

NBER WORKING PAPER SERIES

TAX AUDITS AS SCARECROWS:  
EVIDENCE FROM A LARGE-SCALE FIELD EXPERIMENT

Marcelo L. Bérgholo  
Rodrigo Ceni  
Guillermo Cruces  
Matias Giacobasso  
Ricardo Perez-Truglia

Working Paper 23631  
<http://www.nber.org/papers/w23631>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
July 2017, Revised June 2021

We thank the Uruguay's national tax administration (Dirección General Impositiva) for their collaboration. We thank Gustavo Gonzalez for his support, without which this research would not have been possible. We thank Joel Slemrod for his valuable feedback, as well as that of seminar participants at University of Berkeley, University of Michigan, University of California San Diego, Dartmouth University, Universidad Di Tella, Universidad de la Republica, Universidad de Santiago de Chile, Universidad Católica de Chile, UADE, Corporación Andina de Fomento (Buenos Aires), Banco Central del Uruguay, LACEA 2017, the 2017 NBER Public Economics Fall Meeting, the 2017 RIDGE Public Economics Conference, the 2017 Zurich Center for Economic Development Conference, the 2017 Advances with Field Experiments Conference, the 2018 PacDev Conference, the 2018 AEA Annual Meetings, the 2018 LAGV Conference, the 2018 IIPF Annual Congress, the 2019 LACEA BRAIN Conference and the 2019 National Tax Association meeting. This project benefited from funding by CEF, CEDLAS-UNLP and IDRC. The AEA RCT registration number is AEARCTR-0004593. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Marcelo L. Bérgholo, Rodrigo Ceni, Guillermo Cruces, Matias Giacobasso, and Ricardo Perez-Truglia. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Tax Audits as Scarecrows: Evidence from a Large-Scale Field Experiment

Marcelo L. Bérigolo, Rodrigo Ceni, Guillermo Cruces, Matias Giacobasso, and Ricardo Perez-Truglia  
NBER Working Paper No. 23631

July 2017, Revised June 2021

JEL No. C93,H26,K42

**ABSTRACT**

The canonical model of Allingham and Sandmo (1972) predicts that firms evade taxes by optimally trading off the costs and benefits of evasion. However, there is no direct evidence that firms react to audits in this way. We conducted a large-scale field experiment in collaboration with Uruguay's tax authority to address this question. We sent letters to 20,440 small- and medium-sized firms that collectively paid more than 200 million U.S. dollars in taxes per year. Our letters provided exogenous yet nondeceptive signals about key inputs for their evasion decisions, such as audit probabilities and penalty rates. We measured the effect of these signals on their subsequent perceptions of the auditing process using survey data, as well as on the actual taxes paid using administrative data. We find that providing information about audits had a significant effect on tax compliance but in a manner that was inconsistent with Allingham and Sandmo (1972). Our findings are consistent with an alternative model of risk-as-feelings, in which messages about audits generate fear and induce probability neglect. According to this model, audits may deter tax evasion in the same way that scarecrows frighten off birds.

Marcelo L. Bérigolo  
Instituto de Economía (IECON)  
Universidad de La República  
1926 Gonzalo Ramírez  
Montevideo  
Uruguay  
and IZA  
mbergolo@iecon.ccee.edu.uy

Rodrigo Ceni  
Instituto de Economía (IECON)  
Universidad de La República  
1375 Joaquín Requena  
Montevideo, Uruguay  
rceni@iecon.ccee.edu.uy

Guillermo Cruces  
CEDLAS  
Universidad Nacional de La Plata  
Calle 6 entre 47 y 48  
La Plata, Argentina  
gcruces@cedlas.org

Matias Giacobasso  
Anderson School of Management - UCLA  
110 Westwood Plaza, C 3.10  
Los Angeles, CA 90077  
mggiacobasso@ad.ucla.edu

Ricardo Perez-Truglia  
Haas School of Business  
University of California, Berkeley  
545 Student Services Building #1900  
Berkeley, CA 94720-1900  
and NBER  
ricardotruglia@berkeley.edu

# 1 Introduction

Tax audits are a standard tool that most tax administrations have used throughout history. Audits increase tax revenues directly, because firms caught evading must pay taxes on the hidden income and corresponding penalties. However, with the exception of large taxpayers, these direct revenues are insufficient to make audits cost-effective. Audits play a central role in the deterrence paradigm of tax evasion: the threat of being audited in the future, being caught evading and having to pay penalties deters firms from evading taxes in the present.

Audits may be useful in fostering tax compliance, but there is no direct evidence on how firms react to them. The Allingham and Sandmo (1972) model (hereafter referred to as *AES*) is the canonical model of tax evasion in economics. This model is an application of Becker (1968), in which selfish individuals choose whether to engage in criminal activities based on the trade-off between expected costs and benefits. In *AES*, firms choose the optimal amount of income to hide from the tax authority so that the marginal benefits (i.e., the lower tax burden) equal the marginal costs (i.e., the penalties they will be required to pay if caught). This intuition is so deeply ingrained in economic thought that most economists take it for granted. However, surprisingly little causal evidence exists on whether real firms react to audits in this profit-maximizing fashion (Alm et al., 1992; Luttmer and Singhal, 2014; Slemrod, 2018). In this study, we provide direct tests of the *AES* model based on a high-stakes, large-scale field experiment.

We study small- and medium-sized firms in Uruguay that are subject to the Value Added Tax (VAT), which provides a context in which taxpayers should care about the threat of being audited. This is not always the case: tax agencies can sometimes use third-party reporting to automatically detect and rectify tax evasion regardless of whether the taxpayer is audited or not, thus making the audit threat irrelevant. For instance, the U.S. Internal Revenue Service uses their electronic records to compare the wage amount reported by an individual to the amount reported by the individual's employer. This algorithm automatically rectifies the discrepancies in reporting and sends a notification to the taxpayer with the updated tax amount to be paid. Because the evasion will be caught through the third-party reporting regardless of whether the individual is audited or not, taxpayers should not care about the threat of being audited in such a context (Kleven et al., 2011). On the contrary, in our context of the VAT in a developing country, such automatic cross-checking and rectification does not exist. The VAT paper trail, which consists of non-electronic invoices, can only be scrutinized in the event of an audit.<sup>1</sup> Thus, tax authorities must still rely heavily on the

---

<sup>1</sup>While the VAT requires a paper trail, which is a form of third-party reporting, this paper trail is subject to significant limitations. Most importantly, there is no simple algorithm that automatically detects tax evasion. Moreover, the paper trail breaks down when reaching the consumer (Naritomi, 2019). Finally, firms

threat of audits to discourage VAT evasion (Gomez-Sabaini and Jimenez, 2012; Bergman and Nevarez, 2006).

We collaborated with Uruguay’s Internal Revenue Service (hereafter referred to as ‘IRS’) to conduct a natural field experiment with a sample of 20,440 small- and medium-sized firms that are subject to the VAT. For our study, the IRS mailed four different types of letters with information about audits to the owners of each of these firms.<sup>2</sup> Some of the information contained in each of these letters was randomly assigned, with the goal of testing predictions of *A&S*. Using IRS administrative records, we measured the subsequent effects of the information contained in the letters on the firms’ compliance with the VAT and other tax liabilities in the following year. Additionally, we collaborated with the IRS to conduct a post-mailing survey to capture the effect of this information on these firms’ subsequent perceptions about audits.

The first part of the experimental design, following the seminal work by Slemrod et al. (2001), measures how informing taxpayers about tax enforcement affects their tax compliance. Firms were randomized into four different letter types: *baseline*, *audit-statistics*, *audit-endogeneity*, and *public-goods*. The *baseline* letter type included brief and generic tax information that the IRS often includes in its communications with firms. The *audit-statistics* letter type was identical to the *baseline* letter, with additional information about the probability of being audited and the penalty rate based on tax administration statistics. The relevant hypothesis is that adding the *audit-statistics* message to the *baseline* letter will deter tax evasion and thus increase post-treatment tax payments. We also can compare the effects of this *audit-statistics* message with the effects of other types of messages. The *audit-endogeneity* letter type provided information about a different feature of the auditing process. It was identical to the *baseline* letter, with an additional message about how evading taxes increases the probability of being audited. The *public-goods* letter type was designed to provide a benchmark for a message that might increase tax compliance but does not provide information about tax audits. It was identical to the *baseline* letter, with an additional message describing the social costs of evasion by detailing the set of public goods that could be provided if tax evasion was lower.

We show that, consistent with Slemrod et al. (2001) and the subsequent literature, informing firms about tax enforcement increases their tax compliance. We find that adding the *audit-statistics* message to the *baseline* letter increases tax payments by about 7.0% in the first post-treatment year; in the second post-treatment year the effects are still present,

---

can also collude with each other to tamper with the paper trail (Pomeranz, 2015).

<sup>2</sup>Throughout the paper, for simplicity, we refer to firms’ perceptions and behavior as a shorthand for the perception and behavior of the firms’ owners or managers.

but they are half as large and no longer statistically significant. This effect is economically significant: the estimated average VAT evasion rate in Uruguay is 26% (Gomez-Sabaini and Jimenez, 2012). While the tax base is not necessarily fully comparable, this figure implies that the 7.0% increase equates to a 27% reduction in VAT evasion. The effect of the *audit-statistics* message (increased tax payments 7.0% in the first year) is similar in magnitude to the effect of the other message related to tax audits (7.1%, for *audit-endogeneity*) and larger and more persistent than the effect of the *public-goods* message (5.1% in the first year, but negligible in the second year).

The second and most important part of the experimental design tests the hypothesis that firms react to information about audits as predicted by *AES*. We provide two tests of *AES*. The first test exploits survey data on perceptions about audits. If the *audit-statistics* letter increased average compliance, to be consistent with *AES*, it must be true that this message increased the perceived probability of being audited or the perceived penalty rate. To test this hypothesis, we designed a survey, which was sent months after the firms received the *audit-statistics* and *audit-endogeneity* letters, which elicited perceptions about the probability of being audited and the penalty rate.

The second test of *AES* is based on heterogeneity in the signals provided in the letters. We included exogenous, non-deceptive variation in the information about audit probabilities and penalty rates in the *audit-statistics* letter. To generate this information, we computed the average probabilities and penalty rates using a series of random samples of 50 firms. This sample size was small enough to introduce non-trivial sampling variation in the average probabilities and fines shown to the subjects. Specifically, a given firm could receive a letter saying that the audit probability is 8%, 10%, or 15%, depending on the sample of similar firms chosen for that particular letter. These random variations in probabilities and penalties shown to the firms allow us to test whether firms evade less when they face higher audit probabilities and higher penalty rates, as predicted by *AES*.

The second part of the results suggests that the effects of the *audit-statistics* letter are not consistent with *AES*. The results for the first test, based on the survey data, indicate that the *audit-statistics* message reduced the perceived probability of being audited. According to *AES*, a reduction in the perceived probability of being audited should have reduced tax compliance. On the contrary, we find that the *audit-statistics* message increased average compliance.

The second test shows that, contrary to the *AES* prediction, the effect of the *audit-statistics* message does not change with the signals of audit probability and penalty rates included in the letter. The estimated elasticity of tax compliance with respect to audit probabilities and penalty rates is close to zero and precisely estimated. Moreover, we compare

our experimental estimates to the results from calibrations of *A&S*. We reject the null hypothesis of *A&S* even under conservative assumptions about how much firms learned from the *audit-statistics* message. These results suggest the presence of probability neglect – i.e., that firms react similarly to the threat of being audited regardless of the actual probability of that happening or the penalties involved.

As a complement to the *audit-statistics* treatment arm, we designed a separate treatment arm that created exogenous variation in expected audit probabilities in a more direct way. The *audit-threat* letter type was sent to a separate sample of firms that were pre-selected by the IRS for auditing. We randomly divided this set of firms into two groups, one with a 25% probability of being audited and the other with a 50% probability of being audited. The *audit-threat* letter informed firms of the audit probability that was assigned to them. Consistent with the *audit-statistics* treatment arm, we find probability neglect in the *audit-threat* arm too.

In sum, we find that informing firms about tax audits increased their tax compliance, but this reaction to the information was inconsistent with the optimal reaction predicted by *A&S*. On average firms reduced, rather than increased, their perceived probability of being audited. Furthermore, firms did not react more when facing a higher probability of being audited or a higher penalty rate. These results suggest that firms may comply with taxes because of the threat of being audited but not necessarily in the optimal way as predicted by *A&S*.

This leaves open the question of which alternative model best explains the firms' reactions to audits. Models of salience (Chetty et al., 2009) and prospect theory (Kahneman and Tversky, 1979) can explain some, but not all, findings. Agency issues and information frictions within VAT-paying firms subject to third-party reporting have implications for evasion that differ from the context of individuals paying an income tax, as highlighted in recent models of firm evasion (Kleven et al., 2016), but these issues are not consistent with our evidence either. Instead, our preferred interpretation is based on the model of risk-as-feelings (Loewenstein et al., 2001). The models used for choice under risk are typically cognitive; that is, people make decisions using some type of expectation-based calculus. The risk-as-feelings model proposes that responses to fearsome situations may differ substantially from cognitive evaluations of the same risks. When fear is involved, the responses to risks are quick, automatic and intuitive, and thus neglect the underlying probabilities (Sunstein, 2003; Zeckhauser and Sunstein, 2010). This model of risk-as-feelings can reconcile all of our key findings. Moreover, we present anecdotal and survey evidence indicating that the fear of being audited does indeed play a significant role in tax compliance. We also discuss policy implications for increasing tax capacity.

Our study relates to various strands of literature. First, it belongs to a recent but growing literature that uses field experiments in partnership with tax authorities to study the decisions of individuals to pay taxes. In a seminal contribution, Slemrod et al. (2001) showed that, for a sample of U.S. self-employed individuals, those who were randomly assigned to receive a letter from the Minnesota Department of Revenue with an enforcement message reported higher income in their tax returns. Similar messages about tax enforcement have been shown to have positive effects on tax compliance in other contexts (for recent reviews, see Pomeranz and Vila-Belda, 2018; Slemrod, 2018; Alm, 2019).<sup>3</sup> One standard interpretation in this literature is that taxpayers react to the information about tax enforcement tools and, in line with *AES*, reduce their evasion to re-optimize their behavior. However, there is no direct evidence in favor of or against this interpretation. Our contribution is to fill this gap in the literature.

This paper is closely related to a group of studies testing the predictions of *AES* in a laboratory setting. For example, Alm et al. (1992) conducted a laboratory experiment in which undergraduate students play a tax evasion game. Subjects can hide income from the experimenter, but some subjects are randomly selected to be audited and must pay a penalty if they are caught evading. The authors show that tax compliance in the game increases significantly with audit and penalty rates, but these effects are economically small and smaller than those predicted by optimizing behavior in the context of *AES*. The laboratory experiment setting of Alm et al. (1992) and similar studies have a number advantages, such as their full control over the rules of the game and freedom in the selection of the model parameters. However, these laboratory experiments have two main limitations. First, the subjects are typically undergraduate students playing the tax game for the first time and with no prior experience of paying taxes in the real world. In contrast, subjects in our field experiment are experienced firm owners who have been registered with the tax agency, and thus paying taxes, for an average of 15 years. Second, subjects from laboratory experiments typically pay taxes in the game amounting to less than USD 10. In contrast, subjects in our field experiment paid USD 11,800 per year in taxes on average, which is in the same order of magnitude as the country's GDP per capita.<sup>4</sup> We contribute to this literature in two ways. First, we show that *AES* does not fare substantially better in a natural context with experienced subjects and high stakes. Second, we show that audit threats may still be useful for the tax agency even if the individuals do not react optimally to them.

Our findings also contribute to the more general debate about the determinants of tax

---

<sup>3</sup>The following are some examples: Slemrod et al. (2001); Kleven et al. (2011); Fellner et al. (2013); Pomeranz (2015); Castro and Scartascini (2015); Dwenger et al. (2016); Perez-Truglia and Troiano (2018).

<sup>4</sup>More specifically, the firms in our sample paid an average of USD 7,770 in VAT and USD 4,070 in other taxes in the 12 months before our experiment. In comparison, the GDP per capita in Uruguay was about USD 15,000 in 2015.

compliance. One of the main puzzles in the literature is that evasion rates seem too low, given the low detection probabilities and penalty rates, especially among smaller firms and self-employed individuals (Luttmer and Singhal, 2014). One traditional explanation for this puzzle is based on tax morale: firms and individuals do not evade taxes because they feel morally obliged to comply (Luttmer and Singhal, 2014). Our evidence suggests an alternative explanation for the puzzle: due to the emotional nature of the decision, taxpayers overreact to the threat of audits. In other words, audits may scare taxpayers into compliance in the same way that scarecrows scare birds. Indeed, this interpretation can provide an explanation for the paradox that, despite the low audit probabilities and penalty rates, most taxpayers report the threat of audits as a major reason for why they report their taxable income truthfully (United States Internal Revenue Service, 2018).

The paper is organized as follows. Section 2 discusses the experimental design. Section 3 presents the data sources and discusses the implementation of the field experiment. Section 4 presents the results on the average effect of the *audit-statistics* message, and Section 5 presents the two tests of *A&S*. Section 6 discusses the interpretation of the findings. The final section concludes.

## 2 Experimental Design

Our experiment consisted of a mailing campaign from Uruguay’s IRS with multiple treatment arms and sub-treatments. Rather than comparing firms that received a letter to firms that did not, all of our analyses are based on comparisons between firms that received letters with subtle variations in their content. We can thus control for the potential effects of simply receiving a letter from the tax authority, which might induce compliance on its own – for instance, as a reminder to report taxable income.

The letters consisted of a single sheet of paper with the name of the recipient in the header, the official letterhead of the IRS, and the scanned signature of the IRS General Director. These letters were folded, sealed in an envelope with the official letterhead of the IRS on the outside, and sent by certified mail, which guarantees direct delivery to the recipient, who must sign upon receipt. Panel (a) of Figure 1 presents a diagram with the sample sizes for the different treatment arms that are detailed below.

### 2.1 *Baseline Letter*

The *baseline* letter contained some information about the goals and responsibilities of the tax authority, which the IRS routinely includes in its communications with firms. It explained

that the individual was randomly selected to receive this information, that the letter was for informational purposes only, and that there was no need to reply or to provide any documentation to the IRS. Figure 2 provides a sample of the *baseline* letter, with the addition of a placeholder box with the word “MESSAGE” written inside it.<sup>5</sup> This box was empty in the *baseline* letter but it contained a different message (printed in larger type size and boldface) in each of the other letter types.

## 2.2 *Audit-Statistics* Letter

According to the Allingham and Sandmo (1972) model, we expect risk-averse firms to be interested in information about the audit process, because it helps them optimize their evasion decisions and potentially increase their bottom line.<sup>6</sup> Furthermore, these figures should be particularly valuable in a context where information about audits is limited. For instance, it is easy to find online data about factors potentially relevant for firms’ decision-making, such as prices, inflation and exchange rates. However, information about tax audit probabilities (and, to a lesser extent, actual penalties paid by evading firms) is much harder to come by. Tax authorities seem to prefer to conceal this information.

In the *audit-statistics* letter type, we added the following paragraph to the *baseline* letter that provided information about the audit probabilities ( $p$ ) and penalty rates ( $\theta$ ) for a random sample of firms that were similar to the recipient:<sup>7</sup>

“On the basis of historical information on similar businesses, there is a probability of [ $p\%$ ] that the tax returns you filed for this year will be audited in at least one of the coming three years. If, pursuant to that auditing, it is determined that tax evasion has occurred, you will be required to pay not only the amount previously unpaid, but also a fee of approximately [ $\theta\%$ ] of that amount.”

Note that we communicated the probability that firms will be audited in at least one of the three following years, because IRS experts stated that this was the relevant probability for firms’ decision-making. Uruguay’s tax law indicates that tax audits should cover the previous three years of tax returns and, as a result, the probability that the current year’s tax report

---

<sup>5</sup>For a full-page sample of the letter without this placeholder, see Appendix A.1. For the corresponding samples of the *audit-statistics*, *audit-threat*, *audit-endogeneity* and *public-goods* letter types, see Appendices A.2, A.3, A.4, and A.5, respectively.

<sup>6</sup>We assume that firms in our sample are risk-averse, which is plausible since we deal mainly with small- and medium-sized firms. However, *A&S* has also been generalized to settings with risk-neutral agents (Reinganum and Wilde, 1985; Srinivasan, 1973).

<sup>7</sup>To make the information on the audit probability and the penalty rate clear and salient, we provided all figures as round numbers.

will be audited is roughly equal to the probability that the firm will be audited at least once over the following three years.

In our sample, the average value of  $p$  is 11.7%, and the average value of  $\theta$  is 30.6%. Tax agencies in most countries do not publish data on the values of  $p$  and  $\theta$ , which makes it difficult to compare the Uruguayan case to other contexts. In the United States, for which some comparable data are available, these two parameters are on the same order of magnitude: self-employed individuals face a  $p$  of 11.4% and a base  $\theta$  of 20%.<sup>8</sup>

The goal of this treatment arm was to generate exogenous variation in the firms' perceptions about audit probabilities and penalty rates. Because of legal considerations and other constraints, we could not assign different firms to different sets of information about these factors. We instead induced non-deceptive, exogenous variation in messages that may affect these perceptions by exploiting the sampling variation in statistics about audits and penalties.

More specifically, we divided the firms into five groups of "similar firms," corresponding to the five quintiles of total VAT payments in the fiscal year before our intervention. For each firm, we then drew a random sample of 50 other firms from the same quintile (i.e., similar firms), from which we computed the averages of  $p$  and  $\theta$ . This randomization strategy generated 940 different combinations of  $p$  and  $\theta$ . These estimates of  $p$  and  $\theta$  were unbiased and consistent with the explanation given in a footnote that we included in the letter, meaning that the information provided to recipients was nondeceptive. The footnote explained how we estimated the values of  $p$  and  $\theta$ :

"Estimates are based on data from the 2011–2013 period for a group of firms with similar characteristics, for instance, in terms of total revenue. The probability of being audited was calculated as a percentage of audited firms in a random sub-sample of firms. The rate of the fee was estimated as an average of a random sub-sample of audits."

The values of  $p$  ranged from 2% to 25%, with an average of about 11.7%. The values of  $\theta$  ranged from 15% to 68%, with an average of about 30.6%. Figure 3 presents the audit probability and penalty rate distribution across five groups by firm size (one in each row) and the distribution of the generated within-group parameters. The vertical line denotes the average audit probability or penalty rate based on all members of the group. If we based

---

<sup>8</sup>First, there is a 2.1% probability of being audited in any given year, according to the ratio of returns examined for businesses with no income tax credit and with a reported income between USD 25,000 and 200,000 (Table 9a of IRS, 2014). Each audit covers the previous 3 to 6 years, which implies that the probability that the current year's tax filing will be eventually audited ranges from 5.88% to 11.42%. Second, IRS usually imposes a basic penalty of  $\theta=20\%$ , although the penalties can be higher depending on the specific situation.

our estimates of  $p$  and  $\theta$  included in the letter on the population of firms, every member of the group would have received the same signal (the vertical line). As we computed  $p$  and  $\theta$  from samples of 50 firms, the sampling variation implies that different members of each group received different signals. For example, panel (a.1) of Figure 3 shows that in group 1 (i.e., the first quintile of firms ranked by total VAT payments), the average  $p$  for all group members is 8.2%, whereas the histogram depicts the different signals actually sent to firms within the group. These signals center around the average of 8.2%, but they range anywhere from 2.5% to 25%.<sup>9</sup>

### 2.3 *Audit-Threat Letter*

To complement the evidence from the *audit-statistics* sub-treatment, we implemented an alternative way of randomizing perceptions about audit probabilities using an *audit-threat* letter. We devised a treatment arm that randomly assigned firms to groups with different probabilities of being audited in the following year.

The *audit-threat* letter was identical to the *baseline* letter, with the following additional paragraph:

“We would like to inform you that the business you represent is one of a group of firms pre-selected for auditing in 2016. A [ $X\%$ ] of the firms in that group will then be randomly selected for auditing.”

This *audit-threat* treatment arm was applied to a separate experimental sample, a group of high-risk firms selected by the IRS audit department. The recipients of the *audit-threat* letter thus cannot be compared to recipients of the *baseline* letter. Instead, we randomly assigned the firms in this treatment arm to two groups, one with a 25% probability of being audited in the following year ( $X=25\%$ ) and another with a 50% probability of being audited ( $X=50\%$ ). These messages were non-deceptive: the IRS audit department provided a commitment that they would conduct audits according to these probabilities in the following year.

---

<sup>9</sup>The within-group average  $p$  is different across each of the five different groups, increasing monotonically from 8.2% in the bottom quintile to 13.4% in the top quintile. This implies that some of the variation in the values of  $p$  and  $\theta$  included in the letter was non-random. To estimate the causal effects of the signals  $p$  and  $\theta$ , we must isolate the random variation when analyzing the data. In any case, this aspect of the design is not so important in practice, as most of the variation in signals is indeed due to the sampling variation: for example, regressing  $p$  on the pre-treatment VAT quintile dummies results in  $R^2 = 0.118$ ; and regressing  $\theta$  on the same dummies results in  $R^2 = 0.007$ .

## 2.4 *Audit-Endogeneity Letter*

The *audit-statistics* and *audit-threat* treatment arms conveyed quantitative information about audit probabilities and penalty rates. We also wanted to incorporate into our research design a message about a different aspect of the audit process. Most tax agencies, including Uruguay’s, account for firm characteristics when deciding which ones to audit. They assign higher audit probabilities to firms with higher evasion risk. As a result, evading taxes typically increases the probability of being audited. This factor was incorporated as a special case in *A&S*, in which audit probabilities were determined endogenously. If unsuspecting firms learn about the endogenous nature of their audit probabilities, they should revise their tax evasion decisions and reduce the amount of tax evaded.<sup>10</sup>

We used this insight from economic theory to devise the *audit-endogeneity* message about the nature of the audit process. We asked our counterparts at the IRS to use their evasion-risk scores to divide a small sample of firms into two groups: those suspected of evading taxes and those not suspected of evading taxes. We then computed the difference in audit rates from 2011–2013 between the two groups: the rates were approximately twice as high for the high evasion risk group. We used this information to create the message in the *audit-endogeneity* letter type, which was identical to the *baseline* letter with the addition of the following paragraph:

“The IRS uses data on thousands of taxpayers to detect firms that may be evading taxes; most of its audits are aimed at those firms. Evading taxes, then, doubles your chances of being audited.”

## 2.5 *Public-Goods Letter*

We also devised a treatment arm to provide a benchmark for the effect of messages intended to increase tax compliance without directly mentioning audits. We designed a non-pecuniary message based on the suggestions by the IRS staff and authorities (i.e., on what they expected to be most effective at increasing compliance). This message provided information about the cost of evasion in terms of the provision of public goods, in the spirit of the model of Cowell and Gordon (1988).<sup>11</sup> The *public-goods* letter is identical to the *baseline* letter, with the

---

<sup>10</sup>Konrad et al. (2016) present suggestive evidence of this mechanism in the context of a laboratory experiment, finding that taxpayers who face a situation where suspicious attitudes toward tax officers increase the probability of being audited increase their tax compliance by 80%.

<sup>11</sup>This message is also related to the laboratory experiment from Alm et al. (1992), which finds that one of the reasons for why people decide to pay taxes is their valuation of the public goods provided by means of the tax revenues.

addition of the following message:<sup>12</sup>

“If those who currently evade their tax obligations were to evade 10% less, the additional revenue collected would enable all of the following: to supply 42,000 portable computers to school children; to build 4 high schools, 9 elementary schools, and 2 technical schools; to acquire 80 patrol cars and to hire 500 police officers; to add 87,000 hours of medical attention by doctors at public hospitals; to hire 660 teachers; to build 1,000 public housing units ( $50m^2$  per unit). There would be resources left over to reduce the tax burden. The tax behavior of each of us has direct effects on the lives of us all.”

## 2.6 Survey Design

We designed a survey to be conducted with a sample of owners from our main subject pool several months after they received the letters. The IRS, with the support of the Inter-American Center of Tax Administrations and the United Nations, had previously administered a survey on the costs of tax compliance for small- and medium-sized firms. We collaborated with the tax authority on the design and implementation of a new survey, which included a specific module tailored to our research design. The survey also included seven additional modules, designed by the IRS, about the costs of tax compliance and other topics.<sup>13</sup> We partnered with local and international universities to increase respondent confidence and to highlight the fact that the survey was part of a scientific study, and not an audit or compliance exercise by the IRS.

To further ensure trustworthy responses, the IRS assured potential respondents that the survey responses would remain anonymous and that they could not be traced back to specific individuals or firms. To measure the effect of our experiment on these survey responses, we embedded a code in the survey link to identify the treatment arm of the experiment which the recipient was assigned to. These codes did not uniquely identify any single firm, but they allowed us to link treatment arms and completed questionnaires while maintaining the anonymity of responses.

In our survey module, we assessed whether the *audit-statistics* message shifted the perceptions of the recipients of our letters by means of the following two questions:<sup>14</sup>

---

<sup>12</sup>The content of the message was based on estimates from the following governmental agencies: Administracion Nacional de Educacion Publica (ANEP), CEIBAL, Ministerio de Salud Publica (MSP), Ministerio del Interior (MI), Ministerio de Vivienda, Ordenamiento Territorial y Medio Ambiente (MVOTMA).

<sup>13</sup>A sample of the email with the invitation to participate in the online survey is presented in Appendix A.6.

<sup>14</sup>A screenshot of our survey module is presented in Appendix A.7.

Perceived Audit Probability: “In your opinion, what is the probability that the tax returns filed by a company like yours will be audited at least in one of the next three years (from 0% to 100%)?”

Perceived Penalty Rate: “Let us imagine that a company like yours is audited and that tax evasion is detected. What, in your opinion, is the penalty (in %) as determined by law that the firm must pay in addition to the originally unpaid amount? For example, a fee of X% means that, for each \$100 not paid, the firm would have to pay those original \$100 plus \$X in penalties.”

After each question, we elicited how certain the subject felt about his or her response on a 1–4 scale, from “Not sure at all” (1) to “Very sure” (4).<sup>15</sup>

### 3 Data Sources and Implementation of the Field Experiment

#### 3.1 Institutional Context

Uruguay is a South American middle-income country (the annual GDP per capita was about USD 15,000 in 2015). Our main focus in terms of studying tax evasion is the VAT, which represents the largest tax liability for firms in Uruguay and also the largest source of tax revenue. At the time of the study, the VAT rate was 22%,<sup>16</sup> and VAT revenues accounted for nearly half of the total tax revenues.<sup>17</sup> Uruguay is not atypical in terms of tax evasion. According to estimates from Gomez-Sabaini and Jimenez (2012), evasion of VAT in Uruguay was around 26% in 2008. This is the third-lowest rate among the nine Latin American countries included in the study and comparable to evasion rates in more developed economies. For example, the evasion rate for Italy in 2006 was estimated as 22% (Gomez-Sabaini and Moran, 2014).<sup>18</sup> Uruguay is not an outlier in terms of tax morale either. According to data from the 2010–2013 wave of the World Values Survey, 77.2% of respondents from Uruguay stated that evading taxes is “Never Justifiable,” whereas this proportion is 68.2% for all other

---

<sup>15</sup>For the sake of completeness, we also included a question in the survey to measure the subject’s awareness of the endogeneity of audit probabilities. You can find a screenshot of this question in Appendix A.7.

<sup>16</sup>A small number of products considered basic necessities either had a 10% rate or were exempt from the tax.

<sup>17</sup>Own calculations based on data from the Central Bank of Uruguay and from the Internal Revenue Service. Other sources of tax revenues include the personal income tax, the corporate income tax, and some specific taxes of consumption, businesses and wealth.

<sup>18</sup>Gomez-Sabaini and Jimenez (2012) compute those rates by applying the “indirect” method to estimate tax evasion. This method is based on the comparison of collected VAT with aggregate consumption data from the System of National Accounts (SNA).

Latin American countries on average (weighted by population) and 70.9% for the United States.

In some contexts, tax authorities do not need to rely on audits to mitigate tax evasion. For example, the U.S. Internal Revenue Service uses their electronic records to compare the wage amount reported by a taxpayer to the amount reported by his or her employer. This algorithm automatically rectifies the discrepancies in reporting and sends a notification to the taxpayer with the updated tax amount to be paid. Because the evasion will be caught through third-party reporting regardless of whether the individual is audited or not, the probability of being audited should be irrelevant for taxpayers (Kleven et al., 2011). Instead, we focused on a context in which tax authorities must still rely heavily on the threat of audits to discourage evasion. There is some third-party reporting for the VAT, consisting of the paper trail of invoices for sales and purchases.<sup>19</sup> However, this type of third-party reporting is highly imperfect. Most importantly, there is no automatic cross-checking and rectification of VAT payments. This limitation arises because the paper trail is non-electronic and thus can only be scrutinized by the tax agency in the event of an audit.<sup>20</sup> Additionally, the VAT paper trail has other limitations documented in the literature; for example, it breaks down at the consumer end (Pomeranz 2015; Naritomi 2019). The tax agency has access to other enforcement tools, such as tax withholding, but those other tools have limitations too. As a result, audits are still one of the main ways in which the tax agency can detect tax evasion in our context (Gomez-Sabaini and Jimenez, 2012; Bergman and Nevarez, 2006).

### 3.2 Subject Pool and Randomization

Our experiment was conducted in collaboration with the IRS of Uruguay. As of May 2015, there were 120,142 firms registered in the agency’s database. A sub-sample of 4,597 firms, pre-selected by the IRS, was put aside for the *audit-threat* sample, which we call the secondary experimental sample. We followed a series of criteria to select our main experimental sample from the remaining firms. First, we excluded some firms on request of the IRS. For instance, we excluded very small or very large firms that were subject to special regimes for VAT payments. We also restricted the experimental sample to firms that had made VAT payments in at least three different months in the previous 12 month period and firms with a total

---

<sup>19</sup>Firms may credit VAT paid on input costs (i.e., imports and purchases from their suppliers) against the total sales of goods and services to their costumers (i.e., “tax debit”). They pay VAT to the IRS only on the excess of the total “tax debit” over the tax credit. If the tax credit exceeds the debit, the excess may be carried over for future tax years. While the VAT should in theory be similar in its effects to a retail sales tax, in practice the two types of taxes differ in some significant aspects (Slemrod, 2008).

<sup>20</sup>In other countries, the use of standardized electronic invoicing systems may facilitate and automatize the cross-checking of the VAT trail to detect evasion.

value of at least USD 1,000.<sup>21</sup>

To maximize the impact of our information provision experiment, we did our best to ensure that the letters would be delivered to the firms' owners.<sup>22</sup> Moreover, in very large firms, the effect of the information could be substantially diluted, as it may not reach the owner or the individuals making decisions about tax compliance. Thus, we excluded from our subject pool firms with a total value exceeding USD 100,000 during the previous 12 months.

These criteria left a subject pool of 20,471 firms for the main experimental sample. All firms were randomly assigned to receive one of the four letter types according to the following distribution: 62.5% were assigned to the main treatment arm (*audit-statistics* letter), and 12.5% were assigned to each of the three remaining letter types (*baseline*, *audit-endogeneity*, and *public-goods*).<sup>23</sup> After removing the 19.9% of letters that were returned by the postal service, the final distribution of letter types was as follows: 10,272 received *audit-statistics*; 2,064 received *baseline*; 2,039 received *audit-endogeneity*; and 2,017 received *public-goods* letters (total N = 16,392). The 4,597 firms in the secondary sample were assigned to receive the *audit-threat* letter. Half were randomly assigned to the message of a 25% audit probability, and the other half to the 50% audit probability. After excluding the 12% of letters returned by the postal service, we were left with 2,015 firms in the 25% probability group and 2,033 firms in the 50% probability group (total N = 4,048).

Table 1 allows us to compare the balance of pre-treatment characteristics between firms assigned to the different letter types. Columns (1) through (4) correspond to firms in the main experimental sample. For each characteristic, column (5) presents the p-value of the test of the null hypothesis that the averages for these characteristics are the same across all four letter types. As expected, the differences across letter types are economically small and statistically insignificant. Columns (6) through (8) of Table 1 present a similar balance test for the secondary sample used for the *audit-threat* arm. Again, the characteristics are balanced across the two sub-treatments in the *audit-threat* treatment arm.<sup>24</sup>

---

<sup>21</sup>The sample selection was conducted in May 2015, so this 12 month period spans from April 2014 to March 2015.

<sup>22</sup>In some cases, owners provide the address of external accountants instead of their own or their firm's. We removed from the sample firms that were registered with an accountant's mailing address as their own (the IRS keeps records of addresses for all registered accountants).

<sup>23</sup>The randomization of letter types was stratified by the quintiles of the distribution of VAT payments over the fiscal year before our intervention.

<sup>24</sup>Appendix B.1 provides descriptive statistics for the firms in our subject pool. Moreover, Appendix B.2 shows that the rate of non-delivered letters are mostly balanced across treatments, with only minor and not economically significant differences in missing delivery status for the *public-goods* and *audit-endogeneity* treatment arms with respect to *audit-statistics* and *baseline*.

### 3.3 Outcomes of Interest

The letters were handed to Uruguay’s postal service on August 21, 2015. The vast majority of the letters were delivered to taxpayers during the month of September, and therefore we set September as the last month of the pre-treatment period and October as the first month of the post-treatment period. The main outcome of interest in our study is the total amount of VAT liabilities remitted by taxpayers in the 12 months after receiving the letter.<sup>25</sup> To test for the persistence of our treatment effects, we define a second period of observation between October 2016 and September 2017 (i.e., up to two years after the intervention). Panel (b) of Figure 1 depicts a timeline of the experiment and the data collection.

VAT represented 64.7% of total taxes paid by these firms in the fiscal year that preceded our treatment. The corporate income tax represented 25.4% of total taxes paid, the wealth tax 6.5%, and the personal income tax withholding only 3.3%. In this context of sole proprietorships, micro enterprises and small enterprises, the VAT is the firms’ main tax liability, which is why it is our main focus. Nonetheless, we also obtained data on the other taxes paid by the firms from the IRS, which allows us to assess whether firms effectively changed their overall tax compliance or if they simply substituted the evasion of VAT for that of other taxes.

It should be noted that the firms in our sample are mostly small. On average, the total amount of VAT paid by firms that received the baseline letter in the 12 month pre-treatment period was about USD 7,700, whereas the amount for the corresponding post-treatment period was approximately USD 6,500. This negative trend in VAT payments can be explained by the fact that this sample contains a high share of small firms with a high turnover rate. The size of post-treatment VAT payments varied substantially, ranging from USD 400 at the 10th percentile to USD 16,550 at the 90th percentile.<sup>26</sup>

We can further break down firms’ VAT payments according to their timing. We can observe the date of transfer to the IRS as well as the month for which the payment was imputed. Firms can back-date payments to cover liabilities from previous periods. As firms typically make VAT payments on a monthly basis, they normally cover the current and previous months, which we call concurrent payments. We classified payments covering liabilities incurred two or more months ago as retroactive payments – that is, adjustments for revisions in past liabilities. About 99.4% of firms made at least one concurrent payment in the 12 month pre-treatment period, whereas only 23.8% of firms made at least one retroactive

---

<sup>25</sup>This variable includes all VAT payments, including direct VAT payments and indirect VAT withholdings.

<sup>26</sup>Appendix B.3 presents detailed descriptive statistics about the distribution of pre- and post- treatment payments for firms that received the *baseline* letter type.

payment over this same period.<sup>27</sup>

### 3.4 Survey Implementation

The IRS communicates mainly by postal mail, and therefore it has mailing addresses for all registered firms. It also keeps records of email addresses for a subset of firms that have used their online services. We emailed invitations to all firms in the main experimental sample with a valid email address (N=3,867). While we wanted to roll out the survey shortly after the mail interventions, for reasons beyond our control we were not able to roll it out until May 2016, nine months after the intervention.<sup>28</sup> We find that firms which we invited to the survey were similar in characteristics to the broader set of firms in the main experimental sample.<sup>29</sup>

Our purpose was to elicit the beliefs of firm owners. We did not include email addresses that were repeated more than three times in the full sample, as these most likely belonged to accounting firms representing multiple small- and medium-sized firms. Even after applying this criterion, the IRS records could not ensure that the registered email address belonged to the firms' owners. We thus asked the survey respondent to self-identify as one of the following five types: owner, internal accountant, external accountant, manager, or other employee. From the 3,867 recipients that we invited to participate in the survey, 948 started to answer the survey (response rate of 24.5%).<sup>30</sup> Of these 948, 68.9% self-identified as an owner, 23.5% as a non-owner, and the remaining 7.6% did not provide a response to this question.<sup>31</sup> Our baseline specification excludes respondents who self-identified as non-owners, but the results are similar if we include only those who actively identified as owners.<sup>32</sup> Per an IRS request, respondents could skip as many questions as they wanted. We find that 6.6% and 8.6% of respondents skipped the audit probability and penalty questions respectively, which is comparable to the average rate (6.1%) at which they skipped other questions in the

---

<sup>27</sup>It should be noted that the retroactive payments do not reflect delinquency or outstanding debts to the tax authority. Overall VAT liabilities are computed on a yearly basis, and firms make monthly payments according to their provisional receipts on a pay-as-you-go basis to avoid a large bill at the end of the fiscal year. Thus, retroactive payments reflect changes in past liabilities. For instance, a firm may have "forgotten" to declare a sale in the past and thus need to send a retroactive payment corresponding to the gap between the original and the updated accounting.

<sup>28</sup>While we would have preferred a shorter turnaround between the experiment and the survey, several departments of the tax authority were involved in its design, which created some delays.

<sup>29</sup>Results reported in Appendix B.1.

<sup>30</sup>In this calculation, we require that respondents had answered at least the first two questions of the survey.

<sup>31</sup>The non-owner responses are distributed as follows: 6.1% self-identified as an internal accountant, 8.3% as an external accountant, 2.7% as a manager, 6.3% as other employee.

<sup>32</sup>Results reported in Appendix B.4.2.

survey.<sup>33</sup>

## 4 Results: Average Effect of Messages

### 4.1 Effect of the *Audit-Statistics* Message

Our first set of hypotheses concern whether providing letters with information related to tax enforcement increases tax compliance. We start by describing the effects of our main treatment, the *audit-statistics* message. Findings in previous literature suggest that this message is expected to have a positive effect on tax compliance because it teaches the taxpayer about tax enforcement.<sup>34</sup>

Figure 4 summarizes the raw data before conducting any regression analysis. Panel (a) of Figure 4 corresponds to the effect of the *audit-statistics* message. More precisely, this graph shows the percentage difference in VAT paid between the individuals who were randomly assigned to the *audit-statistics* letter and the individuals who were assigned to the *baseline* letter.<sup>35</sup> This figure shows the difference for each bi-monthly period covering all the months that we have data for, including three pre-treatment years (October 2012 to July 2015) and two post-treatment years (October 2015 to September 2017).<sup>36</sup> By construction, period 0 is defined as the period during which the letters were being delivered (August-September, 2015), which is highlighted in the figure with the vertical dashed line. Following Pomeranz (2015), we normalize by the average difference during the entire pre-treatment period.

Panel (a) of Figure 4 shows that the *audit-statistics* message had economically significant effects on VAT payments. Prior to the delivery of the letters, and due to random assignment, we would not expect to see any differences between individuals who were assigned to one type of letter versus another. In other words, individuals cannot be possibly affected by messages that they have not received yet. As expected, the differences in VAT payments between the

---

<sup>33</sup>This skip rate is the probability of providing a missing answer conditional on reaching that question in the survey. Appendix B.4 presents a series of robustness tests and analyses of differential response rates by treatment group, among other tests.

<sup>34</sup>As long as the enforcement messages do not affect the taxpayers' true income, the changes in the total amount of VAT paid measure changes in tax evasion. However, given the presence of real effects, our estimates provide a lower bound for the impact of information about enforcement on tax compliance. Although real effects are possible in our setting, most of the public finance literature provides evidence that real effects are normally zero or small relative to reporting effects (see for example Saez et al., 2012).

<sup>35</sup>For more details, Appendix B.5 discusses the evolution of the VAT payments for the treatment and control groups separately.

<sup>36</sup>Since a number of firms are required to pay VAT bi-monthly and there is a strong seasonal pattern, we group the data into bins of two months. This is balanced across treatment and control groups and does not affect our results at all. In all the results the amounts are top-coded at the 99.99% percentile to avoid the contamination of the results by outliers.

*audit-statistics* and *baseline* letter recipients hover around zero in the pre-treatment period. After the letter delivery, our hypothesis predicts that there will be a positive wedge between individuals in the *audit-statistics* and *baseline* letter groups. As expected, a positive gap in VAT payments between the two groups arises right after the letters are delivered. More precisely, in the first couple of months after the letter delivery, the difference in VAT payments between the *audit-statistics* and *baseline* letter recipients jumps to 10.4%, and then hovers between 4.0% and 10.5% in the rest of the first post-treatment year. During the second year, the effect becomes weaker over time.

While these results suggest that the *audit-statistics* message increased subsequent VAT payments, we need to use a more formal framework for statistical inference. Since we observe the outcome variable before and after the intervention, we can use that information to reduce the variance of the error term and thus gain statistical power through a difference-in-differences specification which compares treated firms to control firms and the pre-treatment period to the post-treatment period (McKenzie, 2012). With that goal in mind, we follow the econometric specification from Pomeranz (2015). Consider the sample of firms assigned to either the *baseline* letter or the *audit-statistics* letter. Let  $i$  index firms and  $t = \{1, 2\}$  denote time, where  $t = 1$  corresponds to the 12 months pre-treatment and  $t = 2$  corresponds to the 12 months post-treatment. Let  $Y_{it}$  be the outcome variable by taxpayer  $i$  in period  $t$  (e.g.  $Y_{i,2}$  could be the total VAT payments by firm  $i$  in the 12 months post-treatment period).  $D_i^1$  is a dummy variable that takes value 1 if  $i$  was assigned to receive the *audit-statistics* letter and 0 if it was assigned to receive the *baseline* letter instead. Let  $Post_t$  be a dummy variable that takes the value 1 if  $t = 2$  (i.e., after the letters were delivered) and 0 if  $t = 1$ . The regression of interest is the following:

$$Y_{it} = \alpha_0 + \gamma_1 \cdot D_i^1 \cdot Post_t + \alpha_1 \cdot D_i^1 + \alpha_2 \cdot Post_t + \epsilon_{it} \quad (1)$$

The coefficient of interest is  $\gamma_1$ , which measures the differential effect between the *audit-statistics* letter and the *baseline* letter. When the dependent variable is the amount of taxes paid, we estimate a log-linear model, also known as a Poisson regression. The reasons for using this specification are twofold. First and foremost, the Poisson regression allows effects to be proportional – indeed, the coefficients can be readily interpreted as semi-elasticities.<sup>37</sup> Second, the Poisson regression naturally accounts for the bunching at zero of the dependent variable. Note that the Poisson regression can be used for a continuous non-negative variable

---

<sup>37</sup>The Poisson model can be expressed as follows:  $\log(Y_X) = \alpha + \beta X + \varepsilon$ . The effect of a unit change in  $X$  can be re-expressed in log-units of the dependent variable,  $\beta = \log(Y_{X=x+1}) - \log(Y_{X=x})$ . Provided this coefficient is small enough, it can be approximated as a percent-change effect:  $\beta = \log(Y_{X=x+1}) - \log(Y_{X=x}) \approx \frac{Y_{X=x+1} - Y_{X=x}}{Y_{X=x}}$ .

and we do not have to rely on additional functional form assumptions such as equidispersion thanks to the quasi-MLE estimator.<sup>38</sup> In any case, we find that the results are robust to alternative regression models (OLS, Tobit and Probit).<sup>39</sup> Standard errors are always clustered at the firm level.

Table 2 presents the baseline regression results. We start with Panel (a), which corresponds to the comparison between the *audit-statistics* and *baseline* letters. In column (1), the dependent variable corresponds to the effect on VAT paid. The post-treatment coefficient corresponds to the effect during the 12 months after the delivery of the letter (October 2015 to September 2016), as measured by the coefficient  $\gamma_1$  from equation (1) above. The post-treatment coefficient of *audit-statistics* (in column (1) of panel (a)) is positive (0.070) and highly statistically significant (p-value = 0.001). This coefficient is economically significant too: it suggests that the *audit-statistics* message increased VAT payments in the 12 months after the intervention by about 7.0%.<sup>40</sup> To better grasp the magnitude of the effects, we can compare it to some basic benchmarks. The estimated average evasion rate for VAT in Uruguay is 26% (Gomez-Sabaini and Jimenez, 2012), and while the tax base is not necessarily comparable, it provides a benchmark: the 7% increase in VAT payments amounts to a reduction in the evasion rate of 27% ( $= \frac{7.0\%}{26\%}$ ). In terms of previous findings in the literature, the effects of our *audit-statistics* treatment are not directly comparable to those of the audit message from Pomeranz (2015) because the messages differed in content, and because the two studies cover firms from different countries and with different characteristics. Nevertheless, Table 4 from Pomeranz (2015) indicates that the deterrence letter in that study led to an increase in VAT payments of 7.6%, which is similar in magnitude and statistically indistinguishable from the 7.0% effect of our *audit-statistics* message. Moreover, our results are qualitatively consistent with a broader literature that finds effects of messages about enforcement on tax compliance in a variety of contexts: self-employed income in the United States (Slemrod et al., 2001), wage income taxes in Denmark (Kleven et al., 2011), individual TV license fees in Austria (Fellner et al., 2013), individual municipal taxes in Argentina (Castro and Scartascini, 2015), an individual church tax in Germany (Dwenger et al., 2016), and tax delinquencies in the United States (Perez-Truglia and Troiano, 2018).

Table 2 presents a number of robustness checks, which we discuss below. First, we present falsification tests in the spirit of event-study analysis. The pre-treatment coefficients from

---

<sup>38</sup>For more details, see for example Chapter 19 of Wooldridge (2010).

<sup>39</sup>The results from these robustness checks are presented in Appendix B.6.1.

<sup>40</sup>This percent-effect is based on the following approximation:  $\beta = \log(Y_{X=x+1}) - \log(Y_{X=x}) \approx \frac{Y_{X=x+1} - Y_{X=x}}{Y_{X=x}}$ . For simplicity, in the rest of the paper we interpret all of the Poisson coefficients using this same approximation. Alternatively, the exact percent-effect can be calculated exactly using the exponential transformation. For example, the 0.070 coefficient corresponds exactly to a 7.25% effect ( $= 100 \cdot (e^{0.070} - 1)$ ).

Table 2 are estimated with a specification identical to the one used for the post-treatment coefficients, except it uses a “placebo” date for the delivery of the letters: i.e., we pretend that the letters were delivered during August and September of 2014, and estimate the “effects” on the VAT paid in the subsequent 12 months (i.e., October 2014 to September 2015). Since the letters had not really been delivered on that date, we would expect the “effect” of the *audit-statistics* message to be close to zero and statistically insignificant. A finding to the contrary would be suggestive of potential concerns with the specification or the random assignment. As expected, the pre-treatment coefficients for the *audit-statistics* message (column (1) in panel (a) of Table 2) is close to zero (0.009), statistically insignificant (p-value=0.658), and as precisely estimated as the corresponding post-treatment coefficient.

Over time, individuals may forget the information conveyed in the letter, or it may become less salient. Individuals may also update their beliefs and perceptions for other reasons, for instance due to new events such as audits and information campaigns. To assess the persistence of the effects, column (2) in Table 2 replicates the analysis for the second year after the treatment (October 2016 to September 2017). Consistent with the arguments provided above, the effect of the treatment is half as large as the effect during the first year, and is no longer statistically significant. These estimates are consistent with the pattern of effects by quarter depicted in panel (a) of Figure 4, which shows that the effect falls gradually over time and is about half as large in the second year as in the first year. The timing of the effects is furthermore consistent with previous evidence on the effects of tax enforcement messages. For instance, Pomeranz (2015) shows that the effects of her main intervention were also substantially higher in the first 12 months after the intervention, fell substantially thereafter and almost disappeared by the 18th month.<sup>41</sup>

Table 2 also presents results for complementary outcomes. As discussed in the previous section, firms in Uruguay make payments for their current liabilities, but also for taxes that correspond to previous periods because they revise their accounts and correct past mistakes, or because they impute invoices that were not available at the time of the original payment. When firms that engage in tax evasion face a heightened threat of being audited, we can expect them to increase their tax payments (i.e., reduce their evasion) in the future, but we can also expect them to retroactively revise their payments for previous time periods to reduce or eliminate their past evasion. To shed light on this question, columns (3) and (4) in panel (a) of Table 2 split the effects during the first year between retroactive and concurrent payments. The *audit-statistics* message had an economically and statistically significant effect on both retroactive and concurrent payments: the coefficient corresponding to the *audit-statistics* message is 0.383 (p-value=0.006) for retroactive payments and 0.053

---

<sup>41</sup>See for example Figure 2 in Pomeranz (2015).

(p-value=0.012) for concurrent payments.<sup>42</sup>

We have so far established that firms in the *audit-statistics* treatment arms increased their VAT payments compared to recipients of the *baseline* letter. Our analysis focuses on VAT liabilities, which represent the largest fraction of tax payments by firms in our sample. However, our letters referred to taxes in general and did not specifically mention VAT or any other tax in particular. Given the presence of other tax liabilities, the effects we reported on VAT may not represent a net increase in tax payments: firms may increase their evasion (i.e., reduce their payments) of other taxes they are liable for, thereby crowding out payments or substituting evasion to other taxes. On the other hand, firms may need to declare higher income in order to declare higher VAT, and thus they could be obliged to pay more, rather than less, in non-VAT taxes. The results in columns (5) and (6) of Table 2 shed light on these issues. Column (5) presents the effects on all other taxes paid (mostly the corporate income tax). The effects on the payments of other taxes are as economically and statistically significant as those on VAT payments: the *audit-statistics* message had an effect of 8.6% (p-value=0.019) on other tax payments. Column (6) shows that the results are robust if we look at the effect on the sum of VAT and other taxes: the *audit-statistics* message increased this outcome by a statistically significant 7.3% (p-value<0.001).<sup>43</sup>

## 4.2 *Audit-Endogeneity* Message

The first benchmark for the *audit-statistics* message is the *audit-endogeneity* message, which is similar in that it provides information about the enforcement of taxes through audits. Panel (b) of Figure 4 shows the raw evolution of VAT payments in the *audit-endogeneity* treatment relative to the *baseline* treatment. The results suggest that the *audit-endogeneity* message also induced a significant increase in VAT payments, which was similar in timing and magnitude to the effects of the *audit-statistics* message depicted in panel (a) of Figure 4. For a more formal statistical analysis, the regression results are presented in panel (b) of Table 2.<sup>44</sup> The coefficient in column (1) indicates that the *audit-endogeneity* message increased

---

<sup>42</sup>The effect size is larger for retroactive payments than for concurrent payments, but these differences in magnitude have to be taken with a grain of salt because there are large differences in baseline levels between the two outcomes. For example, in the *baseline* letter group, firms paid an average of USD 300 in retroactive payments versus USD 7,370 in concurrent payments in the post-treatment period.

<sup>43</sup>Appendix B.7 presents a finer analysis that breaks down the effects of “other taxes” separately by each of its three components and a series of additional robustness checks, such as alternative estimation methods (Appendix B.6.1), alternative specifications of the dependent variable (Appendix B.6.1) and heterogeneity analysis based on firm characteristics such as size, age and sector (Appendix B.8). Overall, we find that the effects are qualitatively and quantitatively similar across the board.

<sup>44</sup>The results from panels (b) and (c) in Table 2 are based on a regression specification equivalent to the one from equation (1) above which was used to obtain the estimates in panel (a).

subsequent VAT payments by 7.1% (p-value of 0.009).<sup>45</sup> This 7.1% is similar in magnitude and statistically indistinguishable (p-value=0.950) from the corresponding 7.0% effect of the *audit-statistics* message reported in panel (a).<sup>46</sup> The effects of the *audit-endogeneity* message also pass the same robustness checks presented above for the effects of the *audit-statistics* message: for example, the “effect” on the pre-treatment year is close to zero (-0.005) and statistically insignificant (p-value=0.868), and the effects during the second post-treatment year are about half as large as the effects during the first year.

### 4.3 *Public-Goods* Message

Panel (c) of Figure 4 shows the effects of the *public-goods* message on VAT payments. The time series data suggest that the *public-goods* message had a positive effect on tax compliance too, but it dissipates a lot more quickly than the effects of the *audit-statistics* message. The corresponding regression results are presented in panel (c) of Table 2. The *public-goods* message increased VAT payments in the first post-treatment year by 5.1%, and the effect is statistically significant (p-value=0.043). On the other hand, the effects of the *public-goods* message in the second post-treatment year were close to zero (0.4%) and statistically insignificant (p-value=0.906). There is mixed evidence on the effects of moral messages, with evidence that they work in some contexts (Bott et al., 2020; Nathan et al., 2020; Hallsworth et al., 2017) but not in others (Blumenthal et al., 2001; Fellner et al., 2013; Castro and Scartascini, 2015; Dwenger et al., 2016; Meiselman, 2018; Perez-Truglia and Troiano, 2018). The most related study, Pomeranz (2015), included a message of moral suasion that had a positive effect on subsequent VAT payments, although it was not as large as the effect of the deterrence message and statistically insignificant. Our findings on moral suasion fall closest to the findings from the experiment with Norwegian taxpayers reported in Bott et al. (2020): they find that the message of moral suasion increased tax compliance in the short term, but the effects dissipated completely the following year.

## 5 Tests of A&S

The results presented in the previous section are broadly consistent with the evidence in the literature that providing information about audits significantly increases tax compliance.

---

<sup>45</sup>The effect of the *audit-endogeneity* message is somewhat less precisely estimated than the corresponding effect of the *audit-statistics* message, but that is expected due to the differences in sample sizes.

<sup>46</sup>This is an equality test between two coefficients based on the same data but different regressions. To allow for a nonzero covariance between these two coefficients, we estimate a system of seemingly unrelated regressions. In the remainder of the paper, when comparing coefficients from the same data but different regressions, we always use this method.

In this section, we present additional evidence to establish whether the effects of the *audit-statistics* treatment are driven by the *AES* mechanism.

## 5.1 First Test of *AES*: Effects on Perceptions

According to *AES*, the *audit-statistics* message should have a positive effect on tax compliance if it increased the perceived probability of being audited, the perceived value of the evasion fine, or both. We explore this hypothesis by utilizing data from our post-treatment survey. This survey data consists of 365 firms in the *audit-statistics* group and 137 in what we refer to as the pooled control group with individuals who did not receive information related to audits.<sup>47</sup>

Panels (a) and (b) of Figure 5 depict the distributions of perceptions about audit probabilities and penalty rates respectively, as elicited from the survey. The shallow bars with solid borders correspond to the perceptions of firms that received the *audit-statistics* message. The shaded gray bars depict the distribution of perceptions for firms in the pooled control group. The red dashed curve, in turn, corresponds to the distribution of signals sent to firms in the *audit-statistics* letters. The comparison between the shaded bars and the red curve from panel (a) of Figure 5 suggests that, on average, respondents in the control group substantially overestimated the probability of being audited. While our administrative data on audits indicates a probability of being audited of about 11.7%, the mean perception for the control group is 40.7% (p-value<0.001 for the difference). This finding of an overestimation of audit probabilities is consistent with prior survey evidence (Harris and Associates 1988; Erard and Feinstein 1994; Scholz and Pinney 1995).<sup>48</sup> In contrast, the comparison between the shaded bars and the red curve from panel (b) of Figure 5 suggests that the average belief about the penalty rates was unbiased: the actual average penalty computed from administrative data for the experimental sample is 30.6%, while the mean perceived penalty is 30.5% in the control group.

The positive bias in the perceived audit probability could be explained by the availability heuristic bias (Kahneman and Tversky, 1974). According to this model, individuals judge the probability of an event by how easily they recall instances of it. Even though audits are

---

<sup>47</sup>The survey sample size was substantially smaller than that of our experimental sample. To increase the statistical power of our test, we defined this control group by pooling subjects from the *baseline* and the *public-goods* groups, since both received messages with no specific information about audit probabilities or fines. Appendix B.4.2 shows that the results are similar, but less precisely estimated, when we only use recipients of the *baseline* letter for the control group.

<sup>48</sup>However, the prior survey evidence was based on responses from wage earners, for whom the misperception of audit probabilities is mostly inconsequential due to widespread third-party reporting (Kleven et al., 2011). On the contrary, the financial stakes of misperceiving audit probabilities can be substantial in our context.

rare, the fact that they may be salient in memory or discussed by colleagues and the media may induce firms to perceive them to be more frequent than they actually are. Indeed, there is evidence that individuals overestimate the probabilities of a wide range of rare events of a similar nature such as the probabilities of dying in a terrorist attack or in a plane crash (Lichtenstein et al., 1978; Kahneman et al., 1982).

The survey data indicates that the effects of the *audit-statistics* message are inconsistent with the *A&S* predictions. According to *A&S*, if taxpayers overestimate the audit probabilities on average, the *audit-statistics* message should have *reduced* the average tax compliance. On the contrary, the results presented in the previous section show that the *audit-statistics* message *increased* the average compliance.<sup>49</sup> To strengthen this argument, we show that the *audit-statistics* letter indeed had a negative effect on the perceived audit probability.<sup>50</sup> The shallow bars with solid borders in panel (a) of Figure 5 show the distribution of perceptions for respondents who received the *audit-statistics* letter. An inspection of panel (a) of Figure 5 indicates that recipients of the *audit-statistics* letter reported a lower perceived probability of being audited on average, from an average of 40.7% in the pooled control group to an average of 35.2% in the *audit-statistics* group (p-value of the difference 0.03).<sup>51</sup> The same mechanisms behind the *audit-statistics* message are also relevant for the interpretation of the *audit-endogeneity* message: subjects may have found out that the audits were endogenous through the *audit-endogeneity* message and re-optimized their tax evasion accordingly; or they could have had a knee-jerk reaction to any information about audits, even if they knew about the endogeneity component already. Consistent with the evidence on the *audit-statistics* message, we find that the effect of the *audit-endogeneity* message was probably not due to an update of recipients' beliefs, because recipients were already aware of this endogeneity.

---

<sup>49</sup>One caveat for the test presented in this section, and also for the test presented in the section below, is that we are estimating the effects on the average firm. The fact that the average firm does not behave as *A&S* predicts does not imply that none of the firms behave as *A&S* predicts. For instance, it is possible that some firms updated their perceived probability upwards because of the information contained in the letter and increased their tax payments as a consequence. However, the fact that the average effects are so far from the *A&S* prediction suggests that, if anything, the firms behaving as *A&S* predicted must have been a minority.

<sup>50</sup>One caveat to this interpretation is that a reduction in the self-reported probability of being audited could also be caused by an increase in tax compliance due to the endogenous nature of  $p$  with respect to tax evasion.

<sup>51</sup>Additionally, panel (b) of Figure 5 shows that the *audit-statistics* message had a small effect on the perceived penalty rate, decreasing it from an average of 30.5% for the pooled control group to an average of 29.9% for the *audit-statistics* group, with this difference being statistically insignificant (p-value of 0.82). For more details, see also Appendix B.4.3.

### 5.1.1 Addressing Concerns with Survey Data

This first test relies on survey data, and as such it is subject to some common challenges with this type of data. In this section, we discuss and address some of those challenges.

One potential concern is that the responses about audit probabilities and penalty rates mostly reflect measurement errors because the questions were not incentivized. There are several pieces of evidence suggesting otherwise. First, the fact that the survey beliefs changed depending on the information provided in the letters suggests that these responses contained some truthful information. Second, while there is a large positive bias in the perception of audit probabilities, the average perception of the penalty rate (30.5%) is extremely close to the actual probability computed from administrative data (30.6%). The fact that beliefs are so accurate for penalty rates suggests that individuals responded honestly and thoughtfully. Additionally, individuals were better informed about penalty rates than about audit probabilities, which is also consistent with the fact that there is more readily available information about penalty rates: audits are relatively rare events and their probabilities are not advertised and are not common knowledge, whereas evasion penalties are more openly disclosed by the tax agency.

Another potential issue is that respondents were aware that they were misinformed, so they would never have acted on their biased beliefs anyways. Our survey data provides evidence to the contrary. Even though their estimates were substantially off, survey participants reported being confident about their responses. For example, only 16.2% of those in the control group reported being “Not sure at all” about their perceived probability of audit (on a four point scale, ranging from “Not sure at all” to “Very sure”); and a similar share (18.1%) reported being “Not sure at all” about their guess of the penalty rate. Even for the subgroup of individuals from the control group who reported to be “Very sure” about the audit probability, their average belief was, if anything, slightly more biased: they reported a perceived audit probability of 42.1%, which is still substantially higher than the actual probability of 11.7%.

Another concern is that our subjects may be confused about the questions, for instance if they did not understand the definition of an “audit.” Among the 145 responses from the pooled control group, 10.3% of firms reported that they had been audited in the past three years. This share (10.3%) is close to the actual share of firms that were audited (11.7%), thus suggesting that respondents understood the definition of an audit correctly. Moreover, if firms use their own audit history to form their beliefs about auditing probabilities, the firms who had been audited recently should report a higher perceived probability of being audited. Indeed, we found that to be the case: firms who had recently been audited reported a substantially higher average perceived probability of being audited in the future (63.9%) than

firms that had not recently been audited (38.1% – p-value of the difference < 0.001). Likewise, one may worry that subjects have cognitive limitations when responding to questions about percentages and probabilities. However, this should be a minor concern in our subject pool, which is comprised of business owners who should be familiar with fractions and probabilities. While we do not have administrative data to verify this, the anecdotal evidence indicates that this is a highly educated subgroup of the population. And, at the very least, these business owners need some rudimentary arithmetic and understanding of percentages to compute the VAT and other tax liabilities.

Another potential concern is that some respondents may report a probability of 50% as a way of expressing their uncertainty (Bruine de Bruin et al., 2002; Bruine de Bruin and Carman, 2012). Responses of exactly 50% are somewhat common in our data: among individuals in the pooled control group, 41.61% of responses about the perceived audit probability and 13.5% of responses about the penalty rate are exactly equal to 50%. However, our data indicates that most of these responses of exactly 50% are not a product of uncertainty: individuals who provided an answer of 50% are somewhat less confident, but not dramatically less so, than individuals who provided answers different to 50%.<sup>52</sup> Moreover, even if we ignore the 50% responses, the main result would still be robust: individuals in the control group still substantially overestimate the probability of being audited (average perception of 31.5%, compared to the actual probability of 11.7%).<sup>53</sup>

As an additional validation of the survey data, we can measure the effect of the signals about  $p$  and  $\theta$  from the *audit-statistics* sub-treatments. However, we have limited power for this exercise due to the small sample size, as we only have 365 survey responses in the *audit-statistics* group. With that caveat in mind, the survey data suggests that a percentage point increase in the signal about the audit probability provided in the letter increased the perceived audit probability nine months later by 0.397 (SE 0.288) percentage points.<sup>54</sup> While imprecisely estimated and thus statistically insignificant at conventional levels, the magnitude of this point estimate is consistent with the updating in other information-provision experiments (discussed in more detail in Section 5.2.1 below). Most importantly, the true effects of the signals provided in the letter on beliefs were probably stronger than what the above estimates suggest, due to different sources of non-compliance.<sup>55</sup>

---

<sup>52</sup>In the pooled control group, the average certainty for perceived audit probability is 2.05 (in the 1-4 scale from “Not sure at all” (1) to “Very sure” (4)) for individuals who responded with a value of exactly 50%, and 2.57 for individuals who responded with a different value (p-value of difference=0.001). For the responses about the perceived penalty rate, the average certainty is 2.05 for individuals who responded with 50%, and 2.50 for those who responded with another value (p-value of difference=0.06).

<sup>53</sup>For more details, see Appendix B.4.2.

<sup>54</sup>For more details, see Appendix B.4.4.

<sup>55</sup>While we are confident that our certified letters reached the firms’ owners, we cannot be as confident

We further provide an alternative to this test that does not rely on survey data.<sup>56</sup> We assume that firms form prior beliefs about the audit probabilities based on their own exposure to audits, due to the paucity of information available about the auditing process. Take for instance two firms that have been paying taxes for 10 years and, by chance, one of those firms was audited in the past while the other was not. As a result, the firm that was audited in the past will have a higher perceived probability of being audited in the future. The results from this alternative test provide evidence against the *A&S* mechanism too.

## 5.2 Second Test of *A&S*: Heterogeneity with Respect to Signals

The second test is based on the differential effects of the values of the signals provided in the letters. According to *A&S*, the effects of the *audit-statistics* message should be increasing in the signals about the audit probability ( $p$ ) and the penalty rate ( $\theta$ ). The random variation we introduced in the  $p$  and  $\theta$  conveyed in our *audit-statistics* letters allows us to test this hypothesis directly. We first present our estimates of these elasticities, and then compare them with values obtained from calibrations of the *A&S* model.

### 5.2.1 Elasticities with respect to $p$ and $\theta$

For a less parametric look at the data, Figure 6 estimates the effect of the *audit-statistics* message on VAT payments, but broken down by decile of the signals included in the letter.<sup>57</sup> Panel (a) of Figure 6 presents the effect of the *audit-statistics* message by decile of the signal of  $p$  shown in the letter. In the *A&S* framework, we should expect that very low signals of  $p$  should reduce tax compliance (since they most likely reduce the firms' perceived probability of audits), whereas the effect should become larger, and turn positive at some point, as we increase the signal of  $p$ . The coefficients plotted in panel (a) of Figure 6, however, indicate that the effect of the *audit-statistics* letter is not related to the value of  $p$  included in the letter. The coefficients are similar in magnitude for the whole range of values from  $p = 2\%$

---

about whether the owner was the same person who received the email invitation to complete the survey. And while we wanted to conduct the survey shortly after the mailing campaign, we were not able to roll out the survey until nine months after the intervention for reasons beyond our control. In this type of information-provision experiment, the effect of information on beliefs tends to decay substantially in a matter of a few months – recipients may forget the information provided in the letters, or acquire additional information in the meantime. For example, Cavallo et al. (2017) show that the effect of information on beliefs decays by about half in a matter of just three months, and similar findings are reported by Bottan and Perez-Truglia (2017) and Fuster et al. (2018). All these factors will lead to an underestimation of the effects of the letter on beliefs.

<sup>56</sup>For more details, see Appendix B.9.

<sup>57</sup>These results are based on the same specification used for Table 2, except that it also includes dummies for the quintiles of pre-treatment VAT payments as additional controls, from which we drew the sample to calculate  $p_i$  and  $\theta_i$ .

all the way up to  $p = 25\%$ . Moreover, the resulting linear relationship (shown as a dashed red line) has a slope that is close to zero and statistically insignificant. Panel (b) of Figure 6 provides a similar analysis for the heterogeneity by penalty rates ( $\theta$ ) provided in the letter. According to  $A\mathcal{E}S$ , we should expect a positive relationship between the effect of the *audit-statistics* letter and the value of  $\theta$  included in the letter. Panel (b) of Figure 6 shows evidence to the contrary: the coefficients are similar for the whole range of values from  $\theta = 15\%$  to  $\theta = 68\%$ , and the slope is close to zero and statistically insignificant.

Using a more parametric approach, we can quantify the effects of the *audit-statistics* sub-treatments in a way that can be contrasted with the quantitative predictions of  $A\mathcal{E}S$ . For the *audit-statistics* treatment arm, we use the following model:

$$Y_{it} = \alpha_0 + \gamma_p \cdot p_i \cdot Post_t + \gamma_\theta \cdot \theta_i \cdot Post_t + \alpha_1 \cdot p_i + \alpha_2 \cdot \theta_i + \alpha_3 \cdot Post_t + \quad (2)$$

$$+ \sum_{g=2}^5 \alpha_{4,g} \cdot I_{\{i \in g\}} + \sum_{g=2}^5 \alpha_{5,g} \cdot I_{\{i \in g\}} \cdot Post_t + \epsilon_{it}$$

where  $p_i \in (0, 1)$  is the signal about the audit probability included in the letter sent to firm  $i$ , and  $\theta_i \in (0, 1)$  is the signal about the penalty rate included in the letter sent to the same firm. The  $I_{\{i \in g\}}$  variables correspond to a set of dummies for the quintiles of pre-treatment VAT payments, which are the groups from which we drew the sample of “similar firms” to calculate  $p_i$  and  $\theta_i$ . Including these controls ensures that we only exploit the exogenous variation in  $p_i$  and  $\theta_i$  – that is, the heterogeneity due to sampling variation.  $Post_t$  is a dummy variable that takes the value 1 if the observation corresponds to the post-experiment period, and 0 otherwise. Since we are using a Poisson regression model,  $\gamma_p$  and  $\gamma_\theta$  can be directly interpreted as elasticities. For instance, an estimate of  $\gamma_p = 1$  would imply that a 1 percentage point increase in the audit probability conveyed in the letters increased VAT payments by 1%.  $A\mathcal{E}S$  predicts that  $\gamma_p > 0$  and  $\gamma_\theta > 0$  – i.e., that firms’ tax payments are increasing in the perceived probability of audit and evasion penalty rates.

Panel (a) in Table 3 presents the results from the econometric model of equation (2). Column (1) of Table 3 presents estimates of the elasticities of VAT payments with respect to the values of  $p$  and  $\theta$  conveyed in the *audit-statistics* sub-treatments. The elasticity with respect to the audit probability in the first post-treatment year is -0.063 (p-value=0.796). This means that increasing  $p$  by 1 percentage point would decrease VAT payments by a mere 0.063%. The elasticity with respect to the penalty rate is -0.033 (p-value=0.782), which implies that increasing  $\theta$  by 1 percentage point would decrease VAT payments by 0.033%. The estimates are close to zero, statistically insignificant at standard levels and precisely estimated. The precision implies that we can rule out even moderate elasticities: the 95%

confidence interval for the audit probability excludes elasticities above 0.411, and the 95% confidence interval for the penalty rate excludes elasticities above 0.198.

It should be noted that the pre-treatment falsification test does not yield any statistically significant effects, and that the results are similar for the other specifications: for the second post-treatment year (column (2)), by payment timing (columns (3) and (4)), and by type of tax (columns (5) and (6)). The estimates in all cases are close to zero, precisely estimated and statistically insignificant.

One potential confounding factor for the lack of heterogeneity in signals about  $p$  and  $\theta$  is that some subjects might have interpreted the *audit-statistics* message *per se* as a signal that their firms were under the IRS radar, above and beyond the factual information conveyed in the message. We were careful to mitigate this concern in the design of our mailings. For instance, we highlighted the fact that the letter recipients were randomly selected. Nevertheless, some individuals may have ignored or overlooked this cue. Even if the recipients learned something from the receipt of the *audit-statistics* message, there is no reason why they should not learn about the content of the message as well. In other words, the test presented above continues to be valid, as *AES* would still predict that the *audit-statistics* message should have a differential effect depending on the values of  $p$  and  $\theta$ .

To address this concern more directly, we use the *audit-threat* treatment arm, in which the tax agency made an explicit threat to every recipient and thus is not subject to this concern. Panel (d) of Figure 4 depicts the difference in the evolution of VAT payments over time between the two sub-treatments in the *audit-threat* arm, corresponding to audit probabilities of 50% and 25%. We mostly find no systematic difference between the two groups in post-treatment VAT payments. We can also provide a more parametric test based on an econometric model similar to that of equation (2) for firms assigned to the *audit-threat* letter:

$$Y_{it} = \alpha + \tau_p \cdot p_i + \delta \cdot Post_t + \gamma_p (p_i \cdot Post_t) + \epsilon_{it} \quad (3)$$

where  $p_i \in \{0.25, 0.50\}$  is the audit probability included in the *audit-threat* letter sent to firm  $i$ . Again, *AES* predicts  $\gamma_p > 0$ . The results are presented in panel (b) of Table 3. While the estimated coefficient based on *audit-threat* messages implies an elasticity of 0.217 with respect to  $p$  in the first year post-treatment, this estimate is not statistically significant (p-value=0.128) and is economically small. Taking into account the precision and power concerns, the evidence from panel (d) of Figure 4 and panel (b) of Table 3 further reinforces our result. Contrary to the prediction of *AES*, tax compliance does not seem to depend on the probability of being audited, even with a direct and credible threat of an audit (rather than just providing information about audit probabilities).

Taken together, we find robust evidence that firms did not react to the values of  $p$  and  $\theta$  shown in our letters.<sup>58</sup> Since we devoted a large fraction of our subject pool to this treatment arm, these elasticities are precisely estimated. It is not clear, however, whether the estimates are small enough, and precisely estimated enough, to rule out the values of the elasticities predicted by *AES*. We address this question below.

### 5.2.2 *AES* Calibration

For a quantitative test of *AES*, we need to obtain quantitative predictions from *AES*. In this section we present results from different calibrations of the model, and compare the calibration results to the experimental results in the following section.

Let  $Y$  be the total value-added amount and let  $\tau = 0.22$  be the value added tax rate. Let  $E$  be the amount to be under-reported (so  $\tau \cdot E$  is the amount evaded). Each firm has a utility from income given by a Constant Relative Risk Aversion (CRRA) utility function with risk parameter  $\sigma$ . Let  $p$  be the probability that the tax return for a given year will be audited sometime in the future, and  $\theta$  the penalty rate applied over the amount evaded when caught (both of these parameters are defined as in the *audit-statistics* treatment).

Given any reasonable value for the CRRA parameter, the basic *AES* model predicts 100% evasion. As a result, we need to use one of the extensions discussed in the literature to accommodate the 26% evasion rate observed in practice (as estimated by Gomez-Sabaini and Jimenez, 2012). We consider the following extensions: endogenous audit probabilities (Allingham and Sandmo, 1972; Yitzhaki, 1987), third-party reporting and whistle-blowing (Acemoglu and Jackson, 2017), misperceptions about audit parameters (Alm et al., 1992) and social preferences (Luttmer and Singhal, 2014).

The probability of being audited can be broken down as  $p = p_0 + p_1 \frac{E}{Y}$ .<sup>59</sup> The parameter  $p_1 > 0$  represents the endogeneity of the audit process, whereby firms that evade more are more likely to be audited (in the original *AES* model, the audit probability is exogenous so  $p_1 = 0$ ). We can also allow firms to be caught evading due to some non-audit technology such as third-party reporting or whistle-blowing. We represent this with an effective probability of being caught of  $p + \epsilon$ , where the parameter epsilon represents the additional monitoring tool. To allow for misperceptions, we can simply calibrate  $p$  and  $\theta$  to take the average

---

<sup>58</sup>Appendix B.6.2 present a series of robustness checks (alternative specifications based on OLS, Tobit and Probit models, and using an alternative data source for the dependent variable), and the results are similar. An additional robustness check, presented in the same Appendix, shows that results are robust if, instead of estimating the elasticities with respect to  $p$  and  $\theta$  separately, we estimate the elasticity with respect to  $p * \theta$  (i.e., the expected penalty per dollar evaded).

<sup>59</sup>Note that we implicitly assume that, conditional on being audited, evasion is detected in full. In practice, this probability may be smaller than 1. If anything, this would only make the *AES* result more puzzling: firms should be even less worried about being audited.

perceptions reported in the survey (instead of the values calculated from the administrative records of the tax agency). To allow for social preferences, we assume that individuals get some direct utility from paying taxes equal to the fraction  $\alpha$  of the amount paid. This social responsibility parameter  $\alpha$  can take values from 0 to 1, where a higher value denotes higher social responsibility (in the original  $A\mathcal{E}S$ ,  $\alpha = 0$ ).

The optimal evasion choice is given by maximizing the expected utility:

$$\max_{E \in [0, Y]} \frac{1 - p \left( \frac{E}{Y} \right) - \epsilon}{1 - \sigma} \left( Y - \alpha \tau (Y - E) \right)^{1 - \sigma} + \frac{p \left( \frac{E}{Y} \right) + \epsilon}{1 - \sigma} \left( Y - \alpha \tau (Y - E) - (1 + \theta) \tau E \right)^{1 - \sigma} \quad (4)$$

Given a set of parameters, it is straightforward to find the optimal value of  $E$  that solves this maximization problem.<sup>60</sup> Table 4 presents the calibration results. Each row corresponds to a different calibration of  $A\mathcal{E}S$ . The first seven columns correspond to the parameter values, while the last three columns indicate the corresponding predictions:  $\frac{E}{Y}$  is the evasion rate,  $\frac{\partial \log(\tau(Y-E))}{\partial p}$  is the elasticity with respect to the audit probability and  $\frac{\partial \log(\tau(Y-E))}{\partial \theta}$  is the elasticity with respect to the penalty rate.<sup>61</sup>

All the parameters are calibrated so that the predictions always match the average evasion rate ( $\frac{E}{Y}$ ) of 26% (Gomez-Sabaini and Jimenez, 2012). In the first row, we assume a CRRA of 4 and set the audit probability and penalty rate equal to those estimated from the administrative records ( $p_0 = 0.117$  and  $\theta = 0.306$ ). To match the 26% evasion rate, we allow for a non-audit detection rate of  $\epsilon = 0.581$ . The resulting elasticity with respect to  $p$  is 4.548 and the elasticity with respect to  $\theta$  is 3.475. This is the simplest extension to the  $A\mathcal{E}S$  model, and given that its predictions are in the middle range of all our calibrations, we consider this our preferred specification. The remaining rows present results under alternative calibrations of the model. Even though the models are quite different, the predicted elasticities are in the same order of magnitude as our preferred specification.

In the second row, instead of accommodating the evasion rate of 26% by introducing the non-audit detection rate, we assume a social responsibility parameter of  $\alpha = 0.195$ . This different approach yields different predicted elasticities, but they are still in the same order of magnitude: the elasticity with respect to the audit probability is 9.422, and the elasticity with respect to the penalty rate is 1.208. The third row follows a similar specification from the second row, but we augment it by allowing for an endogenous audit probability. We let  $p_0 = p_1 = 0.0896$ , which accommodates two important features of the audit probabilities: the effective audit probability turns out to be equal to the observed average probability of 11.7%,

---

<sup>60</sup>For more details, see Appendix C.

<sup>61</sup>These two elasticities are defined exactly as in the econometric model from equations (2) and (3), to facilitate the comparison between the regression results and the calibrations.

and, consistent with the content of the *audit-endogeneity* message, a firm that does not evade taxes ( $\frac{E}{Y} = 0$ ) would double its audit probability if it decided to evade taxes ( $\frac{E}{Y} = 1$ ). Since this endogeneity parameter is not enough on its own to fit the observed evasion rate, we again rely on the social responsibility parameter to fit the data by setting  $\alpha = 0.229$ . This specification shows that introducing the endogenous nature of the audit probabilities does not change the elasticity with respect to the audit probability by much (it is 4.386, similar to the 4.548 from the first specification), but it does substantially reduce the elasticity with respect to the penalty rate (to 0.592). The fourth row follows a similar specification as in the second row, but we extend it by allowing individuals to have biased perceptions about the audits:  $p_0 = 0.407$  and  $\theta = 0.305$ . These biases would not be enough on their own to fit the observed evasion rate, so we again set the social responsibility parameter to  $\alpha = 0.590$ . Again, we obtain elasticities that are of the same order of magnitude as in the other specifications: the elasticity with respect to the audit probability is 4.021 and the elasticity with respect to the penalty rate is 1.642. The specifications in the second set of four rows are identical to those in the first set of four rows, except that we assume a CRRA of 2 instead of a CRRA of 4. The results indicate that assuming a higher risk aversion leads to elasticities that are even larger in magnitude.

### 5.2.3 Comparison between Experimental Results and the *AES* Calibration

We can test the null hypothesis that the elasticities with respect to  $p$  and  $\theta$  in the main specification of the *audit-statistics* presented in column (1) of Table 3 ( $\gamma_p = -0.063$  and  $\gamma_\theta = -0.033$ ) are equal to those in our preferred *AES* calibration. We can reject the null hypothesis that the elasticity is 4.55 for the audit probability and that it is 3.48 for the penalty rate (both tests with p-values < 0.001). More precisely, the 95% confidence interval for the estimate of  $\gamma_p$  in column (1) of Table 3 is  $[-0.536, 0.411]$ , and the corresponding interval for  $\gamma_\theta$  is  $[-0.264, 0.198]$ : the calibrated elasticities are substantially above these intervals. These calibrated elasticities are also well beyond the confidence intervals for the value of  $\gamma_p$  estimated for the *audit-threat* model  $[-0.062, 0.496]$ , from column (1) of Table 3).

One potential concern with the above comparison is that it is based on the implicit assumption that a letter conveying the message of a 1 percentage point higher signal of  $p$  or  $\theta$  will increase the perception of the parameter by the recipient by 1 percentage point. This is probably a strong assumption: some individuals may not have read the letter in its entirety, they may have not entirely believed in the content of our message, or they may not have updated their prior by the full value of the signal. To establish a benchmark, we can compare our setting with studies of learning from economic variables, such as the inflation rate (Cavallo et al., 2017), the cost of living (Bottan and Perez-Truglia, 2017), and housing

prices (Fuster et al., 2018). These studies find that for each percentage point increase in the feedback given to subjects, the average individual updates their beliefs by about half a percentage point. If we assume this rate of learning, then we should double the elasticities estimated in our regressions before comparing them to the calibrations of *AES*. Under this assumption, we can still reject the null hypothesis that the estimated and the calibrated elasticities of tax compliance with respect to the audit probability are equal (p-value < 0.001 for both  $\gamma_p$  and  $\gamma_\theta$ ).

The effect of the differential values of  $p$  and  $\theta$  in the *audit-statistics* treatments, presented in Section 5.1 above, can provide a direct estimate of the learning rate in our context. The survey data suggests that a percentage point increase in the signal about the audit probability provided in the letter increased the perceived audit probability nine months later by 0.397 percentage points (SE 0.288). Even if it is imprecisely estimated and thus statistically insignificant at conventional levels, it is reassuring that this point estimate suggests a learning rate consistent with other studies of learning. Moreover, we must keep in mind that due to multiple sources of noncompliance, this is probably an underestimate of the true learning rate.<sup>62</sup> Moreover, we can reproduce the analysis under an extremely conservative assumption about the magnitude of the learning rate: even if we assumed that for each percentage point difference in the letter individuals only adjusted their beliefs by one tenth of a percentage point, we would still reject the null hypothesis that the estimated elasticities are equal to those in the *AES* calibration (p-values of 0.032 and 0.001 for the the audit probability and for the penalty rate, respectively).

#### 5.2.4 Comparison to Related Studies

We can also compare our estimated elasticities to those from related studies. There is a literature that uses laboratory experiments to study tax evasion. These experiments often randomize the probability of being audited by the experimenter and the penalties involved. Consistent with our results, those laboratory studies find evidence of probability neglect. For example, Alm et al. (1992) find an elasticity of 0.169 with respect to the audit probability (comparable to our estimate of -0.063), and an elasticity of 0.037 with respect to the penalty rate (comparable to our estimate of -0.033). Indeed, these elasticities are statistically indistinguishable from those obtained in our study (p-values of the differences are 0.338 and 0.555 respectively).

We can also compare our findings to the results of a couple of related field experiments. Dwenger et al. (2016) conducted a field experiment in the context of a local church tax in Germany for which enforcement was extremely lax. While their experiment was not designed to

---

<sup>62</sup>For a discussion of the sources of non-compliance, see footnote 55.

test the *A&S* model, it did include one treatment arm where the message mentioned different audit probabilities ( $p = 0.1$ ,  $p = 0.2$ , or  $p = 0.5$ ). Their results are qualitatively consistent with our finding of probability neglect: the effects of all these probability messages are statistically indistinguishable from each other. Another related experiment, Kleven et al. (2011), included a treatment arm with two different audit probabilities. Consistent with our results, they find economically negligible differences in tax compliance between individuals assigned to different audit probabilities.<sup>63</sup> However, the evidence from Kleven et al. (2011) would still be consistent with *A&S* because their subjects face automatic third-party reporting, which our subjects do not. The authors conducted their experiment with wage earners in a country where tax evasion is automatically detected through third-party reporting regardless of audits. As a result, *A&S* predicts that, consistent with their evidence, wage earners should report their earnings truthfully regardless of the probability of being audited.

## 6 Discussion: Risk-as-Feelings

In this section, we summarize the findings and discuss their potential interpretations and implications.

We present three main findings. First, we documented *increased compliance*: on average, the *audit-statistics* message had a positive effect on tax compliance. Second, we reported *reduced subjective probability*: on average, the *audit-statistics* message decreased the perceived probability of being audited. Third, we documented *probability neglect*: the effect of the *audit-statistics* message did not depend on the audit probability included in the letter or on the firm’s prior belief about this probability. Jointly, these three findings are inconsistent with the predictions of *A&S*. This begs the question of what framework might provide a better fit for these results.

One natural candidate is that of salience (Chetty et al., 2009). In this framework, firms behave as if the probability of detection and the penalty rate are zero unless these parameters are made salient to them. This explanation could reconcile the findings of increased compliance and reduced perceived audit probability: even if firms who were sent the messages adjusted their perceived audit probabilities downwards, they would have behaved as if those probabilities were zero if they had not received those signals. The messages made the audit probabilities salient and thus caused them to not be perceived as zero. However, the salience

---

<sup>63</sup>In one of their treatments, they send letters to individuals stating large audit probabilities of  $p = 50\%$  and  $p = 100\%$ . Compared with a group that did not receive any letter, they find that the letters had a positive and significant effect on declared income and tax liability. The differential effect between these two conditions, while statistically significant, is economically negligible: an increase in the signal of the probability of audit from 50% to 100% increases reported income by 0.025% and taxes paid by 0.05%.

model fails to fit other features of our findings. First, by definition, salience models imply short-lived effects. A reminder about a non-salient tax should affect the behavior of an agent only at the time of receiving the information, but not days or months later. Effectively, salience models predict a rapid decay of the effect of information over time. This prediction contradicts our evidence about the persistent increased compliance from firms that received our *audit-statistics* letter. The effect, while decaying, remained in force for months after the messages were transmitted.<sup>64</sup> Salience models are also inconsistent with our finding of *probability neglect* – making a high audit probability salient should have a stronger impact than making a low audit probability salient.<sup>65</sup>

A second natural explanation can be found in agency issues within the firm: for instance, the person receiving the letter may differ from the person deciding how much tax to evade. This type of agency issues would generate insensitivity to the information received (or at least an attenuation effect). However, we find that firms do react to the information received, but in the opposite direction to that predicted by *A&S* (i.e., information on low audit probabilities reduces, rather than increases, evasion). Moreover, agency and information frictions should be weaker in smaller firms, and those in our experimental sample are small: over 75% of firms in our sample have 5 or fewer employees.<sup>66</sup> Moreover, the heterogeneity analysis indicates no substantial or statistically significant difference for firms below and above the median number of employees (1 to 3 employees versus the rest), which further reinforces our intuition that principal-agent and information frictions might not be the decisive factor at play in our context.<sup>67</sup>

Our preferred interpretation is based on the model of risk-as-feelings (Loewenstein et al., 2001). The models used for choice under uncertainty are cognitive in that agents make decisions using some type of expectation-based calculus. The risk-as-feelings model proposes that responses to fearsome situations may differ substantially from cognitive evaluations of the same risks (Loewenstein and Lerner, 2003).<sup>68</sup> When fear is involved, agents tend to neglect

---

<sup>64</sup>The informational treatment may still increase salience and cause an instantaneous effect with lasting consequences, for instance if it induces a change in the way the firm deals with evasion in transactions. However, such a change would imply a constant effect over time, whereas we find a substantial decline in evasion over the year following the intervention.

<sup>65</sup>Another model from behavioral economics that could be considered here is that of prospect theory (Kahneman and Tversky, 1979; see for example Dharm and al Nowaihi, 2007). This extension of the model, however, is unlikely to explain our findings. For example, prospect theory is unlikely to explain our finding of *probability neglect*: although differences between extremely low probabilities can be ignored under prospect theory, the range of probabilities in our context was far from what is normally considered extremely low (e.g., in the *audit-threat* arm, the probabilities were 25% versus 50%).

<sup>66</sup>More precisely, 29.1% of the firms have a single employee, 46.2% have between 2 and 5 employees, and 15% have between 6 and 10 employees.

<sup>67</sup>Results reported in Appendix B.8.

<sup>68</sup>A related concept, the affect heuristic, corresponds to the quick, automatic and intuitive evaluations of

the cost-benefit calculus and instead make quick, automatic and intuitive responses to the risks. A key prediction of this model is that feelings about risk will be mostly insensitive to changes in probability, which is known in the literature as probability neglect (Sunstein, 2002; Zeckhauser and Sunstein, 2010), or the fear that makes individuals focus on the downside of outcomes and thus ignore the underlying likelihoods. There is evidence of probability neglect in a range of fearsome situations involving electric shocks, arsenic, abandoned hazardous waste dumps, pesticides, and anthrax (Sunstein, 2003; Zeckhauser and Sunstein, 2010).

This model of risk-as-feelings can reconcile our three key findings. The model can reconcile the findings of *increased compliance* and *reduced subjective probability*: even if the perceived probability of an audit decreased among the treated subjects, they may still be scared into paying more taxes because they did not rely on cognitive evaluations of probabilities. The risk-as-feelings model predicts *probability neglect* and thus fits out third finding too.

The model of risk-as-feeling suggests that taxpayers overreact to the threat of audits.<sup>69</sup> This interpretation can provide an explanation for the paradox that, despite the low audit probabilities and penalty rates, most taxpayers still report the threat of audits as a major reason for why they report their taxable income truthfully. For example, a survey by the United States Internal Revenue Service (2018) indicates that 61% of U.S. taxpayers claim that a “fear of audits” exerts a significant influence on their tax compliance decisions.<sup>70</sup> In comparison, audits are perceived to be as strong of a deterrent as third-party reporting: 66% of respondents identified “third-party reporting (e.g., wages, interest, dividends)” as an important factor for their tax compliance. Moreover, there is some direct evidence that, consistent with the risk-as-feelings model, taxpayers have an emotional reaction when thinking of tax audits and the tax authority more generally. Some of the evidence comes from laboratory experiments. For instance, Coricelli et al. (2010) conducted a tax evasion game in the laboratory and measured how emotional arousal affected tax evasion decisions. They showed that the intensity of emotional arousal predicts whether individuals evade and by how much. In a related laboratory study, Dulleck et al. (2016) showed a significant correlation between tax compliance and physiological markers of stress during the tax reporting decision. Instances of fear of tax audits can also be found in the media. For example, a

---

risky situations based on emotions, which might be used as a shortcut for more complex evaluations of risk (Slovic et al., 2004). Borrowing Kahneman’s (2003) terminology for the dual system model of the human mind, emotions might influence the intuitive system.

<sup>69</sup>Indeed, this excessive caution has been documented in other contexts (Loewenstein et al., 2001). For example, a fear of terrorist attacks can make people choose other, more dangerous forms of transport; and a fear of shark attacks can lead to unnecessary legislation (Sunstein, 2002, 2003; Zeckhauser and Sunstein, 2010).

<sup>70</sup>More precisely, 32% of respondents claim that a “fear of audits” exerts “a great deal of an influence” and 29% “somewhat of an influence” on whether they honestly report and pay their taxes.

Washington Post (2016) article claims that “a lot of people are super scared of the Internal Revenue Service” and that its powers “can instill a lot of fear,” and The New York Times (2009) reported cases where the fear of the tax authority is so extreme that it is considered a phobia.

In other areas of public policy, the risk-as-feelings heuristic can be a problem, because it distorts facts and promotes irrational judgment, leading to suboptimal decisions from a pure risk-assessment perspective. For example, Zeckhauser and Sunstein (2010) and others discuss cases involving regulation of nuclear power, vaccines, and other emotion-arousing issues. For tax collection, however, these behavioral biases might have positive implications for the tax authority’s goals. Indeed, there is anecdotal evidence that tax authorities use these tactics to foster tax compliance. In the United States, for example, a disproportionately large number of tax enforcement press releases covering criminal convictions and civil injunctions are released during the weeks immediately preceding Tax Day, presumably to scare taxpayers into preparing compliant returns (Morse, 2009; Blank and Levin, 2010). Some tax experts even claim that the IRS “likes [targeting] celebrities because they get the most bang for their buck in terms of publicity” to “scare the public into complying” (Forbes, 2008).

In particular, the risk-as-feelings framework indicates that vivid imagery can be used to instill fear and bias risk evaluations (Slovic et al., 2004; Zeckhauser and Sunstein, 2010). Coincidentally, tax agencies seem to resort to vivid images in some of their advertising campaigns. A TV advertisement in the United States showed the IRS as “something like poltergeist coming out of a TV set and the world falling apart,” followed by the phrase, “Have you filed your income tax?” (United Press International, 1988). The U.K. tax agency used advertisement campaigns that also rely on fear imagery. One poster features a pair of eyes peeking threateningly through a gash in the paper and reads, “If you’ve declared all your income you have nothing to fear.”<sup>71</sup> This anecdotal evidence suggests that some tax administrations may be leveraging fear for tax collection.<sup>72</sup>

## 7 Conclusions

The canonical model of Allingham and Sandmo (1972) predicts that firms evade taxes by optimally trading off the costs and benefits of evasion, but it is unclear whether real-world

---

<sup>71</sup>This poster is reproduced in Appendix D.

<sup>72</sup>Whether these fear tactics should be used by tax agencies or not, however, is outside the scope of this paper. For example, these tactics may be ethically questionable in the extent to which they rely on deception. Moreover, actively promoting fear could have unintended negative effects, such as imposing negative psychological stress on taxpayers. For a discussion on the ethical and practical issues with the use of communication efforts to increase tax compliance, see for example Morse (2009).

firms react to audits in this way. We designed a large-scale field experiment in collaboration with Uruguay’s tax authority to assess the factors behind firms’ evasion behavior and their reactions to audits. Our findings indicate that firms do increase their tax compliance when informed about the auditing process. However, we do not find this reaction to be consistent with the predictions of *AES*. For example, the information about audits decreased (rather than increased) the perceived probability of being audited; moreover, the effects of our messages about audit probabilities were independent of the signal we conveyed and of the firms’ prior beliefs. Models of salience are consistent with the increased compliance and the reduced perceived perception of audit probabilities that we observed, but they are not consistent with our findings of probability neglect. We argue that all three findings can be reconciled by the risk-as-feelings model, which highlights the role of emotions in decision-making and which predicts that agents might exhibit probability neglect in dreaded or feared situations, like paying taxes.

Our findings also contribute to the more general debate about the determinants of tax compliance. One of the main puzzles in the literature is that evasion rates seem too low. Third-party reporting can explain high compliance for some sources of income, such as wage income (Kleven et al., 2011). However, we would expect much higher evasion rates in other contexts, such as self-employed income, where there is limited third-party reporting and low detection probabilities and penalty rates. One traditional explanation for this puzzle is based on tax morale: firms and individuals do not evade taxes because it is the right thing to do (Luttmer and Singhal, 2014). Our evidence suggests an alternative explanation for the puzzle: due to the emotional nature of the decision, audits scare taxpayers into compliance in the same way that scarecrows scare birds.

We conclude by discussing some policy implications. In the traditional framework of *AES*, the relevant policy lever is the number of audits: the tax agency must find the point at which the marginal cost of an additional audit equals the expected marginal benefit (i.e., higher tax revenues). Our findings suggest that small and medium firms face significant information and optimization frictions when reacting to audits. These frictions introduce new levers for policy-making. For example, tax agencies can decide whether to be transparent about the auditing process,<sup>73</sup> whether to contact taxpayers to remind them of it, and whether to make the costs of being a tax cheat salient and vivid through advertisement campaigns.<sup>74</sup> Indeed, we

---

<sup>73</sup>On the one hand, our evidence indicates that increasing transparency about the audit probability would reduce the average perceived probability of being audited, which could reduce tax compliance. On the other hand, our finding of probability neglect suggests that, in the end, the reduction in perceived audit probability may not affect tax compliance.

<sup>74</sup>For a practical discussion on how to implement this type of policy, including the drawbacks, see Morse (2009). Furthermore, this same principle can be used to improve compliance with other laws. For instance, Dur and Vollaard (2019) provide experimental evidence to show that the salience of law enforcement can be

discussed anecdotal evidence that some tax agencies may already have a working knowledge of these policy levers. For example, some tax agencies seem to avoid transparency about the auditing process while increasing visibility of enforcement actions around tax day. Some even refer to fear in their advertisement campaigns. However, there is no direct evidence on whether these policies effectively increase tax compliance or whether they have unintended effects, such as instigating so much fear in taxpayers that their anxiety and unhappiness trump the positive effects of increased tax revenues. As stated by Alm (2019) in a recent review of the literature, “the role of emotions in tax compliance decisions remains largely unexamined.” Our results highlight the need for more research on probability neglect in the decision to pay taxes. Moreover, additional research should examine the role of emotions on other important economic choices beyond tax compliance.

## References

- Acemoglu, D. and M. O. Jackson (2017). Social Norms and the Enforcement of Laws. *Journal of the European Economic Association* 15(2), 245–295.
- Allingham, M. G. and A. Sandmo (1972). Income Tax Evasion: A Theoretical Analysis. *Journal of Public Economics* 1, 323–338.
- Alm, J. (2019). What motivates tax compliance? *Journal of Economic Surveys* 33(2), 353–388.
- Alm, J., B. Jackson, and M. McKee (1992). Estimating the Determinants of Taxpayer Compliance with Experimental Data. *National Tax Journal* 45(1), 107–114.
- Alm, J., G. H. McClelland, and W. Schulze (1992). Why do people pay taxes? *Journal of Public Economics* 48(1), 21–38.
- Becker, G. S. (1968). Crime and Punishment: An Economic Approach. *Journal of Political Economy* 76(2), 169–217.
- Bergman, M. and A. Nevarez (2006). Do Audits Enhance Compliance? An Empirical Assessment of VAT Enforcement. *National Tax Journal* 59(4), 817–832.
- Bergolo, M., R. Ceni, G. Cruces, M. Giacobasso, and R. Perez-Truglia (2018). Misperceptions about Tax Audits. *AEA Papers and Proceedings* 108, 83–87.
- Blank, J. D. and D. Z. Levin (2010). When Is Tax Enforcement Publicized? *Virginia Tax Review* 30.
- Blumenthal, M., C. Christian, and J. Slemrod (2001). Do Normative Appeals Affect Tax Compliance? Evidence From a Controlled Experiment in Minnesota. *National Tax Journal* 54(1), 125–138.
- Bott, K. M., A. W. Cappelen, E. Ø. Sørensen, and B. Tungodden (2020). You’ve got mail: A randomized field experiment on tax evasion. *Management Science* 66(7), 2801–2819.
- Bottan, N. L. and R. Perez-Truglia (2017). Choosing Your Pond: Location Choices and Relative Income. *NBER Working Paper* (23615).
- Bruine de Bruin, W. and K. G. Carman (2012). Measuring Risk Perceptions: What Does the Excessive Use of 50% Mean? *Medical Decision Making* 32(2), 232–236.

---

used to reduce illegal garbage disposal.

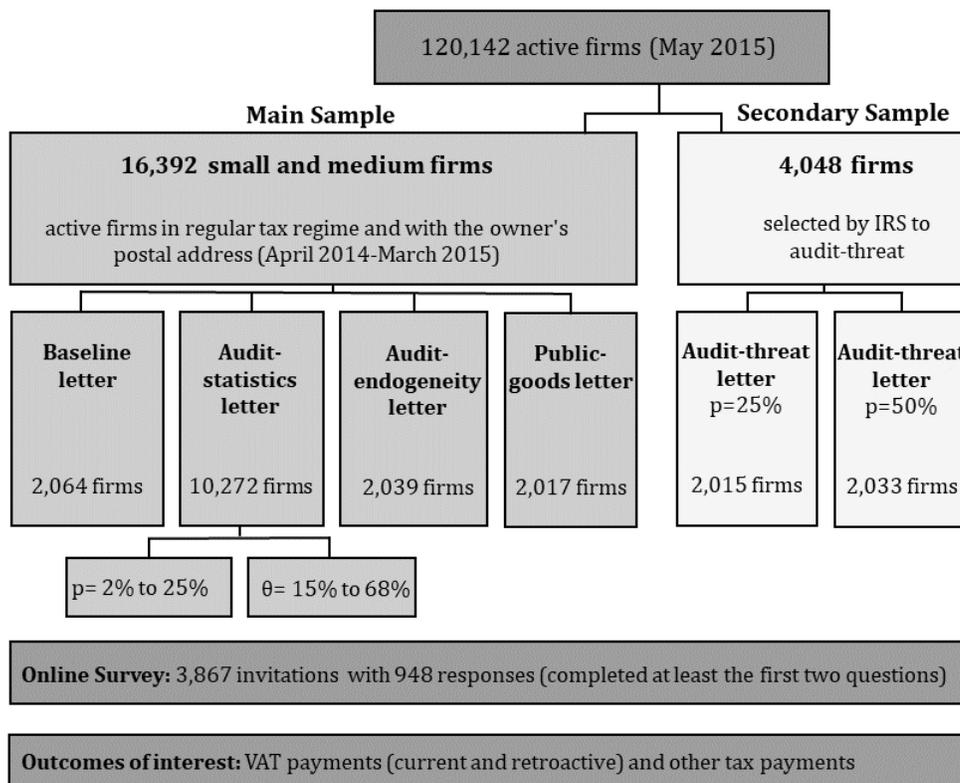
- Bruine de Bruin, W., P. S. Fischbeck, N. A. Stiber, and B. Fischhoff (2002). What number is "fifty-fifty"? Redistributing excess 50% responses in risk perception studies. *Risk Analysis* 22(4), 725–735.
- Castro, L. and C. Scartascini (2015). Tax Compliance and Enforcement in the Pampas Evidence From a Field Experiment. *Journal of Economic Behavior & Organization* 116, 65–82.
- Cavallo, A., G. Cruces, and R. Perez-Truglia (2017, July). Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments. *American Economic Journal: Macroeconomics* 9(3), 1–35.
- Chetty, R., A. Looney, and K. Kroft (2009). Saliency and Taxation: Theory and Evidence. *The American Economic Review* 99(4), 1145–77.
- Coricelli, G., M. Joffily, C. Montmarquette, and M. C. Villeval (2010). Cheating, Emotions, and Rationality: An Experiment on Tax Evasion. *Experimental Economics* 13(2), 226–247.
- Cowell, F. A. and J. P. F. Gordon (1988). Unwillingness to pay: Tax evasion and public good provision. *Journal of Public Economics* 36(3), 305–321.
- Dhami, S. and A. al Nowaihi (2007). Prospect theory versus expected utility theory: Why Do People Pay Taxes? *Journal of Economic Behavior and Organization* 64(1), 171–192.
- Dulleck, U., J. Fooker, C. Newton, A. Ristl, M. Schaffner, and B. Torgler (2016). Tax compliance and psychic costs: Behavioral experimental evidence using a physiological marker. *Journal of Public Economics* 134, 9 – 18.
- Dur, R. and B. Vollaard (2019). Saliency of law enforcement a field experiment. *Journal of Environmental Economics and Management* 93(C), 208–220.
- Dwenger, N., H. Kleven, I. Rasul, and J. Rincke (2016). Extrinsic and Intrinsic Motivations for Tax Compliance: Evidence from a Field Experiment in Germany. *American Economic Journal: Economic Policy* 8(3), 203–232.
- Erard, B. and J. S. Feinstein (1994). The Role of Moral Sentiment and Audit Perceptions in Tax Compliance. *Public Finance* 49, 70–89.
- Fellner, G., R. Sausgruber, and C. Traxler (2013). Testing Enforcement Strategies in the Field: Threat, Moral Appeal and Social Information. *Journal of the European Economic Association* 11(3), 634–660.
- Forbes (2008, July). Pity The Celebrity Taxpayer. *Forbes*.
- Fuster, A., R. Perez-Truglia, and B. Zafar (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *NBER Working Paper*.
- Gomez-Sabaini, J. C. and J. P. Jimenez (2012). Tax structure and tax evasion in Latin America. *Macroeconomics of Development Series* 118.
- Gomez-Sabaini, J. C. and D. Moran (2014). Tax policy in Latin America Assessment and guidelines for a second generation of reforms. *Macroeconomics of Development Series* 133.
- Hallsworth, M., J. A. List, R. D. 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.
- Harris, L. and I. Associates (1988). 1987 taxpayer opinion survey. *Washington, DC: Internal Revenue Service Document*.
- Kahneman, D. (2003). A Perspective on Judgement and Choice: Mapping Bounded Rationality. *American Psychologist* 58(9), 697–720.
- Kahneman, D., P. Slovic, and A. Tversky (Eds.) (1982). *Judgment under uncertainty: heuristics and biases*. Cambridge ; New York: Cambridge University Press.

- Kahneman, D. and A. Tversky (1974). Judgment under Uncertainty: Heuristics and Biases. *Science* 185(4157), 1124–1131.
- Kahneman, D. and A. Tversky (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47(2), 263–291.
- Kleven, H. J., M. B. Knudsen, T. Kreiner, S. Pedersen, and E. Saez (2011). Unwilling or Unable to Cheat? Evidence from a Randomized Tax Audit Experiment in Denmark. *Econometrica* 79(3), 651–692.
- Kleven, H. J., C. Kreiner, and E. Saez (2016). Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries. *Economica* 83, 2016.
- Konrad, K. A., T. Lohse, and S. Qari (2016). Compliance With Endogenous Audit Probabilities. *Scandinavian Journal of Economics*.
- Lichtenstein, S., P. Slovic, B. Fischhoff, M. Layman, and B. Combs (1978). Judged frequency of lethal events. *Journal of experimental psychology: Human learning and memory* 4(6), 551.
- Loewenstein, G. and S. Lerner (2003). The role of affect in decision making. In R. Davidson, K. Scherer, and H. Goldsmith (Eds.), *Handbook of Affective Sciences*. Oxford: Oxford University Press.
- Loewenstein, G. F., E. U. Weber, C. K. Hsee, and N. Welch (2001). Risk as feelings. *Psychological Bulletin* 127(2), 267–286.
- Luttmer, E. F. P. and M. Singhal (2014). Tax Morale. *Journal of Economic Perspectives* 28(4), 149–168.
- McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics* 99(2), 210–221.
- Meiselman, B. S. (2018). Ghostbusting in detroit: Evidence on nonfilers from a controlled field experiment. *Journal of Public Economics* 158, 180 – 193.
- Morse, S. C. (2009). Using Salience and Influence to Narrow the Tax Gap. *Loyola University Chicago Law Journal* 40, 483.
- Naritomi, J. (2019, September). Consumers as tax auditors.
- Nathan, B., R. Perez-Truglia, and A. Zentner (2020). My taxes are too darn high: Tax protests as revealed preferences for redistribution. *NBER Working Paper No. 27816*.
- New York Times (2009, April). A Paralyzing Fear of Filing Taxes. *New York Times*.
- Perez-Truglia, R. and U. Troiano (2018). Shaming tax delinquents. *Journal of Public Economics* 167, 120–137.
- Pomeranz, D. (2015). No Taxation Without Information: Deterrence and Self-Enforcement in the Value Added Tax. *The American Economic Review* 105(8), 2539–2569.
- Pomeranz, D. and J. Vila-Belda (2018). Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities.
- Reinganum, J. F. and L. L. Wilde (1985). Income tax compliance in a principal agent framework. *Journal of Public Economics* 26(1), 1–18.
- Saez, E., J. Slemrod, and S. Giertz (2012). The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of Economic Literature* 50(1), 3–50.
- Scholz, J. T. and N. Pinney (1995). Duty, Fear, and Tax Compliance: The heuristic basis of citizenship behavior. *American Journal of Political Science* 39, 2.
- Slemrod, J. (2008). Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance. *National Tax Journal* 61.

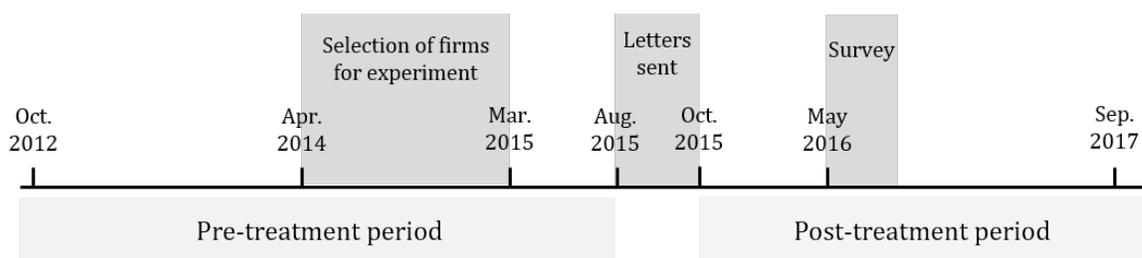
- Slemrod, J. (2018). Tax Compliance and Enforcement. *Journal of Economic Literature Forthcoming*.
- Slemrod, J., M. Blumenthal, and C. Christian (2001). Taxpayer Response to an Increased Probability of Audit: Evidence from a Controlled Experiment in Minnesota. *Journal of Public Economics* 79(3), 455–483.
- Slovic, P., M. L. Finucane, E. Peters, and D. G. MacGregor (2004). Risk as analysis and risk as feelings: some thoughts about affect, reason, risk, and rationality. *Risk Analysis: An Official Publication of the Society for Risk Analysis* 24(2), 311–322.
- Srinivasan, T. N. (1973). Tax Evasion: A Model. *Journal of Public Economics* 2(4), 339–346.
- Sunstein, C. (2002). Probability Neglect: Emotions, Worst Cases, and Law. *Yale Law Journal* 112(1), 61–107.
- Sunstein, C. R. (2003). Terrorism and probability neglect. *Journal of Risk and Uncertainty* 26(2-3), 121–136.
- United Press International (1988). Psychologist takes issue with irs scare tactic. *UPI-United Press International*.
- United States Internal Revenue Service (2018). Comprehensive Taxpayer Attitude Survey (CTAS) 2017 Executive Report. Publication 5296 (Rev. 3-2018) Catalog Number 71353Y, Department of Treasury, Washington, D.C.
- Washington Post (2016, August). That is NOT the IRS Calling You! *The Washington Post*.
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data*. MIT press.
- Yitzhaki, S. (1987). On the Excess Burden of Tax Evasion. *Public Finance Review* 15(2), 123–137.
- Zeckhauser, R. and C. R. Sunstein (2010). Dreadful Possibilities, Neglected Probabilities. In E. Michel-Kerjan and P. Slovic (Eds.), *The Irrational Economist: Making Decisions in a Dangerous World*, pp. 116–123. New York: Public Affairs Press.

Figure 1: Structure of the Field Experiment

a. Samples and Treatment Arms



b. Timeline



Notes: Panel (a) reports the key features of the experimental design. Panel (b) reports the key dates of the field experiment and the survey.

Figure 2: Sample Letter

  
Montevideo, August 20<sup>th</sup> 2015

Mr./Ms. Taxpayer:

The DGI has the authority to perform inspections (see Art. 68 of the tax code) and routine audits of taxpayers on the basis of crosschecks and assessment of data compiled to detect oversights and inconsistency on tax returns as well as pending tax debts.

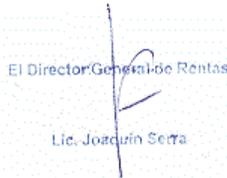
## MESSAGE

The aim of the DGI, and the primary challenge it faces, is to ensure the collection of revenue to sustain life in society. Additionally, its task is to generate a framework of fair and transparent competition where the failure of some to meet their obligations does not have an unfavorable impact on honest taxpayers. In order to meet these goals, inspections are performed in a routine fashion.

Your micro, small, or medium-sized business has been randomly selected to receive this information. *It is solely for your information and its receipt does not require you to present any documentation to the DGI offices.*

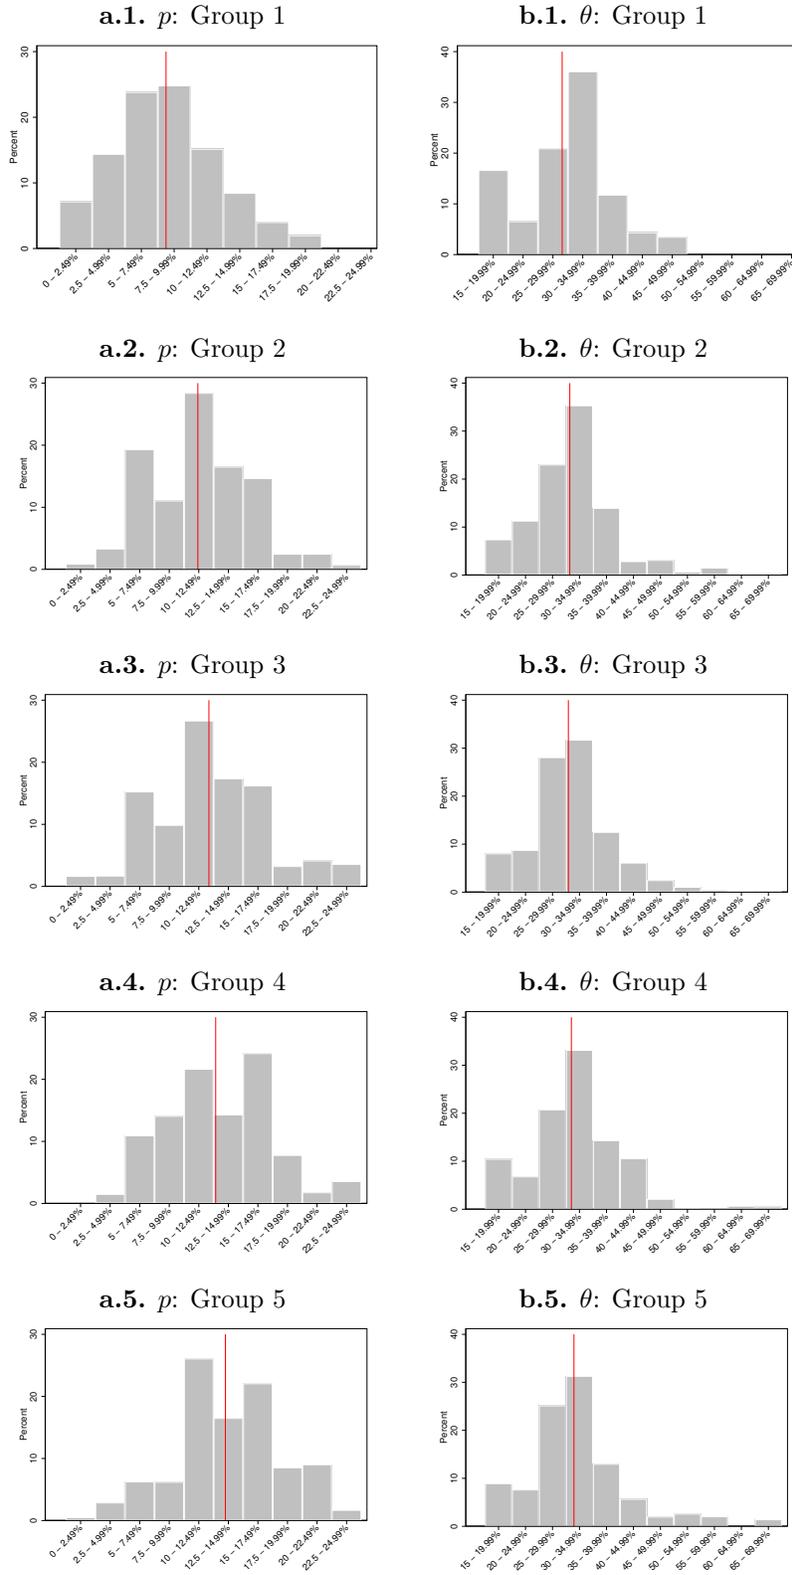
We ask you to comply with your tax obligations for the sake of the country we all want, a more and more developed Uruguay with greater and greater social cohesion.

Sincerely,

  
El Director General de Rentas  
Lic. Joaquín Serra  
Collection and Controls Division  
Internal Revenues Services

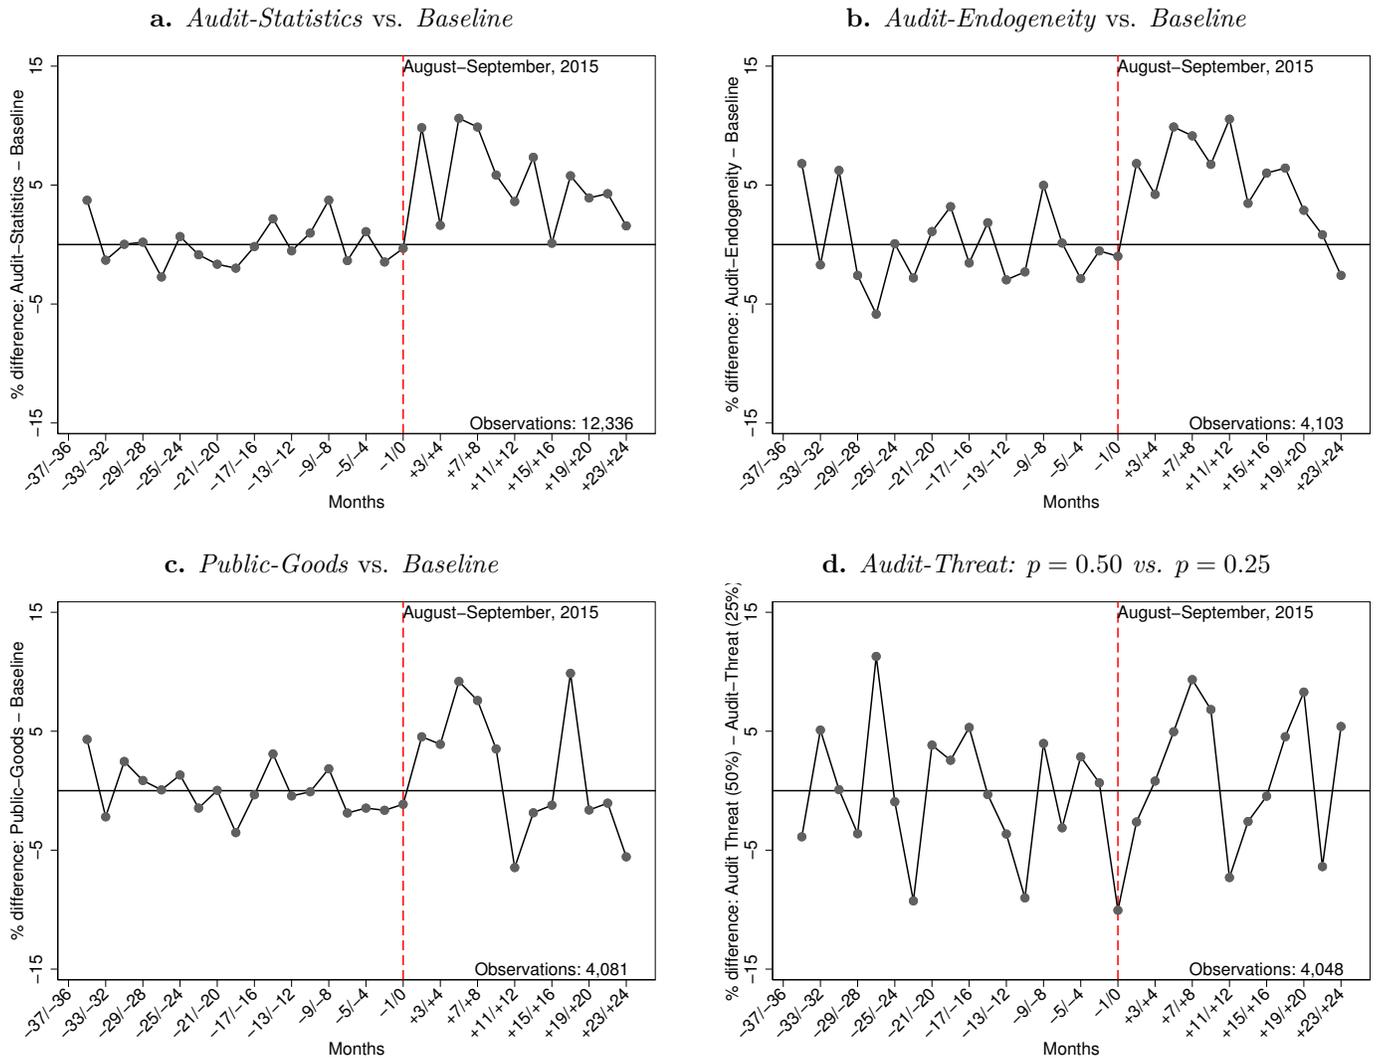
Notes: The *baseline* letter contains information about the goals and responsibilities of the tax authority. In the space with the text MESSAGE, the baseline letter is empty (See A.1 for the full letter). In the *audit-statistics* letter, we added a paragraph to the baseline letter that provided information about the audit probabilities and tax evasion penalty rates (Appendix A.2). In the *audit-threat* letter, we randomly assigned firms to groups with different probabilities (25% and 50%) of being audited in the following year (Appendix A.3). The *audit-endogeneity* letter included information about how evading taxes typically doubles the probability of being audited (Appendix A.4). Finally, the *public-goods* letter included a message providing information about the cost of evasion in terms of the provision of public goods (Appendix A.5).

Figure 3: Distribution of Statistics Shown in *Audit-Statistics* Letters by VAT Payment Quintiles



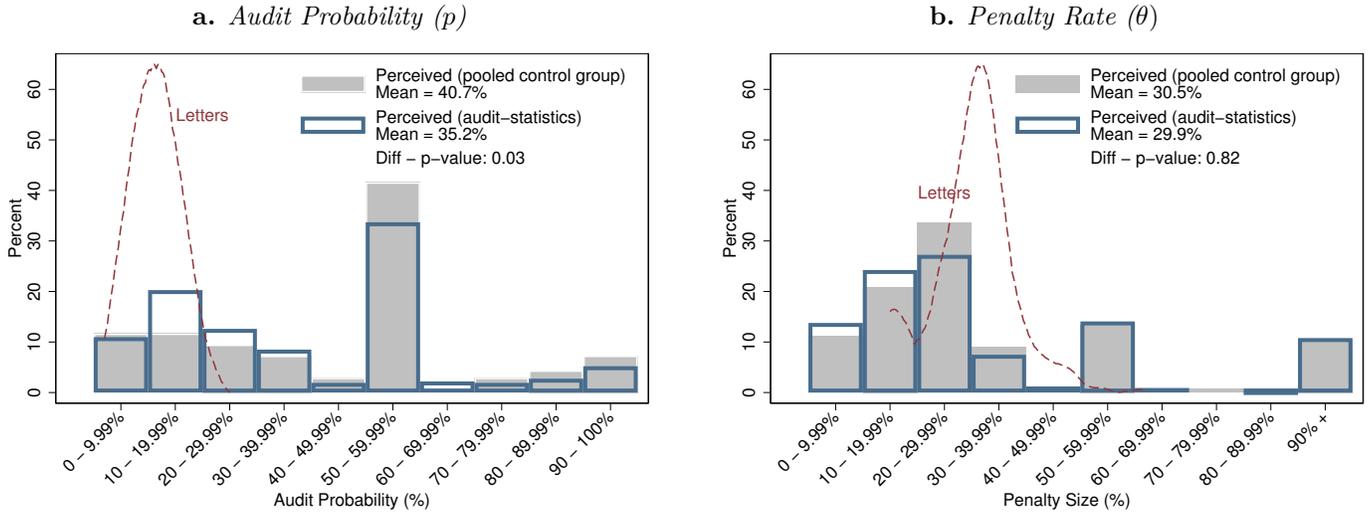
Notes:  $N=10,272$ . These panels show the information provided in the *audit-statistics* letter, including the probability of being audited ( $p$ , in panel (a)) and the penalty rate ( $\theta$ , in panel (b)). Groups 1 through 5 correspond to each of the pre-treatment VAT payment quintiles (group 1 being the bottom quintile and group 5 being the top quintile). In each panel, the red vertical line denotes the average audit probability or penalty rate for all the members in the group.

Figure 4: Effects of *Audit-Statistics*, *Audit-Endogeneity*, *Public-Goods* and *Audit-Threat* Messages on VAT Payments



Notes: These figures plot the percentage difference in bi-monthly total VAT payments between treatment and control groups, normalized by the average pre-treatment percentage difference (i.e. between months -35 and 0) for the same outcome. The data covers the period from October 2012 to September 2017. The months of August and September 2015 – when most of the letters were delivered – are defined as the reference bi-monthly period (and marked with the dashed vertical line). Panel (a) presents the effect of the *audit-statistics* message (i.e., the difference between *audit-statistics* and *baseline* letters), while panel (b) represents the effect of the *audit-endogeneity* message and panel (c) depicts the effect of the *public-goods* message. Panel (d) presents the difference between being assigned a 50% probability of being audited ( $p = 50\%$ ) and a 25% probability of being audited ( $p = 25\%$ ) in the *audit-threat* letters. For each pair of months, VAT payments are top-coded at the 99.99% percentile to avoid the contamination of the results by outliers.

Figure 5: Survey Results: Perception of Audit Probabilities and of Tax Evasion Penalty Rates by Treatment Group

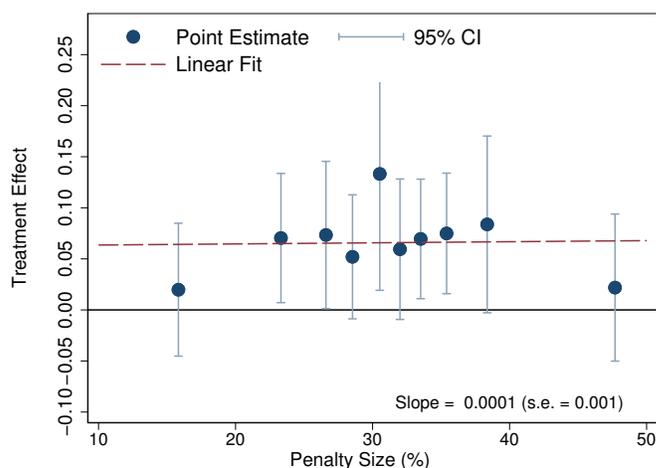
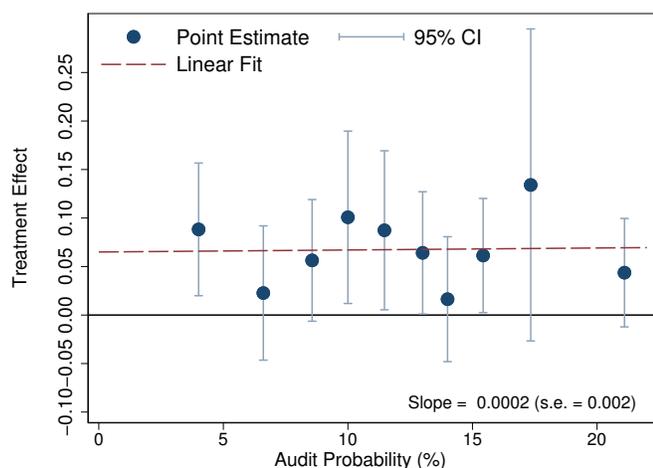


Notes: The histograms are based on the survey responses of individuals who did not self-identify as non-owners in the survey. *Perceived (pooled control group)* (N=137) refers to survey respondents who received the *baseline* (N=69) or the *public-goods* (N=68) letters during the experimental stage (neither of those two letters contained any information about audit probabilities or penalty rates). *Perceived (audit-statistics)* refers to respondents who received the *audit-statistics* letters (N=365). In panel (a) the x-axis represents the probability of being audited; in panel (b) it represents the average penalty rate. We report the mean responses and the p-value of the difference between the two groups. The answers correspond to questions Q2 and Q4 in the survey (see full survey questionnaire in Appendix A.7). The red line represents the density function of the information displayed in the *audit-statistics* letters, measured in the right y-axis (hidden for the sake clarity).

Figure 6: Effect of *Audit-Statistics* vs. *Baseline* by Deciles of  $p$  and  $\theta$

a. *Audit Probability* ( $p$ )

b. *Penalty Rate* ( $\theta$ )



Notes: Panel (a) plots the effect of the *audit-statistics* letter on total VAT payments by decile of  $p$  in the first year post-treatment (October 2015 – September 2016), while panel (b) reports the results from the same regressions by decile of  $\theta$  ( $N=10,272$ ). In both panels, each dot represents the estimated treatment effect for each decile of the parameter considered. These effects are estimated using a regression similar to the one reported in equation (1) but with two differences. First, instead of including a single treatment variable, we include 10 dummy variables, one for each decile of  $p$  or  $\theta$ . These dummies take the value of 1 if the signal specified in the letter belongs to the corresponding decile in the  $p$  or  $\theta$  distribution, and 0 if the signal corresponds to a different decile, or if the firm was assigned to the *baseline* treatment. Second, we include an additional set of dummies for quintiles of the pre-treatment VAT payments, which are the groups from which we drew the sample of “similar firms” to calculate  $p$  and  $\theta$ , and the corresponding interactions with the post-treatment indicator. All effects are depicted with a 95% confidence interval. The results are based on Poisson regressions, so the coefficients can be interpreted directly as semi-elasticities. Confidence intervals are computed with standard errors clustered at the firm level. The dashed line represents the linear fit that results from regressing the treatment effect on the average signal within the decile.

Table 1: Balance of Firm Characteristics across Treatment Groups

	Main Sample					Secondary Sample		
	Audit Statistics (1)	Public Goods (2)	Audit Endogeneity (3)	Baseline (4)	p-value test (5)	Audit Threat (25%) (6)	Audit Threat (50%) (7)	p-value test (8)
Share paid VAT (3 months pre-mailing)	0.925 (0.003)	0.939 (0.005)	0.926 (0.006)	0.928 (0.006)	0.181	0.897 (0.007)	0.891 (0.007)	0.538
Amount of VAT paid (3 months pre-mailing)	1.872 (0.027)	1.963 (0.067)	1.926 (0.069)	1.906 (0.059)	0.557	1.739 (0.097)	1.748 (0.092)	0.950
Years registered with tax agency	15.338 (0.170)	14.746 (0.224)	15.704 (0.538)	15.009 (0.225)	0.268	19.453 (0.285)	19.425 (0.286)	0.944
Share audited between 2013-2015	0.106 (0.004)	0.097 (0.009)	0.089 (0.009)	0.101 (0.009)	0.302	0.134 (0.010)	0.147 (0.010)	0.382
Number of employees	4.814 (0.264)	4.658 (0.538)	4.880 (0.566)	5.089 (0.635)	0.962	4.835 (0.126)	4.880 (0.117)	0.795
Share filed comprehensive tax return in 2013	0.682 (0.005)	0.687 (0.010)	0.691 (0.010)	0.687 (0.010)	0.871	0.999 (0.001)	1.000 (0.000)	0.558
Share no retail goods sector	0.289 (0.004)	0.293 (0.010)	0.283 (0.010)	0.300 (0.010)	0.621	0.431 (0.011)	0.434 (0.011)	0.845
Share retail goods sector	0.218 (0.004)	0.219 (0.009)	0.214 (0.009)	0.227 (0.009)	0.775	0.334 (0.011)	0.322 (0.010)	0.398
Share services sector	0.493 (0.005)	0.488 (0.011)	0.504 (0.011)	0.473 (0.011)	0.232	0.235 (0.009)	0.244 (0.010)	0.482
N	10,272	2,017	2,039	2,064		2,015	2,033	

Notes: Averages for different pre-treatment firm-level characteristics, disaggregated by treatment group and type of sample (robust standard errors are reported in parentheses). The main sample includes all firms selected as described in section 3.2. The secondary sample includes high-risk firms selected by the IRS for the *audit-threat* treatment. The last column of each sample reports the p-value of a test in which the null hypothesis is that the mean is equal for all the treatment groups. Data on the VAT amount and firm characteristics comes from administrative tax records (including monthly payments, annual tax returns and auditing registers). The amount of VAT reported in row 2 is expressed in constant thousands of U.S. dollars as of August 2015.

Table 2: Average Effects of *Audit-Statistics*, *Audit-Endogeneity* and *Public-Goods* Messages on VAT and Other Tax Payments by Time Horizon and Payment Timing

	By Time Horizon		By Payment Timing		By Tax Type	
	First Year (1)	Second Year (2)	Retroactive (3)	Concurrent (4)	Non-VAT (5)	VAT + Non-VAT (6)
<b>a. <i>Audit-Statistics</i> (10,272 firms) vs <i>Baseline</i> (2,064 firms)</b>						
Post-Treatment	0.070*** (0.021)	0.032 (0.027)	0.383*** (0.140)	0.053** (0.021)	0.086** (0.037)	0.073*** (0.020)
Pre-Treatment	0.009 (0.020)	0.004 (0.026)	-0.048 (0.118)	0.012 (0.020)	0.008 (0.043)	0.014 (0.021)
<b>b. <i>Audit-Endogeneity</i> (2,039 firms) vs <i>Baseline</i> (2,064 firms)</b>						
Post-Treatment	0.071*** (0.028)	0.032 (0.036)	0.264* (0.160)	0.061** (0.028)	0.090* (0.054)	0.078*** (0.028)
Pre-Treatment	-0.005 (0.028)	-0.009 (0.035)	0.097 (0.164)	-0.010 (0.028)	0.056 (0.055)	0.017 (0.028)
<b>c. <i>Public-Goods</i> (2,017 firms) vs <i>Baseline</i> (2,064 firms)</b>						
Post-Treatment	0.051** (0.025)	0.004 (0.032)	0.208 (0.170)	0.043* (0.025)	0.067 (0.043)	0.056** (0.024)
Pre-Treatment	-0.003 (0.024)	-0.017 (0.033)	-0.088 (0.163)	0.001 (0.024)	-0.038 (0.054)	-0.015 (0.026)

Notes: \* significant at the 10% level, \*\* at the 5% level, \*\*\* at the 1% level. Standard errors reported in parentheses are clustered at the firm level. Treatment effects are estimated using the difference-in-differences specification reported in equation (1) which compares treated firms to control firms and pre-treatment to post-treatment periods, using yearly aggregated variables. The results are based on Poisson regressions, so the coefficients can be interpreted directly as semi-elasticities. Panel (a) compares the *audit-statistics* message with the *baseline* letter, while panels (b) and (c) replicate the comparison for the *audit-endogeneity* and *public-goods* messages respectively. In the first row of each panel (“Post-Treatment”), the coefficient reported corresponds to a comparison between a post-treatment period and a pre-treatment period. The second row (“Pre-Treatment”) presents a falsification test where two pre-treatment periods are compared. Columns (1) and (2) report the effect of treatment by time horizon. The post-treatment effect reported in column (1) corresponds to the difference-in-differences estimate that compares October 2015 – September 2016 to October 2014 – September 2015. The post-treatment effect reported in column (2) is analogous but uses the second year after the treatment as the post-treatment period (i.e., October 2016 – September 2017). For the falsification tests, column (1) is based on a comparison between October 2014 – September 2015 and October 2013 – September 2014, while column (2) compares October 2014 – September 2015 to October 2012 – September 2013. Columns (3) and (4) present the first year effect of treatment on retroactive (3) and concurrent (4) VAT payments. Columns (5) and (6) report the first year results by type of tax. Column (5) presents the effect of the treatment on other (non-VAT) tax payments, while column (6) reports the effect on the total amount of taxes paid by the firms during the same period. In all cases, we restrict the analysis to firms that effectively received the letter as reported by the postal service.

Table 3: Elasticities of Tax Payments with Respect to Audit Probability and Penalty Rate, *Audit-Statistics* and *Audit-Threat* Sub-Treatments

	By Time Horizon		By Payment Timing		By Tax Type	
	First Year (1)	Second Year (2)	Retroactive (3)	Concurrent (4)	Non-VAT (5)	VAT + Non-VAT (6)
<b>a. <i>Audit-Statistics</i> (10,272 firms)</b>						
Audit Probability (%)						
Post-Treatment	-0.063 (0.242)	0.076 (0.232)	0.009 (1.103)	-0.040 (0.249)	0.109 (0.240)	0.038 (0.208)
Pre-Treatment	0.141 (0.164)	0.018 (0.203)	-1.709 (1.118)	0.229 (0.162)	-0.035 (0.230)	0.063 (0.147)
Penalty Size (%)						
Post-Treatment	-0.033 (0.118)	-0.175 (0.134)	0.928 (0.763)	-0.098 (0.114)	0.061 (0.103)	-0.001 (0.092)
Pre-Treatment	-0.128 (0.108)	-0.163 (0.127)	0.204 (0.524)	-0.145 (0.111)	0.018 (0.119)	-0.078 (0.087)
<b>b. <i>Audit-Threat</i> (4,048 firms)</b>						
Audit Probability (%)						
Post-Treatment	0.217 (0.142)	0.250 (0.175)	-0.347 (0.676)	0.205 (0.209)	0.002 (0.176)	0.233** (0.111)
Pre-Treatment	-0.185 (0.157)	-0.193 (0.171)	-0.432 (0.676)	-0.149 (0.125)	-0.067 (0.148)	-0.257 (0.164)

Notes: \* significant at the 10% level, \*\* at the 5% level, \*\*\* at the 1% level. Standard errors reported in parentheses are clustered at the firm level. Treatment effects are estimated using the difference-in-differences specification reported in equation (2) which compares treated firms that received different signals about  $p$  and  $\theta$ . In all cases, we include an additional set of dummies for quintiles of the pre-treatment VAT payments, which are the groups from which we drew the sample of “similar firms” to calculate  $p$  and  $\theta$ , and the corresponding interactions with the time variable. The results are based on Poisson regressions with variables expressed in percentage terms, so the coefficients can be interpreted directly as elasticities. Panel (a) presents the effect of providing different information regarding  $p$  and  $\theta$  in the *audit-statistics* message. Panel (b) compares the two *audit-threat* messages, i.e. the 50% threat of audit vs. the 25% threat of audit. For example, rows (1) and (3) of panel (a) present the effect of an additional percentage point of  $p$  and  $\theta$  (respectively) in the information included in the letters on post-treatment VAT payments. In the “Post-Treatment” rows, the coefficient reported corresponds to a comparison between a post-treatment period and a pre-treatment period. In the “Pre-Treatment” rows we present a falsification test where two pre-treatment periods are compared. Columns (1) and (2) report the effect of treatment by time horizon. The post-treatment effect reported in column (1) corresponds to the difference-in-differences estimate that compares October 2015 – September 2016 to October 2014 - September 2015. The post-treatment effect reported in column (2) is analogous but uses the second year after the treatment as the post-treatment period (i.e., October 2016 – September 2017). For the falsification tests, column (1) is based on a comparison between October 2014 – September 2015 and October 2013 - September 2014, while column (2) compares October 2014 – September 2015 to October 2012 - September 2013. Columns (3) and (4) present the first year effect of treatment on retroactive (3) and concurrent (4) VAT payments. Columns (5) and (6) report the first year results by type of tax. Column (5) presents the effect of the treatment on other (non-VAT) tax payments, while column (6) reports the effect on the total amount of taxes paid by the firms during the same period. In all cases, we restrict the analysis to firms that effectively received the letter as reported by the postal service.

Table 4: Predicted Elasticities under Different  $AES$  Calibrations

Setup							Predictions		
$\sigma$	$\tau$	$p_0$	$p_1$	$\epsilon$	$\theta$	$\alpha$	$\frac{E}{Y}$	$\frac{\partial \log(\tau(Y-E))}{\partial p}$	$\frac{\partial \log(\tau(Y-E))}{\partial \theta}$
4	0.22	0.117	0	0.581	0.306	1	0.26	4.548	3.475
4	0.22	0.117	0	0	0.306	0.195	0.26	9.422	1.208
4	0.22	0.0896	0.0896	0	0.306	0.229	0.26	4.386	0.592
4	0.22	0.407	0	0	0.305	0.590	0.26	4.021	1.642
2	0.22	0.117	0	0.620	0.306	1	0.26	9.854	6.600
2	0.22	0.117	0	0	0.306	0.169	0.26	18.857	2.111
2	0.22	0.0896	0.0896	0	0.306	0.202	0.26	5.556	0.664
2	0.22	0.407	0	0	0.305	0.534	0.26	8.038	2.829

Notes: Each row corresponds to a different calibration of the extended  $AES$  model presented in Section 5.2.2. The first seven columns correspond to the parameter values. The last three columns correspond to the predictions of the model under those parameter values. The predicted evasion rate ( $\frac{E}{Y}$ ) is always 26% because all the specifications were calibrated to match that rate.