

Disrupting Compliance: The Impact of a Randomized Tax Holiday in Uruguay

Thad Dunning* Felipe Monestier[†] Rafael Piñeiro,[‡]
Fernando Rosenblatt,[§] and Guadalupe Tuñón[¶]

June 7, 2024

*Robson Professor, Department of Political Science, University of California, Berkeley; thad.dunning@berkeley.edu.

[†]Assistant Professor, Departamento de Ciencia Política, Universidad de la República; felipe.monestier@cienciassociales.edu.uy

[‡]Professor, Departamento de Ciencias Sociales, Universidad Católica del Uruguay; rafael.pineiro@ucu.edu.uy.

[§]Senior Lecturer, Department of Politics, The University of Manchester; fernando.rosenblatt@manchester.ac.uk.

[¶]Assistant Professor, Department of Politics and School of International and Public Affairs, Princeton University; tunon@princeton.edu.

Abstract

What is the effect of disrupting tax compliance? In a context of substantial nonpayment, the municipal government in Montevideo, Uruguay sought to reward punctual taxpayers by raffling tax holidays at random to them. Using individual payment records as an unobtrusive outcome measure, we combine natural, field and survey experiments to study the effects. Contrary to the policy's aims, winning a tax holiday reduced tax compliance on average for two years after the interruption of payments. Our evidence suggests that the negative effect arose because the holiday disrupts the habit of paying taxes. Our findings provide broader insights into behavioral foundations of compliance while also suggesting effective policy modifications to address an abiding challenge for governments: the raising of revenue.¹

¹Replication files are available in the JOP Dataverse (<https://dataverse.harvard.edu/dataverse/jop>). The empirical analysis has been successfully replicated by the JOP replication analyst. Our pre-analysis plan was registered on the Evidence in Governance and Politics (EGAP) portal, with amendments (registration ID 20140723AA), now available at OSF (<https://osf.io/39d4r>). The study was approved by the Office for the Protection of Human Subjects at the University of California, Berkeley (Protocol ID 2014-04-6286).

1 Introduction

Interactions between citizens and their governments often involve routinized, recurring behaviors. Repeated interactions may breed habits of citizenship and thus help explain the persistence of broad patterns of civic participation. Yet interventions that create or break habits can also be important forces for change. The power of habit thus implies the possibility of virtuous or vicious cycles in citizen behavior—and raises the question of how habits are formed and the consequences of their disruption.

Consider tax compliance, a critical aspect of citizen-state interaction and key facet of state capacity. Scholars often attribute low rates of compliance to poor monitoring capacity, a lack of “infrastructural power” (Mann 1984; Soifer 2008), and weak state institutions (Levi, 1988; Tilly, 1990; Besley and Persson, 2009). These problems can be especially pronounced in developing countries; yet rich and poor states alike face challenges in eliciting revenue by threatening sanctions. Many Latin American municipalities have experimented with rewards programs that instead offer positive incentives for prompt tax payment.² Beginning in 2004, for example, the municipal government of Montevideo randomly selected taxpayers who had made their tax payments on time and awarded them with a “tax holiday”—a year free of the obligation to pay taxes. Insights from both behavioral economics and classical game theory suggest that promising rewards may have different effects than threatening punishments (Andreoni, Harbaugh and Vesterlund 2003; Balliet and Lange 2014). We thus pre-registered a large, multi-faceted study not only to assess the impact of the policy but also to compare the effects of positive incentives to the more typical negative sanctions for non-compliance.

In addition to rewarding past tax payment, however, Montevideo’s holiday disrupts compliance. One hypothesis we pre-registered, among others, was that by interrupting the habit of tax payment, the holidays might induce a negative effect on compliance among winners after the end of the holiday. Indeed, Montevideo’s policy provides a rare form of randomization that facilitates study of the impact of breaking habits.

In this paper, we use a detailed panel of administrative data on individual tax payments as a reliable, unobtrusive outcome variable to assess the impact of the holidays. First, the randomized policy sets up a

²More than 25 percent of municipalities in Brazil, Colombia, Ecuador, and Uruguay and others in Argentina, Peru, and Mexico now raffle prizes to compliant taxpayers (Online Appendix Figure A1).

natural experiment: by comparing subsequent tax compliance of lottery winners to a sample of eligible non-winners, we can readily estimate the impact of winning a tax holiday. The impact, however, may be broader than this comparison suggests because the prize program might incentivize delinquent taxpayers to pay their taxes promptly to become eligible for the program. Second, we therefore utilized a complementary placebo-controlled field experiment to assess whether informing citizens about the existence of the lottery boosts tax compliance. We designed flyers we mailed in collaboration with the municipal government and stamped them with the municipal logo, which enhanced the realism of the treatments. Finally, we draw on qualitative interviews with municipal officials and tax holiday winners and a survey experiment to assess the effect of the lottery on attitudes about the fairness of the tax system. Placebo outcome tests using data on ineligible taxpayers as well as balance tests on prognostic covariates validate key design assumptions.

We find that far from increasing tax revenues, the tax holiday policy diminished them. Among taxpayers eligible for the lottery, winning the holiday substantially reduced later tax compliance, an effect that lasts on average for up to two years. Moreover, we find no evidence of a compensating incentive effect: while our informational interventions prompted citizens to view their online tax bills, positive information about the lottery has no greater impact on intended or actual compliance than a simple reminder to the placebo control group.

Our evidence suggests that the negative effect arose because the holiday disrupts the habit of paying taxes. (1) We find no negative effects for the one municipal tax for which winners continue to pay a reduced amount of tax during the holiday—and thus for which the habit of payment is not in fact interrupted. (2) The holiday does not decrease compliance among taxpayers signed up for automatic debits, whose payments mechanically restart after the holidays. Only for manual taxpayers does the interruption have a negative impact. (3) The negative effects are also less persistent for taxpayers with a greater “stock” of prior compliance habit. This pattern of effects is consistent with a theory of habit we develop but at odds with several alternative explanations emphasizing broader shifts in attitudes or other behaviors. In the theoretical framework we develop, compliance has feedback effects: paying taxes in one period affects the stock of compliance habit and therefore influences the propensity to pay in later periods. Disrupting compliance behavior can therefore have lasting consequences but these also decay faster among taxpayers with a history of prompt payment.

Our findings provide insights into the behavioral foundations of tax compliance, and they illustrate how habits contribute to patterns of citizen-state interaction outside the more widely studied realm of electoral participation. Gerber, Green and Shachar (2003) show that voters randomly assigned to receive a get-out-the-vote message in one election were significantly more likely to turn out in the next, while Meredith (2009) finds that eligibility to vote in a past presidential election—as determined by whether a voter was just over or just under 18 years old—increases the probability of participation in the subsequent election. Coppock and Green (2015), combining data on millions of voters from several experiments and natural experiments, similarly find persistent effects of voting (see also Aldrich, Montgomery and Wood 2011; Ansolabehere, Hersh and Shepsle 2012). Our research, in contrast to these studies of habit formation, highlights the importance of habit destruction—and suggests that the effects of policy interventions are conditioned by compliance histories and practices.

The results also have implications for broader dynamics relevant to our understanding of state capacity. Social scientists have focused on how development outcomes can be locked-in by strategic equilibria, or by increasing returns and path dependency (Pierson 2000). Our evidence underscores the importance of reserves of habit and suggests how self-reinforcing practices may foster broader patterns of tax compliance. By shaping patterns of tax payment, habit may in turn impact other important outcomes, such as political accountability (Bates and Lien 1985; Ross 2004; Paler 2013).

From a policy perspective, our study also provides insights into concrete ways that governments may boost tax capacity—and potential unintended consequences of programs that seek to do so. Our findings particularly highlight that the delivery mechanism through which positive incentives are provided may matter greatly. Carrillo, Castro and Scartascini 2017 found that when an Argentine municipality rewarded 400 randomly selected compliant taxpayers with the construction of a sidewalk and public recognition in local media, tax compliance increased both among lottery winners and their neighbors (see also Pomeranz and Vila-Belda 2019). Yet such programs do not disrupt the habit of compliance as Montevideo's did. Indeed, after we presented our findings to municipal officials, Montevideo changed its policy to offer cash rebates instead of tax holidays.

We return to broader implications of our findings in the conclusion, after discussing theory and hypotheses (Section 2), describing our empirical strategy (Section 3), and presenting our tests and results (Section 4), discussing evidence for the role of habit and alternative explanations.

2 Why do people comply? Background and theory

Why do people comply—or fail to comply—with their tax obligations? In a standard formalization of the compliance problem, taxpayers compare the utility of evasion to the cost of punishment, discounted by its probability (Allingham and Sandmo 1972). Let y be an asset value, t be the annual tax rate, and z be the annual amount of taxes due; with full nonpayment, $z = ty$. The expected net benefit of evasion is then $z - pc$, where c is the penalty for nonpayment and p is the probability of punishment. In the context of municipal taxes, the cost of punishment c could include (1) fines and interest charges for delayed payments, and ultimately (2) losing one’s house or other property.

To explain systematic non-compliance, researchers often focus on the low value of p , due to the difficulty of sanctioning tax evasion (Bates and Lien 1985). In the context of compliance with municipal taxes in Latin America, governments may know taxpayers’ obligations with certainty, for instance, on the basis of appraised property values; the government’s problem reduces to the apparently easier task of compelling punctual payments—somewhat akin to that faced by a credit-card company collecting debt from delinquent cardholders. Yet, enforcing punishments for non-payment still poses substantial challenges. The political unpopularity of sanctions can be constraining in municipalities, where face-to-face negotiations are common and selective enforcement (or “forbearance,” see Holland 2016) may prevail. Our interviews with officials in several Latin American cities suggest that municipalities only rarely seize and auction properties on which taxes are owed; at most, an embargo is placed on a property so that it cannot be privately sold until debts are cleared (Online Appendix Section A).

Given weak enforcement, the right question may be not why many people do not pay taxes—but why anyone does. To be sure, prize programs such as Montevideo’s may shape the direct material benefit to paying taxes: under Montevideo’s lottery we describe in detail later, if taxpayers have promptly paid their year’s taxes, eligible taxpayers win a year free of tax payments with probability 1/10,000 in any lottery. However, one might expect these strictly material rewards to have a limited impact for a rational, risk-neutral taxpayer. The expected monetary gain from paying a year’s taxes is (without discounting or probability of punishment for nonpayment) $(1/10,000)z - z$.³ Thus, a taxpayer

³The average property value is about US\$36,035 (956,000 pesos), and the average annual property tax is about US\$265 (7,044 pesos). The expected value gross of the tax payment is then US\$265/10000, or about three cents.

has to pay a year's worth of taxes in order to gain, with probability 1/10,000, a year free of tax payments. To be sure, casinos would make no money if gamblers did not take bets with negative expected values—but this does not look like a promising deal for the taxpayer.

Yet, perhaps expressive or material *benefits* of paying taxes can induce compliance. To extend the standard framework, we might allow an expressive benefit b to influence the utility of tax payment. Studies of taxation suggest a broad view of what b might include—for example, the strength of ethnic boundaries (Lieberman 2003), the degree of altruism and fairness (Scheve and Stasavage 2016), or the extent of norms of cooperation that induce “quasi-voluntary” compliance (Levi 1988). Reward programs such as Montevideo's may influence perceptions of the fairness, equity, or transparency of the tax system, which could also shape compliance behavior by boosting b .

2.1 Habit disruption as a cause of non-compliance

This extended framework is still incomplete, however—and leaves unexplained several key features of tax payment data that we explore empirically below. Our pre-analysis plan considered as one among several possible effects that by breaking the habit of compliance, Montevideo's tax holiday could disrupt winners' later punctual tax payments (Mechanism 1B.2 in the PAP, Online Appendix Section G). By “habit,” we do not mean a behavior that is necessarily automatic or unthinking. Social psychologists have developed an understanding of habit as involving repetition of a response under similar conditions, so that the response tends to recur when those conditions occur (Wood and Neal 2007). This leaves open the specific cognitive process through which the repetition of behavior induces habit. The sociologist Pierre Bourdieu's related notion of *habitus* as “a system of acquired dispositions” emphasizes that experience breeds practice, and thus a ‘feel for the game’ that guides action (Bourdieu 1990, 13). At the time we fielded our study, the majority of taxpayers paid their taxes in person at kiosks. They also received bills in advance of each payment period. A compliance habit can form from the repetition of payment under similar conditions, involving the arrival of a tax bill and a trip to the local kiosk. For eligible winners, this is the kind of habit that the tax holiday appears to have disrupted.

Here, we enrich the standard individual compliance model to analyze further the role of habit. Though ex-post (not registered), this extension of the model generates additional implications that can

illuminate and explain results of our pre-registered tests. Importantly, by providing a framework that helps us better to understand the dynamics of habit interruption, the model provides a point of contrast with several alternative hypotheses that we pre-registered and for which we assess evidence later.

Suppose that compliance with taxes is partly habitual. For example, let $\gamma_t = 1$ if a taxpayer complies at payment period t . A framework that accommodates habit formation allows the utility of compliance at t to depend positively on $\theta\gamma_{t-1}$, where $\theta \in (0, 1)$ captures the sensitivity of current tax payment to the previous period's compliance behavior. Note that by a recursive argument, the “stock” of compliance habit at time t is given by $S_t \equiv \sum_{i=1}^{\infty} \theta^i \gamma_{t-i}$.⁴ Thus, the history of payment influences the stock of habit $\theta\gamma_{t-1} + \theta^2\gamma_{t-2} + \theta^3\gamma_{t-3} + \dots$, with more recent behavior weighted more heavily. Then $\gamma_t = 1$ if

$$b_t + lz_t + S_t > z_t - p_t c_t + v_t, \quad (1)$$

and otherwise $\gamma_t = 0$. The left-hand side of (1) gives the benefit of paying taxes: the sum of the expressive reward b_t , an expected material prize lz_t where $l \in (0, 1)$ is the probability of winning the lottery, and the stock of habit S_t . On the right-hand side, non-compliance implies keeping the unpaid taxes z_t but at the expected cost $p_t c_t$. Then the probability of compliance in time t (i.e. $\gamma_t = 1$) is $F(b_t + S_t - (1-l)z_t + p_t c_t)$, where F is the cumulative distribution function of a shock v_t . We add time subscripts to emphasize that these parameters (save for the probability of winning, which is constant in every lottery) may be time-varying. Indeed, they may be functions of the tax holiday policy or of winning the lottery, though for notational simplicity we do not explicitly write them as such. The benefits of tax payment are positively related to the history of past compliance, or the stock of habit.

Importantly, this model of habit generates several implications consistent with our data and findings we present later—but inconsistent with several alternative theories. First, there can be persistent effects of shocks to tax compliance. Thus, exogenous interventions—such as those that switch past compliance ($\gamma_{t-1} = 1$) to non-compliance ($\gamma_{t-1} = 0$)—can produce a negative impact on tax compliance not only in the immediately subsequent period but also beyond it. After all, compliance at period t is a function of compliance at $t-1$; but compliance at period $t+1$ is then a function of compliance at t . An intervention that reduces compliance at $t-1$ and therefore diminishes the stock of compliance habit has “knock-on”

⁴Schmitt-Grohe and Uribe (2008) use a similar approach to study the equity premium puzzle.

effects on compliance at $t + 1$. Disrupting the habit of tax compliance through an intervention like the tax holiday can have lasting effects on compliance.

Second, however, there is naturally a decay in the effect of such habit disruptions. Consider a taxpayer who has always complied prior to $t - 1$. Such a taxpayer has a large positive stock of habit S_t and thus temporary disruptions have a relatively smaller impact on behavior. Also, should the taxpayer comply after the interruption, she will be more likely to continue to do so than those with a smaller reserve of habit. Taxpayers with large S_t are therefore more likely to endogenously recover the habit of paying taxes. Moreover, such a taxpayer may also be likely to have parameter values (p , c , or b) that favor compliance, since the underlying propensity to pay taxes, net of the stock of habit, is also likely to shape compliance behavior—and also therefore feed back into S_t .⁵ Thus, for many future realizations of ν , such a taxpayer will again comply, which will then foster continued compliance.

Finally, habit partitions taxpayers into a set that typically complies and a set that does not, and these compliance histories condition the effects of interventions. For taxpayers with large stocks of past compliance habit (high S_t), the inequality in (1) is more likely to be satisfied, and paying taxes may usually be the preferred option; those for whom S_t approaches zero are more likely not to comply. Moreover, such a partition is self-sustaining due to the feedback effects. To be sure, since compliance is a random variable that depends on ν_t , some taxpayers with high S_t might fail to comply on some occasions. Yet, the effect of an intervention that generates non-compliance at $t - 1$ decays more rapidly for taxpayers who possess a strong prior habit of compliance. Behavioral repetition thus contributes to the emergence and persistence of “good” and “bad” taxpayers.

We illustrate these dynamics with a computer simulation that shows the impact of a three-payment-period interruption of payments analogous to the tax holiday. We compare two types of taxpayers with identical values of the parameters (b , z , p , and c) and thus the same underlying propensity to comply, but who differ in the stock of habit. Thus, we manipulate the history of compliance in the 47 payment periods prior to the disruption, so that “Imperfect Past Compliers” paid taxes with probability 0.4 in each period, and “Perfect Past Compliers” paid with probability 1. This approach allows for construction of a different stock of habit among otherwise identical taxpayers. We then exogenously

⁵We do not subscript as b_i , yet it is natural that the expressive benefit of paying taxes varies across taxpayers.

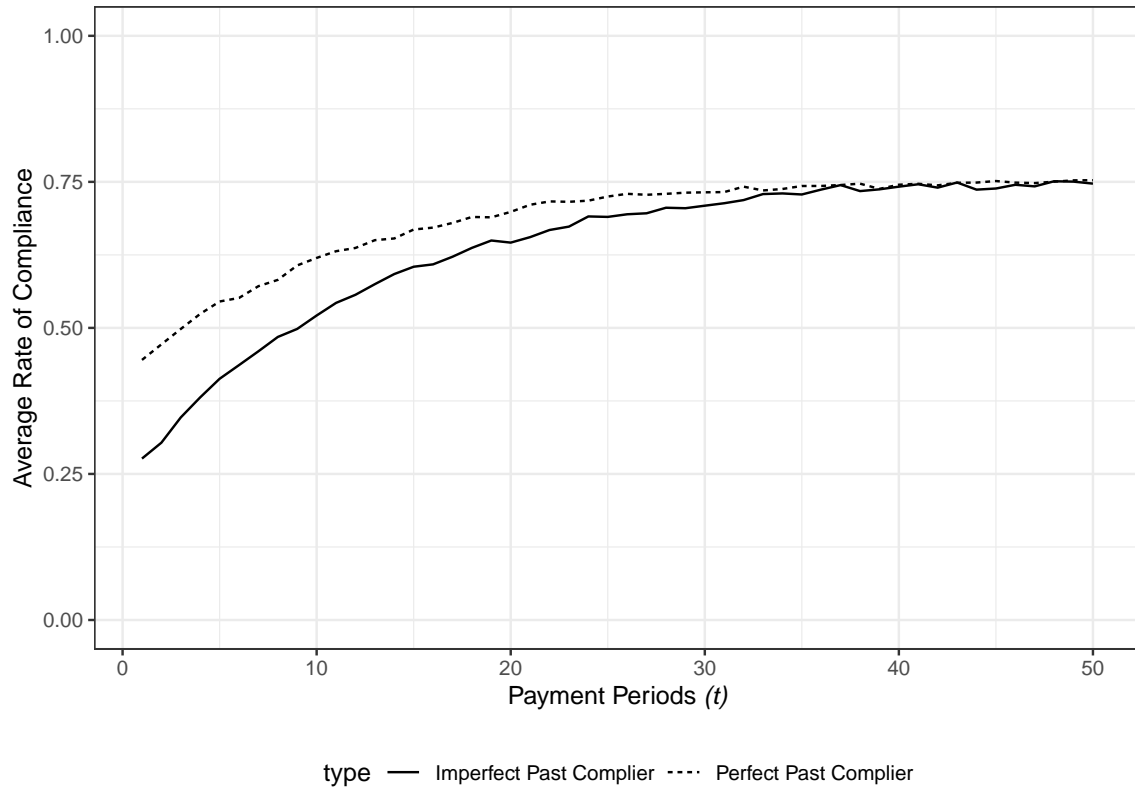
set compliance equal to zero for three payment periods. Finally, we allow tax compliance in every subsequent period to be determined according to equation (1). Thus, the mean-zero random variable v_t is realized and the taxpayer makes a compliance decision, given the parameter values and her stock of habit at time t . We conduct the simulation with 20,000 each of the Imperfect Past Compliers and Perfect Past Compliers (see Online Appendix Figure A22 for the R code). Figure 1 plots the average compliance rate for the two types of taxpayer after the three-period interruption ending at $t = 0$.⁶

The simulation reveals several interesting results. First, exogenous changes to tax compliance behavior—such as interventions like a tax holiday that switch past compliance to non-compliance—can have lasting effects. In this simulation, with $\theta = 0.7$ and thus substantial but not extreme dependence of the stock of habit on preceding payments, it takes many periods for taxpayers to recover to a stable level of average compliance. Second, there is decay in the effect of the disruption over time. Finally, however, the stock of habit matters for the rate of recovery: taxpayers with perfect past compliance and thus a greater stock of compliance habit recover more quickly than those with imperfect past compliance histories. Our framework therefore suggests that exogenous interruptions in the habit of payment can have persistent negative consequences for compliance—but also underscores the importance of the endogenous reserve of habit in shaping compliance behavior. The observable implications of this theory differ from several alternative explanations we consider below.

Testing such a theory of habit is usually very difficult. Choices over tax compliance are likely to be endogenous to attributes and behaviors that covary with habit. Indeed, in our model, taxpayers with greater underlying propensities to pay (net of habitual forces) are likely to pay at greater rates—and therefore build up a greater stock of habit. Thus, habit will be confounded empirically by the underlying propensities. Taxpayers who choose not to comply of their own accord will not provide a valid counterfactual for those who remain compliant. We therefore require an intervention that exogenously switches past compliance to non-compliance. The tax holiday lottery to which we turn in the next section provides just such a randomized natural experiment.

⁶Here $z = 2$, $p = 1$, $c = 1$, $b = 0.05$, $\theta = 0.7$, and $v_t \sim N(0, 1)$; results are qualitatively similar with other values.

Figure 1: **Simulation: The Stock of Habit and the Impact of Interruptions**



The figure shows simulated average tax compliance after a three-period tax holiday.

2.2 Registered hypotheses

Our registered pre-analysis plan (PAP) prospectively outlined a study that would allow us to (1) estimate the impact of one of Latin America’s earliest prize policies designed to reward tax compliance; (2) assess mechanisms through which the policy might work or fail to work; and (3) compare the impact of information about positive and negative inducements (Online Appendix Section G).

Our first and major aim was to assess the impact of winning the lottery, among eligible tax payers (PAP, Hypotheses 1A-C) and also to assess whether providing information about the lottery to ineligible as well as eligible taxpayers affects their attitudes and tax payment behavior (Mechanism 1A.1-1A.2). We hypothesized that the policy might shape compliance not only through the monetary incentive

provided by the tax holiday but also by shifting other parameters of the compliance framework. For example, by allocating benefits through binding public criteria (Stokes et al., 2013), the lottery might boost taxpayers' faith in the equity and transparency of the tax system and increase the expressive benefit of compliance b (Hypothesis 1A). We also registered hypotheses regarding the effects of negative incentives—e.g., of priming the probability p of sanctions c for non-compliance (Mechanism 2A.1).

Finally, we considered mechanisms through which the lottery could lead to no change or even *decrease* tax compliance. Most relevant to our findings, we anticipated that taxpayers who stop paying taxes for a year may fall out of the habit of paying taxes (Mechanism 1B.2). We also pre-specified tests of heterogeneous effects conditional on taxpayers' compliance history (Hypothesis 2C) and costs of compliance (Hypothesis 2D) as well as by type of tax and by whether taxpayers pay manually or automatically (second operationalization of a test of Mechanism 1B.2). To operationalize tests of these hypotheses, in our first amendment to the PAP, we registered code to analyze mock data, a practice pioneered by Humphreys, de la Sierra and van der Windt (2013). We include in the Online Appendix results from running the full set of pre-specified analyses (Section G.4), as well as reproduction code specifically for results presented in the paper and sections B-D in the online appendix (Section F).

3 Research design

The municipal government (*Intendencia de Montevideo*—IM) hoped to reward punctual taxpayers with its tax holiday policy—an especially important goal in a context of previous amnesties for delinquent tax payers. The four taxes eligible for the holiday—property, vehicle, household (“head”), and sewage taxes—comprise the great majority of Montevideo’s revenues: they totaled 92 percent of own revenue in 2021 and 2022, as well as 83 percent of total revenue including central government contributions.⁷

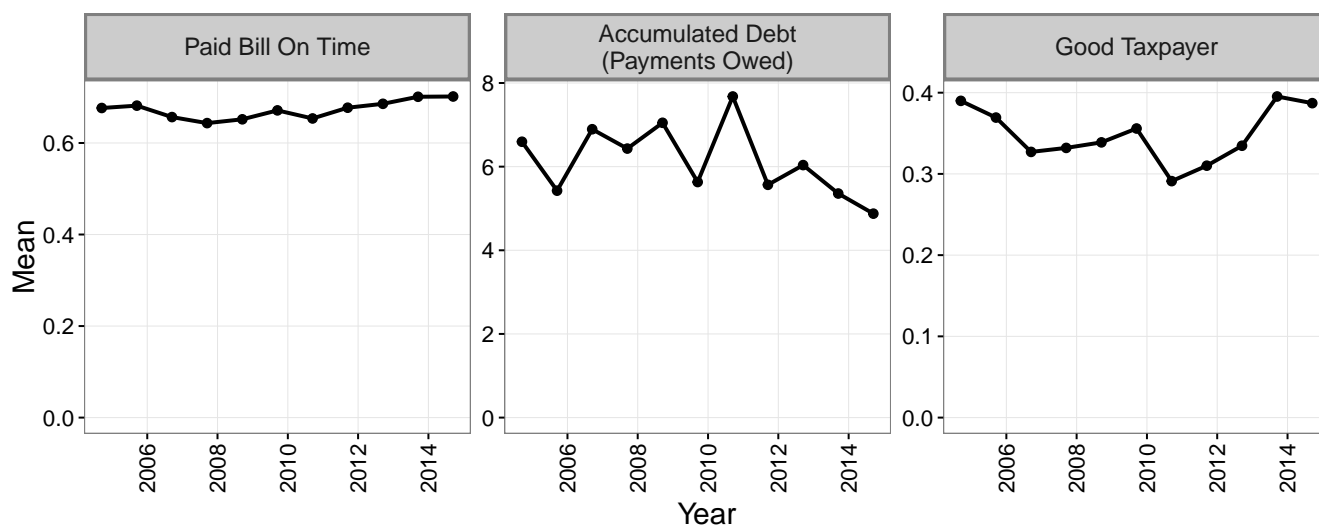
These taxes suffer from pervasive non-compliance. Figure 2 shows payment patterns between 2000 and 2014, based on our random sample of 9,297 property, vehicle, head and sewage tax accounts. The relatively low rate of prompt bill payment (left panel) results in substantial past bills owed per household (middle panel). The municipality also rates only a minority of taxpayers—40 percent or fewer in each year—as “good taxpayers” based on being current on payments over the previous year

⁷See <https://montevideo.gub.uy/sites/default/files/biblioteca/presupuestoin2021-2025.pdf>, accessed March 8, 2023.

(right panel), a designation that establishes eligibility for the tax holiday. (We do not use the term in a normative sense, but the municipality calls its policy “Lottery for Good Taxpayers”).

These averages however mask substantial temporal stability among compliant and non-compliant taxpayers, respectively. Of those classified as good taxpayers at a given moment in time, over 75 percent were thus classified ten payment periods later.⁸ By contrast, only 11 percent of “bad taxpayers” were classified as good taxpayers after ten subsequent periods. Weak monitoring capacity cannot explain these outcomes: with these municipal taxes, the state knows with certainty the amount of tax owed yet still fails to induce prompt compliance among substantial numbers of taxpayers. Reserves of habit, by contrast, could play a role in separating compliant from non-compliant taxpayers. However, the difficult question is whether this reflects any causal effect of habit, since good and bad taxpayers could differ in many ways besides their compliance histories. We therefore turn to our experimental designs to estimate impacts of the tax holiday program and to analyze the role of habit.

Figure 2: **Municipal Tax Compliance Over Time (Montevideo, Uruguay)**



3.1 Natural experiment

The municipal government uses the results of the otherwise unrelated Uruguayan National Lottery to select taxpayers for holidays. Thus, it selects as provisional winners those tax account numbers,

⁸See the first column of Table 4.3, discussed later.

the final four digits of which correspond to the winning number of the relevant National Lottery.⁹ In February 2009, for example, the winning National Lottery number ended in 8662. The municipality thus identified all taxpayer account numbers also ending in 8662, for each of the four types of taxes—property, vehicle, head, and sewage—and screened in eligible taxpayers who owed no past taxes and had paid on time over the previous year.¹⁰ The municipality then sent a letter to the addresses associated with each eligible winning account indicating that taxpayers should come to City Hall to claim their year-long tax holiday. Holidays are granted for the specific tax for which the taxpayer was selected, based on his or her tax-specific account number.¹¹ Head and sewage taxes are paid six times a year (in February, April, June, August, October, and December) and the municipality grants holidays in connection with each payment period; for vehicle and property taxes, taxes are paid and holidays are issued three times a year (in March, July, and November).

To assess the causal impact of payment interruptions, we first identified all taxpayer accounts that were randomly selected in each lottery held since 2004. The treatment group in our natural experiment consists of winning taxpayers who were eligible to claim the tax holiday: these were all “good taxpayers” based on having made timely payments in full in the year prior to the date of the lottery in which their account number was selected. Constructing the control group requires some care: the right counterfactual for winners of each past lottery consists not of currently eligible non-winners but rather those who were eligible to win (but did not) as of each lottery. We therefore randomly generated a four-digit number, different from the winning number, for each lottery since 2004, then screened in all taxpayers whose accounts ended in these numbers and who were eligible as of the corresponding lottery. Our procedure therefore mimicked the random process that created the treatment group.¹² Due

⁹The randomization occurs through the selection of balls from an urn, described in Spanish at http://www.loteria.gub.uy/Juego_Loteria.php and <http://www.montevideo.gub.uy/sorteosBP/pages/sorteosBuenosPagadores.xhtml>. Accessed January 16, 2017; see Online Appendix Figure A2.

¹⁰For each type of tax, as noted earlier, the probability of winning a given lottery is thus $1/10,000$, since $10^4 = 10,000$ is the number of possible final four-digit numbers.

¹¹Taxpayers have different account numbers for each type of tax, and these accounts are not readily linked.

¹²Currently eligible non-winners include those who were eligible and ineligible as of the date of each past lottery,

to the randomization, the treatment and control groups should be balanced up to chance error on all observable and unobservable attributes—save for the presence or absence of the payment interruption. We confirm balance on highly prognostic covariates (e.g. past payment histories) in Appendix Table A1 and in the graphical analysis in Figure 3 in the next section.

Table 3.1 describes this natural experiment. The table depicts the size of treatment and control groups (in bold font, top four rows), distinguishing between four different types of taxes for which tax holidays are awarded. The table shows the numbers assigned to the two groups by virtue of having a winning or non-winning account number. In addition, we have payment data for taxpayers whose account numbers match those in our treatment or control groups but who were ineligible to win as of the corresponding lottery on the basis of their payment histories (plain font in Table 3.1). Although we only use data for eligible taxpayers to estimate treatment effects—since only those taxpayers could claim a tax holiday—we exploit data on ineligible taxpayers for placebo and balance tests. Note that the cells of Table 3.1 are themselves self-weighting random samples of taxpayer accounts in Montevideo. We can thus use them to characterize features of the whole taxpaying population.¹³ Our partners in Montevideo’s tax bureaucracy provided a time-series panel of payment data (2000-2014) for all randomly selected account numbers.

We are interested in the impact of winning the lottery at time $t = 0$ on tax compliance at time $t = 1$ (defined as the payment period when payment obligations resume after the end of the tax holiday) as well as at times $t = 2, 3, 4, \dots$. The average causal effect on compliance at post-treatment period t for the N subjects in the study group depicted in Table (3.1) is $ACE_t = \frac{1}{N} \sum_{i=1}^N Y_{i,t}(1) - Y_{i,t}(0)$, i.e., the difference of average potential tax compliance if eligible taxpayers win the holiday and if they do not. Here, each ACE_t is the effect of treatment assignment, per the intent-to-treat principle (Dunning 2012, 88). Because only 72 percent of notified winners actually claimed the exoneration, we are also interested in the Complier Average Causal Effect at time t ($CACE_t$), which is the ACE divided by

while the treatment group only includes eligible winners at each date. Since potential tax compliance is very plausibly related to eligibility, our procedure avoids bias that this asymmetry could otherwise risk.

¹³For example, we combine eligible and ineligible taxpayers with non-winning numbers for the random sample of 9,297 taxpayer accounts in Figure 2.

the proportion of Compliers in the study group ((Gerber and Green 2012: 39-43).¹⁴ (Here we use “Complier” in the statistical sense of compliance with treatment assignment, rather than compliance with tax obligations).

Table 3.1: **Natural Experiment: Sample Sizes**

Tax	Taxpayer Type	Non-Winning Tax Account Number	Winning Account Number	Study Group Totals
Property	Eligible Taxpayers	1354	1339	2693
Vehicle	Eligible Taxpayers	375	391	766
Sewage	Eligible Taxpayers	404	452	856
Head	Eligible Taxpayers	1041	1007	2048
Property	Ineligible Taxpayers	1225	1211	2436
Vehicle	Ineligible Taxpayers	1924	1899	3823
Sewage	Ineligible Taxpayers	939	915	1854
Head	Ineligible Taxpayers	2062	2083	4145
All Taxes	Eligible Taxpayers	3174	3189	6363
All Taxes	Ineligible Taxpayers	6150	6108	12258

The table depicts winning and non-winning account numbers in our sample. Rows used to estimate treatment effects are in bold; non-boldded rows are used for placebo tests.

The randomization afforded by the natural experiment allows for simple and transparent estimation of these causal effects, either aggregating across all types of taxes and all lotteries or disaggregating by type of tax. We in fact have a series of mini-natural experiments, with random assignment blocked by individual lottery. However, because the winning lottery number is drawn in the same way in each lottery, the probability of assignment is the same in every block. Due to our sampling strategy, in expectation the treatment and control groups are also the same size. We can thus readily estimate each ACE_t with a simple difference of means, in addition to auxiliary strategies such as differences in differences. Consistent estimation of each $CACE_t$ requires an exclusion restriction. Results from our field experiment suggest that information about the tax holiday may not directly influence payment behavior, suggesting that excludability may be satisfied. However, our results are robust to reliance on estimation of the ACE_t , which does not depend on this assumption. We also note that because winners and non-winners are both randomly drawn from the population of taxpayers, estimation of effects in

¹⁴Many of the winners who did not claim exonerations are likely small firms, which are not eligible for the holiday.

our sample of taxpayers also readily estimates effects in the population.

3.2 Field and survey experiments

We also draw on a large field experiment in which we provided randomly selected taxpaying households in Montevideo—including those both eligible and ineligible for the lottery at the time of the experiment—with information about the tax holiday. Specifically, we collaborated with the municipality to design and mail flyers printed with experimentally varied messages. We focused on property taxes both because of the importance of this tax and to eliminate potential sources of heterogeneity that would decrease statistical power. Our baseline reminder serves as a placebo control:

Dear neighbor: We want to remind you that **the second payment of property taxes is due in July**. If you have not received your bill, you can obtain a duplicate copy on our web site (www.montevideo.gub.uy).¹⁵

Our next condition repeats this reminder but adds information about the tax holiday lottery:¹⁶

The municipal government of Montevideo wants to reward good taxpayers. **If you pay on time, you will be automatically entered in a lottery to win a year free of property tax payments**. Lotteries occur every other month of the year in conjunction with the National Lottery. The winners will be duly informed and the results of the lottery will be published on the web site of the city government. **You can be the next winner!**

The experimental realism of our treatments is substantial: when folded for mailing, the municipal logo is visible, and in fact the flyers appear identical to municipal tax bills before being opened.¹⁷

We worked with the municipal bureaucracy to draw a stratified random sample of eligible and ineligible taxpayer accounts, the former being those that could potentially win the tax holiday as of the date of the intervention (Table 3.2).¹⁸ The population from which this sample was drawn should be

¹⁵We reproduce the Spanish-language flyers in the Online Appendix (Figures A11-A17).

¹⁶Flyers vary whether this information emphasized the individual (as here) or social returns of the lottery. Our pre-specified analyses often pool these into a single “reminder + existence of lottery” condition.

¹⁷The flyers were authorized by the municipality and our institutional review board under protocol 2014-04-6286.

¹⁸Per Online Appendix Figure A18, eligible and ineligible taxpayers are quite evenly spread throughout the city.

conceptualized as “all property tax-paying households with bills due in July 2014.” We also verified that none of our sampled taxpayers had actually won a lottery in the past, since our goal was assess the effect of informing taxpayers about the possibility of exoneration. (Note, however, that having won the lottery in the past—even in the immediately preceding year—does not disqualify a taxpayer from winning again). We then randomized sampled taxpayers to intervention groups. All eligible and all ineligible taxpayers had the same probabilities of assignment to each condition, but the probabilities differed across the two subgroups since a larger share of the eligibles were retained as pure controls. We account for these differing assignment probabilities by using inverse probability weights (IPW) for differences of means and also estimate effects of the treatment relative to the placebo control separately for eligible and ineligible taxpayers, the latter being our pre-specified analysis (Online Appendix Section G.5.3).

The treatment groups are statistically balanced on pre-treatment covariates (Online Appendix Table A3). We note that these pre-treatment covariates measure whether taxpayers paid on time in pre-treatment periods and are thus likely to be highly prognostic (i.e. predictive of potential outcomes) and thus provide stronger tests of as-if random assignment. We focus in Table 3.2 on the interventions that are theoretically relevant for alternative explanations, though these are a sub-set of treatments in a larger experiment that also included conditions reminding taxpayers about punishments for non-payment. However, we ensure that results are robust to corrections for multiple comparisons reflecting the full design, per our analysis plan (Online Appendix Figure A19). Our flyers were distributed in phases in June 2014, such that they would arrive approximately 8 days before tax due dates.

For outcome variables in the field experiment, we use two behavioral measures. First, using municipal tax data, we measure whether the account holder accessed his or her Web taxpayer account, for example, to print a duplicate bill (“Web Access”); we interpret this as a measure of intended compliance. Second, we measure whether the account holder paid her tax bill on time (“Paid On Time”). We measure both outcomes over the seven payment periods following our intervention (which includes over two years of data), allowing us to detect any medium-run as well as immediate effects.

Finally, we also use data from a separate survey experiment with respondents in several thousand households to assess whether the lottery acts as a signal about municipal capacity or the equity of the tax system. Thus, we compare respondents who were informed about the lottery—using language very similar to that on our flyers in the field experiment—to those who were told that the municipality “from

Table 3.2: **Field Experiment: Treatment Conditions and Sample Sizes**

Treatment condition	Sample of eligibles (Good taxpayers)	Sample of ineligibles (Bad taxpayers)
Control	N=7,243	N=3,412
Reminder of Taxes Due (Placebo Control)	N=1,532	N=2,080
Reminder + Information About Lottery	N=3,037	N=4,150
Study Group Totals	N=14,784	N=13,862

The table depicts assignment to a subset of treatment conditions in our field experiment (see Online Appendix Table A5 for the full design). We use inverse probability weights (IPW) to account for unequal assignment probabilities for eligible and ineligible taxpayers and also report estimated effects separately for each group. Total N=28,646.

time to time” selects good taxpayers and rewards them with a year free of tax payment. We thereby intend to assess the impact on attitudes of information about a transparent way of allocating prizes for compliance—the lottery—relative to a method with more potential for discretion. We gathered outcome data on five measures of attitudes towards the municipal government and the tax system (for further details, see Online Appendix Sections D and G.3).

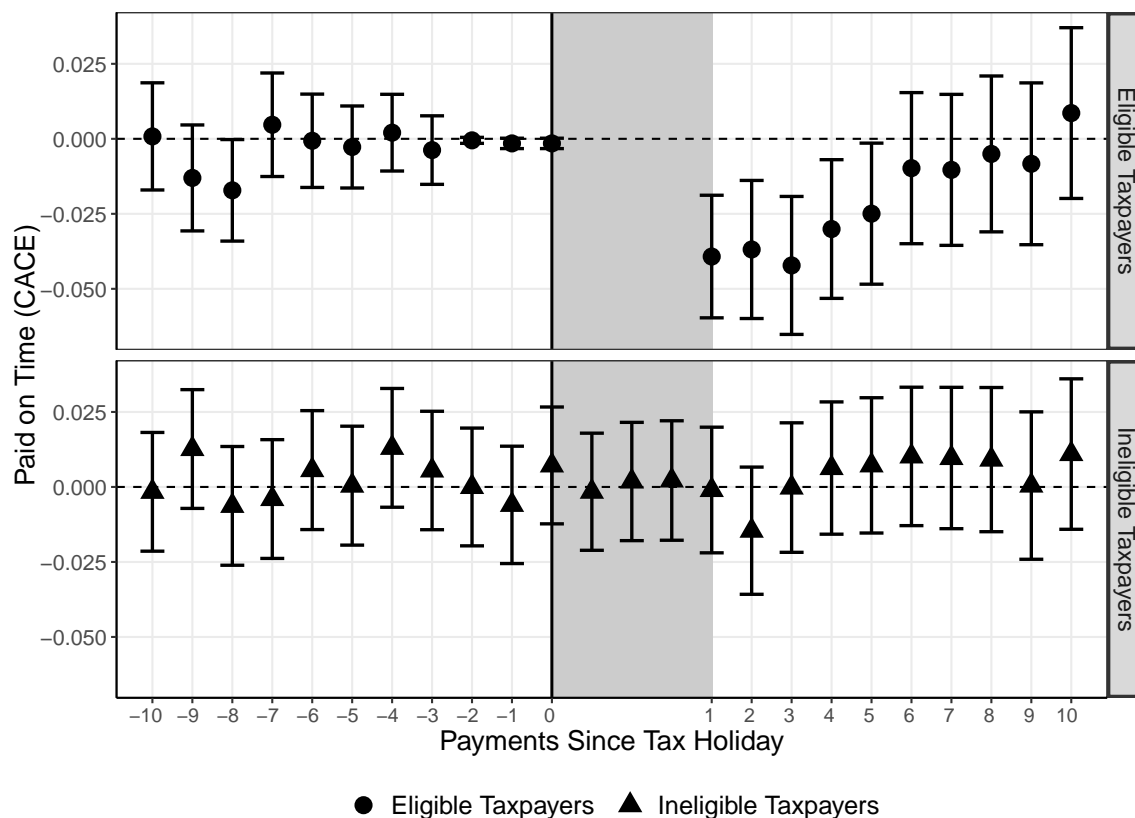
4 Results: the negative effect of tax holidays

Figure 3 graphically depicts estimates of treatment effects in the natural experiment. The dependent variable is the proportion of taxpayers who paid their bill on time at each tax payment period—an appropriate measure both for evaluating the effect of the tax holiday policy and specifically for interrogating the habit mechanism.

The figure’s vertical axis measures the difference in payment rates between taxpayers with winning and non-winning numbers, divided by the estimated proportion of compliers to form each \widehat{CACE}_t . The horizontal axis shows the number of payment periods elapsed before or after the particular lottery in connection with which a taxpayer was sampled into the study group.¹⁹ The number of payment periods per year varies by tax (with three annual payments of property and vehicle taxes, and six of sewage and head taxes), and as we have lottery data from 2004 to 2014, the date of winning varies across taxpayers.

¹⁹We report these results for a longer post-treatment periods in Online Appendix Figure A5.

Figure 3: The Negative Impact of Holidays on Compliance



The figure depicts balance tests, placebo outcome tests, and treatment effects in our natural experiment. The horizontal axis measures tax payment periods before or after the period of the tax holiday (grey vertical strip). The vertical axis shows the estimated complier average causal effect (CACE) for the proportion of taxpayers who paid on time at each payment period. Vertical lines show 95% confidence intervals.

Centering at $t = 0$ allows us to pool analysis across the four types of taxes, though we also disaggregate by tax type in Table 4.3. The grey vertical strip indicates the period of exoneration. Since the holiday can be delayed by failure to claim the exoneration immediately, we use the maximum time elapsed in the treatment group to define the end point of the grey strip. The number 1 on the horizontal axis indicates the first payment due after the holiday period; 2 is the second payment; and so on. Negative numbers indicate pre-treatment periods.

The top panel shows data for taxpayers who were eligible to win the holiday as of the date of the respective lottery, while the bottom panel shows those who were ineligible. Thus, black circles to

the right of the grey strip compare compliance behavior of eligible winners and non-winners after the conclusion of the holiday (top panel). Vertical lines show 95% confidence intervals around these point estimates.²⁰ There is no variation in tax payment behavior among eligible taxpayers at $t = -1$, $t = -2$, and $t = -3$ (and thus the confidence intervals collapse): taxpayers must have paid on time over at least the three previous payment periods to claim the holiday. Table 4.3 reports the \widehat{ACEs} , \widehat{CACEs} and standard errors at post-holiday payment periods 1, 5, and 10, as well as the average for all ten periods.

Tax holidays have a substantial, negative, and persistent effect on subsequent tax compliance. The effect lasts for five payment periods on average, or nearly two years for the property tax. The \widehat{CACE} is -4% percentage points in the first post-holiday payment period, and -3% percentage points pooling over the first ten periods (Table 4.3). Viewed against a control group non-payment rate of 7% in the first post-treatment payment period (a 93% payment rate among these eligible good taxpayers), the first estimate represents a treatment effect of 57 percent. We also find large and persistent effects for our two related primary outcomes, the accumulated number of past payments owed and a dichotomous indicator for whether a given taxpayer account is fully paid as of the due date, as well as persistence of null effects beyond ten post-treatment payment periods (Online Appendix Figures A4-A5).²¹ The significance of these estimated effects survives Bonferroni (1936) and Benjamini and Hochberg (1995) adjustments for comparisons across the primary outcomes (Online Appendix Table A3). Though effects on compliance are negative, they also eventually decay, as we discuss later in our examination of the habit hypothesis.²² We do not find significant effects for one outcome, total accumulated debt owed, perhaps because we could only measure this outcome in one payment period (July 2014); given the persistence but eventual decay in effects on compliance after about two years, we would not expect to

²⁰We estimate the mean and standard error for the groups with winning and losing numbers in each time period separately and use normal approximations for the confidence intervals (since the Ns are large).

²¹The outcome variable in Figure 3 and Table 4.3 measures whether a payment is made on time in each period, rather than whether the taxpayer is fully paid up.

²²We see a significant difference in estimated effects when comparing the first and second post-treatment years, though not the second and third (pp. 20-21 in Appendix Section G4); we also specified analyzing persistence by estimating effects at all post-treatment payment periods as in Figure 3 (p. 37 of original PAP, p. 33 of the first amendment, and p. 25 of second amendment; Online Appendix Sections G1-G3).

see effects on debt in 2014 for taxpayers who won the lottery between 2004 and approximately 2012.²³

Figure 3 also allows for graphical balance and placebo outcome tests that validate key assumptions of our design. First, for the balance tests: randomization implies that prior to $t = 0$, tax payment behavior should be identical in expectation for taxpayers with winning and non-winning numbers. Indeed, Figure 3 shows that differences in pre-treatment compliance are statistically indistinguishable from zero for both eligible and ineligible taxpayers.²⁴ For higher-powered balance tests, we also pool data on eligible and ineligible taxpayers, giving a total of up to 18,621 taxpayer accounts. We then test for balance on a series of prognostic covariates, including tax compliance at several pre-treatment payment periods (Online Appendix, Table A1). None of the p -values approach standard significance levels, indicating that the randomization worked as expected. The public nature of the lottery and its tie to the independent National Lottery also suggests the credibility of the randomization.

Second, the bottom panel of Figure 3 also permits a test of a placebo outcome for which an effect is “known” not to exist (Eggers, Tuñón and Dafoe Forthcoming): the impact of assignment to a winning number for ineligible taxpayers. The municipality does not inform such ineligible taxpayers that their number was drawn and they receive no benefit, so the treatment does not affect any of the parameters of our compliance model. Indeed, we find no discernible difference in post-treatment compliance between ineligible taxpayers with winning and losing lottery numbers (bottom panel). Online Appendix Table A2 reports formal statistical tests which reach the same conclusion. Assignment to a winning account number, absent a tax holiday, does not itself influence compliance behavior. Our field experimental results in section 4.2 further substantiate an exclusion restriction: among eligibles, being informed of the existence of the lottery does not shape behavior. Finally, our evidence in section 4.2 validates the core assumption of non-interference (Cox 1958). In brief, news of the lottery does not substantially travel to non-winners: only 8% of non-winners have heard of the existence of the lottery. In sum, our data are strongly consistent with random assignment and other identifying assumptions of our design.

²³The municipality only gave us anonymized accumulated debt data averaged by treatment category; so here we can estimate effects but do not have covariates and cannot look at variation by year of the lottery.

²⁴Figure 3 reports estimated Complier Average Causal Effects (\widehat{CACE}), but for balance and placebo tests, there is no difference between these and the simple difference of means.

Table 4.3: **Estimated Causal Effects of Tax Holidays**

Post Tax Holiday Payments	Control	\widehat{ACE}	$\widehat{SE}(\widehat{ACE})$	p -value	\widehat{CACE}	$\widehat{SE}(\widehat{CACE})$	p -value
	Group Average						
Payment 1	0.93	-0.03	0.01	0.00	-0.04	0.01	0.00
Payment 5	0.91	-0.02	0.01	0.04	-0.02	0.01	0.04
Payment 10	0.87	0.01	0.01	0.55	0.01	0.01	0.55
Payments 1-10	0.90	-0.02	0.00	0.00	-0.02	0.00	0.00
Payments 1-10 (Property)	0.92	-0.02	0.00	0.00	-0.03	0.01	0.00
Payments 1-10 (Head)	0.91	-0.01	0.00	0.01	-0.02	0.01	0.01
Payments 1-10 (Sewage)	0.94	-0.02	0.01	0.00	-0.03	0.01	0.00
Payments 1-10 (Vehicle)	0.75	0.00	0.01	0.80	0.00	0.01	0.80

The table shows the estimated average causal effects of treatment assignment (\widehat{ACE}) and complier average causal effects (\widehat{CACE}), as well as their standard errors and p -values, for the proportion of taxpayers who paid on time. The first four rows show effects for 1st, 5th, and 10th payments due after the end of the tax holiday, as well as their combination. The final four rows disaggregate the effect by type of tax. Significant p -values survive adjustment for multiple comparisons (Online Appendix Table A3).

Rather than fostering greater tax compliance among winners, winning the tax holiday lottery therefore inhibits it. Note there is a persistent negative effect on compliance—but also attenuation of the effect over time. Thus, the data are consistent with the knock-on effects of habit disruption discussed in subsection 2.1 but also with their eventual decay. This pattern is important for distinguishing habit from other potential mechanisms that might explain our findings, a topic to which we turn in subsection 4.2 after presenting further evidence on the role of habit.


4.1 The role of habit

Is disruption of habit the mechanism that drives the negative impact of tax holidays?

A first piece of evidence comes from variation in effects across types of taxes. Although we were unaware of this when we began our study, with the vehicle tax—unlike property, sewage, and head taxes—winners typically continue to pay a small amount during the putative holiday. The reason is that with the vehicle tax, unlike other taxes, payment is exonerated retroactively, so that the *previous* year’s payments are forgiven. Some taxpayers take the windfall as a refund, while others take it as a credit

against future payments. Yet even those who take credits typically owe minor amounts of vehicle fees in the following year, because the vehicle tax is often increased annually by small amounts. Winning the vehicle tax lottery thus does not typically involve an interruption of the habit of paying taxes.

Figure 4: **Treatment Effects By Type of Tax: Holiday vs. No Holiday**



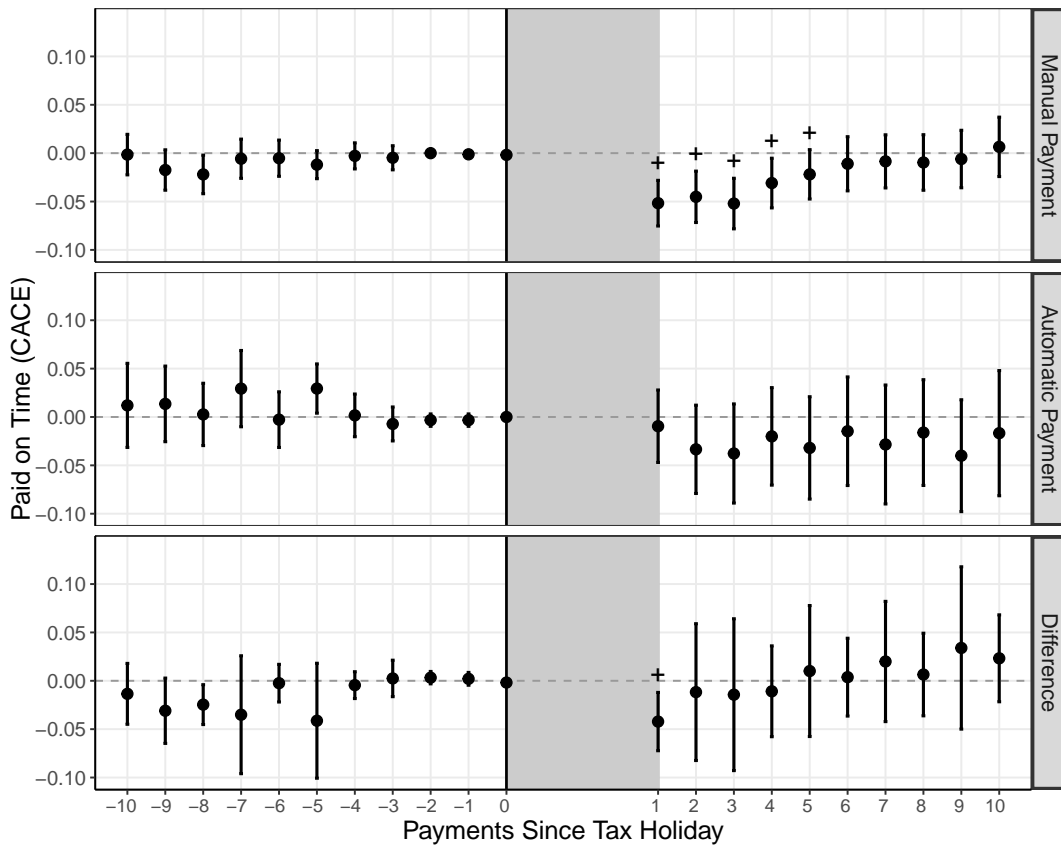
The negative effect only occurs when the payment habit is actually interrupted, in the “Holiday” taxes. A “+” above the confidence intervals indicates estimates that are significant in one-tailed tests. See also Table 4.3 and Online Appendix Figure A9 for pre-specified analysis of heterogeneous effects by type of tax.

If our argument about habit is correct, this should imply weaker or null effects for the vehicle tax. The top two panels of Figure 4 present effect estimates separately for the “holiday” (property, sewage, and head) and “no holiday” (vehicle) taxes, including only eligible taxpayers; the bottom panel presents tests for the difference of the effects.²⁵ Confidence intervals that do not overlap zero indicate statistical

²⁵We report the results in Figures 4-6 with longer post-treatment periods in Appendix Figures A6-A8.

significance of the estimates in two-tailed tests; here and in the next two figures, we also indicate with a “+” estimates that are significant in one-tailed tests.²⁶ Only when the payment habit is interrupted are the effects negative and significant (top panel). However, winning the lottery has no discernible impact in the no-holiday case (middle panel). The difference between these estimated effects is itself statistically significant for the first three post-treatment periods (bottom panel). To be sure, these heterogeneous treatment effects are not conclusive: vehicle taxes may differ from the other taxes in ways that are relevant for the effect of a holiday. However, they are suggestive of the force of habit.

Figure 5: Placebo test: Treatment Effects for Automatic vs. Manual Payers



The negative effect of the holiday only appears for manual taxpayers, not for those in automatic withdrawal programs. The figure depicts effects for the property, head, and sewage taxes. (Registered test of Mechanism 1B.2 in the PAP). A “+” above the confidence intervals indicates estimates that are significant in one-tailed tests.

²⁶Though our PAP discusses directional hypotheses for habit, we did not pre-register the one-tailed tests.

A similar but perhaps more powerful test comes from comparing automatic and manual taxpayers (Figure (5)). Around 21% of eligible taxpayers in our study group had enrolled in automatic payment plans at the date of the relevant lottery.²⁷ For these taxpayers, payments resume automatically after the conclusion of the holiday without any required action. The force of habit therefore could not conceivably operate to generate negative effects for automatic payers, allowing a placebo outcome test. In our original pre-analysis plan (p. 38), we anticipated using survey self-reports on who paid automatically.²⁸ We were in fact able to gather verified data on automatic vs. manual payment from the municipality. The holiday decreases compliance only among manual taxpayers (top panel). For the 1,190 automatic payers of property, head and sewage taxes, we find null estimated effects (middle panel). The difference in the estimated effects is also significant, though only in the first payment period (bottom panel). To be sure, taxpayers who pay automatically could differ from manual payers in various ways that could be related to compliance: manual taxpayers have higher property values.²⁹ Yet, this might bias against finding negative effects only for them, and we do not in any case find substantial heterogeneity of effects by property value (Online Appendix Figure A10).

Our theory in subsection 2.1 has an additional testable implication: compliance should bounce back more quickly for taxpayers with a greater stock of habit (per the simulation in Figure 1).³⁰ Thus, Figure 6 compares effects for “perfect past compliers” who had complied punctually over the 15 periods prior to winning the lottery—and who therefore should had a strong habit stock—to “imperfect past compliers” who were eligible to win the lottery at the relevant date but who had missed at least one payment over the previous 15 periods. The figure suggests this pattern: the negative effects last longer for the marginal taxpayers. Though we cannot reject the null that the effects are the same in the first four post-treatment periods, they are significantly smaller among perfect past compliers in periods five and six (bottom panel), and also in period seven in a one-tailed test.

Our qualitative evidence also lends insight into the nature of the payment interruption—and how it

²⁷This is 20% for property tax, 15% for sewage, and 26% for head; we do not have these data for the vehicle tax.


²⁸See the second operationalization of a test of Mechanism 1B.2 in the original PAP.

²⁹The difference in average property values is 2,088,634 vs. 1,142,571 pesos (2004 data).

³⁰We discussed this as Hypothesis 2C in the original PAP, but there the discussion is not focused on habit; our ex-post theoretical framework in subsection 2.1 suggests a different direction for the relative effect.

disrupts the compliance habit. We conducted 20 qualitative, open-ended interviews with taxpayers who won a property tax lottery in 2014-2015. The vast majority of the winning interviewees paid their tax bill at a local kiosk before winning the lottery and continued to do so after the tax holiday; the method of payment was unchanged for 18 of our 20 interviewees. Many winners recalled receiving a zero-balance bill from the tax authorities during the period of exoneration (and one interviewee showed it to us). In sum, the mode of payment and broader payment context did not change; and winners received notices from the municipal government reminding them that they are taxpayers.

Figure 6: The Stock of Habit: Perfect vs. Imperfect Past Compliers



The negative effect of winning is less pronounced for taxpayers with a greater reserve of payment habit. “Perfect Past Compliers” paid punctually in all 15 payment periods prior to winning the lottery, while “Imperfect Past Compliers” failed to do so in at least one period. Differences in estimated average causal effects are statistically significant at the 0.05 level for the first seven post-treatment payment periods. (Registered test of PAP Hypothesis 2C, though the PAP not discuss that test with respect to the habit mechanism). A “+” above the confidence intervals indicates estimates that are significant in one-tailed tests.

What changed instead is simply the fact of payment. Our evidence suggests that even though the holiday disrupted habits, it did not fundamentally shift taxpayers' vision of themselves as good taxpayers, nor their sense of the "rules of the game" guiding practice (Bourdieu 1990). Three past winners of tax holidays whom we interviewed missed payments after the end of their exonerations. In fact, none of them remembered having missed payments after the holiday.³¹ While habit need not be automatic, delays in payment after the exoneration may not result from strategic considerations, attitudinal shifts (which we explore next), or even intentional action. Whether aware of it or not, these good taxpayers with strong compliance histories appeared simply to get out of the habit of payment; yet they were able to recover the habit eventually.

In sum, our quantitative and qualitative evidence suggests that the negative effect of the tax holiday arose because winning disrupted the habit of paying taxes. The negative effects of disrupting compliance also decayed over time: consistent with our extended theoretical framework, good taxpayers subject to the disruption eventually recovered the habit of tax payment. This pattern is consistent with a habit mechanism—as indicated in our PAP and elaborated in our theoretical framework in subsection 2.1. However, it appears inconsistent with several alternative explanations focused on broader and perhaps more permanent shifts in attitudes and behaviors, as we discuss next.

4.2 Alternative explanations

4.2.1 Informational mechanisms

Our results also help to discount several alternate explanations. One possibility is that the receipt of information about the tax holiday lottery via the notification to winners from the municipality—rather than the interruption of payments—explains the negative effects. As we noted, in a household survey of a probability sample of taxpayers, we found that baseline knowledge of the program is low: only 8% of respondents identified the lottery as a municipal policy that rewards good taxpayers, while only 5% of survey respondents knew someone who had won the lottery.³² Thus, receiving an award letter from the municipality informs many winners of the existence of the policy. In the natural experiment,

³¹Interviewees CTA3112776, CTA356743, and CTA3408155.

³²These percentages are drawn from the placebo control group in our survey experiment (N=412).

the “habit disruption” treatment is therefore bundled with an “information” treatment.

Conceivably, learning about the policy could produce negative effects on future compliance. For example, taxpayers might interpret the new information as a negative signal of municipal enforcement capacity (p in the model) or the attractiveness of paying taxes (b in the model). They might construe the fact that the government holds a lottery to reward good taxpayers as an indicator that it has a hard time eliciting compliance—thus inferring that by complying, they are “suckers” on whom non-compliers are free riding. A priori, of course, many other informational effects might suggest *positive* effects on compliance. For instance, the lottery could boost perceptions of the transparency or equity of the tax system, increasing willingness to pay.

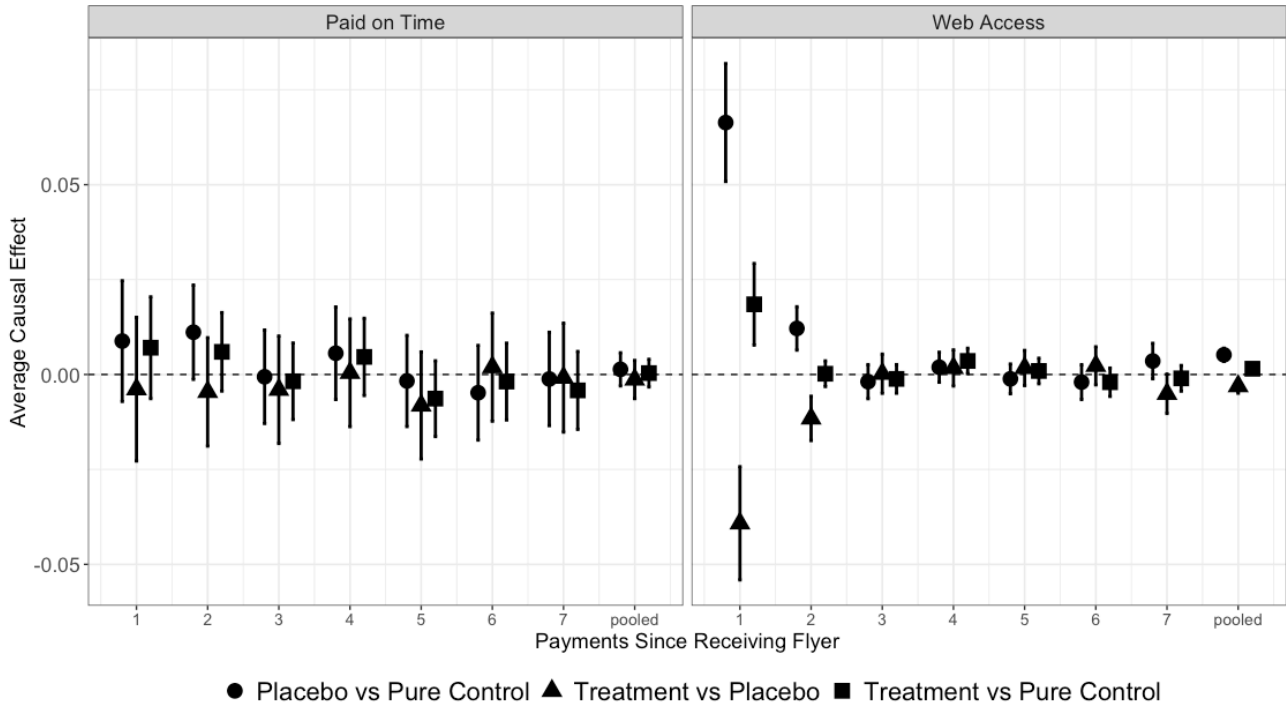
We use our field experiment to assess the effect of information about the tax holiday. As described in section 3.2, treated households were sent a flyer from the municipal government indicating that they could win a tax holiday if they pay taxes on time. In the field experiment (unlike the natural experiment), information about the holiday is not bundled with habit disruption, allowing a cleaner evaluation of the effects of the former. However, this treatment might not activate precisely the same informational mechanisms as a letter indicating taxpayers should claim a holiday (as in the natural experiment). We therefore view the field experiment as providing helpful corroborating evidence, even if it is complicated to rule out alternative mechanisms completely using variation in the treatments.

We find that information has no apparent effect on actual tax compliance (Figure 7). With the Paid On Time outcome measure (left panel), we find no effect of information, either in the first period after the intervention or over the longer period following it. We also find mostly null effects when considering Web Access as the relevant outcome (right panel). The treatment increased Web Access during the first two payment periods but only when compared to the pure control; the placebo treatment boosts Web Access in those periods more than the information treatment does, so the difference between the treatment and the placebo is negative. Information about the tax holiday benefit caused similar patterns of null effects among ineligible and eligible taxpayers considered separately.³³ We note that negative incentives if anything appear more effective: among delinquent taxpayers, a message about individual punishments boosted actual tax compliance by around 4 percentage points, relative to a pure control

³³Online Appendix Figure A21

group that received no flyers.³⁴

Figure 7: **Field Experiment: Effects of Information About the Tax Holiday on Compliance**



The figure depicts the effects of a reminder of an upcoming property tax due date (Placebo), and this reminder plus information about the tax holiday (Treatment), relative to each other and relative to a group that received no flyers (Pure Control). Outcomes measure whether taxpayers paid punctually (Paid On Time) or accessed their online accounts (Web Access). We use inverse probability weights (IPW) for the differences in means to account for unequal probabilities of treatment assignment across eligible and ineligible taxpayers; effects are similar for each group considered separately (Online Appendix Figure A21). The “pooled” columns present estimated average effects over the seven post-treatment periods.

Evidence from the field experiment thus does not support a claim that information about benefits is an alternative explanation for the negative effects of tax holidays on compliance. A positive incentive effect could in principle induce a negative impact of holidays, if winners falsely believe that their probability of winning the lottery (l in the model) a second time is lowered and are induced to comply at lower rates than non-winners by this false presumption. Our household survey data do support the

³⁴Online Appendix Figure A19.

existence of such misperception.³⁵ However, to explain the negative effect of winning a tax holidays on compliance, the lottery itself would have to exert a powerful positive incentivizing effect on behavior; and our field experimental results are not consistent with this.

Perhaps information about the holiday shapes attitudes towards the payment of taxes, however. Our survey experiment lends insights into this question. As shown in Figure 8, we find no effect of information about the lottery on perceptions that (1) the municipal government does a good job, (2) municipal taxes are just, or (3) rewards for good taxpayers are a waste of money. However, being informed about the lottery (4) boosted perceptions of transparency and equity—in particular, decreased agreement with the statement that rewards “go to the same people as always”—and also (5) boosted agreement that it is “worth it to be up to date” on ones taxes.³⁶ Put in terms of our behavioral model, it is possible that the tax holiday policy increases parameters such as b , the expressive benefit of paying taxes; at the least, information about the lottery does not seem to act as a negative signal about municipal capacity. In our study, a positive effect of the lottery on perceptions of the utility of being up to date on payments does not readily serve as an alternative explanation for the *negative* effect of the tax holiday on tax compliance. However, as we discuss in the Conclusion, this is important for generalizability: positive effects of prize programs on expressive benefits could trump negative effects in other contexts.


In sum, informational mechanisms cannot readily explain the negative impact of tax holidays on compliance. Nor can several other alternative explanations that we considered in our pre-analysis plan and discuss in the Online Appendix (subsection B.1). We posited that by buttressing winners against the costs of punishment in case of non-payment (lowering c in our model), “income effects” might explain a negative effect of winning the lottery (Mechanism 1B.1 in the original PAP, p. 37). However, we find little heterogeneity in effects by property value (Online Appendix Figure A10). Nor are other

³⁵Indeed, 42% of respondents thought that the chances of winning the lottery again would be lower for someone who had already one it once—even though the probabilities are in fact independent.

³⁶For the “worth it to be up to date” measure, the estimated effect is a difference of 0.48 points on an 11-point ascending agreement scale, or 0.13 of the standard deviation of 3.699 in the control group, while for the “rewards go to the same people as always” outcome, the estimated effect of 1.05 is .33 of the control-group standard deviation of 3.235. Online Appendix Table A7 shows full results for the survey experiment, including a column showing that these two effects survive pre-specified adjustments for multiple comparisons.

alternative explanations consistent with the negative but decaying pattern of effects we find. The only explanation consistent with the overall pattern of evidence we uncover is that the tax holiday lowered subsequent tax payments by disrupting the habit of compliance.

Figure 8: Survey Experiment: Effects of Information about the Lottery on Attitudes Towards Taxation



The figure shows estimated average causal effects on five measures of attitudes towards taxation in our survey experiment. Outcomes are measured on an eleven-point (0-10) agreement scale, with the exception of “municipal taxes are just,” which is measured on a 4-point scale. Estimated effects are the difference between average responses in those exposed to the “lottery” treatment—where subjects were informed about the tax holiday using language similar to our mailed flyers in the field experiment—and “discretionary” treatments, where subjects were told that the municipality “from time to time” rewards good taxpayers with holidays.

5 Conclusion

Montevideo’s tax holiday program made those who won the exoneration worse taxpayers than they would counterfactually have been, on average. This effect was quite persistent, lasting for five payment

periods after the holiday or nearly two years in the case of property taxes. We find substantial evidence that this effect due to the disruption of the habit of paying taxes, and little or no evidence that alternative explanations can account for the negative effect on compliance.

Our theory and findings thus provide novel insights into the role of habit in political behavior. While scholars have studied habit's influence on voting, political scientists have not systematically explored its impact on many other important modes of citizen-state interaction, such as tax compliance. We also shed light on the consequences of habit disruption. A current shock to a good taxpayer's compliance today can have knock-on effects on future compliance. Yet, our findings also underscore that habits can also be difficult to alter permanently: they tend to be self-reinforcing. The stock of habit appears critical for explaining why some taxpayers persistently comply while others do not.

Our study also sheds light on the role of positive rather than negative incentives in shaping tax compliance. The impact of prize programs elsewhere in Latin America may differ from Montevideo's, since many raffle physical prizes (iPads, cars, and the like) that do not interrupt the habit of tax compliance (Online Appendix Section A). The evidence from our survey experiment suggests that positive incentives could boost some expressive benefits of paying taxes. In other contexts, even with habit interruptions, this effect could dominate negative effects on compliance. However, our field experiment suggests that information about the existence of the prize in Montevideo did not shape compliance behavior, which gives some reason for caution about the broader effects of positive incentives. Scholars could use designs similar as ours to test for the impact of prize lotteries elsewhere.

Finally, habit may influence many other kinds of interactions between citizens and their states. Our research may thus provide insights into sources of weak state capacity. In many Latin American countries, the extension of citizenship preceded the formation of strong states that could enforce compliance. Ensuing habits of non-compliance could have extended across years or generations through a combination of parental transmission and the knock-on effects of practice (Simpser 2016). The good news may be that interventions can build habits as well as destroy them, and the effects of disruptions are persistent but also decay. For policy makers, a lesson may be that larger investments in enforcement today could lessen the need for enforcement tomorrow by forming habits. Yet, lack of attention to habit can have perverse consequences. The failure to construct compliance habits—or their active destruction—may thus have lasting consequences for states' infrastructural power.

Acknowledgements: We are grateful to Lihuen Nocetto for research assistance, to the J-PAL Governance Initiative for generous funding, and to the Departamento de Recursos Financieros in the municipal government of Montevideo, Uruguay for collaboration. For helpful comments, we thank the editor Cameron Thies and the anonymous reviewers, as well as seminar participants at Berkeley, Emory, Harvard, Indiana University-Bloomington, Madison-Wisconsin, McGill, Stanford, Michigan, Princeton, UCLA, the University of Pennsylvania, Yale, and EGAP meetings at the Universidad Diego Portales in Santiago and Columbia University.

References

- Aldrich, John H., Jacob M. Montgomery and Wendy Wood. 2011. "Turnout as a Habit." *Political Behavior* 13:535–563.
- Allingham, M.G. and A. Sandmo. 1972. "Income Tax Evasion: A Theoretical Analysis." *Journal of Public Economics* 1:323–338.
- Andreoni, James, William Harbaugh and Lise Vesterlund. 2003. "The Carrot or the Stick: Rewards, Punishments, and Cooperation." *American Economic Review* 93(3):893–902.
- Ansola-behere, Stephen, Eitan Hersh and Kenneth Shepsle. 2012. "Movers, Stayers, and Registration: Why Age is Correlated with Registration in the U.S." *Quarterly Journal of Political Science* 7:333–363.
- Balliet, D. and P. A. M. Van Lange. 2014. How (and when) reward and punishment promote cooperation: An interdependence theoretical perspective. In *Reward and punishment in social dilemmas*, ed. P. A. M. Van Lange, B. Rockenbach, and T. Yamagishi. Oxford University Press p. 34–51.
- Bates, Robert H. and Da-Hsiang Donald Lien. 1985. "A Note on Taxation, Development, and Representative Government." *Politics & Society* 14:53–70.
- Benjamini, Yoav and Yosef Hochberg. 1995. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society. Series B (Methodological)* 57(1):289–300.
- Besley, Timothy and Torsten Persson. 2009. "The Origins of State Capacity: Property Rights, Taxation and Politics." *American Economic Review* 99(4):1218–44.
- Bonferroni, Carlo E. 1936. "Teoria statistica delle classi e calcolo delle probabilità." *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze* .
- Bourdieu, Pierre. 1990. *In Other Words: Essays Towards a Reflexive Sociology*. Stanford: Stanford University Press.
- Carrillo, Paul E., Edgar Castro and Carlos Scartascini. 2017. "Do rewards work? Evidence from the randomization of public works." Working Paper 794, Inter-American Development Bank, Washington, DC.
- Coppock, Alexander and Donald P. Green. 2015. "Is Voting Habit Forming? New Evidence from Experiments and Regression Discontinuities." *American Journal of Political Science* 60(4):1044–62.

- Cox, D. R. 1958. *Planning of Experiments*. Wiley.
- Dunning, Thad. 2012. *Natural Experiments in the Social Sciences: A Design-Based Approach*. Cambridge: Cambridge University Press.
- Eggers, Andrew, Guadalupe Tuñón and Allan Dafoe. Forthcoming. “Placebo Tests for Causal Inference.” *American Journal of Political Science* .
- Gerber, Alan S. and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation*. W.W. Norton & Co.
- Gerber, Alan S., Donald P. Green and Ron Shachar. 2003. “Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment.” *American Journal of Political Science* 47(3):540–550.
- Holland, Alisha C. 2016. “Forbearance.” *American Political Science Review* 110(2):232–46.
- Humphreys, Macartan, Raul Sanchez de la Sierra and Peter van der Windt. 2013. “Fishing, Commitment, and Communication: A Proposal for Comprehensive Nonbinding Research Registration.” *Political Analysis* 21(1):1–20.
- Levi, Margaret. 1988. *Of Rule and Revenue*. Berkeley: University of California Press.
- Lieberman, Evan. 2003. *Race and Regionalism in the Politics of Taxation in Brazil and South Africa*. New York: Cambridge University Press.
- Mann, Michael. 1984. “The Autonomous Power of the State: Its Origins, Mechanisms, and Results.” *Archives Européennes de Sociologie* 25:185–213.
- Meredith, Marc. 2009. “Persistence in Political Participation.” *Quarterly Journal of Political Science* 4(3):186–208.
- Paler, Laura. 2013. “Keeping the Public Purse: An Experiment in Windfalls, Taxes, and the Incentives to Restrain Government.” *American Political Science Review* 107(4):706–25.
- Pierson, Paul. 2000. “Increasing Returns, Path Dependence, and the Study of Politics.” *The American Political Science Review* 94(2):251–267.
- Pomeranz, Dina and José Vila-Belda. 2019. “Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities.” *Annual Review of Economics* 11:755–81.
- Ross, Michael. 2004. “Does Taxation Lead to Representation?” *British Journal of Political Science* 34(2):229–249.
- Scheve, Kenneth and David Stasavage. 2016. *Taxing the Rich: A History of Fiscal Fairness in the United States and Europe*. Princeton and New York: Princeton University Press and Russell Sage Foundation.
- Schmitt-Grohe, Stephanie and Martin Uribe. 2008. “Habit Persistence.” *New Palgrave Dictionary of Economics*, 2nd ed. . 2nd edition.
- Simpser, Alberto. 2016. “The Culture of Corruption across Generations: An Empirical Study of Bribery Attitudes and Behavior.” Working Paper.

- Soifer, Hillel. 2008. "State Infrastructural Power: Approaches to Conceptualization and Measurement." *Studies in Comparative International Development* 43(3-4):231–251.
- Stokes, Susan, Thad Dunning, Marcelo Nazareno and Valeria Brusco. 2013. *Brokers, Voters, and Clients: The Puzzle of Distributive Politics*. Cambridge University Press.
- Tilly, Charles. 1990. *Coercion, capital, and European states, AD 990*. Cambridge: Basil Blackwell.