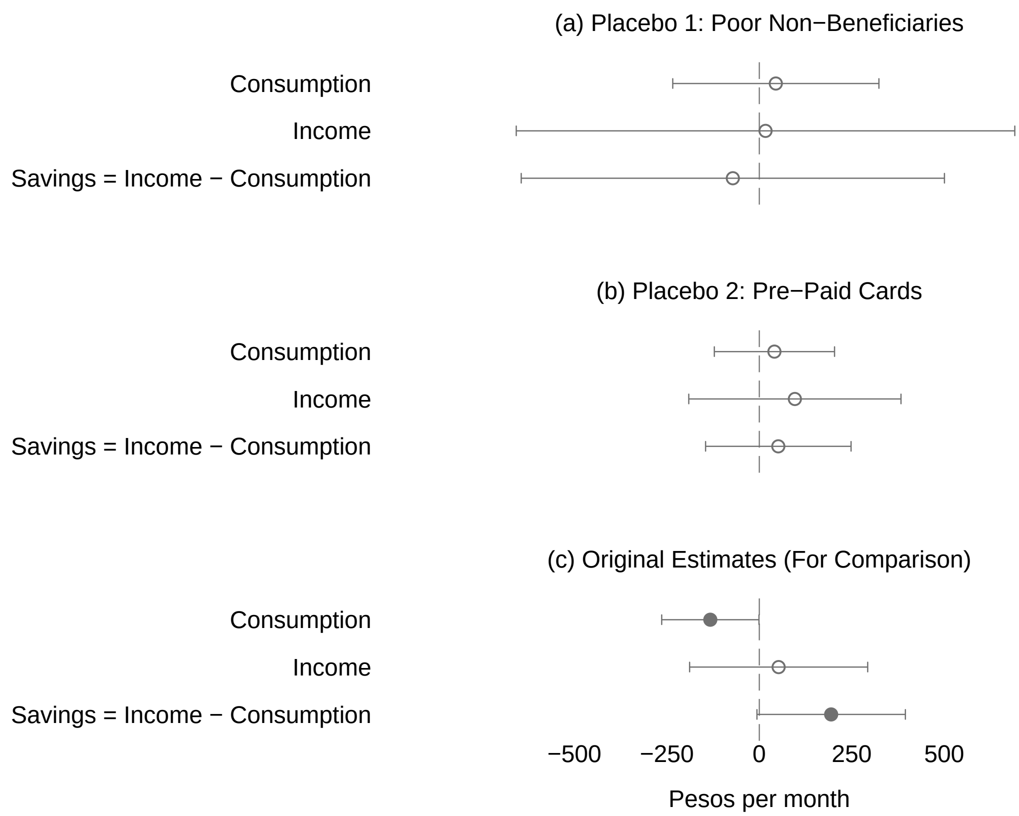

Sources: ENCELURB panel survey combined with administrative data on timing of card receipt and transfer payment histories for each surveyed beneficiary household.
Notes: (a) $N = 1415$ (number of households $= 382$); (b) $N = 8862$ (number of households $= 2300$); (c) $N = 11,275$ (number of households $= 2951$). Whiskers denote 95 percent confidence intervals. Black filled in circles indicate results that are significant at the 5 percent level, gray filled in circles at the 10 percent level, and hollow circles indicate results that are statistically insignificant from 0.

3.10 Conclusion

Although trust in financial institutions is by no means a sufficient condition to enable the poor to save, our findings suggest that it is a necessary condition. A lack of trust in banks could explain why a number of studies have found modest effects of offering savings accounts to the poor, even when these accounts have no fees or minimum balance requirements. Debit cards, a simple technology with high scale-up potential, provided beneficiaries of Mexico's large-scale cash transfer program Oportunidades with a mechanism to monitor banks by checking their balances at any bank's ATM; once beneficiaries built trust in banks, they began to save and their marginal propensity to save increased over time. We find that the observed increase in formal savings represents an increase in overall savings rather than a substitution from other forms of saving, and that beneficiaries reduce consumption of temptation goods, suggesting that saving informally is difficult and the use of financial institutions to save helps solve self-control problems.

The size of the savings effect, at 5% of income after one year with the debit card and 10% after two years, is larger than that of studies on various savings interventions such as subsidizing bank fees, increasing interest rates, and providing commitment savings devices. As a result, interventions that enable account holders to monitor banks and increase their trust in financial institutions may be a promising avenue to enable the poor to save in the formal financial sector. Debit cards and other forms of mobile money, which are simple, scalable technologies that are gaining traction in many developing countries, could thus be a highly effective means of increasing financial inclusion among millions of government cash transfer recipients worldwide.

Appendix A

Formal Axioms and Proofs from Chapter 1

A.1 FI Axioms

Consider pre- and post-fisc incomes $y_i^0, y_i^1 \in \mathbb{R}_+$ for each $i \in S$; denote the vectors of pre- and post-fisc income for these individuals by y^0 and y^1, both ordered by pre-fisc income y_i^0. Now consider income vectors for the same individuals under different pre- and post-fisc scenarios, denoted by x^0 and x^1, both ordered by pre-fisc income x_i^0. The sets of impoverished individuals in scenarios (y^0, y^1) and (x^0, x^1) are denoted $I_y \equiv \{i \in S \,|\, y_i^1 < y_i^0 \text{ and } y_i^1 < z\}$ and $I_x \equiv \{i \in S \,|\, x_i^1 < x_i^0 \text{ and } x_i^1 < z\}$. A measure of FI is a function $f : \bigcup_{n=1}^{\infty} \mathbb{R}_+^n \times \bigcup_{n=1}^{\infty} \mathbb{R}_+^n \times \mathbb{R}_+ \to \mathbb{R}$, which takes as arguments the pre- and post-fisc income vectors and the poverty line.

Axiom 1 (FI Monotonicity)**.** If $y_i^0 = x_i^0$ for all $i \in S$ and there exists $j \in I_y \cup I_x$ such that $y_j^1 > x_j^1$, while $y_k^1 = x_k^1$ for all $k \in I_y \cup I_x \setminus \{j\}$, then $f(y^0, y^1; z) < f(x^0, x^1; z)$.

Axiom 2 (Focus)**.** If $y_i^0 = x_i^0$ and $y_i^1 = x_i^1$ for all $i \in I_y \cup I_x$, then $f(y^0, y^1; z) = f(x^0, x^1; z)$.

Axiom 3 (Normalization)**.** $I_y = \varnothing \Rightarrow f(y^0, y^1; z) = 0$.

Axiom 4 (Continuity)**.** f is jointly continuous in y_i^0, y_i^1, and z.

Axiom 5 (Permutability). $f(y^0, y^1; z) = f(y^0_\sigma, y^1_\sigma; z)$ for any permutation function $\sigma : S \to S$, where $y^0_\sigma \equiv (y^0_{\sigma(1)}, \ldots, y^0_{\sigma(|S|)})$ and $y^1_\sigma \equiv (y^1_{\sigma(1)}, \ldots, y^1_{\sigma(|S|)})$.

Axiom 6 (Translation Invariance). $f(y^0 + \alpha 1_{|S|}, y^1 + \alpha 1_{|S|}; z + \alpha) = f(y^0, y^1; z)$ for all $\alpha \in \mathbb{R}$, where $1_{|S|}$ denotes a vector of ones with length $|S|$.

Axiom 7 (Linear Homogeneity). $f(\lambda y^0, \lambda y^1; \lambda z) = \lambda f(y^0, y^1; z)$ for all $\lambda \in \mathbb{R}_{++}$.

Axiom 8 (Subgroup Consistency). Partition S into m subsets $S_1, \ldots, S_m$, and denote the vectors of pre- and post-fisc incomes for individuals belonging to subset S_a, $a \in \{1, \ldots, m\}$, by y^0_a and y^1_a or x^0_a and x^1_a. If $f(y^0_a, y^1_a; z) < f(x^0_a, x^1_a; z)$ for some $a \in \{1, \ldots, m\}$ and $f(y^0_b, y^1_b; z) = f(x^0_b, x^1_b; z)$ for all $b \in \{1, \ldots, m\} \setminus \{a\}$, then $f(y^0, y^1; z) < f(x^0, x^1; z)$.

A.2 FGP Axioms

Let the sets of pre-fisc poor individuals experiencing fiscal gains under two scenarios be denoted $G_y \equiv \{i \in S \,|\, y^0_i < y^1_i \text{ and } y^0_i < z\}$ and $G_x \equiv \{i \in S \,|\, x^0_i < x^1_i \text{ and } x^0_i < z\}$. A measure of FGP is a function $g : \bigcup_{n=1}^{\infty} \mathbb{R}^n_+ \times \bigcup_{n=1}^{\infty} \mathbb{R}^n_+ \times \mathbb{R}_+ \to \mathbb{R}$, which takes as arguments the pre- and post-fisc income vectors and the poverty line.

Axiom 1′ (FGP Monotonicity). If $y^0_i = x^0_i$ for all $i \in S$ and there exists $j \in G_y \cup G_x$ such that $y^1_j < x^1_j$, while $y^1_k = x^1_k$ for all $k \in G_y \cup G_x \setminus \{j\}$, then $g(y^0, y^1; z) \leq g(x^0, x^1; z)$, with strict inequality if $y^1_j < z$.

The remaining axioms for FI are desirable for a measure of FGP as well, and carry over directly to FGP after replacing f with g, I_y with G_y, and I_x with G_x.

A.3 Proofs

Proof of Proposition 1 We begin with a lemma analogous to one of the propositions in Foster and Shorrocks (1991). To simplify notation, $y_a \equiv (y^0_a, y^1_a)$ for a subset S_a of a partition of S into m subgroups $a = 1, \ldots, m$; similarly, $x_a \equiv (x_a^0, x_a^1)$. We also define vectors $y_{-a}^t \equiv (y_b^t)_{b \neq a \in \{1,\ldots,m\}}$, $t \in \{0, 1\}$ as the vector of pre- or post-fisc incomes of all $i \notin S_a$ (similarly for x_{-a}^t) and $y_{-a} \equiv (y_{-a}^0, y_{-a}^1)$, $x_{-a} \equiv (x_{-a}^0, x_{-a}^1)$.

Lemma. $f(y_a, y_{-a}; z) \geq f(x_a, y_{-a}; z) \Rightarrow f(y_a, x_{-a}; z) \geq f(x_a, x_{-a}; z)$.

Proof. By subgroup consistency, $f(y_a, y_{-a}; z) \geq f(x_a, y_{-a}; z) \Rightarrow f(y_a; z) \geq f(x_a; z)$. (Suppose not. Then $f(y_a; z) < f(x_a; z)$, which by subgroup consistency implies $f(y_a, y_{-a}; z) < f(x_a, y_{-a}; z)$, a contradiction.) $f(y_a; z) \geq f(x_a; z)$ implies either $f(y_a; z) > f(x_a; z)$ or $f(y_a; z) = f(x_a; z)$. In the former case, it immediately follows by subgroup consistency that $f(y_a, x_{-a}; z) \geq f(x_a, x_{-a}; z)$. In the latter case, the implication is shown by contradiction. Suppose that $f(y_a, x_{-a}; z) < f(x_a, x_{-a}; z)$. Then by subgroup consistency we have (since $f(y_a; z) = f(x_a; z)$) $f(y_a, x_{-a}, x_a; z) < f(x_a, x_{-a}, y_a; z)$, which contradicts permutability. □

This lemma shows that a subgroup-consistent and permutable measure of FI is separable by group, using a definition of separability analogous to that used for preferences in the utility literature. Because the lemma can be reiterated within any particular subgroup to further separate individuals in that subgroup, we have that each set of individuals is separable (which is analogous to the "each set of sectors is separable" requirement in Gorman (1968, p. 368)). Hence, from Debreu (1960, theorem 3), there exists a continuous FI function determined up to an increasing linear transformation of the form

$$f(y^0, y^1; z) = \alpha + \beta \sum_{i \in S} \phi_i(y_i^0, y_i^1, z)$$

where ϕ_i is a real-valued function for each $i \in S$. The additional requirement for Debreu's (1960) proof that more than two of the $|S|$ elements of S are essential is satisfied as long as $|S| \geq 3$ and f is non-constant on $[0, z]$, which in turn is implied by monotonicity as long as at least one individual is impoverished.[1]

[1] The assumptions of at least three individuals in society and at least one impoverished individual are

Permutability implies that $\phi_i = \phi_j$ for all $i, j \in S$, so we have $f(y^0, y^1; z) = \alpha + \beta \sum \phi(y_i^0, y_i^1, z)$ where ϕ is a real-valued function. By the focus and normalization axioms:

$$\phi(y_i^0, y_i^1, z) = \begin{cases} \tilde{\phi}(y_i^0, y_i^1, z) & \text{if } y_i^1 < y_i^0 \text{ and } y_i^1 < z \\ 0 & \text{otherwise.} \end{cases} \tag{A.1}$$

By the continuity of f, ϕ and $\tilde{\phi}$ must also be continuous. Consider an individual with $y_i^0 > z$ and $y_i^1 = z$. Since y_i^1 is not less than z, i is not impoverished, so by (A.1), $\phi(y_i^0, y_i^1, z) = 0$. Now consider an alternative situation where $\tilde{y}_i^1 = z - \epsilon$ for a sufficiently small $\epsilon > 0$. In this scenario, $\tilde{\phi}$ cannot be a direct function of y_i^0 or ϕ would be discontinuous at z; instead, $\tilde{\phi}$ must be a direct function of just y_i^1 and z so that an infinitesimal decrease in y_i^1 below z results in an infinitesimal increase in ϕ. By a similar argument, for an individual with $y_i^0 < z$, $y_i^1 = y_i^0$, and $\tilde{y}_i^1 = y_i^0 - \epsilon$, $\tilde{\phi}$ cannot be a direct function of z and instead must directly depend only on y_i^1 and y_i^0 so that an infinitesimal decrease in y_i^1 below $y_i^0 < z$ results in an infinitesimal increase in ϕ.

Given this, we can rewrite $\tilde{\phi}(y_i^0, y_i^1, z) = \tilde{\phi}(\min\{y_i^0, z\}, y_i^1)$. Since $\tilde{\phi}$ is only defined for those who are impoverished (i.e., those for whom $\min\{y_i^0, y_i^1, z\} = y_i^1$), we have

$$\begin{aligned} \tilde{\phi}(y_i^0, y_i^1, z) &= \tilde{\phi}(\min\{y_i^0, z\}, \min\{y_i^0, y_i^1, z\}) && \text{(A.2)} \\ &= \tilde{\phi}(\min\{y_i^0, z\} - \min\{y_i^0, y_i^1, z\}, 0) && \text{(A.3)} \\ &= \big(\min\{y_i^0, z\} - \min\{y_i^0, y_i^1, z\}\big)\tilde{\phi}(1, 0) && \text{(A.4)} \end{aligned}$$

where (A.3) follows from translation invariance and (A.4) from linear homogeneity. Noting

innocuous for any real-world application.

that $\tilde{\phi}(1,0)$ is a constant (that is positive by monotonicity) and denoting it γ, we have

$$\phi(y_i^0, y_i^1, z) = \begin{cases} \big(\min\{y_i^0, z\} - \min\{y_i^0, y_i^1, z\}\big)\gamma & \text{if } i \in I_y \\ 0 & \text{otherwise.} \end{cases}$$

For $i \notin I_y$ we can also write $\phi(y_i^0, y_i^1, z) = (\min\{y_i^0, z\} - \min\{y_i^0, y_i^1, z\})\gamma$ since the non-impoverished are either non-poor before taxes and transfers and non-poor after ($\Rightarrow \min\{y_i^0, z\} = \min\{y_i^0, y_i^1, z\} = z$) or poor before taxes and transfers but do not lose income to the fiscal system ($\Rightarrow \min\{y_i^0, z\} = \min\{y_i^0, y_i^1, z\} = y_i^0$). Therefore $f(y^0, y^1; z) = \alpha + \beta\gamma \sum_{i \in S} (\min\{y_i^0, z\} - \min\{y_i^0, y_i^1, z\})$. By normalization, $\alpha = 0$, which completes the proof. □

Proof of Proposition 2 (a)⇔(b) follows immediately from Proposition 1. For (b)⇔(c), we begin by defining Foster and Rothbaum's (2014) second order downward mobility dominance.

Definition. (y^0, y^1) *second order downward mobility dominates* (x^0, x^1) on $[0, z^+]$ if

$$\int_0^z m(y^0, y^1; c)dc < \int_0^z m(x^0, x^1; c)dc \; \forall \, z \in [0, z^+],$$

where $m(y^0, y^1; z) = |S|^{-1} \sum_{i \in S} \mathbb{I}(y_i^1 < z < y_i^0)$ is Foster and Rothbaum's (2014) *downward mobility curve*, measuring the proportion of the population that begins with income above each poverty line and ends with income below the line.

A sufficient condition for (b)⇔(c) is $f(y^0, y^1; z) \propto \int_0^z m(y^0, y^1; c)dc$. For a given poverty line $z = \hat{z}$, partition the set S into four subsets: $S_1 = \{i \in S \,|\, y_i^1 < y_i^0 < \hat{z}\}$, $S_2 = \{i \in S \,|\, y_i^1 < \hat{z} \leq y_i^0\}$, $S_3 = \{i \in S \,|\, y_i^0 \geq \hat{z}, y_i^1 \geq \hat{z}\}$, $S_4 = \{i \in S \,|\, y_i^0 < \hat{z}, y_i^0 \leq y_i^1\}$. For any subset $S_a \subset S$, denote $f_a(\cdot; z) \equiv \kappa \sum_{i \in S_a} \big(\min\{y_i^0, z\} - \min\{y_i^0, y_i^1, z\}\big)$ and $m_a(\cdot; z) \equiv |S|^{-1} \sum_{i \in S_a} \mathbb{I}(y_i^1 < z < y_i^0)$.

Each $i \in S_1$ experiences downward mobility on the interval $[0, \hat{z}]$ for all $z \in (y_i^1, y_i^0)$ $\Rightarrow$ individual $i \in S_1$ increases $m_1(\cdot; z)$ by $|S|^{-1}$ for $z \in (y_i^1, y_i^0)$ and by zero for $z \leq y_i^1$ and $z \geq y_i^0 \Rightarrow$ individual $i \in S_1$ increases $\int_0^{\hat{z}} m_1(\cdot; c)dc$ by $|S|^{-1}(y_i^0 - y_i^1)$. Summing over all $i \in S_1$, $\int_0^{\hat{z}} m_1(\cdot; c)dc = \sum_{i \in S_1} |S|^{-1}(y_i^0 - y_i^1)$.

$y_i^1 < y_i^0 < \hat{z}\ \forall\, i \in S_1 \Rightarrow f_1(\cdot; \hat{z}) = \kappa \sum_{i \in S_1}(y_i^0 - y_i^1) \Rightarrow$

$$f_1(\cdot; \hat{z}) = \kappa |S| \int_0^{\hat{z}} m_1(\cdot; c)dc. \tag{A.5}$$

Each $i \in S_2$ experiences downward mobility on the interval $[0, \hat{z}]$ for all $z \in (y_i^1, \hat{z}]$, which increases $m_2(\cdot; z)$ by $|S|^{-1}$ for $z \in (y_i^1, \hat{z}]$ and by zero for all other $z \Rightarrow$ individual $i \in S_2$ increases $\int_0^{\hat{z}} m_2(\cdot; c)dc$ by $|S|^{-1}(\hat{z} - y_i^1)$. Summing over all $i \in S_2$, $\int_0^{\hat{z}} m_2(\cdot; c)dc = \sum_{i \in S_2} |S|^{-1}(\hat{z} - y_i^1)$.

$y_i^1 < \hat{z} \leq y_i^0\ \forall\, i \in S_2 \Rightarrow f_2(\hat{z}, \cdot) = \kappa \sum_{i \in S_2}(\hat{z} - y_i^1) \Rightarrow$

$$f_2(\cdot; \hat{z}) = \kappa |S| \int_0^{\hat{z}} m_2(\cdot; c)dc. \tag{A.6}$$

Each $i \in S_3$ does not experience downward mobility on the interval $[0, \hat{z}]$; summing over all $i \in S_3$ and integrating over our domain, we have $\int_0^{\hat{z}} m_3(\cdot; c)dc = 0$. $y_i^0 \geq \hat{z}$ and $y_i^1 \geq \hat{z}\ \forall\, i \in S_3 \Rightarrow$

$$f_3(\cdot; \hat{z}) = \kappa \sum_{i \in S_3}(\hat{z} - \hat{z}) = 0 = \kappa |S| \int_0^{\hat{z}} m_3(\cdot; c)dc. \tag{A.7}$$

Similarly $\int_0^{\hat{z}} m_4(\cdot; c)dc = 0$ because each $i \in S_4$ does not experience downward mobility on $[0, \hat{z}]$. $y_i^0 < \hat{z}$ and $y_i^0 \leq y_i^1\ \forall\, i \in S_4 \Rightarrow$

$$f_4(\cdot; \hat{z}) = \kappa \sum_{i \in S_4}(y_i^0 - y_i^0) = 0 = \kappa |S| \int_0^{\hat{z}} m_4(\cdot; c)dc. \tag{A.8}$$

Given the definitions of $f_a(\cdot; z)$ and $m_a(\cdot; z)$ and that $S = S_1 \cup S_2 \cup S_3 \cup S_4$ and $S_1 \cap$

$S_2 \cap S_3 \cap S_4 = \varnothing$, we have $f(y^0, y^1; z) = \sum_{a=1}^{4} f_a(\cdot; z)$ and $m(y^0, y^1; z) = \sum_{a=1}^{4} m_a(\cdot; z)$. Hence, by (A.5)–(A.8), $f(y^0, y^1; \hat{z}) = \kappa |S| \int_0^{\hat{z}} m(\cdot; c) dc$. This holds for all $\hat{z} \in [0, z^+]$ since the choice of $\hat{z}$ was arbitrary, which completes the proof. □

Proof of Proposition 3 Analogous to the proof of Proposition 1 for FI. □

Proof of Proposition 4 Given in text. □

Appendix B

A Three-Dimensional Model

We now demonstrate that the main forces we identify in our main framework—especially the conflict between benevolence and self-interest—carry through in a model with three line ministers (further generalizations to n ministers are immediate). Having three line ministers does not change the FOC from what we have in Section 2.3.1: $x_j = \max\{r_j - P\beta_j c_j/\pi_j, 0\}$, for $j = 1, \ldots, 3$. Here, we focus only on equilibria in which all agents contribute a positive amount, so, for $j = 1, \ldots, 3$,

$$x_j = r_j - P\beta_j \frac{c_j}{\pi_j}. \tag{B.1}$$

Now, prestige is $P = \sum_{j=1}^{3} \frac{\pi_j x_j}{c_j}$, so, using (B.1), equilibrium prestige is

$$P^* = \frac{\sum_{j=1}^{3} \frac{r_j \pi_j}{c_j}}{1 + \sum_{j=1}^{3} \beta_j}. \tag{B.2}$$

Substitution of (B.2) into (B.1) yields total equilibrium antipoverty expenditures equal to

$$\sum_{j=1}^{3} x_j^* = \sum_{j=1}^{3} r_j - \frac{\left(\sum_{j=1}^{3} \frac{r_j \pi_j}{c_j}\right)\left(\sum_{j=1}^{3} \frac{\beta_j c_j}{\pi_j}\right)}{1 + \sum_{j=1}^{3} \beta_j}. \tag{B.3}$$

The comparison between (B.2) and (B.3) immediately reveals the same conflict between benevolence and self-interest previously explored. For instance, a self-interested finance minister maximizes $\sum_{j=1}^{3} \frac{r_j \pi_j}{c_j}$, by assigning as many resources as possible to the line minister with the largest π_j to c_j ratio. And a benevolent finance minister *minimizes* $\sum_{j=1}^{3} \frac{r_j \pi_j}{c_j}$, by assigning as many resources as possible to the line minister with the smallest π_j to c_j ratio.[1]

Similarly, the decisions of the technical committee follow by-and-large the same principles we have identified in the main model. Examination of (B.2) reveals that a self-interested committee assigns as much weight as possible to the dimension with the largest r_j to c_j ratio; and this results in the same prescription of Section 3.3: use a unidimensional poverty index, if possible. In contrast, a benevolent committee maximizes $\sum_{j=1}^{3} x_j^*$ in (B.3) under the constraint that $\sum_{j=1}^{3} \pi_j = 1$. The FOC of the Lagrangean, after simplifications, is:

$$\frac{d}{d\pi_i}\left[\left(\sum_{j=1}^{3} \frac{r_j \pi_j}{c_j}\right)\left(\sum_{j=1}^{3} \frac{\beta_j c_j}{\pi_j}\right)\right] = \lambda, \quad \text{for } i = 1, 2, 3.$$

Simple differentiation yields

$$-\frac{\beta_i c_i}{\pi_i}\left(\sum_{j=1}^{3} \frac{r_j \pi_j}{c_j}\right) + \frac{r_i \pi_i}{c_i}\left(\sum_{j=1}^{3} \frac{\beta_j c_j}{\pi_j}\right) = \pi_i \lambda,$$

and after summing the above over $i = 1, 2, 3$ we obtain $0 = \lambda$. Then, dividing the FOCs above one by the other we obtain the optimal weight distribution $(\pi_1^*, \pi_2^*, \pi_3^*)$, which equals $\left(\left(1 + \frac{c_3}{c_1}\sqrt{\frac{\beta_3 r_1}{\beta_1 r_3}} + \frac{c_2}{c_1}\sqrt{\frac{\beta_2 r_1}{\beta_1 r_2}}\right)^{-1}, \left(1 + \frac{c_3}{c_2}\sqrt{\frac{\beta_3 r_2}{\beta_2 r_3}} + \frac{c_1}{c_2}\sqrt{\frac{\beta_1 r_2}{\beta_2 r_1}}\right)^{-1}, \left(1 + \frac{c_2}{c_3}\sqrt{\frac{\beta_2 r_3}{\beta_3 r_2}} + \frac{c_1}{c_3}\sqrt{\frac{\beta_1 r_3}{\beta_3 r_1}}\right)^{-1}\right)$.

Focusing for instance on dimension 1, we see that π_1 increases when c_1 increases, c_2 or c_3 decrease, r_1/β_1 increases, and r_2/β_2 or r_3/β_3 increase. Analogously to what determined in the main model, a benevolent committee uses a multidimensional measure in which each dimension's weight increases the harder it is to reduce deprivations for the line minister in

[1] Interestingly, the allocation to the "middle" dimension may remain the same, so the difference between the outcomes of benevolence and self-interest is especially evident when comparing the dimension with the largest π_j to c_j ratio vs. that with the smallest such ratio.

charge of this dimension.

Bibliography

Afkar, R., Jellema, J., Wai-Poi, M., forthcoming. The distributional impact of fiscal policy in Indonesia, in: Inchauste, G., Lustig, N. (Eds.), The Distributional Impact of Fiscal Policy: Experience from Developing Countries. World Bank, Washington, D.C.

Akbas, M., Ariely, D., Robalino, D.A., Weber, M., 2015. How to help the poor to save a bit: Evidence from a field experiment in Kenya. Working Paper. URL: http://sites.duke.edu/merveakbas/files/2014/08/How-to-help-the-poor-.pdf.

Alderman, H., 2002. Do local officials know something we don't? Decentralizatoin of targeted transfers in Albania. Journal of Public Economics 83, 375–404.

Algan, Y., Cahuc, P., 2010. Inherited trust and growth. American Economic Review 100, 2060–2092.

Alkire, S., Foster, J., 2011a. Counting and multidimensional poverty measurement. Journal of Public Economics 95, 476–487.

Alkire, S., Foster, J., 2011b. Understandings and misunderstandings of multidimensional poverty measurement. Journal of Economic Inequality 9, 289–314.

Alkire, S., Foster, J., 2016. Dimensional and distributional contributions to multidimensional poverty. OPHI Working Paper 100. URL: http://www.ophi.org.uk/wp-content/uploads/OPHIWP100_2.pdf.

Alkire, S., Foster, J., Santos, M.E., 2011. Where did identification go? Journal of Economic Inequality 9, 501–505.

Alkire, S., Foster, J.E., Seth, S., Santos, M.E., Roche, J.M., Ballon, P., 2015. Multidimensional Poverty Measurement and Analysis. Oxford University Press, Oxford.

Alkire, S., Santos, M.E., 2014. Measuring acute poverty in the developing world: Robustness and scope of the Multidimensional Poverty Index. World Development 59, 251–274.

Alkire, S., Seth, S., 2013. Selecting a targeting method to identify BPL households in India. Social Indicators Research 112, 417–446.

Alkire, S., Seth, S., 2015. Multidimensional poverty reduction in India between 1999 and 2006: Where and how? World Development 72, 93–108.

Alvarez, F., Lippi, F., 2009. Financial innovation and the transactions demand for cash. Econometrica 77, 363–402.

Amiel, Y., Cowell, F., 1994. Monotonicity, dominance and the Pareto principle. Economics Letters 45, 447–450.

Angelucci, M., Attanasio, O., 2013. The demand for food of poor urban Mexican households: Understanding policy impacts using structural models. American Economic Journal: Economic Policy 5, 146–178.

Angelucci, M., De Giorgi, G., 2009. Indirect effects of an aid program: How do cash transfers affect ineligibles' consumption? American Economic Review 99, 486–508.

Angelucci, M., Karlan, D.S., Zinman, J., 2015. Microcredit impacts: Evidence from a randomized microcredit program placement experiment by Compartamos Banco. American Economic Journal: Applied Economics 7, 151–182.

Angulo, R., 2011. Índice de pobreza multidimensional para Colombia. Departamento Nacional de Planeación presentation.

Aranda, R., Scott, J., 2015. CEQ Master Workbook for Mexico. Mimeo.

Arellano, M., Bover, O., 1995. Another look at the instrumental variable estimation of error-component models. Journal of Econometrics 68, 29–51.

Aristy-Escuder, J., Cabrera, M., Sánchez-Martín, M.E., forthcoming. An analysis of fiscal policy and income redistribution in the Dominican Republic. CEQ Working Paper 37.

Arrow, K., 1972. Gifts and exchanges. Philosophy & Public Affairs 1, 343–362.

Arunatilake, N., Hewawasam, J., Gunasekara, N., forthcoming. The distributional impact of fiscal policy in Sri Lanka, in: Inchauste, G., Lustig, N. (Eds.), The Distributional Impact of Fiscal Policy: Experience from Developing Countries. World Bank, Washington, D.C.

Ashraf, N., 2009. Spousal control and intra-household decision making: An experimental study in the Philippines. American Economic Review 99, 1245–1277.

Ashraf, N., Karlan, D., Yin, W., 2006. Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. Quarterly Journal of Economics 121, 635–672.

Ashraf, N., Karlan, D., Yin, W., 2015. Savings in transnational households: A field experiment among migrants from El Salvador. Review of Economics and Statistics 97, 332–351.

Atkinson, A.B., 1970. On the measurement of inequality. Journal of Economic Theory 2, 244–263.

Atkinson, A.B., 1987. On the measurement of poverty. Econometrica 55, 749–764.

Atkinson, A.B., 2003. Multidimensional deprivation: Contrasting social welfare and counting approaches. Journal of Economic Inequality 1, 51–65.

Atkinson, A.B., Bourguignon, F., 1982. The comparison of multi-dimensioned distributions of economic status. Review of Economic Studies 49, 183–201.

Atkinson, J., de Janvry, A., McIntosh, C., Sadoulet, E., 2013. Prompting microfinance borrowers to save: A field experiment from guatemala. Economic Development and Cultural Change 62, 21–64.

Attanasio, O., Di Maro, V., Lechene, V., Phillips, D., 2013. Welfare consequences of food prices increases: Evidence from rural Mexico. Journal of Development Economics 104, 136–151.

Attanasio, O., Pellerano, L., Reyes, S.P., 2009. Building trust? Conditional cash transfer programmes and social capital. Fiscal Studies 30, 139–177.

Attanasio, O.P., Guiso, L., Jappelli, T., 2002. The demand for money, financial innovation, and the welfare cost of inflation: An analysis with household data. Journal of Political Economy 110, 317–351.

Augsburg, B., de Haas, R., Harmgart, H., Meghir, C., 2015. The impacts of microcredit: Evidence from Bosnia and Herzegovina. American Economic Journal: Applied Economics 7, 183–203.

Azevedo, V., Robles, M., 2013. Multidimensional targeting: Identifying beneficiaries of conditional cash transfer programs. Social Indicators Research 112, 447–475.

Baer, W., Galvão, A.F., 2008. Tax burden, government expenditures and income distribution in Brazil. Quarterly Review of Economics and Finance 48, 345–358.

Baland, J.M., Guirkinger, C., Mali, C., 2011. Pretending to be poor: Borrowing to escape forced solidarity in Cameroon. Economic Development and Cultural Change 60, 1–16.

Banerjee, A., Mullainathan, S., 2010. The shape of temptation: Implications for the economic lives of the poor. NBER Working Paper 15973. URL: http://www.nber.org/papers/w15973.

Banerjee, A.V., Duflo, E., Glennester, R., Kinnan, C., 2015. The miracle of microfinance? Evidence from a randomized evaluation. American Economic Journal: Applied Economics 7, 22–53.

Banfield, E.C., 1958. The Moral Basis of a Backward Society. The Free Press, Glencoe, IL.

Bardhan, P., Mookherjee, D., 2005. Decentralizing antipoverty program delivery in developing countries. Journal of Public Economics 89, 675–704.

Baumol, W.J., 1952. The transactions demand for cash: An inventory theoretic approach. Quarterly Journal of Economics 66, 545–556.

Behrman, J.R., Gallardo-García, J., Parker, S.W., Todd, P.E., Vélez-Grajales, V., 2012. Are conditional cash transfers effective in urban areas? Evidence from Mexico. Education Economics 20, 233–259.

Beneke, M., Lustig, N., Oliva, J.A., 2015. El impacto de los impuestos y el gasto social en la desigualded y la pobreza en El Salvador. CEQ Working Paper 26.

Bergstrom, T., Blume, L., Varian, H., 1986. On the private provision of public goods. Journal of Public Economics 29, 25–49.

Bernheim, B.D., Ray, D., Yeltekin, S., 2015. Poverty and self-control. Econometrica 83, 1877–1911.

Besley, T., Kanbur, R., 1988. Food subsidies and poverty alleviation. The Economic Journal 98, 701–719.

Birdsall, N., 2011. Comment on multi-dimensional indices. Journal of Economic Inequality 9, 489–491.

Blackorby, C., Donaldson, D., 1980. Ethical indices for the measurement of poverty. Econometrica 48, 1053–1060.

Blundell, R., Bond, S., 1998. Initial conditions and moment restriction sin dynamic panel data models. Journal of Econometrics 87, 115–143.

Blundell, R., Bond, S., Windmeijer, F., 2001. Estimation in dynamic panel data models: improving on the performance of the standard GMM estimators, in: Baltagi, B.H., Fomby, T.B., Hill, R.C. (Eds.), Nonstationary Panels, Panel Cointegration, and Dynamic Panels. Emerald Group Publishing, Bingley, UK. volume 15 of *Advances in Econometrics*, pp. 53–91.

Bohnet, I., Herrmann, B., Zeckhauser, R., 2010. Trust and the reference points for trustworthiness in Gulf and Western countries. Quarterly Journal of Economics 125, 811–828.

Bold, C., Porteous, D., Rotman, S., 2012. Social cash transfers and financial inclusion: Evidence from four countries. Consultative Group to Assist the Poor 77, 1–28.

Bossert, W., Chakravarty, S.R., D'Ambrosio, C., 2013. Multidimensional poverty and material deprivation with discrete data. Review of Income and Wealth 59, 29–43.

Bourguignon, F., 2011a. Non-anonymous growth incidence curves, income mobility and social welfare dominance. Journal of Economic Inequality 9, 605–627.

Bourguignon, F., 2011b. Status quo in the welfare analysis of tax reforms. Review of Income and Wealth 57, 603–621.

Bourguignon, F., Chakravarty, S.R., 2003. The measurement of multidimensional poverty. Journal of Economic Inequality 1, 25–49.

Bourguignon, F., Fields, G.S., 1997. Discontinuous losses from poverty, generalized P_α measures and optimal transfers to the poor. Journal of Public Economics 63, 155–175.

Brune, L., Giné, X., Goldberg, J., Yang, D., 2016. Facilitating savings for agriculture: Field experimental evidence from malawi. Economic Development and Cultural Change 64, 187–220.

Bruno, G.S.F., 2005. Approximating the bias of the LSDV estimator for dynamic unbalanced panel data models. Economics Letters 87, 361–366.

Bryan, G., Karlan, D., Nelson, S., 2010. Commitment devices. Annual Review of Economics 2, 671–698.

Bun, M.J.G., Kiviet, J.F., 2006. The effects of dynamic feedbacks on LS and MM estimator accuracy in panel data models. Journal of Econometrics 132, 409–444.

Cabrera, M., Lustig, N., Morán, H., 2015. Fiscal policy, inequality, and the ethnic divide in Guatemala. World Development 76, 263–279.

Carvalho, L.S., Meier, S., Wang, S.W., 2016. Poverty and economic decision-making: Evidence from changes in financial resources at payday. American Economic Review 106, 260–284.

Cavapozzi, D., Han, W., Miniaci, R., 2015. Alternative weighting structures for multidimensional poverty assessment. Journal of Economic Inequality 13, 425–447.

Center for Effective Global Action, 2015. Deep dive: Using pre-paid Visa cards to disburse government-to-person (G2P) payments in the Dominican Republic. Mimeo.

Chakravarty, S.R., 1983. A new index of poverty. Mathematical Social Sciences 6, 307–313.

Chakravarty, S.R., Dutta, B., 1987. A note on measures of distance between income distributions. Journal of Economic Theory 41, 185–188.

Chakravarty, S.R., Mukherjee, D., Renade, R.R., 1998. On the family of subgroup and factor decomposable measures of multidimensional poverty. Research on Economic Inequality 8, 175–194.

Chamon, M., Liu, K., Prasad, E., 2013. Income uncertainty and household savings in China. Journal of Development Economics 105, 164–177.

Chen, S., Ravallion, M., 2010. The developing world is poorer than we thought, but no less successful in the fight against poverty. Quarterly Journal of Economics 125, 1577–1625.

Chowdhury, S., Squire, L., 2006. Setting weights for aggregate indices: An application to the Commitment to Development Index and Human Development Index. Journal of Public Economic Theory 9, 201–219.

Citro, C.F., Michael, R.T., 1995. Measuring Poverty: A New Approach. National Academy Press, Washington, D.C.

CNBV, 2014. Comisión Nacional Bancaria y de Valores comunicado de prensa 094/2014. URL: http://www.cnbv.gob.mx/SECTORES-SUPERVISADOS/SECTOR-POPULAR/CasoFicrea/Comunicado_de_Prensa_Revocaci%C3%B3n_FICREA.pdf.

Cohn, A., Fehr, E., Maréchal, M.A., 2014. Business culture and dishonesty in the banking industry. Nature 516, 86–89.

Collins, D., Morduch, J., Rutherford, S., Ruthven, O., 2009. Portfolios of the Poor: How the World's Poor Live on $2 a Day. Princeton University Press.

CONEVAL, 2010. Metodología para la medición multidimensional de la pobreza en México. CONEVAL, Mexico City.

Conti, G., Heckman, J., Urzua, S., 2010. The education–health gradient. American Economic Review: Papers and Proceedings 100, 234–238.

Corbacho, A., Cibils, V.F., Lora, E., 2013. More than Revenue: Taxation as a Development Tool. Palgrave Macmillan, New York.

Cornes, R., Hartley, R., 2007. Weak links, good shots and other public good games: Building on BBV. Journal of Public Economics 91, 1684–1707.

Cowell, F.A., 1985. Measures of distributional change: An axiomatic approach. Review of Economic Studies 52, 135–151.

Cowell, F.A., Victoria-Feser, M.P., 2002. Welfare rankings in the presence of contaminated data. Econometrica 70, 1221–1233.

Davidson, R., Duclos, J.Y., 2000. Statistical inference for stochastic dominance and for the measurement of poverty and inequality. Econometrica 68, 1435–1464.

Deaton, A., 2011. Measuring development: Different data, different conclusions?, in: Bourguignon, F. (Ed.), Measure for Measure: How Well Do We Measure Development?. Agence Française de Développement, Paris, pp. 11–52.

Debreu, G., 1960. Topological methods in cardinal utility theory, in: Arrow, K.J., Karlin, S., Suppes, P. (Eds.), Mathematical Methods in the Social Sciences. Stanford University Press, Palo Alto, pp. 16–26.

Decancq, K., Fleurbaey, M., Schokkaer, E., 2015. Happiness, equivalent incomes and respect for individual preferences. Economica 82, 1082–1106.

Decancq, K., Lugo, M.A., 2013. Weights in multidimensional indices of wellbeing: An overview. Econometric Reviews 32, 7–34.

Decancq, K., Van Ootegem, L., Verhofstadt, E., 2013. What if we voted on the weights of a multidimensional well-being index? An illustration with Flemish data. Fiscal Studies 34, 315–32.

DeFina, R.H., Thanawala, K., 2004. International evidence on the impact of transfers and taxes on alternative poverty indexes. Social Science Research 33, 311–338.

Demirgüç-Kunt, A., Klapper, L., Singer, D., Oudheusden, P.V., 2015. The Global Findex Database 2014: Measuring financial inclusion around the world. Policy Research Working Ppaer 7255.

Deutsch, J., Silber, J., 2005. Measuring multidimensional poverty: An empirical comparison of various approaches. Review of Income and Wealth 51, 145–174.

Dewilde, C., 2004. The multidimensional measurement of poverty in Belgium and Britain: A categorical approach. Social Indicators Research 68, 331–369.

Duclos, J.Y., 1997. Measuring progressivity and inequality. Research on Economic Inequality 7, 19–37.

Duclos, J.Y., 2008. Horizontal and vertical equity, in: Durlauf, S.N., Blume, L.E. (Eds.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, New York.

Duclos, J.Y., 2009. What is pro-poor? Social Choice and Welfare 32, 37–58.

Duclos, J.Y., Makdissi, P., Wodon, Q., 2008. Socially improving tax reforms. International Economic Review 49, 1505–1537.

Duclos, J.Y., Sahn, D.E., Younger, S.D., 2006. Robust multidimensional poverty comparisons. The Economic Journal 116, 943–968.

Duclos, J.Y., Sahn, D.E., Younger, S.D., 2011. Partial multidimensional inequality orderings. Journal of Public Economics 95, 225–238.

Duclos, J.Y., Tiberti, L., forthcoming. Multidimensional poverty indices: A critical assessment, in: Adler, M., Fleurbaey, M. (Eds.), The Oxford Handbook of Well-Being and Public Policy. Oxford University Press, Oxford.

Duclos, J.Y., Tiberti, L., Araar, A., 2014. Multidimensional poverty targeting. FERDI working paper 102. URL: `http://www.ferdi.fr/sites/www.ferdi.fr/files/publication/fichiers/wp102_ferdi_duclos_and_al_web.pdf`.

Dupas, P., Green, S., Keats, A., Robinson, J., 2016. Challenges in banking the rural poor: Evidence from Kenya's Western Province, in: Edwards, S., Johnson, S., Weil, D.N. (Eds.), Modernization and Development. University of Chicago Press, Chicago. volume 3 of *African Successes*.

Dupas, P., Robinson, J., 2013a. Savings constraints and microenterprise development: Evidence from a field experiment in Kenya. American Economic Journal: Applied Economics 5, 163–192.

Dupas, P., Robinson, J., 2013b. Why don't the poor save more? Evidence from health savings experiments. American Economic Review 103, 1138–1171.

Ebert, U., 1984. Measures of distance between income distributions. Journal of Economic Theory 32, 266–274.

Ebert, U., 1997. Social welfare when needs differ: An axiomatic approach. Economica 64, 233–244.

Ebrill, L.P., Keen, M., Summers, V.P., 2001. The Modern VAT. International Monetary Fund, Washington, D.C.

Engel, E.M.R.A., Galetovic, A., Raddatz, C.E., 1999. Taxes and income distribution in Chile: Some unpleasant redistributive arithmetic. Journal of Development Economics 59, 155–192.

Fellman, J., Jäntti, M., Lambert, P.J., 1999. Optimal tax-transfer systems and redistributive policy. The Scandinavian Journal of Economics 101, 115–126.

Ferraz, C., Finan, F., Moreira, D.B., 2012. Corrupting learning: Evidence from missing federal education funds in brazil. Journal of Public Economics 96, 712–726.

Ferreira, F.H.G., Messina, J., Rigolini, J., López-Calva, L.F., Lugo, M.A., Vakis, R., 2013. Economic Mobility and the Rise of the Latin American Middle Class. The World Bank, Washington, D.C.

Fields, G.S., 2001. Distribution and Development. Russell Sage Foundation and MIT Press, Cambridge.

Fields, G.S., Fei, J.C.H., 1978. On inequality comparisons. Econometrica 46, 303–316.

Fields, G.S., Ok, E.A., 1996. The meaning and measurement of income mobility. Journal of Economic Theory 71, 349–377.

Fields, G.S., Ok, E.A., 1999. Measuring movement of incomes. Economica 66, 455–471.

Foster, J., 2006. Poverty indices, in: de Janvry, A., Kanbur, R. (Eds.), Poverty, Inequality and Development: Essays in Honor of Erik Thorbecke. Kluwer Academic, Norwell, pp. 41–66.

Foster, J., Greer, J., Thorbecke, E., 1984. A class of decomposable poverty measures. Econometrica 52, 761–766.

Foster, J., Rothbaum, J., 2014. The mobility curve: Measuring the impact of mobility on welfare. Working Paper.

Foster, J., Shorrocks, A.F., 1988. Poverty orderings. Econometrica 56, 173–177.

Foster, J., Shorrocks, A.F., 1991. Subgroup consistent poverty indices. Econometrica 59, 687–709.

Foster, J.E., 2007. A report on Mexican multidimensional poverty measurement. OPHI Working Paper 40. URL: http://www.ophi.org.uk/a-report-on-mexican-multidimensional-poverty-measurement/.

Galasso, E., Ravallion, M., 2005. Decentralized targeting of an antipoverty program. Journal of Public Economics 89, 705–727.

Gambetta, D., 1988. Trust: Making and Breaking Cooperative Relations. Basil Blackwell, Oxford.

Gertler, P.J., 2004. Do conditional cash transfers improve child health? Evidence from PROGRESA's control randomized experiment. American Economic Review Papers and Proceedings 94, 336–341.

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

Ghatak, M., 2015. Theories of poverty traps and anti-poverty policies. World Bank Economic Review 29, S77–S105.

Goñi, E., López, J.H., Servén, L., 2011. Fiscal redistribution and income inequality in Latin America. World Development 39, 1558–1569.

Gobierno de la República, 2013. Plan nacional de desarrollo 2013–2018. Government report. URL: http://pnd.gob.mx/wp-content/uploads/2013/05/PND.pdf.

Gorman, W.M., 1968. The structure of utility functions. Review of Economic Studies 35, 367–390.

Grimm, M., 2007. Removing the anonymity axiom in assessing pro-poor growth. Journal of Economic Inequality 5, 179–197.

Guiso, L., Sapienza, P., Zingales, L., 2004. The role of social capital in financial development. American Economic Review 49, 526–556.

Guiso, L., Sapienza, P., Zingales, L., 2008. Trusting the stock market. Journal of Finance 63, 2557–2600.

Harrison, G.W., Rutherford, T.F., Tarr, D.G., 2003. Trade liberalization, poverty and efficient equity. Journal of Development Economics 71, 97–128.

Hassoun, N., Subramanian, S., 2012. An aspect of variable population poverty comparisons. Journal of Development Economics 98, 238–241.

Herrera, H.B., 2014. Pobreza multidimensional en Chile: Una nueva mirada. Ministerio de Desarrollo Social presentation.

Higgins, S., Lustig, N., Ruble, W., Smeeding, T.M., 2015. Comparing the incidence of taxes and social spending in Brazil and the United States. Review of Income and Wealth advance online publication. doi:doi:10.1111/roiw.12201.

Higgins, S., Pereira, C., 2014. The effects of Brazil's taxation and social spending on the distribution of household income. Public Finance Review 42, 346–367.

Hill, R., Inchauste, G., Lustig, N., Tsehaye, E., Woldehanna, T., forthcoming. A fiscal incidence analysis for Ethiopia, in: Inchauste, G., Lustig, N. (Eds.), The Distributional Impact of Fiscal Policy: Experience from Developing Countries. World Bank, Washington, D.C.

Hirschman, A.O., 1984. Against parsimony: Three easy ways of complicating some categories of economic discourse. American Economic Review Papers and Proceedings 74, 89–96.

Hoddinott, J., Skoufias, E., 2004. The impact of PROGRESA on food consumption. Economic Development and Cultural Change 53, 37–61.

Hoynes, H.W., Page, M.E., Stevens, A.H., 2006. Poverty in America: Trends and explanations. Journal of Economic Perspectives 20, 47–68.

Inchauste, G., Lustig, N., Maboshe, M., Purfield, C., Woolard, I., forthcoming. The distributional impact of fiscal policy in South Africa, in: Inchauste, G., Lustig, N. (Eds.), The Distributional Impact of Fiscal Policy: Experience from Developing Countries. World Bank, Washington, D.C.

International Monetary Fund, 2013. Fiscal Monitor: Taxing Times. International Monetary Fund, Washington, D.C. URL: `https://www.imf.org/external/pubs/ft/fm/2013/02/pdf/fm1302.pdf`.

Iyer, R., Puri, M., 2012. Understanding bank runs: The importance of depositor-bank relationships and networks. American Economic Review 102, 1414–1445.

Jack, W., Ray, A., Suri, T., 2013. Money management by households and firms in Kenya. American Economic Review 103, 356–361.

Jack, W., Suri, T., 2014. Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. American Economic Review 104, 183–223.

Jackson, M.O., Yariv, L., 2014. Present bias and collective dynamic choice in the lab. American Economic Review 104, 4184–4204.

de Janvry, A., Emerick, K., Gonzalez-Navarro, M., Sadoulet, E., 2015. Delinking land rights from land use: Certification and migration in Mexico. American Economic Review 105, 3125–3149.

Jaramillo, M., 2014. The incidence of social spending in Peru. Public Finance Review 42, 391–412.

Jaramillo, M., de la Flor, L., Sparrow, B., 2015. Are ethnic groupings invisible for fiscal policy in Peru? An incidence analysis of taxes and transfers on indigenous and non-indigenous Peruvians. Mimeo.

Kakwani, N.C., Pernia, E.M., 2000. What is pro-poor growth? Asian Development Review 18, 1–16.

Kakwani, N.C., Son, H.H., 2008. Poverty equivalent growth rate. Review of Income and Wealth 54, 643–655.

Kanbur, S.M.R., 1987. Measurement and alleviation of poverty: With an application to the effects of macroeconomic adjustment. International Monetary Fund Staff Papers 34, 60–85.

Karlan, D., Linden, L.L., 2014. Loose knots: Strong versus weak commitments to save for education in Uganda? mimeo.

Karlan, D., Mobius, M., Rosenblat, T., Szeidl, A., 2009. Trust and social collateral. Quarterly Journal of Economics 124, 1307–1361.

Karlan, D., Ratan, A.L., Zinman, J., 2014. Savings by and for the poor: A research review and agenda. Review of Income and Wealth 60, 36–78.

Karlan, D., Zinman, J., 2014. Price and control elasticities of demand for savings. Working Paper. URL: http://karlan.yale.edu/sites/default/files/savingselasticities_2014_01_v9.pdf.

Kast, F., Meier, S., Pomeranz, D., 2012. Under-savers anonymous: Evidence on self-help groups and peer pressure as a savings commitment device. Harvard Business School Working Paper 12-060. URL: http://www.hbs.edu/faculty/Publication%20Files/12-060_25eb4ab1-7f56-470b-b3a9-8e039b44118e.pdf.

Kast, F., Pomeranz, D., 2014. Saving more to borrow less: Experimental evidence from access to formal savings accounts in chile. Harvard Business School Working Paper 14-001. URL: http://www.hbs.edu/faculty/Publication%20Files/14-001_2237a8d2-6147-4667-9396-cd5ae9702f03.pdf.

Keen, M., Lockwood, B., 2010. The value added tax: Its causes and consequences. Journal of Development Economics 92, 138–151.

Kiviet, J.F., 1995. On bias, inconsistency, and efficiency of various estimators in dynamic panel data models. Journal of Econometrics 68, 53–78.

Kiviet, J.F., 1999. Expectations of expansions for estimators ina dynamic panel data model: Some results for weakly exogenous regressors, in: Hsiao, C., Lahiri, K., Lee, L.F., Pesaran, M.H. (Eds.), Analysis of panels and limited dependent variable models. Cambridge University Press, Cambridge, UK, pp. 199–225.

Klasen, S., 2008. Economic growth and poverty reduction: Measurement issues using income and non-income indicators. World Development 36, 420–445.

Kling, J.R., Liebman, J.B., Katz, L.F., 2007. Experimental analysis of neighborhood effects. Econometrica 75, 83–119.

Knack, S., Keefer, P., 1997. Does social capital have an economic payoff? a cross-country investigation. Quarterly Journal of Economics 112, 1251–1288.

Kolm, S.C., 1976. Unequal inequalities I. Journal of Economic Theory 12, 416–442.

Kraay, A., 2006. When is growth pro-poor? Evidence from a panel of countries. Journal of Development Economics 80, 198–227.

La Porta, R., de Silanes, F.L., Zamarripa, G., 2003. Related lending. Quarterly Journal of Economics 118, 231–268.

Laibson, D., 2015. Why don't present-biased agents make commitments. Quarterly Journal of Economics 105, 267–272.

Lambert, P.J., 1988. Net fiscal incidence progressivity: Some approaches to measurement, in: Eichhorn, W. (Ed.), Measurement in Economics: Theory and Application of Economics Indices. Springer-Verlag, Heidelberg, pp. 519–532.

Levy, S., Schady, N., 2013. Latin America's social policy challenge: Education, social insurance, redistribution. Journal of Economic Perspectives 27, 193–218.

Llerena Pinto, F.P., Llerena Pinto, M.C., Llerena Pinto, M.A., Perez, G., 2015. Social spending, taxes and income redistribution in Ecuador. CEQ Working Paper 28.

Lopez-Calva, L.F., Lustig, N., Matytsin, M., Popova, D., forthcoming. Who benefits from fiscal redistribution in Russia?, in: Inchauste, G., Lustig, N. (Eds.), The Distributional Impact of Fiscal Policy: Experience from Developing Countries. World Bank, Washington, D.C.

López-Calva, L.F., Ortiz-Juárez, E., 2009. Medición multidimensional de la pobreza en México: Significancia estadística en la inclusión de dimensiones no monetarias. Estudios Económicos número extraordinario, 3–33.

Lustig, N., 2015. The redistributive impact of government spending on education and health: Evidence from 13 developing countries in the Commitment to Equity project, in: Gupta, S., Keen, M., Clements, B., de Mooij, R. (Eds.), Inequality and Fiscal Policy. International Monetary Fund, Washington, D.C.

Maasoumi, E., Racine, J.S., forthcoming. A solution to aggregation and an application to multidimensional well-being frontiers. Journal of Econometrics .

Maasoumi, E., Xu, T., 2015. Weights and substitution deree in multidimensional well-being in China. Journal of Economic Studies 42, 4–19.

Makdissi, P., Wodon, Q., 2002. Consumption dominance curves: Testing for the impact of indirect tax reforms on poverty. Economics Letters 75, 227–235.

Martínez-Aguilar, S., Ortiz-Juarez, E., 2015. CEQ Master Workbook for Chile. Mimeo.

McMillan, J., Woodruff, C., 1999. Interfirm relationships and informal credit in Vietnam. Quarterly Journal of Economics 114, 1285–1320.

Meyer, B., Mok, W.K.C., Sullivan, J.X., 2015. Household surveys in crisis. Journal of Economic Perspectives 29, 199–226.

Mirrlees, J., Adam, S., Besley, T., Blundell, R., Bond, S., Chote, R., Gammie, M., Johnson, P., Myles, G., Poterba, J., 2011. Tax by Design. Oxford University Press, Oxford.

Mitra, T., Ok, E., 1998. The measurement of income mobility: A partial ordering approach. Economic Theory 12, 77–102.

Mullainathan, S., Shafir, E., 2009. Savings policy and decision-making in low-income households, in: Insufficient Funds: Savings, Assets, Credit and Banking Among Low-Income Households. Russell Sage Foundation Press, New York City, pp. 121–145.

Narayan, D., Pritchett, L., 1999. Cents and sociability: Household income and social capital in rural Tanzania. Economic Development and Cultural Change , 871–898.

Nicita, A., 2009. The price effect of tariff liberalization: Measuring the impact on household welfare. Journal of Development Economics 89, 19–27.

Nickell, S., 1981. Biases in dynamic models with fixed effects. Econometrica 49, 1417–1426.

Niehaus, P., Sukhtankar, S., 2013. The marginal rate of corruption in public programs: Evidence from india. Journal of Public Economics 104, 52–64.

Olken, B.A., 2012. Corruption and the costs of redistribution: Micro evidence from indonesia. Journal of Public Economics 90, 853–870.

Parker, S.W., Teruel, G.M., 2005. Randomization and social program evaluation: The case of progresa. Annals of the American Acadamy of Political and Social Science 599, 199–219.

Paz Arauco, V., Gray Molina, G., Yáñez Aguilar, E., Jiménez Pozo, W., 2014. Explaining low redistributive impact in Bolivia. Public Finance Review 42, 326–345.

Permanyer, I., 2014. Assessing individuals' deprivation in a multidimensional framework. Journal of Development Economics 109, 1–16.

Plotnick, R., 1982. The concept and measurement of horizontal inequity. Journal of Public Economics 17, 373–391.

Porta, R.L., Lopez-De-Silanes, F., Shleifer, A., Vishny, R.W., 1997. Trust in large organizations. American Economic Association 87, 333–338.

Porto, G.G., 2006. Using survey data to assess the distributional effects of trade policy. Journal of International Economics 70, 140–160.

Prina, S., 2015. Banking the poor via savings accounts: Evidence from a field experiment. Journal of Development Economics 115, 16–31.

Pritchett, L., Summers, L.H., 1996. Wealthier is healthier. Journal of Human Resources 31, 841–868.

Ravallion, M., 1996. Issues in measuring and modeling poverty. Economic Journal 106, 1328–1343.

Ravallion, M., 1999. Is more targeting consistent with less spending? International Tax and Public Finance 6, 411–419.

Ravallion, M., 2001. Growth, inequality and poverty: Looking beyond averages. World Development 29, 1803–1815.

Ravallion, M., 2011. On multidimensional indices of poverty. Journal of Economic Inequality 9, 235–248.

Ravallion, M., 2012. Why don't we see poverty convergence? American Economic Review 102, 504–523.

Ravallion, M., 2015. The idea of antipoverty policy, in: Handbook of Income Distribution. North-Holland, Amsterdam. volume 2, pp. 1967–2061.

Ravallion, M., Chen, S., 2003. Measuring pro-poor growth. Economics Letters 78, 93–99.

Reynolds, M., Smolensky, E., 1977. Public Expenditures, Taxes and the Distribution of Income: The United States, 1950, 1961, 1970. Academic Press, New York.

Roodman, D., 2009a. How to do xtabond2: An introduction to difference and system GMM in Stata. Stata Journal 9, 86–136.

Roodman, D., 2009b. A note on the theme of too many instruments. Oxford Bulletin on Economics and Statistics 71, 135–158.

Rubalcava, L., Teruel, G., Thomas, D., 2009. Investments, time preferences, and public transfers paid to women. Economic Development and cultural change 57, 507.

Sapienza, P., Zingales, L., 2012. A trust crisis. International Review of Finance 12, 123–131.

Schaner, S., 2015a. The cost of convenience? Transaction costs, bargaining, and savings account use in kenya. Working Paper. URL: http://www.dartmouth.edu/~sschaner/main_files/Schaner_ATM.pdf.

Schaner, S., 2015b. Do opposites detract? Intrahousehold preference heterogeneity and inefficient strategic savings. American Economic Journal: Applied Economics 7, 135–174.

Scott, J., 2014. Redistributive impact and efficiency of Mexico's fiscal system. Public Finance Review 42, 368–390.

Secretaría de Hacienda y Crédito Público, 2010. Presupuesto de egresos de la Federación para el ejercicio fiscal 2010. URL: http://www.apartados.hacienda.gob.mx/presupuesto/temas/pef/2010/.

Sen, A., 1976. Poverty: An ordinal approach to measurement. Econometrica 44, 219–231.

Sen, A., 1981. Poverty and Famines: An Essay on Entitlement and Deprivation. Oxford University Press, Oxford.

Sen, A., 1987. The Standard of Living. Cambridge University Press, Cambridge.

Sen, A., 1999. Development as Freedom. Alfred A. Knopf, New York.

Shimeles, A., Moummi, A., Jouini, N., Lustig, N., forthcoming. Fiscal incidence and poverty reduction: Evidence from Tunisia. CEQ Working Paper 38.

Silber, J., 2011. A comment on the MPI index. Journal of Economic Inequality 9, 479–481.

Subramanian, S., 2006. Social groups and economic poverty: A problem in measurement, in: McGillivray, M. (Ed.), Inequality, Poverty, and Wellbeing. Palgrave Macmillan, London, pp. 143–161.

Suri, T., Jack, W., Stoker, T.M., 2012. Documenting the birth of a financial economy. Proceedings of the National Academy of Sciences 109, 10257–10262.

Tobin, J., 1956. The interest-elasticity of transactions demand for cash. Review of Economics and Statistics 38, 241–247.

Tran, V.Q., Alkire, S., Klasen, S., 2015. Static and Dynamic Disparities between Monetary and Multidimensional Poverty Measurement: Evidence from Vietnam. chapter 15. pp. 249–281.

Tsui, K., 2002. Multidimensional poverty indices. Social Choice and Welfare 19, 69–93.

UNDP, 2010. The Real Wealth of Nations: Pathways to Human Development. Macmillan, New York.

United Nations, 2015. Addis Ababa Action Agenda of the Third International Conference on Financing for Development. United Nations, New York.

Van Kerm, P., 2009. Income mobility profiles. Economics Letters 102, 93–95.

Williamson, O.E., 1993. Calculativeness, trust, and economic organization. Journal of Law and Economics 36, 453–486.

Windmeijer, F., 2005. A finite sample correction for the variance of linear efficient two-step GMM estimators. Journal of Econometrics 126, 25–51.

World Bank, 2013. Financing for Development Post-2015. World Bank, Washington, D.C.

World Bank, 2015. Ethiopia Poverty Assessment 2014. World Bank, Washington, D.C.

Wright, R., Tekin, E., Topalli, V., McClellan, C., Dickinson, T., Rosenfeld, R., 2014. Less cash, less crime: Evidence from the Electronic Benefit Transfer program. NBER Working Paper 19996. URL: `http://www.nber.org/papers/w19996`.

Younger, S.D., Khachatryan, A., forthcoming. Fiscal incidence in Armenia, in: Inchauste, G., Lustig, N. (Eds.), The Distributional Impact of Fiscal Policy: Experience from Developing Countries. World Bank, Washington, D.C.

Younger, S.D., Osei-Assibey, E., Oppong, F., 2015. Fiscal incidence in Ghana. CEQ Working Paper 35.

Zheng, B., 1994. Can a poverty index be both relative and absolute? Econometrica 62, 1453–1458.

Zheng, B., 1997. Aggregate poverty measures. Journal of Economic Surveys 11, 123–162.

Ziliak, J.P., 1997. Efficient estimation with panel data when instruments are predetermined: An empirical comparison of moment-condition estimators. Journal of Business and Economic Statistics 15, 419–431.

Zinman, J., 2009. Credit or debit? Journal of Banking and Finance 33.

Biography

Sean Higgins completed his B.S. *summa cum laude* in Economics from Tulane University in 2011, and was inducted into Phi Beta Kappa as a junior. He continued at Tulane for his Ph.D., where his research has focused on the impact of public policy on poverty. His research has been published in the *Journal of Development Economics*, *Public Finance Review*, and the *Review of Income and Wealth*, and has been funded by the Bill & Melinda Gates Foundation, Fulbright–García Robles Public Policy Initiative, National Science Foundation, and Tulane School of Liberal Arts.