



Essays on Welfare and Taxation

The Harvard community has made this article openly available. [Please share](#) how this access benefits you. Your story matters

Citation	Lagomarsino, Alejandro Luis. 2020. Essays on Welfare and Taxation. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.
Citable link	https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37365923
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

Essays on Welfare and Taxation

A dissertation presented

by

Alejandro Luis Lagomarsino

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

April 2020

© 2020 Alejandro Luis Lagomarsino
All rights reserved.

Dissertation Advisor:
Professor Lawrence Katz

Author:
Alejandro Luis Lagomarsino

Essays on Welfare and Taxation

Abstract

The three essays of my dissertation focus on understanding the labor market and fiscal impacts of key government policies. The first essay, co-authored with Lihuen Nocetto, studies whether a cash assistance program in Uruguay creates a welfare trap. We provide new evidence on the possibility of such a trap exploiting the unique way in which the Uruguayan government re-targeted its main unconditional cash transfer program. With rich administrative longitudinal data and survey data, we provide separate causal estimates of the impact of entry into and of forced exit from the program based on regression discontinuity and dynamic differences-in-differences designs. We focus on three key outcomes: labor supply and formalization of work, human capital investments for children, and take-up of other safety-net programs. Overall, our results suggest that the program does not induce a welfare trap, and that long-term enrollment seems to be associated with strategic behaviors to sustain eligibility status, and not with decreased ability to leave social assistance.

In the second essay, co-authored with Rafael Di Tella and Juan Dubra, we study the causal impact of trust in business elites and trust in government, on preferences for taxation at the top. Using a randomized online survey, and in contrast to previous work, we find that distrust causes an increase in desired taxes on the top 1%.

The third essay focuses on estimating elasticities of taxable sales, profits, and assets with respect to their net-of-tax rates in Peru using bunching evidence at kinks and notches. We find that taxable sales elasticities are sizable at the lowest notch but monotonically decrease as we look at notches associated with higher levels of sales. We find elasticities of taxable profits in line with those found in other countries, and large elasticities of taxable assets.

The latter suggests that the desirability of wealth taxes in developing contexts should be carefully studied.

Contents

Abstract	iii
Acknowledgments	x
Introduction	1
1 Do Cash Assistance Programs Create Welfare Traps?	3
1.1 Introduction	5
1.2 Institutional context	15
1.2.1 The Uruguayan welfare state	16
1.2.2 A primer on Uruguay’s largest unconditional cash transfer program .	19
1.2.3 Eligibility process and fieldwork	22
1.3 Data and summary statistics	27
1.3.1 Data from the Ministry of Social Development	27
1.3.2 Formal labor supply data	30
1.3.3 Education data	31
1.3.4 Data on enrollment in other safety-net programs	32
1.4 Econometric framework	33
1.4.1 Fuzzy regression discontinuity design	33
1.4.2 Dynamic differences-in-differences design	42
1.5 Main results	43
1.5.1 Formal labor supply	44
1.5.2 Education	48
1.5.3 Take-up of safety-net programs	51
1.6 Discussion	55
1.6.1 Material hardship	56
1.6.2 Labor supply and human capital investment responses with survey data	60
1.6.3 Strategic behaviors: complaints, selective deafness, and misreporting	61
1.7 Robustness of the main results	67
1.8 Conclusions	70

2	Meet the Oligarchs: Business Legitimacy and Taxation at the Top	74
2.1	Introduction	76
2.2	Empirical strategy and data	81
2.2.1	Empirical strategy	81
2.2.2	Survey implementation	87
2.3	Results	93
2.4	Supplementary survey	98
2.4.1	First stage: the effect of the treatments on the dimensions of trust	99
2.4.2	Actions vs behaviors, mechanisms	101
2.5	Conclusions	104
3	Uncovering Elasticities with Notches and Kinks: Evidence from Peru	106
3.1	Introduction	108
3.2	A primer on the Peruvian tax system	110
3.2.1	Tax regimes for businesses before 2017	111
3.2.2	The 2017 tax reform	113
3.3	Data and descriptive statistics	116
3.4	Methodology to compute elasticities	118
3.4.1	Pure notches	118
3.4.2	Kinks	120
3.5	Results	122
3.5.1	Taxable sales	122
3.5.2	Taxable profits	126
3.5.3	Taxable assets	127
3.6	Conclusions	128
	References	130
	Appendix A Appendix to Chapter 2	140
A.1	Additional results	140
A.2	Main survey questionnaire	145
A.3	Supplementary survey questionnaire	156

List of Tables

1.1	Characteristics of visited households and individuals at the time of the visit	29
1.2	Balance on observables	38
1.3	Impact of UCT on formal labor supply: fuzzy RD estimates	46
1.4	Impact of UCT on youth enrollment in public schools: fuzzy RD estimates .	50
1.5	Impact of UCT on take-up of housing assistance: fuzzy RD estimates	53
1.6	Impact of UCT on take-up of cash-assistance for formal workers: fuzzy RD estimates	55
1.7	Impact of UCT on re-visit requests: fuzzy RD estimates	63
2.1	Summary statistics	89
2.2	Randomization	91
2.3	Completion of survey by treatment arm	93
2.4	Preferences for taxation	95
2.5	Voting for taxes on the top 1%	102
3.1	Distribution and revenue from tax regimes before and after the 2017 reform	115
3.2	Number of businesses in SUNAT's database	117
3.4	Structural taxable sales elasticities	125
A.1	Taxes: Democrat vs Republican	141
A.2	First stage	143

List of Figures

1.1	Welfare Stigma	10
1.2	Welfare Morale	15
1.3	Steps followed by MIDES to grant or withdraw UCT	24
1.4	Density of household visits by Vulnerability Index (Montevideo)	28
1.5	McCrary test	40
1.6	First stage for pre-visit non-beneficiaries	41
1.7	Impact of UCT on formal labor supply	45
1.8	Dynamic DID estimates: formal employment	47
1.9	Impact of UCT on youth enrollment in public schools	49
1.10	Impact of UCT on completed years of schooling	51
1.11	Impact of UCT on take-up of housing assistance	52
1.12	Impact of UCT on take-up of cash assistance for formal workers	54
1.13	Impact of UCT on food insecurity	57
1.14	Impact of UCT on durable goods consumption	58
1.15	Impact of UCT on housing conditions	59
1.16	Impact of UCT on self-reported employment status and school attendance	61
1.17	Impact of UCT on re-visit requests	62
1.18	Impact of UCT on probability of area re-visit	64
1.19	Impact of UCT on visit attempts before a targeted visit is conducted	66
1.20	Robustness to bandwidth and kernel specification: formal labor supply	68
1.21	Falsification tests using placebo thresholds	70
2.1	Beliefs that taxes on the rich are too low and trust in business and government	79
2.2	Survey design	82
2.3	Preferred tax rate for the top 1%	94
2.4	First stage	100
3.1	NRUS Tax Schedule	112
3.2	Effective tax rate for businesses before and after the 2017 reform	114
3.3	Taxable sales: empirical distributions (NRUS businesses, 2010-16)	123

3.4	Taxable sales: empirical and counter-factual distributions around notches (NRUS businesses, 2010-16)	124
3.5	Taxable profits: empirical distributions (RMT and GR businesses)	127
3.6	Taxable assets: empirical distributions around the kink (RG/RMT businesses, 2010-17)	128
A.1	Beliefs that taxes on the rich are too low and trust in business and government (executive branch of the federal government)	145

Acknowledgments

I am indebted to my advisors-Alberto Alesina, Lawrence Katz, and Stefanie Stantcheva-for their thoughtful advice and continuous support.

I also benefited greatly from the advice of several faculty and non-faculty members (at Harvard and elsewhere), which even if they were not part of my committee, had a tremendous impact on my work. Special thanks within this group to Rafael Di Tella and Juan Dubra for their continuous guidance and support. I would also like to thank Pablo Balán, Gregory Bruich, Raj Chetty, Paola Giuliano, Enrico Di Gregorio, Benjamin Enke, Nathaniel Hendren, Horacio Larreguy, Gabriel Kreindler, Michael Kremer, Nathan Nunn, Gautam Rao, Tatiana Rosá, Michael Thaler, and Clémence Tricaud. Countless other graduate students and faculty members provided interesting conversations and advice, especially during Labor and Public Finance Lunch, Political Economy Lunch, and Development Lunch at Harvard.

For the first chapter of this dissertation, I use de-identified data provided by the Ministry of Social Development in Uruguay. I am very grateful to Luis Lagaxio for his support in making this project possible, as well as to María del Carmen Correa, Lorena Custodio, Guillermo D'Angelo, Richard Detomassi, Martín Moreno, and María Eugenia Oholeguy for their support with the data. Andrés Cruz Labrín provided excellent research assistance for this chapter. For the third chapter, I use de-identified data provided by the tax administration in Peru (SUNAT). I am grateful to Eduardo Mercado and Estuardo Oliver for their support with the data and to Dennis Sanchez for her helpful comments. Luis Villazón provided excellent research assistance.

I acknowledge financial support from the David Rockefeller Center for Latin American Studies, the Institute of Quantitative Social Sciences at Harvard University, the Inter-American Development Bank, and the Molly and Domenic Ferrante Research Fund.

Finally, I also want to express my gratitude for the people that got me through graduate school. The love, continuous support, and encouragement of my family have been my most precious assets during grad school. To my fellow economists and friends – Lisa Abraham,

Edoardo Acabbi, Omar Barbiero, Moya Chin, Paulo Costa, Talia Gillis, Nir Hak, Tzachi Raz, Jonathan Roth, and Natasha Sarin. I truly appreciate having had the opportunity to share this experience with each of you. The same applies to the wonderful group of friends I made from other Harvard departments, especially during the English Language Program – Silvia Canas, Solsiré Cusicanqui, Lukas Maas, Jonathan Romero, and Abdul Wasay.

To my family,

Introduction

My dissertation focuses on understanding the labor market and fiscal impacts of key government policies, with chapters both on cash assistance programs and taxation. Within each of these areas, I have used large-scale administrative datasets and data collected from randomized online survey experiments to answer economically relevant questions. Below, I discuss each specific chapter.

Concerning my first chapter, an extensive body of empirical literature sheds light on whether entering a cash assistance program discourages work and/or induces other behaviors that could foster dependency. However, much less is known about what happens when beneficiaries are forced out of welfare, and both margins are key to assess if cash assistance programs create welfare traps. In this first chapter, I provide the first (as far as I know) causal estimates of the impact of both entry into and forced exit from a cash assistance program to study this topic. I assemble a unique dataset constructed from five different government sources and focus on three key outcomes historically associated with welfare traps: labor supply and formalization of work, human capital investments for children, and take-up of other safety-net programs. I find that formal labor supply of adults drops after entering the program, as critics of cash assistance programs regularly argue. Nevertheless, the paper shows that this does not constitute a welfare trap: long-term recipients that are forced to exit the program increase their formal labor supply.

I find particularly interesting results on education. Youth enrollment in schools drops when households start receiving the cash assistance. This is surprising given that most of the literature either finds null or positive impacts of cash transfers on school enrollment. This is

similar though to what Dahl and Gielen (2018) find in recent work in the Netherlands.

With respect to take-up of other safety-net programs, I find that the program has a negative impact on enrollment in public housing programs and take-up of other types of public cash assistance. This is more suggestive of safety-net program substitution than of increased dependency on multiple programs.

Overall, these results suggest that the program does not induce a welfare trap. While it is true that beneficiaries reduce their formal labor supply on entry, the fact that beneficiaries that are forced to exit welfare increase their formal labor supply suggests that the program does not decrease beneficiaries' ability to find employment. Although we find negative impacts of the program on education enrollment, this does not create a welfare trap, as these impacts do not translate into different educational attainment rates. Take-up of other welfare programs can also be ruled out as a potential driver of dependency, as we find safety-net programs to be substitutes in this context

In my second chapter, we study the impact of two dimensions of trust, namely trust in business elites and trust in government, on preferences for taxation. Using a randomized online survey experiment, we find that distrust causes an increase in desired taxes on the top 1%. This is the opposite to what basically any previous theory would have predicted about this result (especially any theory where voters only care about income as in Meltzer and Richard (1981) seminal paper). We provide a plausible (and up to our knowledge, the first) positive theory of taxation at the top that is orthogonal to preferences for redistribution to interpret our findings.

In my third chapter, we use bunching evidence in Peru to estimate elasticities of taxable sales, profits, and assets with respect to their net-of-tax rates. Peru serves as an interesting laboratory to estimate these elasticities given the number of notches and kinks generated by the Peruvian tax system. We find that taxable sales elasticities are sizable at the lowest notch but monotonically decrease as we look at notches associated with higher levels of sales. We find elasticities of taxable profits in line with those found in other countries, and large elasticities of taxable assets.

Chapter 1

Do Cash Assistance Programs Create Welfare Traps?¹

¹Co-authored with Lihuen Nocetto.

Essay Abstract

A major concern in the design of safety-net programs is the possibility that long-term participation in cash assistance programs inhibits self-sufficiency. We provide new evidence on the possibility of such welfare traps exploiting the unique way in which the Uruguayan government re-targeted its main unconditional cash transfer program. With rich administrative longitudinal data and survey data, we provide separate causal estimates of the impact of entry into and of forced exit from the program based on regression discontinuity and dynamic differences-in-differences designs. We focus on three key outcomes: labor supply and formalization of work, human capital investments for children, and take-up of other safety-net programs. We find that formal labor supply of adults drops after they enter the program. However, this does not constitute a welfare trap: long-term recipients that are forced to exit the program increase their formal labor supply. We find negative impacts on school enrollment but not on completed years of schooling, given the low graduation rates in our population. The program negatively impacts enrollment in public housing programs and take-up of other types of public cash assistance, which is more suggestive of safety-net program substitution than of increased dependency on multiple programs. Overall, these results suggest that the program does not induce a welfare trap, and that long-term enrollment seems to be associated with strategic behaviors to sustain eligibility status, and not with decreased ability to leave social assistance.

1.1 Introduction

A major concern in the design of safety-net programs is the possibility that long-term participation in cash (and near-cash) assistance programs fosters dependency.² An extensive body of empirical literature sheds light on whether entering a cash assistance program discourages work and/or induces other behaviors that could inhibit self-sufficiency. However, much less is known about what happens when beneficiaries are *forced* out of welfare, and both margins are key to answering the question of whether cash assistance programs create welfare traps.³ This study offers an empirical answer to this question by providing separate causal estimates of the impact of entry into and of forced exit from a cash assistance program.

To isolate these causal impacts, we take advantage of the unique way in which Uruguay decided to re-target its main unconditional cash transfer (UCT) program through a proxy-means test for granting and withdrawing the benefit. Through more than 250,000 household visits (during 2012 – 2018) covering roughly one-fifth of the population, the government estimated a socio-economic vulnerability score for each household and determined its eligibility for the program by comparing the score to a pre-specified threshold (i.e., beneficiary households are those whose scores are higher than the pre-specified threshold). This meant that some households that were enrolled in the program prior to the visit lost their benefit, while some households that were not enrolled in the program prior to the visit started receiving the transfer after the visit. The fact that those slightly above the threshold

²Ronald Reagan’s 1986 State of the Union illustrates that this has historically been a concern in the US: “As Franklin Roosevelt warned 51 years ago, standing before this Chamber, he said: ‘Welfare is a narcotic, a subtle destroyer of the human spirit’. And we must now escape the spider’s web of dependency”. The concern is still present nowadays, as reflected in the following passage from the 2020 Budget of the United States: “The Budget proposes commonsense work requirements for the Supplemental Nutrition Assistance Program (SNAP) that would require all able-bodied adult participants to find or train for employment and work toward self-sufficiency”.

³There are several definitions of welfare trap. Cooke (2009) defines it as a “decreasing ability to leave social assistance”. Guzi *et al.* (2014) states that a welfare trap is created when “social benefits are accepted as an alternative to low and insecure earnings”. Shaoan *et al.* (2019) takes a completely different approach and defines it as a “situation where the government sacrifices long-term economic growth and welfare to maintain short-term welfare”. In this paper, “welfare trap” is meant to represent a situation where the receipt of welfare over time inhibits beneficiaries’ ability to be self-sufficient.

become eligible for the program but those slightly below the threshold do not allow us to estimate the impact of the program based on both regression discontinuity and dynamic differences-in-differences designs.

We use rich administrative longitudinal data (up to five years of monthly data for almost a million individuals), and survey data matched from five different government institutions, to follow individuals after they enter/exit the program. We focus on three key outcomes to assess family welfare dependency.

First, we focus on labor supply and formalization of work. Theoretically, the concern is that welfare could discourage work either through an income effect or (for “visible” forms of work) a perception on the part of individuals that their status in the program could be jeopardized if they work or if their earnings surpass a certain threshold. However, the effect could also go in the opposite direction if individuals need the transfer to reach a minimum level of consumption that allows them to be able-bodied workers (the seminal work by Dasgupta and Ray 1986 models the link between malnutrition and unemployment), or if the transfer helps relieve the credit constraint and allows households to invest in small businesses.⁴

Second, we look at human capital investments for children. Theoretically, a cash assistance program could reduce human capital investments under certain scenarios.⁵ Dahl and Gielen (2018) finds that children of parents whose disability insurance benefit was reduced complete more schooling, which the authors interpret as being due to the expectation that they will receive less of a disability insurance benefit when they reach adulthood.⁶ If children invest (at least partially) in human capital to increase the likelihood of getting a formal job as adults, any form of “tax” on formal employment could reduce human capital investments during childhood. Thus one possibility is that children entering

⁴See Baird *et al.* (2018) for a discussion on the channels through which cash transfers could impact adult labor market outcomes.

⁵See Kesselman (1976) for an early theoretical contribution on how welfare could reduce the incentives for human capital accumulation.

⁶“We find intriguing evidence for anticipatory educational investments, consistent with children planning for a future with less reliance on DI” (Dahl and Gielen 2018, p. 3).

the UCT program raise their expectations of receiving the transfer during adulthood as long as they do not participate in the formal labor market (or their formal income is low enough), which could discourage educational investments. Of course, it is also possible that cash assistance programs lead to *more* human capital investments in children, and in fact most theoretical work on this topic has emphasized this possibility. Education is commonly thought of as a normal good (Lazear 1975), and a transfer that relaxes the household's budget constraint could increase enrollment of children in schools for multiple reasons (e.g., because children no longer need to work to support the household, because they are healthier and thus more able to stay in school, or because the transfer enables the purchase of necessary school supplies).

Third, we study whether welfare begets welfare. In particular, we study whether being enrolled in the UCT program impacts take-up of other safety-net programs.⁷ The literature has identified three types of costs that impact take-up of welfare: the stigma associated with it, transaction costs of applying to the program, and costs associated with learning about the program's eligibility and application rules.⁸ If enrollment in UCT lowers any of these costs for other welfare programs, then we could see an increase in take-up of other programs. However, we could see a decrease in take-up if enrolling in the UCT program reduces the marginal benefit of enrolling in other welfare programs. The simple argument would be that UCT induces a positive income effect that makes individuals less willing to bear the costs associated with applying to other welfare programs. Also, decreased take-up could occur if enrolling in UCT mechanically decreases eligibility for other welfare programs.⁹

In addition to these three main outcomes, we look at other dimensions that help us interpret our main results and understand some of their mechanisms. As a proxy for how household expenditures change with the program, we look at how receipt of UCT

⁷We do not study whether there is intergenerational transmission of welfare (i.e., whether children whose parents are enrolled in UCT are more/less likely to enter UCT during adulthood), as our data do not span sufficiently many years to enable us to follow children during adulthood.

⁸See Moffitt (1983) for the seminal theoretical contribution on welfare stigma.

⁹Caruso *et al.* (2019) mentions that it is common for policymakers to define eligibility for multiple programs in such a way that poor families receive at least some coverage.

impacts self-reported food insecurity measures, durable goods consumption, and housing conditions. We also look at strategic behaviors in which people engage to try to enter/stay in the program and that do not alter self-sufficiency. In particular, we look at whether people make a formal request for a re-visit if they are forced to exit the program, whether they are “selectively deaf” when a government official tries to visit them (i.e., whether current participation in the program affects an individual’s willingness to answer the door when a government official shows up), and whether there is misreporting of information gathered at the visit. Documenting these behaviors enables us not only to show that households do want to enroll in this program (despite the welfare stigma) but also helps us to understand why there is “stickiness” in program participation, which is a feature that is seen in this program and that has also been documented in several other contexts.

We find that the formal labor supply of adults (especially those under 40 years of age) drops three years after entering the program. However, beneficiaries that are forced to exit the program increase their formal labor supply three years later. School enrollment (for youth that should be enrolled in secondary school) decreases when entering welfare. Nevertheless, because secondary school graduation rates in the Uruguayan context are low for low-income families, the school enrollment impacts do not translate into differential educational attainment rates. Entry into the program has a negative impact on enrollment in public housing programs (although we note that this effect is not entirely robust) and take-up of other types of public cash assistance. These are more suggestive of safety-net program substitution than of increased dependency on multiple programs. Finally, we document several behaviors that induce stickiness in program participation and that are orthogonal to an individual’s ability to subsist without the cash assistance.

Overall, these results suggest that the program does not induce a welfare trap. If we focus only on entry, we could be tempted to conclude that the program fosters dependency through a reduction in formal labor supply. However, the fact that people that are forced to exit welfare increase their formal labor supply suggests that even if the program decreases employment, it does not decrease beneficiaries’ *ability* to find employment. Although we

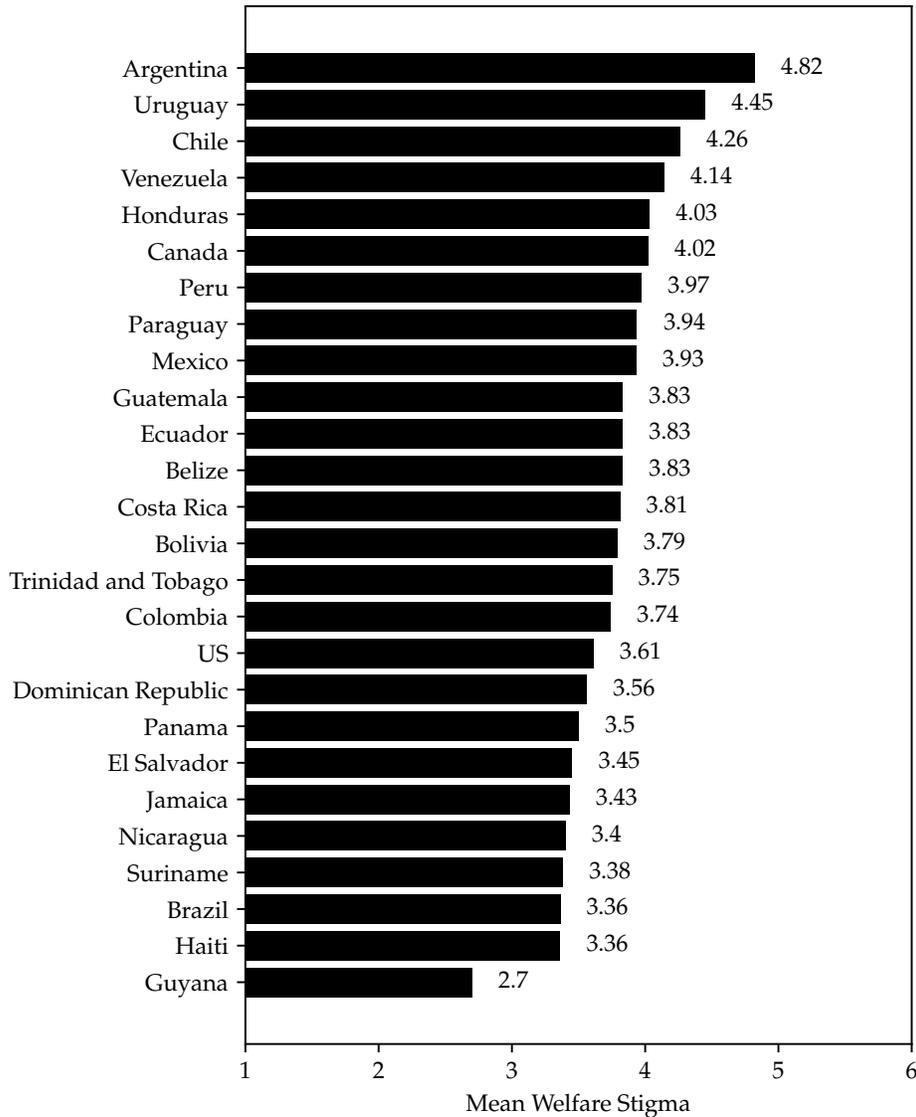
find negative impacts of the program on school enrollment, this does not create a welfare trap, as these impacts do not translate into differential educational attainment rates. Take-up of other welfare programs can also be ruled out as a potential driver of dependency in this context, as we find safety-net programs to be substitutes. Finally, we note that stickiness in program participation (which could be confused with a welfare trap) seems to be mostly associated with strategic behaviors aimed at staying in the program, rather than with a decreased ability to leave social assistance.

One could argue that, perhaps in other contexts, cash assistance programs do create welfare traps. The interesting aspect of finding that our program does not create such a trap is that, if anything, Uruguay is one of the settings where one would most have expected to find evidence of such traps. First, the program we study seems not to be lifting people out of poverty, and households remain in the program for several years.¹⁰ Households enrolled in UCT during September 2018 (the latest date for which data are available) have been receiving the transfer for 5.2 years on average. This could be indicative that the program is inhibiting self-sufficiency.¹¹ Second, it is a setting where beliefs that welfare discourages work are particularly salient, which is probably the main channel associated with the possibility that cash assistance programs foster dependency and create welfare traps. Figure 1.1 shows cross-country beliefs on whether welfare recipients in Latin America, the US, and Canada are lazy. We see that in the sample of 26 economies, Uruguay is surpassed only by Argentina on this measure.

¹⁰The government re-tested a group of households one year (on average) after they were deemed eligible for the program based on to their poverty status; 87% were still not able to exit poverty and remained eligible.

¹¹It could also be the case that individuals are strategically reporting in a way that induces stickiness in program participation but is not necessarily associated with reduced self-sufficiency.

Figure 1.1: Welfare Stigma



Notes: Each bar represents the mean answer by country to the following question: “Some people say that people who get help from government social assistance programs are lazy. How much do you agree or disagree?”. Possible answers are on a scale from 1 (Strongly disagree) to 7 (Strongly agree). Data source is the Americas Barometer (2012).

This paper adds to a large literature on a central topic in labor economics and public finance: the relationship between welfare and the labor market (see Moffitt 2002 or Chan and Moffitt 2018 for a review). Its main contribution is to study this topic with causal estimates of both entry into and forced exit from a cash assistance program. To see why

both margins are key to the study of welfare traps, we note how studying only one margin may not provide sufficient evidence to reach a conclusion. First, suppose that our estimates indicate a reduction in labor supply (or labor income) while an individual is enrolled in the program. This would not be sufficient to conclude that there is a welfare trap. It could be the case that individuals are reducing their earned income while enrolled in the program but their *ability* to generate income remains unchanged, and that earnings would go back up if they exited the program. Second, suppose we find no positive impacts on labor supply when individuals (quasi-randomly) exit the program. Again, it is not possible to reach a conclusion, as we could be dealing with a program that has no negative impacts on labor supply at entry.

To the best of our knowledge, this is the first study to provide causal estimates of both entry into and exit from a cash assistance program. Previous research has mostly relied on either causal estimates of entry into a program on labor supply (or formalization of work), or descriptive studies of characteristics associated with endogenous exit from cash assistance programs (e.g., Hansen 2007; Blank 1989).¹²

Among the first group of studies, there is previous work in the US on the labor market impacts of Food Stamps (Currie 2003 and Hoynes and Schanzenbach 2016 review this literature), currently called the Supplemental Nutrition Assistance Program (SNAP).¹³ While initial work on this topic (Hangstrom 1996; Fraker and Moffitt 1988) shows practically no evidence of work disincentives, recent work has found some evidence, in particular among single women (East 2018; Williamson and Whitmore 2012).¹⁴ Within developing

¹²There are a few papers that report quasi-random variation in exit from welfare (or a reduction in the amount of welfare received), although with quite different programs and contexts than ours (Riddell and Riddell 2014; Dahl and Gielen 2018; Deshpande 2016). There is also work on how recipients of unemployment insurance in the US respond at and before benefit exhaustion (Ganong and Noel 2019; Katz and Meyer 1990).

¹³There is also work on the Temporary Assistance for Needy Families program (TANF), as well as on a permanent and universal cash transfer in Alaska (Jones and Marinescu 2019) and on other major welfare programs in the US (e.g., Baicker *et al.* 2014 on Medicaid, and Chetty *et al.* 2013 and Miller *et al.* 2018 on the Earned Income Tax Credit). In this short review, we highlight SNAP, as this is the program (within the US) most closely related to the UCT program in Uruguay. For a thorough review on means tested transfer programs in the US, see Moffitt (2016a,b).

¹⁴There are also studies in other developed countries (e.g. Mogstad 2012; Autor *et al.* 2019).

countries there is a massive literature on conditional cash transfer (CCT) programs (see Bastagli *et al.* 2019 or Fiszbein and Schady (2009) for a survey of the literature). Banerjee *et al.* (2017) reanalyzes data from seven randomized controlled trials of government-run CCT programs in six developing countries and find no evidence that these discourage work. However, several studies have found that CCT programs discourage formalization of work (e.g., Gasparini *et al.* 2009; Garganta and Gasparini 2015; Alzúa *et al.* 2013; Araujo *et al.* 2017) and argue that the mechanism seems to be the means test.¹⁵ Of particular interest is Bergolo and Cruces (2016), which studies the labor market impacts of a CCT program in Uruguay and finds that it reduces registered employment by 8 percentage points.¹⁶ Amarante *et al.* (2011) studies a previous poverty-alleviation program in Uruguay (*Plan de Atención Nacional a la Emergencia Social*), and finds that it reduces formal labor supply, primarily among women.¹⁷ Our program is different from these in that it is unconditional and has no time limits.¹⁸ Our program is also particular in that it is indirectly means tested for a subset of beneficiaries (i.e., for those also enrolled in Uruguay's main CCT program). This allows us to exploit a variation (i.e., heterogeneous impacts according to whether the UCT is means tested) that can more directly inform on whether the means test is driving formal labor supply responses, as argued in most of the work that finds that cash assistance programs discourage formal employment.

Our results on education enrollment and attainment also contribute to a closely related literature on welfare and human capital investments for children. While studies on the

¹⁵Table 1 in Banerjee *et al.* (2017) summarizes the findings in the literature on labor supply impacts of CCT programs. Despite the work in this area, evidence of disincentives to formal employment in Latin America is still far from conclusive (Bosch and Manacorda 2012).

¹⁶The result is concentrated among adults under 30 years of age, and the authors find evidence that the fall in registered employment is mostly due to a shift to informality and is not a labor supply response.

¹⁷It has also been studied whether this program impacts the incidence of low birth-weight (Amarante *et al.* 2016), political support (Manacorda *et al.* 2011), and teenage school attendance (Amarante *et al.* 2013).

¹⁸The literature on impacts of UCTs on labor market outcomes in developing countries is considerably smaller than that of CCTs. Bastagli *et al.* (2019) reviews the literature on the impacts of non-contributory cash transfers in low- and middle-income countries from 2010 to 2015. Among the 165 studies that were identified, 55% addressed CCTs, and only 25% addressed UCTs. Moreover, most of the UCT studies were in Sub-Saharan Africa, and none were in Latin America.

effects of welfare programs on human capital investment in the US are scarce (Moffitt 2002), there is abundant work on this topic in the context of CCT programs in developing countries (see Saavedra and Garcia 2012 for a meta-analysis of 42 evaluations of CCT programs and their impacts on educational outcomes in developing countries; Medgyesi and Temesváry 2013 reviews effects on human capital accumulation of CCT programs in high-income OECD countries).¹⁹ Once again, our program is a UCT. McIntosh *et al.* (2010) evaluates a cash transfer program in Malawi that has a CCT and a UCT arm and finds that the conditionality effectively plays a role (i.e., the CCT arm outperformed the UCT arm in terms of school enrollment). Thus it is important to understand the role that a program with no conditionalities could play in educational attainment, in particular given that the literature on this topic is quite limited.²⁰

This paper also contributes to a growing literature on welfare take-up.²¹ Keane and Moffitt (1998) studies the labor supply effects of multiple program participation and shows that it is extremely common for households enrolled in a welfare program to participate in more than one program. In 1984, 89% of beneficiaries of Aid to Families with Dependent Children (AFDC) also received Food Stamps and Medicaid, and 42% received a fourth benefit (Keane and Moffitt 1998). This raises the question of whether multiple-program participation is a consequence of certain overlap on eligibility criteria for welfare programs, or whether it is a consequence of some causality going from enrolling in one welfare program to enrolling in one or more additional welfare programs. Baicker *et al.* (2014) studies whether Medicaid increases participation in Food Stamps, TANF, Supplemental Security Income, and Social Security Disability Insurance and finds a positive and statistically significant

¹⁹The exception in the US is Milleri and Sanders (1997). The authors study the impact of Aid to Families with Dependent Children on educational attainment of young women and find effects that are not statistically significant (and 5 of their 6 coefficients have a negative sign).

²⁰Baird *et al.* (2014) present a systematic review of the literature of the effects of cash transfer programs on schooling outcomes and conclude that “.. simply there are too few rigorous evaluations of UCTs” (Baird *et al.* 2014, p.30).

²¹Part of this literature has dealt with understanding the reasons behind incomplete take-up (Finkelstein and Notowidigdo 2019; Kleven and Kopczuk 2011; Bhargava and Manoli 2015) or the intergenerational transmission of welfare (Boschman *et al.* 2019; Dahl *et al.* 2014b; Dahl and Gielen 2018; Antel 1992). There is also work on peer effects in paid paternity leave take-up in Norway (Dahl *et al.* 2014a).

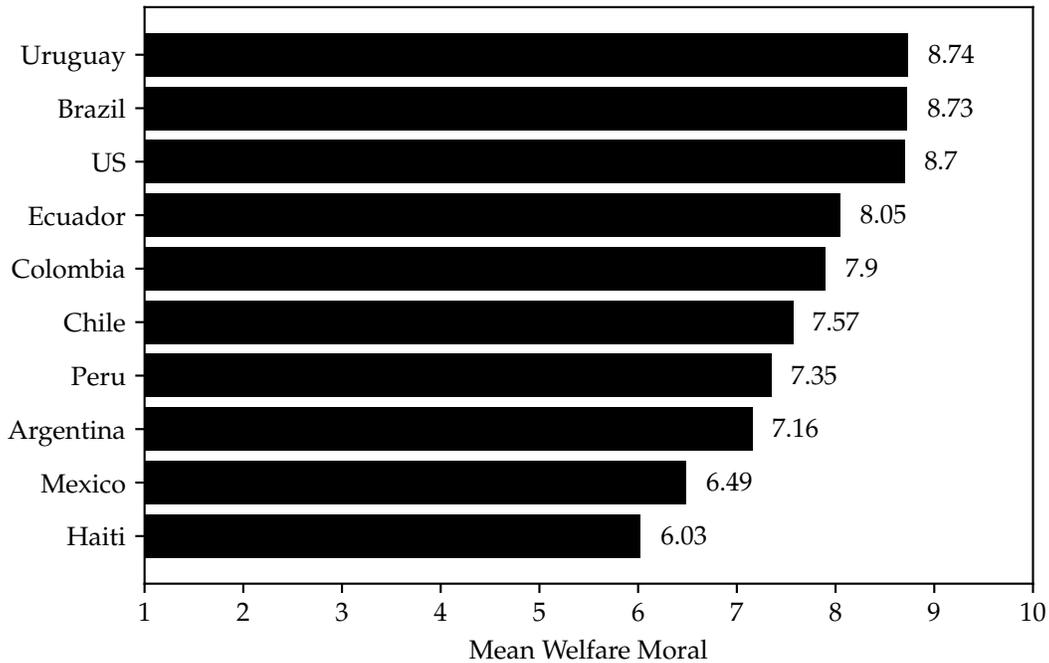
impact on enrollment in Food Stamps.²² The UCT program in Uruguay seems to impact enrollment in the opposite direction (i.e., with a negative impact on enrollment in housing assistance and other cash assistance programs). It would be an interesting avenue of research to understand what drives such differences.²³

Results on government complaints, selective deafness, and misreporting also contribute to the field of forensic economics (see Zitzewitz 2012 for a review). Our results suggest that individuals take actions to stay in the program, which could (perhaps) be considered unethical (e.g., misreporting income, or not answering the door when a government agent shows up, if the individual believes that he will be deemed no longer eligible for the UCT). It is particularly interesting to document these behaviors in a country such as Uruguay where there are high moral standards in regard to participation in welfare, and thus one would not necessarily have expected these behaviors to take place. According to data from the World Value Survey, Uruguay is actually the country in Latin America with the strongest condemnation of claiming government benefits to which one is not entitled (even more so than in the US; see Figure 1.2).

²²Mallar (1982) studies the impact of enrollment in the Job Corps on future receipt of AFDC, General Assistance, Food Stamps, public housing, and unemployment insurance. Overall, he finds significant reductions in receipt of AFDC, General Assistance, Food Stamps, and unemployment insurance.

²³Among the many differences between these programs, Baicker *et al.* (2014) states that Medicaid case workers are instructed (in certain instances) to offer assistance in applying for TANF and SNAP. In Uruguay, this type of assistance is not automatically provided.

Figure 1.2: Welfare Morale



Notes: Each bar represents the mean answer by country to the following question: “Please tell me for each of the following actions whether you think it can always be justified, never be justified, or something in between, using this card: Claiming government benefits to which you are not entitled”. Possible answers are on a scale from 1 (Always justifiable) to 10 (Never justifiable). Data source is the World Value Survey, sixth wave (2010-2014).

The remainder of this paper is organized as follows. Section 1.2 describes the institutional context. Section 1.3 presents the data and some descriptive summary statistics of our sample. Section 1.4 describes the empirical strategy. Section 1.5 presents the main results of the paper. Section 1.6 presents additional evidence to aid in the interpretation of our results and potential mechanisms that drive them. Section 1.7 shows robustness checks of our results. Section 1.8 concludes.

1.2 Institutional context

This section provides background information relevant to understanding this study. It summarizes the information available from governmental reports and from our own conversations with multiple government officials that clarified several aspects of the design and

implementation of the UCT program.

First, we provide an overview of the main components of the Uruguayan welfare state. The objective is to posit the program that we study in the proper context and to understand some of the outcomes that we address later. Second, we describe in detail the UCT program *Tarjeta Uruguay Social*. Third, we explain how the Ministry of Social Development (MIDES) conducts its field work and how the information it gathers is processed to assess eligibility for the UCT program. This is key to understanding our identification strategy.

1.2.1 The Uruguayan welfare state

Uruguay, a small country with 3.3 million inhabitants, is among the most developed countries in South America. With a GDP per capita of USD 17,000 (2018), it is second only to Chile (comparing at PPP), and the United Nations Development Programme situates it in its highest category of human development (“Very high human development”).²⁴ The main pillars of the Uruguayan welfare state are a far-reaching public education system, a generous and solidary social security system (Filgueira 2005), and a health system with broad coverage.²⁵

The Uruguayan government provides free public education at all levels (from elementary school to university) and has established 14 years of mandatory schooling (2 in elementary school, 6 in primary school, and 6 in secondary school).²⁶ Despite practically universal completion of primary school, secondary school completion rates still pose a serious challenge. A report by the National Institute for Educational Assessment and Evaluation (INEEd 2014) shows that only 28% of youths in the 18-20 age group had finished secondary school in 2013.²⁷ Among those enrolled in one of the last three grades at the start of the

²⁴According to The Economist’s 2018 democracy index, it is also the only “full democracy” in South America.

²⁵There are also several housing assistance programs, which are described in the online Appendix B.2.

²⁶The entire system accounted for a public expenditure of 5.1% of GDP in 2018. Some years of secondary school can also be completed at vocational schools.

²⁷This number is 39% if we consider the 21-23 age group.

school year, only 57% passed the grade. Moreover, educational attainment is unevenly distributed. Whereas 64.6% of youths in the 18-20 age group that are in the highest income quintile completed secondary school, only 7% of those in the lowest quintile did.

Another pillar of the system is the safety-net available to formal workers. These workers have to contribute 19.5% of their salary to social security and to a national health plan, and they are eligible for unemployment insurance (which has a maximum duration of six months), an old-age pension (retirement starts at 60 years of age if they contributed to social security for 30 years), paid maternity and paternity leave (98 and 13 days, respectively), access to health care with a public or private provider, and a family monthly allowance if their formal income is below a certain level.²⁸ Nevertheless, there is still a non-negligible share of the population that operates in the informal economy, without this safety net and with significantly lower wages (Amarante and Gómez 2016 finds that in 2014 informal workers earned 30% less than formal workers employed in the same activity). Informality has steadily decreased since 2004, when the share of informal workers reached 40.7%. By 2014 that figure dropped to 23.5% and has remained relatively constant since then (Amarante and Gómez 2016).²⁹

Economic crisis and cash assistance programs. In 2002 a severe economic crisis hit Uruguay which generated a new wave of cash assistance programs to alleviate material hardship.³⁰ Poverty rates reached 23.6% and unemployment 17% (the highest level in 20 years), while real per capita income fell 8% between 2001 and 2002 (Manacorda *et al.* 2011). Capitalizing on the dissatisfaction with the economy and the management of the crisis, in March 2005 a center-left-wing party (*Frente Amplio*) took office. Among the flagship policies of the new government was the creation of an emergency plan (*Plan de Atención Nacional a la Emergencia Social*) that included a cash assistance program targeting the poorest 10% of the

²⁸More details on these benefits and obligations can be found in the online Appendix B.3.

²⁹Despite these levels of informality, Uruguay is among the countries with the lowest rates of informality in Latin America, behind only Chile and Costa Rica (Gasparini and Tornarolli 2009).

³⁰The crisis originated in the banking sector. De la Plaza and Sirtaine (2005) studies the causes and events that led to this crisis.

population.³¹ A newly created Ministry of Social Development (MIDES) was put in charge of its implementation.

By 2008, the crisis was long gone and the government decided to restructure its poverty alleviation programs. The emergency plan was replaced in January 2008 by two new non-contributory cash assistance programs: a CCT program (*Asignaciones Familiares - PE*) and a UCT program (*Tarjeta Uruguay Social*).³² In practice, these two plans served as a continuation of the emergency plan. All families with minors or pregnant women that were previously enrolled in the emergency plan started receiving CCT and UCT.³³

The CCT program targets poor families with children under age 18 and pregnant women, and the amount varies with the number of minors in the household and whether they are enrolled in secondary school.³⁴ The program is managed by the Social Security Administration (SSA), and its selection criterion combines both proxy-means testing and means testing. First, households must complete an application which captures an array of socio-economic data. With these data, the SSA computes a predicted-income score (*Índice de Carencias Críticas*, which we refer to hereinafter as the Vulnerability Index); households for whom this score is above a certain threshold can become eligible for CCT. The second condition they need to meet is to earn less than a certain monthly formal per capita income. The Social Security Administration systematically checks this income for CCT beneficiaries; if a beneficiary household surpasses this threshold for three consecutive months, it loses its

³¹The monthly transfer, which amounted to USD 55, was conditional on school attendance for all children under age 14, and on regular health checkups for pregnant women and all children. However, in practice these conditionalities were never enforced. Families with children and/or pregnant women received an additional transfer ranging from USD 14 to USD 37, depending on the number of children. Manacorda *et al.* (2011) describes the emergency plan in detail.

³²Whenever we refer to the “CCT program”, we mean *Asignaciones Familiares - PE* unless stated otherwise. Similarly, whenever we refer to the “UCT program”, we mean *Tarjeta Uruguay Social*.

³³The transition meant changes in the amount transferred to families. Since the amount of the CCT is based on the number of children (unlike the cash transfer associated with the emergency plan), families with more than two children increased their transfer. Families with two children received approximately their former amount, and families with only one child lowered their monthly transfer by about USD 10.

³⁴There is no cap on the number of beneficiaries per household. The family receives USD 46 for the first child. If the child is enrolled in secondary school, the amount increases by USD 20. For each additional child, those amounts increase at a decreasing rate. A family with N children and n of whom are enrolled in secondary school receives $46 \times N^{0.6} + 20 \times n^{0.6}$ (DINEM 2012).

eligibility for CCT.³⁵

1.2.2 A primer on Uruguay's largest unconditional cash transfer program

Tarjeta Uruguay Social is the main Uruguayan UCT program, and its objective is to “provide food support to people in extreme poverty conditions” (DINEM 2011). Its target population is defined as “those 60,000 households (with or without minors) in a situation of extreme socio-economic vulnerability” (DINEM 2012, p.159), which roughly corresponds to the poorest 5% of households in the country.³⁶

Beneficiary households receive only one magnetic card (which we refer to as the food card), and MIDES, which is the government institution that manages the program, adds money to the food card on a monthly basis.³⁷ The amount depends non-linearly on the number of children comprising each household (up to a maximum of four). In October 2019, a household with two children received the equivalent of 47 current USD.³⁸ Households with no children receive the amount equivalent to one child. Additionally, a household receives double the amount if MIDES concludes that it is within the poorest 30,000 households in the country.³⁹

There are two additional benefits that the UCT provides. First, products paid for with the food card are not taxed with VAT.⁴⁰ Second, for each child under 4 years of age, MIDES

³⁵According to Bergolo and Galvan (2016), by 2014 the maximum monthly formal per capita income was USD 196 for families with two members, and USD 242 for families of more than two.

³⁶When we explain the rules of eligibility, we describe in detail how the concept of “socio-economic vulnerability” is operationalized, and thus give a more precise definition of the target population.

³⁷A household receives these funds on the same day each month, but different households could get the transfer on different days. Depending on the last digit of the national identity number of the card holder, MIDES adds money on day 10, 12, 14, 16, or 18 of each month. We have only monthly level data, so we do not exploit differences in timing of transfer receipt in this study (we also do not have access to the national identity number of the card holder).

³⁸In 2019, a family with one child received USD 31; the amounts for a family with three children and a family with four or more children are USD 60 and USD 84, respectively.

³⁹For brevity, we refer to these households as enrolled in the double-UCT program or beneficiaries of double UCT.

⁴⁰The VAT rate is 22% for most products. Some products, such as flour, rice, bread (not all kinds), cooking

adds an extra amount of 10 USD to the food card (monthly).⁴¹ All transfers are indexed annually by the inflation rate of food (DINEM 2012).

Unlike other Uruguayan cash transfers, the food card is accepted only at “solidary shops”. These are grocery stores and supermarkets that voluntarily signed an agreement with MIDES to accept the food card as a means of payment. These shops are forbidden to sell alcohol, tobacco, or carbonated drinks if the customer is paying with the food card. According to the data provided by the concessionaire of the payment terminals, by 2011 74% of all money transferred was spent on food purchases and 20% was spent on cleaning supplies (CICCA 2012).

The number of solidary shops varied between 700 and 1200 until 2015 (Aguirre *et al.* 2015; MIDES 2016). The network was expanded in 2016. Since then, the number of shops has steadily increased, reaching 2,695 by September 2019. Figure C.5 in the online Appendix shows their distribution across the country. Though solidary shops are concentrated in areas with high population density (in particular, in Montevideo), there seems to be a reasonable coverage of the whole country. As a matter of fact, virtually 100% of the sum of all UCT transfers was spent in 2011, even before the expansion of the network (CICCA 2012).

Origins and re-targeting policy. UCT was launched in January 2008 as a supplementary payment for those 65,000 households with children who were transitioning from the emergency plan. UCT was conceived of as an extra transfer that aimed to warrant food security for extremely poor families. In addition to the families from the emergency plan, 20,000 households were added in May 2009. The latter were former beneficiaries of a food basket program managed by the National Institute of Nutrition (CICCA 2012).

Both the merging of populations from other programs and the vagueness in the definition of the original target population resulted in a controversial internal report on mistargeting. By August 2011, MIDES authorities recognized that the law that created the UCT was not

oil, tea, coffee, and soap, are taxed at 10%.

⁴¹Initially, MIDES provided iron-fortified milk to UCT beneficiaries instead of this extra amount. However, logistical costs and misconduct of sellers paved the way to switching from the milk to an equivalent amount (money-wise) to be added to the food card, starting in May 2016.

sufficiently precise in defining extreme poverty or social vulnerability. In a public report the MIDES division in charge of internal monitoring and assessment (DINEM) stated the following:

“it is desirable to revise and update the definition of the target population (...) it would be of great help to make explicit the UCT objectives and reach (...) It is of the greatest importance for the UCT program to implement a systematized and protocolized procedure in order to grant and withdraw benefits. Moreover, it is important to update information on the beneficiaries.” (DINEM 2011, p. 75)

Absent a precise definition of the target population, DINEM tried to diagnose mistargeting under five different definitions of the target population. In the first scenario, DINEM considered the target population to be all households below the poverty line. Under that definition, they estimated that 63% of poor households were not receiving UCT and that 29% of households receiving it were not poor (i.e., type I and type II errors). A second scenario was to assume that the target population was comprised of households that surpassed a critical threshold according to the proxy-means algorithm used by the CCT program. In that case, 42% of those above the threshold were not receiving UCT, and 64% of beneficiaries did not surpass the threshold (DINEM 2011, pp. 71–73). Under any scenario, mistargeting was a clear concern.

By the end of 2011, a re-targeting policy was set in motion. The target population was then precisely defined, and intensive fieldwork was carried out to refocus the policy. This setting is what provides the quasi-random variation in entry into and exit from the program that we exploit in this study. 65,000 household visits were conducted during 2012. By 2013, all the information collected in those visits was used to re-target benefits. In that year, 32% of those previously receiving UCT were cut from the program (25,167 households). In the same period, 17,874 new households were accepted for receipt of UCT, and this re-targeting policy continued in the following years (DINEM 2012).

1.2.3 Eligibility process and fieldwork

In this subsection, we first describe the proxy-means testing instrument used by MIDES to assess eligibility for the program. Second, we detail how MIDES gathers the necessary information to conduct this test. Third, we comment on the exceptions to this assignment rule.

The Vulnerability Index: a proxy-means test. In 2012 UCT's target population was defined as the 60,000 households in extreme socio-economic vulnerability conditions as measured by the Vulnerability Index (VI). The VI was designed by scholars from an economics research institute of the University of the Republic, initially with the objective of devising a mechanism to select beneficiaries of the CCT program. Given the level of informality among the target population and the possibility of income misreporting, scholars devised an index to predict the probability that a household belongs to the lowest income quintile based on a large set of socio-economic characteristics that were gathered at household visits conducted by MIDES officials (Bergolo and Galvan 2016). Similar targeting mechanisms are applied in other contexts, as proxy-means testing has become the "industry standard" for these types of programs when the poverty status is not directly observable (De Wachter and Galiani 2006).

Specifically, the VI is a highly saturated probit model in which the dependent variable takes value 1 if the household belongs to the lowest quintile of per capita income.⁴² The model was estimated initially in 2009, and re-estimated in 2012 using survey data from the 2011 National Household Survey (NHS), which is a representative sample of the population. Only households below the median per capita income and with minors were included in the estimation sample. Additionally, two different model specifications were defined, based on a territorial criterion: one for the capital city of Montevideo, and another for the rest of

⁴²Although the VI model is confidential, Manacorda *et al.* (2011) lists some of the independent variables used in a previous version of the VI, which are suggestive of the independent variables that are currently used. These are "an indicator for public employees in the household, an indicator for pensioners in the household, average years of education of individuals over age 18, the number of household members, the presence of children by age group (0–5 and 12–17), an indicator variable for whether a member of the household had private health insurance, residential overcrowding, whether the household was renting, toilet facilities (no toilet, flush toilet, pit latrine, other), and a wealth index based on durables ownership (e.g., refrigerator, TV, car, etc.)" (p. 6).

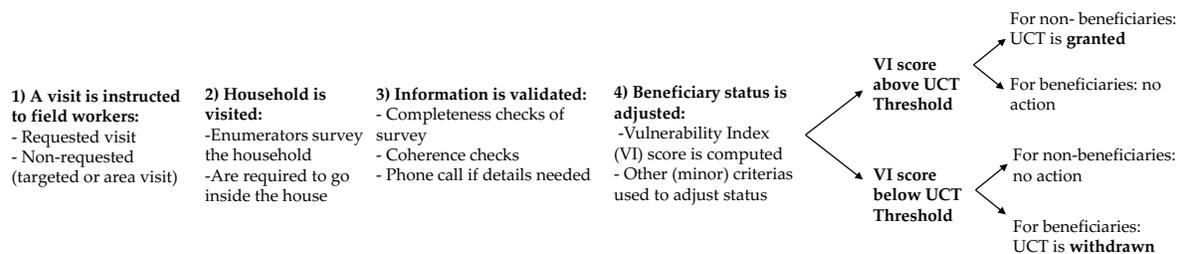
the country.

Even though the VI was designed specifically for CCT (it was used since 2009 for this purpose), in 2012 MIDES decided to use it to define eligibility for UCT. Given that UCT targets the 60,000 most vulnerable households, a threshold had to be defined. VI scores were computed for each household in the 2011 NHS sample, and a value was selected such that only 60,000 households had a score higher than the threshold. The same procedure was used to select the threshold value for double UCT, by predicting the value that corresponds to only the 30,000 households with the most critical values (those closest to 1).⁴³ By 2012, in order to receive UCT a household needed an estimated VI score above 0.62 for Montevideo and above 0.70 for the rest of the country. If a household's VI score was above 0.756 in Montevideo or 0.81 in the rest of the country, then it was eligible for double UCT. Importantly, households never learn their score and are not told which variables enter the algorithm or any specifics on the eligibility thresholds (Bergolo and Galvan 2016).

Household visits: feeding the algorithm. The main way a household can enter/exit the program is by receiving a household visit by a MIDES agent. The agent conducts a survey on site to measure a wide range of variables, including those needed to calculate a household's VI score. By the time that the MIDES report on mistargeting went public, the MIDES administration had decided to conduct extensive fieldwork in order to correct the inclusion and exclusion errors detected (DINEM 2012). This meant both withdrawing UCT from ineligible former beneficiaries, and granting UCT to those eligible. Figure 1.3 shows the steps followed by MIDES that eventually lead to a grant or a withdrawal of an UCT. We explain these steps in what follows.

⁴³One could argue that the VI may be less predictive of household per capita income in the case of households with no children (as the probit was estimated on a sample of households with children). However, to the best of our knowledge, there are no studies of the type I/II errors of the VI as a targeting mechanism for the UCT.

Figure 1.3: Steps followed by MIDES to grant or withdraw UCT



Notes: This figure schematically shows the steps taken by MIDES (since 2012) to grant/withdraw an UCT. It is based on information obtained from regulatory decrees, internal reports elaborated by MIDES (DINEM 2011, 2012), and interviews to MIDES officials.

The first step in order to grant (or withdraw) UCT is the decision of which households to visit. These are chosen with the objective of both reducing leakage (i.e., withdrawing benefits from ineligible beneficiaries) and increasing coverage of the eligible population. We distinguish between two kinds of visits: area visits and targeted visits.

Area visits are conducted to all households in a specific geographical area. All dwellings in the area are intended to be visited by a group of 4 to 7 enumerators and a field supervisor. Particularly at the start of the re-targeting policy, one of the inputs used to decide these areas was the national census conducted in 2011 by the National Institute of Statistics. Satellite images were also used to define these visits. We refer the interested reader to the online Appendix C.1 for a more detailed description of how the census and satellite images were used to decide area visits.

A second type of visit, which we call targeted visits, is used to survey current beneficiaries (whose address is already known), with the objective of checking whether they are still eligible. This mode also comprises family requests to become a new recipient of UCT. The visits are performed by a single enumerator who is provided with a map and a list of 10 to 20 households to visit.⁴⁴

Single enumerators usually work in a 5- to 10-block area surrounded by other enumera-

⁴⁴An example of these maps can be found in the online Appendix, Figure C.6. In case a “targeted” household cannot be located, the household form goes to a “search office” that cross-checks data with other ministries and public services in order to find the correct address.

tors and a supervisor in charge of the entire neighborhood. Supervisors randomly perform visits along with enumerators in order to check the quality of their visits. Specifically, it is required that enumerators go inside houses in order to diminish the probability of misreporting.

Since 2012, most MIDES visits are performed by the Unit of Programs Monitoring (PMU). The PMU is in charge of conducting visits across the country by enumerators from the capital city. Though it is more expensive to send employees from Montevideo to the rest of the country, the authorities decided to do this in an effort to diminish the probability of collusion between local agents and potential beneficiaries. Moreover, a household is not previously notified that a visit is going to be conducted.

The fieldwork can produce different visit outcomes. In order for a visit to be considered successful, enumerators must find an adult capable of answering all questions. In case of a completed visit, the data collected go directly through the selection process described earlier. Besides, if a household is successfully located but could not be visited, either because there were no adults present or because no one at the house answered the door, the “unsuccessful” visit is administratively tagged in order to be considered in a future targeted visit.

Once the visit is performed, it must go through quality control. The first step is a rapid completeness check in the field by a supervisor. Once they are back in the office, reviewers perform different completeness and coherence tests and call the household in case they need to clarify any part of the form. When the quality control process ends, all the data go through a VI algorithm that calculates the household’s VI score. If the VI score is above the UCT threshold (and the household was a non-beneficiary of UCT prior to the visit), then the office in charge of UCT creates the respective UCT account and prints a new magnetic card. Once the food card is printed, a MIDES employee calls the head of the household to ask him or her to go to a MIDES office in order to sign the contract and accept the food card. Money is added to the food card once it is accepted. If the VI score is below the threshold (and the household was a beneficiary of UCT prior to the visit), MIDES stops adding money to the food card; there are minor exceptions to this rule. MIDES officials call the household

to notify them that they will no longer be beneficiaries. Notification may be done by letter as well.⁴⁵ In no case are regional MIDES deputies in charge of notifying beneficiaries of withdrawals. However, both MIDES regional deputies and officials at headquarters usually receive re-visit requests to have the lost benefit restored.

Exceptions to the rule. If a household has a VI score that does not surpass the eligibility threshold, it is still possible to enroll in the program. In February 2012, a special UCT share was created for transsexual people. By July 2013, a lower threshold for acceptance in UCT was instituted for households in which at least one member was enrolled in a priority social program.⁴⁶ Moreover, UCT enrollment was not to be removed for a year after enrollment in that priority social program ended. Another way to receive UCT is through a “special share” managed by regional MIDES deputies to deal with critical situations. These deputies have discretion to allocate a small number of UCTs. Transfers allocated through this mechanism comprise 0.5% of all UCTs and can last only 6 months or until a visit to the beneficiary is conducted. Additionally, in April 2015 MIDES created an index (similar to the VI, called the Complementary Vulnerability Index) that assigns more weight to variables that measure housing vulnerability. Starting in April 2015, if a household’s VI score does not surpass the UCT eligibility threshold, MIDES computes this second index and if it is above a given threshold, then the household is eligible to receive UCT. Moreover, beneficiaries of housing programs were given a grace period of two years in which their benefit is not withdrawn despite losing eligibility by the VI criterion. In 2016, homeless people also became eligible to receive the transfer.

By June 2018, 92% of all UCT beneficiaries had been visited by MIDES agents and had a VI score above the threshold. Of those whose VI is below the threshold, 22% receive UCT because of extreme housing vulnerability (i.e., surpass the threshold for the Complementary Vulnerability Index), 20% are transsexuals, 26% are enrolled in a priority social program,

⁴⁵See grant and withdrawal letters in the online Appendix C.2.

⁴⁶These programs are *Jóvenes en Red*, *Uruguay Crece Contigo*, and *Cercanías*. They provide intense monitoring and counseling via social workers to extremely poor teenagers, children, and families, respectively.

and the others receive UCT via the special shares allocated by the regional office (provisional benefits) and UCT for the homeless (DINEM 2018).

1.3 Data and summary statistics

Our empirical evidence is based on administrative datasets from 5 different governmental sources that offer a comprehensive view of how beneficiaries are affected by entry into and exit from the UCT program.⁴⁷ All our datasets contain information at the individual level which is indexed with a de-identified national identity number that allows us to link the datasets.⁴⁸

1.3.1 Data from the Ministry of Social Development

Our primary dataset comes from MIDES and consists of the responses to the comprehensive questionnaire applied by MIDES agents during the household visits conducted during the period January 2011 to July 2018.⁴⁹ 255,767 households (1,012,170 individuals) were visited during this period, with 35% of these visits taking place in the capital city (Montevideo). Table B.1 in the online Appendix shows the distribution of these visits across time. Whenever we refer to “our sample of individuals”, we mean those that inhabited a visited household (at the actual time of the visit).

Two key variables that we use in all our analysis and that come from this dataset are the date of the visit and the VI score of the household. As will be explained in Section 1.4, having a VI score above a certain threshold is the key instrument that allows us to get quasi-random variation in entry into and exit from the program. Figure 1.4 shows the

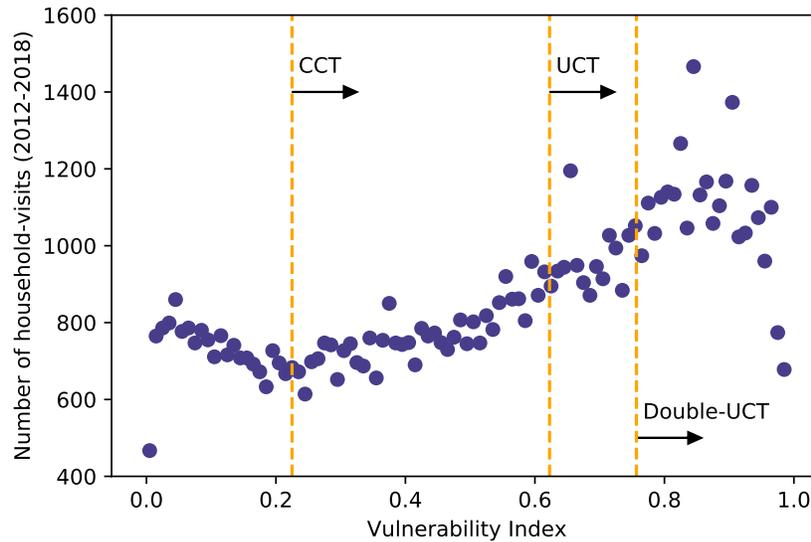
⁴⁷Some of these datasets were provided directly by the respective government organization, while others were accessed through the recently developed integrated system of social analysis (*Sistema de Información Integrada del Área Social - SIAS*).

⁴⁸For all citizens and foreign residents of Uruguay, it is mandatory to have a national identity number. Data from the Americas Barometer (2016) indicate that 99.6% have a national identity number in Uruguay.

⁴⁹See the online Appendix F for a complete translation of all the questions asked at these visits (the original Spanish version is also shown).

distribution of household visits according to their VI score.⁵⁰ We also indicate the thresholds that are used to assess eligibility for CCT, UCT, and double UCT.

Figure 1.4: *Density of household visits by Vulnerability Index (Montevideo)*



Notes: This figure shows the number of households visited by MIDES from January 2011 to July 2018, as a function of their Vulnerability Index score computed on the basis of the visit. From left to right: the first orange line indicates the minimum score needed to gain eligibility for the CCT program, the second line corresponds to the UCT program, and the third line to double UCT. The source of the data is MIDES. Panel *a* shows visits in Montevideo; panel *b* shows visits in the rest of the country.

We also exploit other information gathered at the household visits. In particular, we look at responses on food insecurity questions, durable goods consumption, housing conditions, labor supply, and schooling. Table 1.1 shows some characteristics of the population under study (i.e., individuals living in a visited household).

⁵⁰We show this distribution for visits in Montevideo. Figure B.1 in the online Appendix shows the distribution of visits taking place in the rest of the country.

Table 1.1: *Characteristics of visited households and individuals at the time of the visit (mean values)*

<i>Demographics</i>	
Household size	4.0
Number of minors in the household	2.0
Female	55.6%
Years of education (adults)	7.0
Pregnant (for women, ages 18-40)	4.9%
In Montevideo	34.9%
<i>Safety net programs</i>	
Receives no UCT	69.4%
Receives UCT (not double)	25.0%
Receives double UCT	5.5%
UCT (not double) amount (% household income, all sources)	17.4%
Double UCT amount (% household income, all sources)	30.3%
Beneficiary of cash-assistance for formal workers	2.0%
Participates in a housing program	5.3%
Receives CCT	69.3%
<i>Income and material well-being</i>	
Household income (in USD; all sources)	440.2
Household income (in USD; excludes transfers)	346.5
Reports food insecurity	45.4%
Reports food insecurity for adults	85.2%
Reports food insecurity for minors	47.6%
A household member went to a soup kitchen (last month)	8.0%
<i>Status in the job market (14 years old and older)</i>	
Private sector employee	23.6%
Public sector employee	2.9%
Worker at a cooperative	0.1%
Self-employed	26.6%
Employer (<i>Patrón</i>)	0.1%
Unpaid worker	0.4%
Unemployed	14.5%
Retired (<i>Jubilado</i>)	5.2%
Receives a pension (<i>Pensionista</i>)	5.6%
In charge of household chores	16.0%
Other inactive	5.0%

Notes: This figure shows mean values of different characteristics of our sample measured at the time of the visit. Source of the data on take-up of safety net programs is MIDES (for UCT), the Social Security Bank (for CCT), and SIIAS (for housing and cash assistance for formal workers). Household income from all sources is constructed using households' responses to the survey conducted by MIDES agents and administrative data on UCT and CCT amounts transferred: self reported income from all work sources, plus self-reported income on pensions, plus CCT and UCT amounts transferred to all member of the household. Household income that excludes transfers only considers the self reported income from all work sources. For the rest of the variables, the source of the data is the survey conducted by MIDES agents.

In addition to the responses gathered at household visits, MIDES provided us with three pieces of information. The first of these was a monthly dataset of UCT recipients that includes their national ID number (de-identified) and the amount transferred each month, from August 2009 to September 2018.⁵¹ These data were made available to us for only the individuals in our sample.⁵²

Second, MIDES shared with us a dataset containing the administrative records of re-visit requests with their respective national ID number (de-identified) and the date of the request.

Third, we were provided with information regarding how each household visit form was internally processed at the Ministry. Specifically, we know both the date of the visit and the date that quality control reviewers finished analyzing the visit information and logged it into the system. This allows us to construct a measure of internal processing time for each visit which will be used in our identification strategy. Figure B.2 in the online Appendix shows the average processing time for visits that took place in a given month and year, where the processing time is defined as the number of months elapsed between the date of the visit and the date on which the information gathered at the visit was logged in by the reviewers in the system. The processing time was high at the start of the re-targeting policy, but in mid/late 2013 it was sharply reduced and it is currently around one month.

1.3.2 Formal labor supply data

We have monthly level data (for January 2010 to September 2018) on whether an individual is an active worker in the formal sector, as shown by his contributions to the National Health Insurance Plan (FONASA).⁵³ Since 2008, every formal worker, regardless of his economic activity, must contribute a fixed percentage of his income to FONASA. That

⁵¹Hereinafter, we use the term UCT recipients for individuals that actually receive the transfer (i.e., the card holders) and the term UCT beneficiaries for all individuals that live in a household with a UCT recipient.

⁵²However, this corresponds to practically the entire UCT population. According to an internal report of the Cash Transfers Division of MIDES, by June 2018 92% of all beneficiaries had been visited at least once (and thus were included in our sample).

⁵³McIntosh *et al.* (2017) also uses this variable to measure formal work in Uruguay.

contribution entitles the worker and his nuclear family to choose a public or private health care provider.⁵⁴ There are two exceptions to this rule. First, the military and the police are exempt from this contribution and are not enrolled in FONASA. However, data from the National Household Survey (2013) show that only about 0.5% of workers that belong to a household in the lowest income ventile (i.e., the UCT target population) are in the military or police. Thus we do not believe that any changes we find in our measure of formal employment are due to a switch to or from police/military work. Second, formal employees that earn less than 150 USD per month or work fewer than 13 days a month are not legally bound to contribute to FONASA. Thus if an individual is listed as an active worker in FONASA in one month but not in the next, it could be because he is no longer working in the formal sector at all or because he reduced his participation at the intensive margin (i.e., the number of days worked and/or the formal income earned).⁵⁵

1.3.3 Education data

We gained access to administrative data on public education enrollment, generated by the National Administration of Public Education (ANEP). We accessed individual data on enrollment and grade level for the three different systems that comprise ANEP: elementary and primary; secondary; and vocational. The available data on education are for April 2013 to December 2017.⁵⁶

Although ANEP data do not include information on enrollment in private institutions, it is unlikely that children in our sample attend these institutions. According to data from the National Household Survey, during our period of study 99% of those in the lowest income

⁵⁴Every formal worker must contribute at least 4.5% of his income to FONASA. Additionally, a worker's partner and children are entitled to enroll in FONASA; in that case the worker's contribution increases to 8%.

⁵⁵We gathered data on a second indicator of formal employment for validation: monthly social security contributions to the Social Security Administration. These contributions are restricted to certain activities, while FONASA is not, but within those activities, all workers have to contribute to social security. We have data on this variable for only a subset of our sample (those enrolled in CCT, and only for their months of enrollment). The correlation between our two measures is 0.9.

⁵⁶See the online Appendix B.1 for more details on the frequency of these data.

quintile that attend school do so in the public system.

1.3.4 Data on enrollment in other safety-net programs

There are several programs offered by the state that are designed to act as safety-nets and to support families that are socio-economically vulnerable. We were able to match individual data on enrollment in several social programs that we describe below.

First, we accessed individual administrative records on CCT receipt. Data on CCT transfers are available for January 2012 to September 2018. We can distinguish both the children entitled to the benefit, the adult that receives the transfer, and the total amount transferred. This variable is relevant not as an outcome, but rather as a way of partitioning the sample to test the role of the means test. As described in Section 1.2.1, the CCT is means tested while the UCT is not. However, the UCT is indirectly means tested if households also receive the CCT. Thus looking at heterogeneous impacts of the UCT according to CCT status can serve as a test for the role of the means test.

Second, we matched our data with individual administrative records on beneficiaries of any housing assistance program managed by the Ministry of Housing (July 2012 to September 2018). We note that the dummy is “cumulative” in the sense that if an individual is a beneficiary of a housing program at some point in time, it will still show up as a beneficiary in the future.⁵⁷

Third, we matched our data with records on receipt of a cash transfer for formal employees with dependents.⁵⁸ This benefit is not a universal right, and only people who engage in certain economic activities are entitled to it. The data were recorded monthly and are available for January 2012 to August 2018.

⁵⁷See the online Appendix B.2 for a description of the different housing programs offered by the Ministry of Housing.

⁵⁸This transfer is regulated by law No. 15,084. See the online Appendix B.3 for a description of this program.

1.4 Econometric framework

This study uses a regression discontinuity approach to study how entry into and exit from a cash assistance program impact formal labor supply, schooling, and take-up of other safety-net programs. The main empirical strategy consists of a fuzzy regression discontinuity design (fuzzy RD), and it exploits the fact that the probability of assignment to treatment (i.e., being a beneficiary of UCT) changes discontinuously at a given threshold of the VI. We also estimate a dynamic differences-in-differences model that complements the fuzzy RD estimates. We first specify the regression form of our fuzzy RD and examine the plausibility of its identifying assumptions. Then we do the same for our dynamic differences-in-differences design.

1.4.1 Fuzzy regression discontinuity design

In its most basic setup, a fuzzy RD can be conceived of as an application of instrumental variable regression in which the treatment is instrumented with whether a running variable surpasses a given threshold. In our case, the running variable is the VI score of a household (measured on the basis of the visit), and the treatment is enrollment in the UCT program after the visit. We slightly augment this specification to allow for the fact that surpassing the threshold may have different impacts on beneficiary status depending on the MIDES processing time. As explained in Section 1.2.3, MIDES does not instantly adjust the beneficiary status after a visit. The information gathered at a household visit has to reach MIDES headquarters in the capital city, and a government official (the *reviewer*) has to check and log the information into an internal system before any change in beneficiary status takes place. Specifically, the first stage is the following:

$$UCT_{h,(t+t_0,t+t_f)} = \alpha_0 + \alpha_1 \mathbb{1}[VI_{h,t} > 0] + \alpha_2 \mathbb{1}[VI_{h,t} > 0] \times ProcTime_{h,t} + \alpha_3 ProcTime_{h,t} + f(VI_{h,t}) + \phi X_{i,h,t} + w_{i,h,t} \quad (1.1)$$

Where $UCT_{h,(t+t_0,t+t_f)}$ is the share of months in the range $t + t_0$ to $t + t_f$ that any member

of household h receives UCT (we generally pick $t_0 = 1$ and $t_f = 36$), $VI_{h,t}$ is the VI score of household h that was visited at time t , and $X_{i,h,t}$ denotes additional individual controls for individual i living in household h at time t .⁵⁹ $Proc\ Time_{h,t}$ is a leave-out version of processing time for household h visited at time t (i.e., the mean number of months it took *reviewers* to process household visits that took place at time t , except the one that corresponds to household h).

The RD polynomial is denoted by the letter f ; we choose an asymmetric control function and a polynomial of degree 1 (i.e., a local linear regression). We pick a bandwidth of 0.1 in all our regressions. This is the maximum bandwidth size that we can choose and still satisfy the constraint that we not surpass a threshold where another policy change occurs.⁶⁰ Bandwidths selected using data-driven methods (Calonico *et al.* 2014 and Imbens and Kalyanaraman 2012) generally pick larger bandwidths, so we stick with this smaller bandwidth across all our regressions. Section 1.7 and the online Appendix D.2 shows the robustness of our main results to different bandwidth sizes (including the optimal bandwidth selected with data-driven methods) and to a triangular kernel (as suggested by Gelman and Imbens 2019).

The second stage is the following:

$$Y_{i,h,T} = \beta_0 + \beta_1 U\hat{C}T_{i,h,(t+t_0,t+t_f)} + \beta_2 Proc\ Time_{h,t} + f(VI_{h,t}) + \gamma X_{i,h,t} + \epsilon_{i,h,T} \quad (1.2)$$

Where $Y_{i,h,T}$ is some outcome variable for individual i from household h measured at a time T .⁶¹ In our baseline fuzzy RD results, we look at outcomes three years after the visit ($T = t + 36$). In all cases, we cluster standard errors at the household level.

⁵⁹To improve statistical power, we follow Haushofer and Shapiro (2016) and condition on the pre-visit level of the outcome variable when possible. In addition, we include year and month fixed effects, a dummy for female, a dummy for Montevideo, age, and age squared.

⁶⁰Specifically, we choose a bandwidth of 0.1097. Households not in Montevideo that surpass the UCT threshold by more than 0.1097 duplicate the amount they receive in UCT. In Montevideo they duplicate the amount if they surpass the UCT threshold by more than 0.134. Also, eligibility for some family counseling programs managed by MIDES (such as *Programa Cercanías*) jumps exactly at this level (Perazzo *et al.* 2016).

⁶¹Note that in the first stage the dependent variable does not depend on i , but technically speaking, its fitted value could depend on i if we include individual controls in the regression.

Before discussing identification, we make three slightly technical comments to clarify our unit of observation, how we measure impacts over outcomes measured at a re-visit, and the assumptions that underlie our choice of standard errors.

A comment on the unit of observation. The treatment status (i.e., UCT beneficiary status) in this setting is assigned at the household level. Thus one could consider studying how the UCT impacts the treated household across time. However, this is possible only if the composition of the household is observed in every period. We observe only a snapshot of the household composition at the time of the visit, so technically speaking, we cannot state how the household was impacted, as we cannot follow that entity across time. Also, even if the composition was observed in every period (if the household splits in two, for example), we need ad-hoc assumptions to study the evolution of household-level outcomes. The natural path forward is then to follow individuals exclusively. There are two types of outcomes we look at: outcomes measured at the individual level, either with high frequency or at the time of the visit (e.g., formal employment, enrollment in public schooling) and outcomes measured at the household level only at the time of the visit. For both sets of outcomes, the unit of observation is always the pair individual-visit. This seems natural for individual-level outcomes, but it may not be evident for household-level variables. This is what our estimates capture when looking at household-level outcomes: if someone lives in a household that enters/exits the UCT program, how does that impact the characteristics of that person's household in the future? The impact can be due to a change that takes place in the initial household or because the change in the UCT status induces individuals to move to a different household. It also means that given two individuals i, j living in household h , household-level outcomes for individuals i, j are the same at the time of the initial visit but can be different in future visits.

Measuring impacts on outcomes measured at a re-visit. We can measure the impact of entering/exiting the UCT program on outcomes measured at the household visit for the subset of our sample that was visited more than once. Out of the 773,770 individuals that were visited, 75.7% were visited only once, 19.1% were visited twice, 4.2% three times, and

1.1% more than three times. We use the fuzzy RD regression specification described earlier with the VI score computed on the basis of visit number n to individual i as the running variable, and a certain outcome variable for individual i measured at his next visit ($n + 1$). We use the beneficiary status at the time of the visit $n + 1$ as the endogenous regressor. One concern that this analysis raises is that being revisited may be endogenous to crossing the UCT eligibility threshold at the previous visit, and thus the sample of people that are revisited is differentially selected on whether they crossed (or not) the threshold in the past. Figure B.3 in the online Appendix suggests that this is effectively the case. An individual is ≈ 10 pp less likely to be revisited if he crossed the threshold on a visit (i.e., if he became eligible to receive UCT).

To address this concern, we add the following controls to our specification (in addition to dummies for female and Montevideo): the value of the outcome variable measured in the visit n , age (and its square) at visit $n + 1$, number of months (and its square) between visits n and $n + 1$, and year FE and month FE of visit $n + 1$. The first control is key, as it allows us to rule out the possibility that differences we find in a given outcome are attributed to differential selection on that outcome (time-invariant selection). The other controls allow us to rule out the possibility that differences we find are due to a difference in timing of visits across the groups (above/below the thresholds).⁶²

A comment on standard errors and re-visits. There are two characteristics of our sample that have implications for the relevance of our instrument and our choice of standard errors.

First, 5% of all individuals visited have at least one visit with a VI score that falls within the bandwidth and a future visit with a VI score outside the bandwidth.⁶³ This could imply some reduction in the relevance of our instrument (i.e., as individuals can have multiple visits, the beneficiary status is not determined solely by what happens at a given visit, as it

⁶²What kind of selection problem could still invalidate these results? Suppose we find that the UCT seems to have a negative impact on a variable x at visit $n + 1$. We would need that facing a drop in x during the period between visits n and $n + 1$ is more positively correlated with getting a re-visit of the eligible group at n than with getting a re-visit of the ineligible group at n .

⁶³A VI score falls within the bandwidth if it is within 10.97 pp of the UCT eligibility threshold.

can be changed with future visits).

Second, 4% of all individuals visited have at least two visits with VI scores that fall within the bandwidth. This has the same implication as before but also introduces an additional challenge. Our actual unit of observation is not the individual but rather the pair individual-visit. Thus the same individual (if visited twice with associated VI scores that fall within the bandwidth) appears twice in our regression, and by clustering SE at the household-visit level we are not allowing for positive correlations between the error terms of this individual across visits.⁶⁴ We construct an alternative cluster variable which we call “extended household” and the results remain practically unchanged.⁶⁵

Identification

Internal validity of the fuzzy RD estimates is assessed by checking that the following holds around the threshold: 1) all relevant factors other than the beneficiary status vary smoothly; 2) there is no selective sorting or manipulation; 3) no other policies change discontinuously; and 4) crossing the threshold induces a jump in the probability of being treated (i.e., of being a beneficiary of the UCT).

We check for balance in multiple pre-visit characteristics (or characteristics measured on the basis of the visit at which MIDES surveys the household), both in the entire sample and in the Montevideo/non-Montevideo split. Table 1.2 presents the RD estimate (with no controls, and an asymmetric linear control function) with a pre-visit characteristic as dependent variable centered at the threshold (we consider a bandwidth of 0.1).

Out of 53 coefficients, only 4 are statistically significant (and at the 10% level), so we conclude that the conditional expectation of the potential outcomes seems to be continuous at the threshold. We first look at basic demographics such as percent female in the household,

⁶⁴This concern has been raised in several other event studies (e.g., Jäger and Heining 2019).

⁶⁵The results are available upon request. Each cluster is constructed as follows: 1) pick a household visit h and call that cluster C ; 2) assign all individual-visits $i - v$ in h to cluster C ; 3) assign all individual-visits $i - v'$ to cluster C for $v \neq v'$; 4) with h' denoting the household of $i - v'$, assign all individual-visits $j - v'$ in h' to C ; 5) repeat steps 3 and 4, but with j instead of i , until no more individual-visits get added to C ; 6) go back to step 1, but with an h that is not part of the C constructed in the previous step.

mean age, year of the visit, household size, percent of visits in Montevideo, and women in the 18-40 age group that are pregnant. With the exception of age in Montevideo (which is marginally significant at the 10% level), the coefficients are statistically insignificant. Second, we look at three food insecurity measures, and despite some imbalance in the non-Montevideo sample, the other coefficients do not suggest a systematic bias in this domain. Third, we look at variables related to employment, income, and education. For schooling and formal labor supply, we use exactly the same sample that is used to run our regressions, and it is reassuring to find practically no imbalance in these variables. Fourth, we look at whether there is some imbalance in pre-visit enrollment in safety-net programs (including the baseline values of outcomes we study later), or in previous visit requests, as this is related to an outcome variable we study later (re-visit requests after the initial visit). Finally, we look at whether there is an imbalance in voting in previous Participatory Budgeting Elections in Montevideo.⁶⁶ This variable is interesting, as it captures political participation, which is not captured in our other variables; we do not find that those who are more “politically active” are differentially selected around the threshold.

Table 1.2: *Balance on observables*

	Uruguay	Montevideo	Non-Montevideo
<i>Demographics</i>			
Female	0.002 (0.003)	0.009 (0.006)	-0.001 (0.004)
Age	0.117 (0.126)	0.434* (0.255)	-0.011 (0.144)
Year	0.012 (0.027)	0.037 (0.051)	0.004 (0.032)
Household size	0.016 (0.028)	0.003 (0.055)	0.022 (0.032)
Montevideo	0.003 (0.007)	- -	- -
Pregnant (ages 18-40)	0.003	-0.003	0.005

Continued on next page

⁶⁶These are local elections that take place every three years in the capital city, where citizens of age 16 or over can vote on a number of projects to be funded by the municipal government. See Cabannes (2004) for a systematic analysis of participatory budgeting experiences in Latin America and Europe.

Table 1.2: (Continued) Balance on observables

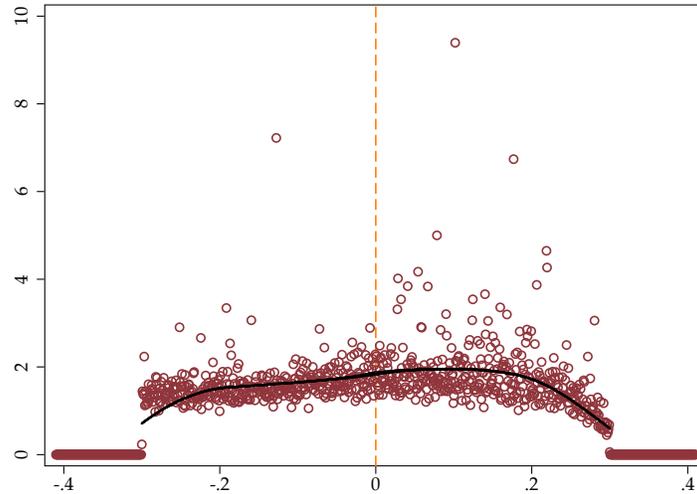
	Uruguay	Montevideo	Non-Montevideo
	(0.004)	(0.006)	(0.005)
<i>Food insecurity</i>			
Food insecurity	-0.002 (0.008)	-0.005 (0.015)	-0.0 (0.01)
Food insecurity, adults	-0.013 (0.008)	0.002 (0.015)	-0.018* (0.01)
Food insecurity, minors	-0.013 (0.011)	0.011 (0.021)	-0.022* (0.013)
<i>Employment, Income and Education</i>			
Formal worker (t-1)	0.003 (0.005)	0.008 (0.012)	0.0 (0.006)
Worker	-0.0 (0.007)	0.002 (0.015)	-0.002 (0.008)
Household income (USD)	9.97 (7.124)	3.625 (14.97)	12.041 (7.929)
In school (t-1)	-0.014 (0.016)	0.003 (0.033)	-0.02 (0.019)
Yrs of education (t-1)	-0.027 (0.028)	0.055 (0.053)	-0.06* (0.034)
<i>Welfare and Political Participation</i>			
CCT beneficiary (t-1)	-0.003 (0.005)	0.001 (0.012)	-0.004 (0.006)
Housing program (t-1)	-0.0 (0.004)	-0.01 (0.009)	0.004 (0.005)
Cash-assistance for formal workers (t-1)	0.0 (0.001)	-0.001 (0.003)	0.0 (0.002)
Voted in Participatory Budgeting	- -	-0.004 (0.008)	- -
Household requested visit	-0.004 (0.012)	0.006 (0.025)	-0.006 (0.014)

Notes: Table shows RD estimates considering an asymmetric linear control function (no additional controls), and a bandwidth of 0.1097 around the threshold that defines UCT eligibility. Standard errors clustered at the household level are shown in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Except it is stated otherwise, all variables are measured at the time of the visit. $t - 1$ corresponds to the month before the visit.

To check for selective sorting or manipulation, we test the continuity of the running variable density function around the threshold, following McCrary (2008). Figure 1.5 shows the results of the McCrary test, with a t-statistic of 0.025, which means we fail to reject the null hypothesis of no discontinuity in the density function of the running variable at the

threshold.⁶⁷

Figure 1.5: McCrary test



Notes: This figure shows the density of household visits (using our full sample) around the UCT eligibility threshold, where a McCrary test (McCrary 2008) was performed. The t-stat is equal to 0.025. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

To the best of our knowledge, there are no other policy changes occurring on either side of the cutoff (and within the bandwidth). The VI is also used by the Social Security Administration to define CCT eligibility and by MIDES to define eligibility to receive double UCT. However, the cutoff for CCT and double UCT in Montevideo (non-Montevideo) are 0.22488131 (0.25648701) and 0.7568 (0.81), respectively. The cutoff for (single) UCT in Montevideo (non-Montevideo) is 0.62260002 (0.70024848), so by picking a bandwidth of 0.1 we are excluding these other policy changes from our estimation sample.⁶⁸

Finally, we show that crossing the UCT threshold induces a discrete jump in the probability of receiving UCT after the visit.⁶⁹ In Figure 1.6 we show the beneficiary status 3, 6, 12,

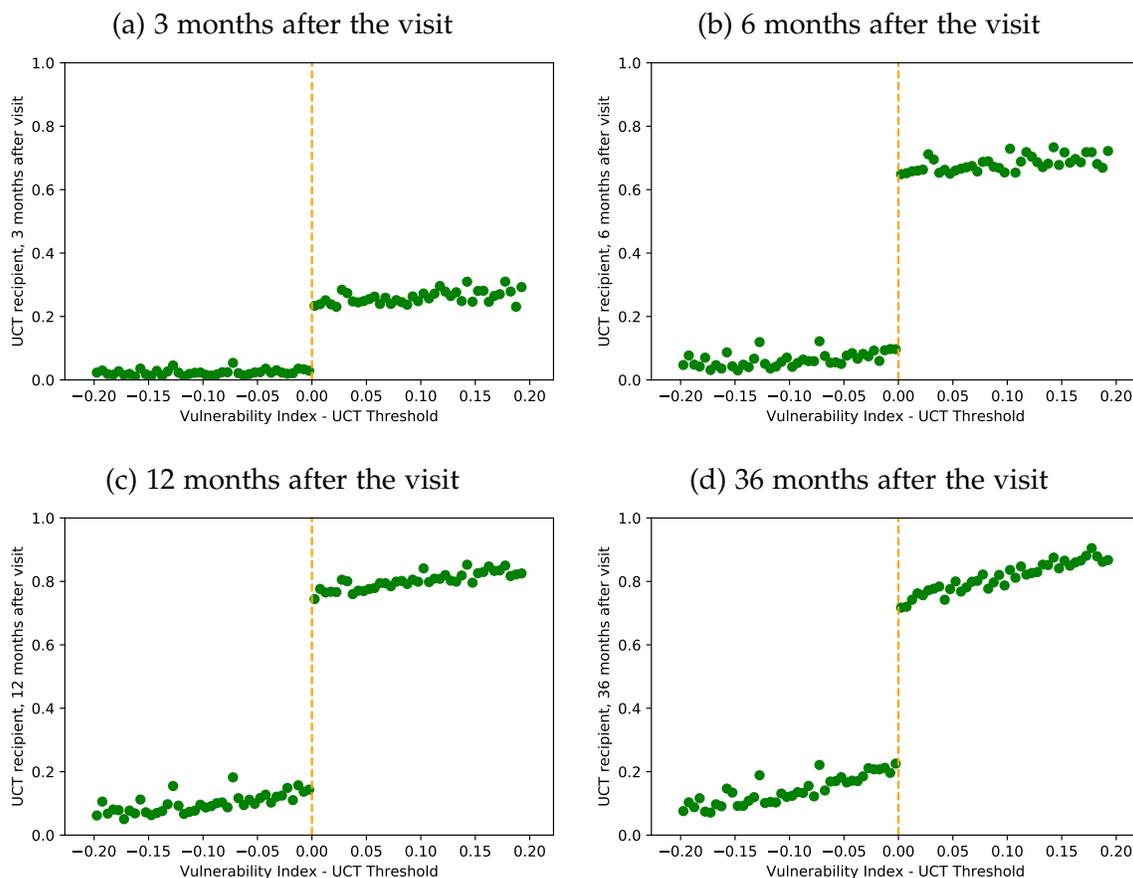
⁶⁷The online Appendix B.6 shows results of the McCrary test splitting the sample into Montevideo and non-Montevideo subsamples. The same conclusions apply (t-stats are -0.27 and 0.19 , respectively).

⁶⁸Another program that uses the VI to define eligibility is *Programa Cercanías*. Our estimates are not confounding the impact of this program, as eligibility is restricted to those with a VI higher than the double-UCT eligibility threshold (Perazzo *et al.* 2016).

⁶⁹We show the first stage using the double-UCT threshold in the online Appendix B.5.

and 36 months after the visit as a function of the VI score for households that were initially not receiving UCT.

Figure 1.6: *First stage for pre-visit non-beneficiaries*



Notes: This figure shows binscatters of the share of households enrolled in the UCT program 3, 6, 12, and 36 months after a visit takes place as a function of their Vulnerability Index scores (normalized to 0 at the UCT eligibility threshold). These binscatters use only the sample of pre-visit non-beneficiaries of UCT. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

Three points are worth mentioning. First, the beneficiary status is not instantly updated after the visit, and it peaks and (partially) stabilizes around the 12th month. Second, although we see a discrete jump (upwards) in the probability of receiving a UCT at the threshold several months after the visit, it does not jump from 0 to 1 (the jump is of size approximately 60 pp one year after the visit). This is due to some exceptions to this

assignment mechanism (see Section 1.2.2), because individuals can get revisited, and also because UCT can be withdrawn if the household is receiving CCT but does not meet CCT's means test at some point in time. Third, this design corresponds to what some authors call a type II fuzzy regression discontinuity design (see Jacob and Zhu 2012), as we have both "crossovers" (i.e., households below the threshold that get treated) and "no-shows" (i.e., households above the threshold that do not get treated).⁷⁰

1.4.2 Dynamic differences-in-differences design

We employ a dynamic differences-in-differences design (DID) when looking at impacts on formal labor supply (this is the outcome for which we have the longest time series). The advantage of this strategy relative to the fuzzy RD is threefold. First, it allows us to measure the impact of going from 2 to 0 transfers (and vice versa), instead of from 1 to 0. This constitutes a stronger treatment. Second, we no longer restrict ourselves to picking observations within the bandwidth and thus consider a larger sample size, which can help reduce some of the noise in our estimates. Third, our dynamic DID specification is orthogonal to our choice of the endogenous regressor and instruments in the fuzzy RD. Thus it can serve as a sort of robustness check that those decisions are not driving our results.⁷¹

To estimate the impact of exiting the program, our treatment group is comprised of those individuals that had a VI score below the UCT threshold and that were initially beneficiaries of UCT. Pre-visit beneficiaries with VI scores between the single-UCT threshold and the double-UCT threshold constitute the control group. To estimate the impact of losing two transfers, we consider only individuals that were (pre-visit) receiving double UCT. The treatment group is comprised of those with a VI score below the single-UCT threshold, and

⁷⁰The same conclusions apply if we look at households that were initially recipients of UCT; see Figure B.4 in the online Appendix.

⁷¹Nevertheless, we perform robustness checks on our fuzzy RD estimates in Section 1.7.

the control group corresponds to those with a VI score above the double-UCT threshold.⁷² This captures the reduced-form impact of crossing the threshold on an outcome of interest when crossing the threshold is associated with a certain change in the beneficiary status in the UCT program.⁷³ Note that the treatment and control groups used to study the impact of losing one transfer never overlap with the groups used to study the impact of exiting from double UCT. Thus if we find similar results in the two specifications (perhaps stronger effects when exiting from double UCT), it is reassuring that we are estimating an exit from the program, as it is not mechanical that the two results go in the same direction. To study the impact of gaining one or two transfers, there is overlap only in the control group.

The regression equation is the following:⁷⁴

$$Y_{i,h,t} = \beta_0 + \gamma_{i,h} + \sum_{k=-l}^{k=L} \beta_k^{Treated} \times Treated_h \times \mathbb{1}[t = k] + \sum_{k=-l}^{k=L} \beta_k \times \mathbb{1}[t = k] + yearFE_{h,t} + monthFE_{h,t} + age_{i,h,t} + age_{i,h,t}^2 + \epsilon_{i,h,t} \quad (1.3)$$

Here t represents the number of months before or after the visit (e.g., $t = 2$ would correspond to the observation taking place two months after the visit to household h).

1.5 Main results

We study how entry into and forced exit from the UCT program impact labor supply and formalization of work, human capital investments for children, and take-up of other

⁷²Analogous definitions are used to measure the impact of gaining single or double UCT. Also, to rule out the possibility that our results are driven by entry into or exit from the CCT program, we include only individuals that were beneficiaries of CCT before the visit and one year after the visit.

⁷³We could do this event study with the date on which the beneficiary status changed as $t = 0$ instead of the visit date. However, we could expect to see changes in behavior starting with the month of the visit if expectations of receiving UCT change with the visit. This is especially relevant in a DID setting (as opposed to an RD setting), as individuals with high enough VI scores could presumably expect that they will receive UCT at some point, and households with low enough VI scores could expect that they will most likely not be part of the UCT program. Thus the visit itself changes expectations on receipt of UCT differentially for the treatment and control groups.

⁷⁴Two regressors are excluded from the regression ($Treated_h \times \mathbb{1}[t = -1]$ and $\mathbb{1}[t = -1]$), which normalize $\beta_{-1}^{Treated}$ to 0.

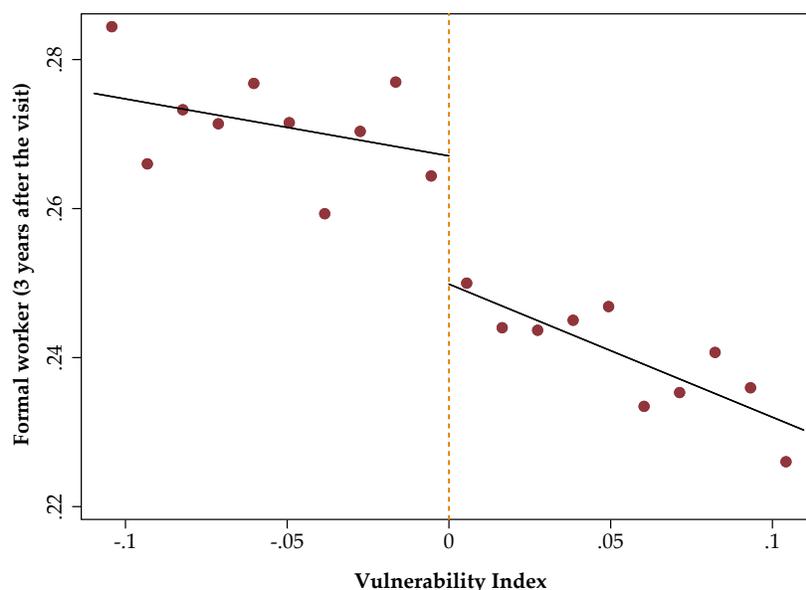
safety-net programs. We first show binscatters to visually inspect the reduced-form result of crossing the UCT threshold on the outcome of interest.⁷⁵ Second, we estimate our fuzzy RD regression where we show results for the “pooled sample” (i.e., not restricting the sample according to pre-visit beneficiary status), as well as for the sample of pre-visit non-beneficiaries and for the sample of pre-visit beneficiaries.

1.5.1 Formal labor supply

We start with what is perhaps one of the oldest topics of interest in labor economics: the relationship between welfare programs and the labor market (Chan and Moffitt 2018). Figure 1.7 shows the reduced-form impact of crossing the eligibility threshold for the UCT on the probability that an individual is formally employed three years after the visit. We focus on individuals in the 18-38 age group three years after the visit whose UCT benefit is (or would be if they receive it) indirectly means tested (i.e., someone in the household is a beneficiary of CCT at the time of the visit), where results are relatively stronger.

⁷⁵We do not comment on the exact magnitude or statistical significance of the results when looking at the binscatters, as this could be misleading. However, we do this when we show the fuzzy RD results. As this is a fuzzy RD and not a sharp design, the actual discontinuity found in the binscatters has to be scaled by the appropriate “first stage” estimate.

Figure 1.7: *Impact of UCT on formal labor supply*



Notes: The variable plotted on the vertical axis is the percentage of individuals within a VI bin that are formally occupied three years after the visit. The horizontal axis shows the Vulnerability Index score computed on the basis of the visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. The sample corresponds to all individuals that were in the 18-38 age group three years after the visit. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, dummy for being formally employed during the month before the visit, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

We see that there seems to be a drop in the probability of formal employment three years after the visit if an individual is on the UCT-eligible side of the threshold, and Table 1.3 confirms this observation. Column 1 shows that receiving the transfer during the three years after the visit decreases formal labor supply at the end of those three years by 2.7 percentage points. This effect is driven by those that were initially not receiving the transfer, which decreased their formal labor supply by 4.4 pp. We see no changes among those that were initially receiving the transfer, although in some sub-samples we do find a mild impact among this population (see the online Appendix B.3).

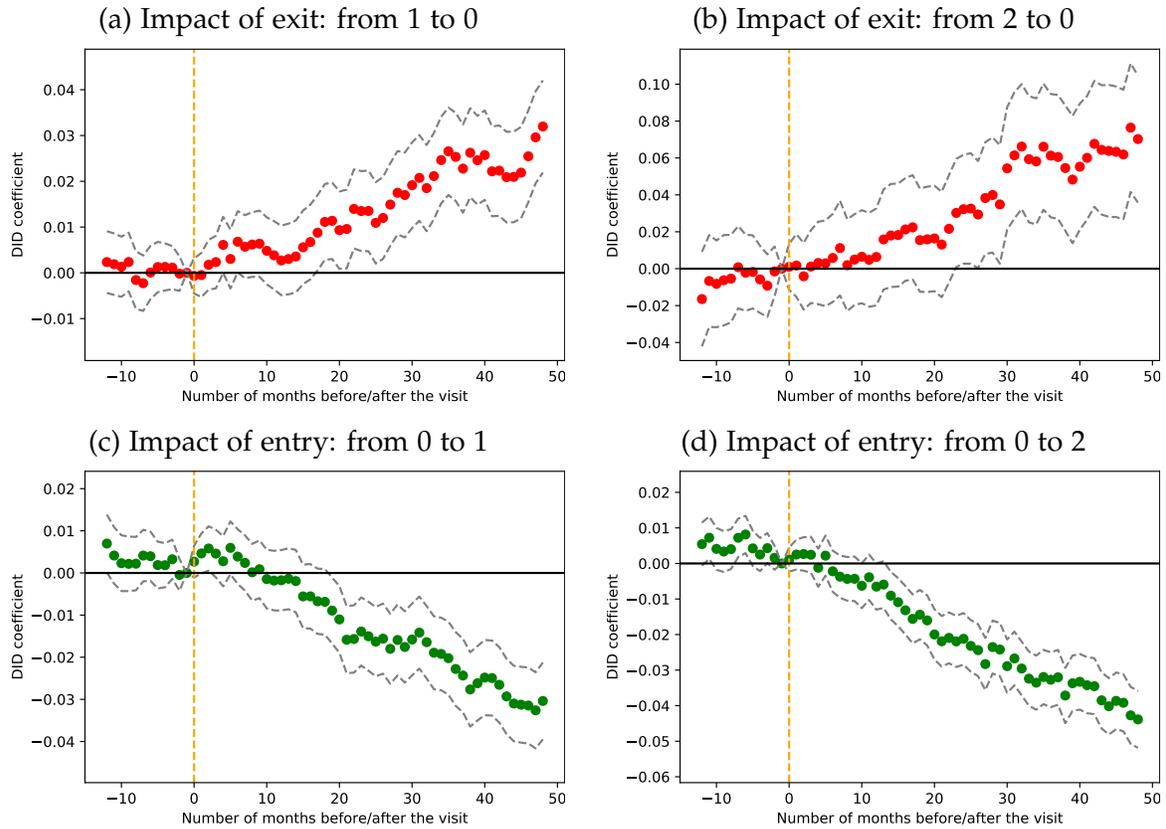
Table 1.3: *Impact of UCT on formal labor supply: fuzzy RD estimates*

	Perc. formally employed 36 mths after		
	Pooled (1)	Non-benef. (2)	Benef. (3)
UCT (1-36 mths after)	-0.027** (0.012)	-0.044*** (0.016)	-0.002 (0.017)
Observations	65885	38517	27368
Mean non - recipients	0.304	0.306	0.247
Bandwidth	0.1097	0.1097	0.1097
RD Polynomial	Linear	Linear	Linear
Controls	Yes	Yes	Yes
Kernel	Rectangular	Rectangular	Rectangular
SE	Cluster	Cluster	Cluster
First stage estimate	0.60	0.56	0.67
F-Stat (First Stage)	8477.0	4551.0	4739.0
P-val: (2) = (3)		0.092	

Notes: Standard errors clustered at the household-visit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bandwidth is set to 0.1097. The dependent variable is a dummy indicating whether the individual is formally employed, 36 months after the visit. Endogeneous regressor is the share of months (within 1-36 months after the visit) in which someone in the household received an UCT from MIDES. Asymmetric and linear RD polynomial considered. Control variables considered are: year FE, month FE, and dummy for Montevideo. Sample consists of individuals that were in the 18-38 age group three years after the visit and that were CCT beneficiaries at the time of the visit. First stage estimate corresponds to the coefficient of the impact of crossing the UCT threshold on the endogeneous regressor (i.e., α_1 in equation 1.1). p-value of the t-test between the difference in the estimates of column 2 vs 3 are presented in the table.

We complement the fuzzy RD estimates with estimates from a dynamic DID model, where we do find statistically significant impacts both of entry into and exit from the program. Panel *a* in Figure 1.8 shows that losing the transfer has a positive impact on formal labor supply (up to +3 pp). Panel *b* shows the impact of a larger loss that we could not measure in our fuzzy RD: initially (pre-visit) receiving double UCT and losing all UCT benefits after the visit. We see basically the same picture, perhaps with slightly higher estimates (as would be expected, given that the loss is larger than in panel *a*, although it could also be related to the pre-trends). Panels *c* and *d* show the impact of gaining a transfer, going from no transfer to either (single) UCT or double UCT. It is reassuring to see that in these cases the impacts on formal labor supply are negative.

Figure 1.8: *Dynamic DID estimates: formal employment*



Notes: This figure shows regression coefficients and associated 95% confidence intervals for the difference between treatment and control individuals k months before/after the visit, that is, the $\beta_k^{Treated}$ from equation (1.3). The coefficient for $k = -1$ is normalized to 0. The outcome variable is a dummy equal to 1 if the individual is formally employed. Besides leads and lags and their interaction with treatment, the specification includes individual-visit FE, age, age squared, year FE, and month FE. Standard errors are clustered at the household level. All regressions include individuals in the 18-38 age group three years after the visit that were beneficiaries of the CCT program, both before and after the visit. Panel (a): The treatment group are those whose VI is below the single-UCT threshold and that were receiving single UCT at the time of the visit; the control group are those with VI between the single-UCT threshold and the double-UCT threshold. Panel (b): The treatment group are those whose VI is below the single-UCT threshold and that were receiving double UCT at the time of the visit; the control group are those with VI above the double-UCT threshold. Panel (c): The treatment group are those whose VI is between the single-UCT threshold and the double-UCT threshold and that were not receiving UCT at the time of the visit; the control group are those with VI below the single-UCT threshold. Panel (d): The treatment group are those whose VI is above the double-UCT threshold and that were not receiving UCT at the time of the visit; the control group are those with VI below the single-UCT threshold).

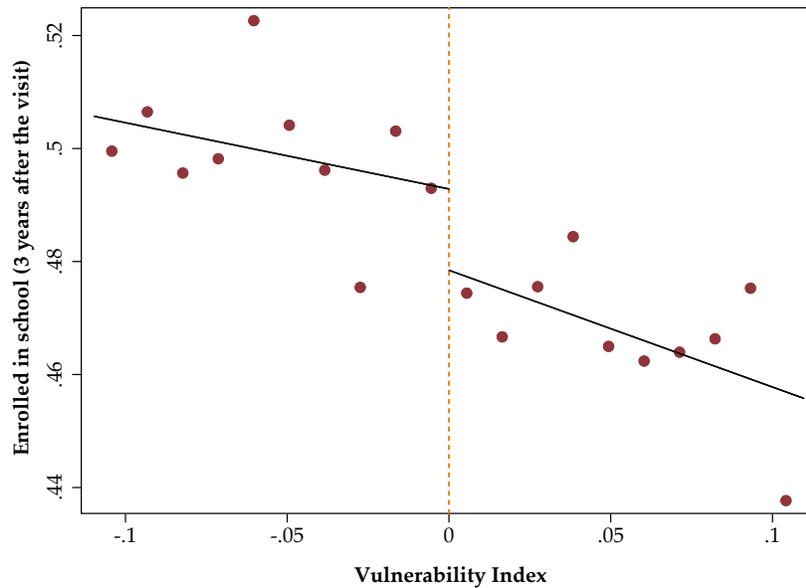
We refer the interested reader to the online Appendix D.1 for results on non-beneficiaries of CCT at the time of the visit and for results on the next 20-year cohort (the 39-59 age group; note that the retirement age starts at 60). We basically find no statistically significant impacts on these groups (with the exception that we do find some impacts on the next 20-year cohort in the DID specification, and in the same direction than the results shown in Figure 1.8). Given that our results on labor supply are being entirely driven by the sample of people whose transfer is means tested (i.e., beneficiaries of CCT), this is suggestive that the mechanism driving the formal labor supply responses is the means testing of the CCT program.

1.5.2 Education

We look at impacts on enrollment in public schools.⁷⁶ As explained in Section 1.2.1, public education in Uruguay can be divided into three subsystems: CEIP, which offers elementary and primary education; CES, which offers secondary education; and CETP, which offers secondary and vocational education. Figure 1.9 shows the probability of being enrolled in any of those three subsystems three years after the visit as a function of the VI score. Our sample corresponds to minors that finished primary school (or are in the last year of primary school) and still need at least three more years of study to finish secondary school. We see that crossing the threshold seems to induce a drop in enrollment rates.

⁷⁶If we look at the poorest 5% of households (which would correspond to the target population of the UCT program) in a nationally representative survey in 2013 (“Encuesta Continua de Hogares INE 2013”), among those minors that go to school, 99% do so in the public system. Thus we interpret our estimates as the impact on schooling overall (not only public schooling).

Figure 1.9: *Impact of UCT on youth enrollment in public schools*



Notes: The variable plotted on the vertical axis is the percentage of individuals within a VI bin that are enrolled in a public school (primary, secondary, and/or vocational) three years after the visit. The horizontal axis shows the Vulnerability Index score computed on the basis of the visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. The sample corresponds to minors that finished primary school (or are in the last year of primary school) and need at least three more years to finish secondary school. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, a dummy for enrollment in public school during the month before the visit, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

Table 1.4 shows the results of the fuzzy RD estimates, which confirm what was apparent (though not entirely clear) in the binscatter. Being enrolled in the program during the three years following the visit reduces the probability of being enrolled in the public education system three years after the visit by 3.3 pp. This effect is driven entirely by those that were initially not receiving the transfer.

Table 1.4: *Impact of UCT on youth enrollment in public schools: fuzzy RD estimates*

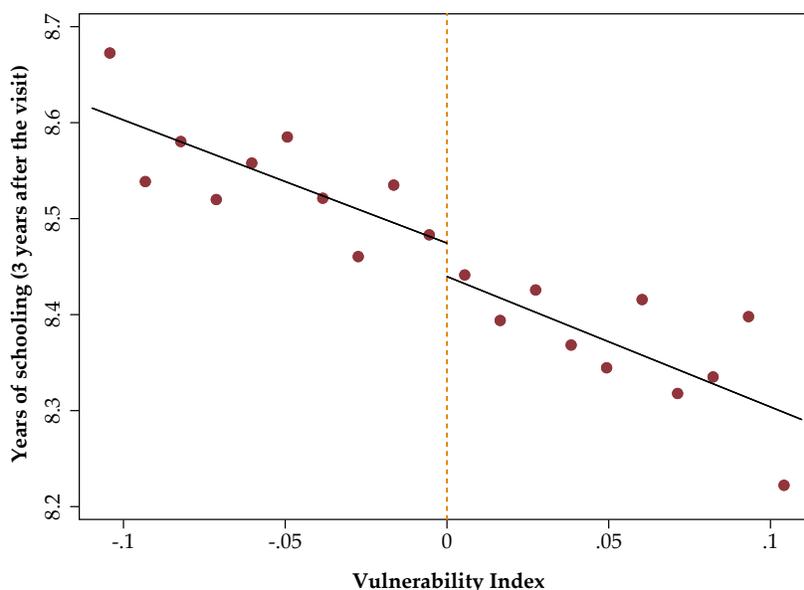
	In public school 36 mths after		
	Pooled (1)	Non-benef. (2)	Benef. (3)
UCT (1-36 mths after)	-0.033* (0.02)	-0.075** (0.032)	0.005 (0.025)
Observations	24132	10721	13411
Mean non - recipients	0.522	0.523	0.502
Bandwidth	0.1097	0.1097	0.1097
RD Polynomial	Linear	Linear	Linear
Controls	Yes	Yes	Yes
Kernel	Rectangular	Rectangular	Rectangular
SE	Cluster	Cluster	Cluster
First stage estimate	0.64	0.56	0.72
F-Stat (First Stage)	3351.0	1245.0	2661.0
P-val: (2) = (3)		0.295	

Notes: Standard errors clustered at the household-visit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bandwidth is set to 0.1097. The dependent variable is a dummy indicating whether the individual is enrolled in any of the educational subsystems of ANEP, 36 months after the visit. Endogeneous regressor is the share of months (within 1-36 months after the visit) in which someone in the household received an UCT from MIDES. Asymmetric and linear RD polynomial considered. Control variables considered are: year FE, month FE, and dummy for Montevideo. Sample consists of minors that either finished primary school or are studying in its last year and still need at least three more years of study to finish secondary school. First stage estimate corresponds to the coefficient of the impact of crossing the UCT threshold on the endogeneous regressor (i.e., α_1 in equation 1.1). p-value of the t-test between the difference in the estimates of column 2 vs 3 are presented in the table.

The results are puzzling. It seems that entering the program decreases human capital accumulation of children. However, we note that passing rates in this context are remarkably low, and so it is not clear whether a decrease in enrollment leads to a drop in completed years of schooling. We have administrative data on the grade in which the individual is enrolled, which allows us to compute the number of completed years of schooling. Figure 1.10 shows that crossing the threshold does not seem to have a clear impact on this measure. Table D.6 in the online Appendix confirms our visual inspection of the binscatter.⁷⁷

⁷⁷While point estimates suggest a drop in completed years of schooling, the effect is non-statistically significant.

Figure 1.10: *Impact of UCT on completed years of schooling*



Notes: The variable plotted on the vertical axis is the mean number of completed years of schooling three years after the visit for individuals within a VI bin. The horizontal axis shows the Vulnerability Index score computed on the basis of the visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. The sample corresponds to minors that finished primary school (or are in the last year of primary school) and need at least three more years to finish secondary school. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, completed years of schooling before the visit, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

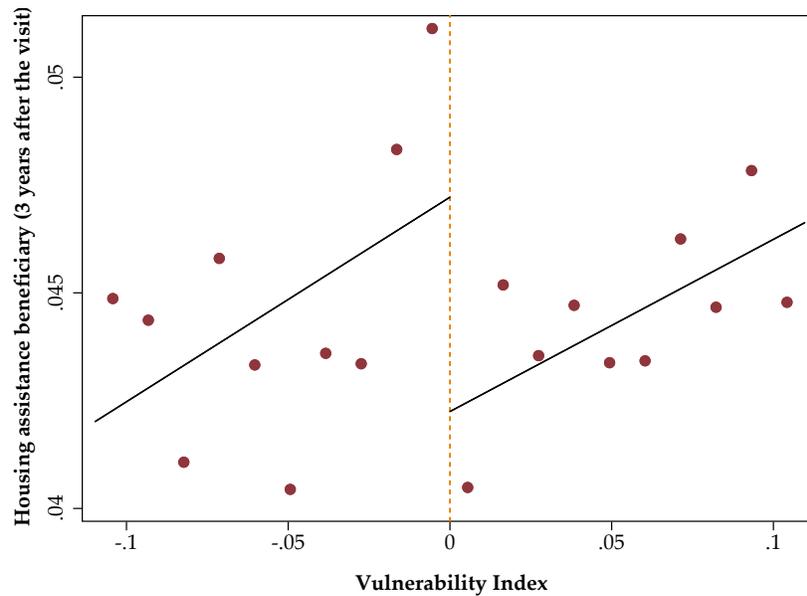
1.5.3 Take-up of safety-net programs

Another form of dependency could take place if receipt of UCT induces individuals to take up more welfare. In particular, we study whether UCT impacts enrollment in housing assistance programs managed by the Ministry of Housing or enrollment in a cash assistance program for formal workers managed by the Social Security Administration. These programs are managed by different institutions that do not use the VI to assess eligibility.

Figure 1.11 shows the reduced-form impact of crossing the UCT eligibility threshold on the probability that an individual is enrolled in a housing assistance program three years after the visit. There seems to be a drop in that probability if an individual is on

the UCT-eligible side of the threshold. However, given the low enrollment rates of this population in housing assistance programs, the binscatter is noisy and it is not clear from just visual inspection whether there is actually a drop in enrollment.

Figure 1.11: *Impact of UCT on take-up of housing assistance*



Notes: The variable plotted on the vertical axis is the percentage of individuals within a VI bin that are enrolled in a housing assistance program offered by the Ministry of Housing three years after the visit. The horizontal axis shows the Vulnerability Index score computed on the basis of the visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

Table 1.5 shows that there is actually evidence that the program induces a drop in enrollment in housing assistance programs. Column 1 shows that receiving the transfer during the three years after the visit decreases enrollment in housing programs at the end of those three years by 1.2 percentage points.⁷⁸

⁷⁸In Montevideo, where results are stronger, we find that both entry (-3.2 pp) and exit (+4.7) have a statistically significant impact on enrollment in housing assistance programs. See Table D.7 in the online Appendix.

Table 1.5: Impact of UCT on take-up of housing assistance: fuzzy RD estimates

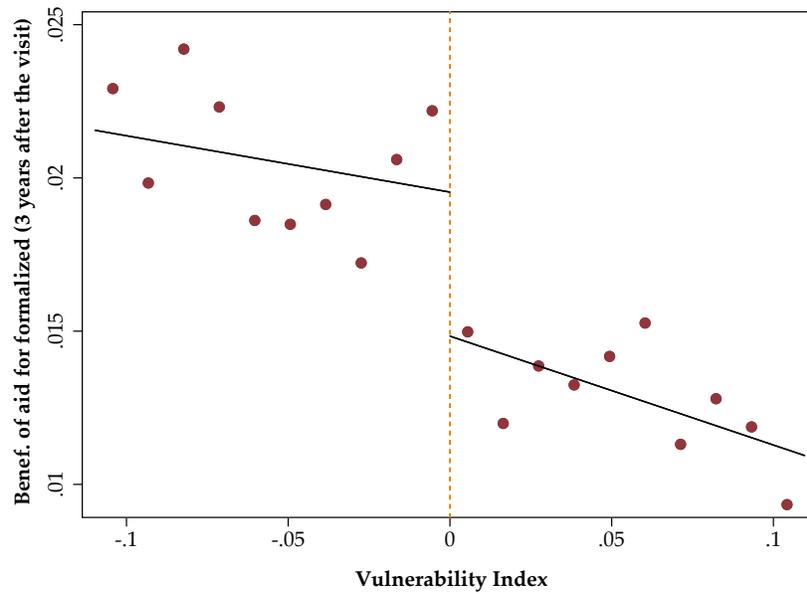
	In a housing program, 3yrs after		
	Pooled (1)	Non-benef. (2)	Benef. (3)
UCT (1-36 mths after)	-0.012** (0.005)	-0.01 (0.007)	-0.014* (0.007)
Observations	74882	38367	36515
Mean non - recipients	0.042	0.041	0.064
Bandwidth	0.1097	0.1097	0.1097
RD Polynomial	Linear	Linear	Linear
Controls	Yes	Yes	Yes
Kernel	Rectangular	Rectangular	Rectangular
SE	Cluster	Cluster	Cluster
First stage estimate	0.60	0.55	0.67
F-Stat (First Stage)	9013.0	4147.0	6225.0
P-val: (2) = (3)		0.538	

Notes: Standard errors clustered at the household-visit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bandwidth is set to 0.1097. The dependent variable is a dummy variable indicating enrollment in a housing assistance program three years after the visit. Endogeneous regressor is the share of months (within 1-36 months after the visit) in which someone in the household received an UCT from MIDES. Asymmetric and linear RD polynomial considered. Control variables considered are: year FE, month FE, and dummy for Montevideo. Sample consists of home owners. First stage estimate corresponds to the coefficient of the impact of crossing the UCT threshold on the endogeneous regressor (i.e., α_1 in equation 1.1). p-value of the t-test between the difference in the estimates of column 2 vs 3 are presented in the table.

Finally, we study whether UCT has an impact on enrollment in a cash assistance program offered to formal workers with dependents.⁷⁹ Figure 1.12 shows the reduced-form impact of crossing the UCT threshold on enrollment in this program for adults living in households with children at the time of the visit. Column 1 in Table 1.6 shows that receiving UCT during the three years after the visit decreases enrollment in this cash assistance program by 0.8 pp (the mean for non-UCT recipients is 2.7 pp, so the effect is non-negligible). This effect is driven entirely by pre-visit non-beneficiaries.

⁷⁹See Section 1.3.4 for details on this program.

Figure 1.12: *Impact of UCT on take-up of cash assistance for formal workers*



Notes: The variable plotted on the vertical axis is the percentage of individuals within a VI bin that were receiving cash assistance for formal workers (in the program managed by the SSA) three years after the visit. The horizontal axis shows the Vulnerability Index score computed on the basis of the visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. The sample corresponds to adults living in households with children at the time of the visit. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, dummy for receipt of cash assistance for formal workers during the month before the visit, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

Table 1.6: *Impact of UCT on take-up of cash-assistance for formal workers: fuzzy RD estimates*

	Beneficiary, 3yrs after		
	Pooled (1)	Non-benef. (2)	Benef. (3)
UCT (1-36 mths after)	-0.008** (0.003)	-0.011** (0.005)	-0.003 (0.004)
Observations	81454	49056	32398
Mean non - recipients	0.027	0.027	0.03
Bandwidth	0.1097	0.1097	0.1097
RD Polynomial	Linear	Linear	Linear
Controls	Yes	Yes	Yes
Kernel	Rectangular	Rectangular	Rectangular
SE	Cluster	Cluster	Cluster
First stage estimate	0.60	0.56	0.68
F-Stat (First Stage)	10189.0	5671.0	5719.0
P-val: (2) = (3)		0.252	

Notes: Standard errors clustered at the household-visit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bandwidth is set to 0.1097. The dependent variable is a dummy variable indicating enrollment in the cash assistance program for formal workers three years after the visit. Endogeneous regressor is the share of months (within 1-36 months after the visit) in which someone in the household received an UCT from MIDES. Asymmetric and linear RD polynomial considered. Control variables considered are: year FE, month FE, and dummy for Montevideo. Sample consists of adults living with minors in the household. First stage estimate corresponds to the coefficient of the impact of crossing the UCT threshold on the endogeneous regressor (i.e., α_1 in equation 1.1). p-value of the t-test between the difference in the estimates of column 2 vs 3 are presented in the table.

1.6 Discussion

The previous section presented our main results but did not discuss some aspects of their interpretation and potential mechanisms. Also, we did not look at other outcomes that could be affected and that could impact our overall assessment of the UCT program (relevant for any policy implications that could be derived from this study). In this section, we first show how UCT positively impacts material hardship in the household. This not only indicates that the program has beneficial effects on the main outcome it is supposed to impact but also suggests a possible mechanism behind the substitution effect we find with housing assistance programs. Second, we exploit our survey data, and the fact that a subset of households are revisited, to re-examine the impacts of the program on labor supply

and education. These results suggest that formal labor supply responses operate mostly through the formality margin, and corroborate the drop we see in school enrollment with administrative data. Third, we show that individuals engage in several strategic behaviors to try to enter or stay in the program.

1.6.1 Material hardship

While in this study we focus on three key outcomes to assess family welfare dependency (i.e., formal labor supply, education, and take-up of other safety-net programs), there are other key outcomes that this program impacts. In particular, we look at its effects on food insecurity, durable goods consumption, and housing conditions.

There are three questions that MIDES asks at the visit that could serve as proxies for food insecurity: *In the last 30 days, did it ever happen that there was not enough food in the household because of lack of money? In the last 30 days, was it impossible for you or another adult in the household to have breakfast, lunch, or something to eat because of lack of money? In the last 30 days, did it ever happen that a child in your household had less to eat than what it was accustomed to because of lack of money?* These questions are similar to the USDA food insufficiency question that has been asked in several surveys since 1977.⁸⁰

Figure 1.13 shows results for the question on food insecurity in the household for minors.⁸¹ We see that crossing the threshold seems to induce a sharp drop in food insecurity.⁸² These results suggest that the transfer is spent, at least in part, on food. However, part of it could also be spent on other goods or services that the household may need.⁸³

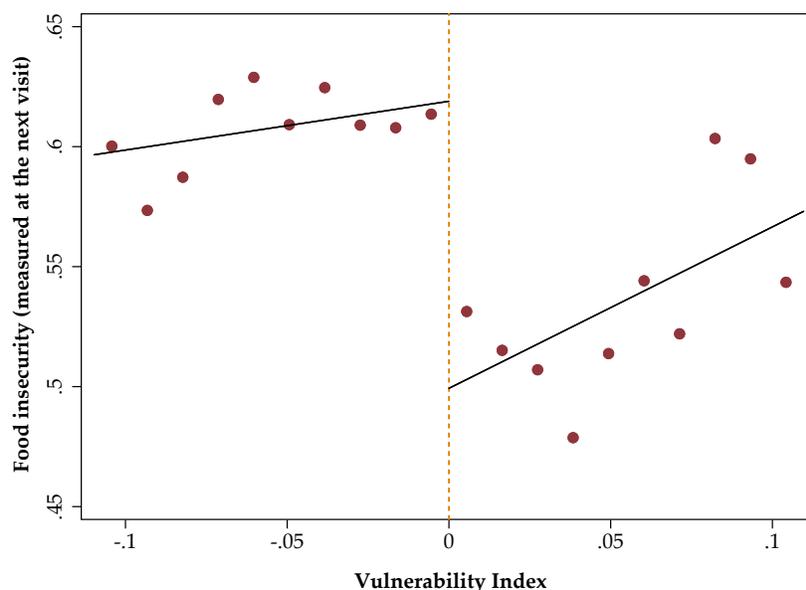
⁸⁰See Gundersen and Oliveira (2001). That question is: “Which of these statements best describes the food eaten in your household in the last month?” Four options are given: “enough of the kinds of food we want to eat”, “enough but not always the kinds of food we want to eat”, “sometimes not enough to eat”, or “often not enough to eat”. Households reporting the third or fourth answer are considered food insufficient.

⁸¹We get similar results by looking at the other two measures. These results can be found in the online Appendix D.3.

⁸²Table D.3 in the online Appendix confirms what is visually apparent: receiving UCT lowers the (self-reported) probability that minors in the household are food insecure by 31 pp.

⁸³We refer to the additional expenditures that the transfer induces, not the actual purchases made with the food card. There is a large literature that studies the marginal propensity to consume food as a result of a cash

Figure 1.13: *Impact of UCT on food insecurity*



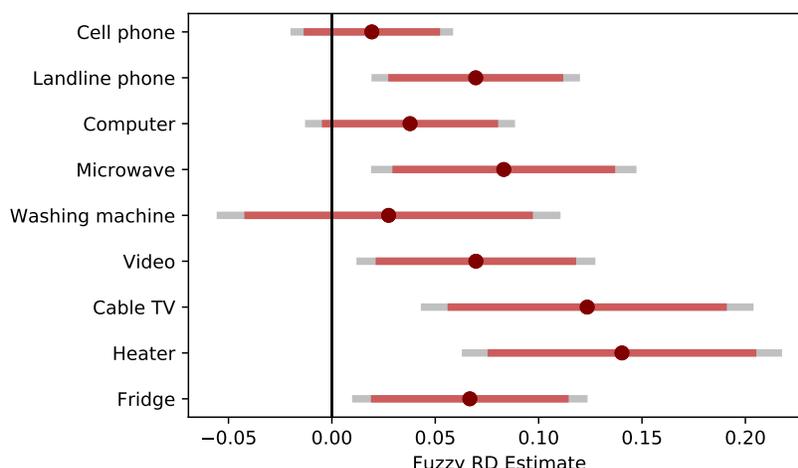
Notes: The variable plotted on the vertical axis is the percentage of individuals within a VI bin that live in a household where minors are food insecure (self-reported) in the next visit ($n + 1$). The horizontal axis shows the Vulnerability Index score computed on the basis of visit n (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are value of the outcome variable at visit n , female, age at visit $n + 1$ (and its square), year of visit $n + 1$ FE, month of visit $n + 1$ FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

Figure 1.14 shows the impact of UCT on durable goods in the household. We see that the UCT increases durable goods consumption in the household across the board (perhaps with the exception of cell phones, computers, and washing machines).⁸⁴

transfer (e.g., Bruich 2014) and finds that it is less than 1, hence it seems plausible that there are impacts on other consumption goods.

⁸⁴The null impact on cell phones is expected, given that practically all households in our sample have cell phones (91%).

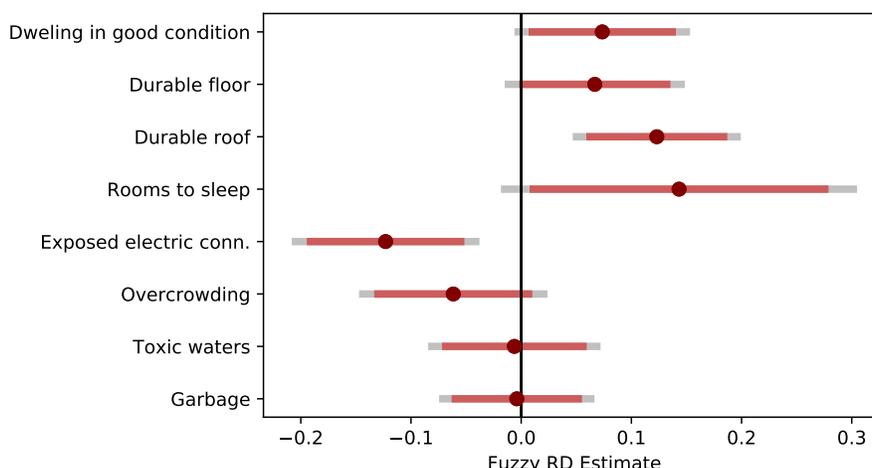
Figure 1.14: *Impact of UCT on durable goods consumption*



Notes: This figure shows fuzzy RD estimates of the impact of UCT receipt on the availability of durable goods in the household at the next visit ($n + 1$), and associated 90% and 95% confidence intervals. The Vulnerability Index score computed on the basis of visit n is used to instrument for beneficiary status at visit $n + 1$ (see “Measuring impacts on outcomes measured at a re-visit” in Section 1.4).

Figure 1.15 shows the impact of receipt of UCT on variables related to the “quality” of the dwelling and two neighborhood characteristics that we use as a sort of placebo check. We see that UCT seems to have a positive impact on the quality of floors, roofs, and the general condition of the dwelling, and a negative impact on undesirable housing conditions such as overcrowding or having an exposed electrical connection. We add two variables that measure not what happens in the household per se but instead in the block where the home is located. Specifically, the household is asked whether there are illegal dump sites or toxic waters in the neighborhood, and we find that the fact that a single household (quasi-randomly) receives the transfer has no impact on these variables.

Figure 1.15: Impact of UCT on housing conditions



Notes: This figure shows fuzzy RD estimates of the impact of UCT receipt on housing conditions at the next visit ($n + 1$), and associated 90% and 95% confidence intervals. The Vulnerability Index score computed on the basis of visit n is used to instrument for beneficiary status at visit $n + 1$ (see “Measuring impacts on outcomes measured at a re-visit” in Section 1.4). *Dwelling in good condition* is a dummy that takes value 1 if the home is in good condition or needs minor repairs, and 0 otherwise. The variable *Durable floor* codes 1 for permanently covered floors and 0 otherwise. *Durable roof* is a dummy variable that takes value 1 if the roof is made primarily of concrete or other durable materials, and 0 otherwise. The variable *Overcrowding* reports whether there are children that sleep in the same bed as adults. *Rooms to sleep* corresponds to the number of rooms used for sleeping. The variable *Exposed electrical connection* records whether there are electrical connections and wiring not embedded in the walls. *Toxic waters* is a dummy that measures the presence of wastewater or accumulation of contaminated water. The variable *Garbage* is a dummy that measures whether there is an accumulation of waste or dumps in the block where the home is located.

The results on material hardship not only indicate that the program positively affects the (material) well-being of beneficiary households but could also suggest a mechanism behind the impacts we found in regard to take-up of housing assistance programs. It would be hard to argue that the reason we see that UCT negatively impacts take-up of housing assistance is that enrollment in UCT increases the costs of enrolling in a housing assistance program (either costs of acquiring information, cost of applying, or stigma). One hypothesis is that receipt of UCT decreases the marginal benefit of enrolling in a housing assistance program. Most of these programs involve relocation, and we see that receipt of UCT improves housing conditions, which in theory should decrease the marginal benefit of moving to a different location.

1.6.2 Labor supply and human capital investment responses with survey data

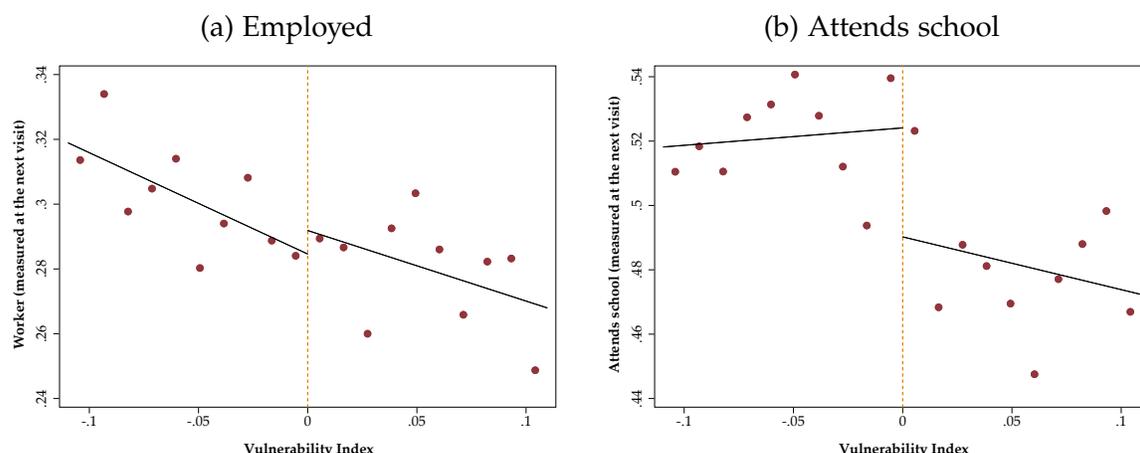
We exploit data on self-reported employment status at future visits to assess whether the impacts we find on the administrative data on formal employment are due to a drop in labor supply or are associated more with a shift to the informal sector. We also look at self-reported school attendance in order to cross-validate the results we find on enrollment in the administrative data.⁸⁵ Figure 1.16 shows the reduced-form impact of crossing the eligibility threshold on self-reported employment and school attendance, as measured in a future visit.

First, there does not seem to be an impact on labor supply. This evidence, together with the fact that formal labor supply responses are driven by beneficiaries whose UCT is indirectly means tested, suggests that the formal labor supply responses we observe in the administrative data are due to a decrease in the formalization of work and probably not associated with decreased employment. Second, we see a drop in self-reported school attendance (the RD coefficient is actually statistically significant at the 5% level). This brings supporting evidence to our school enrollment findings with administrative data.⁸⁶

⁸⁵In both cases, we consider the same sample definitions we used when looking at the impacts on administrative data. However, within these samples we have data (on the self-reported outcome) on only those that were visited at least once after the initial visit.

⁸⁶Although not entirely comparable, results in self-reported data seem to be even stronger than in the administrative data. One reason for this could be that enrollment may not necessarily adjust “instantly” when an individual quits school, and individuals could stay enrolled even if they do not regularly attend school. Self-reported attendance could presumably be more elastic.

Figure 1.16: Impact of UCT on self-reported employment status and school attendance



Notes: In panel *a* the variable plotted on the vertical axis is the percentage of individuals within a VI bin that reported being employed (as a private/public sector employee, a worker at a cooperative, a self-employed person, or an employer) at the next visit ($n + 1$). The horizontal axis shows the Vulnerability Index score computed on the basis of visit n (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. The sample corresponds to all individuals that were in the 18-38 age group three years after visit n . Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are value of the outcome variable at visit n , female, age at visit $n + 1$ (and its square), year of visit $n + 1$ FE, month of visit $n + 1$ FE, and Montevideo FE. The same applies to panel *b*, with two exceptions. First, the outcome variable is a dummy equal to 1 if an individual reports attending school. Second, the sample corresponds to minors that finished primary school (or are in the last year of primary school) and need at least three more years to finish secondary school. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

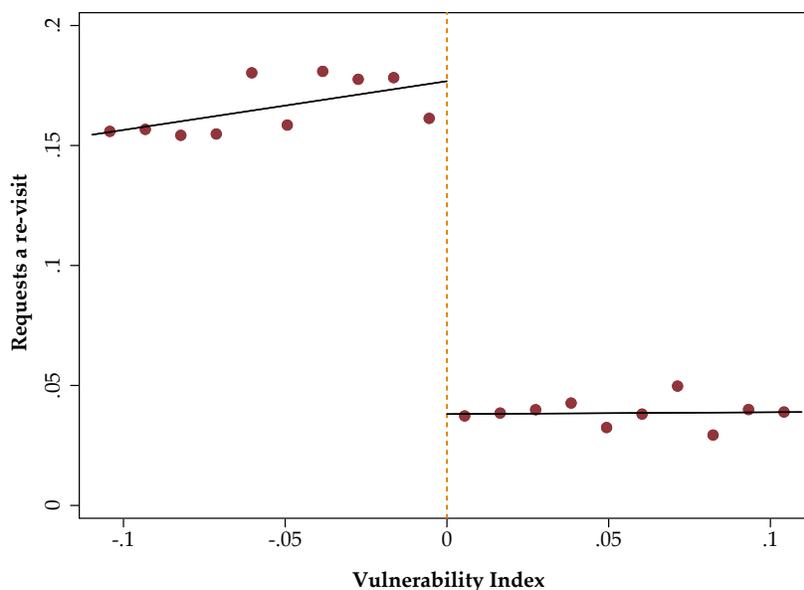
1.6.3 Strategic behaviors: complaints, selective deafness, and misreporting

Individuals seem to take several actions to sustain/gain eligibility. In this section we document three of these. One action is to call or show up at one of the MIDES offices across the country and request a visit to have the household situation assessed/re-assessed. Another is the decision whether to answer the door when a MIDES agent shows up to survey the household. A third one is to misreport information on the MIDES questionnaire.

Figure 1.17 shows the reduced-form impact of crossing the eligibility threshold on the probability that the household requested a re-visit within a year after the visit took place. We consider only visits that took place between January 2015 and December 2017, as data

on visit requests are available for only January 2015 to December 2018. There is a large drop just above the threshold, which suggests that gaining/keeping receipt of UCT is associated with a drop in the probability of requesting a re-visit.

Figure 1.17: *Impact of UCT on re-visit requests*



Notes: The variable plotted on the vertical axis is the percentage of individuals within a VI bin that live in households where a member requested a re-visit within a year after a visit. The horizontal axis shows the Vulnerability Index score computed on the basis of the visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. The sample corresponds to all individuals visited between January 2015 and December 2017. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

Table 1.7 confirms the previous result. On average, being a beneficiary after the visit induces a drop in the probability of requesting a re-visit by 24 percentage points. Moreover, the impact seems to be stronger for those initially enrolled in the program. Put differently, while losing the transfer increases the probability that a household requests a re-visit by 31 percentage points, gaining a transfer reduces the probability of requesting a re-visit by only 23 pp. These differences could be due to how people differentially respond to losses and gains, or it could also be a function of observable and unobservable differences across our

(pre-visit) beneficiary and non-beneficiary populations.

Table 1.7: *Impact of UCT on re-visit requests: fuzzy RD estimates*

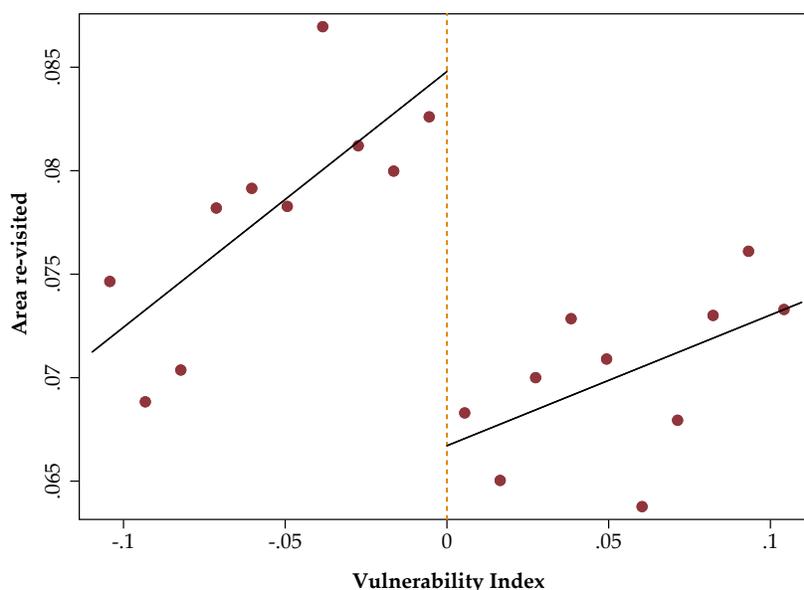
	Household requests a re-visit		
	Pooled (1)	Non-benef. (2)	Benef. (3)
UCT (1-12 mths after)	-0.243*** (0.013)	-0.226*** (0.014)	-0.306*** (0.029)
Observations	27030	21097	5933
Mean non - recipients	0.144	0.144	0.158
Bandwidth	0.1097	0.1097	0.1097
RD Polynomial	Linear	Linear	Linear
Controls	Yes	Yes	Yes
Kernel	Rectangular	Rectangular	Rectangular
SE	Cluster	Cluster	Cluster
First stage estimate	0.61	0.59	0.71
F-Stat (First Stage)	4038.0	3519.0	1070.0
P-val: (2) = (3)		0.021	

Notes: Standard errors clustered at the household-visit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bandwidth is set to 0.1097. The dependent variable is a dummy indicating whether anyone in the household requested a re-visit to MIDES during the year after the visit took place. Endogeneous regressor is the share of months (within 1-12 months after the visit) in which someone in the household received an UCT from MIDES. Asymmetric and linear RD polynomial considered. Control variables considered are: year FE, month FE, and dummy for Montevideo. Sample consists of households visited during 2015-2017. First stage estimate corresponds to the coefficient of the impact of crossing the UCT threshold on the endogeneous regressor (i.e., α_1 in equation 1.1). p-value of the t-test between the difference in the estimates of column 2 vs 3 are presented in the table.

Another behavior we observe is that individuals seem to answer the door in different ways, depending on their initial beneficiary status. We first introduce some indirect evidence of this behavior by looking at area re-visit rates. After that, we show more direct evidence by studying the history of visit attempts for targeted visits.

As explained in Section 1.2.3, whether MIDES decides to visit a household within an area visit should be orthogonal to its characteristics (controlling for the characteristics of its neighbors). Thus we would expect to find the same probability of an area re-visit for households that were visited and registered a VI score right below or above the threshold. However, Figure 1.18 shows that this is not the case. We notice that crossing the threshold induces a sharp drop in the probability of getting an area re-visit in the future.

Figure 1.18: *Impact of UCT on probability of area re-visit*



Notes: The variable plotted on the vertical axis is the percentage of individuals within a VI bin that are revisited in an area visit after a first visit. The horizontal axis shows the Vulnerability Index score computed on the basis of the (first) visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

Our reading of this result is that individuals either spend less time in the house or decide (in a higher proportion) not to open the door when a MIDES official shows up if they are enrolled in the UCT program.⁸⁷ The logic is that beneficiaries may have more to lose than to gain from a re-visit, so they may be resistant to getting re-tested (especially beneficiaries that “just” entered the UCT program in the first place, even if they are not fully aware that their entry was a borderline case). Nevertheless, we acknowledge that, at best, this is only indirect evidence that households are selectively deaf when a MIDES official knocks at the door.

More evidence of this behavior can be found by looking at targeted visits instead. For

⁸⁷If anything, our results on formal labor supply and education would suggest that people enrolled in UCT would be more likely to spend time in the house, so labor supply/education responses should not be driving these differences.

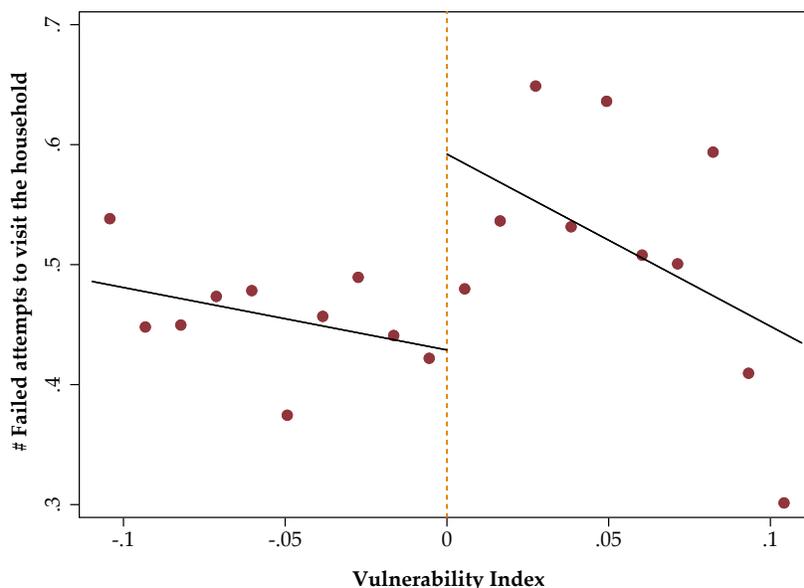
these visits, MIDES has the history of all the visit attempts before the household was successfully surveyed. We do not have the actual dates on which MIDES attempted to visit the household, but we have the number of attempts before each (successful) targeted visit occurred. We first look at the raw data, where we find that among those visited households that were (pre-visit) beneficiaries, the mean number of attempts before the (successful) visit takes place is 0.58. The corresponding number for (pre-visit) non-beneficiaries is 0.42. Thus in a quick comparison among (pre-visit) beneficiaries and non-beneficiaries, we do find that beneficiaries seem to be the harder population to survey.

Of course, the previous comparison does not show causality, and in fact given that we have data on attempts for successful visits but not for visits that never materialized, we cannot conclusively identify the effect. The closest we can get to identifying an effect is as follows: we take all individuals visited on a targeted visit with at least one (successful) previous visit. We run a regression with the number of unsuccessful attempts before the last successful visit to an individual as dependent variable, beneficiary status at the last successful visit as the endogenous regressor, and construct our usual instrument with the VI score from the first visit.

Figure 1.19 shows the mean number of attempts by MIDES before successfully surveying the household as a function of the VI score computed on the basis of the previous successful visit (Table D.8 in the online Appendix shows the corresponding fuzzy RD estimates). We see that being on the right side of the threshold (i.e., eligible for UCT) on a given visit seems to be associated with a higher number of visit attempts for the next visit (among the group of individuals that have at least two visits).

While the three pieces of evidence on selective deafness (i.e., differential area re-visit rates, differences in mean visit attempts in the raw data, and fuzzy RD estimates) do not constitute proof of this behavior, they are suggestive of it.

Figure 1.19: *Impact of UCT on visit attempts before a targeted visit is conducted*



Notes: The variable plotted on the vertical axis is the mean number of visit attempts by VI bin before the targeted visit was successfully completed. The horizontal axis shows the Vulnerability Index score computed on the basis of the previous visit (normalized to 0 at the UCT eligibility threshold). We consider 10 equal-sized bins on both sides of the cutoff. Both the dependent variable and the running variable are residualized on controls before plotting. The controls considered are female, visit attempts for the previous visit, year FE, month FE, and Montevideo FE. Those with Vulnerability Index scores to the *right* of the cutoff are eligible to receive UCT, and those to the *left* of the cutoff are not (with exceptions).

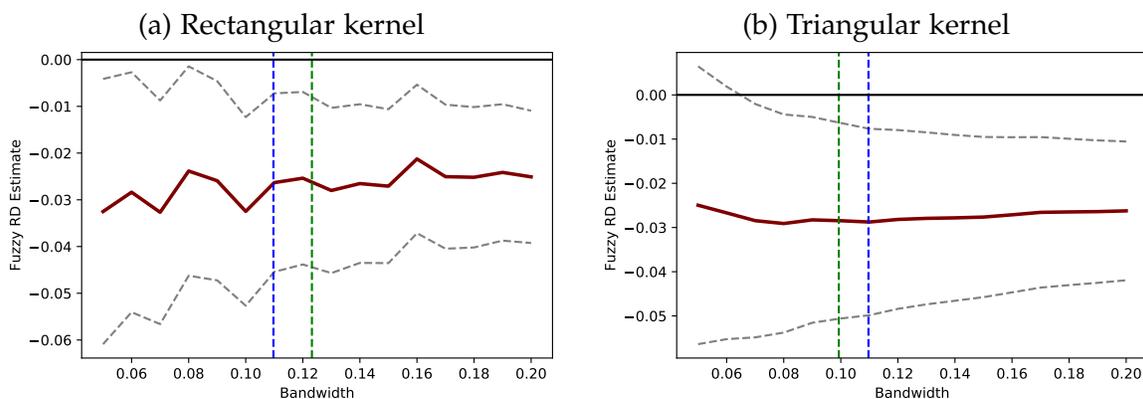
Finally, we study whether individuals misreport information in the interview when asked about their receipt of CCT. We generate a variable equal to 1 if a household is enrolled in CCT according to administrative records from the Social Security Administration but reports that it is not enrolled in CCT during the household visit conducted by MIDES. Figure D.7 in the online Appendix shows that receipt of UCT has no impact on this variable (misreporting) at future visits. However, the binscatters show positive values of 8% around the threshold. Thus while receipt of UCT seems to have no impact on future misreporting, the evidence suggests that households misreport their status in the CCT program.

1.7 Robustness of the main results

Bandwidth size and kernel specification. We investigate the robustness of our fuzzy RD estimates to the choice of bandwidth size and kernel specification (we consider both rectangular and triangular kernels). Assessing the robustness of our results to different bandwidth sizes is especially warranted in this setting, given that our bandwidth (0.1) was not selected with a data-driven procedure. A triangular kernel is a popular choice in RD studies (see Gelman and Imbens 2019), so it also seems natural to check that our results are robust to this kernel specification. We run regression (1.2) for each of our main outcomes with bandwidths in $\{0.05, 0.06, \dots, 0.2\}$. A vertical green line indicates the bandwidth chosen by a data-driven method (“MSERD”; see Calonico *et al.* 2014), and a blue line indicates our baseline bandwidth for reference.

Figure 1.20 shows the robustness of our main formal labor supply result (which corresponds to column 1 in Table 1.3). The estimates maintain similar magnitudes and mostly retain significance for all bandwidths considered, perhaps with the exception of very small bandwidths in panel *b*. However, these bandwidths are quite far away from the data-driven optimal bandwidth in this specification. In the online Appendix D.2 we show these robustness checks for the rest of the main outcomes considered in this study.

Figure 1.20: Robustness to bandwidth and kernel specification: formal labor supply



Notes: This figure shows fuzzy RD estimates ($\hat{\beta}_1$ from equation (1.2)) and 90% confidence intervals, with a dummy variable that takes value 1 if an individual is formally employed three years after the visit as dependent variable. The sample corresponds to individuals that were in the 18-38 age group three years after the visit. Standard errors are clustered at the household level. We show these RD estimates for different bandwidth values: 0.05, 0.06, ..., 0.2. The green vertical line represents the CCT optimal bandwidth (Calonico *et al.* 2014), and the blue line represents the bandwidth chosen in our baseline regressions. Panel *a* uses a rectangular kernel, and panel *b* uses a triangular kernel.

Falsification tests. Even though our results seem to be mostly robust to the choice of bandwidth and kernel specification, visual inspection of the binscatters raises the concern of whether we would have found similar impacts if we had just set the threshold at a different value. Because in several of our outcomes mean values around the threshold are noisy, and because we are dealing with a fuzzy RD (not a sharp RD), binscatters alone hardly tell a conclusive story in our setting. To check that changes in outcomes at the threshold are not a random feature of the data, we estimate the reduced-form impact of crossing a given threshold (i.e., the “gap” between the two regression lines that we see in binscatters) considering all possible thresholds that do not confound effects with other policy changes (i.e., CCT and double UCT).⁸⁸ In all cases, we consider the same bandwidth size (0.1).

⁸⁸We actually take all of these possible thresholds considering values that are 0.01 VI score points away from each other. The maximum value we take is the actual UCT threshold (above that value, the group to the right of the threshold would be confounded with other policy changes), and the minimum is -0.39 (we would confound effects with receipt of CCT if we chose a lower value).

To be clear, this is the regression specification:⁸⁹

$$Y_{i,h,t+36} = \beta_0 + \beta_1 \mathbb{1}[VI_{h,t} > 0] + f(VI_{h,t}) + \gamma X_{i,h,t} + \epsilon_{i,h,T} \quad (1.4)$$

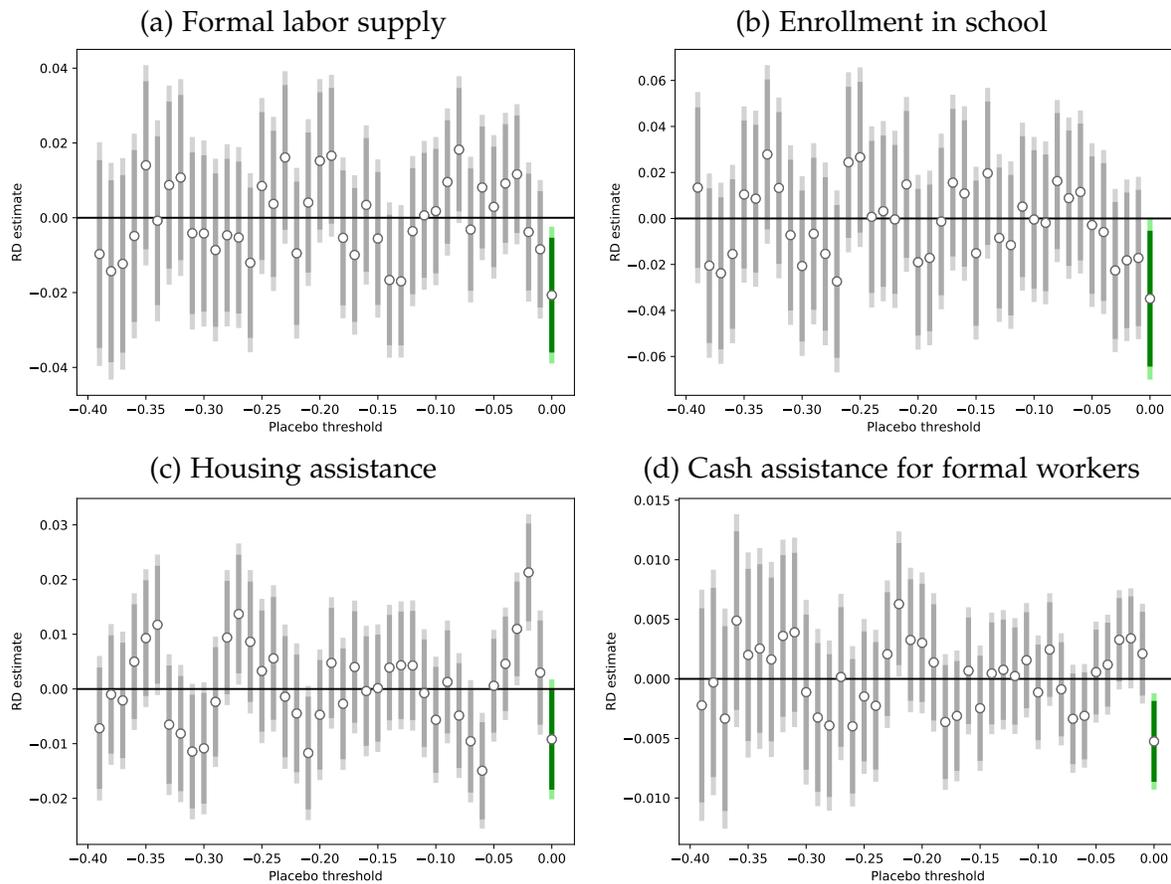
Figure 1.21 shows results of these falsification tests for the main outcomes studied.⁹⁰ We find that the most negative – and the only statistically significant – coefficient (out of the 40 falsification tests we perform for each variable) is the one that corresponds to the “real” (non-placebo) threshold. This holds true for all variables except housing assistance (panel *c*). With respect to this variable, and despite the fact that we see a negative coefficient at the real threshold, this also occurs in 4 of the other falsification tests. Moreover, there is a large positive coefficient at the placebo threshold -0.2 which suggests that the negative coefficient we observe at the real threshold could be directly related, at least in part, to what happens in a small neighborhood of our sample. Nevertheless, we see in Figure D.5 of the online Appendix that even with large bandwidths our result retains significance, so the upward jump right below the threshold does not seem to be the sole explanation of this result.

Overall, we read this evidence as strongly suggestive that our results on formal labor supply, education, and take-up of cash assistance for formal workers are not a random feature of the data, but rather are related to receipt of UCT. The evidence is less clear with respect to housing assistance, and that result should be considered with caution.

⁸⁹See Section 1.4 and the description of equation (1.2) for definitions of the variables.

⁹⁰For formal labor supply and school enrollment, we perform these placebos on the sample of pre-visit non-beneficiaries, which are the samples driving our results on those outcomes.

Figure 1.21: Falsification tests using placebo thresholds



Notes: This figure shows regression coefficients and associated 90% and 95% confidence intervals for the reduced-form impact of crossing placebo thresholds on a given outcome (i.e., the β_1 from equation (1.4)). We consider all placebo thresholds that do not confound effects of other policy changes (i.e., CCT and double UCT) and that are 0.01 VI score points away from each other. In panel *a* the outcome variable is a dummy equal to 1 if the individual is formally employed three years after the visit. In panel *b* the outcome variable is a dummy equal to 1 if the individual is enrolled in school three years after the visit. In panel *c* the outcome variable is a dummy equal to 1 if the individual is enrolled in a housing assistance program offered by the Ministry of Housing. In panel *d* the outcome variable is a dummy variable equal to 1 if the individual receives a cash assistance transfer for formal workers.

1.8 Conclusions

Cash assistance programs have been criticized for inducing behaviors that make individuals more reliant on these transfers and less on their own means to reach a basic standard of living. Most evidence on this topic measures the impact of receipt of cash assistance while

an individual is still enrolled in the program or at some point after he stops receiving the transfer. However, assessing the possibility of welfare traps with this type of evidence could be misleading.⁹¹ This is especially the case if individuals that enter a cash assistance program have no pre-specified exit date (as in many government-run cash assistance programs around the world) and exit is endogenous to their behaviors.

In this paper we assess the existence of welfare traps by providing separate causal estimates of the impact of entry into and forced exit from an unconditional cash transfer program based on both regression discontinuity and dynamic differences-in-differences designs. To identify these effects, we exploit the unique way in which the Uruguayan government decided to re-target its main unconditional cash transfer program through a proxy-means test. Through more than 250,000 household visits (during 2012 – 2018) covering roughly one-fifth of the population, the government estimated a socio-economic vulnerability score for each household and determined its eligibility for the program by comparing that score to a pre-specified value (i.e., if the score is higher than that pre-specified value, the household is eligible for the transfer). We use rich administrative longitudinal data and survey data to focus on three key indicators associated with family welfare dependency: labor supply and formalization of work, human capital investments for children, and take-up of other safety-net programs.

First, we find that labor supply of adults drops 3 pp three years after they enter the program, and that this effect is concentrated among adults under the age of 40. Nevertheless, this does not constitute a welfare trap: long-term recipients that are forced to exit the program increase their formal labor supply by 3 pp. We exploit the fact that the program is indirectly means tested initially for only a subset of the population (80%), and we find that our results are driven by this population. This suggests that the mechanism that drives people to reduce their formal labor supply is the means test per se and not an income effect. This also suggests that our results on formal labor supply probably operate mostly through

⁹¹In programs that have a pre-specified exit date (such as a “one-time” cash transfer), welfare traps can be assessed by looking at the persistence of the effects after the individual or household stops receiving the transfer (see section “Do Cash Transfers Create Dependency?” in Haushofer and Shapiro 2016).

the formality margin and not labor supply. We also do not find statistically significant drops in self-reported labor supply, which considers formal and informal work.

Second, we find that the transfer is associated with a lower (−3 pp) probability of being enrolled in school three years after the visit (for minors that should be enrolled in secondary school three years after the visit). However, because secondary school graduation rates in the Uruguayan context are low for low-income families, the school enrollment impacts do not translate into differential educational attainment rates. The result on enrollment is nevertheless surprising, given that we generally expect cash assistance to increase enrollment in school (or to have a null impact). Dahl and Gielen (2018) also finds that one form of cash assistance (disability insurance) negatively impacts human capital investments in children (albeit in a completely different context). The authors hypothesize that the driver of this result is that government assistance during youth impacts expectations of government assistance during adulthood, which is similar to the scarring effect in Malmendier and Nagel (2011). Testing this hypothesis in our context could be an interesting avenue for future research, given the results we find on school enrollment.

Third, we find that the program has a negative impact on enrollment in public housing programs (−1.2 pp, although this result should be taken with caution) and take-up of other types of public cash assistance (−0.8 pp). This is more suggestive of safety-net program substitution than of increased dependency on multiple programs. A back-of-the-envelope calculation suggests that for every \$100 that the government reduces today in UCT benefits, total public spending is reduced by only \$87 three years later, as a result of safety-net program substitution and a targeting mechanism that allows former beneficiaries to re-enter the program.⁹²

Overall, these results suggest that the program does not induce a welfare trap. While it is true that beneficiaries reduce their formal labor supply on entry, the fact that beneficiaries that are forced to exit welfare increase their formal labor supply suggests that the program does not decrease beneficiaries' *ability* to find employment. Although we find negative

⁹²See the online Appendix E for the details on this estimate.

impacts of the program on education enrollment, this does not create a welfare trap, as these impacts do not translate into different educational attainment rates. Take-up of other welfare programs can also be ruled out as a potential driver of dependency, as we find safety-net programs to be substitutes in this context.

Chapter 2

Meet the Oligarchs: Business

Legitimacy and Taxation at the Top¹

¹Co-authored with Rafael Di Tella and Juan Dubra.

Essay Abstract

We study the causal impact of trust in business elites and trust in government, on preferences for taxation at the top. Using a randomized online survey, we find that our two treatments are effective in changing trust in Major Companies and in Courts/Government. In contrast to previous work, we find that distrust causes an increase in desired taxes on the top 1%. For example, our treatment decreasing trust in business elites causes an increase in desired taxes on the top 1% of 2.4 percentage points (it closes 27% of the Democrat-Republican gap in tax preferences) when trust in government is low; a similar result is obtained for distrust in government.

2.1 Introduction

What determines taxation at the top? The traditional answer in economics, as embodied for example in Meltzer and Richard (1981), emphasizes the role of an unequal distribution of income. A recent paper by Kuziemko *et al.* (2015) explores this idea experimentally by successfully priming subjects on the M-Turk market for tasks with information regarding income inequality and then noting that their desire for taxing the top 1% does not change. In other words, people exhibit an unwillingness to “connect their concern about inequality with government action”. Kuziemko *et al.* (2015) argue that a possible reason is that their treatment emphasizes the severity of inequality, a social problem that the government has been unable to solve. If subjects distrust the government, they are unlikely to think that it will be able to redistribute efficiently. In a supplementary survey, they introduce a novel distrust treatment and find that their primes to distrust the government are effective in decreasing some of the measures of support for higher taxation at the top that they include.

The theory they proposed, connecting trust with taxation, is attractive because trust appears to be an important determinant of preferences for other forms of government intervention. For example, there is a large body of prior work connecting trust and regulation, although the correlation uncovered is negative (i.e., distrust is associated with a larger government; see, for example, Djankov *et al.* 2002). Such a negative correlation is far from natural, prompting Aghion *et al.* (2010) to call it “what is perhaps one of the central puzzles in research on political beliefs: why do people in countries with bad governments want more government intervention”.

To decide between these two conflicting views, we separate trust into two components: trust in businesspeople and trust in the government. We study their correlation with desired taxes with an experimental design borrowed from Kuziemko *et al.* (2015) which allows us to interpret it as causal. We separate trust into these two components because we think it is the way to answer the “central puzzle” raised by Aghion *et al.* (2010). Our conjecture is that countries with bad governments have allowed businesspeople to make money through corruption, so there is a bad opinion of the rich. Taxing them is one (small) way in which

justice can come about.

Our main hypothesis is that trust in business elites and in government have a direct impact on policy preferences because people care about non-monetary dimensions, such as fairness (as in Alesina and Angeletos 2005b; also see Cappelen *et al.* 2007 and Scheve and Stasavage 2012). This contrasts with models where voters only care about income (as in Meltzer and Richard 1981 or Benabou and Ok 2001) so that any impact of trust on policy preferences detected in these models is through its effect on income (e.g., because a more trustworthy government spends more efficiently).

Our experimental design introduces four treatments in which subjects are primed to have a positive vs negative view of businesspeople and of government officials, and then ask them for their desired tax rates. We first show subjects an image and text about a well known businessman, with a description of how they made their money. For the positive priming, we show an image of Bill Gates, and say how he and others “have revolutionized the technology industry. In several other areas, such as biotechnology, entertainment, medical devices, and high-end machinery, US business people have also been at the forefront of innovation”. Or for the negative priming we say “American business people have been involved in some major scandals over the years. Some of the most famous include Bernie Madoff (a Wall Street financier who was able to swindle investors for nearly 20 years) and Ken Lay (the former CEO of failed energy giant Enron who lobbied to obtain regulatory exemptions and government contracts). In several other areas, such as construction and medical supplies, there is also evidence of significant wrongdoing,” accompanied by a picture of Ken Lay. Priming for the case of honest or corrupt officials is similar, and we also randomize the order in which we show the primes (some people see first the business people priming, and some see the government official information). Each subject only sees one business and one government treatment, in a 2x2 design.

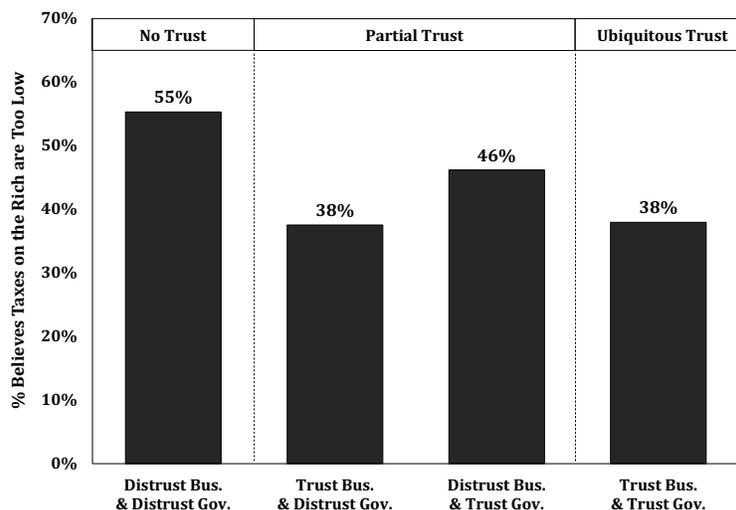
After asking a few related questions, we ask for the subject’s desired tax rate. We find that taxes desired by those primed to have negative views of businesspeople and government officials are larger than those primed to have positive views about at least one of them.

A simple interpretation is that taxes are used as a way to punish oligarchs (undeserving capitalists), when acts of corruption occur, which require corruption on the side of both businesspeople and government officials (our working paper presents a model with those features).

The results are consistent with informal accounts of businesspeople around the world where they are often heavily taxed and perceived to be self-serving, unpleasant, and even corrupt. In the United States (US), where historically they have often been lionized, the public has become increasingly uneasy about the power of business and of the so-called “top 1%”.² This dissatisfaction has been accompanied by a change in the media coverage they receive: during the 2001 dot-com bubble tech entrepreneurs were ubiquitous in the media, whereas after the 2008 financial meltdown bankers appeared to capture the public imagination. In other words, the motivation for our paper is to study if changes of this kind in the beliefs concerning the characteristics of businesspeople affect tax policy preferences. Interestingly, the correlation between these two measures of trust and desired taxes on the rich is particularly interesting in the US. Figure 2.1 uses data from the GSS to report the percentage of respondents that believe taxes on the rich are too low, conditional on their self-reported levels of trust in people running major companies and Congress.

²See Lindblom (1977) for a discussion on the privileged position of American businesses. In 2001, a Gallup poll found that US respondents were evenly split in terms of satisfaction with the size and influence of major corporations. Since 2003, most Americans have been dissatisfied. Dissatisfaction peaked at 67% in 2011. See “Majority of Americans Dissatisfied With Corporate Influence”, Gallup Economy, January 20, 2016.

Figure 2.1: Beliefs that taxes on the rich are too low and trust in business and government



Notes: Each bar represents the percentage of respondents that answered “Much too low” or “Too low” within a given group to the following question: “Generally, how would you describe taxes in America today. We mean all taxes together, including social security, income tax, sales tax, and all the rest. A. First, for those with high incomes, are taxes . . .”. The group “Distrust Bus. & Distrust Gov.” comprises those that answered “Hardly any” to the questions: “I am going to name some institutions in this country. As far as the people running these institutions are concerned, would you say you have a great deal of confidence, only some confidence, or hardly any confidence at all in them? B. Major companies; L. Congress”. Similarly, those in the group “Trust Bus. & Distrust Gov” answered “Hardly any” regarding “Major Companies” and “A great deal” regarding “Congress”. Those in the group “Trust Bus. & Distrust Gov” answered “A great deal” regarding “Major Companies” and “Hardly any” regarding “Congress”. Those in the group “Trust Bus. & Trust Gov” answered “A great deal” regarding “Major Companies” and “Congress”. Source of data is the General Social Survey 1987, 1996, 2006, 2008 and 2016.

These data suggest, broadly, one basic pattern: people who show high levels of trust want lower taxes on the rich than people with high levels of distrust. Of course, these are just correlations so the purpose of our paper is to collect new data that allows us to provide causal estimates. These paint a very similar picture to the one presented in Figure 2.1.

These correlations within the US are similar to other international comparisons. Consider the cases of France and the US. In the US business leaders are more trusted than in France (percentage of individuals that stated they believe CEOs are credible or very credible in the US is 38%, against 23% in France according to the Edelman Trust Barometer 2017). Taxation at the top is also higher in France (according to OECD, in 2016 the top marginal tax rate in

France was 54% and 46% in the US).

Borrowing the ideas used in political science to describe State legitimacy, an alternative label for our trust in business variable is “business legitimacy”, defined as the acceptance of the authority and privileges that emerge from the economic system. Indeed, a key feature of US capitalism is that some very rich people made their money in ways that are known and well regarded by the public. This is less common in other countries, where businesspeople might have power, but often enjoy less social status.³ We focus here on a subset of the rich: business people. Our approach can be interpreted as isolating the role of generalized trust, trust in government officials and trust in business elites. Thus, our paper draws on a large existing literature on trust (see, for example, Knack and Keefer 1995; La Porta *et al.* 1997; Guiso *et al.* 2004, 2008; Fehr and Schmidt 1999; Aghion *et al.* 2010; Algan and Cahuc 2010; Pinotti 2012).⁴ The hypothesis that the attributes of capitalists affect redistribution is also connected to prior work by Scheve and Stasavage (2012) on the dynamics of inheritance taxes and experimental work by Cappelen *et al.* (2018) on how fairness views come about where subjects distinguish between inequalities that are the result of luck and those that are the result of choices.⁵

We interpret lack of trust causing higher taxes as punishment, as in previous work on reciprocity (see, for example, Levine 1998, Rotemberg 2008, and Fong and Luttmer 2007). Of course, there may be other ways in which the public may want to limit the power of business (for example in the form of more regulation, as in Tella and Dubra 2014). More broadly, the idea that fairness can affect the political economy of taxation and government

³On the role and characteristics of a small group of “oligarchs” in controlling a substantial part of the Russian economy, see Guriev and Rachinsky (2005). See also, Akerlof and Romer (1993), La Porta *et al.* (2003), Morck *et al.* (2005) and Khanna and Yafeh (2007). The Economist magazine published a “crony-capitalism” index using data on billionaire wealth in sectors where there is a lot of interaction with State (see “Comparing crony capitalism around the world”, May 5th, 2016).

⁴The importance of trust is emphasized in Banfield (1958), Putnam (1993) and Fukuyama (1995). Arrow (1972) famously asserted “Virtually every commercial transaction has within itself an element of trust [and] much of the economic backwardness in the world can be explained by the lack of mutual confidence”. See also Durlaf and Fafchamps (2004) for a review on social capital and the ways in which it has been measured.

⁵See also Cappelen *et al.* (2018) for an interesting experiment showing that differences in tax preferences between Republicans and Democrats do not arise because of different beliefs about the efficiency consequences of taxes, or about how they affect behavior. Our results help explain the observed difference in preferences.

intervention is present in the theories and correlations discussed in Alesina and Angeletos (2005b,a) and Di Tella and MacCulloch (2009).

Section 2.2 describes the empirical strategy and the implementation of our survey and our M-Turk sample of approximately 9,000 Americans. Section 2.3 presents the main results, while Section 2.4 discusses a very similar, supplementary survey of 3,500 subjects run at a later date to provide information on the “first stage”. Section 2.5 concludes.

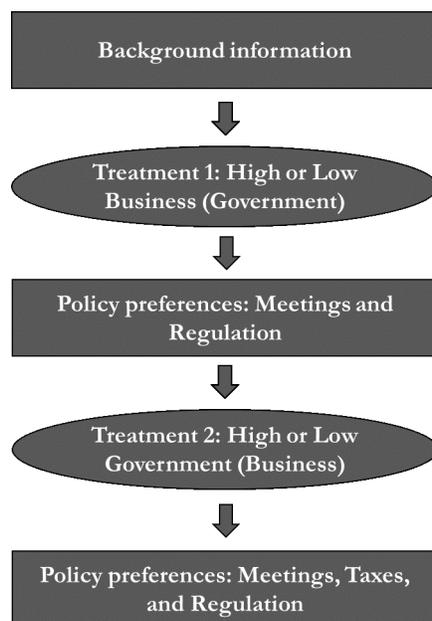
2.2 Empirical strategy and data

2.2.1 Empirical strategy

Figure 2.2 shows the flow of the survey, which is shown in full in Appendix A.2. We start by including some basic questions (age, gender, beliefs about the poor and rich, trust, etc.). Second, we treat our subjects with the first set of reports and questions. Third, we ask them a brief set of questions in order to separate the first treatment from the second. Fourth, we show them a second treatment (if the first treatment was a business treatment, the second treatment is a government treatment and vice versa). Fifth, we ask them a set of questions regarding their policy preferences over taxes.⁶ In the working paper version we also analyze the effect of trust on regulation, and on state capacity, but we do not present this analysis in the current version.

⁶The web link to the survey (this is the website where participants were redirected once they accepted the task in Amazon’s MTurk platform) is: https://hbs.qualtrics.com/SE/?SID=SV_ahE7rZtC1sCrInT

Figure 2.2: *Survey design*



Our empirical strategy is built on a priming design which builds on the assumption that the public's perception of the characteristics of businesspeople, on average, is formed over time through experiences and the accumulation of messages, including those from the media.⁷ We conjecture that business legitimacy stems from the efficiency, and also the honesty of a few, visible, businesspeople, perhaps because the public uses coarse categories (as in Mullainathan *et al.* 2008) or in terms of metaphors (as argued by Johnson and Lakoff 1980). Specifically, the survey includes two treatments (one for government and one for business) disguised as long questions. This is expected to induce subjects to reflect on positive (or negative) aspects of business elites and government in a 2x2 design. More importantly, it is expected to provide an indirect stimulus for related memories (conceptual priming).⁸ We focus on treating subjects with negative views about people (i.e., business

⁷We are agnostic as to the duration of these effects as our empirical strategy only requires a short duration, and for the theory to coincide with cross country evidence we only need a more positive message in the US than in countries with more regulation/taxation. See Hamilton and Zeckhauser (2004) for a study on media coverage of CEOs.

⁸See Berdejó and Chen (2012), for a review. They note "Priming is a cognitive process, in which media information increases temporarily (i.e., primes) the accessibility of knowledge units in the memory of an

leaders or government officials), so we mimic as much as possible the correlations found in the GSS data (Figure 2.1) where trust is expressed as confidence in the *people* running major companies or government. Previous work has also used priming in studies of social preferences (Fong and Luttmer 2007; Chen and Li 2009; Klor and Shayo 2010; Day and Fiske 2017; Ordabayeva and Fernandes 2017) and ideology in the field (Berdejó and Chen 2012).

Our positive business treatment involves priming subjects during the administration of a standard survey with a short statement about the efficiency-honesty of US business leaders, together with a photograph of a well-known business leader (Bill Gates) and some questions about the possible reasons for such high levels of efficiency and honesty in the US. Specifically, the first part introduced the question by explaining, "*American business people are amongst the most successful in the world. Some of the most famous include Bill Gates (founder and CEO of Microsoft) and Steve Jobs, (founder of Apple, NeXT and Pixar), who have revolutionized the technology industry. In several other areas, such as biotechnology, entertainment, medical devices, and high-end machinery, US business people have also been at the forefront of innovation.*" This was followed by the second part: a photograph of Bill Gates with the caption "*Bill Gates, CEO and founder of Microsoft, a company that revolutionized the personal computer industry*". The third part simply asked: "*Why do you think American business people have been so successful?, a. It is due to the system: business people in the US are encouraged to work hard and can gain money and prestige by creating truly good products. b. It is a combination of the system interacting with exceptional individuals, amplified by the availability of capital that allows the successful to expand their business. Or c. It is due to the individuals: there are remarkable business people in the US, who are exceptionally creative and naturally hard working.* Similarly, in the negative version of this treatment subjects read a statement about the involvement of some well-known business leaders in sophisticated economic crimes (more precisely, instances of business leaders capturing the government) accompanied with a photograph of Ken Lay and a set of questions (for a full description of the survey see the Appendix A.2). We call

individual, thus making it more likely that these knowledge units are used in the reception, interpretation and judgment of subsequent external information."

these business treatments the *High Business* and the *Low Business* treatment.

Our government treatment involves another layer of priming during the administration of the survey, this time with a statement, photograph and questions about the possible reasons for the prevalence of honest-efficient government officials in the US (or corrupt government officials in the case of the negative treatment). We call these the *High Government* and the *Low Government* treatment. To rule out order effects, in some cases we first showed individuals the business treatment and then the government treatment, and in some cases the other way around.⁹ There is also an untreated group, which consists of individuals that were presented with no treatment at all (i.e., their survey did not include reports and questions about businesspeople or government officials). We also performed two additional treatments, showing businesspeople or government officials that had been caught for corruption, we call them treatments with punishment, but we do not focus on these treatments.¹⁰

There are three obvious limitations of these treatments. First, we highlight both the honesty and efficiency of business leaders (government officials) in the *High Business* (*High Government*) treatment. In section 4.1 we observe how this increases trust in the business elites (in the government) but cannot assess whether this increase in trust is due to the honesty or the efficiency channel. Second, individuals may distrust business elites and the government for a multiplicity of reasons, so our treatments should be read as emphasizing only a one of many possible dimensions of what contributes to place trust in others. Third, a bribe exchange always involves both the government and business, so it is hard to have “pure” treatments (e.g., involving corrupt business and honest government). Indeed, bribes can occur due to business leaders capturing the government (as emphasized in the *Low Business* treatment), or government officials extorting payments from businesses (as in the

⁹There are no significant “order” effects (results available upon request) so, in this version of the paper we present the pooled treatments. This means that an individual treated first with *High Government* and then with *Low Business* is considered in the same category as an individual treated with *Low Business* first and with *High Government* later.

¹⁰The analysis of these treatments is available upon request.

Low Government treatment). The results in the supplementary survey in section 4.1 are encouraging in this regard because they demonstrate that the business treatments mainly affect trust in business elites, but not in government or other groups/institutions (similarly, government treatments do not affect the view of the business community).

The main regression specification is:

$$\text{Policy preference}_i = \beta_1(\text{High Bus \& High Gov})_i + \beta_2(\text{High Bus \& Low Gov})_i + \beta_3(\text{Low Bus \& High Gov})_i + \gamma X_i + \epsilon_i \quad (2.1)$$

$(\text{High Bus \& High Gov})_i$ is a dummy variable equal to 1 if individual i was treated with *High Business* and *High Government* (0 otherwise), and analogous definitions apply for the other treatments. The omitted group is the group treated with *Low Business* and *Low Government*. We follow Kuziemko *et al.* (2015) and estimate reduced form equations, which is the most conservative approach. While we provide evidence that the treatments only affect the most relevant dimensions of trust, it is impossible to rule out other channels through which the treatment might affect the outcomes of interest (e.g., opinions about inequality).

For the case of trust in business elites (analogous for the case of trust in government), we are interested in the following regression coefficients and linear combination of regression coefficients:

- $\beta_1 - \beta_3$: effect of trust in business elites conditional on *High Government*.
- β_2 : effect of trust in business elites conditional on *Low Government*.
- $\frac{(\beta_1 - \beta_3 + \beta_2)}{2}$: effect of trust in business elites.

The results are similar if we include the untreated group as the omitted category.¹¹ Those in the untreated group were exposed to a somewhat shorter survey, so the comparison would not be valid, as the treatment and untreated group have differences beyond the

¹¹Results are available upon request.

information contained in the surveys. An alternative that we considered was to use a 'placebo' treatment (showing a report about something totally unrelated) but this introduced other concerns. Our approach, studying linear combinations, avoids these problems. It also has other advantages. For example, if we show individuals a negative report about business elites in the US (analogous for government), we cannot be sure that we are actually priming individuals to distrust the business elites as this would depend on their prior. If their prior is that businesspeople are even more corrupt than what was stated in the report, then our treatment may actually prime people to trust the business elites more. What we can be sure is that those primed with *High Business* received a more favorable priming on business elite's efficiency-honesty than those primed with *Low Business*, and so we focus on this comparison. Extreme negative priming could overcome this problem, but as previous work has emphasized (e.g., Day and Fiske 2017), this is not feasible as overly biased report may lead individuals to distrust the survey, feel they are being manipulated, etc. A third possible reason to present linear combinations is that, if we compare our treatments directly with the untreated group we need more power to reach definitive conclusions, than if we compare 'opposite' treatments (i.e., *High Business* vs *Low Business*). This is particularly relevant as we are avoiding overly biased reports that may appear manipulative; or very long reports, with many windows, that could potentially give more power (see Srull and Wyer 1979 that finds a larger effect when the total number of primes is increased) but can create an imbalance in the amount of time and effort that respondents in the treatment and untreated group devote to the survey.

The untreated group is still useful to assess the representativeness of the sample in terms of the questions asked after the two treatments (such as the WVS questions we ask after both treatments). It is also useful as it allows us to express the magnitudes of the impacts of the treatments in terms of the mean of the respective variable for this untreated group (instead of in terms of the mean of either the *High* or *Low* treated group).

2.2.2 Survey implementation

The survey was implemented through Amazon's Mechanical Turk, an Internet-based market for tasks which has well known limitations and advantages (see the review in Horton *et al.* 2011). It has been used to study social preferences (Weinzierl 2014 and Saez and Stantcheva 2016) and several questions in economics, including the effect of peers' wages on job satisfaction (Card *et al.* 2012), or that of inequality on preferences for redistribution (Kuziemko *et al.* 2015), or of reference points on preferences for redistribution (Charité *et al.* 2015).

In our case, MTurk was used to attract subjects by offering a small reward (1 dollar) for taking a brief survey (less than 10 minutes) to "help us learn more about the relationship between politics and government in America". We explained participation was anonymous, we allowed individuals up to 50 minutes to complete the survey and were paid automatically after 8 hours of completing the survey. We followed several steps to ensure high-quality responses. Besides restricting the sample in ways that will be explained below, we recruited only individuals with a Human Intelligence Task approval rate equal to or higher than 80% and we set visibility to "Private" so that only workers that meet this qualification can preview our survey. To check perceptions of bias in our survey, we coded the comments that respondents made at the end of the survey and found that only a small fraction (0.5%) stated that the survey was biased. To discourage respondents from skipping some questions, a pop-up window appeared whenever an individual intended to go to the next window before answering all the questions in the current window. The pop-up indicated the number of questions that were not answered and whether the respondent wanted to continue without answering all the questions. We conducted our main survey on a single wave in late November 2015: 9,217 individuals took it.

We collected data on the time spent by subjects on each of the windows we presented during the survey. Several subjects took far less time than the minimum amount of time required to read the questions. To get potentially meaningful answers we restrict the sample in two ways, apart from 29 individuals who had corrupted data and were therefore not

considered in our sample. First, we consider only individuals that took at least 3 minutes to complete the survey (not considering the time spent in the treatment window; there is also a very short unrelated experiment that was presented after all our survey was completed, which we call the candy experiment. It was not considered when restricting the sample.). Second, among these individuals, we consider only those who spent at least 3 seconds looking at each of the treatment windows (this last condition does not apply to individuals assigned to the untreated group). The total number of observations after applying these two filters is 7,674. We included two treatments where punishment was made salient (of bad businesspeople or officials), and without these observations, our resulting sample includes 5,974 subjects. The analysis of the punishment treatments (available upon request) show no statistically significant differences between these treatments and those where punishment was not made salient. The mean number of minutes spent answering the survey is 7.3 minutes.

Table 2.1: Summary statistics

	All (our sample)	Democrats our sample	Republicans our sample	Kuziemko, et al 2015	WVS 6 th Wave	ACS 2015
<i>Demographics</i>						
Male	43.8%	43.6%	44.7%	42.8%	48.4%	48.6%
Age	34.9	33.8	37.3	35.4	46.5	47.1
White	80.5%	74.8%	93.0%	77.8%	69.8%	74.8%
Black	9.2%	12.5%	1.8%	7.6%	10.4%	12.2%
Hispanic	6.6%	7.7%	4.3%	4.4%	13.4%	15.5%
Asian	6.8%	8.6%	2.9%	7.6%	-	6.2%
Other race	2.6%	2.9%	2.0%	2.6%	-	2.8%
Postgraduate degree	13.3%	14.2%	11.4%	12.6%	11.5%	10.2%
Only college degree	47.4%	47.7%	47.1%	40.7%	24.8%	25.7%
No college degree	39.3%	38.1%	41.5%	46.7%	63.7%	64.1%
Full-time employee	46.7%	47.1%	45.7%	33.2%	42.7%	43.9%
Part-time employee	12.8%	12.6%	13.3%	13.3%	8.8%	16.7%
Self-employed	12.4%	12.1%	12.7%	10.5%	5.1%	7.2%
Unemployed	8.0%	8.6%	6.7%	12.4%	9.4%	3.9%
Student	8.7%	10.0%	5.7%	15.8%	4.7%	3.8%
Not in Labor Force	11.5%	9.5%	15.9%	14.8%	23.8%	31.7%
<i>Political preferences and beliefs</i>						
Trust	4.9	5.0	4.8	-	-	-
Poor didn't make an effort	22.8%	14.7%	40.7%	-	-	-
Rich made an effort	36.9%	27.8%	57.2%	-	-	-
Obama	68.8%	100%	0%	67.5%	-	-
<i>Outcome variables after second treatment (for untreated group)</i>						
Meetings Good	3.8	3.6	4.2	-	-	-
Competiton_Bad	2.6	2.9	2.1	-	2.7	-
More_Gov_Resp	3.9	4.7	2.2	-	4.2	-

Continued on next page

Table 2.1: (Continued) Summary statistics

	All (our sample)	Democrats our sample	Republicans our sample	Kuziemko, et al 2015	WVS 6 th Wave	ACS 2015
Market_Bad	4.1	4.0	4.2	-	-	-
Discretion	31.8%	37.5%	19.5%	-	-	-
No_Discretion_Reg	36.4%	44.1%	21.0%	-	-	-
Tax_1_percent	34.8	37.6	29.0	30.2	-	-
Tax_next9_percent	26.5	28.4	22.9	-	-	-
Tax_next40_percent	17.8	18.2	17.1	-	-	-
Tax_bottom50_percent	9.3	8.9	10.3	-	-	-
High_Fraud	31.9%	33.9%	27.7%	-	-	-
Observations	5974	4085	1856	3746	2138	2,490,616

Notes: Columns 1-3 consider people that spent at least three minutes in the main survey (excluding the treatment windows) and at least three seconds in every treatment (when applicable). The punishment treatments are excluded. Column 4: We consider respondents that took any of the omnibus treatment surveys of Kuziemko, et al. (2015); participants could only choose one ethnicity in this study; variable *Obama* is a variable that equals 1 if individual answered Obama when asked “Who did you support in the presidential election in 2008? If you were not able to vote, just choose the person you wanted to win the election at that time”; for the question on taxes we consider the untreated group of the omnibus treatment surveys (sample size is 1976). Column 5: data source is the 6th wave of the World Value Survey US sample; individuals with employment “Other” are omitted; variables *Competition_Bad* and *More_Gov_Resp* are the same questions used in our study (only that in the WVS answers range from 1-10 so we rescaled these answers to a 0-10 scale). Column 6: data source is the American Community Survey 2015; we consider individuals with 18 years or older.

Our distribution of subjects by state is similar to that of the American Consumer Survey 2015 (the comparison is in the working paper). Note, however, that our survey uses voluntary participants: those who participate may be different from those who do not (even if identical in terms of observables).

Table 2.2 presents the data summarized across treatments. It suggests that, at least with respect to observables, the data are balanced across treatments suggesting a successful randomization.

Table 2.2: Randomization

Variables	Treatment group				
	Untreated Group	High Bus & High Gov	High Bus & Low Gov	Low Bus & High Gov	Low Bus & Low Gov
<i>Demographics</i>					
Male	44.4%	47.3%	43.0%	43.5%	42.0%
	-	0.24	0.51	0.66	0.31
Age	34.4	35.0	34.9	35.0	35.0
	-	0.23	0.24	0.21	0.26
White	80.5%	79.4%	81.0%	80.0%	81.6%
	-	0.60	0.75	0.80	0.56
Black	9.0%	8.6%	10.0%	8.6%	9.5%
	-	0.73	0.43	0.69	0.75
Hispanic	6.5%	7.6%	5.7%	7.0%	6.5%
	-	0.37	0.45	0.60	0.99
Asian	6.8%	7.1%	6.6%	6.8%	6.9%
	-	0.81	0.89	0.99	0.91
Other race	2.8%	2.8%	2.8%	2.3%	2.9%
	-	0.95	0.99	0.44	0.93
Postgraduate degree	15.4%	13.6%	11.9%**	13.4%	13.4%
	-	0.29	0.02	0.18	0.24
Only college degree	45.8%	47.0%	49.4%*	47.1%	45.8%
	-	0.63	0.09	0.56	1.00
No college degree	38.7%	39.4%	38.6%	39.5%	40.7%
	-	0.79	0.96	0.71	0.40
Full-time employee	46.9%	48.3%	47.5%	45.6%	45.2%
	-	0.57	0.79	0.52	0.49

Continued on next page

Table 2.2: (Continued) Randomization

Variables	Treatment group				
	Untreated Group	High Bus & High Gov	High Bus & Low Gov	Low Bus & High Gov	Low Bus & Low Gov
Part-time employee	11.9%	12.2%	11.9%	13.9%	13.4%
	-	0.86	1.00	0.16	0.36
Self-employed	10.6%	12.2%	12.7%	13.2%*	12.0%
	-	0.30	0.12	0.05	0.37
Unemployed	9.5%	10.5%	7.2%*	6.4%***	8.9%
	-	0.53	0.06	0.01	0.66
Student	9.4%	6.9%	9.2%	8.9%	8.1%
	-	0.06*	0.84	0.69	0.34
Not in Labor Force	11.6%	9.9%	11.5%	12.0%	12.4%
	-	0.26	0.94	0.77	0.63
<i>Political preferences and beliefs</i>					
Trust	4.8	5.0	4.9	4.9	4.9
	-	0.12	0.22	0.18	0.19
Poor didn't make an effort	22.8%	24.6%	22.6%	22.4%	22.2%
	-	0.40	0.91	0.83	0.73
Rich made an effort	38.6%	38.2%	36.6%	36.7%	35.0%
	-	0.86	0.32	0.34	0.12
Obama	67.5%	70.4%	67.4%	69.1%	70.5%
	-	0.20	0.94	0.44	0.19
Observations (regression sample)	829	851	1725	1727	842
Observations (unrestricted)	1014	997	2041	2036	1001
Non-response	0.9%	0.7%	0.4%	0.4%	0.7%

Notes: Mean value of the variable is presented in the first row; p-value of the mean differences t-test (with respect to the untreated group) is presented in the second row. *, **, *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. All these statistics are computed using the regression sample. Regression sample corresponds to the sample of people that spent at least three minutes in the survey (not considering the candy experiment and time spent in the treatment windows) and at least three seconds in every treatment (if applicable). Unrestricted sample corresponds to all the individuals (within treatments) that took the survey. Individuals primed with punishment treatments are not included.

Finally, even though observables across treatments are balanced and the attrition of subjects who started the survey but did not finish it was only 2%, we perform two additional

checks to confirm that there is no differential attrition across treatments. First, we run a regression using the full sample of treated subjects, where the dependent variable is whether the individual completed the survey, and we check whether completing it is correlated with being in any treatment group. Column 1 in Table 2.3 shows the results (*Low Business & Low Government* is the omitted group in the regression) where we see that no treatment arm has a differential attrition rate. Another concern is that even if there is no differential attrition in the unrestricted sample, the restrictions we impose on the sample induce different attrition rates (and presumably different distributions of unobservable characteristics). In column 2 we run the same regression as in column 1 but with a dummy variable equal to 1 if the individual enters our restricted sample as our dependent variable: assignment to treatment is uncorrelated with whether an individual entered our restricted sample.

Table 2.3: *Completion of survey by treatment arm*

<i>Regression output</i>	Completed Survey (1)	In Restricted Sample (2)
(β_1) High Business & High Government	-0.000 (0.004)	0.012 (0.016)
(β_2) High Business & Low Government	0.003 (0.003)	0.004 (0.014)
(β_3) Low Business & High Government	0.003 (0.003)	0.007 (0.014)
Observations	6075	6075

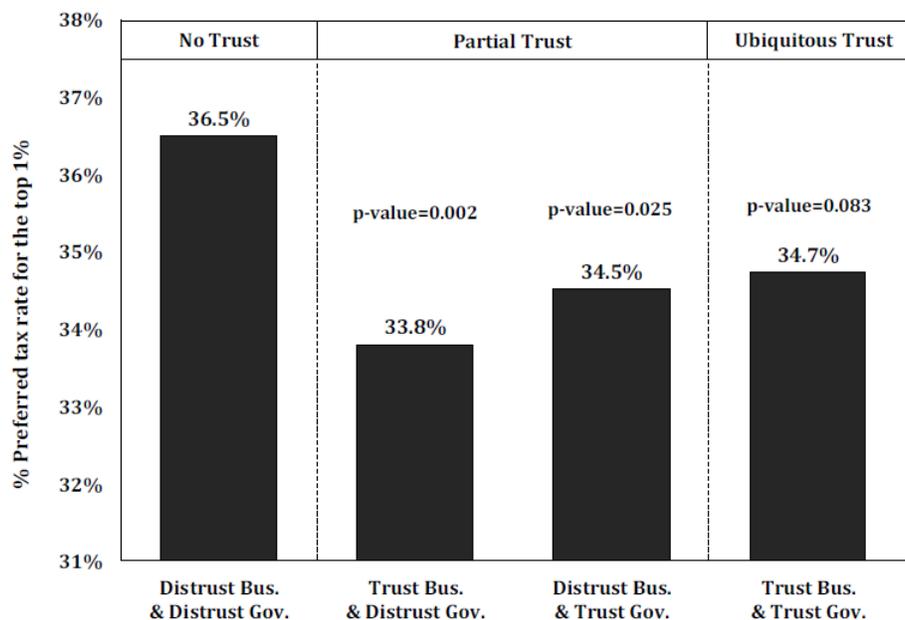
Notes: We show regressions estimates with robust standard errors in parenthesis; a constant term (not shown) is included in every regression. Regarding the treatments, the *Low Business & Low Government* is the omitted group. Dependent variables are: (1) Dummy equal to 1 if the individual got to the page of the last question of the survey. (2) Dummy equal to 1 if the individual spent at least three minutes in the survey (not considering the candy experiment and time spent in the treatment windows) and at least three seconds in every treatment. Respondents assigned to treatments with punishment and the untreated group were not included.

2.3 Results

In Figure 2.1 we used data from the GSS to assess correlations between levels of trust and preferred taxes on the rich. In Figure 2.3 we show the analogue to that figure, using our

experimental results instead.

Figure 2.3: Preferred tax rate for the top 1%



Notes: Each bar represents the mean preferred tax rate on the top 1% for respondents within a given treatment group. Sample of people that spent at least three minutes in the survey (not considering the candy experiment and time spent in the treatment windows) and at least 3s in each treatment. p-values correspond to the difference between the treatment coefficient *Low Business & Low Government* and another treatment group, in a regression where the dependent variable is the preferred tax rate on the top 1% and the only covariates are the four treatment groups (no constant included).

We plot the mean preferred tax rate on the top 1% for each of the four treatment groups: *Low Business & Low Government*, *High Business & Low Government*, *Low Business & High Government*, *High Business & High Government*. The results experimentally replicate the correlation patterns seen in GSS data. In summary, individuals treated with *Low Business & Low Government* want more taxes on the top 1% than any other treatment group, while the rest of the treated groups show preferred top tax rates which are indistinguishable from each other (i.e., the differences in mean preferred tax rates between any two treatment groups, not including *Low Business & Low Government*, are not statistically significant).

In Table 2.4, we study these results in more detail and the impact of the treatments on preferences for taxation over all the income distribution.

Table 2.4: Preferences for taxation

	Tax rate top 1%	Tax rate next 9%	Tax rate next 40%	Tax rate bottom 50%
<i>Panel A: Regression output</i>				
	(1)	(2)	(3)	(4)
Treatments				
(β_1) High Business & High Government	-1.819* (1.009)	-1.644** (0.809)	-0.494 (0.608)	0.257 (0.491)
(β_2) High Business & Low Government	-2.402*** (0.860)	-2.062*** (0.705)	-0.825 (0.522)	-0.441 (0.392)
(β_3) Low Business & High Government	-1.807** (0.864)	-1.540** (0.709)	-0.522 (0.523)	0.109 (0.394)
<i>Other covariates</i>				
Poor didn't make an effort	-4.111*** (0.678)	-3.171*** (0.534)	-0.922** (0.436)	1.092*** (0.386)
Rich made an effort	-7.140*** (0.609)	-4.592*** (0.484)	-1.392*** (0.384)	0.628* (0.328)
Obama	8.511*** (0.640)	5.475*** (0.514)	2.503*** (0.405)	-1.163*** (0.339)
Trust	-0.068 (0.122)	0.046 (0.100)	0.060 (0.081)	0.145** (0.067)
Observations	5097	5097	5095	5086
Untreated group mean	34.693	26.499	17.751	9.332
<i>Panel B: Hypothesis testing over the coefficients</i>				
Effect of Trust in Business Elites				
High Bus – Low Bus	-1.206** [0.0453]	-1.083*** [0.0248]	-0.398 [0.2752]	-0.146 [0.6202]
High Bus – Low Bus High Gov	-0.011 [0.9893]	-0.104 [0.8742]	0.028 [0.9564]	0.148 [0.7365]
High Bus – Low Bus Low Gov	-2.402*** [0.0052]	-2.062*** [0.0034]	-0.825 [0.1138]	-0.441 [0.2608]
Scaled effect	-0.271	-0.354	-0.648	0.318
Effect of Trust in Government				
High Gov – Low Gov	-0.612 [0.3106]	-0.560 [0.2461]	-0.095 [0.7938]	0.403 [0.1720]
High Gov – Low Gov High Bus	0.583 [0.4888]	0.418 [0.5242]	0.331 [0.5177]	0.698 [0.1122]

Continued on next page

Table 2.4: (Continued) Preferences for taxation

	Tax rate top 1%	Tax rate next 9%	Tax rate next 40%	Tax rate bottom 50%
High Gov – Low Gov Low Bus	-1.807** [0.0366]	-1.540** [0.0300]	-0.522 [0.3183]	0.109 [0.7826]
Scaled effect	-0.204	-0.264	-0.410	-0.078

Notes: Panel A presents regressions estimates with robust standard errors in parenthesis; includes demographic controls (gender, age, race, education, and type of employment), plus political variables and pre-treatment beliefs (include relative support for Obama in previous election, attitudes towards the rich and the poor, and general level of trust); a constant term (not shown) is included in every regression. Regarding the treatments, the Low Business & Low Government is the omitted group. *, **, *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. Dependent variables are the following: (1) Preferred tax rate for the top 1%. (2) Preferred tax rate for the next top 9% (1% of households earn more than them, but 90% earn less). (3) Preferred tax rate for the next top 40% (10% of households earn more than them, but 50% earn less). (4) Preferred tax rate for the bottom 50% in the income distribution (poorest). In Panel B we present linear combinations of certain treatment coefficients and p-values (in brackets) for the test of whether these linear combinations are equal to 0. High Bus – Low Bus | High Gov is the difference between the treatment coefficient High Business & High Government and Low Business & High Government; High Bus – Low Bus is the weighted average of High Bus – Low Bus | High Gov and High Bus – Low Bus | Low Gov (analogous for other treatment groups). “Scaled effect” is the result of dividing certain effect by the difference between the untreated group mean of the dependent variable for democrats and republicans. We considered the sample of people that spent at least three minutes in the survey (not considering the candy experiment and time spent in the treatment windows) and at least three seconds in every treatment. Respondents assigned to treatments with punishment were not included.

In column 1, we regress the preferred tax rate on the top 1% on the treatments (and a constant term), a set of demographic controls, political preferences and pre-treatments beliefs. Results are robust to dropping the post treatment variable *Obama* (relative support for Obama in 2012). We first note some interesting correlations. Beliefs about the rich and the poor are strongly correlated with preferred taxes at the top: a more favorable view about the rich, and a less favorable view about the poor correlate with a lower preferred tax rate on the top 1%. Believing that poverty is a result of lack of effort is associated with a preferred tax rate on the top 1% that is 4.1 percentage points lower (which amounts to 48% of the gap between self-identified democrats and republicans on this question). The belief that the rich amassed their wealth because they made an effort is associated with a preferred tax rate on the top 1% that is 7.1 percentage points lower (which amounts to 83% of the gap between self-identified democrats and republicans on this question). Even

controlling for these beliefs, support for Obama in the 2012 election correlates with higher desired tax rates at the top.

We exploit the 2x2 design and study how these two dimensions of trust interact in the determination of policy preferences. Conditional on high levels of trust in government, there is no effect of trust in business elites on taxes at the top, perhaps because subjects interpret the type of corruption primed by *Low Business* as exceptional and broadly under the control of good government. In such circumstances, raising taxes on all the rich might be hard to justify. In contrast, lowering trust in business elites when trust in government is low leads to a significant increase of 2.4 percentage points in the desired tax on the top 1%, which is double the size of the unconditional effect. The mean of the untreated group is 34.7 so the effect is non-negligible: it amounts to 27% of the gap between the average Democrat and Republican in our sample.

As a consequence of the large effects of trust in business conditional on low government, the effect of trust in business elites is negative and significant when averaged across the two government treatments. The causal impact implies an increase of 1.2 percentage points. Since the average desired tax rate at the top for Democrats in our sample is 37.6%, while for Republicans it is 29%, the treatment closes almost 14% of the gap between the two parties.

In contrast, we find that the (average) effect of trust in government on taxes on the top 1% (defined as $\frac{\beta_1 - \beta_2 + \beta_3}{2}$) is negative but not statistically significant. Interestingly, the estimate is similar, both in terms of sign and significance, to what was uncovered by Kuziemko *et al.* (2015) for the effect of distrust on taxes at the top 1%. Our result is driven by the non-significant result we find when we condition on *High Business*. However, lowering trust in government when trust in business elites is low leads to a 1.8 percentage points higher desired tax rate on the top 1%. Our interpretation of these results involves the utility derived by the public from harming players that act unethically. The public wants to punish businesspeople by raising taxes when the businesspeople act unethically, which might only happen when there is low quality of both the official and the businessperson. Put differently, our model suggests that people react to the possibility of bribes being exchanged, which is

most likely to happen when they observe both *Low Business* and *Low Government*.

Note that Kuziemko *et al.* (2015) implement a “distrust treatment” by asking questions designed to prime negative reactions regarding the government. Subjects primed to distrust the government decreased their support for transfer programs to the poor and that their “support for top tax rates generally falls as well (though only some of these effects are significant).” They study four questions. Lower trust in government causes an increase in the desired tax on the top 1% of 0.49 percentage points (with a standard error of 1.326): the sign of this effect is consistent with our data. The coefficients on the other three variables are negative: support for a tax on millionaires; support an increase in the estate tax and support for a petition to senators to increase the estate tax. Only this last result is statistically significant. If taxation at the top is driven by spitefulness and expectations of altruistic behavior, we obtain the opposite prediction, as in our data: more distrust increases taxes.

In columns 2, 3 and 4 we focus on preferences for taxes on other groups in the income distribution. The correlations with the basic set of beliefs are significant and have the expected sign, suggesting subjects paid attention to these other questions as well (e.g., they switch signs for taxes to the poorest half of the population). We find broadly similar effects of our treatments for desired taxes on the next top 9% (which completes the top 10% of the income distribution). When one looks at preferences for taxes over the next top 40%, and the bottom 50% (poorest), the treatments have basically no significant effect. Table A.1 in Appendix A.1 splits the sample in Democrats/Republicans with broadly similar results.

We interpret this evidence as inconsistent with the idea that taxes at the top are centrally determined by efficiency in tax collections, or by a desire to redistribute income regardless of moral entitlements. Instead, the evidence is consistent with reciprocal altruism, where taxes are a way of punishing corrupt businesspeople when the government is weak.

2.4 Supplementary survey

We ran a shorter supplementary survey on a smaller sample during November 2016. The title and description are identical to those of our main survey, with two differences: we paid

\$0.6 to each participant and described the survey as taking 6 minutes approximately. The survey is also identical up to the treatment windows.¹² We then showed an “attention check” question to enhance the accuracy of the responses (as in Alesina *et al.* 2018, who report work by Meade and Craig 2012 showing how these types of questions are helpful both in identifying careless respondents). Only 0.82% of the respondents reported inattention during our survey. The survey then asks the new outcome questions: a multipart question designed to inform the first stage (whether the treatments do in fact move trust in business or government) and a donation decision designed to connect preferences to outcomes.

2.4.1 First stage: the effect of the treatments on the dimensions of trust

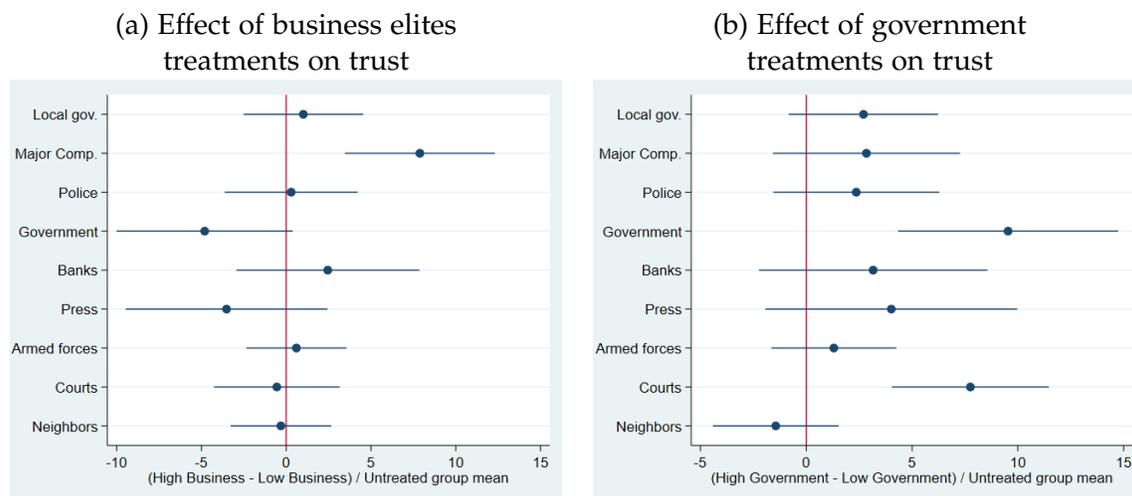
The new (multipart) outcome question first asked subjects their levels of trust in nine organizations: local government, major companies, the police, national government, banks, the press, armed forces, the courts, and their neighbors. This is helpful in assessing the “first stage” and corroborate that our treatments mainly impact the relevant dimensions of trust. We did not include this question in our main survey, and it would not be a good idea to include it, because individuals might feel manipulated (or individuals may be more prone to experimenter demand effects in the following questions) if we showed them a positive/negative text about business and then asked them directly what level of trust they had in business. It is also helpful to learn about the size of the effects in the first stage and to separate the two dimensions of trust.

Figure 2.4 summarizes the results (see Table A.2 for details). It supports our claim that changes in our measure of trust in business elites captures dimensions of trust related to major companies, and not something else (it is the only level of trust that is statistically significantly affected at the 5% level). Also, its impact equals a 7.9% change in the level of trust in Major Companies, measured in terms of deviations from the untreated group

¹²The survey is in Appendix A.3 and https://hbs.qualtrics.com/SE/?SID=SV_3NPCYdpcln07Ix. Besides the untreated group we included the following treatments (always showing the government treatment first): 1) *High Business and High Government*; 2) *High Business and Low Government*; 3) *Low Business and High Government*; 4) *Low Business and Low Government*; 5) *High Business*; 6) *Low Business*.

mean (40% of the gap in this question between Democrats and Republicans in the untreated group). Another take on the size of this effect is that it amounts, approximately, to the mean difference in trust in Major Companies between the US and Spain.¹³

Figure 2.4: *First stage*



Notes: We ran the basic regression specification described in section 2.2 with trust in Local government, Major Companies, The police, The government (in your nation’s capital), Banks, The press, The armed forces, The courts, and Your neighbors as dependent variables. In the left panel we present the following linear combination of coefficients: $\frac{\beta_1 - \beta_3 + \beta_2}{2}$ (as well as their 95% confidence intervals), and we divide this by the untreated group mean of the trust variable used as dependent variable. The same applies for the right panel but with $\frac{\beta_1 - \beta_2 + \beta_3}{2}$ instead. We consider the sample of people that spent at least 1.5 minutes in the supplementary survey (not considering the time spent in the treatment windows) and at least three seconds in every treatment.

With respect to the effect of trust in government, we see that this is associated with changes in the levels of trust regarding the national government and the courts. To get an idea of the magnitude of these results, they correspond to a 9.5% change in the level of trust in the national government and 7.8% in trust in courts, both measured in terms of deviations from the untreated group mean (151% and 268% of the gap in this question between Democrats and Republicans in the untreated group). It also amounts, approximately, to the

¹³This exercise could be carried out in different ways. Here we look at the impact of Trust in business elites on the level of trust in Major Companies measured in standard deviation units. Then we apply this difference to the mean level of trust in Major Companies for the US in the last wave of the WVS (with the standard deviation of the US in that dataset) and find the closest country next to this new level. Note that in the WVS there are only four possible responses, so we coded “A great deal of confidence”=4, “Quite a lot”=3, “Not very much”=2, and “None at all”=1).

mean difference in trust in the National Government between the US and Libya.

We conclude that our business treatment significantly changed trust in major companies without changing other dimensions of trust. It is therefore reasonable to assume that other, more distant, beliefs are also unaffected, so that the treatment isolates the causal effect of changing trust in business on policy preferences. The same is true with respect to our government treatment.

We also compare the effects of the two treatments and note that they impact trust differently. For example, the impact of our two trust treatments on trust in government is different at 1% level of significance (in the expected directions). The same is true with respect to trust in courts. The impact of the two treatments on trust in major companies is different at 12% level of significance (in the opposite direction, as expected).

2.4.2 Actions vs behaviors, mechanisms

The new donation question asked individuals to vote whether they wanted a donation to be made either to Citizens for Tax Justice, or The American Red Cross, or none. We told them that we would donate \$200 to the organization with the highest number of votes (which we did). We explained that Citizens for Tax Justice is “an NGO that seeks to **require the wealthy to pay their fair share**; it is primarily concerned with federal tax policy in the US and its mission is to give ordinary people a greater voice in the development of tax laws” and that The American Red Cross is “an NGO that seeks to **provide humanitarian help**; it is primarily focused on disaster relief and emergency assistance within the US”. Additionally, we included the option of not participating in the voting at all (“I don’t want to vote”). Although this vote is somewhat indirect (for example, it requires some trust in these intermediate organizations) we thought it would provide some useful data regarding subjects’ preferred policies.

Table 2.5 presents the results of a multinomial logit. We are interested in the impact of trust in business elites and trust in government in the log odds ratio of voting for Citizens for Tax Justice relative to voting for The American Red Cross. It shows that the results are

similar to those in the Main Survey but are weaker statistically (possibly due to smaller sample size, or due to attenuation given the intermediate step involving an NGO).¹⁴ It is also possible that the Red Cross is “neutral” and results would have been stronger if we had used a “pro-business” (or anti-tax) NGO. Trust in government has no effect on voting to tax the wealthy when trust in business elites is high, but it is negative when trust in business elites is low. In brief, results from a supplementary survey with a different question (donation to an NGO that aims to increase taxes at the top) are broadly consistent with main results as they suggest that distrust in government causes respondents to vote more for the pro-taxation organization (relative to The American Red Cross).

Table 2.5: *Voting for taxes on the top 1%*

<i>Panel A: Regression output</i>	Dependent variable: $\ln \left(\frac{P(\text{Vote Citizens for Tax Justice})}{P(\text{Vote American Red Cross})} \right)$		
	(1)	(2)	(3)
<i>Treatments</i>			
(β_1) High Business & High Government	0.098 (0.138)	-0.085 (0.139)	-0.063 (0.143)
(β_2) High Business & Low Government	-0.005 (0.137)	-0.031 (0.140)	-0.007 (0.144)
(β_3) Low Business & High Government	-0.272* (0.140)	-0.280** (0.142)	-0.275* (0.148)
<i>Other covariates</i>			
Poor didn't make an effort	-	-	-0.492*** (0.151)
Rich made an effort	-	-	-0.685*** (0.126)
Clinton	-	-	0.760*** (0.120)
Trust	-	-	-0.025 (0.022)
Observations	1,960	1,950	1,946
Untreated group mean	-0.687	-0.690	-0.696
<i>Panel B: Hypothesis testing over the coefficients</i>			

Continued on next page

¹⁴Given that the sample is smaller, and results are a bit weaker, we present the results with no additional controls (column 1), only with demographics controls (column 2), and with the full set of controls (column 3). Results remain unchanged across all these specifications.

Table 2.5: (Continued) Voting for taxes on the top 1%

	Dependent variable: $\ln \left(\frac{P(\text{Vote Citizens for Tax Justice})}{P(\text{Vote American Red Cross})} \right)$		
Effect of Trust in Business Elites			
High Bus – Low Bus	0.084 [0.3906]	0.081 [0.4095]	0.101 [0.3244]
High Bus – Low Bus High Gov	0.174 [0.2148]	0.194 [0.1693]	0.211 [0.1535]
High Bus – Low Bus Low Gov	-0.005 [0.9693]	-0.031 [0.8262]	-0.007 [0.9603]
Effect of Trust in Government			
High Gov – Low Gov	-0.183* [0.0629]	-0.167* [0.0918]	-0.165 [0.1084]
High Gov – Low Gov High Bus	-0.093 [0.4977]	-0.054 [0.6946]	-0.056 [0.6962]
High Gov – Low Gov Low Bus	-0.272* [0.0528]	-0.280** [0.0490]	-0.275* [0.0632]

Notes: Robust standard errors in parenthesis (Panel A). Multinomial logit model estimated, where the dependent variable can take three values: Voted for Citizens for Tax Justice, Voted for The American Red Cross, and did not vote. Regarding the treatments, the Low Business & Low Government group is the omitted group. Column (1) includes no additional controls. Column (2) includes demographic controls (gender, age, race, education, and type of employment). Column (3) includes same demographic controls, plus political variables and pre-treatment beliefs (includes relative support for Clinton in previous election, attitudes towards the rich and the poor, and general level of trust). Untreated group mean reports the mean of the dependent variable for the untreated group in each specification. In Panel B we present linear combinations of certain treatment coefficients and p-values (in brackets) for the test of whether these linear combinations are equal to 0. High Bus – Low Bus | High Gov is the difference between the treatment coefficient High Business & High Government and Low Business & High Government; High Bus – Low Bus is the weighted average of High Bus – Low Bus | High Gov and High Bus – Low Bus | Low Gov (analogous for other treatment groups). *, **, *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. We considered the sample of people that spent at least 1.5 minutes in the supplementary survey (not considering the time spent in the treatment windows) and at least three seconds in every treatment.

The new survey is also designed to help understand the mechanism behind our results. A key element of our intuition is that the public dislikes corruption (it dislikes when the other players act unethically), and that taxing the rich is a way to punish excessively selfish and unethical business people. But in this context, when there is corruption, would increasing taxes to the rich benefit corrupt government officials that might appropriate this revenue? If this were the case, it would generate a trade-off in people's preferences for

taxation: punishing corrupt businesspeople through high taxes would bring about a cost in terms of benefiting another corrupt party (the government official). To test whether this trade-off is present in people's minds, we asked: "Imagine that taxes to the top 1% (richest) of the population increase; what do you think will happen?" We gave them three options (we randomized the order in which the options appeared): "The money will be used to fund an increase in useful government spending" (selected by 45% of respondents), "The money will be wasted without clear benefits for the population" (selected by 40%), and "The money will be appropriated by corrupt government officials" (selected by only 15% of respondents). These are results considering only the untreated group (sample size is 502). The data suggest that the above-mentioned trade-off is not the first answer that comes to mind for a large proportion (85%) of the sample. This reinforces our interpretation that, in some circumstances, higher taxes on the rich are a way to punish business elites that are seen as undeserving.

2.5 Conclusions

We study the causal impact of two dimensions of trust on preferences for taxation using the general design proposed in Kuziemko *et al.* (2015). Our main innovation is using two "trust" treatments "priming" subjects to trust/distrust business elites and to trust/distrust government officials. Our online survey treats almost 9,000 Americans with a combination of photographs of well-known business leaders and of government officials with one-sided descriptions and "questions". In a supplementary survey we observe that these treatments have a large impact on trust in major companies and on trust in the national government, but not on other dimensions of trust.

Our main result is a negative causal impact of trust on preferences for taxation at the top. We find that subjects primed to distrust government and business elites would like a tax of 36.5% on those in the top 1% of the income distribution. This is significantly higher than the rates desired by those primed with other combinations of trust in business elites and government, with point estimates in the range 33.8-34.7 percent. Put differently, distrust

in business elites when trust in government is low causes an increase in desired taxes on the top 1% of 2.4 percentage points. This effect size is reasonable: it amounts to 6.9% of the untreated group mean, while the induced change in trust in Major Companies amounts to 7.9% measured in terms of deviations from the untreated group mean. At high levels of trust in government, desired taxes are essentially unaffected by changes in trust in business elites. Similarly, distrust in government when trust in business elites is low causes an increase in desired taxes on the top 1% of 1.8 percentage points. Again, the effect size is reasonable: it amounts to 5.2% of the untreated group mean, while the induced change in trust in the national government amounts to 9.5% measured in terms of deviations from the untreated group mean. At high levels of trust in business elites there is essentially no discernible impact of changes in trust in government. The results are broadly similar when we study the top 10% of the income distribution, while our “priming” had no effect on preferences for taxing households in other parts of the income distribution.

Our results support models where people demand taxes to punish the “undeserving” rich rather than to redistribute income. This is consistent with the experimental evidence gathered in Fisman *et al.* (2017) and the the work of Scheve and Stasavage (2016) who have argued that significant progressive taxation emerged after mass mobilization for the world wars and its impact on the belief that they implied unequal sacrifice among groups in society. They go on to write, “societies do not tax the rich just because they are democracies where the poor outnumber the rich or because inequality is high. Nor are beliefs about how taxes influence economic performance ultimately decisive. Societies tax the rich when people believe that the state has privileged the wealthy, and so fair compensation demands that the rich be taxed more heavily than the rest”.

The US economic system appears to give the rich several privileges (for example, it involves little redistribution) and confers high status to some of the richest members of society that are widely trusted. The results in this paper suggest that these features of American Exceptionalism are causally connected.

Chapter 3

Uncovering Elasticities with Notches and Kinks: Evidence from Peru¹

¹Co-authored with Rodrigo Azuero and Mariano Bosch.

Essay Abstract

Tax policies in Peru generate several notches and kink points in the choice sets of businesses. With rich administrative data we find clear evidence of bunching at these points. Using this bunching evidence, we estimate elasticities of taxable sales, profits, and assets with respect to their net-of-tax rates. Taxable sales elasticities are sizable at the lowest notch but monotonically decrease as we look at notches associated with higher levels of sales. We find elasticities of taxable profits in line with those found in other countries, and large elasticities of taxable assets.

3.1 Introduction

Since the pioneering work of Saez (2001) there has been an increasing interest in estimating elasticities of taxable income with respect to its net-of-tax rate, given that these has been shown to be a sufficient statistic for computing optimal income tax rates (see Saez *et al.* 2012 for a review). This interest carried over to other areas such as capital taxation (Saez and Stantcheva 2018 shows that the long-run elasticity of capital supply is a sufficient statistic for optimal capital taxation), and corporate income taxation where the tax base could be output or profits, or even both (Best *et al.* 2015). One of the main insights used to derive these elasticities has been to exploit kinks and notches in the choice set of individuals and firms generated by a tax. These kinks/notches have been shown to induce bunching of individuals/businesses around (or at and before) the kink/notch, and under some assumptions they allow us to estimate the elasticity of the tax base to its net-of-tax rate (Saez 2010; Chetty *et al.* 2011; Kleven and Waseem 2013).

This paper follows the methodological approaches developed by Saez (2010) and Kleven and Waseem (2013) to estimate elasticities of taxable sales, profits, and assets in Peru using rich administrative data. Peru serves as an interesting laboratory to estimate these elasticities given the number of notches and kinks generated by the Peruvian tax system (we exploit four notches and two kinks in this study). While there are several studies in developing contexts that estimate the elasticity of taxable income/sales (e.g., Boonzaaier *et al.* 2017; Kleven and Waseem 2013), those focusing on taxable profits are more scarce (e.g., Bachas and Soto 2018) and those focusing on taxable assets are, up to our knowledge, non-existent in the developing world.

We find three main results. First, we find relatively large elasticities of taxable sales when we look at the first notches generated by a tax regime available for micro businesses in Peru. However, these elasticities monotonically decrease as we look at notches associated with higher levels of sales. This monotonic relationship has also been found in other contexts (Kleven and Waseem 2013), so it would be interesting to understand if this is a more general

phenomena, and what model could explain this behavior.²

Second, we find elasticities of taxable profits around 0.2 for small and medium-sized businesses in Peru. Despite Peru being a developing country, these are in line with those found in advanced economies (Devereux *et al.* 2014).

Third, we find that even a 0.4% tax on assets over S/.1 MM (approximately USD 300,000) induces a large and sharp bunching, with an associated elasticity of 0.6. This elasticity is much higher than those estimated in advanced economies (e.g., Seim 2017 estimates an elasticity of taxable wealth in the range [0.09, 0.27]), which suggests that the desirability of wealth taxes in developing contexts should be carefully studied.³

Our paper contributes to a recent bunching literature (Kleven 2016 summarizes and reviews this literature). While our paper makes no theoretical contribution, it adds to the growing empirical evidence on bunching.⁴ Kinks and notches generated by taxes exist all around the globe, but they do not lead to bunching in all contexts (e.g., Bastani and Selin 2014; Bosch *et al.* 2019).

Our results also naturally connect to the literature on taxation and development (see Besley and Persson 2013), as the relatively large elasticities found in this setting could be suggestive of why it is so hard for governments in developing countries to raise taxes.⁵

The remainder of this paper is organized as follows. Section 3.2 describes the Peruvian tax system. Section 3.3 presents the data and some descriptive summary statistics. Section 3.4 describes how we apply the frameworks developed by Saez (2010) and Kleven and Waseem (2013) to estimate elasticities in our setting. Section 3.5 presents the main results of the paper. Section 3.6 concludes.

²Saez (2001) develops a simple tax evasion model to explain why there is bunching at the first kink generated by the US income tax schedule, but not at any other kink.

³Brülhart *et al.* (2016) also study this elasticity for the case of Switzerland and Jakobsen *et al.* (2018) for the case of Denmark.

⁴See Azuara *et al.* (2019) for evidence on bunching of firms around the eligibility threshold of various tax regimes in Peru.

⁵The deeper question would be, of course, why elasticities are so high in this setting. Perhaps the answer could be traced back to a possibly weak endogenous fiscal capacity, although answering this question is out of the scope of this paper.

3.2 A primer on the Peruvian tax system

In 2016, total tax revenue in Peru was approximately 16.1% of its GDP. The average tax to GDP ratio in Latin America and the Caribbean was 22.7% and Peru had the fourth lowest ratio after Guatemala (12.6%), Dominican Republic (13.7%), and Venezuela (14.4%).⁶ Part of the reason why tax levels are low in Peru is due to the high levels of informality observed in the country. Using information from the national household survey for employment (ENAH: “Encuesta Nacional de Hogares”) we find that 59% of the occupied labor force worked in a business that was not formally registered with the national tax authority and 47.3% of employees were working without a contract, which indicates that a large proportion of economic activities in Peru operate under the shadow economy not paying their corresponding tax obligations.

The main sources of tax revenue in Peru are the value added tax, VAT (which represented 50.8% of all tax revenue in 2016), followed by income tax (34.6% of tax revenue), and consumption tax (5.7%). The VAT rate is 18% although some goods and services, mostly related to agriculture, are excluded from the VAT obligations. Additionally, as we will explain in detail below, not all businesses are required to pay VAT. Income tax in Peru is further subdivided into five categories: tax on income from leases (0.01% of all tax revenue), tax on income from sales of assets (1.5%), corporate income tax (16%), personal income tax for independent workers (0.01%), and personal income tax for employees (8%).

Furthermore, there are various corporate income tax regimes in Peru. Such regimes introduce different kinks and notches that we exploit in this paper to estimate elasticities of taxable sales, profits, and assets. In the next subsections, we explain in detail the tax regimes for businesses in Peru before and after 2017. We make this split given that there was a major tax reform in 2017, and some of the notches we exploit are only present in 2010-2016, while some of the kinks we exploit are only present in 2017.

⁶These figures come from OECD (2018).

3.2.1 Tax regimes for businesses before 2017

Before the 2017 reform, businesses in Peru were required to register in one of the three available regimes of corporate income tax: the NRUS⁷, the RER⁸, and the General Regime (GR).⁹ In this study, we exploit the notches generated by the NRUS regime, and a kink in the assets tax schedule for businesses operating under the GR regime. We include the RER regime in this summary for completeness.

The NRUS is a tax regime directed to micro businesses registered as legal persons with monthly sales below S/.30,000 (approximately USD 9,000).¹⁰ Additionally, to be considered eligible for the NRUS a business should have total asset valuation under S/.70,000 (approximately USD 21,000) and should operate in only one location. Under this regime, the value added tax (VAT) and the corporate income tax (CIT) are grouped into a single fee that businesses pay on a monthly basis. Before the 2017 reform, the NRUS had five different categories depending on the level of reported sales. These categories would determine the monthly fee to be paid as specified in Figure 3.1.¹¹ Businesses registered in the NRUS are not required to keep any type of accounting books. They are only required to submit basic sales information to the tax authority via web, mobile phone app, or by submitting a physical copy to authorized banks. In 2016, 38.3% of all businesses registered formally to the tax authorities in Peru were registered in the NRUS.

⁷In Spanish, “Nuevo Régimen Único Simplificado”, New Unique Simplified Regime.

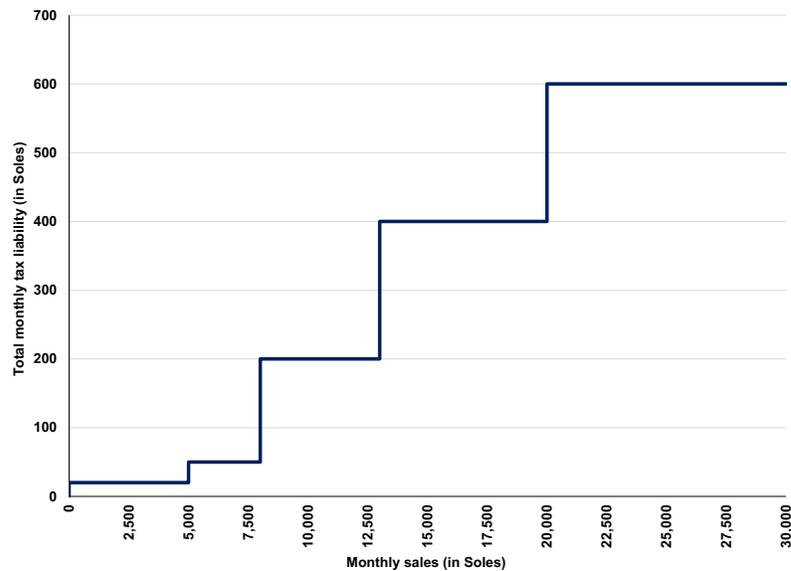
⁸In Spanish, “Régimen Especial de Renta”, Special Tax Regime.

⁹During the year, businesses can switch to a different regime but with some restrictions. If a business starts operating under the RER/GR regime, it can only switch to NRUS regime on January. NRUS businesses can switch to RER/GR any month, and RER can switch to GR any month.

¹⁰Purchases also need to be lower than S/.30,000, although NRUS businesses did not report them until 2017.

¹¹Businesses can pay in one category on a given month, and on another category on a different month, according to their monthly sales.

Figure 3.1: NRUS Tax Schedule



Notes: This figure shows the total monthly tax liability (in Soles) for a businesses operating under the NRUS regime as a function of its monthly reported sales. This tax schedule is valid for the period 2010-2016. Starting in 2017, businesses with monthly sales higher than 8,000 Soles can not operate under this regime.

The RER is a simplified tax regime for medium-sized businesses. To be eligible for this regime, firms should have yearly sales and purchases belows S/.525,000 (approximately USD 157,500), their total asset valuation should not exceed S/.126,000 (approximately USD 37,800), and should have ten or fewer employees. Under this regime, firms are required to pay the VAT as established in the tax code.¹² Additionally, firms pay a 1.5% tax on their total sales that substitutes the corporate income tax. Firms are required to keep accounting books detailing the registry of sales and purchases and are also required to submit monthly declarations with information about their general operations. Businesses registered in the RER are not required to submit annual declarations to the tax authorities, a requirement that before the 2017 reform only applied to businesses under the GR. In 2016, 24.6% of businesses were registered in the RER.

Finally, firms in the GR are required to pay a tax on net-assets, corporate income tax,

¹²The VAT general rate established in Peru corresponds to 18%. However, there are goods and services that are excluded from VAT or that are taxed at different rates such as some type of medications.

VAT, and need to keep five different accounting books detailing their sales, purchases, inventory, and balance sheets. If reported net-assets surpass S/.1,000,000 (approximately USD 300,000), then GR businesses have to pay (annually) 0.4% over the excess net-assets that surpass S/.1,000,000.¹³ This introduces a kink that we exploit to measure the elasticity of taxable assets to its tax rate. Before the 2017 reform, the tax rate for the corporate income tax was set at 28%. The base for the 28% corporate income tax corresponded to income less of allowed deductible expenses, which consisted of all purchases that were related to the main economic activity of the firm. In 2016 36.7% of businesses were registered in the GR.

3.2.2 The 2017 tax reform

In addition to increasing the corporate income tax rate from 28% to 29.5%, the 2017 tax reform established two main changes in the tax code for businesses.

First, it eliminated categories 3, 4, and 5 from the NRUS. This limited businesses that are eligible for this regime to have monthly sales under S/.8,000 (approximately USD 2,400). Because of this, when we exploit the notches created by the NRUS regime, we focus only on the period 2010-2016.

Second, the tax reform introduced a new tax regime for small and medium-sized businesses called the RMT.¹⁴ Firms that want to register under the RMT need to have annual sales under 1,700 UIT, which is equivalent to S/.6,885,000 in 2017 (approximately USD 2 MM).¹⁵ The only difference between the RMT and the GR is in the rate established for the corporate income tax.¹⁶ Under the RMT, profits under 15 UIT (approximately USD 18,000) are taxed at 10% and every unit exceeding this limit is taxed at the GR tax rate of 29.5%. This naturally introduces a kink which we exploit to estimate the elasticity of taxable profits

¹³This tax is called “Impuesto Temporal a los Activos Netos” and is regulated by the legislative decree 797. Businesses can deduct several types of assets when computing their net-assets, such as machinery bought in the last three years, or real estate when it is considered cultural heritage.

¹⁴In Spanish, “Régimen Mype Tributario”.

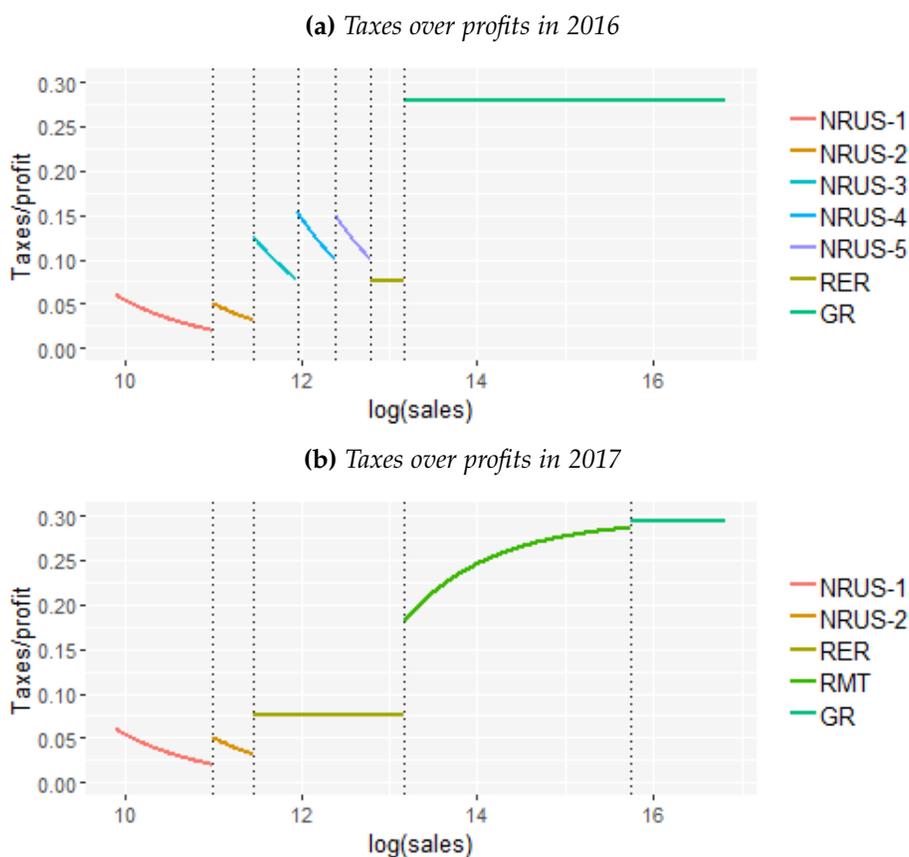
¹⁵UIT stands for “Unidad Tributaria Impositiva” (Taxing Unit).

¹⁶RMT businesses pay the same tax on net-assets than GR businesses, and this tax was not changed in 2017.

to its net-of-tax rate.

To see how all this tax reform effectively changed the tax schedule, in Figure 3.2 we plot the relationship between businesses' sales and the effective rate of taxation, defined as the total tax liability expressed in terms of total profits, under each regime before and after the reform. We simulate the tax liability assuming a profit margin of 20%. Note how the five different categories in the NRUS generated five different discontinuities in the effective rate of taxation before 2017. After 2017, only two categories of the NRUS survived and we observe a transition to the GR represented by the increasing effective tax rates in the RMT.

Figure 3.2: Effective tax rate for businesses before and after the 2017 reform



Notes: We simulate the tax liability for businesses with a profit margin of 20%. VAT liability is not calculated, only the corresponding part for the corporate income tax. Each vertical dashed line represents the sales-purchases limit for each regime.

In Table 3.1 we report the share of businesses in each regime before and after the reform,

as well as the revenue collected by the tax authority from each regime. We note that before the 2017 reform, NRUS was the most common regime for businesses in Peru followed by the GR and the RER. Although there does not seem to be a large disparity in the proportion of businesses registered in each regime, we note that the distribution of revenue is concentrated in businesses registered under the GR. Whereas the total revenue coming from businesses registered in the NRUS and the RER combined did not represent 1% of total government revenue in 2016, corporate income taxes from the GR represented almost 16% of total government revenue. By comparing Tables 3.1a and 3.1b we are able to identify some of the effects of the 2017 reform. First, the proportion of businesses registered under the GR decreased from 36.7% to 14.3%. Second, revenue collected from the GR decreased from S/.16,495 million to S/.15,499 and even the revenue collected from the RMT is not enough to offset the lost revenue from the GR.

Table 3.1: *Distribution and revenue from tax regimes before and after the 2017 reform*

(a) *Tax regimes in 2016.*

Regime	Distribution of businesses	Revenue	Revenue/Gov. income (%)	Revenue/GDP (%)
NRUS	38.3%	185.2	0.18	0.03
RER	24.6%	334.9	0.32	0.05
RMT	-	0	0	0
GR	36.7%	16,495.8	15.94	2.53

(b) *Tax regimes in 2017.*

Regime	Distribution of businesses	Revenue	Revenue/Gov. income (%)	Revenue/GDP (%)
NRUS	36.5%	140.3	0.13	0.02
RER	23.5%	322.9	0.31	0.05
RMT	25.1%	904.1	0.86	0.14
GR	14.3%	15,499.2	14.68	2.53

Notes: Revenue is reported in millions of S/. and does not include VAT paid in each regime.

3.3 Data and descriptive statistics

The main dataset we use comes from the tax administrator of Peru, which is the “Superintendencia Nacional de Aduanas y Administración Tributaria” (SUNAT), and it consists on the universe of corporate tax returns filed during the period 2010-2017 to SUNAT.¹⁷

Businesses have to report either monthly and/or annually to SUNAT, depending on the tax regime under which they operate. In terms of their reporting obligations, NRUS businesses have to declare their sales (and starting in 2017, they also report purchases) to SUNAT monthly through “Formulario 1611”.¹⁸ We only have annualized data of these reports, which as will be explained in Section 3.4.1, it poses some challenges to estimate elasticities of taxable sales (relative to if we had access to the monthly data).

Except for NRUS businesses that do not have to pay the VAT, all businesses have to report monthly to SUNAT their levels of sales, acquisitions, and other information in order to compute their VAT corresponding to that month. They report these figures in PDT 0621 and we have annualized data from this report.

RG and RMT businesses also have to make an annual declaration to SUNAT where they report more detailed level information regarding sales, acquisitions, assets, liabilities, and other information, in order to recompute their corporate tax for that year. We also have access to these reports, and we use the information from these tax returns to estimate the elasticity of taxable profits and assets.

Table 3.2 shows some summary statistics. It is clear from the table below that Peru is undergoing a period of mayor increase in its number of formal businesses, going from 1.3 million businesses in 2010, to 1.8 million in 2017, which represents an increase of 43%.

¹⁷More in particular, we have annualized data from all filings of “Formulario 1611”, annual corporate tax returns filings, and annualized data from PDT 0621.

¹⁸NRUS businesses are the only ones that should complete this form.

Table 3.2: Number of businesses in SUNAT's database (figures expressed in 100,000 of businesses)

	2010	2011	2012	2013	2014	2015	2016	2017
Total (100,000)	12.6	13.5	14.7	15.7	16.6	17.5	18.0	17.9
Regimes (% Total)								
GR	44.2	42.9	41.4	39.7	38.5	37.4	36.7	14.3
RMT	0.0	0.0	0.0	0.0	0.0	0.0	0.0	25.1
RER	17.7	19.5	20.9	22.1	23.2	24.0	24.6	23.5
NRUS	37.5	37.1	37.3	37.7	37.9	38.1	38.3	36.5
Files tax return (% Total)								
Annual return	38.8	37.8	36.8	36.1	35.2	33.6	30.8	33.4
PDT 0621	62.3	62.7	62.6	62.2	62.1	61.9	61.7	62.9
Form 1611	38.3	38.0	38.3	38.6	38.7	39.0	39.1	37.7
Economic activity (% Total)								
Agricultural	2.0	1.9	1.8	1.7	1.6	1.6	1.5	1.5
Fishing	0.3	0.3	0.3	0.3	0.2	0.2	0.2	0.2
Mining	0.5	0.6	0.7	0.7	0.9	0.8	0.7	0.9
Manufacturing	10.0	9.8	9.6	9.4	9.0	8.7	8.6	8.5
Other services	37.6	37.9	38.3	39.0	39.8	40.9	41.5	41.4
Construction	3.7	4.0	4.1	4.3	4.2	4.4	4.5	4.6
Retail	45.8	45.6	45.2	44.8	44.3	43.4	43.0	42.9

Notes: "Total" is the number of businesses that filed any tax return to SUNAT in a given year; "GR" is the number of businesses whose last tax return filed in the year indicated the business should be considered GR (analogous for RMT, RER and NRUS); "Annual return" is the number of businesses that filed an annual corporate tax return in a given year (analogous for PDT 0621 and form 1611); "Agricultural" is the number of businesses that filed any tax return or presented an employment report to SUNAT in a given year and that operate in the agricultural sector according to their CIU (analogous for other sectors). Source of data is SUNAT.

Also, a high percentage of these businesses operate under a "simplified tax regime". Roughly 55% of the firms operate under one of these regimes in 2010 (NRUS or RER) and 85% in 2017 (NRUS, RER or RMT), where a great part of this increase is explained by the introduction of RMT in 2017. This new regime is responsible for the drop in the number of businesses in the GR regime in 2017, as 70% of RMT businesses in 2017 operated under the GR regime in 2016. A remaining 21% corresponds to businesses that entered SUNAT's database in 2017, 7% operated under RER in 2016 and 2% under NRUS.

3.4 Methodology to compute elasticities

This section describes how we estimate elasticities of taxable sales, profits, and assets with respect to their net-of-tax rates. In order to compute these estimates, we closely follow the framework of Kleven and Waseem (2013) when exploiting a pure notch (taxable sales) and that of Saez (2010) when exploiting a kink (taxable profits and assets). Slemrod (2013) summarizes and discusses both approaches.

3.4.1 Pure notches

As described in Section 3.2, businesses enrolled in the NRUS regime have a tax liability that jumps when their reported sales surpass a certain threshold. These are the classical examples of pure notches. Kleven and Waseem (2013) derives a formula to compute the elasticity of taxable income assuming a standard quasi-linear and isoelastic utility function.¹⁹ For the case in which we only consider the possibility of a pure notch, the formula can be simplified as follows:

$$\frac{1}{1 + \frac{\Delta z^*}{z^*}} \left[1 + \frac{\Delta T}{z^*} \right] - \frac{1}{1 + \frac{1}{e}} \left[\frac{1}{1 + \frac{\Delta z^*}{z^*}} \right]^{1 + \frac{1}{e}} - \frac{1}{1 + e} = 0 \quad (3.1)$$

Where z^* is the level of sales where tax liability changes, e is the elasticity, ΔT is the “size of the notch” (i.e., the change in tax liability if a business surpasses the notch), Δz^* is the quantity that determines the bunching region (i.e., there will be bunching at the notch for all individuals that, absent the notch, would have had sales in $(z^*, z^* + \Delta z^*)$).

We cannot get an explicit formula for e , but we can solve for e numerically given an estimate of z^* , ΔT and Δz^* . While z^* is readily observable, this is not the case for ΔT and Δz^* . First, note that if we had monthly level data, then ΔT would simply be the difference between the monthly payment of a certain category of NRUS and the next category (e.g., it would be equal to S/.30 if we are looking at the first notch, given that businesses have to pay S/.20 if they are in the first category of NRUS and S/.50 if they are in the second

¹⁹See Kleven and Waseem (2013) pp. 672-684 for the details of this derivation.

category, as can be seen in Figure 3.1). However, because we only have access to annual data, ΔT would take different values depending on whether we assume that businesses can fully determine how to distribute their sales through the year (which we call the “fully-adjustable” scenario), or not. As an example, if a business’ annual sales surpasses S/.60,000 by an ϵ , it could end up paying between S/.30 and S/.360 because of that ϵ depending on how that ϵ was distributed in time. If the business sells S/.5,000 for eleven months and S/.5,000 + ϵ on the remaining month, then the ϵ ends up increasing the annual tax liability by only S/.30. However, if the business sells S/.5,000 + $\frac{\epsilon}{12}$ every month, the ϵ ends up increasing the tax liability by S/.360 (as the business has to report in the second category of NRUS every month, instead of only on one month). Given the limitations of our current data, we estimate the elasticity of taxable income assuming the two extreme scenarios (i.e., sales can be “fully-adjusted” to minimize annual tax liability, or sales are constant every month).

Regarding Δz^* , we follow the “convergence method” derived Kleven and Waseem (2013) to estimate it.²⁰ We briefly describe this procedure to estimate the earnings response (Δz^*) at the first notch of the NRUS regime. The procedure is analogous for the other three notches generated by the NRUS regime.

First, we estimate a counter-factual density considering NRUS businesses (during 2010-2016) with annual taxable sales between S/.40,000 and S/.80,000 (the discrete change in the tax liability occurs at S/.60,000 which corresponds approximately to USD 18,000). We consider 41 bins of businesses according to their reported annual taxable sales, and define three bins to the left of the threshold as those “visibly affected” by bunching. We consider polynomials up to degree n and run the following regression to estimate the counter-factual density:

$$F_j = \sum_{i=0}^{i=n} \beta_i (y_j)^i + \sum_{i=z_0}^{i=z_h} \gamma_i 1[y_j = y_i] + \epsilon_j \quad (3.2)$$

²⁰Kleven and Waseem (2013) also estimate Δz^* with a second method which we do not pursue in this study. It requires a precise definition of the dominated region (i.e., the region where a business can be strictly better off by producing at the notch). However, this region depends on the ratio of profits over sales for each business and we do not have the information necessary to compute this ratio for NRUS businesses.

Where F_j is the number of businesses with annual taxable sales in the bin j ; y_j is the midpoint of sales for those businesses in bin j ; $[z_0, z_h]$ is the range of excluded bins of businesses in the vicinity of the threshold (and $\Delta z^* = z_h - z^*$). We then construct a fitted series of F_j without considering the dummy variables: $\hat{F}_j = \sum_{i=0}^{i=n} \hat{\beta}_i (y_j)^i$.

Following Kleven and Waseem (2013), z_0 is determined by visually inspecting the distribution and z_h by an iterative procedure in which z_h is chosen in a way such that excess mass before the notch equals missing mass after the notch (i.e., $\sum_{j=z_0}^{j=j^*} (F_j - \hat{F}_j) = \sum_{j>j^*}^{j=z_h} (\hat{F}_j - F_j)$ where j^* is the bin that includes annual sales of S/.60,000).

Once we have an estimate of z_h , we plug in Δz^* , z^* , and ΔT in equation 3.1 and solve numerically for e . Standard errors are estimated using a bootstrap procedure by generating a large number (1,000) of taxable sales distribution from the original distribution.²¹ We compute an elasticity with each of these bootstrapped samples, and the standard error of the elasticity is estimated as the standard deviation of the elasticities computed in the bootstrapped samples.

3.4.2 Kinks

As described in Section 3.2, businesses enrolled in the RMT regime face a marginal tax rate over profits of 10% when profits are lower than 15 UIT (approximately USD 18,000). For profits that surpass this amount, the marginal tax rate is 29.5%. Saez (2010) derives a formula to estimate the elasticity of taxable income exploiting bunching at the kink points of the US income tax schedule, which we apply to our setting.

Saez (2010) derives the following formula assuming a standard quasi-linear and isoelastic utility function:²²

$$b = z^* \left[\left(\frac{1-t_0}{1-t_1} \right)^e - 1 \right] \frac{h(z^*)_- + h(z^*)_+ / \left(\frac{1-t_0}{1-t_1} \right)^e}{2} \quad (3.3)$$

²¹Each of these samples is generated by drawing from replacement from the original distribution and it has the same size (i.e., number of businesses) than the original distribution.

²²See Saez (2010) pp. 185-189 for the details of this derivation.

Where z^* is the kink point, e is the elasticity (in our case, the elasticity of taxable profits with respect to their net-of-tax rate), t_0 is the marginal tax rate before reaching the kink (t_1 is the marginal tax rate after the kink), $h(z^*)_-$ and $h(z^*)_+$ represent the densities of before-tax profits (i.e., if there was no change in the marginal tax rate), and b is our bunching estimate. To solve for e , all the quantities in equation 3.3 are readily available with the exception of $h(z^*)_-$, $h(z^*)_+$ and b .

To estimate these, we need to make an assumption on the “size” of the bunching region by defining a δ such that we expect to see bunching in the region $(z^* - \delta, z^* + \delta)$. Following Saez (2010), we visually inspect the distribution and show the robustness of our results to different values of δ .²³

We follow a slightly different approach to estimate the elasticity of taxable assets, given that in our context the bunching is sharp at the threshold with this variable. In this scenario, using $h(z^*)_-$ and $h(z^*)_+$ to get an estimate of bunching does not seem the best approach and we instead follow Seim (2017) (who builds on Saez 2010 and Chetty *et al.* 2011) to estimate this elasticity.

Analogous to Seim (2017), consider a group of business owners with strictly quasi-concave preferences that have to choose taxable assets on a given year. These business owners are heterogeneous in their preferences, savings and evasion technologies, and are distributed according to some continuous and differentiable cumulative distribution function. With a constant linear tax rate τ , denote the density of business owners with a level of taxable assets z , to be equal to $h(z)$. It is easy to show that the elasticity of taxable assets with respect to its net-of-tax rate can be approximated by:²⁴

²³Chetty *et al.* (2011) and Saez (2010) look at bunching around the cutoff, and not below the cutoff, because in their case bunching did not take the form of a mass point precisely at the cutoff (this is also our case when looking at profits, but not when we look at assets). As Best *et al.* (2015) argues, this is probably because of optimization errors by the agents. Kleven and Waseem (2013) present bunching evidence more consistent with a mass point at the cutoff in Pakistan and thus, estimates excess bunching by looking at firms below the cutoff.

²⁴See Seim (2017), p.405 for the details. It is an approximation because we use $h(z^*)dz^*$ as a proxy for the number of individuals who bunch at the threshold (and the smaller dz^* is, the better the approximation is).

$$\epsilon_{A,\tau} = \frac{B}{h(z^*)z^*} \frac{1 - \tau}{d\tau} \quad (3.4)$$

Where B is the number of businesses that are bunching at the kink z^* . To estimate $\frac{B}{h(z^*)}$ we first estimate a counter-factual distribution of businesses as we did in Section 3.4.1. Then we compute an estimate of B as the number of “excess” businesses in the bunching region (i.e., $\sum_{j=z_0}^{j^*} (F_j - \hat{F}_j)$) and we divide it by the average number of counter-factual businesses in the bunching region to get an estimate of $\frac{B}{h(z^*)z^*}$. Following Bastani and Selin (2014) and Paetzold (2019) we express z^* in units of the bin size that we use to compute our estimate of B .

Standard errors are computed following the same procedure outlined in Section 3.4.1.

3.5 Results

In the following section, we provide estimates of the elasticities of taxable sales, profits, and assets with respect to their net-of-tax rates. The tax regimes available to businesses in Peru generate several kinks and notches in the choice set of businesses which may foster evasion and/or mis-allocation for businesses operating in the vicinity of these discontinuities. These notches and kinks induce bunching of businesses at certain thresholds, which we exploit to compute different elasticities following the methodologies developed by Saez (2010) and Kleven and Waseem (2013).²⁵

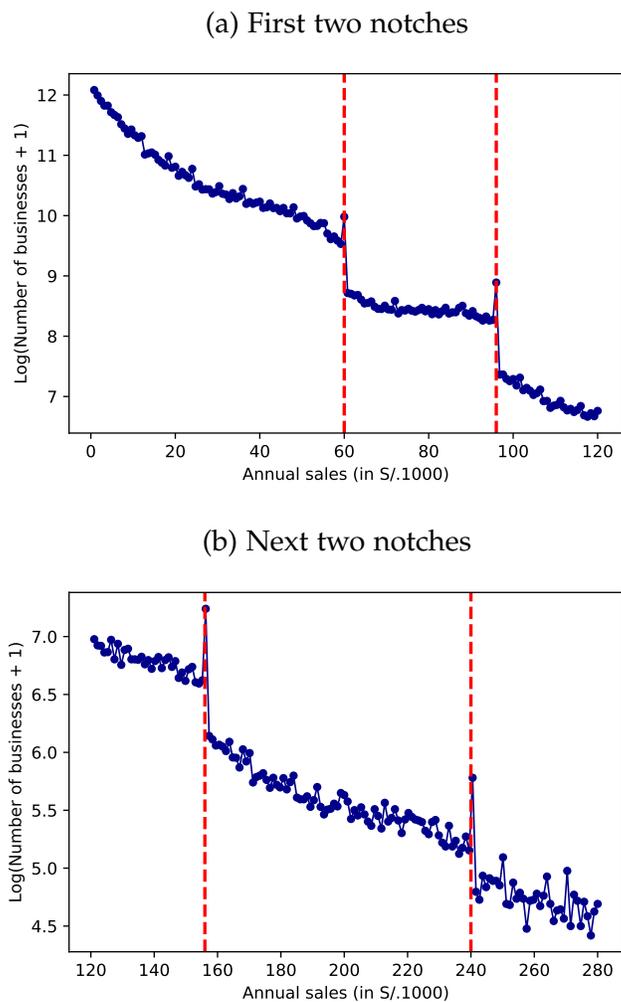
3.5.1 Taxable sales

As described in Section 3.2, NRUS businesses face a tax schedule with four notches. In Figure 3.3 we show the distribution of NRUS businesses according to their reported annual sales, where we see that businesses bunch right at each of the four notches. The bunching

²⁵In every figure where we present an empirical distribution of certain tax base, we show the notch/kink with a vertical red line. The bin that we plot at the notch/kink is always entirely on the tax-favored side of the threshold.

seems not to be diffused around the threshold (as in Saez 2010), but rather quite sharp exactly at the notch.

Figure 3.3: Taxable sales: empirical distributions (NRUS businesses, 2010-16)

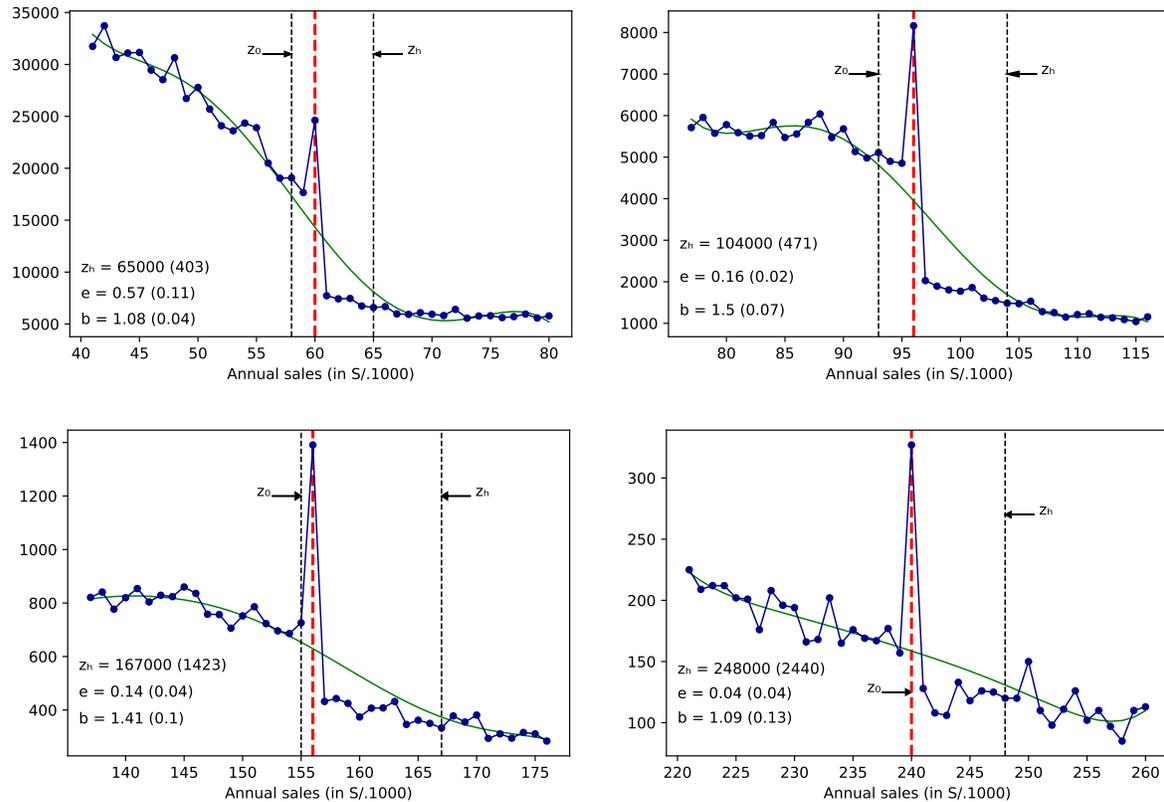


Notes: Each dot represents the log (plus 1) number (in 1000s) of NRUS businesses with annual reported sales that falls in a certain bin, during 2010-2016. We consider 151 bins. The red vertical line represents the maximum annual level of sales for a business staying in the same category of NRUS through the year ($x = 60,000; 96,000; 156,000; 240,000$ Soles). Source of data is SUNAT.

Next, we estimate the elasticity of taxable sales for these micro businesses. Figure 3.4 shows the empirical and counter-factual distribution of businesses in the vicinity of each of the notches generated by the NRUS regime. It also shows our estimates of bunching

(b), z_h , and one of our elasticity estimates.²⁶ Bootstrapped standard errors are shown in parenthesis.

Figure 3.4: Taxable sales: empirical and counter-factual distributions around notches (NRUS businesses, 2010-16)



Notes: Each dot represents the number of NRUS businesses with annual reported sales that falls in a certain bin, during 2010-2016. We consider 41 bins. The counter-factual is computed as described in Section 3.4.1. The red vertical line represents the maximum annual level of sales for a business staying in the same category of NRUS through the year ($x = 60,000; 96,000; 156,000; 240,000$ Soles). Source of data is SUNAT.

Table 3.4 shows how these bunching estimates (column 2) translate into elasticities of taxable sales. Column 3 shows the share of unresponsive individuals in the dominated region which is in line with the values found by Kleven and Waseem (2013) in Pakistan.²⁷ Column

²⁶Following Kleven and Waseem (2013), the bunching estimate is calculated as $b = \sum_{j=z_0}^{j=j^*} (F_j - \hat{F}_j)$ and normalized by the average counter-factual frequency in the dominated range. We consider a “conservative” estimate of dominated range (z^*, z_D), by assuming a ratio of profits over sales of 1 so that the dominated range at the first notch is (60000, 60360), (96000, 97800) at the second notch, etc.

²⁷Kleven and Waseem (2013) finds a^* between 0.512 and 0.861 in their non-rounder sample. a^* is computed

4 shows our estimates of Δz^* and column 5 and 6 show the two “extreme” assumptions we make regarding ΔT . Column 6 and 7 show the elasticity estimates which are a function of the assumption we make regarding ΔT . While it is hard to compare the magnitude of these elasticities to those found in other contexts, given that these seem to vary a lot depending on our assumption of ΔT , a clear picture emerges regarding their relative size. Elasticities seem to be high at the first notch, but they monotonically decrease as we look at notches associated with higher levels of sales (both if we focus on the “fully-adjusted” or “non-adjustment” case). This was also the case in Kleven and Waseem (2013) study in Pakistan, which raises an interesting question of why we observe this pattern.²⁸

Table 3.4: *Structural taxable sales elasticities*

(1) Notch	(2) Bunching b	(3) a^*	(4) Δz^*	ΔT		Elasticity e	
				(5) Fully-adj.	(6) Non-adj.	(7) Fully-adj.	(8) Non-adj.
60,000	1.08 (0.04)	0.6 (0.01)	5000 (396.97)	30	360	7.1 (1.26)	0.57 (0.1)
96,000	1.5 (0.07)	0.56 (0.01)	8000 (473.49)	150	1800	2.25 (0.26)	0.16 (0.02)
156,000	1.38 (0.1)	0.71 (0.04)	11000 (1357.37)	200	2400	1.96 (0.51)	0.14 (0.04)
240,000	1.08 (0.12)	0.82 (0.08)	8000 (2357.23)	200	2400	0.66 (0.49)	0.04 (0.04)

Notes: The table presents four different notches (expressed in Soles) created by the NRUS regime in column 1 and their associated bunching estimates in column 2. Estimates of frictions (share of individuals in the dominated region that are unresponsive to the notch) can be found in column 3. For this statistic, we consider a “conservative” estimation of dominated range, by assuming a ratio of profits over sales of 1. Column 4 shows the earnings response computed with the convergence method explained in Section 3.4.1. Column 5 and 6 shows the two possible values of ΔT considered, depending on whether we assume that sales can be “fully-adjusted” to minimize annual tax liability, or not. Column 7 and 8 shows our estimates of the elasticity of taxable sales. Bootstrapped standard errors are shown in parenthesis. Source of data is SUNAT.

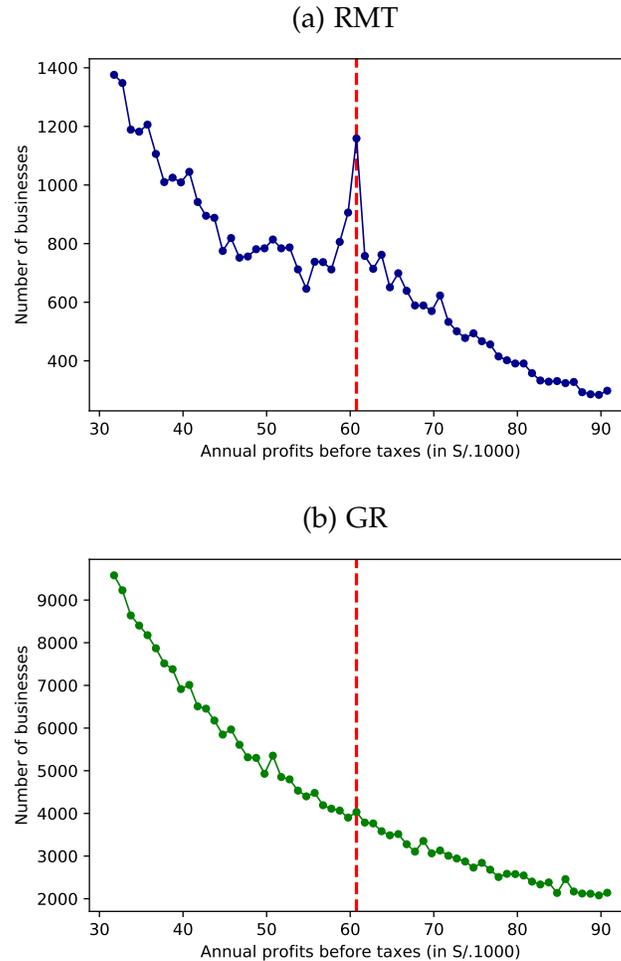
$$\text{as } \frac{\sum_{j>z^*}^{j=z^D} F_j}{\sum_{j>z^*}^{j=z^D} \hat{F}_j}$$

²⁸In Table 2, Kleven and Waseem (2013) show five structural elasticities (computed with the “convergence method”), corresponding to 5 different notches. These are: 1.021, 0.188, 0.171, 0.079, 0.035.

3.5.2 Taxable profits

We look at the distribution of taxable annual profits for RMT and GR businesses around S/. 60,750 (Figure 3.5). As described in Section 3.2, RMT businesses pay 10% on profits if these are less than S/. 60,750, and 29.5% over profits that exceed S/. 60,750. GR businesses pay a flat marginal tax rate on profits of 28% - 29.5% depending on the year. We see that RMT businesses bunch at the kink, while there is no discontinuity in the distribution of GR businesses at the same level of profits. In addition, we note that there seems to be bunching around the threshold and not only at (or before) the kink. This was also the case in Saez (2010), and the argument is that probably businesses cannot perfectly target the kink and end up being slightly above it in some cases. The associated elasticity in this case is 0.24 with a standard error of 0.11.

Figure 3.5: Taxable profits: empirical distributions (RMT and GR businesses)

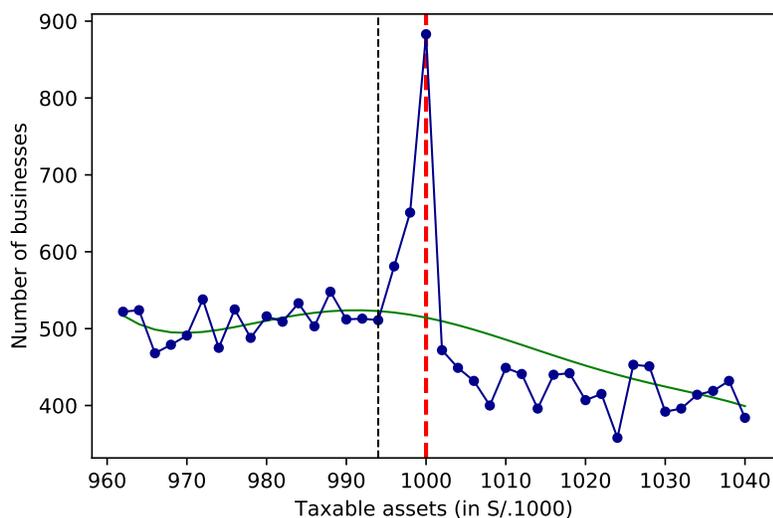


Notes: Each dot represents the number of RMT or GR businesses with annual reported taxable profits that falls in a certain bin. Panel (a) considers data only for 2017 and panel (b) for 2010-2017. We consider 61 bins. The red vertical line represents the level of profits where there is a change in the marginal tax rate for RMT businesses. Source of data is SUNAT.

3.5.3 Taxable assets

We show our estimates of the elasticity of taxable assets for medium sized or big businesses. Figure 3.6 shows the empirical distribution for GR and RMT businesses around the kink where the marginal tax rate over taxable assets changes from 0% to 0.4%.

Figure 3.6: Taxable assets: empirical distributions around the kink (RG/RMT businesses, 2010-17)



Notes: Each dot represents the log (plus 1) number of RG and RMT businesses with annual reported taxable assets that falls in a certain bin, during 2010-2017. We consider 41 bins. The red vertical line represents the level of taxable assets where there is a change in the marginal tax rate. The vertical dashed black lines represent the lower bound of the bunching region. Source of data is SUNAT.

We see that there does not seem to be bunching all around the threshold, but rather below and at the kink. This region can be found between the black and red dashed lines, and the associated elasticity in this case is 0.56 with a standard error of 0.15.

3.6 Conclusions

In this study, we apply the frameworks of Kleven and Waseem (2013) and Saez (2010) to estimate elasticities of taxable sales, profits, and assets with respect to their net-of-tax rates in Peru. We use rich administrative data provided by the Peruvian tax administration (SUNAT) for the period 2010-2017 and find three main results.

First, we find sizable elasticities of taxable sales when we look at the first notches of NRUS, which monotonically decrease as we look at notches associated with higher levels of sales. This monotonic relationship is also apparent in Kleven and Waseem (2013); Saez (2010) also finds that people seem to be more responsive to the first kink of the US income tax schedule than to other kinks. It could be interesting to compute these estimates in other

contexts to see how general this pattern is, and what could be the underlying model that explains this behavior.

Second, we find elasticities of taxable profits for small and medium-sized business in Peru around 0.2, which are aligned to those found in advanced economies.

Third, we find that even a 0.4% tax on assets over S/.1 MM (approximately USD 300,000) induces a large and sharp bunching, which means that the elasticity of taxable assets is quite high. Compared to Seim (2017), who estimates the elasticity of taxable wealth in Sweden, our estimate of 0.6 is twice the size.²⁹ This result suggests that the desirability of wealth taxes in developing contexts should be carefully studied.

²⁹As a sanity check, to see if these large differences make sense, we compare the distribution of businesses around the kink in Peru and Sweden. In Sweden (see Figure 1 in Seim 2017), the marginal tax rate over net-assets changes from 0% to 1.5% at the kink and the number of businesses at the kink are 15% higher than those in the counter-factual. In Peru, the marginal tax rate changes from 0% to 0.4% and the number of businesses at the kink are 140% higher than those in the counter-factual.

References

- AGHION, P., ALGAN, Y., CAHUC, P. and SHLEIFER, A. (2010). Regulation and distrust. *The Quarterly Journal of Economics*, **125** (3), 1015–1049.
- AGUIRRE, F., BLANCHARD, P., BORRAZ, F. and SALDAIN, J. (2015). ¿Los beneficiarios del programa Tarjeta Uruguay Social accederían a mejores precios de ampliarse el conjunto de comercios donde pueden comprar? *Mimeo*.
- AKERLOF, G. A. and ROMER, P. M. (1993). Looting: The Economic Underworld of Bankruptcy for Profit. *Brookings Papers on Economic Activity*, **1993** (2), 1–73.
- ALESINA, A. and ANGELETOS, G.-M. (2005a). Corruption, inequality, and fairness. *Journal of Monetary Economics*, **52** (7), 1227–1244.
- and — (2005b). Fairness and Redistribution. *The American Economic Review*, **57** (3), 415–426.
- , STANTCHEVA, S. and TESO, E. (2018). Intergenerational Mobility and Preferences for Redistribution. *American Economic Review*, **108** (2), 521–554.
- ALGAN, Y. and CAHUC, P. (2010). Inherited trust and growth. *American Economic Review*, **100** (5), 2060–2092.
- ALZÚA, M. L., CRUCES, G. and RIPANI, L. (2013). Welfare programs and labor supply in developing countries: Experimental evidence from Latin America. *Journal of Population Economics*, **26** (4), 1255–1284.
- AMARANTE, V., FERRANDO, M. and VIGORITO, A. (2013). Teenage School Attendance and Cash Transfers: An Impact Evaluation of PANES. *Economía*, **14** (1), 1–33.
- and GÓMEZ, M. (2016). Diferenciales de ingreso entre trabajadores formales e informales en Uruguay, 2001-2014. *Revista de Economía*, **23** (1), 71–86.
- , MANACORDA, M., MIGUEL, E. and VIGORITO, A. (2016). Do cash transfers improve birth outcomes? Evidence from matched vital statistics, and program and social security data. *American Economic Journal: Economic Policy*, **8** (2), 1–43.
- , —, VIGORITO, A. and ZERPA, M. (2011). Social Assistance and Labor Market Outcomes: Evidence from the Uruguayan PANES. *IDB Technical Note 453*.
- ANTEL, J. J. (1992). The Intergenerational Transfer of Welfare Dependency: Some Statistical Evidence. *The Review of Economics and Statistics*, **74** (3), 467–473.

- ARAÚJO, M. C., BOSCH, M., MALDONADO, R. and SCHADY, N. (2017). The Effect of Welfare Payments on Work in a Middle-Income Country. *IDB Working Paper Series No. 830*.
- ARROW, K. J. (1972). Gifts and Exchanges. *Philosophy & Public Affairs*, **1** (4), 343–362.
- AUTOR, D., KOSTØL, A., MOGSTAD, M. and SETZLER, B. (2019). Disability Benefits, Consumption Insurance, and Household Labor Supply. *American Economic Review*, **109** (7), 2613–2654.
- AZUARA, O., AZUERO, R., BOSCH, M. and TORRES, J. (2019). Special Tax Regimes in Latin America and the Caribbean. *IDB Working Paper Series No. 970*.
- BACHAS, P. and SOTO, M. (2018). Not(ch) Your Average Tax System: Corporate Taxation under Weak Enforcement. *World Bank Policy Research Working Paper 8524*.
- BAICKER, K., FINKELSTEIN, A., SONG, J. and TAUBMAN, S. (2014). The Impact of Medicaid on Labor Market Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment. *American Economic Review: Papers & Proceedings*, **104** (5), 322–328.
- BAIRD, S., FERREIRA, F. H., ÖZLER, B. and WOOLCOCK, M. (2014). Conditional, unconditional and everything in between: a systematic review of the effects of cash transfer programmes on schooling outcomes. *Journal of Development Effectiveness*, **6** (1), 1–43.
- , MCKENZIE, D. and ÖZLER, B. (2018). The Effects of Cash Transfers on Adult Labor Market Outcomes. *World Bank Policy Research Working Paper No. 8404*.
- BANERJEE, A. V., HANNA, R., KREINDLER, G. E. and OLKEN, B. A. (2017). Debunking the Stereotype of the Lazy Welfare Recipient: Evidence from Cash Transfer Programs. *The World Bank Research Observer*, **3** (2), 155–184.
- BANFIELD, E. C. (1958). *The moral basis of a backward society*.
- BASTAGLI, F., HAGEN-ZANKER, J., HARMAN, L., BARCA, V., STURGE, G. and SCHMIDT, T. (2019). The Impact of Cash Transfers: A review of the evidence from low- and middle-income countries. *Journal of Social Policy*, **48** (3), 569–594.
- BASTANI, S. and SELIN, H. (2014). Bunching and non-bunching at kink points of the Swedish tax schedule. *Journal of Public Economics*, **109**, 36–49.
- BENABOU, R. and OK, E. A. (2001). Social mobility and the demand for Redistribution: The POUM Hypothesis. *The Quarterly Journal of Economics*, **116** (2), 447 – 487.
- BERDEJÓ, C. and CHEN, D. L. (2012). Priming Ideology? Electoral Cycles Without Electoral Incentives Among Elite U . S . Judges. *Mimeo*, (February).
- BERGOLO, M. and CRUCES, G. (2016). The Anatomy of Behavioral Responses to Social Assistance When Informal Employment Is High. *IZA Discussion Paper No. 10197*.
- and GALVAN, E. (2016). Intra-Household Behavioral Responses to Cash Transfer Programs: Evidence from a Regression Discontinuity Design. *IZA Discussion Paper Series*, (10310).

- BESLEY, T. and PERSSON, T. (2013). Taxation and Development. In *Handbook of Public Economics*, vol. 5, Elsevier B.V., pp. 51–110.
- BEST, M. C., BROCKMEYER, A., KLEVEN, H. J., SPINNEWIJN, J. and WASEEM, M. (2015). Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan. *Journal of Political Economy*, **123** (6), 1311–1355.
- BHARGAVA, B. S. and MANOLI, D. (2015). Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment. *American Economic Review*, **105** (11), 3489–3529.
- BLANK, R. M. (1989). Analyzing the length of welfare spells. *Journal of Public Economics*, **39**, 245–273.
- BOONZAAIER, W., HARJU, J., MATIKKA, T. and PIRTTILA, J. (2017). How Do Small Firms Respond to Tax Schedule Discontinuities? Evidence from South African Tax Registers. *VATT Working Papers* 85.
- BOSCH, M. and MANACORDA, M. (2012). Social Policies and Labor Market Outcomes in Latin America and the Caribbean: A Review of the Existing Evidence. *Centre for Economic Performance, London School of Economics and Political Science*.
- BOSCH, N., JONGEN, E., LEENDERS, W. and MÖHLMANN, J. (2019). Non-bunching at kinks and notches in cash transfers in the Netherlands. *International Tax and Public Finance*, **26**, 1329–1352.
- BOSCHMAN, S., MAAS, I., KRISTIANSEN, M. H. and VROOMAN, J. C. (2019). The reproduction of benefit receipt: Disentangling the intergenerational transmission. *Social Science Research*, **80**, 51–65.
- BRUICH, G. A. (2014). The Effect of SNAP Benefits on Expenditures: New Evidence from Scanner Data and the November 2013 Benefit Cuts. *Mimeo*.
- BRÜLHART, M., GRUBER, J., KRAPF, M. and SCHMIDHEINY, K. (2016). Taxing Wealth: Evidence from Switzerland. *NBER Working Paper* 22376.
- CABANNES, Y. (2004). Participatory budgeting: a significant contribution to participatory democracy. *Environment and Urbanization*, **16** (1), 27–46.
- CALONICO, S., CATTANEO, M. D. and TITIUNIK, R. (2014). Robust data-driven inference in the regression-discontinuity design. *Stata Journal*, **14** (4), 909–946.
- CAPPELEN, A. W., HAALAND, I. K. and TUNGODDEN, B. (2018). Beliefs about Behavioral Responses to Taxation. *Mimeo*, pp. 1–21.
- , HOLE, A. D., SORENSEN, E. and TUNGODDEN, B. (2007). The pluralism of fairness ideals: An experimental approach. *American Economic Review*, **97** (3), 818–827.
- CARD, D., MAS, A., MORETTI, E. and SAEZ, E. (2012). Inequality at work: The effect of peer salaries on job satisfaction. *American Economic Review*, **102** (6), 2981–3003.

- CARUSO, M., GALIANI, S. and WEINSCHELBAUM, F. (2019). Poverty alleviation strategies under informality: evidence for Latin America. *NBER Working Paper No. 2633*.
- CHAN, M. K. and MOFFITT, R. (2018). Welfare Reform and the Labor Market. *Annual Review of Economics*, **10**, 347–381.
- CHARITÉ, J., FISMAN, R. and KUZIEMKO, I. (2015). Reference points and redistributive preferences: experimental evidence. *NBER Working Paper*, (21009).
- CHEN, Y. and LI, S. (2009). Group identity and social preferences. *The American Economic Review*, **99** (1), 431–457.
- CHETTY, R., FRIEDMAN, J. N., OLSEN, T. and PISTAFERRI, L. (2011). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: evidence from danish tax records. *The Quarterly Journal of Economics*, **126** (2), 749–804.
- , — and SAEZ, E. (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings. *American Economic Review*, **103** (7), 2683–2721.
- CICCA (2012). *Informe Tarjeta Uruguay Social*. Tech. rep.
- COOKE, M. (2009). A Welfare Trap? The Duration and Dynamics of Social Assistance Use among Lone Mothers in Canada. *Canadian Review of Sociology*, **46** (3), 179–206.
- CURRIE, J. (2003). U.S. Food and Nutrition Programs. In R. A. Moffitt (ed.), *Means-Tested Transfer Programs in the United States*.
- DAHL, B. G. B., LØKEN, K. V. and MOGSTAD, M. (2014a). Peer Effects in Program Participation. *American Economic Review*, **104** (7), 2049–2074.
- DAHL, G. B. and GIELEN, A. C. (2018). Intergenerational spillovers in disability insurance. *NBER Working Paper No. 24296*.
- , KOSTOL, A. R. and MOGSTAD, M. (2014b). Family Welfare Cultures. *The Quarterly Journal of Economics*, pp. 1711–1752.
- DASGUPTA, P. and RAY, D. (1986). Inequality as a Determinant of Malnutrition and Unemployment. *The Economic Journal*, **96** (384), 1011–1034.
- DAY, M. V. and FISKE, S. T. (2017). Movin on Up? How Perceptions of Social Mobility Affect Our Willingness to Defend the System. *Social Psychological and Personality Science*, **8** (3), 1948550616678454.
- DE LA PLAZA, L. and SIRTAINÉ, S. (2005). An Analysis of the 2002 Uruguayan Banking Crisis. *World Bank Policy Research Working Paper No. 3780*.
- DE WACHTER, S. and GALIANI, S. (2006). Optimal income support targeting. *International Tax and Public Finance*, **13** (6), 661–684.
- DESHPANDE, M. (2016). Does Welfare Inhibit Success? The Long-Term Effects of Removing Low-Income Youth from the Disability Rolls. *American Economic Review*, **106** (11), 3300–3330.

- DEVEREUX, M. P., LIU, L. and LORETZ, S. (2014). The Elasticity of Corporate Taxable Income: New Evidence from UK Tax Records. *6* (2), 19–53.
- DI TELLA, R. and MACCULLOCH, R. (2009). Why Doesn't Capitalism Flow to Poor Countries? *Brookings Papers on Economic Activity*, pp. 285–321.
- DINEM (2011). *Informe MIDES: Evaluación y seguimiento de programas 2009-2010*. Tech. rep.
- (2012). *Informe MIDES: Seguimiento y evaluación de actividades y programas 2011-2012*. Tech. rep.
- (2018). *Evolución de la TUS y presupuesto consolidado a junio 2018*. Tech. rep.
- DJANKOV, S., LA PORTA, R., LOPEZ-DE SILANES, F. and SHLEIFER, A. (2002). The Regulation of Entry. *Quarterly Journal of Economics*, **117** (1), 1–37.
- DURLAF, S. N. and FAFCHAMPS, M. (2004). Social capital.
- EAST, C. N. (2018). Immigrants' labor supply response to Food Stamp access. *Labour Economics*, **51**, 202–226.
- FEHR, E. and SCHMIDT, K. M. (1999). A Theory Of Fairness, Competition, and Cooperation. *Quarterly Journal of Economics*, **114** (3), 817–868.
- FILGUEIRA, F. (2005). Welfare and Democracy in Latin America: The Development, Crises and Aftermath of Universal, Dual and Exclusionary Social States. *UNRISD Working Paper*.
- FINKELSTEIN, A. and NOTOWIDIGDO, M. J. (2019). Take-up and targeting: experimental evidence from SNAP. *The Quarterly Journal of Economics*, pp. 1505–1556.
- FISMAN, R., GLADSTONE, K., KUZIEMKO, I. and NAIDU, S. (2017). Do Americans Want to Tax Capital? Evidence from Online Surveys. *NBER Working Paper Series*.
- FISZBEIN, A. and SCHADY, N. (2009). *Conditional Cash Transfers: Reducing Present and Future Poverty*. Washington DC: World Bank Policy Research Report.
- FONG, C. M. and LUTTMER, E. F. P. (2007). What Determines Giving to Hurricane Katrina Victims? Experimental Evidence on Income, Race, and Fairness. *National Bureau of Economic Research Working Paper Series*, **No. 13219** (2008), 64–87.
- FRAKER, T. and MOFFITT, R. (1988). The effect of food stamps on labor supply. *Journal of Public Economics*, **35** (1), 25–56.
- FUKUYAMA, F. (1995). *Trust: the social virtues and the creation of prosperity*. New York: Free Press.
- GANONG, P. and NOEL, P. (2019). Consumer Spending During Unemployment: Positive and Normative Implications. *American Economic Review*, **109** (7), 2383–2424.
- GARGANTA, S. and GASPARINI, L. (2015). The impact of a social program on labor informality: The case of AUH in Argentina. *Journal of Development Economics*, **115**, 99–110.

- GASPARINI, L., HAIMOVICH, F. and OLIVIERI, S. (2009). Labor informality bias of a poverty-alleviation program in Argentina. *Journal of Applied Economics*, **12** (2), 181–205.
- and TORNAROLLI, L. (2009). Labor Informality in Latin America and the Caribbean: Patterns and Trends from Household Survey Microdata. *Desarrollo y Sociedad*, **63**, 13–80.
- GELMAN, A. and IMBENS, G. (2019). Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs. *Journal of Business & Economic Statistics*, **37** (3), 447–456.
- GUISSO, L., SAPIENZA, P. and ZINGALES, L. (2004). The role of social capital in financial development. *The American Economic Review*, **94** (3), 526–556.
- , — and — (2008). Trusting the Stock Market. *The Journal of Finance*, **63** (6), 2557–2600.
- GUNDERSEN, C. and OLIVEIRA, V. (2001). The Food Stamp Program and Food Insufficiency. *American Journal of Agricultural Economics*, **83** (4), 875–887.
- GURIEV, S. and RACHINSKY, A. (2005). The Role of Oligarchs in Russian Capitalism. *Journal of Economic Perspectives—Volume*, **19** (1—Winter), 131–150.
- GUZI, M., REPUBLIC, C., CENTRAL, C. and LABOR, E. (2014). An Empirical Analysis of Welfare Dependence in the Czech Republic. *Journal of Economics and Finance*, **64** (5), 407–432.
- HAMILTON, J. T. and ZECKHAUSER, R. (2004). Media Coverage of CEOs: Who? What? Where? When? Why? *Workshop on the Media and Economic Performance*, pp. 1–59.
- HANGSTROM, P. A. (1996). The Food Stamp Participation and Labor Supply of Married Couples: An Empirical Analysis of Joint Decisions. *Journal of Human Resources*, **31** (2), 383–403.
- HANSEN, J. (2007). Human Capital and Welfare Dynamics in Canada. *The B.E. Journal of Economic Analysis & Policy*, **7** (1), Article 27.
- HAUSHOFER, J. and SHAPIRO, J. (2016). The short-term impact of unconditional cash transfers to the poor: experimental evidence from Kenya. *The Quarterly Journal of Economics*, pp. 1973–2042.
- HORTON, J. J., RAND, D. G. and ZECKHAUSER, R. J. (2011). The online laboratory: Conducting experiments in a real labor market. *Experimental Economics*, **14** (3), 399–425.
- HOYNES, H. and SCHANZENBACH, D. W. (2016). US Food and Nutrition Programs. In R. A. Moffitt (ed.), *Economics of Means-Tested Transfer Programs in the United States, Volume I*.
- IMBENS, G. and KALYANARAMAN, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies*, **79** (3), 933–959.
- INEED (2014). *Informe sobre el estado de la educación en Uruguay 2014*. Montevideo: INEEEd.
- JACOB, R. and ZHU, P. (2012). A Practical Guide to Regression Discontinuity. *Mimeo*.
- JÄGER, S. and HEINING, J. (2019). How Substitutable Are Workers? Evidence from Worker Deaths. *Mimeo*.

- JAKOBSEN, K., JAKOBSEN, K., KLEVEN, H. J. and ZUCKMAN, G. (2018). Wealth taxation and wealth accumulation: theory and evidence from Denmark. *NBER Working Paper 24371*.
- JOHNSON, M. and LAKOFF, G. (1980). *Metaphors We Live By*. Chicago: University of Chicago Press.
- JONES, D. and MARINESCU, I. (2019). The Labor Market Impacts of Universal and Permanent Cash Transfers: Evidence from the Alaska Permanent Fund. *Mimeo*.
- KATZ, L. F. and MEYER, B. D. (1990). The impact of the potential duration of unemployment benefits on the duration of unemployment. *Journal of Public Economics*, **41**, 45–72.
- KEANE, M. and MOFFITT, R. (1998). A Structural Model of Multiple Welfare Program Participation and Labor Supply. *International Economic Review*, **39** (3), 553–589.
- KESSELMAN, J. R. (1976). Tax effects on job search, training, and work effort. *Journal of Public Economics*, **6**, 255–272.
- KHANNA, T. and YAFEH, Y. (2007). Business Groups in Emerging Markets: Paragons or Parasites? *Journal of Economic Literature*, **XLV** (June), 331–372.
- KLEVEN, H. J. (2016). Bunching. *The Annual Review of Economics*, **8**, 435–464.
- and KOPCZUK, W. (2011). Transfer Program Complexity and the Take-Up of Social Benefits. *American Economic Journal: Economic Policy*, **3**, 54–90.
- and WASEEM, M. (2013). Using notches to uncover optimization frictions and structural elasticities: theory and evidence from Pakistan. *The Quarterly Journal of Economics*, **128**, 669–723.
- KLOR, E. F. and SHAYO, M. (2010). Social identity and preferences over redistribution. *Journal of Public Economics*, **94** (3-4), 269–278.
- KNACK, S. and KEEFER, P. (1995). Institutions and Economic Performance: Cross Country Tests Using Alternative Institutional Measures. *Economics and Politics*, **7** (3), 207–227.
- KUZIEMKO, I., NORTON, M. I., SAEZ, E. and STANTCHEVA, S. (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. Online Appendix. *American Economic Review*, **105** (4), 1478–1508.
- LA PORTA, R., LOPEZ-DE-SILANE, F., SHLEIFER, A. and VISHNY, R. W. (1997). Trust in Large Organizations. *AER Papers and Proceedings*, **87** (May), 333–338.
- , LOPEZ-DE-SILANES, F. and ZAMARRIPA, G. (2003). Related Lending. *The Quarterly Journal of Economics*, (February), 231–268.
- LAZEAR, E. (1975). Education: Consumption or Production. *NBER Working Paper No. 104*.
- LEVINE, D. K. (1998). Modeling Altruism and Spitefulness in Experiments. *Review of Economic Dynamics*, **1** (3), 593–622.
- LINDBLOM, C. E. (1977). *Politics and markets: the world's political economic systems*.

- MALLAR, C. (1982). Evaluation of the Economic Impact of the Job Corps Program: Third Follow-Up Report. *Mathematica Policy Research Report*.
- MALMENDIER, U. and NAGEL, S. (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking? *Quarterly Journal of Economics*, **126** (1), 373–416.
- MANACORDA, M., MIGUEL, E. and VIGORITO, A. (2011). Government Transfers and Political Support. *American Economic Journal: Applied Economics*, **3** (1), 1–28.
- MCCRARY, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, **142** (2), 698–714.
- MCINTOSH, C., BAIRD, S. and ÖZLER, B. (2010). Cash or Condition? Evidence from a Cash Transfer Experiment. *World Bank Policy Research Working Paper No. 5259*.
- , MONESTIER, F., PIÑEIRO, R., ROSENBLATT, F. and TUÑÓN, G. (2017). The Impacts of a Randomized Housing Policy in Uruguay. *EGAP Pre-analysis plan*.
- MEADE, A. W. and CRAIG, S. B. (2012). Identifying careless responses in survey data. *Psychological Methods*, **17** (3), 437–455.
- MEDGYESI, M. and TEMESVÁRY, Z. (2013). Conditional cash transfers in high-income OECD countries and their effects on human capital accumulation. *GINI Discussion Paper No. 84*.
- MELTZER, A. H. and RICHARD, S. F. (1981). A Rational Theory of the Size of Government. *The Journal of Political Economy*, **89** (5), 914–927.
- MIDES (2016). *Memoria Anual 2016: Ministerio de Desarrollo Social*. Tech. rep.
- MILLER, C., KATZ, L. F., AZURDIA, G., ISEN, A., SCHULTZ, C. and ALOISI, K. (2018). *Boosting the Earned Income Tax Credit for Singles: Final Impact Findings from the Paycheck Plus Demonstration in New York City*. Tech. rep.
- MILLERI, A. and SANDERS, S. G. (1997). Human capital development and welfare participation. *Carnegie-Rochester Conference Series on Public Policy*, **46**, 1–43.
- MOFFITT, B. R. (1983). An Economic Model of Welfare Stigma. *American Economic Review*, **73** (5), 1023–1035.
- MOFFITT, R. A. (2002). Welfare programs and labor supply. *Handbook of Public Economics*, **4**, 2394–2430.
- (2016a). *Economics of Means-Tested Transfer Programs in the United States*, vol. 1. Chicago: Univ. Chicago Press.
- (2016b). *Economics of Means-Tested Transfer Programs in the United States*, vol. 2. Chicago: Univ. Chicago Press.
- MOGSTAD, M. (2012). Are Lone Mothers Responsive to Policy Changes? Evidence from a Workfare Reform in a Generous Welfare State. *Scandinavian Journal of Economics*, **114** (4), 1129–1159.

- MORCK, R., WOLFENZON, D. and YEUNG, B. (2005). Corporate Governance, Economic Entrenchment, and Growth. *Journal of Economic Literature*, **43** (September), 655–720.
- MULLAINATHAN, S., SCHWARTZSTEIN, J. and SHLEIFER, A. (2008). Coarse thinking and persuasion. *The Quarterly Journal of Economics*, **123** (2), 577–619.
- OECD (2018). *Revenue Statistics in Latin America and the Caribbean 1990-2016*.
- ORDABAYEVA, N. and FERNANDES, D. (2017). Similarity focus and support for redistribution. *Journal of Experimental Social Psychology*, **72** (March), 67–74.
- PAETZOLD, J. (2019). *How do taxpayers respond to a large kink? Evidence on earnings and deduction behavior from Austria*, vol. 26. Springer US.
- PERAZZO, I., SALAS, G. and VIGORITO, A. (2016). *Evaluación de impacto del Programa Cercanías*. Tech. rep.
- PINOTTI, P. (2012). Trust, Regulation and Market Failures. *Review of Economics and Statistics*, **94** (3), 650–658.
- PUTNAM, R. D. (1993). *Making democracy work: Civic traditions in modern Italy*. N.J.: Princeton University Press.
- RIDDELL, C. and RIDDELL, W. C. (2014). The pitfalls of work requirements in welfare-to-work policies: Experimental evidence on human capital accumulation in the Self-Sufficiency Project. *Journal of Public Economics*, **117**, 39–49.
- ROTEMBERG, J. J. (2008). Minimally acceptable altruism and the ultimatum game. *Journal of Economic Behavior and Organization*, **66** (3-4), 457–476.
- SAAVEDRA, J. E. and GARCIA, S. (2012). Impacts of Conditional Cash Transfer Programs on Educational Outcomes in Developing Countries: A Meta-Analysis. *RAND Working Paper No. 921-1*.
- SAEZ, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies*, **68** (1), 205–229.
- (2010). Do Taxpayers Bunch at Kink Points? *American Economic Journal: Economic Policy*, **2** (3), 180–212.
- , SLEMROD, J. and GIERTZ, S. H. (2012). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, **50** (1), 3–50.
- and STANTCHEVA, S. (2016). Generalized Social Marginal Welfare Weights for Optimal Tax Theory. *American Economic Review*, **106** (1), 24–45.
- and — (2018). A simpler theory of optimal capital taxation. *Journal of Public Economics*, **162**, 120–142.
- SCHEVE, K. and STASAVAGE, D. (2012). Democracy, war, and wealth: Lessons from two centuries of inheritance taxation. *American Political Science Review*, **106** (1), 81–102.

- and — (2016). *Taxing the Rich*. Princeton University press.
- SEIM, D. (2017). Behavioral Responses to Wealth Taxes: Evidence from Sweden. *American Economic Journal: Economic Policy*, **9** (4), 395–421.
- SHAOAN, H., YAN, C. and RUI, L. (2019). Welfare Rigidity, the Composition of Public Expenditure and the Welfare Trap. *Social Sciences in China*, **40** (3), 110–129.
- SLEMROD, J. (2013). Buenas notches: lines and notches in tax system design. *eJournal of Tax Research*, **11** (3), 259–283.
- SRULL, T. K. and WYER, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, **37** (10), 1660–1672.
- TELLA, R. D. and DUBRA, J. (2014). Anger and Regulation. **116** (3), 734–765.
- WEINZIERL, M. (2014). The promise of positive optimal taxation: Normative diversity and a role for equal sacrifice. *Journal of Public Economics*, **118**, 128–142.
- WILLIAMSON, H. and WHITMORE, D. (2012). Work incentives and the Food Stamp Program. *Journal of Public Economics*, **96** (1-2), 151–162.
- ZITZEWITZ, E. (2012). Forensic Economics. *Journal of Economic Literature*, **50** (3), 731–769.

Appendix A

Appendix to Chapter 2

A.1 Additional results

Table A.1: Taxes: Democrat vs Republican

		Tax rate top 1%		Tax rate next 9%		Tax rate next 40%		Tax rate bottom 50%	
		Democrat	Republican	Democrat	Republican	Democrat	Republican	Democrat	Republican
<i>Panel A: Regression output</i>									
<i>Treatments</i>									
(β_1) High Business & High Government	-1.959 (1.250)	-1.553 (1.813)	-1.701* (0.990)	-1.279 (1.482)	-0.470 (0.725)	-0.007 (1.152)	0.515 (0.568)	0.367 (1.002)	
(β_2) High Business & Low Government	-1.899* (1.085)	-3.253** (1.503)	-1.698* (0.873)	-2.747** (1.271)	-0.423 (0.634)	-1.609* (0.943)	-0.234 (0.471)	-0.759 (0.731)	
(β_3) Low Business & High Government	-2.257** (1.082)	-0.841 (1.548)	-1.878** (0.869)	-0.779 (1.310)	-0.664 (0.627)	-0.144 (0.974)	-0.111 (0.455)	0.715 (0.783)	
Observations	3528	1589	3528	1589	3526	1589	3520	1586	
Untreated group mean	37.581	29.049	28.375	22.854	18.159	17.079	8.888	10.256	
<i>Panel B: Hypothesis testing over the coefficients</i>									
<i>Effect of Trust in Business</i>									
High Bus – Low Bus	-0.800 [0.2901]	-1.982* [0.0612]	-0.760 [0.2041]	-1.623* [0.0598]	-0.114 [0.7974]	-0.735 [0.2654]	0.195 [0.5736]	-0.553 [0.3394]	
High Bus – Low Bus High Gov	0.298 [0.7771]	-0.712 [0.6328]	0.176 [0.8292]	-0.500 [0.6675]	0.194 [0.7564]	0.137 [0.8818]	0.626 [0.2220]	-0.348 [0.6982]	
High Bus – Low Bus Low Gov	-1.899* [0.0803]	-3.253** [0.0306]	-1.698* [0.0518]	-2.747** [0.0308]	-0.423 [0.5048]	-1.609* [0.0881]	-0.234 [0.6190]	-0.759 [0.2996]	
<i>Effect of Trust in Government</i>									
High Gov – Low Gov	-1.159 [0.1255]	0.429 [0.6849]	-0.940 [0.1163]	0.344 [0.6898]	-0.355 [0.4241]	0.729 [0.2698]	0.319 [0.3591]	0.920 [0.1123]	
High Gov – Low Gov High Bus	-0.060 [0.9544]	1.699 [0.2392]	-0.003 [0.9971]	1.467 [0.1910]	-0.047 [0.9401]	1.602* [0.0728]	0.749 [0.1549]	1.125 [0.1877]	

Continued on next page

Table A.1: (Continued) Taxes: Democrat vs Republican

	Tax rate top 1%		Tax rate next 9%		Tax rate next 40%		Tax rate bottom 50%	
	Democrat	Republican	Democrat	Republican	Democrat	Republican	Democrat	Republican
High Gov – Low	-2.257**	-0.841	-1.878**	-0.779	-0.664	-0.144	-0.111	0.715
Gov Low Bus	[0.0370]	[0.5871]	[0.0308]	[0.5522]	[0.2896]	[0.8822]	[0.8074]	[0.3610]

Notes: Robust standard errors in parenthesis. *, **, *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. A constant term (not shown) is included in every regression. Regarding the treatments, the Low Business & Low Government group is the omitted group. Democrat corresponds to the sample of people that supported Obama or leaned towards Obama in the 2012 election (Republican are those that supported or leaned towards Romney). In Panel B we present linear combinations of certain treatment coefficients and p-values (in brackets) for the test of whether these linear combinations are equal to 0. We considered the sample of people that spent at least three minutes in the survey (not considering the candy experiment and time spent in the treatment windows) and at least three seconds in every treatment.

Table A.2: First stage

	Local gov. (1)	Major Comp. (2)	Police (3)	Gov. (4)	Banks (5)	Press (6)	Armed forces (7)	Courts (8)	Neighbors (9)
<i>Panel A: Regression output</i>									
Treatments									
(β_1) High Business & High Government	0.204 (0.140)	0.485*** (0.143)	0.159 (0.173)	0.203 (0.161)	0.234 (0.162)	0.020 (0.168)	0.134 (0.150)	0.403*** (0.149)	-0.112 (0.135)
(β_2) High Business & Low Government	0.108 (0.139)	0.411*** (0.144)	0.127 (0.170)	-0.214 (0.162)	0.241 (0.164)	-0.051 (0.171)	0.055 (0.154)	0.104 (0.153)	-0.014 (0.139)
(β_3) Low Business & High Government	0.201 (0.144)	0.183 (0.145)	0.251 (0.170)	0.402** (0.164)	0.271 (0.165)	0.246 (0.169)	0.105 (0.151)	0.569*** (0.153)	-0.086 (0.138)
Observations	1,960	1,960	1,960	1,960	1,960	1,960	1,960	1,960	1,960
Untreated group mean	5.482	4.526	5.972	4.297	4.169	3.948	7.020	5.592	6.382

<i>Panel B: Hypothesis testing over the coefficients</i>									
Effect of Trust in Business Elites									
High Bus – Low Bus	0.055 [0.574]	0.357*** [0.000]	0.017 [0.884]	-0.207* [0.070]	0.102 [0.373]	-0.139 [0.246]	0.042 [0.690]	-0.031 [0.771]	-0.02 [0.836]
High Bus – Low Bus High Gov	0.002 [0.986]	0.303** [0.036]	-0.092 [0.584]	-0.199 [0.212]	-0.037 [0.818]	-0.227 [0.177]	0.029 [0.841]	-0.166 [0.254]	-0.026 [0.848]
High Bus – Low Bus Low Gov	0.108 [0.437]	0.411*** [0.005]	0.127 [0.456]	-0.214 [0.187]	0.241 [0.141]	-0.051 [0.765]	0.055 [0.720]	0.104 [0.497]	-0.014 [0.918]
Effect of Trust in Government									
High Gov – Low Gov	0.148 [0.133]	0.129 [0.207]	0.141 [0.237]	0.410*** [0.000]	0.132 [0.251]	0.159 [0.185]	0.092 [0.387]	0.434*** [0.000]	-0.092 [0.342]

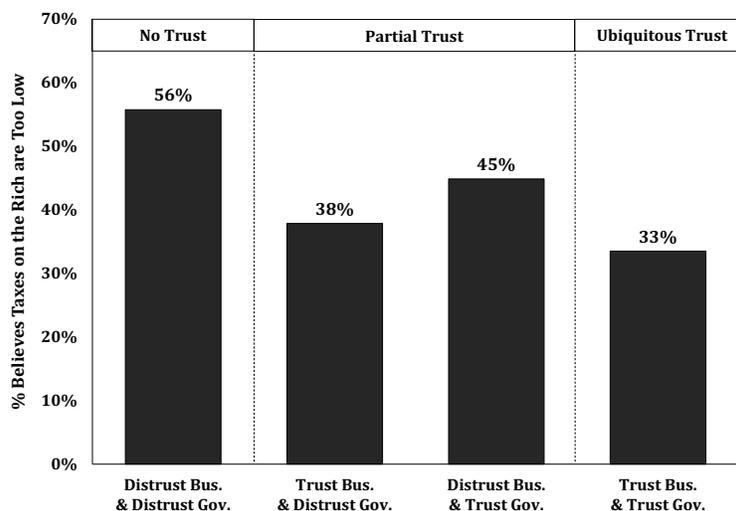
Continued on next page

Table A.2: (Continued) First stage

	Local gov. (1)	Major Comp. (2)	Police (3)	Gov. (4)	Banks (5)	Press (6)	Armed forces (7)	Courts (8)	Neighbors (9)
High Gov – Low Gov	0.095	0.075	0.032	0.417***	-0.007	0.071	0.078	0.298**	-0.098
High Bus	[0.479]	[0.604]	[0.850]	[0.008]	[0.963]	[0.676]	[0.597]	[0.041]	[0.470]
High Gov – Low Gov	0.201	0.183	0.251	0.402**	0.271	0.246	0.105	0.569***	-0.086
Low Bus	[0.163]	[0.207]	[0.140]	[0.014]	[0.102]	[0.145]	[0.487]	[0.000]	[0.533]

Notes: Robust standard errors in parenthesis. Panel A: the dependent variables are the level of trust that the respondent has on each group/organization (scale: 0-10). Regarding the treatments, the *Low Business & Low Government* group is the omitted group. In Panel B we present the linear combination of certain treatment coefficients and p-values (in brackets) for the test of whether these linear combinations are equal to 0. High Bus – Low Bus | High Gov is the difference between the treatment coefficient High Business & High Government and Low Business & High Government; High Bus – Low Bus is the weighted average of High Bus – Low Bus | High Gov and High Bus | Low Gov (analogous for other treatment groups). *, **, *** indicate statistical significance at the 10%, 5% and 1% levels, respectively. We considered the sample of people that spent at least 1.5 minutes in the supplementary survey (not considering the time spent in the treatment windows) and at least three seconds in every treatment.

Figure A.1: Beliefs that taxes on the rich are too low and trust in business and government (executive branch of the federal government)



Notes: Each bar represents the percentage of respondents (unweighted) that answered “Much too low” or “Too low” within a given group to the following question: “Generally, how would you describe taxes in America today.. We mean all taxes together, including social security, income tax, sales tax, and all the rest. A. First, for those with high incomes, are taxes . . .”. The group “Distrust Bus. & Distrust Gov.” is formed by those that answered “Hardly any” to the questions: “I am going to name some institutions in this country. As far as the people running these institutions are concerned, would you say you have a great deal of confidence, only some confidence, or hardly any confidence at all in them? B. Major companies; E. Executive branch of the federal government”. Similarly, those in the group “Trust Bus. & Distrust Gov.” answered “Hardly any” regarding “Major Companies” and “A great deal” regarding “Executive branch of the federal government”. Those in the group “Trust Bus. & Distrust Gov.” answered “A great deal” regarding “Major Companies” and “Hardly any” regarding “Executive branch of the federal government”. Those in the group “Trust Bus. & Trust Gov.” answered “A great deal” regarding “Major Companies” and “Executive branch of the federal government”. Source of data is the General Social Survey 1987, 1996, 2006, 2008 and 2016.

A.2 Main survey questionnaire

You are being asked to take part in a survey being done by a group of researchers from Harvard University that will help us learn more about the relationship between politics and government in America.

The survey will take you about 10 minutes. Please select the link below to complete the survey. At the end of the survey, you will receive a code to paste into the box below to receive credit for taking our survey.

If you have any questions, please contact us at rditella@hbs.edu. The survey is anonymous, and no one will be able to link your answers back to you. Please do not include your name or other information that could be used to identify you.

Survey link:

Code:

1. Gender

- (a) Male
- (b) Female
- (c) I'd prefer to supply my own response:

2. Age

3. Race (select all that apply)

- (a) White
- (b) Black
- (c) Hispanic or Latino
- (d) Asian
- (e) Other

4. In which state do you currently reside?

- Alabama
- Alaska
- Arizona
- Arkansas
- California
- Colorado

- Connecticut
- Delaware
- District of Columbia
- Florida
- Georgia
- Hawaii
- Idaho
- Illinois
- Indiana
- Iowa
- Kansas
- Kentucky
- Louisiana
- Maine
- Maryland
- Massachusetts
- Michigan
- Minnesota
- Mississippi
- Missouri
- Montana
- Nebraska
- Nevada
- New Hampshire
- New Jersey

- New Mexico
- New York
- North Carolina
- North Dakota
- Ohio
- Oklahoma
- Oregon
- Pennsylvania
- Puerto Rico
- Rhode Island
- South Carolina
- South Dakota
- Tennessee
- Texas
- Utah
- Vermont
- Virginia
- Washington
- West Virginia
- Wisconsin
- Wyoming
- I do not reside in the United States

5. Which category best describes your highest level of education?

- (a) Eighth Grade or less

- (b) Some High School
- (c) High School degree/ GED
- (d) Some College
- (e) 2-year College Degree
- (f) 4-year College Degree
- (g) Master's Degree
- (h) Doctoral Degree
- (i) Professional Degree (JD, MD, MBA)

6. What is your current employment status?

- (a) Full-time employee
- (b) Part-time employee
- (c) Self-employed or small business owner
- (d) Unemployed and looking for work
- (e) Student
- (f) Not in labor force (for example: retired, or full-time parent)

7. Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?

- (a) Need to be very careful (0)
- (b) Most people can be trusted (10)

8. Please think about poor people in the US: Do you think they are poor mainly because (choose the most important reason)

- (a) they did not make an effort
- (b) they lacked opportunities
- (c) they were unlucky

9. Please think about rich people in the US: Do you think they are rich mainly because (choose the most important reason)

- (a) they made an effort
- (b) they were born into rich families
- (c) they stole money
- (d) they had good luck

TREATMENT 1

10. Government officials regularly have private meetings with business people to discuss matters of mutual interest. Some argue that such meetings are helpful because they allow the exchange of useful information between government and business and the design of more efficient regulation for complex areas. Critics, on the other hand, argue that these meetings are harmful because they create the opportunity for undue influence, lobbying and the exchange of bribes. In your view, what goes on at these meetings?

- (a) Mainly exchange of bribes for favors (0)
- (b) Mainly exchange of useful information (10)

11. There are some recent proposals to increase government regulations on firms in the US. How likely is it that you would support these type of proposals?

- (a) Very unlikely (0)
- (b) Somewhat unlikely (3-4)
- (c) Somewhat likely (6-7)
- (d) Very likely (8-10)

12. Here are some things the government might do for the economy. Please show which actions you are in favor of and which you are against. (0=Strongly against, 2-3=Against, 5=Neither in favor nor against, 7-8=In favor, 10=Strongly in favor)

(a) Control of wages by law

(b) Control of prices by law

TREATMENT 2

13. Going back to the topic of meetings (between government officials and business people), in the political arena we can find a wide range of views. Some politicians argue strongly in favor of these meetings while others argue strongly against them. Which type of politician are you more likely to support?

(a) A politician that is against allowing these meetings (0)

(b) A politician that is in favor of allowing these meetings (10)

14. Now I'd like you to tell me your views on two issues. How would you place your views on this scale? 0 means you agree completely with the statement on the left; 10 means you agree completely with the statement on the right; and if your views fall somewhere in between, you can choose any number in between

(a) Competition is good. It stimulates people to work hard and develop new ideas (0)

(b) Competition is harmful. It brings out the worst in people (10)

(a) People should take more responsibility to provide for themselves (0)

(b) The government should take more responsibility to ensure that everyone is provided for (10)

15. I'm going to read off one thing that people sometimes say about a democratic political system. Could you please tell me if you agree strongly, agree, disagree or disagree strongly?

(a) In democracy, the economic system runs badly (0=Disagree strongly, 3-4=Disagree, 6-7=Agree, 10=Agree strongly)

16. Some people think it is better to give discretion to policymakers to decide how much regulation to impose on the different sectors of the economy (e.g., how much regulation to impose on banks, on energy companies, etc). What do you think?

(a) Yes, I think it is a good idea to leave them discretion to decide on the proper amount of regulation for each sector

(b) No, I don't want them to have discretion; I prefer the economy to have less regulation overall

(c) No, I don't want them to have discretion, I prefer the economy to have more regulation overall

17. Now we would like to ask you about the income tax rates¹ that you think different people should pay. The income tax rate is the percentage of your income that you pay in federal income tax. For example, if you earn \$30,000 and you pay \$3,000 in income taxes, your income tax rate is 10%. Please use the sliders below to tell us how much you think each of the following groups should pay as a percentage of their total income.

(a) The top 1% (richest)

(b) The next 9% (1% of households earn more than them, but 90% earn less)

(c) The next 40% (10% earn more than them, but 50% earn less)

(d) The bottom 50% (poorest)

18. What was the role of fraud during the 2008 financial crisis in the US? Most analysts agree that there was a bubble as a result of excessive risk-taking in financial markets. But those analysts differ in the extent to which they believe fraudulent practices were involved. Which comes closest to your opinion?

¹We consider only the Federal income tax, which is a tax on household income. If you receive a regular paycheck, this tax is automatically taken out of your pay. When you file a federal tax return each year, you calculate the exact amount you owe, and you get a tax refund from the federal government if you paid more than you owe. To keep things simple, we do not include other taxes such as social security taxes, state income taxes or sales taxes.

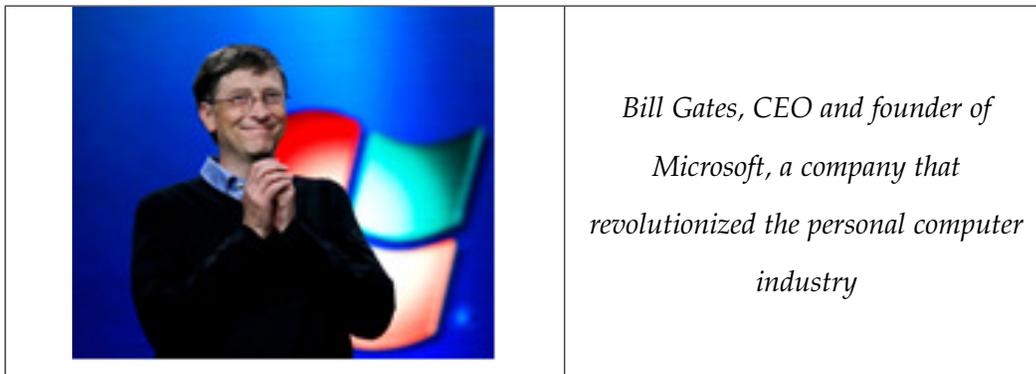
- (a) There was some fraud but this did not cause the crisis.
- (b) There was a lot of fraud, but there was so much risk-taking that the crisis would have happened anyway.
- (c) There was a lot of fraud and it was a central cause of the crisis.

19. In the last election, where did you stand politically?

- (a) Supported Obama
- (b) Center (but leaning Obama)
- (c) Center (but leaning Romney)
- (d) Supported Romney

Treatment High Business

American business people are amongst the most successful in the world. Some of the most famous include Bill Gates (founder and CEO of Microsoft) and Steve Jobs, (founder of Apple, NeXT and Pixar), who have revolutionized the technology industry. In several other areas, such as biotechnology, entertainment, medical devices, and high-end machinery, US business people have also been at the forefront of innovation.



Why do you think American business people have been so successful?

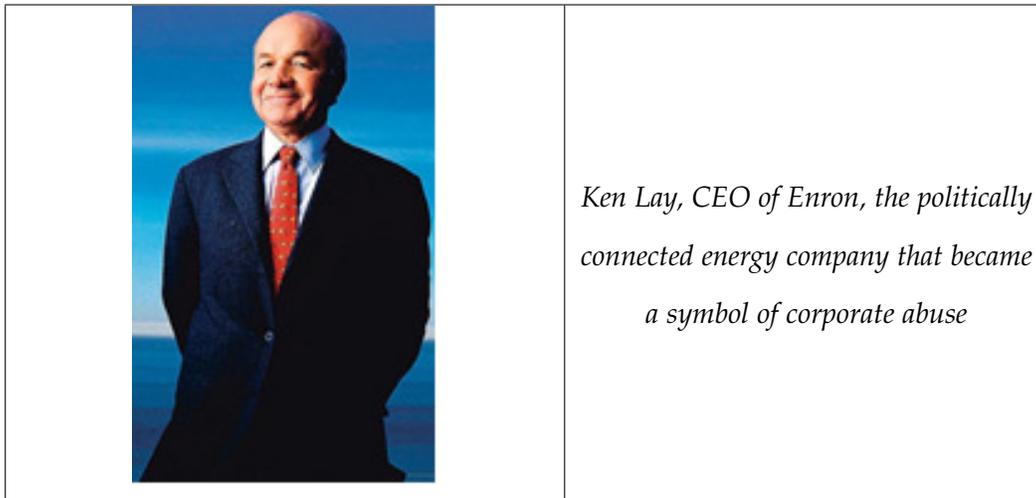
1. It is due to the system: business people in the US are encouraged to work hard and can gain money and prestige by creating truly good products.

2. It is a combination of the system interacting with exceptional individuals, amplified by the availability of capital that allows the successful to expand their business.
3. It is due to the individuals: there are remarkable business people in the US, who are exceptionally creative and naturally hard working.

Treatment Low Business

American business people have been involved in some major scandals over the years.

Some of the most famous include Bernie Madoff (a Wall Street financier who was able to swindle investors for nearly 20 years) and Ken Lay (the former CEO of failed energy giant Enron who lobbied to obtain regulatory exemptions and government contracts). In several other areas, such as construction and medical supplies, there is also evidence of significant wrongdoing.

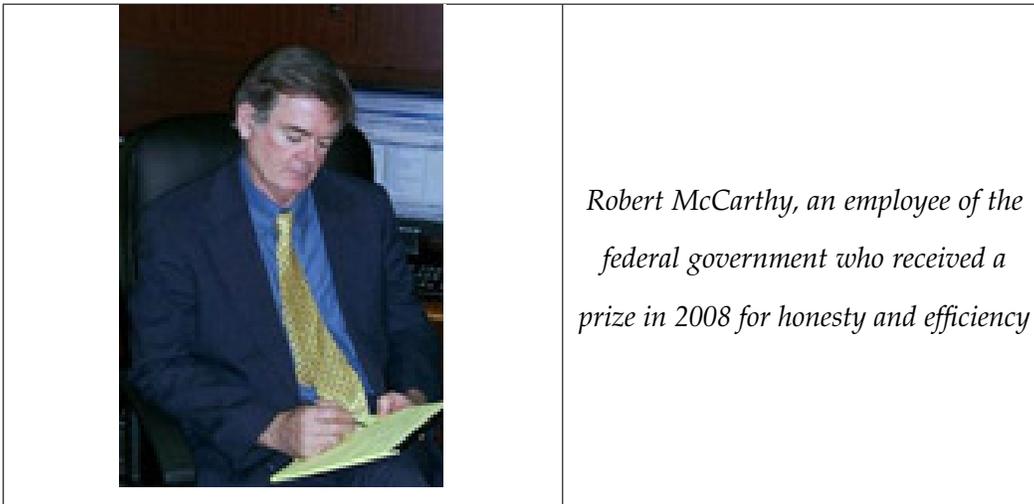


Why do you think there has been so much wrongdoing in American business?

1. It is due to the system: business people in the US are encouraged to focus on profits and can gain lots of money by getting favors from regulators and politicians.
2. It is a combination of the system interacting with greedy individuals, amplified by the availability of capital that allows the dishonest to hide their actions.
3. It is due to the individuals: there are business people in the US that are remarkably greedy and dishonest.

Treatment High Government

American policymakers and institutions of government are some of the most successful in the world. There are several famous cases of government officials who are dedicated and honest (one example is Robert McCarthy who helped improve the administration of two large federal programs). The US government is consistently ranked as one of the most honest and efficient in the world (for example, according to indices constructed by the World Bank, the US is one of the top countries in terms of Regulatory Quality and Control of Corruption).



Why do you think the US government is so much more efficient and honest than the governments in other countries?

1. It is a question of incentives: officials in the US can have a long and well-rewarded career in government by being honest and efficient. The temptations are not worth their while.
2. It is due to the existence of independent checks: the American judiciary system has a long tradition of protecting the rule of law and combating corruption.

Treatment Low Government

American policymakers and institutions of government have been involved in some major scandals over the years. There are several famous cases of government officials

involved in major corruption scandals (one example is Sal DiMasi, who had a long career in government in spite of extorting bribe payments from several businesses, including one business owned by IBM). There are several other examples of significant wrongdoing in government.



Why do you think so much wrongdoing takes place in American government?

1. It is a question of incentives: government officials in the US can gain large amounts of money extracting payments from firms that want to comply with all existing regulations. The temptations are just too profitable.
2. It is due to the lack of effective checks: the legal system has so many loopholes that corrupt officials can defend themselves in very effective ways.

A.3 Supplementary survey questionnaire

You are being asked to take part in a survey being done by a group of researchers from Harvard University that will help us learn more about the relationship between politics and government in America.

The survey will take you about 6 minutes. Please select the link below to complete the survey. At the end of the survey, you will receive a code to paste into the box below to receive credit for taking our survey.

If you have any questions, please contact us at rditella@hbs.edu. The survey is anonymous, and no one will be able to link your answers back to you. Please do not include your name or other information that could be used to identify you.

Survey link:

Code:

1. Gender

- (a) Male
- (b) Female
- (c) I'd prefer to supply my own response:

2. Age

3. Race (select all that apply)

- (a) White
- (b) Black
- (c) Hispanic or Latino
- (d) Asian
- (e) Other

4. In which state do you currently reside?

- Alabama
- Alaska
- Arizona
- Arkansas
- California
- Colorado

- Connecticut
- Delaware
- District of Columbia
- Florida
- Georgia
- Hawaii
- Idaho
- Illinois
- Indiana
- Iowa
- Kansas
- Kentucky
- Louisiana
- Maine
- Maryland
- Massachusetts
- Michigan
- Minnesota
- Mississippi
- Missouri
- Montana
- Nebraska
- Nevada
- New Hampshire
- New Jersey

- New Mexico
- New York
- North Carolina
- North Dakota
- Ohio
- Oklahoma
- Oregon
- Pennsylvania
- Puerto Rico
- Rhode Island
- South Carolina
- South Dakota
- Tennessee
- Texas
- Utah
- Vermont
- Virginia
- Washington
- West Virginia
- Wisconsin
- Wyoming
- I do not reside in the United States

5. Which category best describes your highest level of education?

- (a) Eighth Grade or less

- (b) Some High School
- (c) High School degree/GED
- (d) Some College
- (e) 2-year College Degree
- (f) 4-year College Degree
- (g) Master's Degree
- (h) Doctoral Degree
- (i) Professional Degree (JD, MD, MBA)

6. What is your current employment status?

- (a) Full-time employee
- (b) Part-time employee
- (c) Self-employed or small business owner
- (d) Unemployed and looking for work
- (e) Student
- (f) Not in labor force (for example: retired, or full-time parent)

7. Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people?

- (a) Need to be very careful (0)
- (b) Most people can be trusted (10)

8. Please think about poor people in the US: Do you think they are poor mainly because (choose the most important reason)

- (a) they did not make an effort
- (b) they lacked opportunities
- (c) they were unlucky

9. Please think about rich people in the US: Do you think they are rich mainly because (choose the most important reason)

- (a) they made an effort
- (b) they were born into rich families
- (c) they stole money
- (d) they had good luck

TREATMENT 1

10. Government officials regularly have private meetings with business people to discuss matters of mutual interest. Some argue that such meetings are helpful because they allow the exchange of useful information between government and business and the design of more efficient regulation for complex areas. Critics, on the other hand, argue that these meetings are harmful because they create the opportunity for undue influence, lobbying and the exchange of bribes. In your view, what goes on at these meetings?

- (a) Mainly exchange of bribes for favors (0)
- (b) Mainly exchange of useful information (10)

11. There are some recent proposals to increase government regulations on firms in the US. How likely is it that you would support these type of proposals?

- (a) Very unlikely (0)
- (b) Somewhat unlikely (3-4)
- (c) Somewhat likely (6-7)
- (d) Very likely (8-10)

12. Here are some things the government might do for the economy. Please show which actions you are in favor of and which you are against. (0=Strongly against, 2-3=Against, 5=Neither in favor nor against, 7-8=In favor, 10=Strongly in favor)

(a) Control of wages by law

(b) Control of prices by law

TREATMENT 2

13. Before proceeding to the next set of questions, we want to ask for your feedback about the responses you provided so far. It is vital to our study that we only include responses from people who devoted their full attention to this study. This will not affect in any way the payment you will receive for taking this survey. In your honest opinion, should we use your responses, or should we discard your responses since you did not devote your full attention to the questions so far?

(a) Yes, I have devoted full attention to the questions so far and I think you should use my responses for your study.

(b) No, I have not devoted full attention to the questions so far and I think you should not use my responses for your study.

14. I am going to name nine organizations/groups. For each one, could you tell me how much confidence you have in them: (0= none at all, 3-4= not very much confidence, 6-7= quite a lot of confidence, 10= a great deal of confidence)

(a) Local government

(b) Major Companies

(c) The police

(d) The government (in your nation's capital)

(e) Banks

(f) The press

(g) The armed forces

(h) The courts

(i) Your neighbors

15. At the end of this survey we are going to donate \$200 to charity and we would like you to vote for the organization that should receive the money. The organization with the highest number of votes among the respondents of this survey will receive \$200. There is only a small number of people taking the survey so please take your time to decide. You will be informed of the results within a week.

(a) I vote for Citizens for Tax Justice (an NGO that seeks to **require the wealthy to pay their fair share**; it is primarily concerned with federal tax policy in the US and its mission is to give ordinary people a greater voice in the development of tax laws).

(b) I vote for The American Red Cross (an NGO that seeks to **provide humanitarian help**; it is primarily focused on disaster relief and emergency assistance within the US).

(c) I don't want to vote.

16. Imagine that taxes to the top 1% (richest) of the population increase; what do you think will happen?

(a) The money will be appropriated by corrupt government officials.

(b) The money will be wasted without clear benefits for the population.

(c) The money will be used to fund an increase in useful government spending.

17. In the last election, where did you stand politically?

(a) Supported Clinton

(b) Center (but leaning Clinton)

(c) Center (but leaning Trump)

(d) Supported Trump

18. In previous presidential elections, where did you typically stand?

- (a) Voted republican
- (b) Leaned republican
- (c) Switched depending on the election
- (d) Leaned democrat
- (e) Voted democrat
- (f) Don't know