

Tax Information Avoidance

David Hagmann

Alex Rees-Jones

Jeffrey Yang

Harvard Kennedy School

The Wharton School and NBER

The Wharton School

November 2, 2018

Preliminary and Incomplete Draft: Do not cite or circulate

Abstract: While economists typically model information as having purely instrumental use and thus weakly positive value, accounting for psychological response can result in conscious willingness to pay to avoid aversive information. We test for this type of avoidance to information about tax incentives. We present an experiment in which subjects face a piece-rate wage under a progressive tax, and must make decisions about how much labor to supply and whether to make tax-incentivized donations of their proceeds to charity. Subjects have opportunities to receive tax information provided at regular intervals, and incentivized measures of their willingness to pay or willingness to accept such information are elicited. We find that substantial fractions of respondents are willing to pay to avoid information on opaque tax systems, and that this phenomenon can blunt the behavioral response to tax incentives.

JEL Classification: D9, H2.

Keywords: behavioral public finance, information avoidance, taxation, labor supply, charitable giving.

* We thank Ben Lockwood for helpful comments throughout the project. We thank the Wharton Behavioral Lab, the Wharton Dean's Grant, and the Pension Research Council/Boettner Center for Pension and Retirement Research Center for financial assistance.

Many of our most consequential decisions are influenced by the taxes that we face. When considering two employment opportunities in different states, for example, workers must compare not only the wage and benefits packages, but also the different tax rates. When deciding whether or not to work an extra shift, workers must consider how much of the extra wage will make it to their pockets post-tax. When comparing different options for retirement savings, individuals must account for potentially different tax consequences of different savings options. When deciding whether to donate to charity, buy a house, or to pursue any other tax-incentivized behavior, knowledge of the relevant deductions or credits is necessary for calculating the after-tax price of these activities.

While information about tax law is clearly necessary for optimal decision making, a growing body of evidence suggests that taxpayers are imperfectly informed about key tax parameters (see, e.g., Fujii and Hawley 1988; Blaufus et al. 2015; Gideon 2017; Rees-Jones and Taubinsky 2018). Recent research confirms an intuition many have developed through personal experience: paying taxes is highly aversive (Sussman and Olivola 2011; Kessler and Norton 2016), and many people leave money on the table to avoid filling out additional tax forms (Benzarti 2018).

In this project, we seek to connect these phenomena to a well-developed literature in economics and psychology on *information avoidance* (for a review, see Golman, Hagmann, and Loewenstein 2017). The key contribution of this literature is acknowledgment of the non-instrumental, psychological effects of information. When modeling Bayesian decision makers, economists often treat information as solely valuable through its assistance in making optimal decisions. Crucially, it has no direct effects on utility. This formulation abstracts from the possibility that bad news might be psychologically aversive, might generate dread, and might lead

to new paths of behavior that a time-inconsistent decision maker might wish to avoid in the short term (even if optimal from a time-consistent perspective). To take an archetypal example from this literature, Dorsey, Oster, and Shoulson (2013) show that patients at 50% risk of Huntington's disease are surprisingly averse to resolving the uncertainty with genetic testing. In a similar vein, Ganguly and Tassoff (2018) demonstrate that undergraduate experimental participants are often eager to avoid learning the results of common STD tests, and some are willing to pay to stop from learning the results. While it is clear that an idealized Bayesian decision-maker would want the results of tests like these, one can sympathize with a less-than-perfect decision-maker that dreads receiving such information.

We design an online experiment to test if similar forces are relevant for the consideration of information about taxes. In this experiment, workers on Amazon Mechanical Turk complete a piece-rate task. In some cases, this task is compensated with a transparent wage. In others, respondents' wage is subject to a perhaps-less-transparent progressive tax schedule. We establish that many respondents have incorrect beliefs about the taxes they face after each round of tasks, generated by a failure to remember or attend to the tax schedule. Between rounds of task completion, we elicit a measure of respondents' willingness to pay for information that helps resolve uncertainty about tax consequences, and elicit this in a way that allows respondents to communicate a need to be paid to accept this information. At the end of the experiment, we provide respondents with the opportunity to donate some of their earnings to the American Red Cross, treating charitable giving as an archetypal example of a tax incentivized behavior. We additionally randomly vary whether the incentive to donate is fully transparent versus implemented through a perhaps-opaque tax deduction, and similarly elicit willingness to pay or accept for information relevant to that deduction.

We document three groups of findings. First, we find a substantial fraction of respondents would only accept tax information in cases where they would be compensated to do so, despite its positive instrumental value. We code such respondents as information avoidant. When elicited at the end of a round of tasks, subjects were information avoidant in 24.7% of cases, with 39.2% of subjects' information avoidant in at least one round. Second, we test if the propensity to be information avoidant is correlated with underlying ability to complete our piece-rate task, and with whether information about the tax was purely bad news about wage levels versus if it could communicate a positive redistributive subsidy. Both correlations could arise if tax information avoidance were driven by motivated reasoning in this domain; however, we find no statistically significant evidence of their presence. Finally, we find that 27.0% of subjects avoid information about the tax deduction granted for charitable donation, and furthermore that total charitable giving is lower when incentivized through a tax deduction versus a transparent subsidy. These results suggest that tax information avoidance provides a channel for “moral wiggle room,” allowing subjects to justify failure to donate by an underestimate of the subsidy that they choose to fail to correct.

The paper proceeds as follows. Section I describes our experimental design. Section II presents the main results. Section III concludes by discussing a variety of tax behaviors that may be partially explained by our results.

I. Experimental Design

Subjects were recruited online via Amazon.com's Mechanical Turk to participate in an experiment for a base payment of \$0.10. In addition to the base payment, participants had an opportunity to earn a bonus payment based on their performance in experiment's real effort task.

The experiment consists of two stages. In the first stage, subjects complete 5 rounds of a real effort task, in which they earn a piece-rate wage for completing “slider tasks”, which involve moving a slider (see Fig. 1 for an example of the slider task) to a specified location. Subjects’ earnings from the real effort task may be subject to a tax, depending on the treatment. Between each round, subjects have the opportunity to learn about their wage and tax structure (if applicable). Prior to the second stage of the experiment, one of the five rounds is selected at random to be incentivized. In the second stage, subjects are given the opportunity to give up a portion of their earnings to donate a specified amount to the American Red Cross. Prior to their decision, subjects have the opportunity to learn about how the decision to donate will impact their earnings. In the first stage, subjects are randomly assigned to one of three treatments: the *Quadratic Tax*, *Demogrant Tax*, and *Random Wage* treatment, and in the second stage subjects are randomly assigned to either the *Baseline* or the *Transparent* treatment, making for a total of 6 treatments.

The experimental procedure is as follows. Subjects are first given instructions describing the real effort task and their compensation structure. In the *Quadratic Tax* treatment, subjects earn a piece-rate wage of 30 cents per slider completed, and pay a tax of $0.5 * (\text{correct sliders})^2$ cents. On top of their wages and taxes, subjects in the treatment receive a lump-sum payment of \$2.00 for completing the experiment. In the *Demogrant Tax* treatment, the compensation structure is identical, and differs only in the framing of the tax. In this treatment, the \$2.00 lump-sum payment is bundled with the tax, which is given as $0.5 * (\text{correct sliders})^2 - 200$ cents. In the *Random Wage* treatment, subjects pay no taxes and their known piece-rate wage (in cents) is randomly drawn from {5, 10, 15, 20, 25, 30}. In this treatment, subjects also receive a \$2 lump-sum payment. After reading the instructions, subjects answer questions about the task as part of a

comprehension check. The experiment ends for subjects that fail to answer the questions correctly; the remaining subjects proceed to the real effort task.

In each of the five rounds of the real effort task, subjects can complete as many sliders as they choose within a 1 minute time limit, up to a maximum of 48 sliders. After each round, subjects are asked to supply their beliefs about their marginal and total wages earned in the round (all treatments) and their marginal and total taxes paid in the round (*Quadratic Tax* and *Demogrant Tax* treatments only). After supplying their beliefs, subjects are given the opportunity to learn about their actual earnings in the round through a mechanism that elicits their willingness to pay for the information. Subjects are posed a sequence of questions that ask whether or not they are willing to receive the information for a price, so as to narrow down the subject's willingness to pay (see Fig. 2). After subjects answer the questions, a price for the information (in cents) is randomly drawn from $\{-45, -30, -15, 0, 15, 30, 45\}$, and subjects who elected to receive the information at an equal or lower price pay receive the information at the drawn price; otherwise, the subject does not receive the information. In the *Quadratic Tax* and *Demogrant Tax* treatments, willingness-to-pay is elicited for both wage and tax information, and in the *Random Wage* treatment only willingness-to-pay for wage information is elicited.

Once subjects complete all 5 rounds, they learn which round was selected to be incentivized, and are given a summary of their earnings in that round, including the total amount spent or received to acquire information throughout the first stage.

In the second stage of the experiment, subjects are given the option to give up one slider's worth of wages in order to donate that amount scaled by a conversion rate c to the American Red Cross. The conversion rate is known and randomly drawn from $\{1.25, 1.5, 2, 2.25, 2.5, 2.75\}$. If the subject was assigned to the *Quadratic Tax* and the *Demogrant Tax*

treatments in the first stage, their donation is tax-advantaged: if they choose to donate, the marginal tax paid on the last slider completed will be refunded. Prior to their decision, subjects have the opportunity to learn about how the decision will impact their earnings, and are given the same sequence of questions as in Fig. 2 to determine their willingness to pay for the information.

The two possible treatments for the second stage, the *Baseline* and *Transparent* treatments, are identical in payoff structure but differ in the information given to the subject. In the *Baseline* treatment, subjects are not told the explicit dollar value they would be giving up if they chose to donate, only that they would give up a slider's worth of wages (and be refunded the marginal tax, if applicable) in order to donate the specified amount to the American Red Cross. In the *Transparent* treatment, subjects are given the same explanation as in the *Baseline* treatment, but are also told how their donation will affect their earnings in dollar terms, including the effect of any tax refunds, if applicable.

II. Main Results

A total of 2,287 subjects were recruited for the study. 66.0% of subjects passed the comprehension check, and the proportion of passing subjects does not differ significantly between first stage treatment groups ($F=1.01$, $p=.363$). Accounting for both passing the initial comprehension check and later attrition, we are left with 537, 460, and 512 subjects in the *Quadratic Tax*, *Demogrant Tax*, and *Random Wage* first stage treatments, respectively, closely matching our preregistered sample size target of 1,500. Of these 1,509 subjects, 751 and 758 subjects were assigned to the *Baseline* and *Transparent* second stage treatments, respectively. Sessions lasted approximately 26 minutes on average, and the average bonus was \$6.14, which was paid on top of the base payment of \$0.10.

See Table 1 for summary statistics of the number of sliders completed in the real effort task. Subjects completed 15.49 sliders on average across all rounds and treatments. We do not find a significant difference in labor provision across the three first-stage treatments, and subjects tend to complete more sliders on average in the later rounds. We do not find evidence that labor provision is sensitive to variation in wages in the *Random Wage* treatment (See Fig. 3). See Table 2 for summary statistics of the proportion of subjects that chose to donate to charity. On average, 46.3% of subjects chose to donate.

We preregistered our sample construction, dependent variables and hypotheses of interest, and main analyses on aspredicted.org. We summarize those tests below.

Establishing the Presence of Information Avoidance

A chief question of interest is whether subjects avoid information about their earnings when given the opportunity to obtain the information; recall that such opportunities arise after each round in the first stage of the experiment, and prior to the donation decision in the second stage. Here, we say a subject avoided information if they elected to not receive the information at a negative price—that is, they chose not to receive the information when they would receive some positive payment for doing so. We find that a material proportion of subjects avoid information about their earnings. Across all instances in the first stage where earnings information was offered, subjects avoided wage and tax information 27.0% and 24.7% of the time, respectively, and 42.7% and 39.2% of subjects avoided wage and tax information at least once across the 5 rounds, respectively. See Fig. 4 and 5 for a summary of subjects' willingness to pay for the wage and tax information. Rates of information avoidance were higher on average in later rounds (see Tables 4 and 5). While avoidance of wage information has been previously

documented in real-effort tasks (see Huck, Szech, and Wenner 2018), we believe these findings constitute the first experimental evidence of deliberate avoidance of tax information.

In the second stage, where subjects have the opportunity to learn about how the decision to donate will impact their earnings, 27.0% of subjects avoid information. See Fig. 6 for a summary of subjects' willingness to pay for the earnings information.

Establishing the Correlates of Information Avoidance in Labor Supply

A natural question is whether tax information avoidance is correlated with the framing of the tax. In the *Quadratic Tax* treatment, the tax is framed solely as a loss to the subject's earnings, whereas in the *Demogrant Tax* treatment, the tax is framed as a gain initially, that decreases as the subject completes more sliders. Given the tendency to seek good news and avoid bad news (Golman, Hagmann, and Loewenstein 2017), one might expect that subjects are more likely to avoid information about their taxes in the *Quadratic Tax* treatment than in the *Demogrant Tax* treatment. We also examine whether information avoidance is correlated with individual ability to complete the slider task, which we measure by the amount of sliders completed in the first task. Because labor supply in rounds following the first are influenced by endogenous factors such as whether earnings information was revealed to the subject, we chose to incorporate labor supply in the first round only into our measure of ability.

Table 6 shows this analysis, which uses data only from the subjects in the *Quadratic Tax* and *Demogrant Tax* treatments. Here, each observation is an instance where subjects are offered information about their taxes; all standard errors are clustered at the subject level. Columns (1) and (2) report interval regression estimates with elicited willingness to pay for tax information as the dependent variable (see Fig. 2 for an explanation of how bounds on willingness to pay are

inferred from the subject's choices in our elicitation mechanism). As regressors we have a dummy for the tax framing and the number of sliders completed in round 1, our measure of individual ability. Columns (3) and (4) report OLS estimates for a binary dependent variable coding whether tax information was avoided against the same set of regressors.

In general, we do not find evidence that tax information avoidance is correlated with tax framing or individual ability. In regressions involving either of our information avoidance measures, coefficient estimates for the tax framing dummy and our measure for individual ability are not statistically significant, though the estimates for the tax framing dummy are in the hypothesized direction.

Establishing the Correlates of Information Avoidance in Charitable-Giving Incentives

We now focus on the second stage of the experiment: the donation decision. Recall that prior to donating, subjects are given the opportunity to obtain information on how their donation decision will impact their earnings. Because the donation decision is tax-advantaged for subjects in either of the first-stage tax treatments, one might expect that subjects in these treatments may prefer to avoid learning about the extent to which their donation decision will be tax advantaged, in order to maintain “moral wiggle room” to justify a decision not to donate. However, because the extent of the tax benefit is explicitly supplied to subjects in the *Transparent* treatment, we should expect that this particular motive for information avoidance to be diminished in this treatment versus the *Baseline* treatment. As such, one might expect information avoidance to be less prevalent in the *Transparent* treatment than in the *Baseline* treatment.

Table 7 shows this analysis, which includes only subjects in the *Quadratic Tax* and *Demogrant Tax* treatments. Columns (1) and (2) report interval regressions estimates of the

subject's willingness to pay for information against a second-stage treatment dummy and various controls, including the extent to which the donation decision is tax-advantaged (given by the marginal tax paid on the last slider the subject completed) and the donation amount (determined by the randomly drawn conversion rate). Here we find that while the estimates for the treatment effect are negative, as hypothesized, the effect is not statistically significant. Columns (3) and (4) report OLS regressions for the binary measure of information avoidance against the same regressors as in columns (1) and (2); again, we find that the estimates for the treatment effect are consistent with the hypothesized direction but is not statistically significant.

While we do not find evidence that information avoidance is correlated with the second-stage treatment, we do find that the likelihood of donating is correlated with treatment, as well as with whether or not the subject avoided the earnings information prior to the donation decision. Table 8 reports OLS estimates for a binary dependent variable coding whether the subject donated against the same regressors as in Table 7. To test whether subjects who avoided information would be more or less likely to donate, a binary regressor coding whether the subject avoided information is also included.

Subjects are on average 14.0 percentage points more likely to donate in the *Transparent* treatment vs. the *Baseline* treatment, a statistically significant effect size that holds when financial incentives for donation are included as controls. Note that motivated reasoning is not the exclusive explanation for this discrepancy, as the alternative explanation is that subjects are more likely to donate when there is less uncertainty about how donating will affect the subject's final payout. However, several elements of our results support the former interpretation. Importantly, subjects have the opportunity to learn about how the donation decision will affect their payout at no cost to themselves, and can even obtain the information with an additional

bonus payment. As a result, the extent to which we see information avoidance among subjects along with the treatment effect is inconsistent with the interpretation that subjects donate less due to greater uncertainty, as subjects have the opportunity to resolve this uncertainty. In fact, we find that subjects who avoided information were *less* likely to donate (column 4 of table 8), consistent with the motivated reasoning interpretation, which posits that subjects avoid the information because it allows them the moral wiggle room to justify a decision not to donate.

III. Discussion

The results of our experiment highlight a stark, but intuitive, fact: thinking about taxes is aversive (Sussman and Olivola 2011; Kessler and Norton 2016), and this aversion can lead people to choose not to think about them. While surely not the sole explanation, such results provide a compelling account of a variety of anomalous tax behaviors. Taxpayers are ill informed of their tax rates, despite the large incentives present to learn them (Fujii and Hawley 1988; Blaufus et al. 2015; Gideon 2017; Rees-Jones and Taubinsky 2018). Taxpayers often leave substantial tax benefits on the table (Chetty, Friedman and Saez 2013; Bhargava and Manoli 2015; Benzarti 2018). Taxpayers often demonstrate inelastic response to opaque tax incentives, at times even moving opposite of the incentivized direction (Miller and Mumford 2015). Finally, taxpayers often appear to fail to attend to sales taxes, even though they are aware of the tax rate and capable of the relevant computations (Chetty, Looney, Kroft 2009; Goldin and Homonoff 2013; Feldman, Goldin, and Homonoff 2018; Taubinsky and Rees-Jones 2018). In sum, evidence of failure to collect or attend to tax information abounds. We suggest a simple explanation for this phenomenon: despite the instrumental value, taxpayers actively avoid this information.

References

- Benzarti, Youssef, “How Taxing is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs.” NBER Working Paper No. 23903.
- Blaufus, Kay, Jonathan Bob, Jochen Hundsdoerfer, Christian Sielaff, Dirk Kieseewetter, and Joachim Weimann, “Perception of Income Tax Rates: Evidence from Germany,” *European Journal of Law and Economics*, 2015, 40(3), 457–478.
- Bhargava, Saurabh and Day Manoli, “Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment,” *American Economic Review*, 2015, 105 (11), 3489– 3529.
- Chetty, Raj, John N Friedman, and Emmanuel Saez, “Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings,” *American Economic Review*, 2013, 103 (7), 2683–2721.
- Chetty, Raj, Adam Looney, and Kory Kroft, “Salience and Taxation: Theory and Evidence,” *American Economic Review*, 2009, 99, 1145-1177.
- Feldman, Naomi, Jacob Goldin, and Tatiana Homonoff, “Raising the Stakes: Experimental Evidence on the Endogeneity of Taxpayer Mistakes,” *National Tax Journal*, 2018, 71(2), 201-230.
- Fujii, Edwin T and Clifford Hawley, “On the Accuracy of Tax Perceptions,” *The Review of Economics and Statistics*, 1988, 70(2), 344–347.

- Gideon, Michael, “Do Individuals Perceive Income Tax Rates Correctly?” *Public Finance Review*, 2017.
- Goldin, Jacob and Tatiana Homonoff, “Smoke Gets in your Eyes: Cigarette Tax Salience and Regressivity,” *American Economic Journal: Economic Policy*, 2013, 5 (1), 302–336.
- Golman, Russell, David Hagmann, and George Loewenstein. 2017. “Information Avoidance.” *Journal of Economic Literature* 55 (1): 96–135.
- Huck, Steffen, Nora Szech and Lukas Wenner, “More Effort with Less Pay: On Information Avoidance, Optimistic Beliefs, and Performance,” Working Paper, 2018.
- Kessler, Judd and Michael Norton, “Tax Aversion and Labor Supply,” *Journal of Economic Behavior and Organization*, 2016, 124, 43-54.
- Liebman, Jeffrey B. and Richard Zeckhauser, “Schmeduling,” Working Paper, 2004.
- Miller, Benjamin and Kevin Mumford, “The Salience of Complex Tax Changes: Evidence From the Child and Dependent Care Credit Expansion,” *National Tax Journal*, 2015, 68 (3), 477–510.
- Mirrlees, J.A., “An Exploration in the Theory of Optimum Income Taxation.” *The Review of Economic Studies*, 1971, 38(2), 175-208.
- Rees-Jones, Alex and Dmitry Taubinsky, “Measuring ‘Schmeduling’,” NBER Working Paper No. 22884, 2018.
- Saez, Emmanuel, “Do Taxpayers Bunch at Kink Points?” *American Economic Journal: Economic Policy*, 2010, 2 (3),180–212.

Sussman, Abigail and Christopher Olivoa, "Axe the Tax: Taxes are Disliked more than Equivalent Costs," *Journal of Marketing Research*, 2011, 48, 91-101.

Taubinsky, Dmitry and Alex Rees-Jones, "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment," *Review of Economic Studies*, 2018, 85(4): 2462-2496.