

Learning Spillovers in Conditional Welfare Programs: Evidence from Brazil*

Fernanda Brollo Katja Maria Kaufmann Eliana La Ferrara

This version: January 2020

Abstract

We study spillovers in learning about the enforcement of *Bolsa Familia*, a program conditioning benefits on children’s school attendance. Using original administrative data, we find that individuals’ compliance responds to penalties incurred by their classmates and by siblings’ classmates (in other grades/schools). As the severity of penalties increases with repeated noncompliance, the response is larger when peers are punished for “higher stages” than the family’s, consistent with learning. Individuals also respond to penalties experienced by neighbours who are exogenously scheduled to receive notices on the same day. Our results point to social multiplier effects of enforcement via learning.

JEL codes: I24, I25, I38, O15

Keywords: CCT programs, enforcement, learning, spillovers, Brazil

*Fernanda Brollo, University of Warwick, CAGE, CEPR, Coventry, CV4 7AL, United Kingdom, f.brollo@warwick.ac.uk; Katja Kaufmann, Mannheim University, briq, CESifo, IZA, L7, 3-5, 68161 Mannheim, Germany, kaufmann@vwl.uni-mannheim.de; Eliana La Ferrara, Bocconi University, IGIER and LEAP, via Roentgen 120122 Milano, Italy, eliana.laferrara@unibocconi.it. We thank the editor and three anonymous referees, Manuela Angelucci, Josh Angrist, Felipe Barrera-Osorio, Esther Duflo, Caroline Hoxby, Robert Jensen, Leigh Linden, Lance Lochner, Marco Manacorda, Alex Mas, Karthik Muralidharan, Phil Oreoupoulos, Johannes Rincke, Jeff Smith, Michele Tertilt, Chris Udry, conference and seminar participants at RES2019, Dartmouth, EUDN, Harvard, Michigan, MIT, NBER, Namur, PSE, TSE, Uppsala, UPF and Yale for helpful comments. Giulia Zane, Simone Lenzu and Emanuele Colonnelli provided excellent research assistance. Kaufmann gratefully acknowledges financial support from the elite program of the Baden-Wuerttemberg Foundation and from the German Research Foundation (DFG) through CRC TR 224 (project C01).

1 Introduction

The enforcement of rules is key in many economically relevant contexts: tax payment, environmental regulations, conditional welfare programs – to name just a few. Enforcement is costly, as it may involve monitoring, policing, auditing and punishing. But enforcement is needed, as compliance behaviour typically depends on the (perceived) probability of punishment in case of noncompliance and the penalties associated with it. The importance of this issue is underscored by the seminal theoretical contributions of Becker (1968) and Stigler (1970).

One question of particular interest is how social interactions influence individuals' compliance behaviour and whether there are spillovers of enforcement. Peer groups can serve as important information transmission networks, which could reinforce or offset the direct effects of enforcement, leading to a compliance rate that is substantially different from what would otherwise be expected.¹ It is therefore critical to understand if enforcement creates spillovers, how large these are, what are the underlying channels and what are the relevant social networks in which spillovers materialize. The existence of such spillovers has obvious implications for the cost effectiveness and the optimal design of enforcement activities.

This paper studies the importance of spillovers in the enforcement of one of the world's largest conditional welfare programs: the *Bolsa Familia* program (BFP) in Brazil. BFP provides a monthly stipend to poor families that is conditional on every school-aged child attending at least 85% of school days every month. The program

¹Indeed, information transmission in informal networks is important in settings as diverse as commercial advertising (“word-of-mouth”), the transmission of information about vaccines, micro-finance and welfare programs (e.g., Banerjee, Chandrasekhar, Duflo and Jackson, 2019). Evolutionary psychologists have also argued that “gossiping gives humans the ability to spread valuable information to very large social networks” (Dunbar, 2003).

is enforced through a system of up to five ‘warnings’ or penalties, which gradually increase in severity with subsequent instances of noncompliance. We investigate whether, how and why an individual’s likelihood of compliance is affected by the fact that her peers experience the enforcement of program conditions.

Monitoring and enforcement had been weak to non-existent until 2007, when the government made a concerted effort to strengthen it (see details in the Online Appendix). It is therefore plausible that, during the first years of the program, beneficiaries formed a prior that conditionality would not be strictly enforced, and only in later years noncompliant families started receiving punishments. This suggests scope for learning about the probability of enforcement both from one’s own experience and from peers. In fact, Brollo, Kaufmann and La Ferrara (2018) find that individuals change their behaviour in response to the penalties that *they themselves* receive, showing that the ex ante perceived probability of enforcement was smaller than one. In this paper we focus on learning *from peers*.

Using original administrative data on compliance behaviour and enforcement of program conditions, we show that a larger fraction of classmates being punished for noncompliance increases individual school attendance. This is true not only for penalties received by the individual’s own classmates, but also for penalties received by classmates of her siblings, who are in other grades or schools.

To investigate the underlying mechanisms, we exploit the fact that families progress through different ‘warning stages’ and receive increasingly severe penalties. When we distinguish between peers who get penalized at a lower warning stage than the family’s own and peers who receive warnings for higher stages, the latter induce a larger increase in individual compliance. This is consistent with

learning because those warnings convey *new* information on the likelihood that the government implements higher order punishments.

Lastly, we show that individuals also respond to *neighbours* experiencing enforcement. In particular, we exploit quasi-random variation in the timing when different beneficiaries are scheduled to cash-in transfers and receive notifications of punishment in a given month. Notifications of penalties are received in the last ten working days of the month and the exact date depends on the last digit of the household head's Social Identification Number (NIS), which is essentially random. We find that individuals respond to penalties of neighbours who are exogenously scheduled to receive notifications on the same day at the local service point, even after controlling for the overall penalties received among neighbours in that month. Again, the response depends on neighbours' punishment stage relative to the family's own. Classmates and neighbours are therefore two important sources of information leading to a social multiplier of enforcement.

In terms of magnitude, we find that a one-standard deviation increase in the fraction of a child's classmates who are warned in a given month leads to a decrease in the child's probability of not complying with school attendance by 1.2% in the month following the warnings and 2.4% in the month after that. The effects are similar when we consider neighbours: if the fraction of families who use the local service point and experience enforcement on the same day increases by one standard deviation, a child's likelihood of noncompliance decreases by 1.5 (3)% in the following month (the month after that). These estimates are very similar when the effects of classmates and neighbours are estimated separately or together, suggesting that the two types of social interaction constitute two largely separate

ways of information transmission. Note that these magnitudes can be considered a lower bound for the total spillovers that arise from enforcement, because they are intention-to-treat effects since the actual interaction or communication patterns are not known, and because we net out the ‘direct’ effects coming from school attendance of peers.

It is worth highlighting what characterizes our paper in terms of richness of our data and methodological approach. In the context of enforcement and compliance with rules, good data are notoriously difficult to find. Our setting allows us to (i) use administrative data which overcomes measurement problems related to self-reported information; (ii) use high frequency data on compliance which helps identifying effects due to timing; and (iii) investigate the role of different types of social interaction (siblings, classmates, neighbours).

Methodologically, while we study social interactions, we are not analyzing ‘peer effects’ in the standard sense. Instead of estimating how an individual’s compliance behaviour is affected by peers’ compliance, we analyze how it is affected by peers’ experience of enforcement. This helps us get around one of the challenges that typically arise in the peer effects literature, namely the reflection problem (Manski, 1993). Two remaining challenges are correlated unobservables and endogenous group membership, and we address them using the following strategies.

To identify effects arising from the interaction between classmates, we rely on the fact that there is quasi-random variation in the delay of warnings and on a strategy which is related in spirit to the approach of partially overlapping peer groups, but adjusted to our context.² In particular, we test if a child’s compliance is affected

²See, e.g. Bramoullé, Djebbari, and Fortin (2009), De Giorgi, Pellizzari, and Redaelli (2010) and Lee, Liu and Lin (2010).

after her peers are warned but not before, and if compliance is affected not only by the penalties received by a child's own classmates (which could reflect correlated shocks at the classroom-grade level), but also by penalties experienced by her siblings' classmates (in different grades and schools).

To identify effects of social interactions between neighbours, we make use of quasi-random variation in group membership generated by the day on which individuals receive notice of penalties, as explained above. By comparing the results based on these two different identification strategies and finding that they are consistent, we provide further evidence on the validity of our approaches.

The effects we estimate have important policy implications. Policymakers could consider strategies that inform people at the school or neighbourhood level about how many beneficiaries have been experiencing enforcement, thereby directly creating learning spillovers. This would allow them to use the tool of enforcement more effectively by generating a bigger impact from every executed instance of enforcement, which is critical given the large costs of enforcement. Also, while our paper focuses on enforcement spillovers in the context of conditional cash transfers, our analysis has broad relevance for a number of welfare programs that embed conditionality. This is important, as governments around the world increasingly rely on 'conditional' programs in areas as diverse as unemployment and social assistance benefits, maternity grants, child support and support for asylum seekers.³

Our paper contributes to several strands of literature. First, there is a literature that estimates behavioural responses to program enforcement: Banerjee, Glenner-

³Welfare programs can have important externalities on society by affecting individual incentives to exert effort and their likelihood to rely on the safety net in the future (e.g., Dahl and Gielen, 2018). Incorporating conditions into welfare programs has been seen as one way to address externalities and improve the political and social support for such programs.

ster and Duflo (2008) in the area of monitoring public workers; Lochner (2007) in the area of crime; Olken (2007) in the area of corruption; Carillo, Pomeranz and Singhal (2017), Dwenger, Kleven, Rasul und Rincke (2016) and Kleven et al (2011) on tax enforcement; and Black, Smith, Berger and Noel (2003) and Van den Berg, van der Klauuw and van Ours (2004) for enforcement in the context of unemployment insurance schemes.⁴

The literature on enforcement *spillovers* is more limited, with the notable exceptions of Rincke and Traxler (2011) and Drago, Mengel and Traxler (2017) on spillovers of the enforcement of TV license fees in Austria; and Pomeranz (2015) and Boning et al. (2018) on spillovers of tax enforcement among firms. Compared to these papers, we study a very different context: the largest conditional cash transfer program in the world. To the best of our knowledge, ours is the first paper to estimate enforcement spillovers with respect to conditionality in welfare programs (and not to tax evasion). This makes a significant difference, because in our setting individuals are substantially poorer and less educated, and enforcement is not the norm (program conditions were not enforced for several years in Brazil). While in principle all individuals (and firms) should have the means for paying the taxes they owe, it is not obvious that poor households have a margin to react when the decision is about school attendance of their children and when what they risk is losing a cash transfer (as opposed to, for example, large fines or going to jail). Also, our analysis contributes to the above literature by investigating the relative importance of different peer groups, by studying dynamic consequences of enforcement (in our context the decision to comply has to be taken every month) and by studying

⁴Since we do not have direct measures of *beliefs* about enforcement, we infer them from individuals' behavior, akin to the approach in Chetty, Friedman and Saez (2013).

information transmission through different types of signals (e.g., different enforcement/punishment stages).

Second, our work relates to a large literature on conditional cash transfers (CCTs). Various studies have shown that CCTs have significant medium to long-term effects, helping to break the intergenerational transmission of poverty.⁵ Scholars have also compared conditional and unconditional (UCT) transfer programs, with Schady and Araujo (2006) and Baird, McIntosh and Özler (2011) finding stronger effects for conditional ones, while Benhassine et al. (2015) finding similar effects of the two types, provided that the unconditional one is ‘labeled’ as education support. Our contribution differs from the CCT versus UCT comparison. In a world where conditionality nominally exists but program recipients anticipate that it will not be enforced, one may fail to find a difference in the impact of the two types of programs. But this does not imply that conditionality, if enforced, would be ineffective. Our paper is precisely an attempt to understand how people update their beliefs about enforcement after seeing peers experience it, and change their behaviour in response.

Like us, DeBrauw and Hoddinott (2011) focus on the enforcement of program conditions and find that it has direct effects on the compliance behaviour of beneficiaries in the context of the Mexican Progresa program. They exploit the fact that

⁵Long-term effects of CCT programs have been shown for example for Colombia (see, e.g., Barrera-Osorio, Linden and Saavedra, 2019); for Mexico (see, e.g., Fernald, Gertler and Neufeld, 2009, and Parker, Rubalcava and Teruel, 2012); for Nicaragua (see, e.g., Barham, Macours and Maluccio, 2017); and for the United States (see, e.g., Aizer, Eli, Ferrie and Lleras-Muney, 2016). In addition there is a large literature on the direct short- to medium-term effects of CCTs, such as, for Brazil, Bourguignon, Ferreira and Leite (2003), De Janvry, Finan and Sadoulet (2011) and Bursztyn and Coffman (2012) who study Bolsa Escola, the predecessor of Bolsa Familia in Brazil, and Bastagli (2009), De Brauw et al (2015a,b), Chioda, de Mello and Soarez (2016) and Rios (2017) who focus on Bolsa Familia.

some beneficiaries who received transfers did not receive the forms needed to monitor school attendance and use matching and fixed effect methods to show that the absence of these forms significantly reduced compliance.⁶ Aside from the difference in context and empirical methodology, our work investigates if the benefits of enforcement are limited to such direct effects, or if there are spillover effects which could imply substantially larger benefits of enforcement.⁷

Lastly, our analysis is related to the literature on social interactions and learning about program features, including Aizer and Currie (2004) on network effects in the utilization of publicly-funded prenatal care; Alatas et al (2016) on network structure and the aggregation of information in a community-based targeting program; Banerjee et al (2013, 2019) on information transmission in the context of a microfinance program; Duflo and Saez (2003) on the transmission of information on retirement plans; and Dahl, Loken and Mogstad (2014) on peer effects via learning in the context of paternity leave.

The remainder of the paper is organized as follows. In section 2 we provide background information on Bolsa Familia, while in section 3 we present the data and descriptive statistics. In sections 4 and 5 we discuss our empirical strategy and results, respectively. Section 6 concludes.

⁶There are a few papers which analyze the enforcement of eligibility criteria, i.e. errors of type one and two in terms of families included in or excluded from the program (e.g., Cameron and Shah, 2014), but there is very little evidence on the enforcement aspect of *conditionalities*.

⁷Papers that analyze peer effects in the context of conditional cash transfer programs but not enforcement spillovers are, for example, Barrera-Osorio, Bertrand, Linden and Perez-Calle (2011) for Colombia; Angelucci et al (2010), Angelucci and De Giorgi (2009) and Bobonis and Finan (2009) for the Mexican program Progresa/Oportunidades.

2 Background Information on Bolsa Familia

The Bolsa Familia Program (BFP) reaches around 14 million Brazilian families, that is 60 million poor people (equivalent to about 30% of the Brazilian population) with an annual budget of over 24 billion reais (USD 11 billion, about 0.5% of GDP). Thus BFP reaches nearly three times as many people and is about three times as large in terms of budget as the well-known conditional cash transfer program Progres/Oportunidades.

BFP was launched by the Brazilian president Inácio Lula da Silva in 2003. Monitoring and enforcement of conditionalities had been weak to non-existent in the first years of the program and was strengthened from 2007 onwards (see Online Appendix B for further details on the history and the targeting of the program).

Benefits. BFP provides two types of benefits, a “base” and a “variable” transfer that depend on family composition and income. The base benefit is provided only to families in extreme poverty (with a monthly per capita income of up to R\$60 –approximately US\$30), regardless of their demographic composition. All families with a per capita income of up to R\$120 receive a “variable” benefit which depends on the number of children in the family (capped at three to avoid incentivizing fertility) and on whether the mother is pregnant or breast-feeding. During our sample period the base benefit amounts to R\$ 60 per month, and the variable benefit to R\$ 20. Benefit amounts and eligibility thresholds are periodically adjusted for inflation. To illustrate the magnitude of the program transfers, a family with three children that is right at the threshold to be classified as extremely poor (family income of R\$ 300) would receive monthly transfers of around R\$120, which amounts to 40% of their total family income.

Conditionality. BFP transfers are conditional on all age-relevant family members complying with requirements in terms of school attendance. Each school-aged child between 6 and 15 has to attend at least 85% of school days each month (absence due to health reasons is justified and does not count). If one or more children fail to meet this requirement in a given month, the family is affected for the whole amount of the transfer, i.e. also for quotas that pertain to other children. This element of “joint responsibility” for children in the same family is a unique feature of BFP, e.g., compared to other well-known CCTs such as Progres/Oportunidades.⁸

Penalties. The consequences of noncompliance vary depending on the historical record of compliance of each family. In the first case of noncompliance the family receives a warning without any financial repercussion. With the second instance of noncompliance, the family receives a second warning and benefits are blocked for 30 days, after which the family receives the accumulated benefit of the previous and the current month. The third and fourth warnings lead to a loss of benefits for 60 days each time and these benefits are never recovered. Finally, after the fifth warning the benefit is canceled and the family loses eligibility. According to the general rules, the family can re-apply to the program 18 months later. The 18 months rule is also used to “reset” a family warning history, e.g., if a family is in warning stage 2 but then complies with the rules for 18 consecutive months, their

⁸In 2008 a new subprogram, Benefício Variável Jovem (BVJ), was added to BFP to increase school attendance of teenagers aged 16 and 17. BVJ has its own set of conditions and warning system and does not affect the family’s warning stage in the main program (i.e. for the children aged 6 to 15). We only have data for a few months on BVJ, hence we focus on children aged 6 to 15. BFP also has conditions related to health behavior, such as health check-ups for pregnant women and vaccinations for children below age 5. Based on our data, those conditions are rarely enforced and we do not observe individual responses to enforcement of those conditions (in contrast to having monthly data on children’s school attendance). Therefore we focus on children aged 6 to 15 and analyze responses to enforcement of school attendance conditions.

warning history goes back to stage zero.

Timing of penalties. Each month families can withdraw their transfer money at a local service point with a Bolsa Familia “electronic benefit card”.⁹ In case of noncompliance, at the time of withdrawal the family receives a warning message which specifies the penalty that the family incurs. To avoid bottlenecks at the local service point when families pick up their BFP transfers, each family is assigned a pre-specified date each month starting from which the family can withdraw the transfer money (and contextually receive the warning). The exact date in each month is determined by the last digit of the social security number of the legally responsible adult of the family, which is basically random, as we will show. In particular, the ten digits (0,1,...,9) correspond to the ten weekdays in the last two weeks of the month.¹⁰ We will exploit this randomness to generate a quasi-exogenous peer group of neighbours cashing-in their benefits and receiving warnings on the same day and to thereby identify the causal impact of neighbours experiencing enforcement.

Information about the BFP program. Families are well informed about transfer amounts, conditionalities and the existence of penalties: these aspects are widely and regularly publicized on TV, radio and newspapers in Brazil, and are spelled out in a booklet issued to each beneficiary family (*Agenda de Compromissos*). In case of noncompliance, at the time of withdrawal of the transfer money,

⁹Local service points are either Caixa branches or Caixa “correspondents”. Caixa Econômica Federal, a savings and credit union, is the government agency responsible for transferring BFP benefits to the beneficiary families. Caixa has a large network of banking correspondents, which are commercial establishments with a different business focus, used to expand access to remote and particularly poor areas (Kumar et al., 2006).

¹⁰Every family receives a benefit calendar with the information on when they can pick up the transfer in the different months.

the family receives a message that reports the family's warning stage, the month of noncompliance to which the warning refers, the name(s) of the child(ren) who failed to comply and which type of warning the family might receive in the next instances of noncompliance.

An example of the first paragraph of a warning message for a family receiving their first warning is: "Your family has not complied with the conditionalities of Bolsa Familia for the first time. At this moment the payment will not be blocked. But if you fail to comply again your benefit payment may be blocked, suspended or even cancelled."¹¹ Note that the text does not say that the payment will be suspended with probability one, hence there is a degree of uncertainty about the actual enforcement of conditionality which is embedded in the system. This uncertainty is further reinforced by the experience of lax enforcement in the early years of the program.¹²

Implementation. The implementation of this conditionality scheme involves different actors. First of all, children's attendance is recorded by school teachers.¹³ The school sends the attendance lists of students to the municipality, reporting the exact fraction of school days attended in case attendance was below 85%, otherwise only reporting that the student complied. Each municipality collects the lists and sends them to the Ministry of Education (MEC), which determines whether the

¹¹In addition to this short paragraph, the warning message reports which child failed to meet the attendance requirement and in which month, and repeats the general rules of BFP in a clear and salient manner. It briefly mentions again the program conditions (i.e. of at least 85% of school attendance of all school aged children) and which instance of noncompliance may lead to which type of penalty.

¹²In 2008/2009 85-90% of instances of noncompliance were enforced. The presence (or lack of) enforcement, however, is not the source of variation we use in our analysis. We rely on quasi-random variation in the timing of warnings and on overlapping peer groups, as discussed in detail in section 4.

¹³We will discuss the concern of teacher underreporting in the robustness section 5.4.

family as a whole complied or not in a given month, i.e. whether all children between 6 and 15 attended at least 85% of school days. MEC sends a detailed report to the Ministry of Social Development (MDS), which establishes which warning the family should receive in case of noncompliance and whether the family is entitled to the transfer for that month (based on the warning stage reached). The MDS sends this information to the Caixa Econômica Federal, a savings and credit union that transfers the benefit amount to the account of the family if the family is entitled to receiving the transfer for that month.

These different steps of the process involve time and lead to a ‘delay’ in terms of the month in which the warning is received compared to the month of noncompliance (for an illustration, see Appendix Figure A.1.¹⁴ There is substantial variation in the extent of delay, which we make use of in some of our robustness checks. During our sample period the delay between missed attendance and receipt of the warning ranges from 2 to 6 months, with a mean of 3.8 and a median of 3 months.¹⁵

3 Data

We make use of a unique dataset that we assembled combining several sources of administrative data of the Brazilian Ministry of Social Development (MDS) and the Brazilian Ministry of Education (MEC). The first source is the household registry (*Cadastro Unico*) held by MDS, which contains socio-economic characteristics of

¹⁴If a family fails to comply once more between the last instance of noncompliance and before receiving the corresponding warning, the family will receive the additional warning for this later, i.e. in the order that the noncompliance occurred. Also the warning message states to which month of noncompliance the warning message refers.

¹⁵As we will show, the variation in delay is orthogonal to household and child characteristics and is entirely driven by time and area effects (see Appendix Table A.1).

all BFP beneficiary households.¹⁶ We have information on the universe of households enrolled in the program during calendar years 2008 and 2009 in the Northeast of Brazil, one of the poorest and largest regions in the country, comprising 30% of the Brazilian population, and with more than half of its inhabitants living in poverty. To conduct our analysis we extracted a 10% random sample, yielding a total of 478,511 households in the program.¹⁷ However, for the construction of the peer variables we use the full population of classmates and neighbours.

A second dataset from MDS contains monthly records of school attendance during 2008-09 for each child monitored for the attendance conditionality, i.e., each child aged 6 to 15. We know whether the child complied with the attendance condition in a given month, and in case of noncompliance we know the exact fraction of days attended below the threshold of 85%.

A third dataset from MDS contains monthly information on penalties (“warnings”), i.e. on the five different types of penalties of increasing severity discussed in section 2, which allows us to create the complete warning history of each family. In particular, we know in which month the family received a warning, the month in which attendance failed to meet the threshold and that gave rise to that particular warning, and the warning stage in which the family is at any given point in time.

The fourth administrative dataset we use is payroll data, which contains monthly information on the benefits that the family received and whether the benefit was blocked or suspended in a given month. All four of these datasets can be linked

¹⁶For each household member the *Cadastro Unico* comprises information on age, gender, race, marital status, education and employment status and for the household as a whole it reports house property, garbage collection, electricity etc.

¹⁷To improve our precision in estimating the effects of enforcement, i.e. penalties, we over-sampled households that received a penalty at least once during our sample period, and use regression weights to correct for this.

through the social security number (*Número de Identificação Social* - NIS) of the child or of the household member legally responsible for the child.

Importantly, our data allows us to investigate three different types of social interaction (“peer groups”): siblings, classmates and neighbours. With the help of the household registry data we can identify siblings using the identifier of the mother. Moreover, these data allow us to identify neighbours who are BFP beneficiaries, since it contains information on the eight digit zip code (*Código de Endereçamento Postal* - CEP) of the neighbourhood where the household lives.

To identify classmates, we complement the above administrative records on BFP with the School Census compiled by the Ministry of Education, which contains information on all children who are enrolled in a given school, grade and class. We merge the School Census with the BFP administrative data using the NIS code of the child when available in the School Census and otherwise based on area code, school code, grade, full name and date of birth of the child.¹⁸ This allows us to identify the peers who are in the same school, grade and class for each child and who are also recipients of BFP.

In addition to being able to identify an individual’s siblings, classmates and neighbours, we can also compile their (and their families’) attendance, warning and transfer history, so that we can analyze whether an individual is affected by her peers experiencing enforcement.

[Insert Table 1]

Table 1 presents summary statistics of the main variables of interest. A full set of summary statistics for all variables (including controls) is provided in Appendix

¹⁸We are able to match 68% of individuals, 41% based on their NIS.

Table A.2. The outcome variable we use in our analysis is whether the individual fails to comply with BFP rules, i.e., the child attends less than 85% of school days in a given month. We refer to this variable as “noncompliance”. Panel A of Table 1 shows summary statistics on variables related to individuals’ classmates, while panel B refers to the classmates of individuals’ siblings and panel C to neighbours.

Table 1 shows that an individual’s likelihood of noncompliance is 5.1% in any given month. The average child who is a BFP recipient in our sample has 17 other BFP recipients in her class, 52 in the same grade and 279 in the same school. About 1.4% of one’s classmates who are BFP recipients (‘peers’) receive a warning message in a given month (for any of the five warning stages).

Panel B of Table 1 focuses on the sample of individuals with at least one sibling and shows that the maximum fraction of siblings’ classmates who are warned is 1.4%, while the maximum fraction warned of classmates of siblings at other schools (of the opposite gender) is 1.3 (0.9)%.

Panel C of Table 1 shows that the average individual lives in a neighbourhood with 4042 other individuals who are BFP beneficiaries, of which a tenth, i.e. 404, are scheduled to cash-in their transfer –and receive a penalty if applicable– on the same day of the month. About 0.6% of neighbours are warned and 0.6% are warned on the same day.¹⁹

¹⁹The two means are the same as the neighbours scheduled to receive notice on the same day are essentially a random subset of the neighbours.

4 Empirical Strategy

In this section we discuss our empirical strategies to identify whether there are spillover effects from the enforcement of program conditions on peers. In particular, we consider three types of social interactions: among siblings, classmates and neighbours.

To identify spillover effects from siblings and classmates we rely on a strategy which is related to the one of partially overlapping networks.²⁰ We exploit the fact that the classmates of an individual's siblings are not her direct peers and are therefore not affected by shocks that directly hit the family (e.g. illness) or that are specific to the child's grade. In particular, we analyze whether an individual's compliance behaviour is affected not only by the penalties that her direct classmates receive, but also by penalties experienced by her siblings' classmates in other grades or schools (while controlling for her own and her direct classmates' penalties, which accounts not only for direct shocks specific to the individual's class and school but also for shocks at a higher geographic level).

To identify spillover effects from neighbours, we make use of quasi-random variation in the set of neighbours who cash in the benefit on the same day (because they have the same last digit of the NIS code). This creates exogenous peer groups conditional on living in the same neighbourhood.

Finally, we discuss how we shed light on the underlying mechanisms behind the spillover effects, in particular how we test for the importance of information

²⁰For a theoretical discussion of this strategy, see Bramoullé, Djebbari and Fortin (2009), De Giorgi, Pellizzari and Redaelli (2010) and Lee, Liu and Lin (2010). For examples of recent applications, see, e.g., Nicoletti, Salvanes and Tominey (2018) who investigate the importance of family spillovers by instrumenting siblings' behaviour via the behaviour of siblings' neighbours.

transmission and learning (i.e., belief updating).

4.1 Identification: Learning from Classmates

Children who are in the same class as one’s own (or the parents of these children) are likely to be key sources of information regarding the implementation of the program, e.g., by reporting that they have received a penalty because they failed to attend school. To explore these effects we rely on individual children as the unit of observation and construct different sets of peers, comprising the child’s own classmates and the classmates of a child’s siblings.

Identifying the impact of the warnings received by classmates is a non-trivial issue. To discuss our empirical strategy it is useful to start from a simple specification, which we will not estimate but which helps highlight identification challenges. Consider the following model:

$$Y_{iht} = \alpha PEERWARN_{iht} + \sum_{k=1}^5 \gamma_k WS_{ht}^k + \delta \mathbf{X}_{iht} + \mathbf{D}_t + \mathbf{D}_i + \epsilon_{iht} \quad (1)$$

where i denotes the child, h the household, t the month; Y is an indicator for whether child i failed to comply in month t , WS^k is a set of dummies to denote if the household is in warning stage k ; \mathbf{X} is a vector of child level (time varying) controls including age dummies and number of brothers and sisters in different age brackets (6-10, 11-15, 16-18); \mathbf{D}_t denotes month and year fixed effects; \mathbf{D}_i denotes child fixed effects and ϵ is the error term. Controlling for the warning stage WS_{ht}^k that the family reached is important in a regression explaining the decision of noncompliance, because it directly determines the costs of the next instance of

noncompliance with program conditions. The inclusion of the child fixed effects allows us to control for time-invariant unobservables (among which could be the child's ability or intrinsic motivation to attend school, etc.). This implies that in our analysis we exploit variation within individuals over time.

The key regressor of interest in (1) is $PEERWARN_{iht}$, which is the fraction of i 's classmates who receive any type of warning in month t .²¹ Suppose we found –as we do– a negative correlation between a child's noncompliance in a given month and the fraction of peers who are warned ($\widehat{\alpha} < 0$). Before we can interpret this correlation as learning we need to address several identification challenges.

Correlated shocks. The first threat to identification stems from correlated shocks that may directly affect a student and her peers, thus inducing a correlation between $PEERWARN_{iht}$ and ϵ_{iht} in equation (1). Consider for example an economic shock leading to an increase in the opportunity cost of schooling in the area where individual i lives. In response to such a shock, both individual i and her peers would be more likely not to comply and thus more likely to receive a warning.²²

If the shock was persistent, it would induce individual i and her peers to still not comply several months later, when the warning is received. However, this type of mechanism would generate a *positive*, not negative, correlation between a child's noncompliance and her peers' warnings ($\widehat{\alpha} > 0$), while what we find in the data is a negative correlation ($\widehat{\alpha} < 0$).

If, on the other hand, the shock was mean-reverting, this would lead to a nega-

²¹Here we pool the five different types of penalties that classmates receive, while further below we investigate the effect of different types of penalties that peers incur to learn more about the underlying mechanisms.

²²To address the concern that the arrival of an individual's own warning might be correlated with the arrival of her peers' warnings, we always control for own warnings and analyze if peers' warnings have an additional effect.

tive correlation between a child's noncompliance and her peers' warnings and could thus be confounded with the learning interpretation we offer. To address this concern, we include the lead of peers' warnings as a falsification test, i.e., we add $PEERWARN_{ih,t+1}$ among the controls in equation (1). Our reasoning exploits the time lag from the moment in which noncompliance occurs and the moment in which the warning arrives (as discussed in section 2, this delay has a median value of 3 months, a mean of 3.8 months and the delay is not correlated with individual or family characteristics but only with when the noncompliance took place, see Appendix Table A.1). If a child's attendance increases because of mean-reversion after the initial instance of noncompliance, there is no reason why the attendance would start reverting to the mean with the exact same delay as the warnings of her peers, which vary between 2 and 6 months. In the case of mean-reversion, one would typically expect changes in attendance to start occurring *before* the arrival of the warnings, which implies that we should find a negative and significant coefficient on $PEERWARN_{ih,t+1}$. Failing to find an effect of the lead variable would be hard to reconcile with the interpretation of mean-reverted shocks: it would mean that the initial shock, which led to noncompliance, starts reverting exactly, say, four months afterwards (when peers happen to receive the warning) but not three months afterwards.

An additional strategy we employ to deal with grade or school-specific correlated shocks is to exploit warnings received not only by individual i 's own classmates, but by the classmates of i 's *siblings*. Since siblings typically have different ages, the effect of warnings received by siblings' peers should not reflect class or grade-specific shocks. This should be particularly true for a second specification

we use, in which we focus on peers of siblings in *other schools* (being in different schools is mostly due to siblings' age difference). The warnings received by students of those schools should not be correlated with shocks experienced by child *i*'s own school, conditional on the warnings of the child's family and of her own peers, which also controls for shocks at a higher geographic level (see, e.g., Nicoletti, Salvanes and Tominey, 2018).

On the other hand, warnings received by siblings' peers contain information about the quality of BFP enforcement. We include the *maximum* fraction of peers who got warned across child *i*'s siblings, as the strongest signal of enforcement.²³

Conventional peer effects. A separate concern relates to direct or 'conventional' peer effects. Brollo, Kaufmann and La Ferrara (2018) show that once an individual receives a warning, she reacts by reducing the likelihood of noncompliance, i.e., increasing attendance. This means that when a child's classmates receive warnings, they will attend school more in response to their own warnings. The child may then start attending more because she observes her peers doing so. This response would imply a spillover effect of enforcement on peers' school attendance, but cannot necessarily be interpreted as learning about enforcement: it might be due to learning about the benefits of schooling or to the individual preferring to attend school with more peers.

To identify "learning about enforcement" we pursue two approaches. First, we directly control for the fraction of child *i*'s classmates who fail to comply with program conditions, to analyze if *warnings* of *i*'s peers have an independent effect

²³The more families in a class receive a warning, the more likely it is that this becomes a topic of discussion among children and parents of beneficiary households.

on i 's likelihood to fail.²⁴ Second, we analyze whether the probability that child i fails to comply decreases when her *siblings'* peers (possibly at a different school) get warned. Since those peers are not child i 's classmates (or not even in child i 's school), the direct effect of their school attendance should not be important. On the other hand, warnings received by siblings' peers contain relevant information about the strictness of enforcement. Lastly, we show further evidence of learning, as opposed to imitation, by making use of the relationship between individual's own and peers' warning stage, as discussed at the end of this section.

Sibling peer effects. While the exercise of analyzing the effect of *siblings'* peers can help to rule out peer effects between classmates as the main mechanism, a remaining concern might be sibling peer effects, i.e. the fact that an individual might respond to warnings of a sibling's classmates not because of the information content of the warning, but because the sibling responds to her classmates' increased attendance and the individual in turn responds to her sibling's increased attendance. Such sibling peer effects are likely to be stronger between siblings of the same gender, since they tend to interact more (see, e.g., Dunifon et al (2017)). Therefore we test whether the effects of siblings' classmates' warnings are less strong when the sibling is of the opposite gender as the individual herself.

Timing. Families learn about the fact that they have been punished when they go to the local Caixa branch to cash in their transfer. They are scheduled to do so over the last ten weekdays of the month, that is, on average five weekdays before the end of the month. Thus families have little time to respond to classmates' warnings

²⁴This may lead to an underestimate of the "true" learning effect, if part of the learning about enforcement happens through observing one's peers' higher attendance.

in the same month when they are received. For this reason we include a variable for classmates' warnings in the month following the warning and in addition one month later (i.e. as a first- and second-order lag), since it is ex-ante unclear how persistent the effects are. The variable for classmates' warnings in the same month serves as a placebo test to rule out mean-reverting shocks.

To sum up, the three specifications we estimate are the following. The benchmark specification to estimate the effect of classmates' warning is:

$$Y_{iht} = \sum_{n=-1}^1 \alpha_n PEERWARN_{ih,t-1+n} + \sum_{k=1}^5 \gamma_k WS_{ht}^k + \delta \mathbf{X}_{iht} + \mathbf{D}_t + \mathbf{D}_i + \epsilon_{iht} \quad (2)$$

To net out 'direct', or conventional, peer effects we add the fraction of noncompliant classmates:

$$Y_{iht} = \sum_{n=-1}^1 \alpha_n PEERWARN_{ih,t-1+n} + \sum_{n=-1}^0 \beta_n PEERFAIL_{ih,t+n} + \sum_{k=1}^5 \gamma_k WS_{ht}^k + \delta \mathbf{X}_{iht} + \mathbf{D}_t + \mathbf{D}_i + \epsilon_{iht} \quad (3)$$

Finally, to estimate spillovers from siblings' classmates we estimate:

$$Y_{iht} = \sum_{n=-1}^1 \alpha_n PEERWARN_{ih,t-1+n} + \sum_{n=-1}^1 \beta_n MaxPEERWARN_{ih,t-1+n} + \sum_{k=1}^5 \gamma_k WS_{ht}^k + \delta \mathbf{X}_{iht} + \mathbf{D}_t + \mathbf{D}_i + \epsilon_{iht} \quad (4)$$

where $PEERWARN_{iht}$ is defined as above, that is the fraction of i 's classmates who receive any type of warning in month t . $PEERFAIL_{iht}$ is the fraction of i 's

classmates who are BFP recipients and fail to comply in month t . $MaxPEERWARN_{iht}$ is the maximum fraction of classmates who got warned in month t among i 's siblings. This variable is constructed alternatively from all of i 's siblings or from siblings who attend a different school than i or from siblings who are of the opposite gender.

Given the panel nature of our data, we estimate equations (2), (3) and (4) and all other regressions using a linear probability model and clustering the standard errors at the household level (and –as a robustness check– at the school level when analyzing the role of classmates, and at the neighbourhood level when analyzing the role of neighbours).

Information transmission. To provide further evidence on information transmission, we analyze how the effect of classmates' warnings varies depending on the peers' warning stage *relative* to the family's own warning stage. We expect that warnings received by peers who are in a lower warning stage than the family's own should carry relatively less information, because the family has already experienced first hand that the government enforces the program up to that level. On the other hand, warnings of a level *higher* than one's own carry new information, because the family could still hold a prior that the government punishes up to some point, but will not go through with more severe and costly punishments. Peers' warnings of the same stage as the family's could have an effect –albeit lower than that of higher warning stages– if families perceive the likelihood of enforcement to be correlated across warning stages with a correlation that is stronger between consecutive warning stages (e.g. because people believe that the higher the warning stage, the less likely that the government will go through with more severe and

costly punishments).

We estimate the following equation separately for families in warning stages $l = 1, \dots, 4$

$$Y_{iht} = \sum_{k=1}^5 \alpha_k PEERWARN_{iht-1}^k + \delta \mathbf{X}_{iht} + \mathbf{D}_t + \mathbf{D}_i + \epsilon_{iht} \quad (5)$$

where the variables are defined as before, except that $PEERWARN_{ih,t}^k$ is included separately for different warning stages $k = 1, \dots, 5$. Since the regression specification includes individual fixed effects and conditions on a family's warning stage, we analyze how a given individual in a specific warning stage is affected by her classmates receiving different types of warnings.

4.2 Identification: Learning from Neighbours

While classmates are an important source of information about the strictness of enforcement, another potential source are adults who are BFP beneficiaries and who live in the same neighbourhood. We hypothesize that household heads who experience penalties may interact directly, and not only through their children in schools, and focus on a dimension of interaction for which we can find exogenous variation.

Consider the interaction among adults living in the same neighbourhood. These adults and their families may experience common shocks and hence their children's attendance behaviour may be correlated. However, as we discussed in section 2, the day of the month starting from which households are scheduled to cash-in their transfer at the local Caixa branch is determined by the last digit of the social security

number (NIS) of the designated household member.

First, we provide evidence that the last digit of the household head's social security number is as good as random. Appendix table A.3 shows that this last digit is uncorrelated with our main outcome variable, namely "noncompliance". Appendix table A.4 displays summary statistics of individual and family characteristics for each of the ten digits: differences are small (at the third or at most second decimal digit) and mostly insignificant.

Given the exogeneity of the last digit of the NIS, the set of families living in a given zip code scheduled to receive their transfer (and warning, if applicable) on the same day is also exogenous. Why would this be relevant from an information point of view? Households receive warnings when they pick up the transfer from the local Caixa service point.²⁵ Designated BFP beneficiaries (94% of whom are women) who visit the same Caixa point on the same day are more likely to meet and communicate about penalties on the day when any information about enforcement is most salient. We thus test whether compliance of a household is affected by warnings received by other households that live in the same zip code *and can cash the transfer on the same day*, conditional on warnings received by all households in the neighbourhood.²⁶

We estimate the following specification:

²⁵Kumar et al (2006) report that in the early 2000's Caixa started expanding their banking correspondent network to reach even remote locations, and that by 2005 there was roughly one Caixa correspondent per zip code.

²⁶Families are assigned a date starting from which they can pick up their benefits, but they can choose to do so at a later date. While the actual date is endogenous (and might depend on friends' decisions), the scheduled date is random and thus creates an exogenous peer group. Anyway the two are likely to be strongly correlated since the actual date cannot be before the scheduled one, and because poor people are often cash constrained. Given that we use the scheduled and not the actual date, our estimates are intention-to-treat effects and thus lower bounds.

$$\begin{aligned}
Y_{iht} = & \sum_{n=-1}^1 \alpha_n BANKWARN_{ih,t-1+n} + \sum_{n=-1}^1 \beta_n neighbourWARN_{ih,t-1+n} \\
& + \sum_{k=1}^5 \gamma_k WS_{ht}^k + \delta \mathbf{X}_{iht} + \eta neighbourS_{iht} + \mathbf{D}_t + \mathbf{D}_i + \epsilon_{iht}
\end{aligned}$$

where all variables are defined as in equation (1), except for the following. We include as controls the number of BFP beneficiaries in i 's zip code in month t ($neighbourS_{iht}$) and $neighbourWARN_{iht}$, which denotes the fraction of BFP beneficiaries living in i 's zip code who receive a warning in month t . The key regressor of interest is $BANKWARN_{iht}$, which is the fraction of BFP beneficiaries living in i 's zip code who receive a warning in month t among those whose families can pick up the transfers at the local Caixa 'bank' on the same day. By construction, this variable is not based on a peer group which is chosen endogenously.

In terms of timing, we include the key regressors, i.e. the fraction of BFP neighbours who receive a warning and the fraction of BFP neighbours who receive a warning on the same day, as a first- and second-order lag, since warnings are received at the end of the month, so that there is very little time to respond in that month (as discussed in the previous section). Similarly to our analysis of classmates, we conduct the falsification test of whether the contemporaneous variable of $BANKWARN$ is significantly different from zero or not to rule out the importance of correlated shocks and mean reversion.

As discussed in section 4.1, given the panel structure of our data we estimate a linear probability model and cluster the standard errors at the household level and –as a robustness check– at the zip-code (neighbourhood) level.

Moreover, we provide further evidence on information transmission by conducting a test analogous to equation (5) in the case of classmates, that is we analyze whether warnings received by peers who are in a lower warning stage than the family's own have a smaller effect than warnings of a level *higher* than one's own.

We estimate the following equation separately for families in warning stages $l = 1, \dots, 4$

$$\begin{aligned}
Y_{iht} = & \sum_{k=1}^5 \alpha_k \text{BANKWARN}_{iht-1}^k + \sum_{k=1}^5 \beta_k \text{neighbourWARN}_{iht-1}^k \\
& + \delta \mathbf{X}_{iht} + \eta \text{neighbourS}_{iht} + \mathbf{D}_t + \mathbf{D}_i + \epsilon_{iht}
\end{aligned} \tag{6}$$

where the variables are defined as before, except that the main variable of interest, BANKWARN_{iht}^k (as well as the control variable $\text{neighbourWARN}_{iht}^k$), is included separately for the different warning stages $k = 1, \dots, 5$.

5 Results

In this section we present our main results. We start by examining the effect of warnings received by a child's own classmates, then move to warnings received by classmates of a child's siblings, and finally consider warnings received by neighbouring families.

5.1 Learning from Classmates

Classmates. Table 2 reports our main results on the effect of classmates' experiences of enforcement on individuals' compliance behaviour.

[Insert Table 2]

The coefficient on the fraction of classmates who are warned (i.e., receiving any type of penalty) in column 1 is -0.011 for the lagged variable, i.e. the month immediately following the warnings, and -0.023 for the second lag (i.e. the month after that), both significant at the 1% level. This indicates that after more classmates receive a warning, an individual's noncompliance decreases. In other words, the child is more likely to attend school as required by the program, the more of her classmates experience enforcement. To assess the magnitude of the effect in relation to the mean: we find that a one-standard deviation increase in the fraction of a child i 's classmates warned decreases child i 's noncompliance by 1.2% in the month immediately following the warning and by 2.4% in the month after that. We consider these magnitudes as a lower bound for the total spillovers arising from enforcement for two reasons. First, since we do not know who interacts or communicates with whom, these are intention-to-treat effects. Second, by controlling for the fraction of classmates noncompliant (and its lag) we net out the 'direct' peer effects coming from classmates' increased attendance in response to their own warnings.

Appendix table A.5 presents the same results as table 2, but displays in addition the coefficients on the included (time-varying) controls including the family's own warning stage, which we control for in all specifications.

In column 2 we include the fraction of classmates warned as contemporaneous

variable (i.e. referring to the month at the end of which a child's classmates are warned) to conduct a falsification test and rule out the importance of correlated shocks and mean reversion. As explained in section 4.1, if mean reversion were driving our results, we would expect a negative coefficient on the contemporaneous variable of classmates' warnings, since it is extremely unlikely that the shocks that led to the initial noncompliance would revert to the mean in the month immediately after the peers received warnings (which happens between 2 to 6 months after the noncompliance) but not one month earlier. We find instead a precisely estimated zero, lending support to our interpretation.

In column 3 we control for the fraction of classmates who fail to meet the attendance threshold and find that both the contemporaneous and the lagged variable are positively correlated with own noncompliance. At the same time, we still find a negative effect of classmates' warnings that is equally significant and comparable in size to that of the previous specifications. In particular, the coefficient on the fraction of classmates who are warned is now -0.010 for the lagged variable and -0.020 for the two-month lag. Again, this is consistent with our learning interpretation and suggests that 'conventional' peer effects operating through peers' attendance are not driving our results.

In Table 2 we cluster standard errors at the household level. Results are identical when we cluster at the school level instead (see Table A.14 in the Online Appendix).²⁷

Siblings' classmates. In table 3 we focus on warnings received by siblings'

²⁷In Appendix Table A.8 we estimate a specification that does not control for own warning stage. The results are very similar - if anything the estimated effects are larger. Results are also unchanged if we control for the fraction of classmates warned over the number of classmates noncompliant (Appendix Table A.11).

peers. For this analysis we naturally need to restrict the sample to children who have at least one sibling in the age range of BFP conditionality. It turns out that a large fraction of our original sample satisfies this condition, but in column 1 we report our benchmark estimates for the reduced sample to show that the magnitude of the coefficients is virtually the same as in column 1 of table 2.

[Insert Table 3]

Column 2 of table 3 shows that not only an individual's own peers but also her siblings' peers matter for her attendance decision: controlling for the warnings of the child's own classmates, an increase in the fraction of siblings' classmates warned reduces a child's likelihood of noncompliance the month following the warnings and two months after the warnings. The coefficient on the fraction of siblings' classmates warned is -0.005 for the lagged variable and -0.008 for the second lag (significant at the 5 and 1% level, respectively) compared to a lagged (second lag) effect of -0.008 (-0.019) on the fraction of own classmates warned. In terms of magnitude, if the fraction of i 's siblings' classmates warned in a given month increased by one standard deviation, child i 's noncompliance would decrease by 0.6% in the same month and 1% in the following month, relative to the mean.

The results remain significant and comparable in size for the second lag when we exclusively focus on siblings going to a different school (column 3), while the coefficient on the lagged variable remains negative but insignificant. The findings in columns 2 and 3 lend further support to our learning interpretation, as the peers who receive warnings in columns 2 and 3 are attending different grades or schools, are exposed to different shocks and are no direct peers of child i .²⁸

²⁸Shocks that are common to child i 's class and her siblings' peers (in another grade or school)

Finally, we try to assess whether the above results could be due to sibling peer effects instead of information transmission. While we cannot provide a direct test for this, we start from the prior that siblings of the same gender are typically more likely to interact and have common hobbies, for example. So if peer effects were driving the results, we should find weaker effects of sibling's classmates' warnings when the sibling is of the opposite gender than the individual herself. In column 4 of table 3 we focus on siblings of the opposite sex and we find coefficients that are very similar to those previously estimated (and if anything somewhat larger).

In Appendix table A.15 we present results clustering standard errors at the school level. Results are unchanged.²⁹

Figure 1 provides a visual representation of the estimated effects of warnings to own and siblings' classmates, showing coefficients and standard errors from specification (2) in table 3.

[Insert Figure 1]

Informational content of different warnings. To provide an additional piece of evidence that the underlying mechanism of the spillover effects is information transmission and belief updating about enforcement, we analyze how the effect of classmates' warnings varies depending on the peers' warning stage *relative* to the family's own warning stage (as discussed at the end of section 4.1). In particular, we expect warnings of a level *higher* than one's own to carry more information,

are controlled for by including a variable for the fraction warned of child i 's direct classmates, as discussed in section 4.1.

²⁹In Appendix Table A.9 we show that the results are similar if we do not control for own warning stage. If anything, the estimated effects are larger. Results are also unchanged if we control for the fraction of classmates warned over the number of classmates noncompliant (Appendix Table A.12).

because they refer to government enforcement for stages that the household has not yet reached.

[Insert Table 4]

In Table 4 we distinguish between classmates who receive warning 1, 2, ...5 (the variables listed by row) and estimate four different regression equations, conditional on whether a family is currently in warning stage 1, 2, 3 or 4 (the different columns). Clearly, all stages higher than one's own should be relevant, but the next stage might be particularly important for at least two reasons. First, in our data, the higher the warning stage, the fewer the households that reach it. This implies that more households receive a warning for the next stage, compared to subsequent warnings. Second, if people are present-oriented, the closer in the future is the punishment, the more they care about it.

The estimates in Table 4 lend support to our hypothesis that warnings received by peers in higher warning stages have a stronger impact on school attendance than warnings for the same or lower warning stage.³⁰

For example, in column 1 the coefficient on peers' warnings for stage 2 is twice as large as that for stage 1. In column 2 the coefficient on peers' warnings for stage 3 is about one and a half times as large as for stage 2 (similarly for column 3). Only in column 4 we find that while the coefficient of peers' warnings for stage 4 is significant, the one for stage 5 is not, which is likely due to the fact that extremely

³⁰The coefficients on lower warning stages than the household's own stage are never significant. The negative coefficient on warnings of the same stage can be due to the fact that people perceive the likelihood of enforcement to be correlated across warning stages with a stronger correlation between consecutive warning stages (e.g. because people believe that the higher the warning stage, the less likely that the government will go through with more severe and costly punishments, as discussed at the end of section 4.1).

few families reach warning stage 5 at all.

Figure 2 displays the estimated coefficients (and standard errors) of classmates' warnings for different stages, separately by the household's own warning stage, from table 4.

[Insert Figure 2]

Overall, these results corroborate our interpretation of learning about enforcement, as it would be difficult to find alternative explanations that produce the asymmetric pattern that we uncover. In particular, such an asymmetry would not be consistent with 'peer effects' in the standard sense, that is, with imitation. In the case of imitation an individual's response might depend on the type of warning the peers receive (since that determines the peers' attendance behaviour), but conditional on the peers' type of warning, it should not depend on the family's own warning stage.

Another alternative mechanism could be that peers' warnings act as a reminder that noncompliance is punished – an 'attention' or 'salience' channel. While we do not aim to rule out salience altogether, we think this is unlikely to drive our effects for two reasons. First, the existence of program conditions is constantly reminded to beneficiaries, hence it should be relatively salient throughout, and not only when peers receive warnings. In fact, the information is communicated to beneficiary families when they register for the program via a simple leaflet and explanations by social workers at that time; it is then repeatedly stressed in the media (TV, radio, newspapers)³¹; and social workers themselves often go door-to-door to reiterate

³¹For example, Lindert and Vincensini (2010) analyzed over 6,500 articles from six Brazilian newspapers over 2001-2006 and report that nearly half the articles they found on BFP discussed program conditions and punishments.

the importance of compliance.³² Also, before benefits are suspended, beneficiaries receive a warning slip directly from the ATM where they withdraw the transfer, and this slip contains again a reminder of the conditionalities.

The second reason why we think that salience is unlikely to drive our results is that, if peer warnings simply served as reminders that there are program conditions, any type of warning could fulfill this function. Instead we have shown in Table 4 that there is a systematic asymmetry in effects, that is, peers who receive lower level warnings have less of an impact.

5.2 Learning from Neighbours

In the previous section we have shown the importance of classmates in terms of transmitting information about the strictness of enforcement and leading to important multiplier effects of enforcement. Another potential source of information are adults who are BFP beneficiaries from the same neighbourhood and who may interact directly, and not only through their children in schools.

Of course, neighbours are an endogenous peer group and are likely to experience common shocks and hence their children's attendance behaviour may be correlated. However, as discussed in sections 2 and 4.2, the day of the month starting from which households can cash-in their transfers at the local service point (and receive warnings, if applicable) is as good as random. Thus the group of neighbours who go to the local service point on the same day is exogenous conditional

³²In a case study of BFP, Wong (2017) reports on page 18 that “social workers assigned to specific catchment areas in Belo Horizonte often go door-to-door in poorer neighbourhoods to check that BFP recipients are meeting the program's conditionalities. (...) And in Sapeacu, community health agents make daily home visits to remind families about meeting the BFP's health and education conditionalities.”

on living in the same neighbourhood. We hypothesize that household heads (94% of whom are women) who visit the same service point on the same day may be likely to meet and communicate about warnings, and we test whether compliance of a household is affected by warnings received by neighbours *who can cash the transfer on the same day*, after controlling for the warnings received by all households in the neighbourhood.

[Insert Table 5]

Table 5 reports the results. Column 1 shows that there is indeed a strongly significant negative relation between the fraction of BFP beneficiaries from the same zip code who are warned on the same day (*BANKWARN*) and individual noncompliance. The coefficients on the the lagged variable and the second lag are -0.028 and -0.055 , respectively. The corresponding effect sizes, relative to the mean, are as follows: an increase of one standard deviation in the fraction of neighbours warned on the same day leads to a decrease of a child's probability of noncompliance by 1.5% in the month immediately following the warnings and 3% in the month after that.

In column 2 we include *BANKWARN* as contemporaneous variable (i.e. referring to the month at the end of which warnings take place) to conduct a falsification test and rule out the importance of correlated shocks and mean reversion. The coefficient on this variable is a precisely estimated zero (with a coefficient of 0.0004), lending support to our learning interpretation.

In column 3 we control for the fraction of BFP beneficiaries from the same neighbourhood who are warned in a given month (*neighbourWARN*). The fraction

warned *on the same day* (*BANKWARN*) is exogenous when controlling for the overall fraction of neighbours warned within that month, as discussed above. The coefficient of the lag of *neighbourWARN* is positive, while that on its second lag is negative, which might be due to correlated shocks that have some persistence but mean-revert at some point.

More importantly, controlling for *neighbourWARN* does not change our key result, which is that a larger fraction warned *on the same day* significantly decreases an individual's likelihood of noncompliance. The point estimates of -0.035 and -0.028 in column 3 imply that a one-standard deviation increase in the fraction of BFP neighbours warned on the same day decreases a child's probability of noncompliance by 1.9% in the month immediately following the warnings and by 1.5% in the month after that.

In Appendix table A.16 we present results clustering standard errors at the zip-code level and we find that all our conclusions are unchanged. In particular, the coefficients on the fraction of neighbours warned on the same day, and on its lag remain significant at the 1% level.³³

We also analyzed how the effect of other households' warnings varies depending on their warning stage relative to the family's own stage in the case of neighbours. We find that families in warning stage 1 –who constitute 70% of the sample– respond to neighbours receiving a first warning, but their response is more than three times stronger for neighbours receiving their second warning (see Appendix Table A.17). For the 30% of families in higher warning stages, the pattern is less

³³In Appendix Table A.10 we omit from the controls the dummies for own warning stage. The results are virtually unchanged. Results are also unchanged if we control for the fraction of classmates warned over the number of classmates noncompliant (Appendix Table A.13).

clear. Thus, analogous to our analysis in section 5.1, for the majority of families, neighbours in higher warning stages than the household's have a stronger impact on compliance, consistent with our learning interpretation.

5.3 Learning from Classmates and Neighbours

In this section we include both classmates and neighbours simultaneously to test whether the effect of one type of peer is partly captured by the other type of peer, and to shed light on the relative importance of the different types of social interaction.

In column 4 of table 5 we therefore include among the controls the first and second lag of the fraction of classmates warned, which leaves the coefficients on neighbours' warnings unchanged. Also the coefficient on the lagged fraction of classmates warned is unchanged compared to column 1 of Table 2, while the second lag of this variable is somewhat smaller but still negative and significant. This suggests that neighbours and classmates constitute two relatively independent sources of information about enforcement. The estimated effects of classmates and neighbours warnings from column 4 are also shown in Figure 3.

[Insert Figure 3]

The estimates in column 4 of table 5 imply that an increase of one standard deviation in the fraction of BFP neighbours warned on the same day decreases a child's probability of noncompliance by 1.9% in the month following the warnings and by 1.5% in the month after that. A one-standard deviation increase of the fraction of BFP classmates warned decreases a child's probability of noncompliance by 1% in the month following the warnings and 1.2% the month after that. The

magnitude of the response is thus very similar across the two types of peers.

Column 5 of table 5 displays the last specification where we add warnings of siblings' peers. Again the coefficients on warnings of neighbours, of classmates and of siblings' classmates are very similar to those obtained in the separate specifications (see table 3).

To summarize, the fact that we find consistent effects for two different types of peers that are based on two entirely different identification strategies (and sources of variation), strengthens our confidence in the validity of our approach and interpretation. Moreover, our results suggest that the main mechanism behind these effects is information transmission.

5.4 Robustness: Behavioural Response vs Teachers' Reporting

One remaining interpretation issue is the extent to which the behavioural responses we identify correspond to an actual increase in attendance, and not simply to more lenient reporting by teachers.³⁴

It is conceivable that when a family receives a penalty, parents may try to convince teachers to be less strict in registering non-attendance in the future. Even in this case our hypothesis that families learn about the strictness of enforcement would be valid: it would be the action taken by the family that would differ (e.g., persuading teachers instead of sending children to school).

However, in this paper we have investigated whether individuals respond to penalties received *by others*. In particular, we found that child *i*'s noncompliance

³⁴Unfortunately, we cannot directly test for misreporting, differently from Linden and Shastri (2012) who rely on external monitors' verification in a sample of Indian schools.

decreases when the classmates of i 's siblings get warned or when i 's neighbours receive warnings. Since the teacher of child i would not know about those children's warnings, our results are more likely interpreted as affecting actual school attendance as opposed to misreporting.³⁵

Also, children have several different teachers who need to register attendance: bribing or convincing all of them may not be easy, especially because stakes are relatively high for teachers and school principals who may be caught falsifying records.³⁶

Finally, a comprehensive control system is in place to monitor the implementation process and the compliance with program conditions.³⁷ In particular, the Social Controls Councils (SCCs) established in all municipalities are charged with overseeing the monitoring of school attendance and formally requesting the blocking of benefits where this is justified. For this purpose, SCCs regularly conduct random audits of schools.³⁸ Moreover, since 2006 MDS monitors municipal implementation quality using a system of financial incentives (performance-based administrative cost subsidies).

Taken all the evidence together, our results suggest that families respond to peers experiencing enforcement and that at least an important part of their behavioural

³⁵It seems difficult for families to ask teachers to falsify attendance records simply because someone else in another school has been warned. It is important to note that teachers who are generally lenient and do not (or rarely) register nonattendance would not generate the pattern we see, which is a change in recorded attendance in response to peers' (such as siblings' classmates' or neighbours') receipt of warnings.

³⁶Teachers earn a relatively high salary and it is unclear that parents from poor households could convince teachers (or offer a sufficiently large amount of money to induce teachers) to falsify records and risk punishment by the school principal, and/or complaints by other students or teachers.

³⁷For details, see Online Appendix C and Lindert et al (2007).

³⁸The SCCs themselves are monitored in random audits of municipalities conducted by the General Controllers Office on a monthly basis.

response should correspond to an actual increase in terms of school attendance.³⁹

6 Conclusions

In this paper we study the implementation of the large-scale conditional cash transfer program “Bolsa Familia” (BFP) in Brazil. This program conditions transfers to poor families on children’s school attendance: when families fail to comply with the requirement, they receive a series of warnings and financial penalties.

We analyze how people respond to peers experiencing the enforcement of program conditions and show that enforcement creates spillover effects. In particular, individuals respond to warnings received by classmates, classmates of siblings (in other schools and of other gender) and neighbours. We try to shed light on the underlying mechanisms and show that the effects are not simply due to direct peer effects, but to information transmission. For example, we find that individuals respond more strongly to peers experiencing penalties of higher order (i.e., “warning stages” that they themselves have not yet experienced), as these penalties convey more information about future enforcement. Importantly, we find consistent effects for two different types of peers –classmates and neighbours– that are based on two entirely different identification strategies and this strengthens our confidence in the validity of our approaches.

What do we learn from our findings? First, our results show that it is important

³⁹A different way in which teachers could play a role is to act as intermediaries in information transmission. Note that teachers are not informed about a family’s warnings, unless they learn about them directly from the family or the child. If classmates get warned and communicate this to their teacher, who in turn reminds children of the possible consequences of noncompliance, then we see this as yet another channel of information transmission (in addition to children and parents and neighbours talking to each other).

to take spillover effects into account in a cost-benefit calculation, since monitoring and enforcement are very costly (administrative costs, political costs, welfare costs of taking away transfers from poor families). This aspect seems particularly important for developing countries, as they might lack administrative capacity or political will to strictly enforce the rules.⁴⁰

Second, we shed light on mechanisms, distinguishing spillovers via information transmission from those via imitation. This is policy relevant for the following reason. In light of the dual goals of conditional cash transfer programs (increasing children's school attendance *and* supporting poor families) and given the large welfare costs of penalizing poor families by taking away (part of) their transfer, it is critical to design the enforcement system as effectively as possible. Ideally, one would want to credibly threaten enforcement and reach high compliance levels with as little *actual* enforcement as possible. Our findings that relate spillover effects to information transmission suggest that governments might be able to improve compliance –holding constant the level of enforcement– by giving credible information about enforcement activities at the school or neighbourhood level.

To conclude, our paper shows that learning spillovers in enforcement are important in the context of one of the world's largest conditional cash transfer programs and that they can have important policy implications. While the results may not immediately generalize to other welfare programs, our analysis suggests strategies for estimating such spillover effects in the context of conditional programs.

⁴⁰Brollo, Kaufmann and La Ferrara (2019) study the electoral costs of enforcement and the political incentives for manipulating the implementation of BFP conditionality in Brazil.

Fernanda Brollo: University of Warwick, CAGE and CEPR

Katja Maria Kaufmann: Mannheim University, briq, CESifo and IZA

Eliana La Ferrara: Bocconi University, IGIER and LEAP

References

Aizer, A. and Currie, J. (2004). 'Networks or Neighbourhoods? Correlations in the Use of Publicly-Funded Maternity Care in California', *Journal of Public Economics*, vol. 88, pp. 2573–85.

Aizer, A., Eli, S., Ferrie, J. and Lleras-Muney, A. (2016). 'The Long-Run Impact of Cash Transfers to Poor Families', *American Economic Review*, vol. 106(4), pp. 935–71.

Alatas, V., Banerjee, A., Chandrasekhar, A.G., Hanna, R. and Olken, B.A. (2016). 'Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia', *American Economic Review*, vol. 106(7), pp. 1663–1704.

Angelucci, M. and Giorgi, G.D. (2009). 'Indirect Effects of an Aid Program: How do Cash Transfers Affect Ineligibles' Consumption?', *American Economic Review*, vol. 99(1), pp. 486–508.

Angelucci, M., Giorgi, G.D., Rangel, M. and Rasul, I. (2010). 'Family Networks and School Enrollment: Evidence from a Randomized Social Experiment', *Journal of Public Economics*, vol. 94(3-4), pp. 197–221.

Baird, S., McIntosh, C. and Özler, B. (2011). 'Cash or Condition? Evidence from a Cash Transfer Experiment', *The Quarterly Journal of Economics*, vol. 126(4), pp. 1709–53

Banerjee, A., Glennerster, R. and Duflo, E. (2008). 'Putting Band-Aid on a Corpse: Incentives for Nurses in the Indian Public Health Care System', *Journal of the European Economic Association*, vol. 6(2-3), pp. 487–500.

Banerjee, A., Chandrasekhar, A.G., Duflo, E. and Jackson, M.O. (2013). ‘The Diffusion of Microfinance’, *Science*, vol. 341(6144).

Banerjee, A., Chandrasekhar, A.G., Duflo, E. and Jackson, M.O. (2019). ‘Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials’, *The Review of Economic Studies*, vol. 86(6), pp. 2453–90.

Barham, T., Macours, k. and Maluccio, J. (2017). ‘Are Conditional Cash Transfers Fulfilling Their Promise? Schooling, Learning, and Earnings After 10 Years’, Discussion Paper No. DP11937, CEPR.

Barrera-Orsorio, F., Bertrand, M., Linden, L.L. and Perez-Calle, F. (2011). ‘Improving the Design of Conditional Transfer Programs: Evidence from a Randomized Education Experiment in Colombia’, *American Economic Journal: Applied Economics*, vol. 3(2), pp. 167–195.

Barrera-Orsorio, F., Linden, L.L., and Saavedra, J.E. (2019). ‘Medium- and Long-Term Educational Consequences of Alternative Conditional Cash Transfer Designs: Experimental Evidence from Colombia’, *American Economic Journal: Applied Economics*, vol.11(3), pp. 54–91.

Bastagli, F. (2009). ‘Conditionality in Public Policy Targeted to the Poor: Promoting Resilience?’, *Social Policy & Society*, vol. 8(1), pp. 127–140

Becker, G.S. (1968). ‘Crime and Punishment: An Economic Approach’, *Journal of Political Economy*, vol. 76(2), pp. 169–217.

Benhassine, N., Devoto, F., Duflo, E., Dupas, P. and Pouliquen, V. (2015). ‘Turning a Shove into a Nudge? A “Labeled Cash Transfer” for Education’, *American Economic Journal: Economic Policy*, vol. 7(3), pp. 86–125.

- Black, D., Smith, J., Berger, M. and Noel, B. (2003). ‘Is the Threat of Reemployment Services More Effective than the Services Themselves? Evidence from Random Assignment in the UI System’, *American Economic Review*, vol. 93(4), pp. 1313–1327.
- Bobonis, G.J. and Finan, F. (2009). ‘Neighbourhood Peer effects in Secondary School Enrollment Decisions’, *The Review of Economics and Statistics*, vol. 91(4), pp. 695–716.
- Boning, W.C., Guyton, J., Hodge II, R.H., Slemrod, J. and Troiano, U. (2018). ‘Heard it Through the Grapevine: Direct and Network Effects of a Tax Enforcement Field Experiment’, Working Paper 24305, NBER.
- Bourguignon, F., Ferreira, F. and Leite, P.G. (2003). ‘Conditional Cash Transfers, Schooling and Child Labor: Micro-simulating Brazil’s Bolsa Escola Program’, *The World Bank Economic Review*, vol. 17(2), pp. 229–254
- Bramoullé, Y., Djebbari, H., and Fortin, B. (2009). ‘Identification of Peer Effects Through Social Networks’, *Journal of Econometrics*, vol. 150, pp. 41–55.
- Brollo, F., Kaufmann, K.M. and La Ferrara, E. (2018). ‘Enforcement of Program Conditions and Individual Compliance in Welfare Programs’, mimeo.
- Brollo, F., Kaufmann, K.M. and La Ferrara, E. (2019). ‘The Political Economy of Program Enforcement: Evidence from Brazil’, *Journal of the European Economic Association*
- Bursztyn, L. and Coffman, L.C. (2012). ‘The Schooling Decision: Family Preferences, Intergenerational Conflict, and Moral Hazard in the Brazilian Favelas’, *Journal of Political Economy*, vol. 120(3), pp. 359–97.

- Cameron, L. and Shah, M. (2014). ‘Mistargeting of Cash Transfers, Social Capital Destruction, and Crime in Indonesia’, *Economic Development and Cultural Change*, vol. 62(2), pp. 381–415.
- Carrillo, P., Pomeranz, D. and Singhal, M. (2017). ‘Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement’, *American Economic Journal: Applied Economics*.
- Chetty, R., Friedman, J. and Saez, E. (2013). ‘Using Differences in Knowledge Across Neighbourhoods to Uncover the Impacts of the EITC on Earnings’, *American Economic Review*, vol. 103(7), pp. 2683–2721.
- Chioda, L., de Mello, J.M.P. and Soares, R. (2016). ‘Spillovers from Conditional Cash Transfer Programs: Bolsa Familia and Crime in Urban Brazil’, *Economics of Education Review*, vol. 54, pp. 306–20
- Dahl, G.B. and Gielen, A.C. (2018). ‘Intergenerational Spillovers in Disability Insurance’, Working Paper 24296, NBER.
- Dahl, G.B., Loken, K.V. and Mogstad, M. (2014). ‘Peer Effects in Program Participation’, *American Economic Review*, vol. 104(7), pp. 2049–74.
- DeBrauw, A. and Hoddinott, J. (2011). ‘Must Conditional Cash Transfer Programs be Conditioned to be Effective? The Impact of Conditioning Transfers on School Enrollment in Mexico’, *Journal of Development Economics*, vol. 96(2), pp. 359–70
- DeBrauw, A., Gilligan, D., Hoddinott, J. and Roy, S. (2015a). ‘The Impact of Bolsa Familia on Education’, *World Development*, vol. 70(6), pp. 303–16.
- DeBrauw, A., Gilligan, D., Hoddinott, J. and Roy, S. (2015b). ‘Bolsa Familia

and Household Labor Supply’, *Economic Development and Cultural Change*, vol. 63(3), pp. 423–57.

De Giorgi, G., Pellizzari, M. and Redaelli, S. (2010). ‘Identification of Social Interactions through Partially Overlapping Peer Groups’, *American Economic Journal: Applied Economics*, vol. 2(2), pp. 241–75.

De Janvry, A., Finan, F. and Sadoulet, E. (2011). ‘Local Electoral Incentives and Decentralized Program Performance’, *The Review of Economics and Statistics*, vol. 94(3), pp. 672–85.

Drago, F., Mengel, F. and Traxler, C. (2015) ‘Compliance Behaviour in Networks: Evidence from a Field Experiment’, IZA DP No. 9443.

Duflo, E. and Saez, E. (2003). ‘The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment’, *Quarterly Journal of Economics*, vol. 118(3), pp. 815–42

Dunbar, R.I.M. (2004). ‘Gossiping in an Evolutionary Perspective’, *Review of General Psychology*, vol. 8(2), pp. 100–10.

Dunifon, R., Fomby, P. and Musick, K. (2017). ‘Siblings and Children’s Time Use in the United States’, *Demographic Research*, vol. 37, pp. 1611–24.

Dwenger, N., Kleven, H.J., Rasul, I. and Rincke, J. (2016). ‘Extrinsic and Intrinsic Motivations for Tax Compliance: Evidence from a Field Experiment in Germany’, *American Economic Journal: Economic Policy*, vol. 8(3), pp. 203–32.

Fernald, L., Gertler, P. and Neufeld, L. (2009). ‘10-Year Effect of Oportunidades, Mexico’s Conditional Cash Transfer Programme, on Child Growth, Cognition, Lan-

guage, and Behaviour: a Longitudinal Follow-up Study”, *Lancet*, vol. 371, pp. 828–37.

Kleven, H.J., Knudsen, M.B., Kreiner, C.T., Pedersen, S. and Saez, E. (2011). ‘Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark’, *Econometrica*, vol. 79 (3), pp. 651–92.

Kumar, A., Nair, A., Parsons, A. and Urdapilleta, E. (2006). ‘Expanding Bank Outreach through Retail Partnerships Correspondent Banking in Brazil’, Working Paper No. 85, The World Bank.

Lee, L.F., Liu, X. and Lin, X. (2010). ‘Specification and Estimation of Social Interaction Models with Network Structures’, *The Econometrics Journal*, vol. 13, pp. 145–76.

Linden, L. and Shastri, K. (2012). ‘Grain Inflation: Identifying Agent Discretion in Response to a Conditional School Nutrition Program’, *Journal of Development Economics*, vol. 99(1), pp. 128–38.

Lindert, K., Linder, A., Hobbs, J. and de la Brière, B. (2007). ‘The Nuts and Bolts of Brazil’s Bolsa Família Program: Implementing Conditional Cash Transfers in a Decentralized Context’, SP Discussion Paper No. 0709, The World Bank.

Lindert, K. and Vincensini, V. (2010). ‘Social Policy, Perceptions and the Press: An Analysis of the Media’s Treatment of Conditional Cash Transfers in Brazil’, SP Discussion Paper No. 1008, The World Bank.

Lochner, L. (2007). ‘Individual Perceptions of the Criminal Justice System’, *American Economic Review*, vol. 97(1), pp. 444–60.

- Manski, C. (1993). ‘Identification of Endogenous Social Effects: The Reflection Problem’, *Review of Economic Studies*, vol. 60, pp. 531–42.
- Nicoletti, C., Salvanes, K.G. and Tominey, E. (2018). ‘The Family Peer Effect on Mothers’ Labour Supply’, *American Economic Journal: Applied Economics*, vol. 10(3), pp. 206–34
- Olken, B.A. (2007). ‘Monitoring Corruption: Evidence from a Field Experiment in Indonesia’, *Journal of Political Economy*, vol. 115(2), pp. 200–49.
- Parker, S., Rubalcava, L. and Teruel, G. (2012). ‘Do Conditional Cash Transfer Programs Improve Work and Earnings among its Youth Beneficiaries? Evidence after a Decade of a Mexican Cash Transfer Program’, Unpublished.
- Pomeranz, D. (2015). ‘No Taxation Without Information: Deterrence and Self-Enforcement in the Value Added Tax’, *American Economic Review*, vol. 105(8), pp. 2539–69.
- Rinke, J. and Traxler, C. (2011). ‘Enforcement Spillovers’, *The Review of Economics and Statistics*, vol. 93(4), pp. 1224–34.
- Rios, J. (2017). ‘Welfare Analysis of Transfer Programs with Jumps in Reported Income: Evidence from the Brazilian Bolsa Família’, mimeo, Stanford University.
- Schady, N. and Araujo, M.C. (2006). ‘Cash Transfers, Conditions, School Enrollment, and Child Work: Evidence from a Randomized Experiment in Ecuador’, World Bank Policy Research Working Paper 3930, The World Bank.
- Shimeles, A., Gurara, D.Z. and Woldeyes, F. (2017). ‘Taxman’s Dilemma: Coercion or Persuasion? Evidence from a Randomized Field Experiment in Ethiopia’, *American Economic Review*, vol. 107(5), pp. 420–24.

Stigler, G.J. (1970). 'The Optimum Enforcement of Laws', *Journal of Political Economy*, vol. 78(3), pp. 526–36.

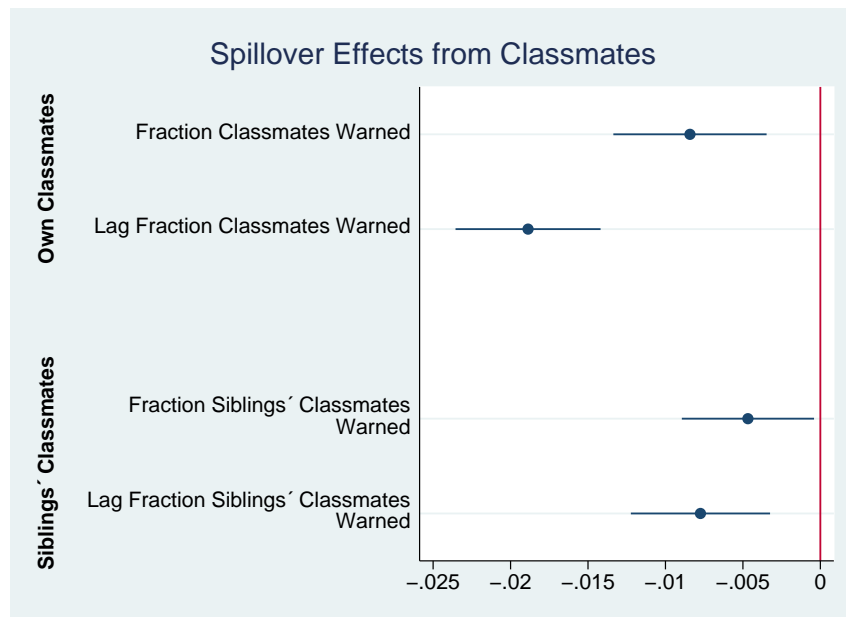
Van den Berg, G.J., Van der Klaauw, B. and Van Ours, J.C. (2004). 'Punitive Sanctions and the Transition Rate from Welfare to Work', *Journal of Labor Economics*, vol. 22, pp. 211–41.

Wong, J. (2017). 'Reaching the Hard to Reach: A Case Study of Brazils Bolsa Familia Program ', mimeo, The Reach Project, University of Toronto.

Figures and Tables


Figures

Figure 1: *Spillover Effects of Own and Siblings' Classmates*



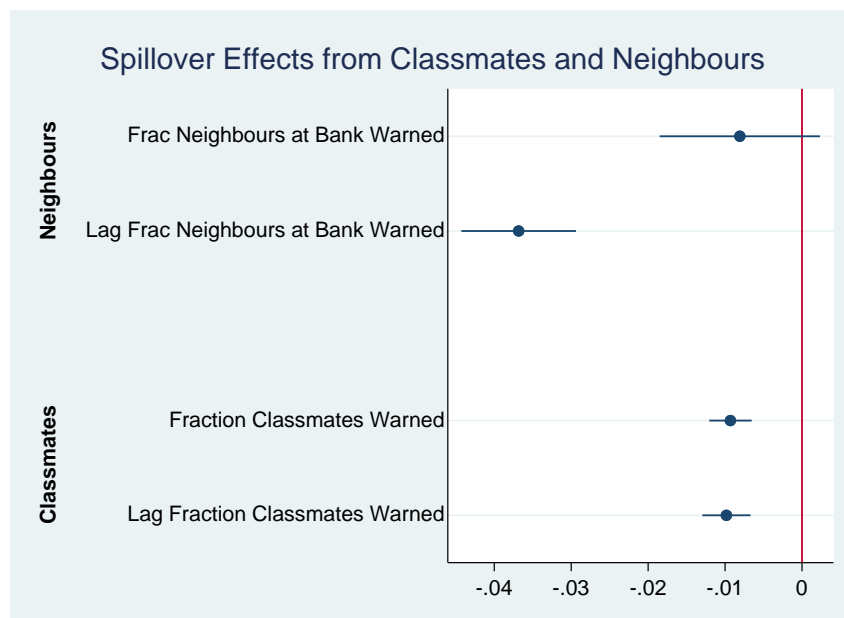
Note: This figure provides a visual representation of the estimated effects of warnings to own and siblings' classmates, showing coefficients and standard errors from specification (2) in table 3.

Figure 2: *Spillover Effects of Classmates by Own Warning Stage*



Note: This figure provides a visual representation of the estimated effects of warnings corresponding to different stages (separately by the household's own stage), showing coefficients and standard errors from table 4.

Figure 3: *Spillover Effects of Neighbours and Classmates*



Note: This figure provides a visual representation of the estimated effects of warnings of classmates versus neighbours, showing coefficients and standard errors from specification (4) in table 5.

Tables

Table 1: *Summary Statistics*

Variable	Observations	Mean	Std. Dev.
Noncompliance	5,903,731	0.051	0.220
A. Peer Warnings: Classmates			
Fraction of Classmates Warned	5,903,731	0.014	0.059
No of Peers in School	5,903,731	279.4	223.0
No of Peers in Grade	5,903,731	52.0	50.5
No of Peers in Class	5,903,731	16.5	6.8
B. Peer Warnings: Siblings' Classmates			
Max Frac of Siblings' Classmates Warned			
All Siblings	4,894,515	0.014	0.058
Siblings in Other Schools	4,894,515	0.013	0.058
Siblings of Opposite Gender	4,894,515	0.009	0.047
C. Peer Warnings: Neighbours			
Frac of Neighbours Warned on Same Day	4,797,361	0.006	0.026
Frac of Neighbours Warned	4,797,361	0.006	0.017
No of Beneficiaries in Neighbourhood	4,797,361	4042.0	3891.0

Table 2: *Effect of Classmates' Warnings*

Dependent Variable: Sample	Noncompliance in a Given Month		
	<i>All Children</i>		
	Benchmark	Placebo: Contemp. Effect	Control for Peers' Noncompl.
	(1)	(2)	(3)
Lag of Fraction of Classmates Warned	-0.0111*** (0.002)	-0.0120*** (0.002)	-0.0100*** (0.002)
Second Lag of Frac of Classmates Warned	-0.0230*** (0.002)	-0.0235*** (0.002)	-0.0199*** (0.002)
Frac of Classmates Warned		0.0001 (0.002)	
Fraction of Classmates Noncompliant			0.2811*** (0.021)
Lag of Frac of Classmates Noncompl			0.0954*** (0.009)
Controls	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
No. Obs.	5,903,731	5,619,182	5,903,731
R-squared	0.28	0.27	0.28

Notes: Robust standard errors in parentheses which are clustered at the household level (see Online Appendix table A.14 for clustering at the school level). *** p<0.01, ** p<0.05, * p<0.1. Included controls are dummies for the family's warning stage, dummies for the individual's age, number of brothers and sisters in the age categories 6 to 10, 11 to 15, and 16 to 18, month and year dummies. All specifications include individual fixed effects.

Table 3: *Effect of Siblings' Classmates' Warnings*

Dependent Variable: Sample	Noncompliance in a Given Month			
	Benchmark	<i>Children with Siblings</i>		
		Siblings' Classmates		
		All Siblings	Siblings: Other Schools	Siblings: Other Gender
(1)	(2)	(3)	(4)	
Lag of Max Frac of Siblings' Classmates Warned		-0.0047** (0.002)		
Second Lag Max Frac of Siblings' Classmates Warned		-0.0077*** (0.002)		
Lag of Max Frac of Siblings' Classmates Warned (Only Sibs at Other School)			-0.0013 (0.002)	
Second Lag Max Frac of Siblings' Classmates Warned (Only Sibs at Other School)			-0.0062** (0.003)	
Lag of Max Frac of Siblings' Classmates Warned (Only Sibs of Other Gender)				-0.0067** (0.003)
Second Lag Max Frac of Siblings' Classmates Warned (Only Sibs of Other Gender)				-0.0073** (0.003)
Lag of Fraction of Classmates Warned	-0.0106*** (0.002)	-0.0084*** (0.003)	-0.0099*** (0.003)	-0.0083*** (0.002)
Second Lag of Fraction of Classmates Warned	-0.0226*** (0.002)	-0.0189*** (0.002)	-0.0191*** (0.003)	-0.0202*** (0.002)
Controls	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
No. Obs.	4,894,515	4,894,515	4,894,515	4,894,515
R-squared	0.28	0.28	0.28	0.28

Notes: Robust standard errors in parentheses which are clustered at the household level (see Online Appendix table A.15 for clustering at the school level). *** p<0.01, ** p<0.05, * p<0.1. Included controls are dummies for the family's warning stage, dummies for the individual's age, number of brothers and sisters in the age categories 6 to 10, 11 to 15, and 16 to 18, month and year dummies. All specifications include individual fixed effects.

Table 4: *Information Content of Classmates' Warnings*

Dependent Variable: Conditional on:	Noncompliance in Given Month			
	Own Warning Stage			
	1	2	3	4
Lag of Frac Classmates Warned (WS 1)	-0.0123*** (0.003)	-0.0139 (0.009)	-0.0119 (0.016)	0.0318 (0.028)
Lag of Frac Classmates Warned (WS 2)	-0.0242*** (0.007)	-0.0202** (0.009)	-0.0177 (0.033)	-0.0690 (0.044)
Lag of Frac Classmates Warned (WS 3)	-0.0187 (0.012)	-0.0310* (0.017)	-0.0510** (0.021)	-0.0388 (0.051)
Lag of Frac Classmates Warned (WS 4)	-0.0197 (0.022)	-0.0391 (0.026)	-0.0786* (0.043)	-0.0591** (0.027)
Lag of Frac Classmates Warned (WS 5)	0.0285 (0.036)	0.0852 (0.070)	0.0384 (0.088)	0.1156 (0.123)
Controls	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
No. Obs.	1,988,235	637,777	232,023	88,496
R-squared	0.46	0.52	0.52	0.54

Notes: Robust standard errors in parentheses (clustered at the household level). *** p<0.01, ** p<0.05, * p<0.1. Included controls are: fraction of classmates noncompliant, age dummies, number of brothers and sisters in the age categories 6 to 10, 11 to 15, and 16 to 18, month and year dummies.

Table 5: *Effect of Neighbours' and Classmates' Warnings*

Dependent Variable:	Noncompliance in Given Month				
	Benchmark (1)	Placebo: Contemp. Effect (2)	Controls for Neighbours Warned (3)	Neighbours and Classmates (4)	Neighbours Classmates, Siblings (5)
Lag of Frac Neighbours Warned on Same Day	-0.0276*** (0.004)	-0.0276*** (0.004)	-0.0349*** (0.006)	-0.0347*** (0.006)	-0.0347*** (0.006)
Second Lag Frac Neighbours Warned on Same Day	-0.0549*** (0.003)	-0.0548*** (0.003)	-0.0284*** (0.004)	-0.0276*** (0.004)	-0.0273*** (0.004)
Frac Neighbours Warned on Same Day		0.0004 (0.004)			
Lag of Fraction Classmates Warned				-0.0102*** (0.002)	-0.0064*** (0.002)
Second Lag of Fraction Classmates Warned				-0.0120*** (0.002)	-0.0076*** (0.002)
Lag of Max Frac of Siblings' Classmates Warned					-0.0096*** (0.002)
Second Lag Max Frac of Siblings' Classmates Warned					-0.0103*** (0.002)
Lag of Frac Neighbours Warned			0.0160* (0.009)	0.0154* (0.009)	0.0152* (0.009)
Second Lag of Frac Neighbours Warned			-0.0612*** (0.007)	-0.0568*** (0.007)	-0.0553*** (0.007)
Controls	Yes	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Observations	4,797,361	4,797,361	4,797,361	4,797,361	4,797,361
R-squared	0.21	0.21	0.21	0.21	0.21

Notes: Robust standard errors in parentheses which are clustered at the household level (see Online Appendix table A.16 for clustering at the zip-code level). *** p<0.01, ** p<0.05, * p<0.1. Included controls are the fraction of neighbours warned (and its lag), the number of BFP neighbours, dummies for the family's warning stage, dummies for the individual's age, number of brothers and sisters in the age categories 6 to 10, 11 to 15, and 16 to 18, month and year dummies. All specifications include individual fixed effects.