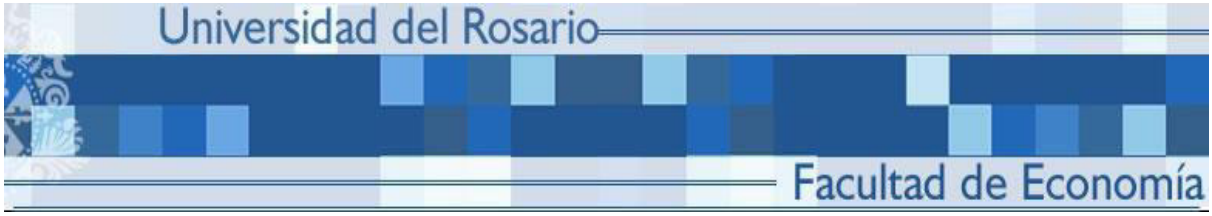

Can Facebook Ads and Email Messages Increase Fiscal Capacity? Experimental Evidence from Venezuela

Jorge Gallego
Federico Ortega

SERIE DOCUMENTOS DE TRABAJO

No. 224

Noviembre de 2018

Can Facebook Ads and Email Messages Increase Fiscal Capacity? Experimental Evidence from Venezuela*

Jorge Gallego
Universidad del Rosario

Federico Ortega
UCAB

November 19, 2018

Abstract

Tax compliance is an important challenge in weakly institutionalized countries, in which citizens do not trust the State and prefer to evade taxation. However, e-government strategies may improve fiscal capacity, as the transaction costs of compliance are reduced and more information from taxpayers is gathered and exploited. Can compliance be increased, and hence fiscal capacity strengthened, using online communication strategies that exploit these tools and sources of information? We perform a randomized field experiment in the capital of Venezuela, Caracas, to determine if online strategies, namely email reminders and targeted Facebook advertisements, can increase tax compliance. We vary the mechanism used to approach taxpayers to test if more direct and personalized methods, such as email messages, are more effective than general advertisement tools, such as Facebook ads. Moreover, our design allows us to test potential complementarities between these strategies thus boosting the capacity of the local government to increase compliance. We find that these strategies are cost-effective methods for increasing tax revenues, but that the effects vary across different types of taxpayers.

Keywords: Tax compliance, Randomized controlled trial, Fiscal capacity, Online strategies

JEL Codes: C93, H26, H71, O12

*We would like to thank Ricardo Salazar and the whole SEDAT team, as well as CAF's Directorate of Impact Evaluation and Policy Learning, for their support throughout this project. We received generous comments from Adriana Camacho, Rachid Laajaj, Stanislao Maldonado, Lesbia Maris, Jan Pierskalla, Mounu Prem, Daniel Ortega, as well as seminar participants at CAF, Universidad del Rosario, Universidad de los Andes, MPSA, and APSA 2018. Armando Flores, Maria Andreina Cantele and Maria Paula Medina provided superb research assistance. As usual, all errors remain ours.

1. Introduction

One of the most important components of state capacity is the ability of governments to tax citizens (North and Weingast, 1989; Besley and Persson, 2009, 2013). Tax revenues may be spent in public goods and investments aimed at improving the economy's productivity (Barro, 1990), raising the growth rate and minimizing the efficiency losses imposed by taxes (Barro and Sala-i-Martin, 1991). However, tax evasion is a common problem in developing countries (Besley and Persson, 2014), where citizens, for different reasons, do not fulfill their obligations and prefer to evade. In Latin America, for instance, tax evasion totals an estimated US $340,000 million, i.e. 6.7% of the regional GDP (CEPAL, 2016). The problem is even more serious in places where taxpayers distrust the state and where institutions are so weak that deterring evasion proves to be extremely difficult, if not impossible. Recent research shows that revenues raised by local governments themselves, such as property taxes, as opposed to transfers made by the central government, tend to be spent more efficiently and have more impact on public service delivery and citizen welfare (Gadenne, 2015; Martinez, 2017). Consequently, understanding which channels are more effective in increasing compliance and how to strengthen fiscal capacity is a major challenge for scholars and policymakers.

Recent literature on this issue has explored the behavioral features underlying tax compliance, both from a theoretical and an empirical perspective. From a theoretical point of view, the canonical model of evasion assumes that individuals weigh up the benefits of not paying taxes, with the probability of being caught and the costs of being punished (Allingham and Sandmo, 1972). However, given that the predictions of this model do not reconcile the empirical fact that people do pay taxes on a large scale and under different circumstances, several theoretical alternatives have been proposed. Such alternatives incorporate modifications on preferences and utility, new behavioral assumptions, or social interactions (Myles and Naylor, 1996; Yaniv, 1999; Snow and Warren, 2005; Dhami and al Nowaihi, 2007; Traxler, 2010).

From an empirical point of view, a recent wave of experiments has focused on strategies to increase tax compliance. Most of these papers have studied what type of messages, information, or behavioral incentives, are more effective to encourage citizens to pay their taxes (Slemrod et al., 2001; Ortega and Sanguinetti, 2013; Ortega et al., 2013; Torgler, 2004; del Carpio, 2014; Hallsworth et al., 2017; Castro and Scartascini, 2015; Gallego and Ortega, 2018). The effects of deterrence, reciprocity, peer effects, and transaction costs,

among other factors of interest, on tax compliance have been analyzed through lab, survey, and field experiments, with varying results. However, the role played by the method used to convey compliance messages has been placed under much less scrutiny, despite the fact that the total cost of any moral suasion strategy strongly depends on this factor (with Ortega and Scartascini (2015) being a notable exception). This fact is somewhat surprising, given that the literature in other fields—e.g. voter turnout and financial inclusion—shows that the delivery mechanism matters (Green et al., 2013). Obviously, any cost-benefit analysis of a strategy aimed to reduce evasion will depend on the method used by the tax authority to approach taxpayers.

Moreover, the amount of information collected by governments on taxpayers grows every day, and can be used to encourage them to comply. Email addresses are valuable not only because they allow authorities to send reminders, but also because they enable the identification of citizens on social media outlets such as Facebook or Twitter when combined with other information, such as complete names, and telephone numbers, among others. Hence, targeted advertisement, like the one used for merchandising and electoral campaigns (Brookman and Green, 2014), can also be used to remind people of their tax obligations. Additionally, email messages and social media ads, coupled with reliable systems, are effective methods to reduce the transaction costs associated with paying taxes, as they can include links to online payment platforms where citizens may quickly fulfill their obligation. Consequently, in this paper, we want to understand if new delivery mechanisms, such as social media ads, complement the identified effects that more traditional online strategies, such as email reminders, have on tax compliance (Ortega and Scartascini, 2015).

We perform a field experiment in the Venezuelan municipality of Sucre, which forms part of the metropolitan area of Caracas, to determine if online strategies, such as email reminders and social media ads, are effective in increasing tax compliance. In collaboration with the tax authority of this municipality, we gather information of taxpayers that owe the vehicle tax. This obligation needs to be fulfilled annually but is recalculated every two months to adjust for the interest rate, and both individuals and firms are charged as a function of the vehicle(s) they possess. However, compliance rates are quite low and the tax authority tends not to take advantage of the detailed contact information of taxpayers that it has. Our intervention, that targets a total of 6183 debtors of this tax, is composed of four treatment arms, including the control group. In the case of the first treatment group, on September 13th 2017, we sent an email message to a total of 1546 randomly selected taxpayers, reminding them of their obligation. The communication included some personal information, like the amount owed and the characteristics of the vehicle, and a link that

redirected to an online payment platform. This was a deterrence message, in the sense that it emphasized on the legal consequences of evading.

As stated previously, email messages are useful not only to send direct reminders, but also because they can be used to target accounts on social media. Consequently, for our second treatment arm, we used a targeted campaign on Facebook, encouraging citizens to comply and pay their taxes before the deadline. The campaign lasted a total of 15 days, starting on September 13th as well, and we paid for it in order to maximize the probability of exposing each of the 1546 taxpayers in this treatment group to the ads. As opposed to email reminders, these ads were not personalized—it is not possible to do this on Facebook—but they were also deterrent as they reminded users about the payment deadline, and included the link to facilitate online payments. Given the differences between the two treatments, where the email message was more direct and personalized while the campaign was more general and lasted for a longer period of time, we expected the former to be more effective than the latter. But rather than considering both strategies as substitutes, the interesting question is whether they complement each other or not. For this purpose, we included a third treatment group, composed of 1546 subjects that received both an email reminder and exposure to the social media ads. Consequently, compared to the control group, it is interesting to determine if the effect of the Email & Facebook treatment is higher than the effect of solely the Email condition.

All of our treatments have positive and significant effects on subjects' behavior, both on the probability of making any payment, but also on the amount paid. Debtors receiving only an email reminder, for instance, were 7.7 percentage points more likely of making a payment and paid 58% more, compared to the control group. Facebook ads, as expected, had a lower—albeit significant—effect on both outcomes. Subjects in the 'Facebook ad' treatment group were 1 percentage point more likely to make any payment and paid 8% more than the control group. In accordance with this result, we find that Facebook ads slightly boosted the effect of email reminders. Subjects in the combined treatment group were 9.5 percentage points more likely of making any payment and paid 70% more than the control group. Even though these effects are small in magnitude, the low cost of social media ads—we spent around US $60 in total—make them cost-effective.

Additionally, we explore some of the mechanisms behind these effects. First, we find that email reminders are highly effective among natural persons, as opposed to legal persons like firms which are less sensitive to this type of communication strategy. This may be a consequence of the fact that the credibility of the deterrence messages is lower among firms.

However, and this is more surprising, we find that Facebook ads and their combination with email reminders are more effective on firms. It is important to consider that in the case of Sucre in 2017, legal persons are mostly small enterprises run by few people. Consequently, it sounds reasonable to assume that social media ads are visualized by the owners or someone taking important decisions within these firms. We hypothesize that, given that the ads were displayed on a daily basis during two weeks and just before the deadline, they served as reminders of the tax calendar to small business owners. This was more helpful for firms which tend to be busier and have more obligations than natural persons.

We also test for differential effects in terms of the size of the debt, the number of cars, the timing of payments and the methods used to pay. We find that the combined strategy of using email reminders and Facebook ads at the same time, is more effective to mobilize certain taxpayers who do not respond to the individual treatments. The data also shows that email reminders tend to have a sharp instantaneous effect, while the impact of ads takes a while to materialize. This is consistent with the idea that email messages scare and deter, as they are personalized, contain private information, and are more intrusive, while social media ads serve as a tax calendar. Concerning the payment method, our treatments have the largest effects on the probability of paying online, but still have positive and significant effects on payments at the bank or at the tax authority's office. Hence, we boost compliance not only because our strategies reduce transaction costs, but also because we are able to activate other behavioral aspects of decision-making. We can safely conclude that online strategies affect offline behavior as well.

We believe that our paper represents an important contribution to the literature on tax compliance, for at least five reasons. First, most of recent experimental research has been conducted in developed countries—especially the U.S., and little is known about the effects of this type of intervention in extremely weak institutional settings like Venezuela in 2017. Essentially, it is crucial to understand if these strategies are effective in developing countries, because fiscal capacity is lower in such contexts. Second, most of the experimental research has focused on the content of tax-compliance messages, while few studies are centered on the role of the delivery mechanism employed by the tax authority. Third, the few papers that compare the effects of different delivery mechanisms have put more emphasis on offline strategies—letters, visits, phone calls, etc.—than on online methods. Only email reminders have been included in these analyses, while the role of internet advertisements has been systematically ignored, despite the fact that private firms and governments increasingly rely on them on a daily basis.

Fourth, our experiment represents a novel example of how a government may use the information provided by e-government tools to strengthen itself. Virtual offices—i.e. tax-payment websites, like the one used by the municipality of Sucre in Venezuela, are cost-effective strategies that not only cut the amount of red-tape in administrative processes. They also provide valuable information to the authorities that may be used to improve the targeting of public communication. Finally, the experimental design and the nature of our sample frame allows us to explore interesting heterogeneous treatment effects, that have not been frequently analyzed in the literature. In particular, we are able to test if our interventions have differential effects across natural and legal persons. Moreover, to deal with the strong sample requirements implied by an analysis of heterogeneous treatment effects at such level—especially because in this case the sample of firms may not be big enough and some of the effects can be small—we rely on randomization inference techniques, something that has not been explored by this literature.

The paper is composed of six sections, including this introduction. In section 2, we discuss why online strategies may have effects on tax compliance. Section 3 describes the context in which the experiment was conducted. In section 4, we explain in detail the experimental design and the empirical strategy employed to determine the effect of our interventions on tax compliance. Section 5 reports the results of the experiment, both in terms of main and heterogeneous treatment effects. We conclude in Section 6.

2. Online Strategies and Tax Compliance

The recent wave of randomized controlled trials on tax enforcement have mainly focused on the effects that different types of behavioral appeals have on citizens' levels of compliance (for an overview see (Slemrod, 2018)). Many of these studies have been conducted in developed countries (Slemrod et al., 2001; Torgler, 2004; del Carpio, 2014; Hallsworth et al., 2017; Meiselman, 2018; Boning et al., 2018), but given the high salience of this problem in low-income nations, recent insightful research has been conducted in developing contexts as well, albeit at a lower scale (Ortega and Sanguinetti, 2013; Ortega et al., 2013; Castro and Scartascini, 2015; Gallego and Ortega, 2018). Most of these interventions seek to change taxpayers' beliefs about the enforcement environment in general, for instance in terms of the penalties associated to evasion and the probability of being caught. Social norms and moral sentiments, such as civic duty, equity, or reciprocity, have also guided recent behavioral interventions. To a lesser extent, some studies have focused on changing taxpayers' perceptions on the costs of complying (Gallego and Ortega, 2018; Meiselman,

2018). In most cases, the evidence suggests that deterrence messages are more effective than mentioning social norms or the costs of complying (Castro and Scartascini, 2015; Gallego and Ortega, 2018; Meiselman, 2018).

Most of these experiments have used letters as the vehicle to approach taxpayers. In contrast with other fields, the literature on tax compliance has not fully studied the role that different delivery methods play when encouraging subjects to behave in a certain way. Scholars in the so-called Get Out the Vote (GOTV) literature, have thoroughly compared alternative channels such as letters (Gerber and Green, 2000; Gerber et al., 2003), phone calls (Gerber and Green, 2001; Cardy, 2005), visits (Michelson, 2005; Alvarez et al., 2010), email (Stollwerk, 2006; Nickerson, 2007), text messages Allison and Strauss (2009); Malhotra et al. (2011), radio ads (Panagopoulos and Green, 2010), among others. It has been found that more personalized and intrusive strategies, such as visits and text messages, are more effective to mobilize voters (Green et al., 2013). However, the costs of each form of intervention should be taken into account in order to get a better sense of the cost-effectiveness of each method.

Ortega and Scartascini (2015) and Boning et al. (2018) are among the few studies that investigate the effects that different modes of intervention may have on tax compliance. In an experiment conducted in Colombia, Ortega and Scartascini (2015) compare the effectiveness of three delivery mechanisms: letter, email, and visits by a tax inspector.[1] All three modes generate positive effects on compliance, and there are sizable differences across methods, with visits having the largest effects. However more effective they are, visits are also more expensive inducing tax authorities to weigh up the costs and benefits of each strategy. It is important to note that this experiment represents the first instance where online methods such as email reminders are studied, and that positive and significant effects are found. Similarly, in their analysis of the direct and network effects of contacting firms in the U.S., Boning et al. (2018) compare the impact that letters and personal visits have on firm behavior. Again, visits are much more effective than letters, to the point that the former have network effects while the latter do not.

With the exception of the email messages analyzed by Ortega and Scartascini (2015), online strategies have not been studied at all. In fact, to the best of our knowledge, ours is the first study that tests the effect of social media ads on tax compliance. This is somewhat surprising for various reasons. First, as e-government strategies become more popular, authorities acquire more personal data and a vast amount of information that may be used

[1] In an accompanying paper, the authors introduce phone calls as an alternative delivery method.

to target taxpayers in order to encourage them to fulfill their obligations. This data is useful because email reminders may be designed to include personal information, which according to the literature should drive larger effects. But at the same time, private information may be used to individually target subjects through social media ads—although with some imprecision—something that is difficult or impossible with traditional media such as radio or TV ads.

Second, it is surprising that scholars have not explored the effects of online strategies on tax compliance, despite the fact that they have become an increasingly common form of mass communication and that its impacts on other dimensions of human behavior have been thoroughly analyzed. Most existing research has assessed the effects of internet advertisement on various forms of online behavior, such as online purchasing, click behavior, installation of software applications, and online donations to charities (Goldfarb and Tucker, 2011; Bakshy et al., 2012; Lewis et al., 2012; Ryan, 2012; Aral and Walker, 2011; Lacetera et al., 2016; Brookman and Green, 2014). Few studies assess the effects of online ads on offline behavior. Brookman and Green (2014), for instance, study the effects of online advertisement on political persuasion. The authors find that Facebook ads have no effects on candidate name recognition nor candidate favorability, casting serious doubts on the effectiveness of these strategies, at least in the realm of politics.

One may ask: why is it worth studying the effects of online strategies on tax compliance? If the goal of contacting taxpayers is to change their beliefs about the system, either in terms of sanctions, the probability of being caught, or the costs of complying, the strategy used to approach them matters. First, because some methods are more reliable than others to establish an effective contact with taxpayers. For example, compared to social media ads, emails may be a more direct method for contacting individuals. In modern days, they may even be more effective than letters or phone calls. Second, if the goal is to change beliefs, some channels are better than others because they allow for the creation of more personalized messages that could have more influence on beliefs. Emails may include personal information, such as the name of the debtor, the amount owed, the characteristics of the vehicle or property being charged, etc. It is not possible to include individual information in an internet advertisement. Third, if the goal of the approach is to create awareness about the payment deadlines, social media ads may be more effective because online campaigns can be continuously implemented during a predetermined time window. In other words, ads may serve as reminders of deadlines. In contrast, it may not be a good idea to annoy taxpayers and saturate them with high-frequency emails on the same subject. Finally, when online payments are available—as in the case of Sucre—both email reminders and social

media ads may be useful as they can redirect citizens to the platforms where payments may be completed.

In sum, most RCTs on tax compliance have focused on the content of messages, while very few have studied the role of the delivery mechanism. The role of online communication strategies has not been thoroughly analyzed, despite the fact that governments have access to vast amounts of information on taxpayers, which may be used to design and implement targeted campaigns. Additionally, these mechanisms tend to be cheap and cost-effective. Rather than seeing alternatives like emails and internet ads as substitutes, these communication devices can complement each other in a combined strategy aimed at increasing fiscal capacity.

3. Background

3.1. Venezuela's Political and Economic Turmoil

The experiment was conducted during the second semester of 2017, in moments of high turbulence in Venezuela. Despite having the largest oil reserves in the world, both from an economic and a political perspective, this year was one of the most challenging in the history of the country. For its fourth consecutive year, there was a decrease in its GDP, with a 31.9% contraction with respect to 2013 (CEPAL, 2017). Due to the fact that the central government spent more than what it was able to raise, hyperinflation reached historic peaks, as the monetary base grew at levels above a 400% rate. Prices kept rising at unprecedented levels, 2017 being the fourth year in a row with a three-digit inflation rate, and the second one with a rate above the 300% level. In terms of currency, the Bolivar experienced a depreciation of 396% with respect to the U.S. dollar which is commonly used as a reference for the official exchange rate, but such depreciation topped more than 1600% according to parallel market rates.

With such unprecedented levels of hyperinflation, economic stagnation, and unemployment, the period in which the experiment took place was characterized by high levels of social unrest and street protests. Additionally, this humanitarian crisis led millions of Venezuelans to migrate to neighboring countries, especially to Colombia, Brazil, Peru, and Ecuador. The UN's International Organization for Migration calculated that by the end of 2017, about 1,6 millions of Venezuelans were living outside of the country. Face with this conjunction of issues, Nicolas Maduro's regime experienced political instability, constant critiques from

the opposition, and low levels of popular support. In fact, in a desperate move, the regime installed a controversial National Constituent Assembly, in theory to draft a new constitution, but in practice to weaken the National Assembly which was controlled by opposition parties. Naturally, such context represents a burden for national and local governments that need credibility and legitimacy in order to increase fiscal revenues, even in the case of small taxes, such as the vehicle tax. Not surprisingly, general levels of tax compliance in the country were quite low in 2017.

3.2. The Municipality of Sucre and its Tax Authority

Sucre is a Venezuelan municipality of 700,000 inhabitants in the state of Miranda and is one of the five municipalities that composes the metropolitan area of Caracas. From 2008 to 2017, and consequently at the time of the experiment, Sucre was governed by opposition non-Chavista parties. In 2018, however, the Partido Socialista Unido de Venezuela, came to the mayor's office, after a Chavista hardliner, Jose Vicente Rangel, won the local election held on December 2017. The Decentralized Service of Tax Collection (SEDAT, by its acronym in Spanish) is the mayoral office in charge of raising different types of tributes in Sucre. These payments represent the main sources of revenue for the municipality, and include industry and commerce taxes, property taxes, and vehicle taxes. In this study, we focus on the vehicle tax,[2] which is updated every two months and has to be paid by both individuals and firms. Baseline levels of compliance indicate a surprisingly high 70% rate of evasion. In cooperation with the authorities of Sucre, and with the support of CAF-Development Bank of Latin America, we had access to the database of taxpayers, which included information on debts, compliance, vehicle characteristics, and other personal data like home and email addresses, phone numbers, etc.

In Sucre, every person possessing a vehicle has to register it at SEDAT's virtual office. Consequently, the tax authority digitally collects important pieces of information on taxpayers, that may be exploited to increase compliance. The authority calculates the tax rate for every vehicle, that depends on the year of the car, its type, and the usage it has. Taxpayers have three methods for complying with their obligation: 1) Pay online with a credit card at SEDAT's virtual office; 2) Print the bill from the virtual office and pay directly at any of the five authorized banks; 3) Pay directly at SEDAT's office using a debit or a credit card. This is an interesting setup to study how the use of individual-level private information can enhance fiscal state capacity through inexpensive mechanisms such as the

[2] In an accompanying paper we focus on property taxes as well (Gallego and Ortega, 2018).

online strategies used in this experiment—especially given the hard economic times and the political instability afflicting the regime.

4. Experimental Design

4.1. The Treatments

As part of a project carried out with the Directorate of Impact Evaluation and Policy Learning at CAF-Development Bank of Latin America, and in alliance with SEDAT, we aim to evaluate the effect of online strategies to approach debtors of the vehicle tax and encourage them to fulfill their obligation. For this purpose, we designed and implemented a randomized control trial at the taxpayer level in Sucre, Venezuela. Interventions take place on behalf of the municipality of Sucre and its tax authority. The experiment started on September 13 of 2017, interventions ended on the 27th, and we analyze taxpayer behavior until October 31. Figure A.1 in the Online Appendix presents the timeline of the experiment. A pre-analysis plan of this study, including hypotheses and estimation strategies, was registered at the Evidence on Governance and Politics repository (see https://egap.org).

The experiment is composed of three treatment arms and a control group, for a total 6183 unique taxpayers that are randomly assigned to each group. Note that the analysis is conducted at the taxpayer level, but each debtor may have more than one car, and consequently, more than one account at SEDAT's virtual office. We unify all the accounts of a same person into a single unit of analysis, to avoid having taxpayers receiving multiple treatments. In the first treatment arm, subjects received a deterrence email message stressing the fact that they better comply with their obligation in order to avoid legal sanctions.[3] The message included personal information of the taxpayer, like her name, the size of the debt, some characteristics of the vehicle, and the account number. It also included some general information, like the payment deadline, a link to FAQs, and a link that redirected subjects to a platform for online payments. Logos of SEDAT and the Municipality of Sucre were included as headers, to enhance the credibility of the message. Figure 1 shows an example of an email reminder sent to a hypothetical taxpayer.

[3] We use a deterrence message, as opposed to other alternatives appealing to other forms of moral suasion, given the results of previous related experiments (Gallego and Ortega, 2018; Meiselman, 2018).

Our second treatment arm is a Facebook campaign. Using taxpayers' information, we paid for targeted advertisement using Facebook's "create an audience" tool.[4] In this way, we created a 20-day engagement campaign and we optimized the number of unique daily impressions, in order to reach as many of our targeted users as possible.[5] Introducing personal information of users allows us to increase the probability of reaching our targets. Consequently, when configuring the campaign, we introduced data such as the names, email addresses, and phone number of subjects assigned to this treatment group. Naturally, 100% compliance is not guaranteed through this procedure—so we will estimate Intention to Treat Effects (ITT), but according to Facebook numbers, we were able to reach a relatively high number of users—a total of 2132 unique subjects, while our goal was to approach 3092 taxpayers.[6] We discuss this issue in more detail in Section 5. The promoted Facebook post is less personalized than the email message and includes the following information: deadline and invitation to pay the vehicle tax, link to SEDAT's virtual office, and a deterrence image stressing the importance of compliance. Logos of SEDAT and the municipality were included as well. Figure 2 shows the nature of the Facebook post. These posts are less personalized than email reminders, as Facebook advertisement tools allow to target a set of individuals, but do not allow to individualize the messages sent to each one of them.

Finally, the third treatment arm is composed of subjects receiving both the email message and exposure to the Facebook campaign. We included this third group to test whether social media campaigns boost the effects of email reminders, assuming that the latter are more effective than the former. Admittedly, sending email reminders should be the default for tax authorities as it is a cost-effective mean of communication on tax compliance. However, what guides our interest here is whether on top of sending emails it is worth paying for targeted social media ads. Naturally, the experiment includes a fourth group that is not approached in any way and that serves as the control group. Hence, members of the control group represent the status quo in Sucre, which implies receiving no form of communication from the municipality. Block randomization was conducted at the taxpayer level, as we stratified on whether the subject was a natural or a legal person. To optimize balance on observables, as in Castro and Scartascini (2015) we randomized 100 times and selected the trial that maximized balance among covariates (i.e. minimized the maximum t-statistic of coefficients in a regression of each covariate on treatment statuses).

[4] Given the configuration of the advertisement campaign, occasionally some users might have been reached on Instagram as well, a social media platform that is integrated in many ways with Facebook.

[5] Facebook's *optimization for ad delivery* tool offers three optimization options: 1) number of clicks; 2) number of impressions; and 3) number of unique daily impressions. We chose the third one because our goal was to treat as many of our target users as possible.

[6] Although see Eckles et al. (2018) for some limitations of this technique.

Figure 1: Example of an Email Reminder

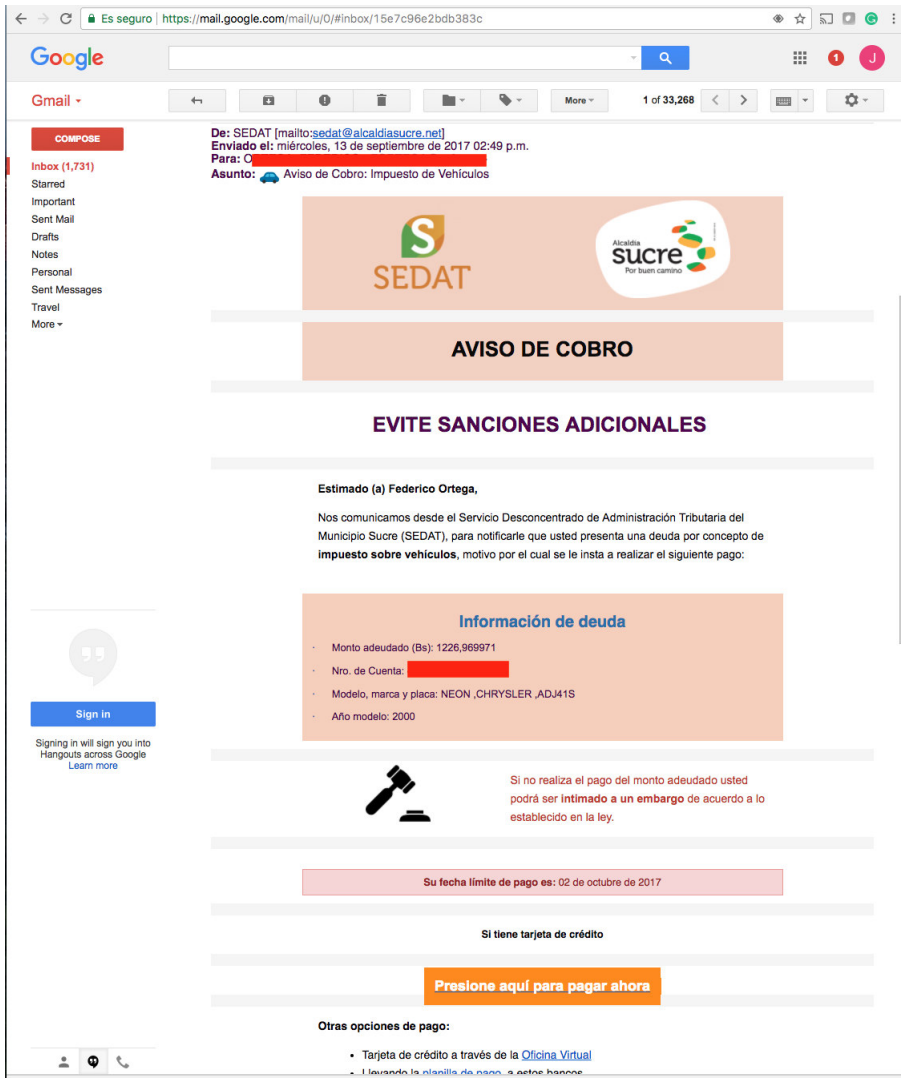

Figure 2: Example of a Facebook Ad

4.2. Empirical Strategy

We want to determine if online strategies are effective mechanisms for increasing tax compliance. As such, one of our goals is to determine if taxpayers receiving an email reminder or being targeted by a Facebook ad are more likely to pay than citizens in the control group. Our first hypothesis is that both methods are effective: the probability that subjects in the Email or Facebook groups pay their debt is greater than the probability to pay for the control group. However, if we were to compare these two groups, we hypothesize that email reminders are more effective than Facebook ads. The reason is that email messages are personalized and contain private information, regarding the amount owed, the type of car, etc. Studies in other areas, such as the Get Out the Vote (GOTV) experiments referenced above (Gerber and Green, 2015; Green et al., 2013), reveal that more personalized messages are more effective compared to general notifications. We think that in the case of tax compliance the same reasoning holds. Therefore, we hypothesize that *the probability that subjects in the Email treatment group pay their debt should be higher than the probability to pay in the Facebook treatment group.*

However, these two strategies do not have to be considered as strategic substitutes. Ultimately, both are cost-effective strategies that can be employed when the tax authority has enough information on taxpayers. On the one hand, sending messages can be done for free, discounting the time needed to send them, and this process can be automatized using specialized software. On the other hand, Facebook ads are relatively cheap compared to other advertisement strategies. In this case, we spent less than US $60 in order to target about 3000 subjects during 20 days. In the context of Venezuela in 2017, given the difficult problems of hyperinflation and exchange rates, this sum may be considered high. But for most other countries, it may be considered a cheap intervention. Consequently, one of the main goals of the paper is to test if, conditional on receiving an email reminder, the effect is large on a person that is exposed to social media ads as well. Email messages may deter people from evading, while Facebook ads, in addition, may serve as reminders of the tax schedule. If this is the case, we hypothesize that *the probability that subjects in the Email & Facebook treatment group pay their tax should be higher than the probability to pay for the Email treatment group or the Facebook treatment group.*

Consequently, we estimate linear probability models of the form:[7]

[7] To account for the dichotomous nature of the payment dummy, in the Online Appendix we report the results of analogous logistic models.

$$Payment_i = \alpha + Email_i\beta_1 + FB_i\beta_2 + (Email\&FB)_i\beta_3 + Natural_i + \varepsilon_i \quad (1)$$

where $Payment_i$ is a dummy variable indicating if subject i pays her taxes or not and $Email_i$, FB_i, and $Email\&FB_i$ indicate the treatment group that she belongs to. Also, $Natural_i$ is a dummy indicating if the observation corresponds to a natural or a legal person, and is included to account for the blocked randomization procedure that was used. In some specifications, we also include a battery of taxpayer-level covariates, that include the baseline debt, the number of accounts possessed by the subject, and certain characteristics of the vehicle. In every case, we estimate heteroskedastic robust standard errors. Section A in the Online Appendix reports descriptive statistics of the variables used in this study.

Alternatively, we estimate models in which the dependent variable corresponds to the natural log of the amount paid by the taxpayer,[8] in order to express the effects of our interventions in percentage terms. Moreover, in order to study some of the channels at play, we estimate these models using other dependent variables, such as payment method chosen by taxpayers that fulfill their obligation (online, at the bank, or at the SEDAT office), or the timing of such payment. These tests are important to determine, for instance, if online strategies may also affect offline behavior, like paying at the bank. Also, if it is true that email messages scare subjects, while Facebook ads serve as reminders of the tax schedule, the timing of the payments should differ across groups.

Naturally, we also want to test for any heterogeneous effects at the type-of-person level. We hypothesize that the effects of email reminders, Facebook ads, or both, may differ depending on whether the taxpayer is a natural or a legal person. For instance, on the one hand, it may be the case that the threat contained in the email may be more credible for natural persons, compared to firms. Firms interact more frequently with tax authorities, and consequently their beliefs may be very different—they may know better the real sanctioning capacity of SEDAT. On the other hand, if social media ads serve as reminders of the tax calendar, this information may be more valuable for firms, which may be more prone to ignore deadlines as they are busier and tend to have more obligations. To test for potential heterogeneous treatment effects on this dimension, we estimate models of the form:

$$\begin{aligned} Payment_i = {} & \alpha + Email_i\beta_1 + FB_i\beta_2 + (Email\&FB)_i\beta_3 + Natural_i \\ & + (Email_i \times Natural_i)\delta_1 + (FB_i \times Natural_i)\delta_2 + (Email\&FB_i \times Natural_i)\delta_3 + \varepsilon_i \end{aligned}$$

[8] We use $\log(x+1)$ transformations, given that in most cases the amount paid is 0.

where the coefficients of interest are δ_1, δ_2, and δ_3, as they allow us to compare the effect of our interventions across natural and legal persons. Positive values of any of these coefficients implies that the respective treatment effect is larger for natural persons, and vice versa. One possible drawback of this analysis is that, at the subgroup level, it may be underpowered due to sample size considerations, especially in the case of firms. For this purpose, our analysis of heterogeneous treatment effects relies on *randomization inference* p-values, which are the result of simulations in which each treatment condition is reassigned (Young, 2017). Randomization inference is recommended in situations where the sample size is not big enough or when the outcome variable does not follow a normal distribution (Gerber and Green, 2012; Imbens and Rubin, 2015). In our case, the problem is that the outcome variable is highly skewed—the majority of debtors keep evading, in such a way that assuming normality seems incorrect. Hence, conventional p-values may be misleading when drawing inferences for subgroups like firms, given that in such cases the samples sizes are much smaller.

5. Results

We now present the main results of the experiment. First, we show that the treatment and control groups are comparable, in the sense that balance on observables is achieved. This condition is true for both the pooled dataset, as well as for each stratum that we used during the randomization stage. Second, we report the main treatment effects of the experiment. We estimate both Intention to Treat (ITT) effects as well as Compliers Average Causal Effects (CACE) (Gerber and Green, 2012), that reveal the impact of our Email, Facebook, and Email & Facebook treatment conditions on the probability of debt payment and on the amount paid. Next, we test for heterogeneous effects of our treatments, disaggregating between natural and legal persons. Fourth, we analyze whether the effects of our treatments are mediated by important covariates, such as the size of the debt and the number of accounts (vehicles) of taxpayers. Finally, we shed some light on the mechanisms at play, by analyzing the effects of treatment conditions on both the timing and the methods of payment.

5.1. Covariate Balance

Table 1 reports balance statistics for the group of covariates[9] on which baseline information was available: pre-treatment debt, number of accounts of the taxpayer,[10] type of person (natural or legal), and six dummies identifying the type of vehicle.[11] For each covariate, we ran a regression where the covariate acts as the outcome variable, against treatment indicators (Email, Facebook, and Email & Facebook) that act as regressors. Taxpayers in the control group are used as the benchmark for comparison. First, this is done for the pooled sample of observations that includes both natural and legal persons. Balance is satisfactory, as the treatment status is generally not correlated with the covariates of interest. Only two exceptions are found amongst the 27 correlations that we test for using this method: in the case of type 1 and type 2 vehicles which are correlated with the Facebook treatment condition. One correlation is significant at the 10% level which we deem within the acceptable standards and is most likely the result of chance. Note that our sample is composed of 6183 taxpayers, which correspond to 5518 and 665 natural and legal persons, respectively.

Table 2 reports the results of an equivalent exercise, but this time disaggregating the sample between natural and legal persons. It is important to underscore that balance is achieved at these levels as well, as we test for the existence of heterogeneous treatment effects across the type of person. Again, the results show that in general balance is achieved between treatment and control groups. Panel A in Table 2 shows that the imbalance found in two variables in the pooled case is driven by natural persons, as the same significant coefficients are found in such case. In the case of legal persons, as shown in Panel B, only one significant correlation is found, for type 3 vehicles and the Email & Facebook treatment. From this analysis, we conclude that balance is satisfactory. Nonetheless, we include specifications both without and with covariates for the main effects that we estimate in the next subsection.

[9] This is the full set of covariates shared by SEDAT.

[10] A single taxpayer may have multiple accounts if she owns more than one vehicle. As we said before, we merged the original dataset so that the unit of analysis is the individual taxpayer and not the tax account.

[11] Cars, vans, motorcycles, transport vehicles, cargo vehicles, and others.

Table 1: Covariate Balance: Pooled Data

	(1) Baseline Debt	(2) Number of Accounts	(3) Natural Person	(4) Type 1 Vehicle	(5) Type 2 Vehicle	(6) Type 3 Vehicle	(7) Type 4 Vehicle	(8) Type 5 Vehicle	(9) Type 6 Vehicle
Email	0.0256	0.0249	-0.000577	0.00412	0.00438	-0.00844	0.00257	9.63e-17	-0.00263
	(0.0262)	(0.124)	(0.0111)	(0.0173)	(0.0149)	(0.00794)	(0.00559)	(0.000647)	(0.00875)
Facebook	0.00721	0.0411	0.0000695	0.0391**	-0.0247*	-0.00456	-0.00131	0.000647	-0.00910
	(0.0262)	(0.124)	(0.0111)	(0.0173)	(0.0149)	(0.00794)	(0.00559)	(0.000647)	(0.00875)
Email&Facebook	-0.0211	0.0320	0.0000695	0.00930	-0.00209	-0.00198	0.00193	0.000647	-0.00781
	(0.0262)	(0.124)	(0.0111)	(0.0173)	(0.0149)	(0.00794)	(0.00559)	(0.000647)	(0.00875)
Constant	7.545***	1.507***	0.893***	0.626***	0.227***	0.0550***	0.0239***	-6.85e-17	0.0680***
	(0.0186)	(0.0875)	(0.00788)	(0.0122)	(0.0106)	(0.00561)	(0.00395)	(0.000458)	(0.00619)
N	6183	6183	6183	6183	6183	6183	6183	6183	6183

Notes: Standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office. *Natural* is an indicator of natural versus legal person. *Type* 1-*Type* 6 are indicators for the type of vehicle. OLS estimation used in every specification. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

Table 2: Covariate Balance: Type of Debtor

	(1) Baseline Debt	(2) Number of Accounts	(3) Type 1 Vehicle	(4) Type 2 Vehicle	(5) Type 3 Vehicle	(6) Type 4 Vehicle	(7) Type 5 Vehicle	(8) Type 6 Vehicle
Panel A: Natural Persons								
Email	0.0219	-0.00508	0.0131	-0.00290	-0.00580	0.00145	1.91e-17	-0.00580
	(0.0232)	(0.0364)	(0.0178)	(0.0157)	(0.00826)	(0.00617)	(0.000513)	(0.00658)
Facebook	-0.00143	-0.000221	0.0466***	-0.0349**	-0.00366	-0.000744	2.00e-17	-0.00727
	(0.0231)	(0.0364)	(0.0178)	(0.0157)	(0.00826)	(0.00617)	(0.000513)	(0.00658)
Email&Facebook	-0.0244	-0.0140	0.0104	-0.00959	0.00721	0.00288	0.000725	-0.0116*
	(0.0231)	(0.0364)	(0.0178)	(0.0157)	(0.00826)	(0.00617)	(0.000513)	(0.00658)
Constant	7.470***	1.305***	0.658***	0.229***	0.0500***	0.0261***	-1.68e-17	0.0370***
	(0.0164)	(0.0257)	(0.0126)	(0.0111)	(0.00584)	(0.00437)	(0.000363)	(0.00465)
N	5518	5518	5518	5518	5518	5518	5518	5518
Panel B: Legal Persons								
Email	0.0526	0.262	-0.0680	0.0646	-0.0305	0.0119	-3.89e-17	0.0220
	(0.130)	(1.089)	(0.0519)	(0.0480)	(0.0271)	(0.00846)	(0.00425)	(0.0517)
Facebook	0.0794	0.386	-0.0241	0.0602	-0.0120	-0.00602	0.00602	-0.0241
	(0.130)	(1.090)	(0.0520)	(0.0480)	(0.0272)	(0.00847)	(0.00426)	(0.0518)
Email&Facebook	0.00628	0.416	-5.64e-16	0.0602	-0.0783***	-0.00602	-4.04e-17	0.0241
	(0.130)	(1.090)	(0.0520)	(0.0480)	(0.0272)	(0.00847)	(0.00426)	(0.0518)
Constant	8.172***	3.187***	0.361***	0.211***	0.0964***	0.00602	3.38e-17	0.325***
	(0.0919)	(0.771)	(0.0368)	(0.0340)	(0.0192)	(0.00599)	(0.00301)	(0.0366)
N	665	665	665	665	665	665	665	665

Notes: Standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office.*Type* 1-*Type* 6 are indicators for the type of vehicle. OLS estimation used in every specification. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

5.2. Email Reminders, Facebook Ads, and Tax Compliance

In this section, the main treatment effects of our intervention are reported. First, we graph the response rate of each treatment group. Figure 3 plots the average compliance rate of taxpayers, across treatment groups. It is clear that the interventions including an email reminder were the most effective and that the effect was higher when this strategy was complemented by Facebook ads. As expected, the pure effect of these online advertisements is much lower. Nonetheless, the compliance rate is higher when only Facebook ads are used compared to the control group. Naturally, it is important to determine if all these differences are statistically significant.

Table 3 corroborates these intuitions through a series of linear regressions in which a payment dummy is regressed on the treatment indicators. Columns 1-4 report the coefficients of interest, that account for Intention to Treat (ITT) estimators, as it is the case that not all subjects assigned to the Email conditions read the messages, nor all subjects assigned to a Facebook condition see the ads. The results reveal that all of our treatment conditions have positive and significant effects on the probability of paying the debt and that this result is robust to the inclusion of covariates. In the case of subjects receiving an email reminder, the payment probability is 7.7 percentage points higher, compared to the control group. Section C in the Online Appendix shows that equivalent results are found if we use Probit models instead of linear probability estimations. Additionally, section D shows that the significance of these effects remains unchanged even if we conduct corrections for multiple comparisons, using the Benjamini and Hochberg (1995) procedure.

The effect on tax compliance is even higher in the case of subjects receiving email reminders and Facebook ads. Compared to the control group, subjects in this treatment condition are on average 9.5 percentage points more likely to pay. An F test of difference between coefficients reveals that this 1.8 percentage difference between the two effects (Email & Facebook - Email) is marginally insignificant (p-value of 0.1015). However, we do find a positive and significant effect for the condition in which subjects are solely exposed to Facebook ads. Compared to the control group, subjects in this group are on average 1.1 percentage points more likely to make some payment.

In the case of email messages, we do have information on which subjects received and opened their reminders. Hence, at least for subjects in the Email and Email & Facebook treatments, we can estimate the Causal Average Compliers Effect (CACE) (Gerber and Green, 2012) of the assignation to any treatment including an email reminder. Unfortunately, Facebook

Figure 3: Compliance Rate by Treatment Status

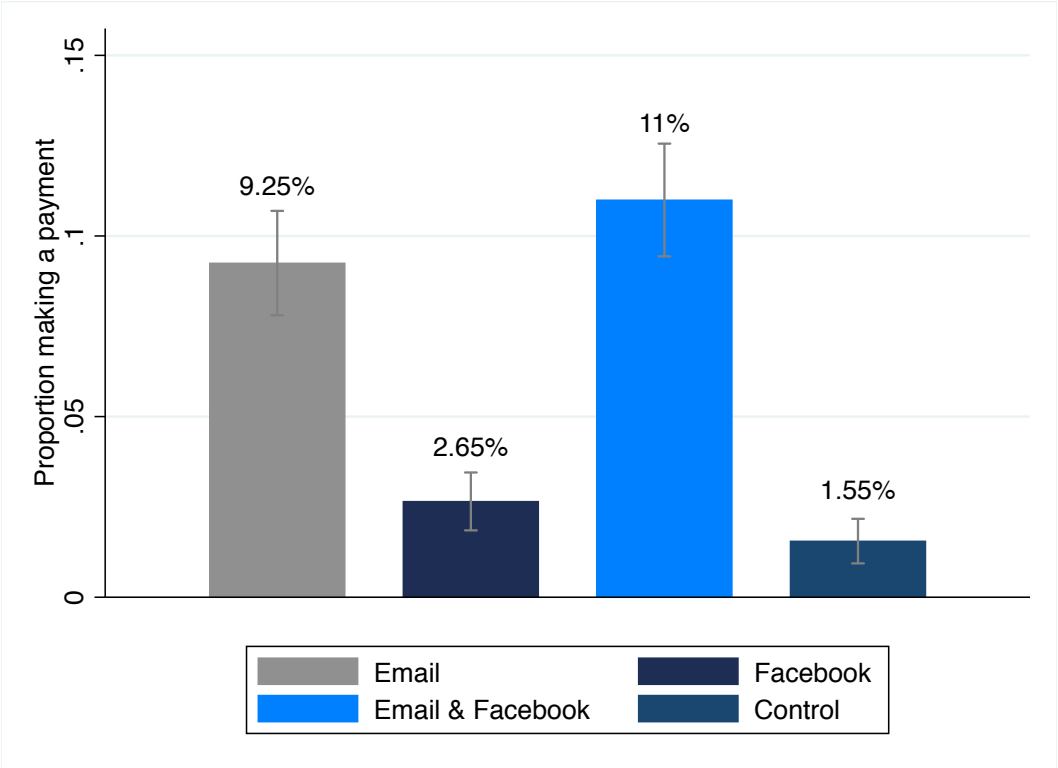

does not provide information on which subjects were truly exposed to the ads, so for treatments involving these ads, the best we can do is to calculate ITT effects. However, non-compliance to the Facebook-ad exposure would downwardly bias our estimates of both the Facebook and the Email & Facebook treatment effects: subjects in these groups incorrectly not exposed to the ads would probably have lower payment probabilities, while subjects in the Email and control conditions incorrectly exposed to ads would be more likely to pay. Therefore, our estimates of both effects are probably lower bounds of the real effects.

Columns 5-8 in Table 3 report the results of 2SLS estimations in which the allocation to the Email and Email & Facebook conditions are used as instruments of whether the subjects in these conditions read the reminders. The effects are quite substantial in this case. Subjects complying with the Email treatment are 20 percentage points more likely of paying some or all of their debt. In the case of subjects assigned to the Email & Facebook treatment, the probability is 26 percentage points higher. More importantly, an F test of difference between coefficients shows that this difference in the effects of the Email & Facebook and the Email conditions is significantly different from 0 at a 5% significance level (p-val=0.046). Section E in the Online Appendix reports the results of the first stages of these models. It is safe to conclude that email messages represent a very effective tool for increasing tax compliance of debtors. In the case of Facebook ads, the effect is smaller, as expected. Email messages are more intrusive, personalized and include private information, while Facebook ads are impersonal and just provide general information. Moreover, it is unclear how much compliance is achieved in terms of exposure to the ads—our estimates for such treatment are ITT. However, given that the behavioral mechanisms underlying these strategies can be quite different, Facebook ads may be more effective for certain types of taxpayers.

In Table 4 we report the results of equivalent models, in which we use the natural logarithm of the amount paid as the outcome variable. Hence, coefficients should be interpreted as the percentage difference in the amount paid when comparing each treatment condition and the control group. The first row of the Table shows that for any model, with or without covariates, and either for ITT or CACE estimations, receiving an email or both receiving an email and being exposed to Facebook ads has positive and significant effects on the amount paid to Sucre's tax authority. Columns 1 and 2 reveal, for instance, that subjects in the Email group pay on average 58% more than taxpayers in the control group. The effect is higher—153% more—among compliers. The effect of the Facebook treatment, once more, is weaker but positive and significant. Subjects in this group pay 8% more than people in the control condition. Finally, the largest effects are driven by the Email & Facebook treatment, with coefficients of 70% and 191%, respectively.

Table 3: Treatment Effects: Payment Indicator

	(1) Payment Indicator	(2) Payment Indicator	(3) Payment Indicator	(4) Payment Indicator	(5) Payment Indicator	(6) Payment Indicator	(7) Payment Indicator	(8) Payment Indicator
Email	0.0769***	0.0769***	0.0769***	0.0767***	0.204***	0.204***	0.204***	0.203***
	(0.00799)	(0.00799)	(0.00800)	(0.00800)	(0.0204)	(0.0204)	(0.0204)	(0.0205)
Facebook	0.0110**	0.0110**	0.0109**	0.0109**	0.0110**	0.0110**	0.0109**	0.0110**
	(0.00509)	(0.00509)	(0.00508)	(0.00508)	(0.00509)	(0.00509)	(0.00507)	(0.00508)
Email&Facebook	0.0944***	0.0945***	0.0944***	0.0944***	0.259***	0.259***	0.259***	0.259***
	(0.00852)	(0.00851)	(0.00851)	(0.00851)	(0.0224)	(0.0224)	(0.0224)	(0.0224)
Natural	-0.0679***	-0.0665***	-0.0650***	-0.0657***	-0.0696***	-0.0698***	-0.0680***	-0.0678***
	(0.0130)	(0.0135)	(0.0135)	(0.0139)	(0.0125)	(0.0128)	(0.0129)	(0.0132)
Constant	0.0762***	0.0607	0.0783*	0.0718	0.0777***	0.0793**	0.0999**	0.0894***
	(0.0127)	(0.0435)	(0.0468)	(0.0548)	(0.0123)	(0.0398)	(0.0429)	(0.0167)
Baseline Debt	N	Y	Y	Y	N	Y	Y	Y
No. Accounts	N	N	Y	Y	N	N	Y	Y
Type of Vehicle	N	N	N	Y	N	N	N	Y
Estimator	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS	2SLS
N	6183	6183	6183	6183	6183	6183	6183	6183

Notes: Heteroskedastic robust standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Natural* is an indicator of natural versus legal person. *Payment Indicator* is a dummy indicating whether the taxpayer paid her debt or not. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office. *Type* 1-*Type* 6 are indicators for the type of vehicle. OLS estimation used in columns 1-4. 2SLS estimation used in columns 5-8. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

Table 4: Treatment Effects: Amount Payed

	(1) Log Amount Paid	(2) Log Amount Paid	(3) Log Amount Paid	(4) Log Amount Paid	(5) Log Amount Paid	(6) Log Amount Paid	(7) Log Amount Paid	(8) Log Amount Paid
Email	0.585***	0.582***	0.583***	0.581***	1.548***	1.542***	1.543***	1.540***
	(0.0618)	(0.0616)	(0.0617)	(0.0617)	(0.158)	(0.158)	(0.158)	(0.158)
Facebook	0.0854**	0.0847**	0.0843**	0.0839**	0.0854**	0.0848**	0.0843**	0.0849**
	(0.0403)	(0.0403)	(0.0401)	(0.0401)	(0.0402)	(0.0402)	(0.0401)	(0.0401)
Email&Facebook	0.706***	0.708***	0.708***	0.708***	1.939***	1.945***	1.942***	1.943***
	(0.0652)	(0.0652)	(0.0652)	(0.0652)	(0.172)	(0.172)	(0.172)	(0.172)
Natural	-0.598***	-0.522***	-0.509***	-0.513***	-0.611***	-0.546***	-0.532***	-0.517***
	(0.107)	(0.106)	(0.107)	(0.110)	(0.104)	(0.102)	(0.103)	(0.106)
Constant	0.657***	-0.194	-0.0504	-0.112	0.668***	-0.0535	0.112	0.271
	(0.105)	(0.308)	(0.324)	(0.390)	(0.102)	(0.294)	(0.312)	(0.371)
Baseline Debt	N	Y	Y	Y	N	Y	Y	Y
No. Accounts	N	N	Y	Y	N	N	Y	Y
Type of Vehicle	N	N	N	Y	N	N	N	Y
Estimator	OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS	2SLS
N	6183	6183	6183	6183	6183	6183	6183	6183

Notes: Heteroskedastic robust standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Natural* is an indicator of natural versus legal person. *Log Amount Paid* is the natural log of the amount paid plus one. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office. *Type* 1-*Type* 6 are indicators for the type of vehicle. OLS estimation used in columns 1-4. 2SLS estimation used in columns 5-8. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

5.3. Heterogeneous Effects: Type of Taxpayer

We now test for heterogeneous treatment effects across natural and legal persons. Our sample is composed of 5518 natural persons and 665 legal persons. Given that we may be underpowered to test for heterogeneous effects across these two groups, in particular in the case of firms, we rely on randomization inference p-values (Young, 2017). This is done in order to determine if any differential impacts are significantly different from 0. For this purpose, we estimate our basic models for the payment probability and the natural log of the amount paid, but including the interactions of each of our treatment indicators with the natural person dummy.

We hypothesize that email reminders, which theoretically should deter evasion more than Facebook ads, are more effective on individuals. Our claim is that for firms, that interact more frequently with the tax authority, the threat of punishment for not paying the vehicle tax should be less credible than for ordinary citizens. We also postulate that the mechanism that drives the (small) effects of Facebook advertisement is the reminder effect they have on firms concerning the tax calendar, especially the payment deadline. Firms tend to be busier and have more obligations, leading them to wait until the deadline to proceed to payment. Then, we hypothesize that this strategy is more effective on firms. In Sucre most firms are small, so it is reasonable to assume that social media targeting at this level is feasible and effective—ads probably reach proprietors, CEOs, or accountants. Indeed, for each small business there is a person that can be targeted online.

Table 5 confirms our first hypothesis. The analysis of significance is based on the randomization inference p-values included inside the brackets, for the reasons described above. In every model, both for the payment probability and the amount paid, the coefficient of the interaction of the Email condition and the natural person dummy is positive and significant. The effect on the payment probability of getting an email reminder is more than 3 percentage points higher for natural persons, compared to firms. In the case Facebook ads, the result is less clear. No significant differences are found for the payment probability, but negative and significant coefficients exist for the amount paid. This result suggest that ads are more effective (at least in terms of the amount paid) on firms. Moreover, in the case of the Email & Facebook intervention, in every specification, the coefficient of the corresponding interaction is negative and significant. It seems fair to conclude that the combination of both strategies is more effective on firms' tax compliance rather than on citizens' tax compliance.

Table 5: Heterogeneous Effects: Natural vs. Legal Persons

	(1) Payment Indicator	(2) Payment Indicator	(3) Payment Indicator	(4) Payment Indicator	(5) Log Amount Paid	(6) Log Amount Paid	(7) Log Amount Paid	(8) Log Amount Paid
Email	0.0474	0.0473	0.0470	0.0463	0.422	0.417	0.414	0.410
	(0.0331)	(0.0331)	(0.0331)	(0.0331)	(0.277)	(0.276)	(0.277)	(0.277)
Facebook	0.0181	0.0179	0.0175	0.0172	0.180	0.171	0.168	0.166
	(0.0310)	(0.0310)	(0.0308)	(0.0308)	(0.259)	(0.259)	(0.257)	(0.258)
Email&Facebook	0.108***	0.108***	0.108***	0.105***	0.860***	0.859***	0.854***	0.835***
	(0.0368)	(0.0368)	(0.0369)	(0.0369)	(0.301)	(0.301)	(0.302)	(0.302)
Natural	-0.0703***	-0.0690***	-0.0679***	-0.0694***	-0.575***	-0.502***	-0.493***	-0.502***
	(0.0210)	(0.0214)	(0.0215)	(0.0219)	(0.172)	(0.173)	(0.174)	(0.176)
Email×Natural	0.0331***	0.0331***	0.0335***	0.0341***	0.182***	0.185***	0.188***	0.192***
	(0.0340)	(0.0340)	(0.0341)	(0.0341)	(0.284)	(0.283)	(0.283)	(0.284)
	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0020]	[0.0020]	[0.0000]	[0.0000]
Facebook×Natural	-0.00793	-0.00778	-0.00737	-0.00708	-0.106*	-0.0972*	-0.0938*	-0.0922
	(0.0313)	(0.0313)	(0.0311)	(0.0311)	(0.261)	(0.261)	(0.259)	(0.260)
	[0.2650]	[0.2630]	[0.3010]	[0.3250]	[0.0510]	[0.0760]	[0.0980]	[0.1130]
Email&FB×Natural	-0.0157**	-0.0156**	-0.0150**	-0.0121*	-0.172***	-0.169***	-0.164***	-0.142**
	(0.0377)	(0.0377)	(0.0379)	(0.0379)	(0.308)	(0.308)	(0.309)	(0.309)
	[0.0300]	[0.0260]	[0.0340]	[0.0830]	[0.0020]	[0.0020]	[0.0020]	[0.0100]
Constant	0.0783***	0.0631	0.0809	0.0759	0.636***	-0.117	-0.117	-0.117
	(0.0209)	(0.0466)	(0.0495)	(0.0570)	(0.171)	(0.417)	(0.417)	(0.417)
Baseline Debt	N	Y	Y	Y	N	Y	Y	Y
No. Accounts	N	N	Y	Y	N	N	Y	Y
Type of Vehicle	N	N	N	Y	N	N	N	Y
N	6183	6183	6183	6183	6183	6183	6183	6183

Notes: Heteroskedastic robust standard errors are shown in parentheses. Randomization inference p-values are shown in brackets for heterogeneous treatment effects. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Natural* is an indicator of natural versus legal person. *Payment Indicator* is a dummy indicating whether the taxpayer paid her debt or not. *Log Amount Paid* is the natural log of the amount paid plus one. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office. *Type* 1-*Type* 6 are indicators for the type of vehicle. OLS estimation used in columns 1-4. 2SLS estimation used in columns 5-8. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

5.4. Heterogeneous Effects: Size of the Debt and Number of Accounts

Examining the effects of our interventions across taxpayers in terms of the size of their debts and the number of accounts they possess at SEDAT yields interesting results. To some degree, both variables serve as imperfect proxies of income,[12] as the size of the vehicle tax—hence the amount owed—depends on the characteristics of the car, such as its commercial value. In addition, the number of accounts reflects the number of vehicles that a person or firm owns.[13] Consequently, in this section we test for heterogeneous treatment effects along these dimensions. For *each* treatment condition, we estimate logistic models of the probability of payment as a function of the baseline debt and the number of accounts. Then, we calculate the predicted values of such probabilities for the relevant range of values of this couple of predictors. Figure 4 plots heat maps of these predicted probabilities as a function of different combined values of the baseline debt and the number of accounts. Panel A plots the heat map for subjects in the Email treatment condition, panel B for the those in the Facebook condition, while panel C corresponds to taxpayers in the Email & Facebook condition.

Several observations are worth noting. First, in all three treatment groups—and the same is true if we pool the data and construct an equivalent heat map—the predicted payment probability is higher for individuals with a low baseline debt and a high number of accounts. This is made evident by the fact that the heat is higher in the southeastern region of the map in all three panels. It is reasonable to conclude that these interventions are particularly helpful to encourage compliance of small debtors. Second, despite the common pattern depicted by the maps, interesting differences appear between treatment groups. The coldest of all three maps is the one corresponding to the Facebook treatment group, in particular for the region of subjects with just a few number of accounts—which is the modal case. This is not surprising, as we have shown that this treatment condition produces the lowest, albeit significant, impact on tax compliance. In contrast, the Email & Facebook treatment produces the warmest map. In fact, a close inspection of the map shows that this specific condition is the most effective in encouraging small debtors with a small number of accounts to comply with taxes. In comparison, the Email treatment condition does not seem to provide as much incent. As such, this analysis of heterogeneous treatment effects reveals that targeted Facebook campaigns may be an effective complement of a more traditional

[12] Unfortunately, we did not have access to taxpayers' incomes.

[13] For this reason, there are subjects with a large number of accounts in our dataset. A firm may possess a large number of cars.

online strategy—namely sending email reminders—in order to encourage compliance from debtors that otherwise would forego their obligation.

5.5. Timing and Payment Method

In this subsection, we study two additional dimensions that might help understand the mechanisms underlying the effects of online strategies on tax compliance: the timing of payments and the payment methods employed by taxpayers to fulfill their obligations. First, concerning the timing of payments, heterogeneous treatment effects are expected, for at least two reasons: first, email reminders were sent on a single day, at the beginning of the experiment—September 13, 2018. Facebook ads, instead, were posted during a period of two weeks starting that same day, and it is unclear when subjects viewed the ads. In addition, the behavioral traits that we expect to be triggered by each strategy can differ as a function of time. Email messages are personalized, include private information of taxpayers, and, in this case, include a deterrence message. The goal is to pressure debtors into proceeding with an immediate payment. Conversely, Facebook ads cannot be personalized and simply provide general information, including the October 2nd deadline. Even though they include coercive signals, one might expect that ads may simply serve as reminders of the tax schedule, in such a way that debtors who respond to them do not proceed to payment immediately.

Figure 5 corroborates some of these intuitions. It plots the cumulative fraction of debtors that fulfill their obligations, as a function of time, disaggregated by treatment condition. Several observations are worth noting here. First, right after email reminders were sent, the distance between the fraction of payers in the groups that did and did not receive an email starts to grow. The treatment effect is immediate for some taxpayers and continues to grow until it stabilizes. Second, the distance between the groups that were exposed and not exposed to the social media ads starts to grow about two weeks after the campaign was launched and just before it ended. Therefore, it seems that accumulation of time and exposition is necessary for the Facebook treatment to have some sort of effect. It is interesting to note that a similar pattern can be observed when comparing the Email & Facebook versus Email groups, or the Facebook versus control groups.

Figure 4: Size of the Debt and Number of Accounts

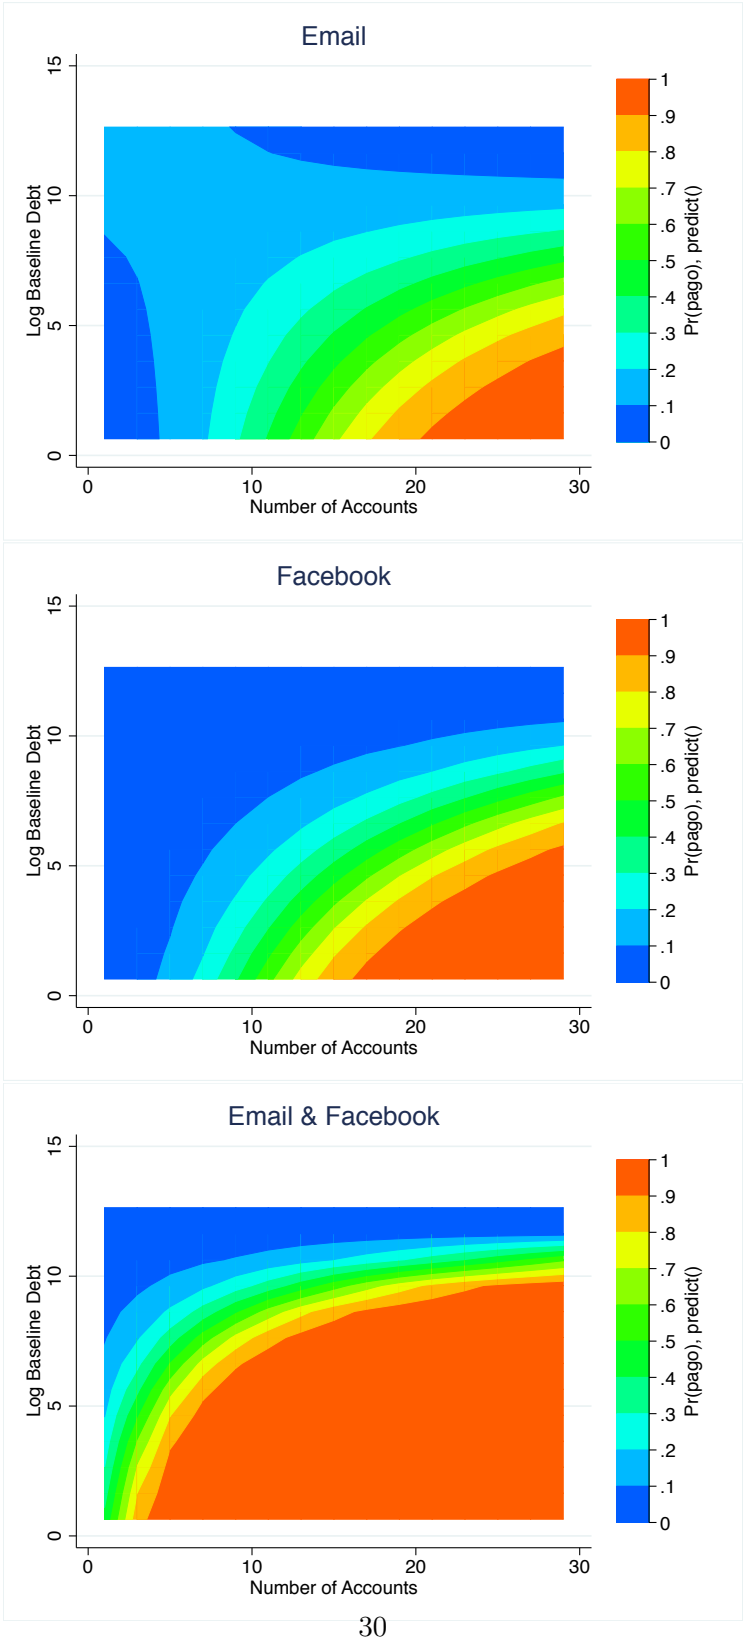

Figure 5: Payment Timing across Treatments

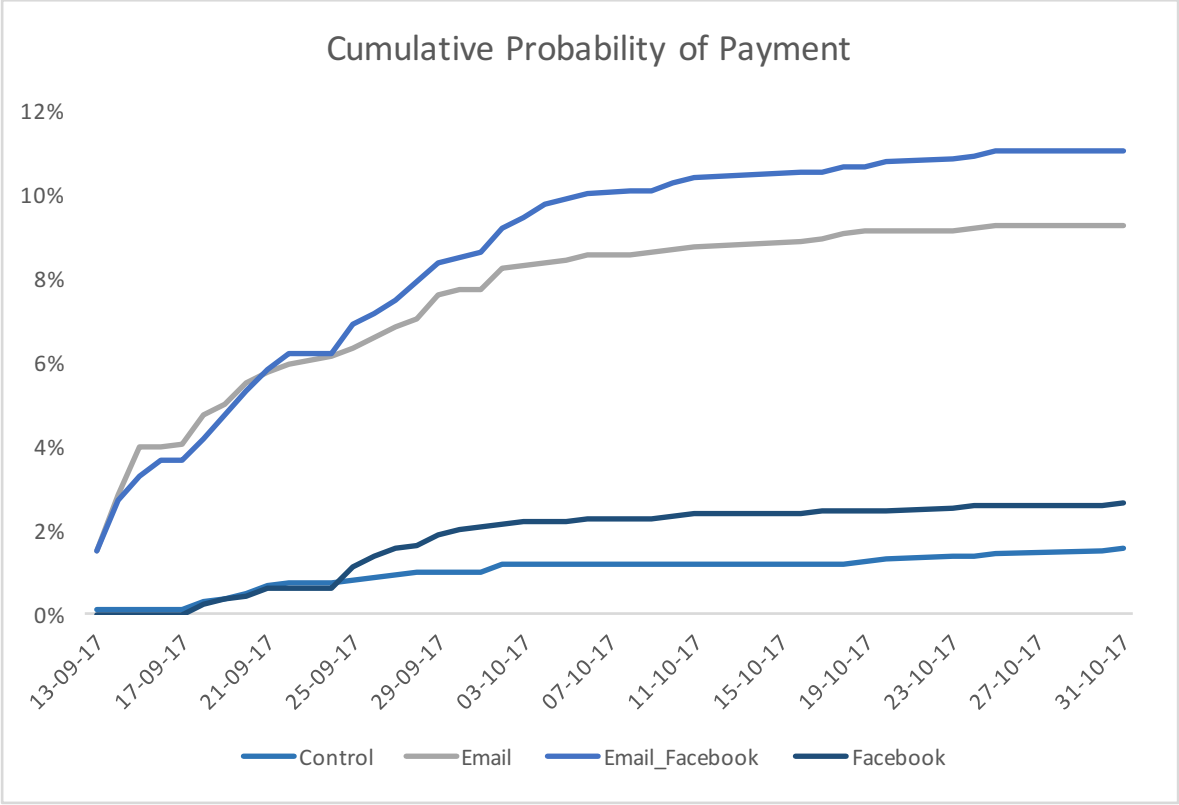

Table 6: Effects on Timing of Payments

	(1) One Week or Before	(2) One Week or Before	(3) Two Weeks or Before	(4) Two Weeks or Before	(5) Three Weeks or Before	(6) Three Weeks or Before	(7) Payment Anytime	(8) Payment Anytime
Email	0.0511***	0.0509***	0.0601***	0.0600***	0.0730***	0.0728***	0.0769***	0.0767***
	(0.00608)	(0.00607)	(0.00694)	(0.00694)	(0.00763)	(0.00763)	(0.00799)	(0.00800)
Facebook	-0.000649	-0.000899	0.00647	0.00615	0.0103**	0.0101**	0.0110**	0.0109**
	(0.00233)	(0.00233)	(0.00403)	(0.00402)	(0.00466)	(0.00464)	(0.00509)	(0.00508)
Email&Facebook	0.0505***	0.0504***	0.0686***	0.0684***	0.0893***	0.0892***	0.0944***	0.0944***
	(0.00604)	(0.00604)	(0.00726)	(0.00724)	(0.00815)	(0.00815)	(0.00852)	(0.00851)
Natural	-0.00874	-0.0132	-0.0439***	-0.0527***	-0.0591***	-0.0583***	-0.0679***	-0.0657***
	(0.00766)	(0.00870)	(0.0109)	(0.0122)	(0.0123)	(0.0132)	(0.0130)	(0.0139)
Constant	0.0123*	-0.0154	0.0489***	0.0191	0.0651***	0.0492	0.0762***	0.0718
	(0.00734)	(0.0349)	(0.0105)	(0.0466)	(0.0120)	(0.0537)	(0.0127)	(0.0548)
Covariates	N	Y	N	Y	N	Y	N	Y
N	6183	6183	6183	6183	6183	6183	6183	6183

Notes: Heteroskedastic robust standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Natural* is an indicator of natural versus legal person. *One Week or Before* indicates if a payment was done during the first week after the intervention started. *Two Weeks or Before* indicates if a payment was done during the first two weeks after the intervention started. *Three Weeks or Before* indicates if a payment was done during the first three weeks after the intervention started. *Payment Anytime* indicates if a payment was done anytime. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office. *Type* 1-*Type* 6 are indicators for the type of vehicle. OLS estimation used in columns 1-4. 2SLS estimation used in columns 5-8. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

Table 6 corroborates what the graph exhibits. For this purpose, we construct four dependent variables based on the payment timing of subjects: One Week or Before indicates whether the taxpayer paid during the first week of the experiment; Two Weeks or Before if she paid during the first two weeks, and so on; while *Anytime* is equivalent to the payment indicator used previously when we explored the main treatment effects. The first row of the Table shows that the effect of email reminders increases over time, as it has an important and significant effect during the first week—about 5 percentage points in the beginning, while it grows to 7.7 percentage points by the end. Only Facebook ads have no effects whatsoever during the first week, but by the third week they have a one percentage point impact. Also, during the first two weeks, the effects of the Email and the Email & Facebook treatments are very similar, but they start to diverge as time passes by.

Finally, we explore whether the treatments have differential effects on the payment methods chosen by taxpayers. At the moment of the interventions, debtors had three options in order to fulfill their obligation: online payment using a credit card, physical payment at their bank, or physical payment at the SEDAT office. Naturally, the second and third options represent costly actions that increase the transactions costs of tax compliance. In fact, both in our email messages and Facebook ads, links were included that redirected taxpayers to online platforms where payments could be made, in order to maximize compliance. Consequently, we expect to find differential effects at the level of the payment method, but it is interesting to determine if our treatments have impacts beyond reducing the transaction costs of compliance.

We reestimate our basic models, with binary variables indicating whether debtors paid online, at an affiliated bank, or at the SEDAT office specified as dependent variables. Rows 1 and 3 in Table 7 reveal that treatments including email reminders had positive impacts on the three forms of payment, but that the magnitude of the effects is higher in the case of online payments. Nonetheless, the effect is still important enough to encourage some debtors to pay at a bank or at the SEDAT offices. Row 2 is quite interesting. The positive effect of the Facebook treatment is driven by online payments exclusively: these ads are not enough to encourage a significant number of taxpayers to physically fulfill their obligation at an office. Needless to say, increasing online payments is something good for the tax authorities, as many transaction costs are reduced or avoided if this alternative becomes the favorite method to fulfill the obligations.

Table 7: Effects on Payment Methods

	(1) Online Payment	(2) Online Payment	(3) Bank Payment	(4) Bank Payment	(5) Office Payment	(6) Office Payment
Email	0.0414***	0.0412***	0.0103***	0.0104***	0.0259***	0.0257***
	(0.00552)	(0.00551)	(0.00313)	(0.00315)	(0.00526)	(0.00526)
Facebook	0.0103***	0.0101***	0.00324	0.00336	-0.00259	-0.00258
	(0.00340)	(0.00340)	(0.00230)	(0.00225)	(0.00316)	(0.00316)
Email&Facebook	0.0433***	0.0433***	0.0162***	0.0161***	0.0356***	0.0356***
	(0.00562)	(0.00562)	(0.00363)	(0.00362)	(0.00578)	(0.00579)
Natural	-0.00774	-0.0113	-0.0528***	-0.0480***	-0.00704	-0.00637
	(0.00742)	(0.00781)	(0.00903)	(0.00953)	(0.00691)	(0.00774)
Constant	0.0108	-0.0335	0.0497***	0.0978**	0.0153**	0.0159
	(0.00714)	(0.0309)	(0.00858)	(0.0414)	(0.00724)	(0.0248)
Covariates	N	Y	N	Y	N	Y
N	6183	6183	6183	6183	6183	6183

Notes: Heteroskedastic robust standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Natural* is an indicator of natural versus legal person. *Online Payment* indicates if an online payment took place. *Bank Payment* indicates if an in-person payment at the bank took place. $Office\ Payment$ indicates if an in-person payment at SEDAT's office took place. *Baseline Debt* is the natural log of the baseline debt plus 1. $Number\ of\ accounts$ is the number of accounts registered at the virtual office. $Type$ 1-$Type$ 6 are indicators for the type of vehicle. OLS estimation used in columns 1-4. 2SLS estimation used in columns 5-8. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

6. Conclusion

In this paper, we have presented the results of a field experiment implemented in Venezuela in collaboration with Sucre's tax authority. Our results suggest that online strategies are effective methods of increasing tax compliance. Not surprisingly, email reminders are more effective than social media ads, as they consist in a more intrusive, direct, and personalized way to approach citizens. Indeed, personal information can be included in such messages, a procedure that the literature has shown to be highly effective. However, the interesting question here is whether ads boost the effect of email reminders. In fact, we find that debtors exposed to social media ads, in addition to receiving the email message, are marginally more likely to pay compared to those that received an email without being exposed to the ads. The difference is small, but given the low cost of these targeted advertisement strategies, it may be cost-effective to implement them.

To the best of our knowledge, this is the first paper that systematically studies the effects of social media advertisements on tax compliance. It is somewhat surprising that scholars and tax authorities have ignored these methods to approach taxpayers, despite the fact that elsewhere, both in the private and the public sectors, online strategies have become popular because they allow very specific micro-targeting of populations. Moreover, the surge of e-government tools, such as the virtual office used by Sucre's tax authority, provides an important amount of private information on taxpayers, which can be used to guide these communication strategies.

Our study is also useful to uncover the differential effects that these interventions have across types of taxpayers. Little is known about the differential responses of firms and natural persons to interactions with the tax authority. We find that a more intrusive and personalized strategy, such as an email reminder, is more effective with natural persons. In contrast, despite being counterintuitive, social media ads may have stronger effects on firms. At least, this is true for small businesses like the ones approached in our intervention. We hypothesize that the reason for this is that ads serve as a reminder of the tax schedule and the deadlines associated to the process. This is enhanced by the fact that the ads were displayed during a range of time, as opposed to the one-shot email intervention. This reminder may be more helpful for entrepreneurs, which are probably busier and have more obligations to fulfill and tend to comply at higher rates. Methodologically, it is interesting that our results are supported by randomization inference techniques. Indeed, we use these for several reasons: the skewedness of the main outcome variable, and the perhaps insufficient number of firms we have to statistically power the analysis of heterogeneous treatment effects.

Finally, it is important to recognize that the case of Venezuela in 2017 is very particular, given the economic and political turmoil faced by the country, potentially undermining the external validity of our results. While this concern is justified, we think that one of the contributions of the study is that it was conducted in a developing country facing an economic and a political crisis, while most RCTs on tax compliance have been conducted in developed countries, with few exceptions. Therefore, our paper enhances our understanding of what type of strategies may be useful to increase fiscal capacity where it is most needed. Ultimately, the problem of evasion is worse in developing and rentier states where governments lack legitimacy. The good news is that we do find positive effects on compliance of the online strategies used in the experiment. Furthermore, the magnitudes of the effects are not substantially different from what has been found in other contexts and with other

types of interventions. To some extent, the Venezuelan situation should serve as a lower bound of what can be achieved in another context.

References

Allingham, M. and A. Sandmo (1972). Income Tax Evasion: A Theoretical Analysis. *Journal of Public Economics 1*(3), 323–338.

Allison and A. Strauss (2009). Don't Forget to Vote: Text Message Reminders as a Mobilization Tool. *American Journal of Political Science 53*, 787–804.

Alvarez, M., A. Hopkins, and B. Sinclair (2010). Mobilizing Pasadena Democrats: Measuring the Effects of Partisan Campaign Contacts. *Journal of Politics 72*, 31–44.

Aral, S. and D. Walker (2011). Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks. *Management Science 57*(9), 1623–1639.

Bakshy, E., D. Eckles, R. Yan, and I. Rossen (2012). Social Influence in Social Advertising: Evidence from Field Experiments. Proceedings of the 13th ACM Conference on Electronic Commerce.

Barro, R. (1990). Government Spending in a Simple Model of Endogeneous Growth. *Journal of Political Economy 98*, 103–125.

Barro, R. and X. Sala-i-Martin (1991). Public Finance Models of Economic Growth. *Review of Economic Studies 59*, 645–661.

Benjamini, Y. and Y. Hochberg (1995). Tags controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal of the Royal Statistical Society 57*(1), 289–300.

Besley, T. and T. Persson (2009). The Origins of State Capacity: Property Rights, Taxation, and Politics. *American Economic Review 99*(4), 1218–1244.

Besley, T. and T. Persson (2013). Taxation and Development. In *Handbook of Public Economics*. Elsevier.

Besley, T. and T. Persson (2014). Why do developing countries tax so little? *Journal of Economic Perspectives 28*(4), 99–120.

Boning, W., J. Guyton, R. Hodge, J. Slemrod, and U. Troiano (2018). Heard it through the Grapevine: A Randomized Experiment Assessing Direct and Network Effects of Tax Treatments Strategies. Working Paper, NBER.

Brookman, D. and D. Green (2014). Do Online Advertisements Increase Political Candidates' Name Recognition or Favorability? Evidence from Randomized Field Experiments. *Political Behavior 36*(2), 263–289.

Cardy, E. (2005). An Experimental Field Study of the GOTV and Persuasion Effects of Partisan Direct Mail and Phone Calls. *The Annals of the American Academy of Political and Social Science 601*, 28–40.

Castro, L. and C. Scartascini (2015). Tax Compliance and Enforcement in the Pampas: Evidence from a Field Experiment. *Journal of Economic Behavior and Organization 116*, 65–82.

CEPAL (2016). Estudio Económico para Amética Latina y el Caribe 2016. Technical report, Comisión Económica para América Latina y el Caribe.

CEPAL (2017). Estudio Económico para Amética Latina y el Caribe 2017. Technical report, Comisión Económica para América Latina y el Caribe.

del Carpio, L. (2014). Are the Neighbors Cheating? Evidence from a Social Norm Experiment on Property Taxes in Peru. Working Paper, Princeton University.

Dhami, S. and A. al Nowaihi (2007). Why Do People Pay Taxes? Prospect Theory versus Expected Utility Theory. *Journal of Economic Behavior and Organization 64*(1), 171–192.

Eckles, D., B. Gordon, and G. Johnson (2018). Field Studies of Psychologically Targeted Ads Face Threats to Internal Validity. *PNAS 115*(23), E5254–E5255.

Gadenne, L. (2015). Tax Me, but Spend Wisely? Sources of Public Finance and Government Accountability. Technical report, Institute of Fiscal Studies.

Gallego, J. and F. Ortega (2018). Fwd: Re: Don't Forget to Pay your Taxes. Email Reminders and Tax Compliance in Venezuela. Working Paper, Universidad del Rosario.

Gerber, A. and D. Green (2000). The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment. *American Political Science Review 94*, 653–663.

Gerber, A. and D. Green (2001). Do Phone Calls Increase Voter Turnout? A Field Experiment. *Public Opinion Quarterly 65*, 75–85.

Gerber, A. and D. Green (2012). *Field Experiments: Design, Analysis, and Interpretation.* W.W. Norton.

Gerber, A. and D. Green (2015). *Get Out the Vote. How to Increase Voter Turnout.* Brookings Institution Press.

Gerber, A., D. Green, and M. Green (2003). Partisan Mail and Voter Turnout: Results from Randomized Field Experiments. *Electoral Studies 22*, 563–579.

Goldfarb, A. and C. Tucker (2011). Search Engine Advertising: Channel Substitution when Pricing ads to Context. *Management Science 57*(3), 458–470.

Green, D., M. McGrath, and P. Aronow (2013). Field Experiments and the Study of Voter Turnout. *Journal of Elections, Public Opinion and Parties 23*(1), 27–48.

Hallsworth, M., J. List, R. Metcalfe, and I. Vlaev (2017). The Behavioralist as Tax Collector: Using Natural Field Experiments to Enhance Tax Compliance. *Journal of Public Economics 148*, 14–31.

Imbens, G. and D. Rubin (2015). *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press.

Lacetera, N., M. Macis, and A. Mele (2016). Viral Altruism? Charitable Giving and Social Contagion in Online Networks. *Sociological Science 3*(11), 202–238.

Lewis, R., D. Reiley, and T. Schreiner (2012). Ad Attributes and Attribution: Large Scale Field Experiments Measure Online Consumer Acquisition. Working Paper, Google.

Malhotra, N., M. Michelson, T. Roger, and A. Valenzuela (2011). Text Messages as Mobilization Tools: The Conditional Effect of Habitual Voting and Election Salience. *American Politics Research 39*, 664–681.

Martinez, L. (2017). Sources of Revenue and Government Performance: Evidence from Colombia. Working Paper, University of Chicago.

Meiselman, B. (2018). Ghostbusting in Detroit: Evidence on Nonfilers from a Controlled Field Experiment. *Journal of Public Economics 158*, 180–193.

Michelson, M. (2005). Meeting the Challenge of Latino Voter Mobilization. *The Annals of the American Academy of Political and Social Science 601*, 85–101.

Myles, G. and R. Naylor (1996). A Model of Tax Evasion with Group Conformity and Social Customs. *European Journal of Political Economy 12*(1), 49–66.

Nickerson, D. (2007). Does Email Boost Turnout? *Quarterly Journal of Political Science 2*, 369–379.

North, D. and Weingast (1989). Constitutions and Commitmment The Evolution of Institutions Governing Public Choice in Seventeenth-Century England. *Journal of Economic History 49*(4), 803–832.

Ortega, D., L. Ronconi, and P. Sanguinetti (2013). Reciprocity and Willingness to Pay Taxes, Evidence from a Survey Experiment in Latin America. Working Paper, CIAS.

Ortega, D. and P. Sanguinetti (2013). Deterrence and Reciprocity Effects on Tax Compliance: Experimental Evidence from Venezuela. Working Paper, CAF.

Ortega, D. and C. Scartascini (2015). Don't Blame the Messenger. A Field Experiment on Delivery Methods for Increasing Tax Compliance. Working Paper, CAF.

Panagopoulos, C. and D. Green (2010). Spanish-Language Radio Advertisements and Latino Voter Turnout in the 2006 Congressional Elections: Field Experimental Evidence.

Political Research Quarterly 64, 588–599.

Ryan, T. (2012). What Makes Us Click? Demonstrating Incentives for Angry Discourse with Digital-Age Field Experiments. *Journal of Politics 74*(4), 1138–1152.

Slemrod, J. (2018). Tax Compliance and Enforcement. Working Paper, NBER.

Slemrod, J., M. Blumenthal, and C. 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.

Snow, A. and R. S. Warren (2005). Ambiguity about Audit Probability, Tax Compliance, Taxpayer Welfare. *Economic Inquiry 43*(4), 865–871.

Stollwerk, A. (2006). Does E-mail Affect Voter Turnout? An Experimental Study of the New York City 2005 Election. Working Paper, Yale University.

Torgler, B. (2004). Moral suasion: an alternative tax policy strategy? evidence from a controlled field experiment in switzerland. *Economics of Governance 5*(3), 235–253.

Traxler, C. (2010). Social Norms and Conditional Cooperative Taxpayers. *European Journal of Political Economy 26*(1), 89–103.

Yaniv, G. (1999). Tax Compliance and Advance Tax Payments: A Prospect Theory Analysis. *National Tax Journal 52*(4), 753–764.

Young, A. (2017). Channelling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results. Working Paper, LSE.

Appendix

A. Summary Statistics

Table A.1 reports basic descriptive statistics of the outcomes, treatment variables, and covariates used in this study. It is interesting to note that the overall proportion of subjects paying their obligation is quite low: only 6% comply with their obligation. This low proportion explains the sharp difference between the average baseline debt and the average amount paid.

Table A.1: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Baseline Debt (in logs)	7.548	0.73	0.631	12.914	6183
Number of Accounts	1.532	3.437	1	197	6183
Natural Person	0.892	0.31	0	1	6183
Type 1 Vehicle	0.639	0.48	0	1	6183
Type 2 Vehicle	0.222	0.415	0	1	6183
Type 3 Vehicle	0.051	0.221	0	1	6183
Type 4 Vehicle	0.025	0.155	0	1	6183
Type 5 Vehicle	0	0.018	0	1	6183
Email	0.25	0.433	0	1	6183
Facebook	0.25	0.433	0	1	6183
Email&Facebook	0.25	0.433	0	1	6183
Payment Indicator	0.061	0.24	0	1	6183
Amount Paid (in logs)	0.46	1.819	0	10.937	6183
Online Payment	0.028	0.164	0	1	6183
Bank Payment	0.01	0.1	0	1	6183
Office Payment	0.024	0.152	0	1	6183
Payment $\leq$ 1 Week	0.03	0.17	0	1	6183
Payment $\leq$ 2 Week	0.044	0.204	0	1	6183
Payment $\leq$ 3 Week	0.055	0.229	0	1	6183

B. Timeline of the Experiment

Figure A.1 depicts the timeline of the experiment.

Figure A.1: Timeline of the Experiment

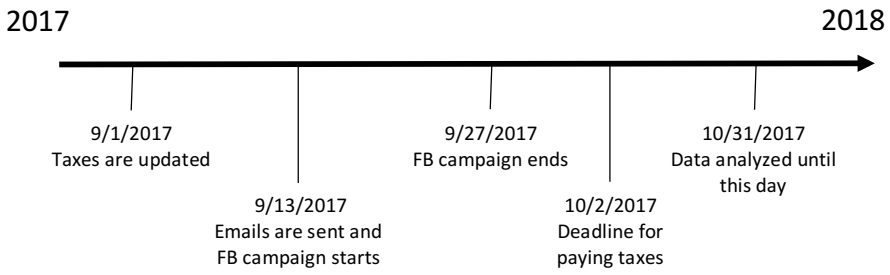

C. Probit Models

We estimate the main treatment effects of our interventions on the probability of fulfilling the obligations using Linear Probability Models (see Tables 3 and 4). In this section, we reestimate those effects, but using Probit and IVProbit models, to account for the fact that the outcome variable is dichotomous. Table A.2 shows that these models yield similar results.

Table A.2: Probit Models of Treatment Effects

	(1) Payment Indicator	(2) Payment Indicator	(3) Payment Indicator	(4) Payment Indicator	(5) Payment Indicator	(6) Payment Indicator	(7) Payment Indicator	(8) Payment Indicator
Email	0.866***	0.866***	0.866***	0.866***	1.738***	1.736***	1.737***	1.736***
	(0.0917)	(0.0917)	(0.0917)	(0.0918)	(0.221)	(0.221)	(0.222)	(0.223)
Facebook	0.234**	0.233**	0.231**	0.232**	0.226**	0.226**	0.223**	0.226**
	(0.105)	(0.105)	(0.105)	(0.105)	(0.100)	(0.100)	(0.100)	(0.101)
Email&Facebook	0.964***	0.965***	0.963***	0.964***	2.053***	2.054***	2.048***	2.047***
	(0.0907)	(0.0907)	(0.0906)	(0.0907)	(0.216)	(0.216)	(0.217)	(0.218)
Natural	-0.501***	-0.489***	-0.484***	-0.492***	-0.550***	-0.544***	-0.538***	-0.517***
	(0.0722)	(0.0778)	(0.0778)	(0.0821)	(0.0730)	(0.0778)	(0.0780)	(0.0845)
Constant	-1.757***	-1.895***	-1.799***	-1.892***	-1.631***	-1.698***	-1.573***	-1.358***
	(0.104)	(0.330)	(0.338)	(0.392)	(0.0753)	(0.313)	(0.322)	(0.368)
Baseline Debt	N	Y	Y	Y	N	Y	Y	Y
No. Accounts	N	N	Y	Y	N	N	Y	Y
Type of Vehicle	N	N	N	Y	N	N	N	Y
Estimator	Probit	Probit	Probit	Probit	IVProbit	IVProbit	IVProbit	IVProbit
N	6183	6183	6183	6181	6183	6183	6183	6181

Notes: Heteroskedastic robust standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Natural* is an indicator of natural versus legal person. *Payment Indicator* is a dummy indicating whether the taxpayer paid her debt or not. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office. *Type* 1-*Type* 6 are indicators for the type of vehicle. Probit estimation used in columns 1-4. IVPobit estimation used in columns 5-8. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.

D. Multiple Comparisons Corrections

In this section, we present the results of the multiple comparisons analysis conducted for the main results of the experiment. In our case, different causes may lead to false discoveries, as we have multiple treatments, outcomes, estimators, as well as covariate adjustments. Hence, in this context of multiple comparisons, the probability of committing Type I error is bigger and consequently it is reasonable to adjust our p-values. To this purpose, we follow the *False Discovery Rate* approach, which is less stringent than techniques based on *Family-Wise Error Rates* such as the Bonferroni correction. In particular, we implement Benjamini-Hochberg's procedure (Benjamini and Hochberg, 1995).

Under this method, we order our m p-values and rank them, from the lowest to the highest. Then, we calculate the Benjamini-Hochberg critical value (BH value hereafter) for each p-value, defined as $BH = \alpha\frac{k}{m}$, where α is the critical significance level employed (usually 0.05) and k is the p-value's ranking. Finally, we determine the highest p-value that satisfies $p_k \leq BH$. All tests with p-values smaller than the one just described, are declared significant. We implement this method to the estimations reported on Tables 3 and 4, which correspond to out main models of the treatment effects. In total, we have 16 models and three coefficients of interest, one per treatment arm. Hence, we analyze a total of 48 p-values, as a consequence of different combinations of outcomes, specifications, estimators, and treatments.

Tables A.3 and A.4 report the results of Benjamini-Hochberg corrections, resembling Tables 3 and 4. For every p-value, we calculate the BH critical value and determine wether we reject the null hypothesis of absence of treatment effect. Note that even for the highest p-value of all models, such magnitude is lower than its correspondent BH value. Therefore, in every case we can reject the null hypothesis. Consequently, even after taking into account the multiple comparisons problem, we find evidence that our online strategies had positive and significant effects on tax compliance.

Table A.3: Benjamini-Hochberg Multiple Comparisons Corrections

		(1) Payment Indicator	(2) Payment Indicator	(3) Payment Indicator	(4) Payment Indicator	(5) Payment Indicator	(6) Payment Indicator	(7) Payment Indicator	(8) Payment Indicator
Email	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	BH value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	Reject H0	1	1	1	1	1	1	1	1
Facebook	P-value	0.031	0.031	0.031	0.032	0.031	0.031	0.031	0.030
	BH value	0.035	0.035	0.035	0.042	0.035	0.035	0.035	0.034
	Reject H0	1	1	1	1	1	1	1	1
Email&FB	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	BH value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	Reject H0	1	1	1	1	1	1	1	1
Baseline Debt		N	Y	Y	Y	N	Y	Y	Y
No. Accounts		N	N	Y	Y	N	N	Y	Y
Type of Vehicle		N	N	N	Y	N	N	N	Y
Estimator		OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS	2SLS

Table A.4: Benjamini-Hochberg Multiple Comparisons Corrections (continued)

		(1) Log Amount Paid	(2) Log Amount Paid	(3) Log Amount Paid	(4) Log Amount Paid	(5) Log Amount Paid	(6) Log Amount Paid	(7) Log Amount Paid	(8) Log Amount Paid
Email	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	BH value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	Reject H0	1	1	1	1	1	1	1	1
Facebook	P-value	0.034	0.035	0.036	0.037	0.034	0.035	0.035	0.034
	BH value	0.043	0.046	0.049	0.05	0.043	0.046	0.046	0.043
	Reject H0	1	1	1	1	1	1	1	1
Email&FB	P-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	BH value	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
	Reject H0	1	1	1	1	1	1	1	1
Baseline Debt		N	Y	Y	Y	N	Y	Y	Y
No. Accounts		N	N	Y	Y	N	N	Y	Y
Type of Vehicle		N	N	N	Y	N	N	N	Y
Estimator		OLS	OLS	OLS	OLS	2SLS	2SLS	2SLS	2SLS

E. First Stage Regressions

In this section, we report the results of the first-stage regressions of models in which we estimate the Compliers Average Causal Effect (CACE) for email reminders (see columns 5-8 in Tables 3 and 4). Remember that we have information on who read our reminders—hence, on who complies with these treatments. Therefore, we can use treatment assignment to the Email and Email&Facebook conditions as instruments for compliance. Table A.5 reports the results of the first stages of such models. The Table shows, as expected, that assignment to the Email condition is a strong predictor of whether a subject reads an email in that group. Similarly, assignment to the Email&Facebook condition strongly predicts whether someone in that group reads its associated reminder. Hence, the instruments are not weak.

Note that in the case of Facebook ads, unfortunately we cannot include a similar instrument, as we do not have individual information on who saw the advertisements. Hence, our identifying assumption in this case is that exposure or not to social media ads is orthogonal to any other characteristics that may affect compliance. We argue, however, that any non-compliance in exposure to the Facebook ads would downwardly bias our estimations of the effects of the Facebook and Email & Facebook conditions. On the one hand, subjects in these two groups incorrectly *not* exposed to the ads, would be less likely to pay their taxes. On the other hand, subjects in the Control and the Email groups incorrectly exposed to the ads would be more likely to pay. Hence, our estimates of the effects of treatments including ads should be lower bounds of the potential effects of these interventions.

Table A.5: First Stage Results of CACE Models

	(1) Email Complier	(2) Email Complier	(3) Email Complier	(4) Email Complier	(5) Email & FB Complier	(6) Email & FB Complier	(7) Email & FB Complier	(8) Email & FB Complier
Email	0.378***	0.378***	0.378***	0.377***	0.00000141	-0.0000864	-0.0000962	-0.000143
	(0.0123)	(0.0123)	(0.0123)	(0.0123)	(0.0000279)	(0.000149)	(0.000155)	(0.000470)
Email&Facebook	-0.000000361	0.000123	0.000172	0.000137	0.364***	0.364***	0.364***	0.364***
	(0.0000579)	(0.000186)	(0.000206)	(0.000478)	(0.0122)	(0.0122)	(0.0122)	(0.0122)
Facebook	-0.000000361	-0.0000427	-0.0000218	-0.000386	-0.000000170	-0.0000254	-0.0000109	-0.000182
	(0.0000579)	(0.000159)	(0.000176)	(0.000524)	(0.0000273)	(0.0000968)	(0.000112)	(0.000514)
Natural	0.00520	0.00950	0.00885	0.00390	0.00245	0.00501	0.00456	-0.00104
	(0.00989)	(0.0104)	(0.0104)	(0.0109)	(0.00983)	(0.0104)	(0.0105)	(0.0109)
Constant	-0.00464	-0.0525	-0.0602	-0.118***	-0.00218	-0.0307	-0.0361	-0.104**
	(0.00883)	(0.0362)	(0.0369)	(0.0432)	(0.00877)	(0.0392)	(0.0406)	(0.0483)
Baseline Debt	N	Y	Y	Y	N	Y	Y	Y
No. Accounts	N	N	Y	Y	N	N	Y	Y
Type of Vehicle	N	N	N	Y	N	N	N	Y
N	6183	6183	6183	6183	6183	6183	6183	6183

Notes: Heteroskedastic robust standard errors are shown in parentheses. *Email*, *Facebook*, and *Email&Facebook* indicate each treatment condition, respectively. *Natural* is an indicator of natural versus legal person. *Email Complier* indicates whether a subject in the Email condition reads the reminder. *Email & FB Complier* indicates whether a subject in the Email&FB condition reads the reminder. *Baseline Debt* is the natural log of the baseline debt plus 1. *Number of accounts* is the number of accounts registered at the virtual office. *Type* 1-*Type* 6 are indicators for the type of vehicle. OLS estimation used in columns 1-4. 2SLS estimation used in columns 5-8. * is significant at the 10% level, ** is significant at the 5% level, *** is significant at the 1% level.