

Is Paying Taxes Habit Forming?

Evidence from Uruguay

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

This draft: October 16, 2016

Prepared for presentation at the Comparative Politics Colloquium, University of California, Berkeley, October 20, 2016. We are grateful to Lihuen Nocetto for research assistance and to the J-PAL Governance Initiative for generous funding. We received helpful comments on this project from participants at EGAP meetings at Columbia University and the Universidad Diego Portales (Chile), the Political Geography Conference at Harvard, and seminars at Emory, Indiana University Bloomington, McGill, Princeton, Stanford, UCLA, University of Pennsylvania, University of Wisconsin-Madison, and Yale. The pre-analysis plan for this project was registered with EGAP (www.e-gap.org/design-registrations) as well as the American Economic Association (www.socialscienceregistry.org) in July 2014. The study was approved by the Office for the Protection of Human Subjects at the University of California, Berkeley (Protocol ID 2014-04-6286).

^{*}University of California, Berkeley.

[†]Universidad de la República, Montevideo, Uruguay

[‡]Universidad Católica, Montevideo, Uruguay

[§]Universidad Diego Portales, Santiago, Chile

[¶]University of California, Berkeley and University of Pennsylvania.

Abstract

Interactions between citizens and their states often involve routinized, repeated behaviors—and repetition may itself breed habits of citizenship that exert a causal influence on behavior. This implies the possibility of virtuous or vicious cycles in civic participation, as well as broader development outcomes. Yet, it is usually difficult to separate habit from confounding explanations for repeated behaviors. We study a policy in Montevideo, Uruguay that randomly assigns tax holidays, or year-long interruptions of payments, to punctual taxpayers; the program is designed both to reward and induce tax compliance, a critical aspect of citizen-state interaction and a key facet of state capacity. We find that far from fostering compliance, the tax holiday inhibits it: winning the lottery results in a substantial and persistent reduction in compliance after the end of the holiday. We use field and survey experiments to disentangle informational mechanisms from the effects of habit disruption. Our findings have both social-scientific and policy implications.

1 Introduction

Interactions between citizens and the state often involve routinized, repeated behaviors—and this repetition may itself breed habits of citizenship, with independent consequences for civic participation. Habits can also impact development outcomes. Social scientists have long recognized that countries may be locked into development traps, which may emerge from strategic behavior or from increasing returns and path dependency.¹ However, habit may be an important contributor. By the same token, interventions that disrupt habits can be critical forces for change. The power of habit therefore implies the possibility of virtuous or vicious cycles in civic participation—and raises the question of how habits are formed and how they may be disrupted.

Consider the case of tax compliance, a critical aspect of citizen-state interaction and a key facet of state capacity as well as political development. Scholars often attribute low rates of compliance, especially in developing countries, to weak monitoring capacity and the inability of states to penetrate society and compel tax payment.² Yet habit could also sustain high or low tax compliance, in developed or developing countries alike. In data we gathered for this project, we found persistent non-compliance with municipal tax bills in Montevideo, Uruguay, a Latin American country thought to possess a relatively capable and efficacious state. The average tax account in Montevideo owes around six past due tax payments; only around 65 percent of tax bills are paid on time; and the municipality classifies just 40 percent of property taxpayers as “good taxpayers,” based on having paid taxes punctually over the previous year.³ Weak monitoring capacity alone cannot explain these outcomes. For example, with municipal property taxes, the state knows with certainty the amount of tax owed on the basis of appraised values, yet still fails to induce prompt compliance among substantial numbers of taxpayers.

Habit, by contrast, may play a role in shaping such outcomes. While some taxpayers are intermittent compliers, there is also substantial structural stability, with many “good taxpayers” paying taxes punctually year after year while similarly situated people are repeatedly non-compliant. The difficult question is whether this reflects any causal effect of habit, or whether other attributes account for all of the difference between compliant and non-compliant taxpayers. Empirical research on the habit of

¹On path dependency, see Pierson (2000), Collier and Collier (1991), or Acemoglu and Robinson (2012).

²See e.g. Besley and Persson (2009) or Mann (1984).

³The data are from Figure 1, discussed later.

paying taxes is not well developed, perhaps because it is exceedingly difficult to generate as-good-as-random assignment of past tax compliance. Yet, understanding habit's impact on tax payment is important for both social-scientific and policy purposes.

We show in this article that states—and social scientists—ignore the importance of habit at their peril. Faced with the reality of low compliance, and the apparent difficulty of using negative incentives such as threats of punishment to elicit payments, many municipalities in developing countries have experimented with rewards programs that offer positive incentives for prompt tax payment. For example, upwards of 25 percent of municipalities in Brazil, Colombia, Ecuador, and Uruguay, and many localities in Argentina, Peru, and Mexico, now offer rewards programs such as prize lotteries to good taxpayers.⁴ Montevideo was one of the earliest innovators of such policy in Latin America: since 2004, and across four types of municipal taxes, the municipality has raffled a year free of tax payments to eligible good taxpayers who have been punctual compliers over the previous year. These tax holidays both reward and may incentivize future compliance by good and bad taxpayers alike; they may also boost perceptions that the tax system is equitable and transparent.

However, such holidays also interrupt the habit of tax payment. Indeed, the lotteries provide a rare form of randomization that allows study of the impact of disrupting tax compliance. Using a detailed panel of administrative records on individual tax payments as well as household survey data, we compare lottery winners to a randomly selected control group of eligible non-winners. This natural experiment allows us to assess how winning a tax holiday influences subsequent tax compliance, using unobtrusive outcomes that are not subject to social desirability bias from self-reports and that are measured over more than a decade. Placebo outcome tests with data on ineligible “winners” as well as balance tests on pre-treatment covariates validate key assumptions of our design. We also draw from extensive qualitative interviews, both with municipal officials throughout Latin America and with voters who won the tax holiday lottery.

We find that far from promoting tax compliance, Montevideo's tax holiday inhibits it: among eligible taxpayers who would claim the tax holiday, winning the lottery results in up to a substantial reduction in the proportion who pay on time, an effect that lasts for several years. This consequently also reduces the proportion of eligible good taxpayers and increases accumulated tax debt. Our in-

⁴See Section 2.

terpretation that this is due to interruption of habit is bolstered by the fact that we do not find such negative effects for the one municipal tax for which winners continue to pay some small amount of tax during the holiday—and for which the payment habit is thus not in fact interrupted. Moreover, we find no effect on taxpayers signed up for automatic debits, whose payments mechanically restart after the holidays: only for manual taxpayers does the interruption have a negative impact. Effects are also less pronounced for taxpayers with a greater “stock” of prior compliance habit.

The main alternative explanations for our findings are informational. Our survey data suggest that the municipal government has poorly advertised the tax holiday: although the program has existed for a decade, only about 8% of Montevidean taxpayers could identify the tax lottery as a municipal policy that rewards good taxpayers. Along with disrupting the payment habit, winning the lottery thus also informs many winners about the existence of the rewards program. To explain the negative impact of winning on subsequent compliance, taxpayers would have to interpret that information as a negative signal about municipal capacity or about the payment behavior of other taxpayers. Alternately, winners might falsely believe that their probability of winning the lottery a second time is lowered, which might lower future compliance if the lottery indeed has a positive incentive effect. Winning the lottery might also lead to income or substitution effects, in which taxpayers spend the additional revenue gained from the tax holiday on consumption goods in a way that makes it more difficult to resume good taxpaying once the holiday expires. Finally, the tax holiday could have a behavioral effect by breaking a taboo against non-compliance among good taxpayers or by crowding out intrinsic incentives for compliance.

To assess these alternative explanations, we also conducted a field experiment with a series of informational interventions. In collaboration with the municipal government, we mailed flyers stamped with the municipal logo—which appear very much like tax bills themselves—reminding taxpayers of an impending due date for taxes and providing additional experimentally varied messages. Along with a pure control group that received no flyers, we assigned households to receive a mere reminder (the placebo control group); a reminder plus information about the existence of the lottery; and a reminder plus information about fines and sanctions for non-payment. We then track the impact of our treatments using administrative data on subsequent tax payment and on whether taxpayers logged in to their online accounts, which we take as a measure of intended compliance. By providing information about the lottery to both eligible and ineligible taxpayers without actually awarding the tax holiday, this large

field experiment allows us to separate informational effects from the impact of habit disruption.

We find no evidence that informational explanations can account for our results. Information about the lottery has no greater impact on intended or actual compliance than a simple reminder to a placebo control group that taxes are due. While many taxpayers do incorrectly perceive that winning the lottery produces a lower probability of winning again, our evidence is not consistent with a powerful positive incentive effect of the lottery that would explain the negative effect of winning on compliance. We also find some evidence from a supplementary survey experiment that information about the lottery positively impacts attitudes towards the fairness of taxes. It is therefore unlikely that information about the lottery acts as a negative signal. We also emphasize that while alternative explanations such as the breaking of a taboo against non-compliance might suggest more permanent effects, the effects of habit interruptions may eventually decay. Indeed, this is what we find, suggesting that the force of habit is the most plausible explanation for our findings.

Our study makes several contributions to the study of the habitual bases of interactions between citizens and states, as well as the causes of tax compliance. The influence of habit has been studied most extensively in the case of political participation. For example, 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 vote in the next election as well. 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. If voting is partly habitual, interruptions of habit can also have substantial consequences. For example, disruptions caused by residential mobility and re-registration requirements may account for turnout disparities between young and old voters in the United States (Aldrich, Montgomery and Wood, 2011; Ansolabehere, Hersh and Shepsle, 2012). In domains other than voting, the idea simply that habit leads to recurrence is less well explored both theoretically and empirically, in part because habit is usually so difficult to separate from confounding factors that may result in recurrence. Nor have scholars, including those studying voting, convincingly isolated the impact of habit disruption, rather than habit formation.

In settings where taxpayers face a low probability of punishment for non-compliance, the key question may not be why people do not pay their taxes—but why any taxpayers do comply. Our research

provides insight into this important puzzle, similar in its essence to the paradox of voting.⁵ While expressive or material benefits of paying taxes could play some role in sustaining compliance, the simple repetition of compliant behavior can itself breed future payment. This also suggests that public policies that are inattentive to habit—such as Montevideo’s temporary exoneration of tax obligations as a reward for good compliance behavior—can have perverse consequences. We return to these interpretations in the conclusion, after discussing the theory that elucidates our results (Section 2); describing our empirical strategy, including natural and field experimental designs, and presenting our main findings (Section 3); and discussing evidence for the habit mechanism as well as considering alternative explanations (Section 4).

2 Why do people comply? Taxes, incentives, and habit

By “habit” we do not mean a behavior that is necessarily automatic or unthinking. Social psychologists have developed a specific 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. In Uruguay, the majority of taxpayers pay their taxes in person at local kiosks.⁶ They also receive tax bills in advance of each payment period (which continue to arrive even during tax holidays, though showing zero balance). Thus, a tax 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 or to City Hall. We argue that habit may elucidate several features of tax compliance, in a way that decision-theoretic models emphasizing the costs of evasion or the benefits of tax payment cannot.

Why, then, do people comply with taxes? In a standard formalization of the compliance problem, taxpayers compare the utility of evasion to the cost of punishment discounted by its probability.⁷ For example, let y be an asset value, t be the annual tax rate, and z be the unpaid annual amount of taxes due; with full nonpayment, $z = ty$. The expected utility of full nonpayment in any year is thus $z - pc$, where c is the penalty for nonpayment and p the probability of punishment. In the context of the municipal

⁵E.g., Downs (1957).

⁶These are called “Red de Cobranzas.” A smaller proportion pay by automatic bank debit.

⁷E.g., Allingham and Sandmo (1972).

taxes we study in this paper, 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.⁸ In developing countries in particular, states may have limited capacity—what Mann calls “infrastructural power”—to penetrate society and monitor and discipline non-compliers.⁹

Yet, even where administrative capacity is strong and monitoring problems are negligible, enforcing punishments for non-payment can pose challenges. Our interviews with officials in several Latin American cities suggest that governments 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. The most typical outcome of a legal process against a debtor is renegotiation, in which delinquent taxpayers agree to a discounted payment plan. A tax official in Tigre, Argentina notes “In the last 20 years, we have never auctioned either a commercial or residential property. In general, we end up with an agreement.”¹⁰ General amnesties for delinquent taxpayers, particularly in times of economic crisis, are also highly prevalent. For example, Montevideo's city government issued 11 amnesties at different points between 1997 and 2013.¹¹ The political unpopularity of taxes can also be especially salient in municipalities, where face-to-face negotiations with delinquent taxpayers is common and selective enforcement (or “forbearance,” see Holland 2016) may prevail. As one municipal official put it, “proximity means that a neighbor can approach the administration to justify why he doesn't pay. By contrast, at the national level nobody will pay him any mind.”¹²

Thus, even in what might seem like a best-case situation for inducing compliance—strong administrative capacity and limited monitoring problems—the probability of punishment for non-compliance can be very low. In the case of municipal property taxes, bureaucrats know appraised values and

⁸E.g. Bates and Lien (1985).

⁹Mann (1984), also Soifer (2008).


¹⁰Daniel Chillo, Secretario de Ingresos Públicos; all translations ours. Of course, negotiations happen in the shadow of legal proceedings: “the majority of the legal actions of City Hall fall by the wayside because [negotiations] were successful” (Geraldo Cruz, federal deputy and former mayor of Embu das Artes, state of São Paulo, Brazil says). Renegotiations also do not result in full compliance: instead, they “generate a new problem for you. Now the person doesn't only have pay the monthly rate but also the quota from the agreement” (Carlos Maisterrena, Sub-Director de la Administración Fiscal Municipal, Ciudad de Paraná, Entre Ríos, Argentina).

¹¹That includes various amnesties in 1997, 2000, 2001, 2002, 2003, 2004, 2005, 2008, 2009, 2010, and 2013.

¹²Chillo, note 9.

can identify taxpayers’ obligations with certainty.¹³ Enforcement becomes an apparently easier problem, that of cajoling taxpayers to fulfill their known tax obligations promptly.¹⁴ Non-compliance is nonetheless endemic. Our simple random sample of property, vehicle, head and sewage tax accounts in Montevideo shows that between 2000 and 2014, the municipal government has classified under 40% taxpayers as good taxpayers at any moment in time, based on being current on payments over the previous year (right panel of Figure 1). The compliance rate per tax bill due is around 70% (middle panel), which results in growing average delinquency over time (left panel).

Figure 1: Non-Compliance Over Time: Taxes In Montevideo, Uruguay



The figure is based on a simple random sample of 9,297 tax accounts in Montevideo, Uruguay. The left panel plots the proportion of tax bills paid on time in each year; the middle panel depicts average accumulated debt (past tax bills due); and the right panel shows the proportion of “good taxpayers,” a municipal measure that establishes eligibility for the tax holiday lottery we study in this paper.

Given the low probability of serious punishment for evasion, the provocative question may not be why some people don’t pay taxes—but why anyone does. One natural possibility is that the material or expressive benefits of paying taxes induce compliance. Latin American municipal governments from Salta, Argentina to Perúbe, Brazil to Miraflores, Peru now raffle prizes—from televisions to new

¹³With the exception of Castro and Scartascini (2015) most existing experimental literature on tax compliance focuses on income taxes (e.g. Blumenthal et al. (2001); Slemrod, Blumenthal and Christian (2001); Kleven et al. (2011)), VAT (Pomeranz (2013)), or very specific fees (Fellner, Sausgruber and Traxler, 2013).

¹⁴In this sense, the problem facing the state is akin to that of a credit-card company faced with non-payment by consumers.

cars and houses to discounted payments—to taxpayers who pay their taxes on time. Our 10% random sample of municipalities in Latin American countries found that 79% of Uruguayan municipalities offer such a reward for good taxpayers, along with an estimated 24% of Brazilian municipalities (134/558 in our sample); incentive programs are also prevalent in Argentina, Bolivia, Colombia, Ecuador, Mexico, and Peru.¹⁵ These programs constitute an important recent policy innovation in developing countries.

Many of our interviewees believe that such incentive programs induce greater compliance. As an official in Tigre, Argentina noted, “We have a compliance rate of 85% with the Municipal Service Tax today, whereas when we started [prize lotteries] in 2009, it was at 68%.”¹⁶ While perhaps an example of “casual” inference, the theory of change appears plausible: prizes may not only sustain compliance among good taxpayers but also induce bad taxpayers to bring their accounts up to date. An Argentine official noted that “one of the conditions [to participate] was not to be in a payment plan. What did people say? ‘Make me a payment plan with six quotas, I’ll finish the sixth and enter [the lottery].’”¹⁷ A former Brazilian mayor noted, “The effect of [starting the lottery] meant that by the second year, many indebted people went to look for payment plans. After the results . . . were shown, a lot of municipalities adopted these policies.”¹⁸ In Montevideo, the center-left government of the Frente Amplio initiated its lottery in the context of an amnesty for many delinquent taxpayers, following the economic crisis of 2002.¹⁹ As municipal officials told us, the economic crisis generated a dilemma: how to lower the burden for those under dire circumstances while at the same time counteracting negative perceptions of forgiveness for non-compliance through amnesties. The lotteries were their answer.²⁰

From the perspective of standard frameworks, however, the material benefits of rewards programs alone appear unlikely to boost tax payment. Under Montevideo’s lottery policy, good taxpayers win a year free of tax payments with probability 1/5,000 in any tax payment period. The expected utility

¹⁵We used Web searches, interviews with municipal bureaucrats and politicians, and ancillary sources. See the online Appendix.

¹⁶Chillo, note 9.

¹⁷Carlos Maisterrena, Sub-Director de la Administración de Fiscal Municipal, Ciudad de Paraná, Entre Ríos, Argentina; translation ours.

¹⁸Beto Tricoli, Prefeito de Atibaia (2001-2008), Brazil.

¹⁹For discussion, see http://historico.elpais.com.uy/Suple/LaSemanaEnElPais/04/10/29/lasem_ciud_118264.asp; http://www.montevideo.com.uy/notnoticias_66228_1.html; and <http://www.180.com.uy/articulo/14284>. Accessed May 2, 2015.

²⁰In October 2013, the municipality announced a renewed amnesty for certain bad taxpayers, underscoring the difficulties of cracking down on non-compliance. There have been amnesties in 2004, 2008, and 2010, among other years. See also URLs in previous note.

of paying the full year's taxes in one year is then (without discounting) $(1/5,000)z - z$.²¹ In other words, a taxpayer has to pay a year's worth of taxes z in order to gain, with probability $1/5,000$, a year free of tax payments z in the following year. Casinos would make no money if gamblers did not take bets with negative expected values; yet this does not look like a promising deal for the taxpayer.²² Of course, prize lotteries such as Montevideo's may possibly affect non-material or expressive benefits of paying taxes, e.g. by influencing perceptions of the fairness, equity, or transparency of the tax system. As one Brazilian mayor put it, "the best weapon [against non-compliance] is transparency."²³ For example, in one Argentine municipality raffle, "there are two city councilors and two functionaries from other departments of the municipality present. And in the same lottery appears the winner from the monthly lottery immediately prior—so we do the drawing of lots and then give a TV to the person who won the preceding lottery."²⁴ The public nature of this lottery also underscores an additional theme, recognition.²⁵ One official emphasized "we take a photo and put the program on the webpage of the municipality, we publish a list of the taxpayers among whom we do the lottery . . . we do not seek to give [the winning taxpayer] value but rather recognition."²⁶ Another interviewee noted, "every time we award prizes we make a kind of show. In some way this helps us to create a tax-paying culture... It is not worth it to hit unless you also caress those you need to caress."²⁷

To extend the standard framework, let b therefore be an expressive benefit of paying taxes. Under Montevideo's policy, the expected net utility of non-compliance is then $z - (1/5,000)z - b - pc$, and tax payment occurs whenever $pc > (\frac{4,999}{5,000})z - b$ —which is satisfied more easily the larger is b . By boosting both monetary and expressive benefits of paying taxes, the incentive programs cropping up throughout Latin American might thus induce greater compliance.

²¹The average appraised property value is US\$36,035 (956,000 Uruguayan pesos) and the annual value of property taxes is over US\$265 (7,044 Uruguayan pesos), which is non-trivial. However, the *expected* value of winning any lottery is then US\$265/5,000, or about five US cents. We estimate these values from the control group in our natural experiment.

²²For discussion of positive and negative incentives such as lotteries and state amnesties, see (Parle and Hirlinger, 1986; Loftus, 1985; Alm, McKee and Beck, 1990; Falkinger and Walther, 1991; Alm et al., 1990). According to Smith and Stalans (1991), the costs of implementing governing programs that give taxpayers positive material incentives to comply with tax laws may outweigh the benefits.

²³Cruz, footnote 11.

²⁴Chillo, note 9.

²⁵Previous studies suggest that social incentives and peer effects may be an effective way to improve compliance; see e.g. Chetty, Mobarak and Singhal (2014).

²⁶Chillo, note 9.

²⁷Carlos Maisterrena, Sub-Director of Fiscal Administration, Ciudad de Paraná, Entre Ríos, Argentina.

2.1 Habit as a cause of compliance

However, this extended framework is also incomplete—and leaves unexplained several key features of tax payment data that we explore empirically.

Suppose instead that the tendency to comply with taxes is partly habitual.²⁸ For example, let $\gamma_t = 1$ if a taxpayer fails to comply with taxes at time t and otherwise equals zero. A framework that accommodates habit formation gives the expected utility of choosing $\gamma_t = 1$ as $z - (1/5,000)z - b - pc + \theta\gamma_{t-1}$. Here, $\theta \in (0, 1)$ captures the sensitivity of current tax payment to past compliance behavior. Note that by a recursive argument, the “stock” of habit at time t is given by $S_t = \sum_{i=1}^{\infty} \theta^i \gamma_{t-i}$.²⁹ Thus, tax payment could therefore become nearly automated through behavioral repetition. Completing the argument, one could suppose that compliance is not a deterministic function of these parameters but is also affected by some mean-zero random noise v_t : thus, $\gamma_t = 1$ if $z - (1/5,000)z - b - pc + S_t + v_t > 0$. Note that in this formalization, the benefits of tax payment today are positively related to past compliance.

This framework generates several implications consistent with the data on tax compliance we present below. First, behavioral repetition readily generates types of taxpayers who typically comply and those who do not. For taxpayers with large “stocks” of past non-compliance (high S_t), non-compliance may usually be the preferred option; while those for whom S_t approaches zero are more likely to comply. To be sure, if compliance is a random variable, some types with high S_t might nonetheless comply on some occasions, depending on the realization of v_t . Yet habit can contribute to the emergence and persistence of “good” and “bad” taxpayers, in the language of the municipal policy we study empirically.

Second, exogenous changes to tax compliance behavior—such as interventions that switch past compliance to non-compliance—can have lasting effects. After all, compliance at time $t + 1$ is a function of compliance at t ; but compliance at t is a function of compliance at $t - 1$. Disrupting the habit of tax compliance may therefore imply less tax compliance not only in the immediately following period but beyond it. Thus, there can be “knock-on” effects of past shocks to tax compliance.

Finally, however, there is naturally a “decay” in the effect of such habit disruptions. Consider a

²⁸Our pre-analysis plan considered this possibility; see Dunning et al., July 23, 2014. “Pre-Analysis Plan—Positive vs. Negative Incentives for Compliance: Evaluating a Randomized Tax Holiday,” p. 7 and 37 (Mechanism 1B.2).

²⁹Similar approaches to modeling the stock of habit is used in economics to study the equity premium puzzle.

taxpayer who has always complied until time $t - 1$ and therefore has stock $S_t = 0$. Since she has always complied, such a taxpayer is also likely to have parameter values (p , c , or b) that favor compliance.³⁰ Switching past compliance $\gamma_{t-1} = 0$ to non-compliance, $\gamma_{t-1} = 1$, makes compliance less attractive at time t and can have knock-on effects beyond that, as described in the previous paragraph. However, for many realizations of v_t , such a taxpayer will again comply—which, given parameter values that favor compliance, will tend to foster continued compliance at $t + 1$ and subsequent periods. These implications are important because they contrast with alternative explanations for our findings.

Thus, winning the lottery could break the habit of payment and lead winners to pay less reliably than winners for some period of time after winning. In terms of our formalization, by rewarding good taxpayers with a year's interruption in their tax payment obligations, the tax holiday lottery that we study empirically switches past compliance, $\gamma_{t-1} = 0$, to non-compliance, $\gamma_{t-1} = 1$, and therefore also affects the stock of habit. Our framework thus suggests that such an incentive program could have quite perverse impacts on future compliance. We turn next to the empirical evidence for this claim.

3 The negative effect of tax holidays: a natural experiment

To study the effects of the tax holiday, we use the design of the lottery as a natural experiment.³¹ The municipal government uses the results of Uruguay's National Lottery to select taxpayers for holidays. Thus, it selects as provisional winners of tax holidays those account numbers, the final four digits of which correspond to the winning number of the relevant National Lottery.³² 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, across four types of taxes for which one taxpayer would have different account numbers—property, vehicle, head, and sewage³³—then screened in those

³⁰For ease of presentation, we have not subscripted these parameters; yet it is natural that e.g. the expressive benefit of paying taxes b varies across taxpayers.

³¹Snow (1855), Freedman (1991), Angrist and Krueger (2001), Dunning (2012).

³²The randomization occurs through the selection from balls from an urn, as described in Spanish at http://www.loteria.gub.uy/Juego_Loteria.php. For an example of posted lottery results, see http://www.loteria.gub.uy/ver_resultados.php?vdia=21&vmes=6&vano=2013. Winning taxpayer numbers are posted at <http://www.montevideo.gub.uy/sorteosBP/pages/sorteosBuenosPagadores.xhtml>. Accessed May 2, 2015.

³³In Spanish, these are, respectively, “Contribución Inmobiliaria,” “Patente de Rodados,” “Tasa General Municipal,” and “Tarifa de Sanamiento.”

taxpayers who owed no past taxes and had paid on time over the previous year.³⁴ To contact winners, the municipality sends a letter to the addresses associated with each of these accounts indicating that winners should come to City Hall to claim their year-long tax holiday.³⁵

To construct the study group for our natural experiment, we first identified all taxpayer accounts that were randomly selected in each lottery since 2004. Our treatment group consists of winning taxpayers who were eligible to claim the tax holiday: they were all “good taxpayers” in the year prior to the date of the lottery in which their account number was selected. (Our use of the term should not be understood in a normative sense. The municipality calls its policy “Lottery for Good Taxpayers”). Constructing the appropriate control group requires some care: the right counterfactual consists not of currently eligible non-winners but rather sets of taxpayers who were eligible to win as each past lottery, but whose account numbers were not randomly selected in the tax holiday lottery.³⁶ To create the control group, we 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 “good taxpayers” as of the date of the corresponding lottery. Our procedure therefore mimicked the random process that created the treatment group. We requested data from our municipal partners for ineligible as well as eligible taxpayers with the selected account numbers. Although we only use data for eligible taxpayers to estimate treatment effects (since only those taxpayers could potentially claim a tax holiday lottery), we exploit data on ineligible taxpayers for placebo tests. Our partners in Montevideo’s tax bureaucracy provided a time-series panel of payment data (2000-2014) for all randomly selected account numbers.

Table 3.1 depicts the size of treatment and control groups (in bold font), distinguishing between the different types of taxes.³⁷ In addition, we have payment data for ineligible taxpayers whose account

³⁴A screenshot off the municipality’s website showing the list of the winning “good taxpayers” as reported by the municipality can be found in the supplementary appendix.

³⁵For head and sewage taxes, the municipality grants such holidays on a bimonthly basis six times a year (in February, April, June, August, October, and December); for vehicle and property taxes, it issues holidays three times a year (in March, July, and November).

³⁶If we were instead to use currently eligible taxpayers as a control group, we would mix taxpayers who were eligible and ineligible as of the date of each past lottery. Since the treatment group only includes eligible winners as of the date of each lottery, this asymmetry would risk bias, if potential outcomes are related to eligibility status (which seems extremely plausible since the outcome is tax compliance).

³⁷Here we in fact have a series of mini-natural experiments, in which each lottery generates a treatment group of winners and a control group of non-winners. Thus, the random assignment is effectively blocked by individual lottery; however, the probability of winning is the same in every block.

numbers would have made them eligible for tax holidays, were they up-to-date on their payments; and ineligible taxpayers whose account numbers match those in our control group. Note that the cells of Table 3.1 are themselves random samples from the population of taxpayer accounts in Montevideo, as of the date of each lottery. We can thus use these random samples to characterize features of the taxpaying population. For example, we combine eligible (3,189) and ineligible (6,108) taxpayers in the non-winning group for the random sample of 9,297 accounts used in Figure 1.

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

Tax	Taxpayer Type	Winning Tax Account Number	Non-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 the sample size of lottery winners and non-winners in the natural experimental study group. Rows used to estimate treatment effects are in bold; non-bolded rows are used for placebo tests.

To estimate the effect of winning a tax holiday, we compare payment behavior of taxpayers in the treatment and the control groups after the holiday. Figure 2 shows pooled effects across all four types of taxes.³⁸ The horizontal axis shows the number of non-holiday payment periods elapsed before or after the particular lottery, in connection with which a taxpayer won a tax holiday (in the treatment group) or was sampled into the control group. For each taxpayer, we center the horizontal axis at $t = 0$, the tax payment period at which the taxpayer won or “lost” the lottery and after which the exoneration began for winners. (Note that non-winners in our control group did not “lose” the lottery in any meaningful competitive sense, since—as we show later—the vast majority are unaware of the existence of the lottery; their account number was simply not selected for the tax holiday). The dependent variable

³⁸This pooled analysis was our pre-registered protocol posted at <https://www.socialscienceregistry.org/trials/452> and <http://egap.org/design-registration/registered-designs/> (design 84: 20140723). Complete results produced using the code in our registered mock analysis will be available in an online Appendix.

in the analysis is the proportion of taxpayers who paid their bill on time at each tax payment period. The grey vertical strip indicates the period of exoneration; thus, the first estimate to the right of the strip is for the first payment due after the holiday, the second estimate is the second payment due, and so forth.³⁹ The top panel depicts effects for eligible taxpayers—those who had been prompt in tax payments for the previous year before the date of the lottery—while the bottom panel shows differences across winning and losing account numbers for ineligible taxpayers.

Before moving to treatment effects, note that Figure 2 permits graphical balance and placebo outcome tests that validate our design. For payment periods less than or equal to zero, the figure shows balance on a highly prognostic covariate: pre-treatment tax compliance. Thus, among both eligible and ineligible taxpayers, the compliance behavior of winning account numbers is statistically indistinguishable from losing account numbers—just as the lottery’s random assignment of winning numbers should imply. Note also that for eligible taxpayers, there is no variation in tax payment behavior (and thus the confidence intervals collapse) in the three payment periods prior to the relevant lottery: to be eligible at the date of a particular lottery, taxpayers must have paid on time over at least the three previous payment periods.⁴⁰ Confidence intervals in earlier periods and for ineligible taxpayers suggest that with winning and losing accounts are statistically balanced with respect to past compliance behavior. To test formally if the data are consistent with random assignment, Appendix Table 1.1 reports balance tests for pre-treatment tax compliance at different payment periods, as well as an indicator for being retired at the date of the lottery and the 2004 property value in pesos. To generate higher-powered tests, we use data on both eligible and ineligible taxpayers, since both types are randomized to the winning and non-winning lottery numbers. None of the p -values approach nominal significance levels.⁴¹

Figure 2 also permits a placebo outcome test: the “effect” of assignment to a winning number for ineligible taxpayers, who were not eligible for and thus did not receive any benefit from having a winning lottery number.⁴² As the bottom panel of the figure suggests, there is no statistically discernible

³⁹For each tax, we use the maximum period of the exoneration, which is sometimes delayed by failure to claim the exoneration immediately. The number of payment periods per year varies by tax, per note 36.

⁴⁰The requirement is compliance over the previous year; that makes three payments of property and vehicle taxes and six payments of sewage and head taxes, per note 36.

⁴¹Geographically, the treatment and control groups are also intermingled evenly throughout the municipality (see Appendix Figures 1 and 2).

⁴²A placebo outcome test validates the assumptions of a design by assessing whether there is an effect where one is “known” not to exist. On placebo outcomes, see Dafoe and Tuñón (2014).

difference in post-treatment compliance between ineligible taxpayers with winning and losing lottery numbers. This evidence also supports an important “exclusion restriction”: assignment to a winning lottery number (rather than to the treatment of a tax holiday) did not itself influence compliance behavior.⁴³ Appendix Table 1.2 reports formal statistical tests which reach the same conclusion. Thus, the data from the natural experiment are strongly consistent with random assignment to treatment conditions and with other identifying assumptions of our design.⁴⁴

Finally, Figure 2 shows estimates of treatment effects. Thus, the black circles to the right of the grey strip (top panel) compare post-holiday compliance behavior of winning and losing account numbers, among eligible taxpayers. Vertical lines show 95% confidence intervals around these point estimates.⁴⁵ The figure shows estimated complier average causal effects (CACEs) to account for the fact that not all taxpayers claimed exoneration, for example, because they were corporations not physical persons.⁴⁶ (This uses “complier” in the statistical sense of compliance with treatment assignment, rather than compliance with tax obligations). Here, we divide the estimated effect of treatment assignment by the proportion of taxpayers with winning numbers who claimed the exoneration. Table 3.2 reports the estimated average causal effect (ACE) of treatment assignment (i.e., intent-to-treat analysis), along with standard errors and estimated CACEs, at post-holiday payment periods 1, 5, and 10, as well as the average for all ten periods. It also breaks down the latter estimate by type of tax. (We return to cross-tax variation in the next section).

The conclusion from both the figure and table are clear: there is a large, *negative*, and persistent effect of tax holidays on subsequent tax compliance, which lasts at least four payment periods (or over a year). In the Appendix we present effects for two related measures: the proportion of “good taxpayers” at each payment period and the accumulated number of payments owed. There the effects are even larger and more persistent, in part because failure to pay taxes on time in one period leads to ineligibility for the lottery (classification as a “bad taxpayer”) for the subsequent year. Also, because taxpayers

⁴³See Gerber and Green 2012: 39-43 or Dunning 2012: 118-121.


⁴⁴In section 4, we discuss another core assumption: non-interference. In brief, news of the lottery does not substantially travel from winners to non-winners: only 8% of non-winners have heard of the existence of the lottery.

⁴⁵These are computed for each time period separately: we estimate the mean and standard error for the groups with winning and losing numbers at each time period, and use normal approximations for the confidence intervals (since the N s are large).

⁴⁶Some physical persons might not also go to City Hall to claim the tax holiday when notified they won. Averaging across taxes, only 36% of eligible taxpayer accounts claim the tax holiday. Formally, this is a situation of “single crossover” from the treatment to the control group (Dunning 2012: 136-143).

may stay current on payments while not paying off older debt, the overhang of missed payments lasts substantially longer—for over three years. This has important policy and welfare implications we discuss later, but for purposes of assessing effects on habit, the “paid on time” measure in Figure 2 is most theoretically relevant.

Figure 2: **Natural Experiment: Balance Tests, Placebo Outcomes, and Treatment Effects**



The figure depicts balance tests, placebo outcome tests, and treatment effects in the natural experiment. The horizontal axis measures tax payment periods before or after the tax holiday. The grey vertical strip indicates the period of the tax holiday (the treatment). Comparisons between winners and non-winners to the left of zero—the date at which each taxpayer won or could have won a particular tax holiday lottery—test for balance on pre-treatment tax compliance. Post-treatment comparisons between ineligible taxpayers with winning and losing numbers (bottom panel) provide a placebo outcome test, since none of these taxpayers actually won a tax holiday. Finally, post-treatment comparisons among eligible taxpayers estimate the treatment effects of the tax holiday (top panel). The vertical axis shows the complier average causal effect (CACE) for the proportion of taxpayers who paid on time, at each payment period. Compliers are taxpayers with winning account numbers who would claim the exoneration. Vertical lines show 95% confidence intervals. See also Table 3.2.

In sum, rather than fostering greater tax compliance, winning the tax holiday lottery inhibits it.

We see both a persistent negative effect on compliance and also attenuation of the effect over time.

These features of the data are consistent with knock-on effects of habit disruption discussed in Section 2 but also with their eventual decay, and are important for distinguishing habit from other potential mechanisms that might explain our findings, a point to which we turn next.

Table 3.2: **Natural Experiment: Average Causal Effects**

	Control Group Average	\widehat{ACE}	$\widehat{SE}(\widehat{ACE})$	p -value	\widehat{CACE}
Post Tax Holiday Payment 1	0.93	-0.03	0.01	0.00	-0.04
Post Tax Holiday Payment 5	0.91	-0.02	0.01	0.04	-0.02
Post Tax Holiday Payment 10	0.87	0.01	0.01	0.55	0.01
Post Tax Holiday Payments 1-10	0.90	-0.02	0.01	0.00	-0.03
Post Tax Holiday Payments 1-10 (Property)	0.92	-0.03	0.01	0.00	-0.03
Post Tax Holiday Payments 1-10 (Head)	0.91	-0.01	0.01	0.15	-0.02
Post Tax Holiday Payments 1-10 (Sewage)	0.93	-0.02	0.01	0.04	-0.04
Post Tax Holiday Payments 1-10 (Vehicle)	0.76	-0.00	0.02	0.98	-0.00

The table shows the estimated average causal effects of treatment assignment (i.e., of a winning lottery number) for eligible taxpayers, denoted \widehat{ACE} . The estimated standard error is $\widehat{SE}(\widehat{ACE})$. The outcome variable is the proportion of taxpayers who paid on time. The first three rows show the effect for the 1st, 5th, and 10th payments due after the end of the tax holiday. The fourth row shows the effect for the first ten post-holiday payments combined. The final four rows disaggregate the effect on the combined outcome by type of tax. The final column shows the estimated complier average causal effect (\widehat{CACE}).

4 Is habit the mechanism?

Is the disruption of habit the mechanism that drives our results? Good students who are given a year off of homework may be less diligent when made to do homework again. Is a similar force to blame for the negative effects of the holiday on tax compliance?

We take several approaches to answering this question. First, we assess the direct evidence for this proposition, showing that the negative impact only holds in settings where the habit of paying taxes is plausibly interrupted. Second, we show evidence from field and survey experiments that casts substantial doubt on the plausibility of alternative explanations. No single piece of evidence is dispositive in this inquiry; however, our results as a whole strongly corroborate the habit hypothesis.


One first piece of evidence comes from variation in effects across types of taxes, already shown in

rows 5-8 of Table 3.2. Although we were unaware of this when we began our study, with the vehicle tax—unlike property, sewage, and head taxes—winning taxpayers typically continue to pay a small amount of taxes after winning the lottery. 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 within the following year, because the vehicle tax is often increased annually, or because inflation of the nominal value of the vehicle fee (as opposed to the nominal value of the credit) leads to the need to continue making small payments. In sum, winning the vehicle tax lottery does not in fact involve a disruption in the habit of paying taxes to the same degree as other municipal taxes.

If our argument about habit is correct, we should therefore expect weaker or null effects for the property tax. Figure 3 graphically disaggregates the effect of tax holidays by type of reward. The figure presents data only for eligible taxpayers, to focus on treatment effects rather than placebo tests; as in Figure 2, we present estimated complier average causal effects for the proportion of taxpayers who pay on time at each payment period. As the figure suggests, the effects for the taxes without a tax holiday reward (the vehicle tax) are null, while the effect for the taxes with a tax holiday reward are similar to those in figure 3 (see also Table 3.2).⁴⁷ To be sure, these results are not dispositive: vehicle taxes may differ from other taxes in ways that are relevant for the effect of the tax holiday. However, the results are strongly suggestive of the influence of habit, since we see null effects for the one tax where payment is not interrupted.

⁴⁷The online Appendix shows the effect graphically for each of the four taxes separately.

Figure 3: Treatment Effects By Type of Tax: Holiday vs. No Holiday



The top panel shows treatment effects for property, sewage, and head taxes, for which winning the lottery leads to interruptions in tax payment (“Holiday.”) The bottom panel shows treatment effects for property tax, where the habit of payment is not in fact interrupted (“No Holiday”).

A similar but perhaps stronger test comes from the fact that around 21% of the taxpayers in our natural-experimental treatment and control groups had enrolled in automatic payment plans at the date of the relevant lottery.⁴⁸ For them, the tax holiday also does not interrupt the habit of payment. Moreover, after the conclusion of the holiday, payments resume without any action on the part of the taxpayer.⁴⁹ The forces of habit therefore could not conceivably operate to generate negative compliance effects for this group of automatic payers. As a placebo outcome test, we can therefore assess whether there is any effect of winning on automatic payers—that is, a group for which there is a “known” null effect.

As Figure 5 confirms, the only detectable negative effect of the tax holiday appears among manual taxpayers: when the habit of paying is not interrupted, the lottery does not induce non-compliance. Our

⁴⁸The proportions are 20% for property tax, 15% for sewage, and 26% for head; we do not have data on automatic payment for the vehicle tax. For the property tax, we have missing data on type of payment for 7 of the 2,693 eligible property tax payers described in Table 3.1.

⁴⁹Only in a few cases where there is some problem with the payment scheme would the taxpayer need to become involved.


focus on habit suggests that we should only find significant treatment effects for manual taxpayers—and that is exactly what we find. To be sure, taxpayers who pay automatically could differ from manual payers in various ways that could be related to compliance. One might think, for example, that automatic payers tend to be of higher socioeconomic status and therefore less prone to the negative effects of habit disruption. Surprisingly, the property values of manual taxpayers are in fact substantially greater, by a factor of almost 2.⁵⁰ In addition, while municipal data do not allow us to measure precisely which taxpayers are physical persons and which are firms, 38% of good taxpayers who pay automatically claimed the exoneration while just 26% of manual payers did so. This suggests a greater prevalence among the automatic payers of physical persons—whom we might otherwise expect to be more prone to the negative impact of holidays, if firms are more organized about paying their taxes.⁵¹ It is also worth noting that our tests for automatic payers also have lower statistical power than our tests for manual payers, which could conceivably account for the null effects. Notwithstanding these caveats, for the property, head and sewage taxes, we have over 1190 automatic payers from the 5,597 good taxpayers described in Table 3.1 with which to estimate effects.⁵² This is therefore an important further piece of evidence that corroborates the habit hypothesis.

⁵⁰The difference is 2,088,634 vs. 1,142,571 pesos, on average (2004 data).

⁵¹However, it is also the case that automatic payers are more likely to pay the entire year of taxes in advance (41% versus 28%) and thus could in general be “better compliers” than manual payers. For our entire sample of property taxpayers (both good and bad), about 10% had automatic payment in 2004; this increased marginally over time to 13% in 2014. So good taxpayers are more likely to have automatic payment plans, as one would expect.

⁵²Note that we do not have data on the type of payment for the vehicle tax.

Figure 4: **Placebo test: Treatment Effects for Automatic Payers**




The figure depicts the effect of winning the holiday for property, head, and sewage taxes. The top panel shows effects for all taxpayers; the middle panel effects for those enrolled in automatic debit at the time of the lottery; and the bottom panel effects for those who paid manually. As in Figures 2 and 3, values to the right of the grey strip estimate treatment effects, while those to its left provide graphical balance tests. There is no discernible effect for automatic taxpayers, whose habit is not interrupted by the tax holiday.

A final piece of evidence stems from the comparison of effects of taxpayers who have complied punctually over the 15 periods prior to winning the lottery—and who therefore have a strong “stock” of habit—with those “marginal” eligible taxpayers who have not always previously complied. Our model in section 2.1 suggests that the tax holiday should have a negative effect on future compliance for both types—but that it should be less for always-compliant taxpayers, whose stock of previous habit should allow them to bounce back more quickly from the interruption. Figure ?? shows exactly this pattern: effects are more pronounced and last longer for the marginal taxpayers. Differences in the effects are themselves statistically significant for seven payment periods after the resumption of tax obligations.

In sum, our evidence is strongly consistent with the claim that the disruption of a habit drives the

negative effect of tax holidays. What type of habit interruption does the holiday generate? We conducted 20 qualitative, open-ended interviews with taxpayers who won a property tax lottery in 2014-2015 to gain insight into this question. Virtually without exception, interviewees paid their tax bill at a local kiosk before winning the lottery and continued to do so after the tax holiday.⁵³ Thus, the method of payment did not change for 18 of our 20 interviewees. Many interviewees also recalled receiving a tax bill from City Hall showing a zero balance, during the period of exoneration. Thus, the context of payment did not change, and winners received notices from the municipal government reminding them that they are taxpayers. What changed is the fact of payment: the tax holiday interrupted the actual compliance behavior of winners, while leaving other elements of the tax payment context fixed.

Figure 5: **The Stock of Habit: Always-Compliant Versus Marginal Taxpayers**



The figure compares effects for taxpayers who complied punctually in every period prior to winning the lottery (“Always Compliant”) and those who had failed to pay taxes on time at some point during this period (“Marginal Taxpayers”). Taxpayers in both groups had complied punctually over the three periods prior to the relevant lottery, as required to be eligible for the holiday. The differences in effects in the top and bottom panel are statistically significant at standard levels for the first seven payment periods after the end of the holiday.

⁵³Only one paid continuously by automatic debit (CTA 592794). Another had changed the method of payment but attributed that to being retired and not wanting to leave the house for security reasons (CTA 572594).

4.1 Alternative explanations

4.1.1 Informational mechanisms

We turn finally to possible alternative explanations. The most important is informational. Winning the tax holiday provides a year's exoneration and therefore disrupts the habit of paying taxes. Yet, receiving an award letter from the municipality also informs many taxpayers of the existence of the policy. In our 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 know someone who has won the lottery.⁵⁴

Is it possible that information about the policy, rather than the tax holiday itself, exerts negative effects on future compliance? Taxpayers might interpret the new information as a negative signal of municipal capacity or the attractiveness of paying taxes. Perhaps taxpayers 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 a *positive* effect on compliance; for example, the lottery could boost perceptions of the transparency or equity of the tax system, thereby increasing willingness to pay.⁵⁵ These are important possibilities too, though they would tend to counteract the overall negative effect of the tax holiday on compliance.

To evaluate these alternatives, we 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 lottery. We then compare the subsequent tax payment behavior of this group to a randomly assigned placebo control group that received only a reminder, as well as a pure control group. These are a sub-set of interventions from a larger experiment that also included conditions reminding taxpayers about punishment for non-payment.⁵⁶ In this paper, we focus only the interventions that are theoretically relevant for alternative explanations. However, we ensure that all our results are robust to corrections for multiple compar-

⁵⁴These percentages are drawn from the placebo control group in our field experiment described below, whom we did not inform about the existence of the lottery (N=412).

⁵⁵We discussed both positive and negative informational mechanisms in our pre-analysis plan.

⁵⁶Several studies have found that messages increasing the salience of penalties can increase compliance; see Slemrod, Blumenthal and Christian (2001); Kleven et al. (2011); Castro and Scartascini (2015). In our study, that is not the case.

isions reflecting the full design, per our pre-specified analysis plan.⁵⁷ Our analysis of the effects of furnishing taxpayers with information on the tax holiday—without awarding of an actual interruption of payments—allows us to unbundle informational effects from the impact of habit disruption.

Specifically, we collaborated with the municipal government to design and mail flyers printed with different messages for our intervention groups. We focus on property taxes here both because of the importance of this tax and to eliminate potential sources of heterogeneity that would decrease statistical precision. Our baseline reminder serves as a placebo control condition:

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).⁵⁸

Our next condition repeats this baseline reminder but adds information about the existence of the tax holiday lottery:⁵⁹

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. Indeed, the municipality sometimes prints messages to taxpayers on the inside of tax bills.⁶⁰ The experience of receiving our flyers stamped with the municipal government’s logo would thus be similar to the experience of receiving a tax bill on which the municipality prints encouragements to pay taxes.⁶¹

⁵⁷We also present complete results in the online Appendix.

⁵⁸The bold text is as in the original Spanish; see the online Appendix for the Spanish-language flyers.

⁵⁹We varied whether this information emphasized the individual or social returns of the lottery; the quoted text is from the individual condition. As also specified in our pre-analysis plan, our registered analyses often pool these into a single “reminder + existence of lottery” condition, as we do here.

⁶⁰For our experiment, we opted to send separate flyers that look like tax bills due to logistical considerations: rigorously ensuring that messages printed on each actual tax bill would follow treatment assignment proved infeasible.

⁶¹The design of the flyers and the messages printed on them were authorized by the municipal government, and our experimental protocol was approved by Berkeley’s institutional review board.

To create our study group for the experiment, we worked with the municipal bureaucracy to draw a random sample of administrative tax payment records.⁶² The population from which our household survey and administrative data samples are drawn should be conceptualized as “all property tax-paying households with bills due in July 2014.”⁶³ We drew a stratified random sample of eligible and ineligible taxpayers, the former being those who could take part in the tax holiday lottery as of March 2014, on the basis of punctual compliance over the previous year.⁶⁴ 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.⁶⁵

We then randomized these sampled taxpayers to intervention groups. Table 4.3 shows the sample sizes for each treatment condition. Balance tests show that the treatment groups are statistically balanced on pre-treatment covariates.⁶⁶ Our flyers were distributed in phases in June, such that they would arrive approximately 8 days before tax due dates, which differ slightly for different households.⁶⁷ We use two behavioral tax compliance measures—for which we have full coverage—to evaluate the effects of our treatments. The first is whether the account holder accessed his or her Web account, for example, to print a duplicate bill (we also obtained these data from the municipality); we call this Intended Compliance. Second, we measure whether the account holder Paid On Time, both in July 2014—the first payment period immediately following our intervention—and on average over the period July 2014 to July 2016.

Does information about the lottery affect compliance behavior among eligible or ineligible taxpayers? Figure 6 compares the compliance of those sent a reminder of the tax due date, plus information

⁶²In detail, we set a seed, sequentially generated random numbers of the same length as taxpayer account ID numbers, and then drew the tax records corresponding to those account numbers.

⁶³For example, we filtered out certain tax payers who pay their property tax bills for the whole fiscal year in March or November (and thus would not conceivably be influenced by our field experimental intervention in June-July 2014), and taxpayers who are not in fact issued bills.

⁶⁴Note that the average value of the “good taxpayer” variable for property tax is around 53%; however, we ended up with a somewhat larger group of good taxpayers, because initially sampled accounts for bad taxpayers had higher rates of taxpayers with invalid addresses or due dates, to whom the municipality did not in fact issue bills. One might imagine that bad taxpayers are concentrated in poorer, outlying areas of Montevideo. In fact, per Appendix Figure 10, good and bad taxpayers are evenly spread throughout the city.

⁶⁵Note, however, that having won in the past—even in the immediately preceding year—does not disqualify a taxpayer from winning again.

⁶⁶See the online Appendix.

⁶⁷See our pre-analysis plan and amendments thereto for implementation details.

Table 4.3: **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
Total	N=14,784	N=13,862

The table depicts assignment to a subset of treatment conditions in our field experiment. Here, total N=28,646. See the Appendix for the full design.

about the lottery, to those sent only a reminder (the placebo control group).⁶⁸ As the figure shows, there is no difference between the information group and the placebo control group for the Paid On Time measure, either in the first period after the intervention or for the longer period following it. As for the Intended Compliance measure, eligible taxpayers (but not ineligible ones) appear to access their accounts less in the reminder group in the period immediately following the intervention, but this apparent effect disappears over the longer period where we have a very precisely estimated null effect—suggesting that the effect is a statistical fluke or a very short-run impact. (It is possible that the greater amount of text on the flyer with lottery information actually lessens the impact of the reminder itself). We verify as well that these null effects are not due to inattention on the part of taxpayers to the flyers themselves: relative to the pure control group, both the placebo control and the lottery information treatments elevate the measure of Intended Compliance, and they did so both for eligible and ineligible taxpayers. We also see an effect among ineligible taxpayers for the Paid on Time measure. However, the placebo control treatment has at least as large an effect as the treatment. Information about the lottery appears to have little to no impact on compliance.⁶⁹

Nor does the lottery appear to act as a negative signal about municipal capacity or the equity of the tax system. Indeed, in a separate survey experiment conducted with respondents in several thousand households, we compared respondents who were informed about the lottery using language very similar to that printed on our flyers, to those who were told that the municipality “from time to time” selects good taxpayers and rewards them with a year free of tax payment. The former treatment was therefore

⁶⁸This comparison of treatment to placebo control is our pre-registered analysis; we also discuss comparisons to the pure control group.

⁶⁹See the online Appendix.

intended 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.⁷⁰ We gathered outcome data on five measures of attitudes towards the municipal government and the tax system. As shown in Figure 7, here we find no effect of information about the lottery on perceptions that the municipal government does a good job, municipal taxes are just, or that rewards for good taxpayers are a waste of money, we find that being informed about the lottery boosted perceptions of transparency and equity—in particular, decreased agreement with the statement that rewards “go to the same people as always.” Information about the lottery also boosted agreement that it is worth it to be up to date on ones taxes, even though our behavioral evidence suggests no positive effect on compliance. In sum, information about the lottery does not seem to act as a negative signal about municipal capacity—and informational mechanisms do not appear to account for the impact of the holiday on compliance.

4.1.2 Other alternatives

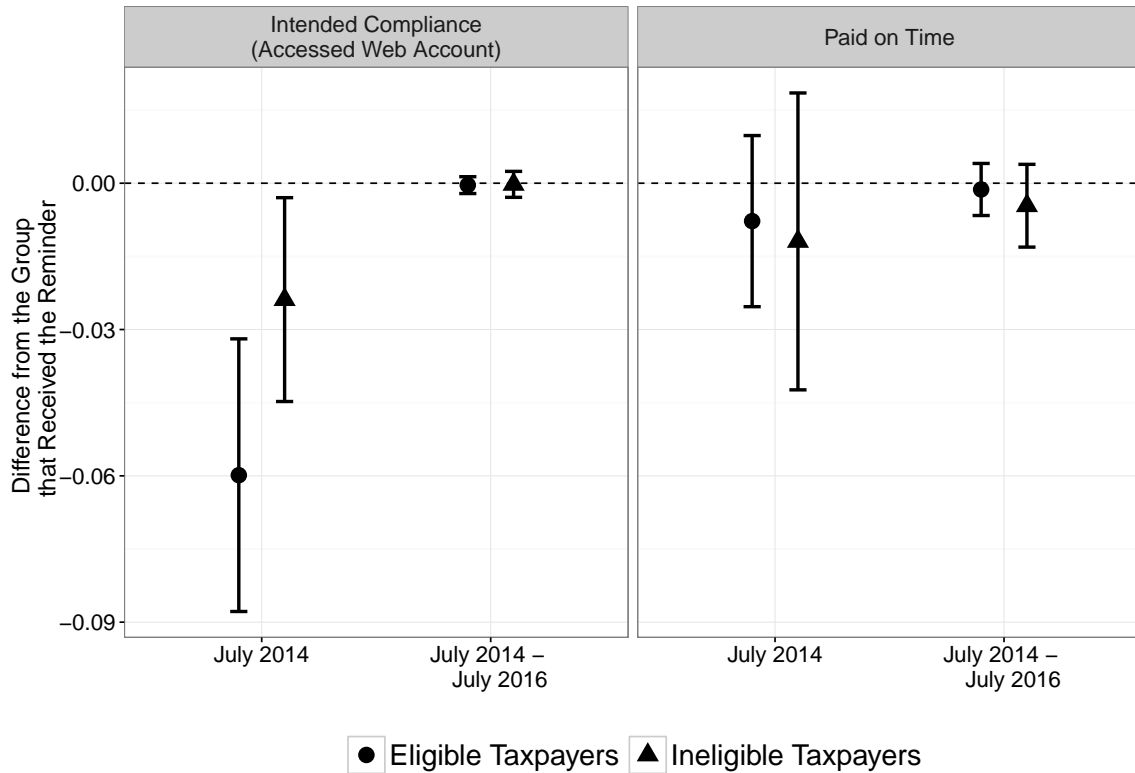
A different set of alternative explanations might focus on the impact of winning the lottery itself—rather than the tax holiday to which winning entitles eligible taxpayers. For example, winners might infer (incorrectly) that having won the lottery once, their probability of winning the lottery a second time is lowered, which could lead to lower compliance among winners if the lottery indeed exerts a positive incentive effect. Our household survey data do support the existence of such misperception.⁷¹ Yet, to explain the negative effect of winning on future compliance, the lottery itself would have to exert a powerful positive incentive: otherwise, lottery winners wouldn’t be induced to comply at *lower* rates than non-winners by the false presumption that their chances of winning again are lowered. And the lottery does not appear to have a positive incentivizing effect, as we have just shown.

⁷⁰Logistical problems in the execution of our survey caused substantial delays and also did not allow us to reach the full study group in the field experiment. We therefore rely on the survey experiments for evidence about the impact of information on attitudes, while we use the field experiment and administrative data to study actual behavior.

⁷¹Indeed, 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.

Figure 6: **Field Experiment: Effects of Information About the Holiday on Tax Compliance, Relative to Reminder Group**

(Eligible and Ineligible Taxpayers)



The figure depicts the effects of providing a reminder about the property tax bill due date and information about the existence of the tax holiday lottery, relative to a placebo control group that received only a reminder. We present estimates separately for taxpayers who were eligible to win the lottery at the time of intervention and those who were ineligible but who would become eligible if they brought their payments up to date during a maximum of one year. Here, we pool informational treatments that emphasized the individual or social benefits of the tax lottery. Intended Compliance measures whether the account holder accessed his or her Web account, for example, to print a duplicate bill (we also obtained these data from the municipality); Paid On Time measures whether the taxpayer paid the bill punctually. Both outcomes are measured in July 2014—the first payment period immediately following our intervention—as well as on average, over the period July 2014 to July 2016.

Another theory could be that the lottery gives winners additional income for a year, which generates substitution or income effects that last beyond the holiday period. For example, taxpayers may find it difficult to revert back to a lower consumption level at the end of the holiday and therefore do not fulfill their tax obligations. This appears plausible as a theoretical matter. We also asked property

tax holiday winners in qualitative interviews what they did with the additional income after winning the lottery; many did mention using this to cover other costs. However, a powerful piece of evidence against this conjecture is that with the vehicle tax, winners also receive an important temporary shock to their income when they win. Yet, as explained previously, they go on paying some small amount of vehicle tax, so their habit is not disrupted. If substitution or income effects rather than habit disruption explain our finding, we should observe a negative effect for winners of the vehicle tax as well—yet we instead observe a null rather than negative effect (per Figure 3). Note also that per our qualitative interviews, the importance of the rebate varied across taxpayers, however. One spoke fondly of going out to tea or to lunch; but another said “it didn’t change my life at all, it was enough to buy a good pair of shoes.”⁷² We might thus expect any substitution effect to be more important for different sorts of taxpayers. However, we have not found evidence for heterogeneous effects by income or past payment history.⁷³

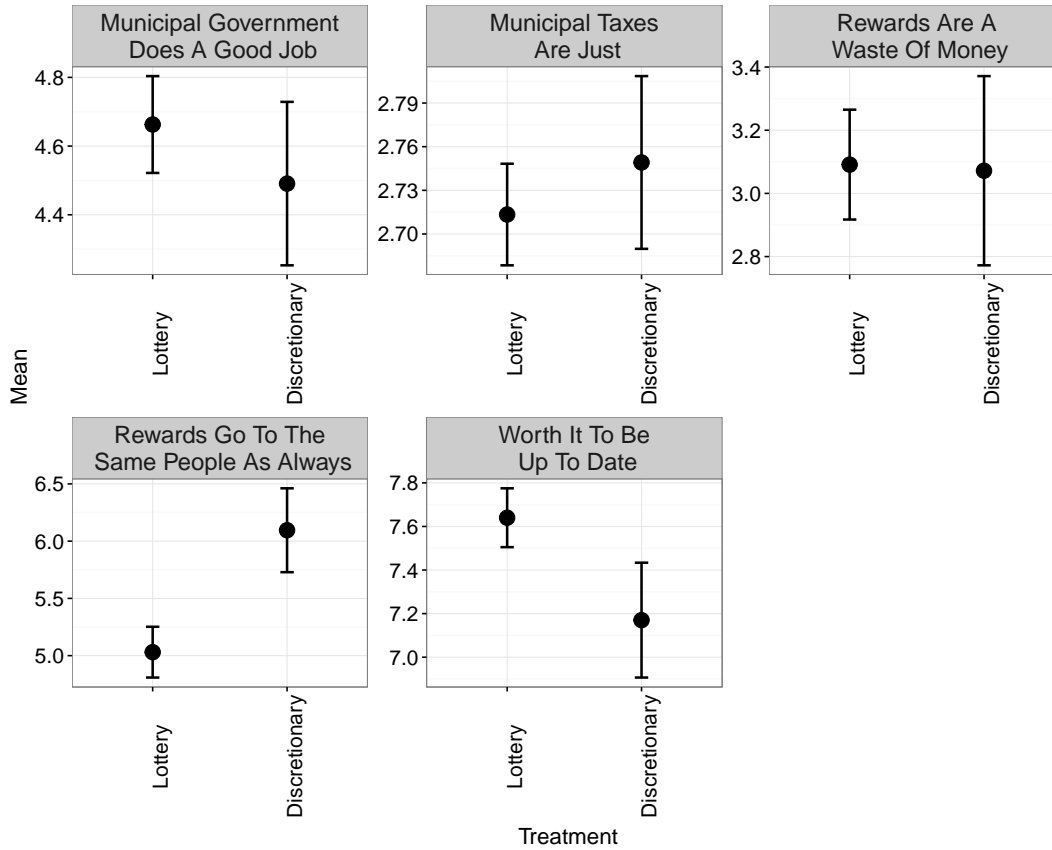
Finally, the holiday could have other kinds of behavioral effects. For example, by breaking a behavioral “taboo” against non-compliance, experiencing the exoneration may make good taxpayers less willing to pay. This explanation is not unrelated to habit, though not exactly the same. Yet, we would expect the effect of such a shock to be quite persistent: breaking the taboo might lead to a more or less permanent reduction in the willingness to pay taxes, rather than producing an effect that is substantial but that decays over time—as we instead find in our data. By emphasizing the material reward for paying taxes and thus stimulating “extrinsic” incentives, the lottery might instead crowd out “intrinsic” incentives, such as a sense of civic duty.⁷⁴ However, we would also expect such crowding out of intrinsic incentives to be more or less permanent, once taxpayers know about the rewards policy. Moreover, we would also expect this effect to become operative once taxpayers learn about the tax holiday—yet we see no longer-term negative effect of information in our field experiment.

⁷²Interviewees CTA 512794 and CTA 334095, respectively.

⁷³Analyses reported in the online Appendix.

⁷⁴Gneezy and Rustichini (2000)

Figure 7: Survey Experiment: Effects of Information about the Lottery on Attitudes Towards Taxation



The figure shows effects in our survey experiment on five measures of attitudes towards taxation. Respondents in the “lottery” group were informed about the reward lottery using language similar to that printed on our mailed flyers. Respondents in the “discretionary” group were instead told that the municipality “from time to time” selects good taxpayers and rewards them with a year free of tax payment. Differences between point estimates in the top panel are not statistically significant; those in the bottom panel are significant.

Thus, these alternative informational and behavioral explanations cannot explain key features of our data—while habit disruption does. Put in terms of our behavioral model, it does not appear that the tax holiday greatly shapes parameters such as b , the expressive benefit of paying taxes; and if it does improve perceptions of the transparency or equity of the tax system, that change is not sufficient to counteract the negative impact of the holiday on compliance behavior. Instead, the holiday seems to disrupt taxpayers’ “stock” of compliance habit, with important negative consequences for the flow of subsequent compliance. This effect can be severely disruptive, yet it also can eventually recede—which may only underscore that even if disruptions matter, habits can also be difficult to al-

ter permanently. Our evidence therefore indicates the importance of the forces of habit, in a highly consequential arena for state capacity as well as political development.

5 Conclusion

Habit may play an important role in a wide range of political phenomena—especially, participation in civic acts such as voting, protest, or the payment of taxes. Social psychologists have recognized the importance of repetition of a behavior in a similar context for engendering continuity of behaviors and outcomes. While voting scholars have given more attention to habit, its influence on many important modes of citizen-state interaction is substantially underexplored.

For policy makers, lack of attention to the forces of habit can lead to perverse consequences. Municipal governments throughout Latin America have developed positive incentive schemes such as randomized lotteries to boost tax compliance. Montevideo, one of the pioneers of such policies, has offered good taxpayers a chance to forego their tax payments for one year, in the hopes of both rewarding and inducing compliance. Yet, our findings suggest the tax holiday program is a net revenue loser, not only because of foregone tax payments during the tax holiday—but because by disrupting the habit of payment, the holiday actually depresses future compliance. Our field experiment also suggests no offsetting benefit of incentivizing tax payments among bad taxpayers. Of course, there can be other rationale for maintaining such programs, such as engendering positive attitudes among citizens towards the transparency or equity of the tax system. However, policy makers ignore the importance of habit disruption at their peril. Indeed, after we presented findings from our study to bureaucrats in Montevideo, the municipality shifted towards a policy of cash rebates for eligible lottery winners, rather than tax holidays.⁷⁵ Many tax lotteries elsewhere in Latin America do not feature holidays, so they may not be subject to the particular pathology we identify here, yet their impacts have not been rigorously assessed using strong designs and a combination of data sources, as we do here.

From the perspective of social science, the findings underscore the importance of habit as a force for political action, in a highly consequential realm of citizen-state interaction. Social scientists have long

⁷⁵See a description in Spanish at <http://www.montevideo.gub.uy/institucional/noticias/nuevo-plan-de-reconocimiento-a-contribuyentes-y-regularizacion-de-adeudos>, accessed September 15, 2016.

recognized the ways in which countries may be locked into high or low development paths, especially with respect to outcomes such as state capacity. Some of these traps emerge in equilibrium from strategic behavior, while others are due to increasing returns or related sources of path dependency.⁷⁶ The idea simply that habit leads to recurrence is less well explored both theoretically and empirically, in part because habit is usually so difficult to separate from confounding factors that may result in recurrence. Where the forces of habit have been better studied—as in the literature on voting—research has focused on how engaging in a particular behavior such as voting may engender a persistent habit.

Yet, political scientists have been less attentive to the consequences of habit disruption. Many policies and interventions can involve such interruptions: soldiers are given leaves of absence, students take gap years, professors are granted sabbaticals. These policies may have many offsetting benefits (especially sabbaticals), but they could also have largely unremarked negative impacts through the channel of habit. Even research on the negative effects of summer vacations on scholarly performance tends to focus on the differences across children in opportunities to learn during the break, rather than the impact of habit disruption itself.⁷⁷

The good news, perhaps, is that the effects of habit disruptions are persistent but also decay, and interventions can build habits as well as destroy them. If left unchecked by countervailing forces—in our study, the fact that taxpayers who were given exonerations were reliable tax compliers to begin with—habit disruptions could have long-lasting consequences, however. States and social scientists alike should therefore consider the impact of habit in generating vicious as well as virtuous cycles in civic participation.

⁷⁶Pierson (2000); Collier and Collier (1991); Acemoglu and Robinson (2012).

⁷⁷Cooper and Greathouse (1996)

References

- Acemoglu, Daron and James Robinson. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. Crown Business.
- 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.
- Alm, James, Michael McKee, Betty R Jackson and Tax Analysts. 1990. *Alternative Government Approaches for Increasing Tax Compliance*. Tax Analysts.
- Alm, James, Michael McKee and William Beck. 1990. "Amazing grace: Tax amnesties and compliance." *National Tax Journal* pp. 23–37.
- Angrist, Joshua D. and Alan B. Krueger. 2001. "Instrumental Variables and the Search for Identification: From Supply and Demand to Natural Experiments." *Journal of Economic Perspectives* 15(4):6985.
- Ansolabehere, 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.
- Bates, R.H. and D-H. L. Lien. 1985. "A Note on Taxation, Development, and Representative Government." *Politics & Society* 14:53–70.
- Besley, T. and T. Persson. 2009. "The Origins of State Capacity: Property Rights, Taxation and Politics." *American Economic Review* 99(4):1218–44.
- Blumenthal, Marsha, Charles Christian, Joel Slemrod and Matthew G Smith. 2001. "Do normative appeals affect tax compliance? Evidence from a controlled experiment in Minnesota." *National Tax Journal* pp. 125–138.
- Castro, Lucio and Carlos Scartascini. 2015. "Tax compliance and enforcement in the Pampas: Evidence from a field experiment." *Journal of Economic Behavior and Organization* 116(C):65–82.
- Chetty, Raj, Mushfiq Mobarak and Monica Singhal. 2014. "Increasing Tax Compliance Through Social Recognition." *Policy Brief*.
- Collier, Ruth Berins and David Collier. 1991. *Shaping the Political Arena: Critical Junctures, the Labor Movement, and Regime Dynamics in Latin America*. Princeton University Press.
- Cooper, Harris, Barbara Nye James Lindsay and Scott Greathouse. 1996. "The Effects of Summer Vacation on Achievement Test Scores: A Narrative and Meta-Analytic Review." *Review of Education Research* 66(3):227–268.
- Dafoe, A. and G. Tuñón. 2014. "Placebo Tests for Causal Inference." *International Studies Association Conference, Toronto, Canada. March 27, 2014*.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. Harper and Row, N.Y.

- Dunning, Thad. 2012. *Natural Experiments in the Social Sciences: A Design-Based Approach*. Cambridge University Press.
- Falkinger, Josef and Herbert Walther. 1991. "Rewards versus penalties: on a new policy against tax evasion." *Public Finance Review* 19(1):67–79.
- Fellner, Gerlinde, Rupert Sausgruber and Christian Traxler. 2013. "Testing enforcement strategies in the field: Threat, moral appeal and social information." *Journal of the European Economic Association* 11(3):634–660.
- Freedman, David. 1991. "Statistical Models and Shoe Leather." *Sociological Methodology* 21:291–313.
- 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.
- Gneezy, Uri and Aldo Rustichini. 2000. "A Fine Is a Price." *The Journal of Legal Studies* 29:1–17.
- Holland, Alisha C. 2016. "Forbearance." *American Political Science Review* 110(2):232–46.
- Kleven, Henrik Jacobsen, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen and Emmanuel Saez. 2011. "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark." *Econometrica* 79(3):651–692.
- Loftus, Elizabeth F. 1985. "To file, perchance to cheat." *Psychology Today* 19(4):34–39.
- 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.
- Parle, William M and Mike W Hirlinger. 1986. "Evaluating the use of tax amnesty by state governments." *Public Administration Review* pp. 246–255.
- Pierson, Paul. 2000. "Increasing Returns, Path Dependence, and the Study of Politics." *The American Political Science Review* 94(2):251–267.
- Pomeranz, Dina. 2013. No taxation without information: Deterrence and self-enforcement in the value added tax. Technical report National Bureau of Economic Research.
- Slemrod, Joel, Marsha Blumenthal and Charles Christian. 2001. "Taxpayer response to an increased probability of audit: evidence from a controlled experiment in Minnesota." *Journal of public economics* 79(3):455–483.
- Smith, Kent W and Loretta J Stalans. 1991. "Encouraging Tax Compliance With Positive Incentives: A Conceptual Framework and Research Directions." *Law & Policy* 13(1):35–53.
- Snow, John. 1855. *On the Mode of Communication of Cholera*. London: John Churchill, New Burlington Street, England.
- Soifer, Hillel. 2008. "State Infrastructural Power: Approaches to Conceptualization and Measurement." *Studies in Comparative International Development* 43(3-4):231–251.