

Using Divide and Conquer to Improve Tax Collection: Theory and Laboratory Evidence

Sylvain Chassang^{*,†} Lucia Del Carpio Samuel Kapon
Princeton University INSEAD Princeton University

June 28, 2022

Abstract

We consider a government collecting taxes from a large number of tax-payers using limited enforcement capacity. Under random enforcement, limited capacity results in multiple equilibria: if most agents comply, limited enforcement is sufficient to dissuade individual misbehavior; if most agents do not comply, enforcement capacity is overstretched and fails to dissuade misbehavior. In settings without behavioral frictions, prioritized enforcement strategies can implement high collection as the unique rationalizable outcome. Motivated by a field implementation opportunity, we investigate both theoretically and experimentally the extent to which this insight extends to environments with incomplete information and bounded rationality.

KEYWORDS: tax collection, government capacity, divide and conquer.

1 Introduction

A government’s ability to achieve goals such as tax collection, low crime rates, or environmental protection depends on its capacity to enforce mandated behavior on agents that

*Chassang: chassang@princeton.edu, Del Carpio: lucia.delcarpio@insead.edu, Kapon: kapon@princeton.edu.

†We are indebted to Alessandra Casella, Pedro Dal Bo, and Catherine Hafer for thoughtful discussions that lead to substantial revisions of this manuscript. We are grateful to Vishal Ashvinkumar, Kathleen Ngangoue, Paula Onuchic, and Tianzan Pang for early feedback on our experiment design. The paper greatly benefited from discussions with Dilip Abreu, Abhijit Banerjee, Daniel Diermeier, Guillaume Frechette, Marina Halac, Navin Kartik, John Klopfer, Alessandro Lizzeri and Ariel Rubinstein. We thank seminar participants at Chicago Booth, LSE, MIT, Northwestern, Princeton, PSE, SITE (2021), ThReD (2021), the University of Alabama, UBC, and UPenn for very helpful comments.

refuse to comply. In many cases, government capacity is limited: a government has the manpower to enforce rules only for a small number of non-compliers. For instance, in the United States the IRS has the capacity to audit under 1% of tax returns every year.¹ In addition damages for non-compliance are often limited, ruling out enforcement through large probabilistic punishments à la Becker (1968).² In the case of the IRS, penalties for negligent underreporting of income amount to 20% of unpaid taxes. Such limited enforcement capacity can lead to multiple equilibria. If most agents comply with government policy, then limited enforcement is sufficient to dissuade isolated agents from misbehaving. If many agents do not comply, overstretched enforcement capacity has a minimal impact on incentives and behavior. This paper seeks to better understand the extent to which divide-and-conquer enforcement strategies can help select a high compliance equilibrium in the presence of realistic frictions.

Motivated by a field implementation opportunity, described in our companion paper Del Carpio et al. (2022),³ we study the problem of a government entitled to collect an amount of taxes-due D from each of N agents. The government is able to forcefully collect the amount D but doing so is costly in terms of time and resources.⁴ The difficulty is that: (i) the government is able to perform at most αN forceful collections, with $\alpha \in (0, 1)$; (ii) upon collection, the maximum amount of damages the government can claim is D . Instead of forcefully collecting taxes, the government can offer agents to settle their taxes by paying a given price P . Agents who settle are not collected on. Enforcement capacity is spent on forceful collection from non-complying agents. The government’s main policy instrument

¹This varies by income bracket, from under .5% to roughly 5%. See [IRS statistics](#) for updated numbers.

²In the US, the Eighth Amendment affords protections against excess punishment in order to limit the scope for abuse by the state itself.

³In randomized controlled trial [AEARCTR-0007305](#) we partner with a district of Lima, Peru to evaluate different tax collection strategies on a sample of roughly 13,000 taxpayers delinquent in paying property taxes.

⁴In the context of our field application, forceful collection can take the form of garnishment of bank accounts, as well as physically seizing assets. Both are costly in terms of money and time, and require a lengthy due process. In the context of property taxes, it is reasonable to assume that the government knows the amount of taxes-due, but is unsure about the tax-payer’s ability or willingness to pay them. In the context of income taxes, the government may also be uncertain about the amount of taxes due. We study this extension in Appendix A and show that our analysis extends essentially as is.

is to commit to an enforcement rule, i.e. an order in which to enforce taxes against non-compliant agents.

We first establish benchmark results in a static frictionless environment in which agents are capable of settling with probability 1. We contrast *uniform* enforcement rules, in which a number αN of agents are randomly chosen from non-complying agents, with *prioritized* enforcement rules, in which agents are given common-knowledge priorities ahead of settling and non-complying agents face collection in order of priority. For any given settlement price $P \in [\alpha D, D]$, uniform random enforcement leads to multiple equilibria: a high settlement equilibrium in which all agents settle, and minimal enforcement capacity effectively dissuades individual deviations; a low settlement equilibrium in which agents choose not to settle, thereby weakening incentives to comply. In contrast, under prioritized enforcement, there is a unique, high settlement equilibrium. Regardless of the overall behavior of tax-payers, incentives are always tightly focused on a small group of marginal agents: it is dominant for the αN agents with the highest priority to settle.⁵ This initiates an unraveling process: it is a best response for the $2\alpha N$ agents with the highest priority to settle, and so on.

Our next set of results considers the impact of frictions on the effectiveness of prioritized enforcement. Specifically, we assume that with exogenous probability q an agent is simply unable to settle.⁶ Agents unable to settle interrupt the unraveling argument described above. We show that as the number of agents N gets large, there is essentially a unique equilibrium. With probability approaching 1, the share α/q of agents with the highest priority settle (if they can), while remaining agents do not. When capacity is binding ($\alpha < q$) per capita tax revenue is proportional to $(1 - q)D/q$. Denominator q reflects the equilibrium cost of frictions: expended capacity cannot be redeployed to incentivize other taxpayers to settle. We

⁵This relates to the point made by Eeckhout et al. (2010) that when government capacity is limited, random public crackdowns may be more effective than the thinly spread incentives provided by uniform enforcement. Our setting is different because once agents comply, the enforcement capacity needed to ensure compliance can be recycled to induce other agents to comply.

⁶This could be because tax-payers suffer a liquidity shock, or face a personal crisis preventing them to attend to their obligations. Alternatively, q may correspond to the probability of an administrative error, for instance charging a previous resident for property taxes even though they have moved.

study direct mechanism design in a setting where taxes owed D and frictions q are heterogeneous among agents $i \in [0, 1]$. For N large, prioritized enforcement mechanisms approach the highest performance attainable under partial Bayesian implementation in any mechanism. The optimal priority rule ranks agents according to individualized score $(1 - q_i)D_i/q_i$, reflecting the trade-off between expected collection amount $(1 - q_i)D_i$, and expected capacity expenditure q_i . In contrast to the results of Halac et al. (2020), incomplete information about priorities does not increase collection – this is because the settlement price P that can be extracted from agents increases linearly with their expectation of enforcement. We show in Appendix A that the analysis extends nearly as is to the context of income taxes where the tax-payer has private information about the amount of taxes owed.

To investigate design under bounded rationality, we extend our static game-form and allow settlement decisions to take place over time. We study the impact of providing dynamic information about the settlement behavior of others on settlement behavior.⁷ We show that under rationalizability as our solution concept, information design is essentially irrelevant. Provided settlement offers get worse over time, then as the number of agents N gets large, for any information structure, the share α/q of agents with the highest priority settle as soon as they are able to with probability approaching 1. In contrast, information design is not irrelevant under boundedly rational solution concepts. We focus on implementation in non-obviously dominated strategies (Li, 2017), and show that by providing agents updated information about their effective rank (reflecting the settlement behavior of others ahead of them) it is possible to ensure that tax revenue under any non-obviously dominated strategy profile approaches the second best tax revenue. By revealing the behavior of those ranked higher over time, information makes settlement sequentially dominant, and allows time to replicate the work of many iterations of rationality. This echoes the point made by Glazer and Rubinstein (1996) that the extensive form can provide a guide on how to solve an underlying normal form game.

⁷A recent empirical literature evaluates how information about their neighbors' behavior affects taxpayers' behavior (Del Carpio, 2014, Castro and Scartascini, 2015, Dwenger et al., 2016), highlighting the importance of norms in enforcing compliance.

Implementation in non-obviously dominated strategies also allows us to identify a complementarity between information and incentives to pay early: the second best is approached in all non-obviously dominated strategies only if agents receive updated rank information *and* have strict incentives to pay taxes as early as possible. If no incentives to settle early are provided, then agents can delay their settlement behavior, thereby stopping the flow of information. This echoes the importance of activity rules in simultaneous ascending auctions emphasized in Milgrom (2000).

With field implementation in mind we complement our theoretical analysis with laboratory experiments. They largely validate qualitative features of our analysis: prioritized enforcement improves collection rates over random enforcement, but only when players receive updated information about their effective rank. Although implementation in non-obviously dominated strategies gets *many* important qualitative predictions right, it misses certain aspects of behavior that are likely to matter in the field. Along the lines of quantal response equilibrium (McKelvey and Palfrey, 1995), players sometimes play obviously dominated strategies when the incentives not to are relatively mild. This causes information to trickle down more slowly and reduces the impact of information provision. As a result there is a motive to increase incentives to settle early, even though this is suboptimal under obviously dominant implementation. This is partly compensated by the fact that agents appear to apply multiple rounds of elimination of dominated strategies which speeds up the production of information. We delineate possible implications of these observations for design.

Our work contributes to the mechanism design literature that seeks to implement desirable outcomes in solution concepts less demanding than Bayes Nash equilibrium. Divide-and-conquer schemes play an important role in the work of Abreu and Matsushima (1992) on virtual implementation in rationalizable strategies. Divide-and-conquer schemes also play an important role in the literature on contracting with externalities, including Segal and Whinston (2000), Spiegler (2000), Segal (2003), Winter (2004), Dal Bó (2007) and more recently Eliaz and Spiegler (2015), and Halac et al. (2019, 2020). This literature often emphasizes

the cost of implementing outcomes in rationalizable strategies due to the fact that players must be compensated for potential strategic uncertainty. This trade-off does not appear in our enforcement context: the costly part of incentives tends to remain off-of-the equilibrium path. In contrast, we study the impact of behavioral frictions that perturb the iterated best-reply rationale making divide-and-conquer schemes theoretically attractive.

We also contributes to the extensive literature on the economics of tax-compliance summarized in Slemrod (2019). The workhorse Becker-Allingham-Sandmo model of tax compliance (Becker, 1968, Allingham and Sandmo, 1972) is a single agent model in which a tax payer weighs the pros and cons of paying taxes or not, possibly including mental payoffs associated with tax-morale, and compliance challenges (Luttmer and Singhal, 2014, De Neve et al., 2021). We depart from this view and argue that under limited enforcement capacity, compliance becomes a strategic interaction across tax-payers, which may exhibit multiple equilibria. This echoes Slemrod (2019)'s concern that experiments which explore the impact of direct threats of enforcement on compliance, such as (Kleven et al., 2011), would not be credible to tax-payers if implemented at scale. We study how to best deploy limited enforcement capacity in the presence of incomplete information, and bounded rationality.

Beyond implementation in rationalizable strategies, a growing body of work recognizes the importance of designing mechanisms suitable for boundedly rational agents. Crawford and Iriberri (2007) shows that a level- k model can explain overbidding in experimental auctions. Mathevet (2010) and Healy and Mathevet (2012) emphasize design steps that can be taken to ensure a mechanism is supermodular, leading to stable learning dynamics. De Clippel et al. (2019) study implementation using k iterations of best reply as a solution concept. Closely related to our analysis of dynamic information provision, Glazer and Rubinstein (1996) show that extensive forms play an important role as a guide to play, simplifying the computational burden on players. Li (2017) introduces obviously strategy-proof mechanisms, a class of mechanisms suitable for boundedly rational agents, and shows that different extensive form implementations of the same normal form game can vary in their suitability for boundedly

rational players.⁸

Most importantly, this paper hopes to inform and stimulate the application of divide-and-conquer mechanisms in real-life settings.⁹ While divide-and-conquer plays an important role in the theory of mechanism design, it has received relatively limited experimental and empirical attention. Sefton and Yavaş (1996) and Katok et al. (2002) implement the static Abreu and Matsushima (1992) mechanism (relying on many rounds of iterated dominance) and an extensive-form variant (relying on many rounds of backward induction). Consistent with the criticism of Glazer and Rosenthal (1992), experimental play in both mechanisms does not correspond to rationalizable behavior. However, findings depend on the underlying strategic environment. In coordination games, Abreu and Matsushima (1992) mechanisms are unable to select Pareto and risk dominated equilibria, but they can help select Pareto efficient but risk dominated equilibria. We provide evidence that divide-and-conquer may be effective in our tax collection context. More importantly, we identify design features that enhance the effectiveness of divide-and-conquer: providing agents updated rank information, and ensuring they have strong enough incentives to comply early.

The paper is structured as follows. Section 2 sets up a simple model of tax collection. It establishes benchmark results clarifying the value of common-knowledge enforcement priorities in a static setting with fully responsive agents. Section 3 clarifies the impact of non-response frictions on the effectiveness of enforcement priorities and establishes bounds on tax-revenue under any mechanism. Section 4 casts compliance decisions in a dynamic context, and studies the value of dynamic information provision under rationalizable and

⁸This theoretical literature builds on a large body of experimental evidence. Kagel et al. (1987) documents that prices in a second-price sealed-bid auction exceed those predicted by dominant strategy play, but that prices in an ascending clock auction are well predicted by dominant strategy play. Kagel and Levin (2001) documents a similar observation in multi-unit demand auctions.

⁹While we are unaware of published formal evaluations of divide-and-conquer mechanisms in the field, we are hopeful that this class of mechanisms can make a real difference in governments' ability to leverage limited capabilities. Operation Ceasefire (Braga et al., 2001, Kennedy, 2011, 2012), a multi-city homicide reduction program explicitly prioritizes the assignment of law enforcement capabilities to homicides in the order in which they are committed, thereby dissuading gangs to initiate gang wars. The analysis of Chassang et al. (2022) suggests that although it met with varying success, Operation Ceasfire had significant option value. This gives us reasonable hope that appropriately designed implementations of divide-and-conquer can be effective in practice.

non-obviously dominated strategies. Section 5 reports findings from laboratory experiments and interprets them from the perspective of design. Section 6 discusses limits of our analysis and possible applications other than tax collection. Appendix A extends the analysis to the case of income taxes where tax-payers have private information about the amount of tax they owe and report taxable income to the tax collection authority. Proofs are presented in the main text, or included in Appendix B. A subset of player instructions is included in Appendix C. Appendix D reports additional experimental findings.

2 Benchmark Model

We first study a benchmark static model with frictionless agents. Sections 2.1 and 2.2 setup the model, and highlight the value of enforcement priorities: they help select a unique equilibrium. Section 2.3 discusses our key assumptions and motivates them in the context of our field application.

2.1 Framework

N agents indexed by $i \in I \equiv \{1, \dots, N\}$ each owe a principal a fixed amount D . The agents and the principal are all risk-neutral. The principal can potentially collect D from an agent by force, but the principal has limited enforcement capacity.¹⁰ Specifically, the principal can forcefully collect from only $\alpha N \geq 1$ agents with $\alpha \in (0, 1)$. Forceful collection allows the principal to collect amount D but does not impose additional punishments. We focus on a specific class of collection mechanisms, which we later show approaches optimal collection under any mechanism as N grows. The principal can make settlement offers and commit to an enforcement schedule according to the following extensive-form game:

- (i) The principal gives each agent the possibility to settle by paying a fixed price P .

¹⁰This could be because forceful collection requires resources (e.g. physically seizing assets is difficult), or because due process steps must be taken (e.g. formal audits may be required, and their conclusions may be litigated).

Agents who settle are spared from forceful collection.

- (ii) Agents simultaneously decide whether or not to settle and pay price P .
- (iii) The principal forcefully collects D from agents who do not settle, according to a complete order \prec over I .

We consider two possible enforcement priorities \prec :

- Random priorities \prec_R : agents are drawn sequentially ex post (i.e. period (iii)), with uniform probability and without replacement;
- Common knowledge priorities \prec_{CK} : the ordering is specified ex ante (i.e. period (i)) and is common knowledge among players. For simplicity, we assume that agents are ranked in descending order of their index $i \in \{1, \dots, N\}$ (i.e. agent 1 has the highest priority).

Payoffs and solution concept. We denote by $s_i \in \{0, 1\}$ agent i 's decision to settle for the principal's offer. The principal's total payoff is

$$\Pi \equiv \frac{1}{N} \sum_{i \in I} s_i P.$$

Note that payoffs exclude the proceeds from forceful collection. This simplifies computations and reflects the fact that the net benefits of forceful collection may be ambiguous: enforcement costs may be well above the amount that can be legally collected from agents. This also clarifies that the value of enforcement comes from incentive provision rather than direct revenue.

This section and the next use both Bayes Nash Equilibrium and rationalizability as solution concepts. We consider non-obviously dominated strategies (Li, 2017) in Section 4.

2.2 The Value of Enforcement Priorities

The following results clarify the value of prioritized enforcement: it selects a high collection equilibrium as the unique rationalizable strategy profile; in contrast, random enforcement induces multiple equilibria involving both high and low collection levels.

Proposition 1 (multiple equilibria under random enforcement). *Under random enforcement order \prec_R , for any settlement price $P \in [\alpha D, D]$, there exists a Nash equilibrium such that all agents settle, and a Nash equilibrium such that all agents refuse to settle.*

Proof. Consider $P \in [\alpha D, D]$, and assume that all agents settle. Then a deviator who refuses to settle faces enforcement with probability 1. Since $P \leq D$, it is indeed individually optimal for a tax-payer to settle.

Assume now that all agents refuse to settle. Then in equilibrium, an agent faces enforcement with probability α , incurring expected cost αD . Since $P \geq \alpha D$, it is individually optimal for an agent not to settle. ■

Proposition 2 (divide-and-conquer). *For any settlement price $P \in [\alpha D, D]$, under common-knowledge enforcement order \prec_{CK} , a unique strategy profile survives iterated elimination of dominated strategies: all agents settle.*

In other words, first-best collection can be approximated arbitrarily well in rationalizable strategies.

Proof. We show that for every settlement offer $P < D$, it is iteratively dominant for all agents to settle, so that the principal raises tax revenue NP . The proof is by induction on the priority of agents. The induction hypothesis is that in all strategy profiles that survive k -iterations of elimination of dominated strategies, all agents with priority higher than k choose to settle. The induction hypothesis holds for $k = 1$ since the highest priority agent faces collection with probability 1 in the event they do not settle. In turn, if the hypothesis holds for $k \geq 1$, then an agent of rank $k + 1$ that does not comply is audited with probability

1. Hence, it is iteratively dominant for an agent of rank $k + 1$ to comply, which establishes the induction step. ■

The key assumption here is that players know their own rank, and that this is common knowledge. Whether or not agents know the rank of others does not affect Proposition 2.

In the absence of frictions, prioritized enforcement is extremely effective: an enforcement capacity $\alpha N = 1$ is enough to induce any arbitrary number of agents to settle. This is driven by the fact that enforcement capacity is used only off-of-the-equilibrium path and can be efficiently reassigned if tax-payers settle. Sections 3 and 4 study how realistic frictions perturb the effectiveness of prioritized enforcement by pushing enforcement back on the equilibrium path.

We note that there is no cost to implementation in rationalizable strategies compared to Bayesian implementation. Prioritized enforcement achieves first-best collection and limited capacity turns out to be non-binding in this frictionless environment.

2.3 Motivation and key assumptions

Before we start introducing frictions of interest, we find it useful to motivate our modeling assumptions in the context of our field application: collecting quarterly property taxes in a district of Lima, Peru.

Capacity constraints. In our candidate application, capacity is limited by due process and budgetary constraints. Direct collection measures (such as garnishing bank accounts) are only allowed after in-person warnings and the initiation of a legal procedure to collect. Practically, the city can take direct collection steps on less than 10% of delinquent tax-payers. Additionally, capacity increases must be approved by budgetary committees focused on cost controls, and are often rejected, even if they are arguably revenue enhancing.

From a theory perspective, we note that our analysis would be largely unchanged if the principal could freely choose auditing capacity at a variable cost, provided that the cost is paid regardless of whether audits happen or not (collection agents must be hired, trained,

and paid even if there are no delinquent tax-payers).

Limited punishment. Our model does not allow for dissuasive punishments in the style of Becker (1968). The analysis would not be significantly changed if limited punishments were applied. This is equivalent to setting a settlement price P strictly less than the amount D that could in principle be forcefully collected. However, our results would be changed if arbitrarily high punishments were available. Such punishments could compensate for very limited enforcement capability.

In practice, there are limits to legitimate levels of punishments. In the US, the Eighth Amendment limits the punishments that both federal and state governments can apply. In the case of tax collection, the maximum penalty that the IRS can apply in case of under-reported income is 20%. In our candidate application, the penalties for late payment are limited to moderate fixed administrative costs, on the order of 10% of the average delinquent amount.

Sample size. Our analysis will focus on environments where the number of tax-payers N is large. This is true in our candidate application: we seek to raise property taxes from roughly 13,000 delinquent taxpayers. This reduces concerns over collusion and side transfers between tax payers, but potentially raises concern about the validity of rationalizability as a solution concept: in the prioritized enforcement mechanism described above, many rounds of iterated elimination of dominated strategies would be necessary to induce low priority tax-payers to settle. This increases the importance of bounded rationality considerations.

Uncertainty about agents. Because our candidate application concerns property taxes, there is little uncertainty about the amount of taxes due D (observed heterogeneity in tax-due changes little to the analysis, see Section 3). However there may be some uncertainty about the identity of the tax-payer responsible for payment, in case records have not been appropriately updated. In addition, there is some uncertainty regarding whether a tax-payer

is aware that taxes are due, and financially capable of paying them. In principle, it may be possible to predict tax-payer behavior using extensive covariates observable by the state. This includes taxes due, employment records, past repayment behavior, and so on.

In other tax collection settings, for instance income tax, there may be uncertainty about the amount of taxes that people owe, and a formal audit may be included as part of the enforcement process. In Section 3 we consider environments in which tax-payers have private information about whether or not they can repay taxes. Appendix A shows that the analysis extends without difficulty to environments in which tax-payers have private information about the amount of taxes they really owe.

Why first play is important. As Chen and Ledyard (2010) point out, dominance solvable implementation is attractive because it exhibits good learning properties. However, it is possible that learning may require many iterations of play. In a public policy setting, the political goodwill needed to experiment with novel institutional designs is a limited resource, and good policies may be rapidly abandoned if they fail initially, even if it is because of out-of-equilibrium play. For this reason, we think that ensuring institutions behave well on first play is a first-order design objective. This is why our analysis emphasizes implementation under solution concepts weaker than rationalizability.

Other applications. We consciously specialize our model to our application of interest: the collection of known property taxes. It motivates some of our modeling choices, and clarifies the relevant dimensions of inquiry. However, it is clear that the ideas we develop apply to a broader set of contexts where governments are seeking to discipline several agents using limited resources. We describe possible applications to law enforcement and organizational change in Section 6. Appendix A covers the enforcement of unknown taxes.

3 Tax Collection with Non-Compliance Frictions

Proposition 2 suggests that even with very limited enforcement capacity ($\alpha N = 1$), prioritized enforcement can ensure a high compliance equilibrium. However, the argument relies on a high degree of confidence that non-compliance by higher ranked agents will not exhaust the principal's enforcement capacity. We now consider a variant of the game introduced in Section 2 in which agents are exogenously and independently unable to settle with probability q . This friction naturally reduces the effectiveness of prioritized enforcement.

We study the impact of frictions in two steps: first, we characterize behavior under random and prioritized enforcement; second, we allow for heterogeneous agents and show that enforcement using common knowledge priorities approaches optimal collection in Bayes Nash equilibrium under any mechanism.

3.1 Behavior under random and prioritized enforcement

Random enforcement. As in the case without frictions, random enforcement leads to multiple equilibria.

Let us denote by $x \wedge y$ the minimum of x and y .

Proposition 3. *Pick $\epsilon > 0$. Under random enforcement order \prec_R , there exists \underline{N} large enough such that for all $N > \underline{N}$*

(i) *For any settlement price $P \in \left[\alpha D, (1 \wedge \frac{\alpha}{q} - \epsilon) D \right]$, there exists a Nash equilibrium in which all agents settle their taxes if they can, and a Nash equilibrium in which no agents settle their taxes.*

(ii) *For all $P > (1 \wedge \frac{\alpha}{q} + \epsilon) D$, it is dominant for agents not to settle their taxes.*

Proof. We begin with point (i). Consider first the strategy profile in which all agents choose to settle their taxes if they can. As N grows large, the law of large numbers implies that with probability approaching 1, a share of tax-payers approaching q is exogenously unable

to settle their taxes. Hence it follows that for any $\epsilon > 0$, as N gets large, the expected payoff from not settling taxes is less than $-(1 \wedge \frac{\alpha}{q} - \epsilon)D$. Since $P \leq (1 \wedge \frac{\alpha}{q} - \epsilon)D$ this implies that there is indeed an equilibrium in which all tax-payers choose to pay their taxes if they can.

Consider now a strategy profile in which none of the agents settle their taxes. Then an agent's probability of being audited is equal to α . Since $P \geq \alpha D$, it is indeed optimal for agents not to settle their taxes.

We now turn to point (ii). Regardless of the number of agents who choose to settle if they can, for any $\nu > 0$, with probability approaching 1, the share of tax-payers who are unable, or choose not to settle is greater than $q - \nu$. This implies that for any $\epsilon > 0$, and regardless of the strategies of other players, for N large enough, a tax-payer's probability of audit is less than $1 \wedge \frac{\alpha}{q} + \epsilon$. This implies that whenever $P > (1 \wedge \frac{\alpha}{q} + \epsilon)D$ it is dominant for players not to settle their taxes. ■

Note that the per capita tax raised in the high settlement equilibrium approaches $(1 - q) \times (1 \wedge \frac{\alpha}{q}) \times D$. We now show that prioritized enforcement achieves the same collection revenue while resolving the issue of equilibrium multiplicity.

Prioritized enforcement. It is convenient to index an agent with rank $i \in I$ by her *scaled rank* $\rho = i/N$. This facilitates the statement of asymptotic results as the number N of agents grows large.

Proposition 4. *Consider prioritized enforcement order \prec_{CK} . Fix a settlement price $P \in (0, D)$ and $\epsilon > 0$. For N large enough, under all rationalizable strategy profiles,*

- (i) *agents with rank $\rho > \frac{\alpha}{q} + \epsilon$ do not settle;*
- (ii) *agents with rank $\rho < \frac{\alpha}{q} - \epsilon$ settle if they are able to.*

Proof. We first show (i), i.e. that for N large enough, it is dominant for agents with rank

$\rho > \frac{\alpha}{q} + \epsilon$ not to settle. For any given strategy profile, let us denote by

$$A(\rho) \equiv \frac{1}{N} \sum_{i=1}^{\rho N} \mathbf{1}_{s_i=0}$$

the realized share of agents with rank less than ρ who do not settle. By the Law of Large Numbers, with probability approaching 1 as N gets large, for any strategy profile, $A(\frac{\alpha}{q} + \epsilon) > \alpha$. This implies that uniformly over strategy profiles, an agent with rank $\rho > \frac{\alpha}{q} + \epsilon$ faces enforcement with probability approaching 0. Hence, for N large enough, it is dominant for such an agent to refuse any settlement offer $P > 0$.

Now consider (ii), the case of agents with rank $\rho < \frac{\alpha}{q} - \epsilon$. We define the sequence $\rho_K \equiv \alpha \sum_{k=1}^K (1-q)^{k-1}$. Note that ρ_K converges to $\frac{\alpha}{q}$ as K gets large.

For any $K \in \mathbb{N}$ and $\nu > 0$, we establish the following hypothesis $H_{K,\nu}$:

- (i) uniformly over strategy profiles surviving K iterated elimination of dominated strategies, with probability 1 as N gets large, $A(\rho_K - K \times \nu) \leq q\rho_K$.
- (ii) as N gets large enough, for all strategy profiles surviving K iterated elimination of dominated strategies, agents with rank $\rho < \rho_K - K \times \nu$ settle if they can.

Consider the case where $K = 1$. Since $\rho_1 = \alpha$, it is dominant for all players with rank $\rho \leq \rho_1$ to comply if they can. Since the exogenous non-compliance rate is q , it follows that with probability 1 as N gets large, $A(\rho_1 - \nu) \leq q\rho_1$.

We now show that $H_{K,\nu}$ implies $H_{K+1,\nu}$. Indeed, since $A(\rho_K - K \times \nu) \leq q\rho_K$ with probability approaching 1, this means that the spare enforcement capacity that can be used on agents with rank greater than $\rho_K - K \times \nu$ is greater than $\alpha - q\rho_K$ with probability approaching 1. This implies that in all strategy profiles surviving $K+1$ iterations of iterated elimination of dominated strategies, non-complying agents with rank less than $\rho_K - K \times \nu + \alpha - q\rho_K = \rho_{K+1} - K \times \nu$ get audited with probability approaching 1. Hence, all strategy profiles surviving $K+1$ rounds are such that agents with rank less than $\rho_{K+1} - K \times \nu$

settle. By the Law of Large Numbers, this implies that $A(\rho_{K+1} - (K + 1)\nu) \leq q\rho_{K+1}$ with probability 1 as N gets large.

To conclude, observe that we only need a fixed number of induction steps to establish point (ii). Consider \underline{K} large enough that $\rho_{\underline{K}} \geq \frac{\alpha}{q} - \frac{\epsilon}{2}$, and set $\nu = \frac{\epsilon}{2\underline{K}}$. The induction hypothesis $H_{\underline{K},\nu}$ implies point (ii). ■

It is worth noting that although prioritized enforcement achieves the same revenue as the high settlement equilibrium under random enforcement, it implements different settlement behavior. Under prioritized enforcement a share $1 \wedge \frac{\alpha}{q}$ of tax payers pay the full amount of tax due D if they are able to. In the high settlement equilibrium under random enforcement all tax payers pay an amount $(1 \wedge \frac{\alpha}{q})D$, if they are able to. This raises potential equity concerns, which we revisit in Section 5.¹¹

When the friction rate q depends on the settlement price P via an increasing function $Q(P)$, Proposition 4 can be viewed as characterizing a second stage given a chosen price P . The optimal P can then be chosen according to the following corollary.

Corollary 1 (endogenous frictions). *As N gets large, it is approximately optimal to make a settlement offer P solving*

$$\max_P P \times (1 - Q(P)) \times \min \left\{ \frac{\alpha}{Q(P)}, 1 \right\}. \quad (1)$$

The principal's problem is related to that of a monopolist selling compliance certificates and facing demand curve $Q(P)$. The monopolist's revenue correspond to the first two factors $P \times (1 - Q(P))$. Factor $\min \left\{ \frac{\alpha}{Q(P)}, 1 \right\}$ reflects the cost of incentive provision for the principal and compliance externalities across tax-payers: non-compliant agents not only fail to settle, but they reduce the incentives of lower priority agents to settle.

¹¹Note that the principal's preferred equilibrium under random enforcement exploits the principal's knowledge of the friction rate q : the settlement offer is equal to $D \times (1 \wedge \frac{\alpha}{q})$. Prioritized enforcement requires no such knowledge.

3.2 Optimal collection with heterogeneous agents

The analysis so far is limited for two reasons:

- First, we don't allow for heterogeneity among agents: they all owe the same taxes D , and have the same prior likelihood of being unable to settle q .
- Second, we focus on a specific class of mechanisms. Is it possible to improve revenue by letting agents send messages to the principal? Would incomplete information priorities be useful as in Halac et al. (2020)? Would it be beneficial to split the population of agents in multiple subgroups, with independent enforcement priorities?

To make progress on these questions, we now assume that taxes due D_i and friction rate q_i are indexed by tax-payer identity $i \in \{1, \dots, N\}$. In addition, the collection costs may vary across agents: forceful collection against agent i consumes $\lambda_i \in [0, \bar{\lambda}]$ units from the principal's total enforcement capacity αN .

We now show that as the population N gets large, prioritized enforcement achieves under rationalizable strategies the upper bound for tax revenue in Bayes Nash equilibrium under any mechanism.

Without loss of generality we consider direct, truthful, and obedient mechanisms in which: tax-payers send a message $m_i \in \{0, 1\}$ revealing whether they are capable of making payments; the government then sends price offers $P_i \in [0, D_i]$ and settlement recommendations $\hat{s}_i \in \{0, 1\}$; the government implements an enforcement action $a_i \in \{0, 1\}$, with $a_i = 1$ denoting forceful collection. Note that settlement offers P_i , recommendations \hat{s}_i , and enforcement actions a_i are correlated random variables across tax-payers. In particular, realized enforcement actions must satisfy the capacity constraint $\sum_{i=1}^N a_i \lambda_i \leq \alpha N$.

Proposition 5 (upper-bound on equilibrium revenue). *Under any mechanism, in Bayes Nash equilibrium, expected tax revenue is bounded above by*

$$\max \left\{ \sum_{i=1}^N \delta_i (1 - q_i) D_i \mid (\delta_i)_{i \in \{1, \dots, N\}} \in [0, 1]^N \text{ such that } \sum_{i=1}^N \delta_i q_i \lambda_i \leq \alpha N \right\}. \quad (2)$$

Proof. Under a truthful and obedient equilibrium, conditional on submitting a message $m_i = 1$, the expected utility of tax-payer i is bounded above by $-\mathbb{E}[P_i \widehat{s}_i | m_i = 1]$. Since a tax-payer can always choose to submit messages $m_i = 0$ and take settlement decision $s_i = 0$, it follows from incentive compatibility that for any tax-payer i ,

$$-\mathbb{E}[P_i \widehat{s}_i | m_i = 1] \geq -\mathbb{E}[a_i D_i | m_i = 0]. \quad (3)$$

Because of capacity constraints, it must be that $\sum_{i=1}^N a_i \lambda_i \leq \alpha N$. This implies that

$$\sum_{i=1}^N q_i \lambda_i \mathbb{E}[a_i | m_i = 0] \leq \alpha N. \quad (4)$$

Together (3) and (4) imply that

$$\sum_{i=1}^N q_i \lambda_i \mathbb{E} \left[\frac{P_i \widehat{s}_i}{D_i} \mid m_i = 1 \right] \leq \alpha N. \quad (5)$$

In turn total expected revenue is equal to $\sum_{i=1}^N (1 - q_i) \mathbb{E}[P_i \widehat{s}_i | m_i = 1]$. Let $\delta_i \equiv \mathbb{E} \left[\frac{P_i \widehat{s}_i}{D_i} \mid m_i = 1 \right] \in [0, 1]$. In equilibrium, expected collection is equal to $\sum_{i=1}^N \delta_i (1 - q_i) D_i$. Condition (5) implies that weights $(\delta_i)_{i \in \{1, \dots, N\}}$ satisfy

$$\sum_{i=1}^N \delta_i q_i \lambda_i \leq \alpha N.$$

This concludes the proof. \blacksquare

Problem (2) is a linear optimization problem with a single constraint. The marginal benefit of increasing agent i 's probability of settlement δ_i is $(1 - q_i) D_i$ while the marginal shadow cost is $\mu q_i \lambda_i$ where μ is the Lagrangian multiplier associated with the capacity constraint. Hence it is optimal to set $\delta_i = 1$ for all agents such that $(1 - q_i) D_i / q_i \lambda_i > \mu$ and $\delta_i = 0$ for all agents such that $(1 - q_i) D_i / q_i \lambda_i < \mu$.

We show in Appendix A that Proposition 5 extends nearly as is when taxes-due are

privately observed by agents but are uncertain to the principal. In any mechanism, the government can achieve no better collection than bound (2), where friction rates q_i depend on take-it-or-leave-it settlement price offers P_i , optimally chosen by the principal.¹²

Attaining the upper-bound using prioritized enforcement. Set settlement offers $P_i = D_i - \nu$ for ν small. We assign tax-payer i a score z_i defined by

$$z_i \equiv \frac{(1 - q_i)P_i}{q_i \lambda_i},$$

and consider the prioritized enforcement mechanism with prices P_i , and enforcement conducted in decreasing order of score z . For simplicity we also index agents in decreasing order of score.

Proposition 6 (heterogeneous agents). *For any $\epsilon > 0$, as N gets large, all rationalizable strategy profiles are such that*

- (i) *tax-payers with rank j such that $\frac{1}{N} \sum_{i \leq j} q_i \lambda_i \leq 1 \wedge \alpha - \epsilon$ settle;*
- (ii) *tax-payers with rank j such that $\frac{1}{N} \sum_{i \leq j} q_i \lambda_i \geq 1 \wedge \alpha + \epsilon$ do not settle;*
- (iii) *as N gets large, with probability 1, aggregate revenue approaches*

$$\max \left\{ \sum_{i=1}^N \delta_i (1 - q_i) P_i \mid (\delta_i)_{i \in \{1, \dots, N\}} \in [0, 1]^N \text{ such that } \sum_{i=1}^N \delta_i q_i \lambda_i \leq \alpha N \right\}. \quad (6)$$

Proof. The proof of points (i) and (ii) is essentially identical to that of Proposition 4.

The proof of point (iii) follows from points (i) and (ii), as well as the fact that the solution to (6) takes the form $\delta_i = 1$ for all $i < i^*$ and $\delta_i = 0$ for all $i > i^*$, with i^* such that

¹²Bound (2) turns out to be a tight worst-case scenario for collection where the distribution of taxes-due of player i must have mean $(1 - q_i)D_i$. Intuitively, a tax-payer who knows that they either owe 0 with probability q_i or D_i with probability $1 - q_i$ can be viewed as the perfectly informed counterpart of a tax-payer expecting to owe $(1 - q_i)D_i$. More informed agents extract more rents than less informed ones.

$$\sum_{i < i^*} q_i \lambda_i \leq \alpha N \text{ and } \sum_{i \leq i^*} q_i \lambda_i \geq \alpha N. \quad \blacksquare$$

Since $\nu > 0$ can be made arbitrarily small, this implies that prioritized enforcement mechanisms achieve the revenue bound of Proposition 5: allowing for more general mechanisms does not generate additional revenue. In addition, we achieve implementation in rationalizable strategies at no cost. This shows that under rationalizability, neither information design, nor splitting the overall population in subgroups increase expected revenue.

The major caveat to these observations is that they hold when we require that the agents' behavior be rationalizable. This is a demanding requirement in games with many players. We now turn to boundedly rational solution concepts, and emphasize the value of designing an appropriate extensive form.

4 Dynamic Settlement and Bounded Rationality

In this section, we embed the agents' decision to comply or not in time. This allows us to study the value of providing agents with information about the behavior of others. We show that such information does not increase revenue or significantly change behavior under rationalizability. In contrast, information design can play an important role under boundedly rational solution concepts.

We consider the following variant of the static game introduced in Section 3. Time $t \in [0, 1]$ is continuous. Each agent $i \in \{1, \dots, N\}$ becomes able to settle (i.e. wakes up) according to a Poisson process with intensity $-\log(q)$, so that by time $t = 1$, an expected share $1 - q$ of agents have woken up. If an agent is able to settle at date t , she is able to settle at all further dates $t' \in (t, 1]$. Settlement decisions are irreversible. For simplicity, we assume that agents' are homogeneous with respect to tax due D , and friction q .¹³ We denote by $s_{i,t} \in \{0, 1\}$ the agent's compliance status at time t . Once date $t = 1$ is reached, the principal investigates non-compliers according to common knowledge enforcement priority

¹³The analysis extends essentially as is to the case of heterogeneous agents.

\prec_{CK} .

The principal can commit to a deterministic settlement schedule $(P_t)_{t \in [0,1]}$.

Definition 1. *We say that a price schedule $(P_t)_{t \in [0,1]}$ is strictly increasing if $P_0 > 0$, $P_1 < D$, and for all $\epsilon > 0$, there exists $\eta > 0$ such that for all $t \in [0, 1]$,*

$$P_{t+\epsilon} - P_t \geq \eta.$$

The principal can also commit to arbitrary information policies over the past settlement behavior of agents. Specifically, in each period t , given a history of settlement decisions $h_t = (s_{i,t'})_{i \in \{1, \dots, N\}, t' < t}$ each agent i obtains a signal $x_{i,t}$ measurable with respect to h_t . This may include revealing the entire set of agents who have settled, revealing the highest rank of agents that have settled, or any other statistic of history h_t .

4.1 Design under common knowledge of rationality

We begin by showing that information design is essentially irrelevant under rationalizability. We know from Propositions 5 and 6 that it cannot increase revenue. In fact, when settlement prices are strictly increasing, it essentially does not affect settlement times.

Proposition 7 (irrelevance of design). *Take as given a strictly discounted price schedule, and an information policy. For any $\epsilon > 0$, as N becomes large, under any rationalizable strategy profile,*

- (i) *with probability approaching 1, an agent with rank $\rho < \frac{\alpha}{q} - \epsilon$ settles within a delay ϵ of being able to settle;*
- (ii) *with probability approaching 1, an agent with rank $\rho > \frac{\alpha}{q} + \epsilon$ does not settle.*

The proof, closely related to that of Proposition 4, is contained in Appendix B.

4.2 Design under bounded rationality

We now clarify the value of information design under bounded rationality. We choose to work with obvious dominance (Li, 2017), since it's a particularly weak solution concept, that has proven a useful tool to identify the fine differentiating properties of seemingly equivalent mechanisms (Akbarpour and Li, 2020, Pycia and Troyan, 2021).

Let us denote by $\sigma_i : h_i \mapsto s_i \in \{0, 1\}$ a feasible strategy of player i .¹⁴ We denote by σ_{-i} strategy profiles by players other than i , and by ω the underlying moves of nature. We denote by $u_i(\sigma_i, \sigma_{-i}, \omega | h_i)$ the realized payoff of agent i given history h_i , their own behavior σ_i , the behavior of others σ_{-i} , and realized moves of nature ω (here corresponding to agents' random wake up time).

Definition 2. *We say that a strategy σ_i obviously dominates a strategy σ'_i if and only if, for every history h_i potentially on the equilibrium path, at which strategies σ_i and σ'_i first differ,*

$$\sup_{\sigma_{-i}, \omega} u_i(\sigma'_i, \sigma_{-i}, \omega | h_i) \leq \inf_{\sigma_{-i}, \omega} u_i(\sigma_i, \sigma_{-i}, \omega | h_i).$$

A strategy σ_i is not obviously dominated if there is no strategy that obviously dominates it.

We say that agents receive no information if signals $x_{i,t}$ are constant and uninformative. We say that settlement prices are constant if for all t , $P_t = P_0$. We say that agents receive updated rank information if they receive signals

$$x_{i,t} = \frac{1}{N} \left(1 + \sum_{j < i} 1 - s_{j,t} \right),$$

i.e. agents learn their updated rank, once the settlement behavior of agents ranked higher is taken into account. The scaling factor $1/N$ simplifies notation but is otherwise irrelevant.

¹⁴Feasible strategies must be such that: if history h'_i follows h_i , then $s_i(h_i) = 1 \Rightarrow s_i(h'_i) = 1$; for any history h_i such player i has not woken up, $s_i(h_i) = 0$.

Proposition 8. (i) *Assume that players receive updated rank information and settlement prices are strictly increasing. For any $\epsilon > 0$, and all non-obviously-dominated strategy profiles $(\sigma_i)_{i \in I}$, with probability 1 as N goes to infinity, all agents with rank $\rho \leq 1 \wedge \frac{\alpha}{q} - \epsilon$ settle their taxes if they can.*

(ii) *If players receive no information, or if settlement prices are constant, then there exist non-obviously-dominated strategy profiles such that no agent with rank $\rho > \alpha$ settles their taxes.*

Using non-obviously dominated strategies as a solution concept allows us to clarify the value of information design in ways that rationalizability is silent on. In addition, it identifies a qualitative complementarity between providing updated rank information and providing incentives to settle early.

A possible concern is that implementation in non-obviously dominated strategies presumes that a small difference in payoffs is enough to ensure that players never play an obviously dominated strategy. In practice however, it is plausible that players may take time to react even once they receive information clarifying that it is obviously dominant for them to settle. Such delay would slowdown the flow of information, and significantly reduce the effectiveness of the mechanism studied in Proposition 8.

Ultimately, what is the right solution concept is an empirical question. To make progress on the issue we turn to experimental evidence from the lab.

5 Experimental Evidence

Our analysis so far has highlighted how the effectiveness of various mechanisms may vary as a function of the solution concept. The goal of our experimental investigation is to gauge the impact of various aspects of design under realistic play, and thereby assess the ability of different solution concepts to inform design in the field. Specifically, we seek to answer the following questions:

- (i) What is the relative performance of random versus prioritized enforcement? What are equity implications?

If the high settlement equilibrium is selected, then random enforcement may achieve the same revenue as prioritized enforcement, while also improving equity. If the high settlement equilibrium is not always selected then prioritized enforcement may deliver higher revenue, and implications about equity are ambiguous.

- (ii) Does updated rank information increase settlement?

If it does, this strengthens the case for non-obviously dominated strategies as a solution concept.

- (iii) Do agents use obviously dominated strategies? How important is delay in reaction time?

If there is delay in settlement time, then analysis under non-obviously dominated play overestimates the effectiveness of updated rank information.

- (iv) To what extent do people exhibit higher levels of rationality than that assumed by non-obviously dominated play?

Carefully exploiting higher levels of rationality could reduce the impact of delay in reaction time.

5.1 Design

Baseline game. Our main experiment was run on Amazon Mechanical Turk (MTurk) from August to October of 2021. Because of the difficulty of simultaneously recruiting sufficiently many reliable players (Kapon, 2022), and to allow multiple treatments to be run at the same time, we set the number of agents N to 10. To ensure that our previous analysis applies even though N is not large, as well as to make inferences about strategic thinking as crisp as possible, we set friction rate q to 0. Another set of experiments run with friction rate $q = 20\%$ yielded quantitatively similar results (see Appendix D for details).

We implemented the dynamic settlement game of Section 4, with the experimenter playing the role of the principal, and recruited participants playing the role of agents. All agents received an initial endowment of 100 points and owed the same amount $D = 100$. In our three main treatment arms, the initial settlement price was set to $P_0 = 89$, and increased

linearly over time up to $P_1 = 91$. In a fourth treatment arm, the initial settlement price was set to $P_0 = 80$ and increased to $P_1 = 91$. Time $t = 1$ corresponded to 45 seconds.

The principal's enforcement capacity was set to $\alpha = 10\%$, so that the principal can physically collect taxes from a single agent. To reduce sampling variation, the players were able to settle at some time randomly drawn without replacement from the set of 10 equidistant points between 5 seconds and 36 seconds.¹⁵

Treatments. We implemented three main treatments corresponding to different enforcement policies and different information structures. Under these three treatments, the initial settlement price was set to $P_0 = 89$, with a final settlement price at $P_1 = 91$.

In the **random** enforcement treatment, participants were not informed of the order in which enforcement would occur, and did not receive information about the settlement behavior of others. Players were simply made aware of when it was possible for them to settle, and at what price.

The other two main treatments implemented a prioritized enforcement rule, in which participants were informed of their enforcement priority, but received different additional information over time:

- In the **priority+no-info** treatment, players were given no information about the realized settlement of others.
- In the **priority+info** treatment, players were informed of their *real time effective rank*, i.e. their updated rank after taking into account settlement by other players.

Finally, a fourth **priority+info+stakes** treatment replicated the **priority+info** treatment but increased the incentives for fast settlement by setting initial settlement price to $P_0 = 80$ and final settlement price to $P_1 = 91$.

Protocol. The experiment design was filed with the AEA RCT registry under ID number AEARCTR-0004802. The experiment was programmed in oTree (Chen et al., 2016) and

¹⁵The buffer at the beginning was to ensure that any minor latency issues in the software would not impede play, while the buffer at the end ensured that a player had time to respond to being able to settle.

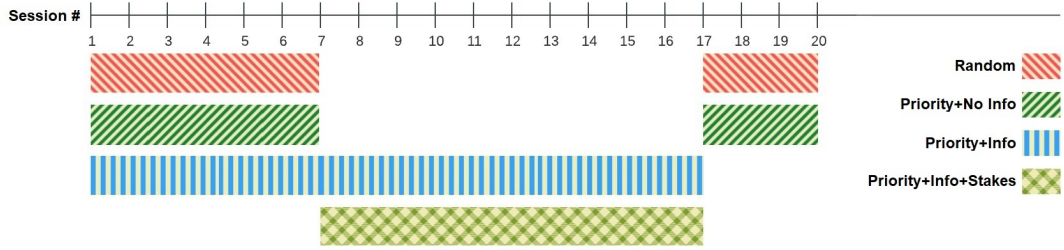


Figure 1: treatment overlap across sessions

experimental instructions were conveyed to players through their browser. Screenshots of instructions are reproduced in Appendix C.

Because of the difficulty of recruiting many MTurk users to play simultaneously (Kapon, 2022), we did not implement all four treatments jointly at all times. Instead we implemented overlapping joint sessions along the lines described by Figure 1. When we compare different treatment outcomes, we focus on the subset of overlapping sessions for the relevant treatments.¹⁶ Participants played the collection game 5 times. The first collection game did not count towards participants’ final payoff. Points earned in the last four collection games were averaged across games, and converted to cash at the rate of USD 8 for 100 points. Players were *not* reallocated across different treatments over time.

Participants earned a USD 3.5 fee for showing up at a pre-announced time. The experiment began once the required number of participants arrived. Participants earned between USD 0 and USD 8 from their play in the collection game, with mean total earnings at approximately USD 6. Participants played for an average of 25 minutes. Participants were selected from a pool of US adults over 18 years old, with an MTurk approval rate over 98% and who had completed at least 10 tasks on MTurk.

¹⁶Specifically, we ran 7 sessions, each with 30 participants randomly assigned to one of three treatments: random, priority+no-info, priority+info. To understand the role of steeper incentives to settle early, we ran 10 sessions with 20 participants randomly assigned to either priority+info or priority+info+stakes. Finally, we ran 3 sessions with 20 participants randomly assigned to random or priority+no-info. Altogether, we ran 10 sessions of each treatment, except for priority+info, of which we ran 17.

5.2 Findings

5.2.1 Is prioritized enforcement effective and when?

Mean settlement by treatment. Table 1 displays results from regressing settlement rates and tax revenue on treatment status for the 7 overlapping sessions of treatments **random**, **priority+no-info**, and **priority+info**. Treatment **random** is the omitted category.

Table 1: Settlement rates and revenues across treatments.

	settlement rate	tax revenue (per person)
constant	0.443	39.86
priority+no-info	0.068 (0.271)	6.109 (0.359)
priority+info	0.318 (0.000)	28.72 (0.000)
Observations	840	840

Two-sided p-values in parentheses. Standard-errors are clustered at the (treatment, session) level.

Three observations are immediate. First, players do not play the high settlement equilibrium under random enforcement: roughly 44% of players settle, compared to a 100% theoretical bound under the high settlement equilibrium.

Second, while the **priority+no-info** treatment increases settlement rates and revenues, it fails to implement full settlement by a large margin. It improves settlement rates by 6.8pp (or 15.3%).¹⁷

Third, the **priority+info** treatment does a much better job of reducing the distance to full settlement. It increases settlement rates by 31.8pp (or 71.8%). Effects on revenues are similar.

Altogether, these findings show that in our context, non-obviously dominated play appears to be a much better suited solution concept than either selecting the high settlement equilibrium, or rationalizability.

¹⁷The effect is significant at the 10% level if we use the 10 overlapping sessions of the **random** and **priority+info**, with a magnitude of 7pp.

Distributional effects. The distribution of group-level settlement rates is also instructive. Figure 2 plots the c.d.f. of group-level settlement rates, computed at the (session, treatment, round) level, by treatment.

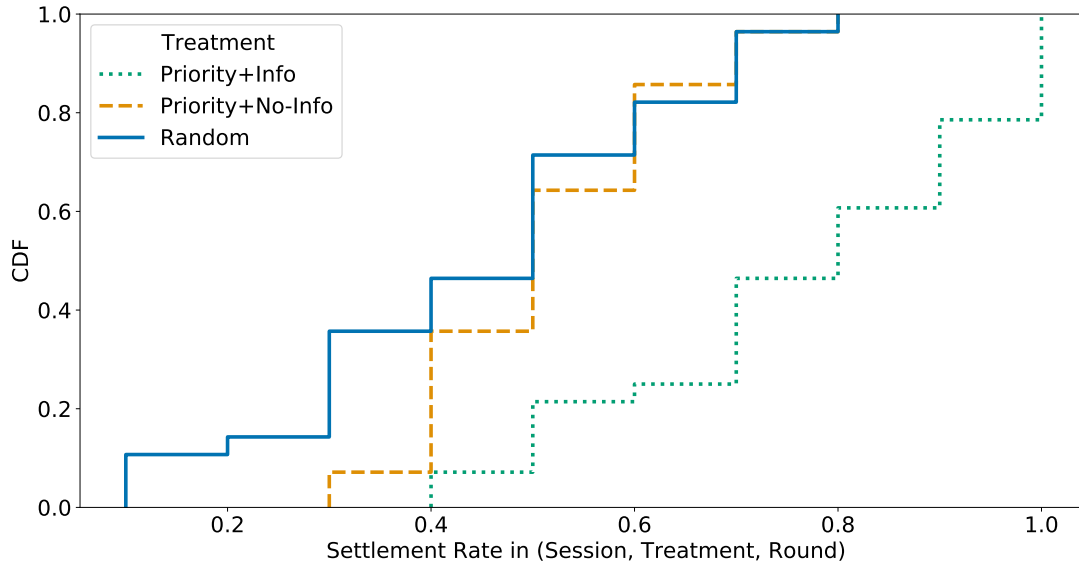


Figure 2: Cumulative distribution function of settlement rate by treatment.

Two facts are noteworthy. First, the **priority+info** treatment induces a first-order stochastic dominance (FOSD) increase in settlement rates. In addition, although the mean impact of **priority+no-info** over **random** is small, **priority+no-info** does seem to effectively reduce the left tail of outcomes. In data from the 10 overlapping sessions between the two treatments, it raises the 20th percentile of settlement rates from 30% to 40% (p-value 0.057). This can be viewed as an improvement in the equity of taxation across groups. Intuitively this finding makes sense since settling is dominant for at least one player under **priority+no-info**, while a settlement rate of 0 is an equilibrium under random enforcement.

Information does not delay settlement. A final outcome of interest to stakeholders is the timing of tax payments. In practice, speeding up the payment of taxes reduces governments' need to use short-term debt to finance their operations. In addition, earlier

information about likely tax revenues lets governments plan expenditures better.

In principle, giving players more information may increase settlement delay: the anticipation of receiving information creates an option value for waiting since it may reveal that other players are in fact not settling. For this reason, although Proposition 7 suggests that providing information should not delay settlement under rationalizability, it is not in fact obvious that it is the case in practice. It is however borne out in the data.

We define settlement delay as the difference between the time at which a player wakes up and the time at which they settle. Figure 3 illustrates the distribution of settlement delay. The settlement delay of players who choose not to settle is set to 45 seconds. Figure 3 shows that **priority+info** induced a FOSD decrease in delay over both other treatments.

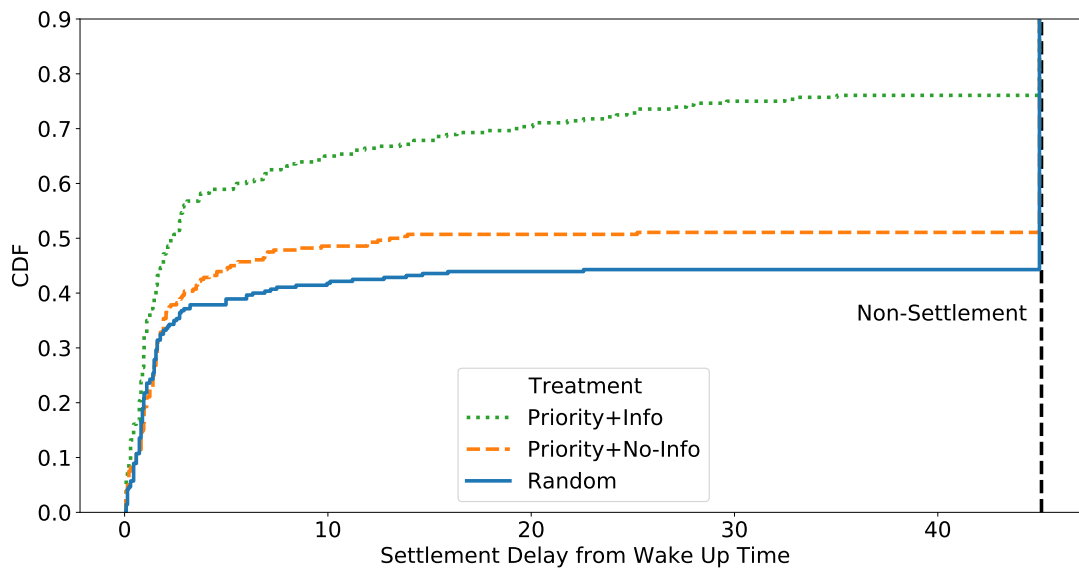


Figure 3: CDF of settlement delay

5.2.2 The Limits of Non-Obviously Dominated Play

The evidence so far suggests that non-obviously dominated play is a useful solution concept for design in our context. However, there remains some gap between theoretical and experimental performance. We are concerned by the assumption that players immediately take

the obviously dominant action, once they receive the information that they are next in line for enforcement. This seems likely to be false in practice, and important for Proposition 8 to hold.

A thought experiment. To illustrate the issue, consider a version of our game with $N = 10$, $\alpha = 10\%$, under the **priority+info** treatment. For simplicity, assume that all players wake up at time $t = 0$, and set $t = 1$ at 50 seconds. If we assume, as under non-obviously dominated play, that players take the obviously dominant action whenever possible, then players will iteratively settle very fast, resulting in a settlement rate of 100%.

Now imagine that it takes 10 seconds for a player to get their act together and settle once they realize it is obviously dominant for them to do so, and that they do not start getting their act together until settlement becomes obviously dominant. Then within the 50 seconds allocated to collection, only 5 players will end up settling.

Note that if delay is an issue, a slightly higher level of rationality could make a difference on ultimate settlement rates. For instance, if players start getting their act together once settling is the only action that survives two rounds of best response, then in the example above, 10 players will have time to settle within 50 seconds, in five pairs.

The possibility of delay in optimization suggests new design steps:

- (i) It may be valuable to increase incentives for early settlement, in particular by lowering the initial settlement price P_0 . This would be suboptimal under non-obviously dominated play, but could increase revenue if delay in settlement is a constraint.
- (ii) Engaging players' higher level rationality maybe useful, even though we remain well short of rationalizability.

Our experiments provides some preliminary facts informing these design questions.

The impact of higher stakes on settlement behavior. The `priority+info` and `priority+info+stakes` treatments were specifically designed to investigate the value of raising stakes for fast settlement. By decreasing the initial settlement price from $P_0 = 89$ to $P_0 = 80$, incentives for fast settlement were effectively multiplied by 5. The maximum benefit from settlement was multiplied by two.

Bottom-line outcomes suggest little impact of incentives for fast settlement on settlement rates and revenue. Over the 10 relevant overlapping sessions, `priority+info` exhibited a settlement rate of 70.8%, and a mean per capita revenue of 63.8. The `priority+info+stakes` treatment induced a settlement rate of 73.5% and a mean per capita revenue of 63.0. The two-sided p-values for the differences are respectively 0.39 and 0.79.

This aggregate result masks differences in the timing of settlement across treatments. Figure 4 plots the c.d.f. of settlement delays (i.e. the difference between wake up time and settlement time) across treatments. Non-settlers are assigned a delay of 45 seconds. Conditional on eventually settling, mean settlement delays were respectively 3.33 and 5.20 for high vs. low incentives to settle fast (p-value 0.001).¹⁸

The reason delay in settlement does not end up affecting ultimate collection rates is that the total time available for settlement (45 seconds) was sufficiently large not to be a binding constraint on settlement behavior. If we had set $t = 1$ to 10 seconds, it is likely we would have observed differences in the ultimate settlement rate.

This partly speaks to the difficulty of extrapolating from the lab to the field: in the lab, players have little demands on their time other than to play the game selected for them. Our experience implementing in the field suggests that agents have many demands on their time preventing them from taking swift action. In addition, assembling liquid assets needed for payment may also cause delays. This means that taking steps to reduce reaction time may be quite important in the field. Offering convenient payment options, as well as financing

¹⁸Similarly for players that reach a history at which they are awake and have effective rank 1, the mean delay in settlement (among ultimate settlers) from the first time they reached such a history was respectively 1.55 seconds and 2.27 seconds with and without incentives to settle fast (p-value 0.077).

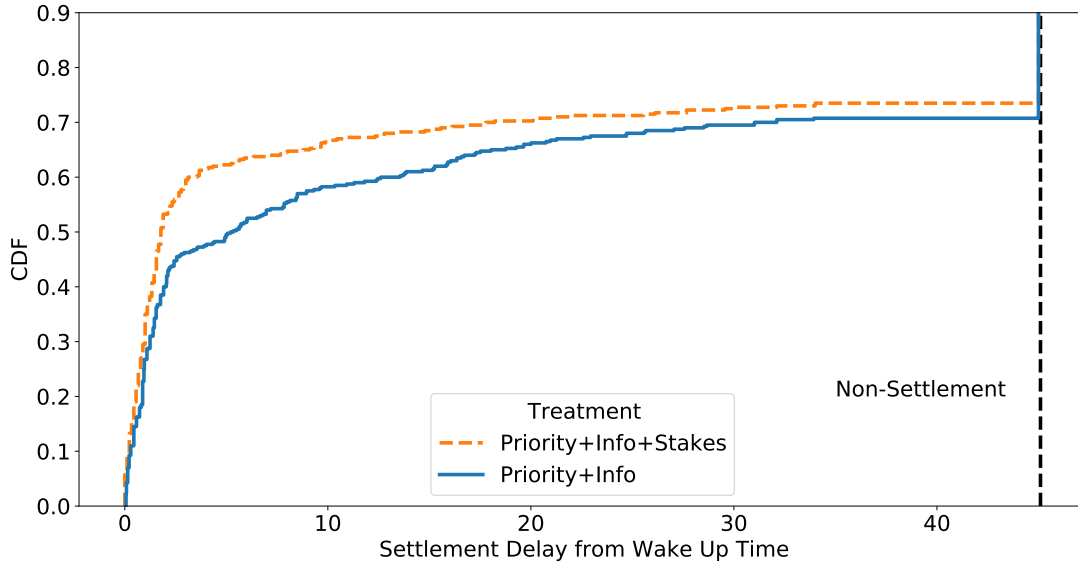


Figure 4: Settlement delay by treatment.

options seems particularly valuable.

Players exhibit higher order rationality, and the game-form matters. Non-obviously dominated play makes weak assumptions on the players' rationality. It does not assume more than one round of elimination of dominated strategies. Experimental play shows that this is a conservative assumption. Figure 5 illustrates the distribution of effective rank at settlement among settlers for the `priority+no-info` and `priority+info` treatments. The mean effective rank at settlement among settlers were respectively 3.21 and 3.04. Respectively 53.8% and 47.0% of players settled with an effective rank at settlement greater than 2.

Interestingly, the degree of higher order rationality exhibited by players appears to depend on the game form. Table 2 reports OLS estimates of the impact of treatment on the difference between effective rank at settlement and effective rank at wake across our three prioritized

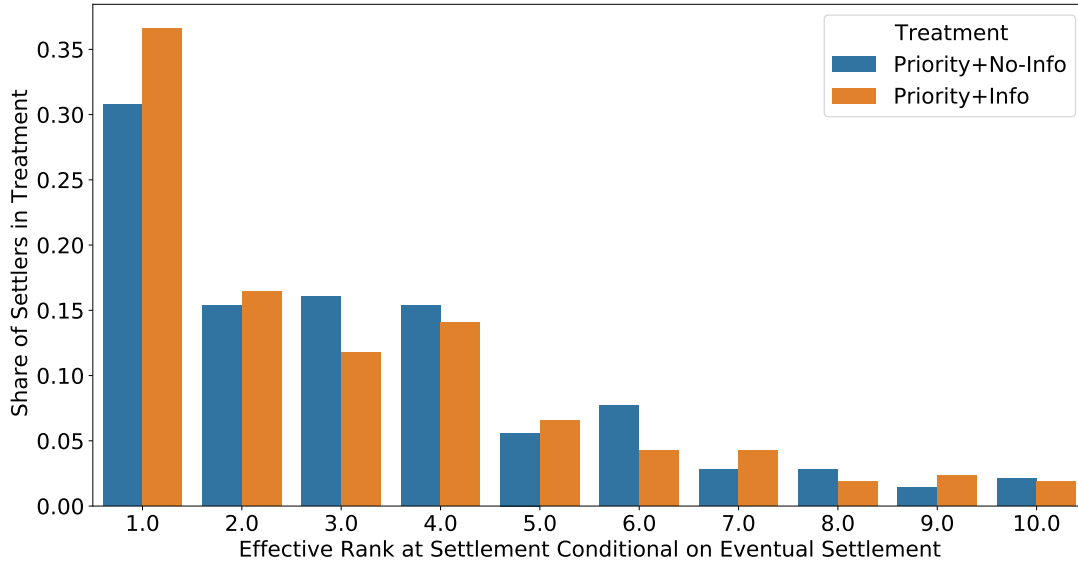


Figure 5: Effective rank at settlement.

enforcement treatments:¹⁹

$$\text{effective rank at settlement} - \text{effective rank at wake} \sim \quad (7)$$

$$1 \oplus \text{is}[\text{priority+info}] \oplus \text{is}[\text{priority+info+stakes}] \oplus \text{wake up time},$$

where $\text{is}[A]$ denotes a dummy variable indicating treatment A , and the omitted category is priority+no-info .

In words, effective rank information appears to reduce the number of iterated best-replies agents perform before making settlement choices.

¹⁹Because this regression focuses on ultimate settlers, and the three treatments induce different settlement rates, this regression is potentially affected by selection bias. For instance if individuals less likely to settle only do so once it is dominant for them. To account for this, results reported in Table 2 focus on players whose initial rank is less than 4, since such players tend to settle at similar rates across treatments (respectively 0.738, 0.871 and 0.869 for priority+no-info , priority+info and $\text{priority+info+stakes}$). The qualitative findings are robust if we consider players whose initial rank is less than 3, or the entire set of players who ultimately settle.

Table 2: Difference in effective rank between wake up and settlement times.

constant	-0.140 (0.002)
priority+info	-0.156 (0.000)
priority+info+stakes	-0.108 (0.004)
wake up time	0.006 (0.005)
Observations	494

Two-sided p-values in parentheses. Standard-errors are clustered at the (treatment, session) level.

5.3 Speculative implications for design

Concerns that non-obviously dominated play underestimates the cost of delay, and the observation that players seem to exhibit higher levels of rationality have implications for design. A systematic analysis is beyond the scope of this paper, but brief speculation helps clarify why these fine aspects of solution concepts are policy relevant.

Prioritized enforcement in small groups. Delayed reaction, and higher level thinking may provide a rationale for applying prioritized enforcement on smaller subgroups, rather than on the entire group. For instance if $\alpha = 10\%$, and $N = 100$, so that capacity $\alpha N = 10$, it may in practice be more effective to apply prioritized enforcement to 10 groups of 10 agents, each with a capacity of 1, than to the entire group with a capacity of 10.²⁰

Such a design is helpful if agents use a heuristic strategy along the following lines. At any history, agents eliminate obviously dominated actions for only one player: the player with the highest priority ahead of them. Agents then settle whenever it is obviously dominant to do so, conditional on the remaining set of strategies.

Assume that $q = 0$, and that time $t = 1$ is set to 50s. In addition, assume there is a 10s delay in reaction time. Then a single enforcement group with a capacity of 10 would lead to agents settling in sets of 11, at 10s intervals, leading to 55 tax-payers settling by $t = 1$. If instead players are split in 10 subgroups of 10, in each subgroup, agents settle in pairs at

²⁰If capacity is unused in one group, it can be recycled to an other group that needs it.

10s intervals, leading to 100 tax-payers settling by $t = 1$.

p-dominant threats. Another way to exploit plausible refinements of non-obviously dominated play is to assume that players settle whenever it is p -dominant to do so, for $p \in (0, 1)$ (Morris et al., 1995). When $q = 0$, so that there are no moves of nature, settling is obviously dominant if it is 0-dominant.

Consider an environment with $q = 0$, $N = 10$, $\alpha = 10\%$, and $t = 1$ set at 50s. We generalize prioritized enforcement mechanisms to the following class: in every period t , the mechanism designer issues direct threats to a subgroup of k players; players in the threatened group then simultaneously choose whether to settle or not. The mechanism designer enforces randomly among threatened agents who do not settle. Assume that it takes 10s for the mechanism designer to issue a threat, and for agents to react. Finally, assume that threatened tax-payers settle whenever it is $\frac{1}{2}$ -dominant for the to do so within the stage coordination game generated by the threat.²¹

Prioritized enforcement with updated information about rank corresponds to setting $k = 1$ in the mechanism above: a single player is threatened at a time. Under such a mechanism by $t = 1$, 5 players will have time to settle. Since we haven't assumed that settlement amount P impacts delay, setting $P = D$ raises the maximum tax-revenue, $5D$.

Now consider the mechanism in which settlement amount P is set to $\frac{3}{4}D$, and $k = 2$. P is chosen so that settling is $\frac{1}{2}$ -dominant in the two-player coordination game generated by a threat. As a result players settle in pairs at 10s intervals, so that all 10 end up settling. This yields tax-revenue $7.5D$.

Whether these subtle design considerations matter in practice, and how to model them in a satisfactory systematic way suitable for design is left for future investigation.

²¹In this coordination game, settling is $\frac{1}{2}$ -dominant if it is a best-response to settle whenever other players all settle with probability at least $1/2$.

6 Discussion

6.1 Limits of our analysis

Collusion. Because divide-and-conquer schemes exploit externalities between agents, they are fragile to collusive arrangements between agents. For instance, the entire group may ask the first αN agents to refuse to settle, and compensate them for the cost of enforcement. It may be possible to take design steps making such collusion more difficult. For instance, although players need to know their own rank, rank information need not be made public. In the context of our field application, the large number of households involved (more than 13,000), and the fact that they are anonymous, makes collusion less concerning.

Fairness. In principle, the use of prioritized enforcement may raise fairness concerns. In particular, if agents are heterogeneous, the priority score $(1 - q_i)D_i/q_i\lambda_i$ may result in regressive tax enforcement. Agents who owe a relatively small amount D_i , but consume a low amount of enforcement capacity $q_i\lambda_i$ in expectation may be given a higher priority than agent who owe a large amount D_i but also consume a large amount of resources $q_i\lambda_i$ in expectation. This turns out not to be a problem in our field application since the estimated propensity to settle $1 - q_i$ is increasing in agents' tax due amount D_i . Tax-payers who owe more taxes therefore get higher priority. More generally, the scoring rule may have to be constrained to reflect fairness objectives, at some loss in expected revenue.

Commitment. Divide-and-conquer mechanisms rely on commitment power: the principal must take enforcement actions in the way announced to agents, even if it is not ex post efficient to do so. In our application of interest, forceful collection against median-income taxpayers is likely ex post suboptimal: it induces both direct economic costs, and indirect political costs that are more important to the principal than the amounts actually collected. In practice, it may be necessary for the principal to signal its commitment power by diligently following through on collection threats, especially early in the implementation of the

mechanism.

Learning. Our analysis emphasizes the value of boundedly rational solution concepts to inform design in the field. It is possible that rationalizability or Bayes Nash equilibrium perform better if agents have the opportunity to learn how to play the game. However, in the context of public policy implementation, it is essential that policies perform well on the first try. Policies that deliver poor short-term results are unlikely to last. For this reason, we think that it is important from a public policy perspective to use solution concepts that successfully predict play by inexperienced agents.

6.2 Other applications

Our goal has been to better understand the practical challenges of implementing divide-and-conquer mechanisms in the field. While our model is designed to speak to a specific application, we believe the insights gained can be exploited to optimize the use of limited enforcement capacity in many other settings. We provide two examples.

Law enforcement. During crime waves, police forces may not have the resources needed to successfully prosecute a significant share of crimes. The resulting feeling of impunity may in turn increase the crime rate. A remarkable experiment, Operation Ceasefire, sought to break the circular logic of this high crime equilibrium by using divide-and-conquer. Operation Ceasefire was initiated in Boston in the mid-90s during a wave of gang related homicides and extended to a large number of cities in the US (Braga et al., 2001, Kennedy, 2011, 2012). As part of the policy, gangs were made aware of the fact that the police were capable of associating homicides to the responsible gang with great precision, even though bringing together actionable evidence valid in a court of justice was much more difficult. Second, the police departments and the relevant District Attorney made a commitment to allocate a significant share of their resources towards making life difficult for the *first few gangs* suspected of committing a homicide. This is effectively an enforcement priority based on the

time at which a murder is committed. Chassang et al. (2022) argue that Operation Ceasefire had significant option value: change, when successful, had durable positive effects.

Organizational change. We believe that prioritized enforcement may help fight widespread corruption or misbehavior (including discrimination, verbal abuse, and sexual harassment) in organizations. If misbehavior is widespread, it may be effectively impossible to fire all misbehaving agents without seriously crippling the organization: in corrupt settings, a large share of government employees may be accepting bribes; in firms, a large number of managers may be guilty of misconduct. Prioritized enforcement provides a way to initiate organizational change without firing all misbehavers. Say that the organization can tolerate firing only 2% of its workforce. That essentially corresponds to setting $\alpha = 2\%$ in our analysis. Then, agents can be investigated according to a suitable priority rule. The first 2% found to be in violation of the organization’s rules are fired. The value of this mechanism is that it makes only minimal, realistic threats, yet can potentially ensure full compliance.

Appendix

A Extension to income tax

This appendix briefly outlines how to extend the model of Section 3 to an income tax setting in which tax payers have private information about the amount of taxes $D_i \leq \bar{D}$ they would owe following a formal audit. Based on observables, the principal has a prior density f_i (with c.d.f. Q_i) over the actual tax due D_i for tax-payer i . The tax-payer knows D_i . Draws of D_i are independent across tax-payers. For simplicity, we assume that

$$\frac{1 - Q_i(D_i)}{f_i(D_i)}$$

is decreasing in $D_i \in [0, \bar{D}]$.

In this context the collection action taken $a_i \in \{0, 1\}$ is really an audit decision. As in Section 3 the capacity constraint is that the total audit costs $\sum_{i=1}^N \lambda_i a_i$ must be less than αN .

As in Section 3 the government can commit to any direct mechanism in which:

- each tax-payer i reports an amount of tax-due $m_i \in [0, \bar{D}]$;
- the government recommends a payment \hat{P}_i to each tax-payer i ;
- each tax-payer i chooses an actual payment P_i ;
- the government implements a feasible audit profile as function of messages, recommendations and actual payments.

For simplicity, we assume as in the rest of the paper that the principal maximizes collection from tax-payers who don't require an audit:

$$\sum_{i=1}^N (1 - a_i) P_i.$$

Proposition A.1 (upper-bound on equilibrium revenue). *Under any mechanism, in Bayes Nash equilibrium, expected tax revenue is bounded above by*

$$\max \left\{ \sum_{i=1}^N \delta_i (1 - Q_i(P_i)) P_i \left| \begin{array}{l} (P_i, \delta_i)_{i \in \{1, \dots, N\}} \in ([0, \bar{D}] \times [0, 1])^N \\ \text{such that } \sum_{i=1}^N \delta_i Q_i(P_i) \lambda_i \leq \alpha N \end{array} \right. \right\}. \quad (8)$$

Bound (8) corresponds to bound (2) with a friction rate q_i set to $Q_i(P_i)$ for optimally chosen settlement prices P_i : intuitively, tax-payers get a take-it-or-leave-it price offer P_i and endogenously refuse to pay whenever $D_i \leq P_i$.

Importantly, conditional on an optimal choice of prices $(P_i)_{i \in \{1, \dots, N\}}$, an analogue of Proposition 6 also holds: bound (8) is asymptotically attained by setting optimal settlement

prices P_i , and implementing a prioritized enforcement scheme using score

$$z_i \equiv \frac{(1 - Q_i(P_i))P_i}{\lambda_i Q_i(P_i)}.$$

Note that while optimizing over $(\delta_i)_{i \in \{1, \dots, N\}}$ in (8) is immediate, optimizing over $(P_i)_{i \in \{1, \dots, N\}}$ may be computationally demanding.

Proof. Consider a Bayes Nash equilibrium of a direct mechanism. A feasible auditing policy must satisfy the following constraint in expectation:

$$\mathbb{E} \left(\sum_{i=1}^N a_i \right) \leq \alpha N.$$

Consider a given tax-payer i with equilibrium audit probability $\mathbb{E}(a_i) = \bar{\alpha}_i$. Because the audit constraint in expectation is a relaxation of the ex post feasibility constraint, expected collection from i is lower than the highest expected collection from i under any individual collection mechanism such that $\mathbb{E}(a_i) \leq \bar{\alpha}_i$.

Let us denote by $\bar{a}_i(D_i)$ the audit probability of a tax-payer that discloses tax-due D_i , and asked to make a payment $P_i(D_i)$. The expected payoff of a tax-payer with true tax-due D_i , reporting tax-due D'_i is

$$-P_i(D'_i) - \bar{a}_i(D'_i)(D_i - P_i(D'_i)).$$

Observing that the payoff of a tax-payer with tax-due 0 is 0, incentive compatibility and the usual application of the envelope theorem yields the payoff formula

$$P_i(D_i)(1 - \bar{a}_i(D_i)) = \int_0^{D_i} \bar{a}_i(D) dD - \bar{a}_i(D_i)D_i.$$

This implies that the expected collection from tax payer i is bounded above by

$$\begin{aligned} \max_{\bar{a}_i} \int_0^{\bar{D}} \left[\int_0^{D_i} \bar{a}_i(D) dD - \bar{a}_i(D_i) D_i \right] f_i(D_i) dD_i \\ \bar{a}_i \text{ s.t. } \int_0^{\bar{D}} \bar{a}_i(D_i) f_i(D_i) dD_i \leq \bar{\alpha}_i \end{aligned} \quad (9)$$

Letting $\mu \geq 0$ denote the Lagrange multiplier on the auditing constraint, and applying Fubini's equality, this means that the audit policy \bar{a}_i solving (9) solves

$$\max_{\bar{a}_i} \int_0^{\bar{D}} \bar{a}_i(D_i) [1 - Q_i(D_i) - (D_i + \mu) f_i(D_i)] dD_i$$

Since $\frac{1-Q_i(D_i)}{f_i(D_i)}$ is decreasing in D_i it follows that an audit policy \bar{a}_i^* solving (9) will take a threshold form: there exists D_i^* such that for all $D_i > D_i^*$, $\bar{a}_i^*(D_i) = 0$, while for all $D_i < D_i^*$, $\bar{a}_i^*(D_i) = 1$. In turn, for all $D_i > D_i^*$, $P_i(D_i) = D_i^*$. In other terms the optimal individual taxation policy is a posted settlement price. If the tax-payer accepts, then no audit takes place. If the tax-payer refuses, then an audit takes place with probability 1.

This implies that collection under any mechanism is bounded above by

$$\begin{aligned} \max \left\{ \sum_{i=1}^N (1 - Q_i(P_i)) P_i \mid (P_i)_{i \in \{1, \dots, N\}} \text{ such that } \sum_{i=1}^N Q_i(P_i) \lambda_i \leq \alpha N \right\} \\ = \max \left\{ \sum_{i=1}^N \delta_i (1 - Q_i(P_i)) P_i \mid (P_i, \delta_i)_{i \in \{1, \dots, N\}} \in ([0, \bar{D}] \times [0, 1])^N \text{ such that } \sum_{i=1}^N \delta_i Q_i(P_i) \lambda_i \leq \alpha N \right\} \end{aligned}$$

where the point of the last equality is to highlight that as in the case of Proposition (5), given prices P_i , the optimal policy offers all tax-payers with score

$$z_i \equiv \frac{(1 - Q_i(P_i)) P_i}{\lambda_i Q_i(P_i)}$$

greater than some threshold z^* a take-it-or-leave-it settlement offer at price P_i , under the threat of audit if they do not accept, while tax-payers with scores z_i less than z^* are not

audited even if they do not settle. ■

B Proofs

Proof of Proposition 7. We first establish point (ii). Consider an agent with rank $\rho > \frac{\alpha}{q} + \epsilon$. The number of agents with rank $\rho' < \rho$ who cannot settle is greater than αN with probability approaching 1 as N becomes large. This implies that the payoff from never settling approaches 0 as N gets large, or, using Landau notation, is of order $o(1)$. Denote by **Settles** the event that the agent settles at some point. The agent's expected payoff is bounded above by $-P_0 \times \text{prob}(\text{Settles})$. By revealed preferences, we must have $o(1) \leq -P_0 \times \text{prob}(\text{Settles})$, which implies that $\text{prob}(\text{Settles}) = o(1)$. Hence, a single round of rationality is sufficient to establish point (ii).

Let us turn to point (i). We proceed by induction. Let $A(\rho) \equiv \frac{1}{N} \sum_{i=1}^{\rho N-1} (1 - s_{i,t=1})$ denote the ultimate share of agents with rank less than ρ who have not settled their taxes at time $t = 1$. For $K \in \mathbb{N}$, let $\rho_K \equiv \alpha \sum_{k=0}^K (1 - q)^k$. Our induction hypothesis at $K \in \mathbb{N}$ is that for all $\epsilon > 0$, and for all agents with rank $\rho \leq \rho_K - \epsilon$, the probability that the agent ultimately settles if possible approaches 1 as N becomes large.

Consider first agents with rank $\rho \leq \alpha$. Those agents know they will face enforcement with probability one if they do not settle. Since the price schedule is strictly discounted, their best response is to settle immediately. This establishes the induction hypothesis for $K = 0$.

We now show that the induction hypothesis at $K - 1$ implies the induction hypothesis at K . We establish in passing that it also implies vanishing delays. Pick $\epsilon > 0$ and consider an agent with rank $\rho < \alpha \sum_{k=0}^K (1 - q)^k - \epsilon$. The induction hypothesis at $K - 1$ implies that with probability approaching 1, all agents with rank ρ strictly below $\alpha \sum_{k=0}^{K-1} (1 - q)^k$ ultimately settle. This implies that the mass of enforcement actions $A(\rho_{K-1})$ converges to $q\rho_{K-1}$ as N grows large. This means that the spare enforcement capacity (scaled by $1/N$)

that can be assigned to agents with rank $\rho \geq \rho_{K-1}$ is asymptotically equal to $\alpha - q\rho_{K-1}$. Since $\rho_K = \rho_{K-1} + \alpha - q\rho_{K-1}$, it follows that with probability 1 as N becomes large, $A(\rho_K - \epsilon) \leq \alpha$ with probability approaching 1. Since enforcement is almost certain for such players it is intuitive that they should settle with very little delay with probability approaching 1. Under strict discounting any amount of delay is costly, and not settling is almost certainly a losing proposition. We now provide a formal argument.

Regardless of the agent's strategy, she is unable to settle with probability q , leading to a payoff approaching $-qD$ with probability 1. Since this component of payoffs is independent of the player's strategy, we focus on payoffs conditional on the event that the agent is ultimately able to settle.²² Let $\neg\text{Settle}$ denote the event that the agent never settles, Delay denote the event that the agent settles but with a delay greater than ϵ , and let t^* denote the first date at which the agent is able to settle.

By settling immediately, the agent is able to guarantee herself a payoff equal to $\mathbb{E}[-P_{t^*}]$. The payoff from the agent's subjectively optimal strategy is bounded above by

$$\mathbb{E}[-P_{t^*} \times (1 - \mathbf{1}_{\text{Delay}} - \mathbf{1}_{\neg\text{Settle}})] - \mathbb{E}[P_{t^*+\epsilon} \mathbf{1}_{\text{Delay}}] - \mathbb{E}[D \mathbf{1}_{\neg\text{Settle}} \mathbf{1}_{A(\rho_K - \epsilon) \leq \alpha}].$$

By optimality, this implies that,

$$\mathbb{E}[-P_{t^*}] \leq \mathbb{E}[-P_{t^*} \times (1 - \mathbf{1}_{\text{Delay}} - \mathbf{1}_{\neg\text{Settle}})] - \mathbb{E}[P_{t^*+\epsilon} \mathbf{1}_{\text{Delay}}] - \mathbb{E}[D \mathbf{1}_{\neg\text{Settle}} \mathbf{1}_{A(\rho_K - \epsilon) \leq \alpha}].$$

Observe that $\text{prob}(\neg\text{Settle and } A(\rho_K - \epsilon) \leq \alpha) \geq \text{prob}(\neg\text{Settle}) - \text{prob}(A(\rho_K - \epsilon) > \alpha)$.

Since prices are strictly discounted, this implies that there exists $\eta > 0$ such that

$$\begin{aligned} \mathbb{E}[-P_{t^*}] &\leq \mathbb{E}[-P_{t^*} \times (1 - \mathbf{1}_{\text{Delay}} - \mathbf{1}_{\neg\text{Settle}})] - \mathbb{E}[(P_{t^*} + \eta) \mathbf{1}_{\text{Delay}}] \\ &\quad - \mathbb{E}[(P_{t^*} + \eta) \mathbf{1}_{\neg\text{Settle}}] + D \times \text{prob}(A(\rho_K - \epsilon) > \alpha) \end{aligned}$$

²²In other terms, payoffs conditional on being able to settle at some point are an affine transformation of unconditional payoffs.

$$\Rightarrow \eta[\text{prob}(\text{Delay}) + \text{prob}(\neg\text{Settle})] \leq D \times \text{prob}(A(\rho_K - \epsilon) > \alpha).$$

As we noted above, the induction hypothesis at $K - 1$ implies that $\text{prob}(A(\rho_K - \epsilon) > \alpha)$ goes to 0 as N gets large. Since $\eta > 0$ is fixed independently of N , this proves that the induction hypothesis holds at K and that delay also vanishes as N gets large. This concludes the proof. ■

Proof of Proposition 8. We begin with point (i). Assume that players receive updated rank information and that prices are strictly increasing. Any non-obviously-dominated strategy σ_i must be such that whenever $x_{i,t} \leq \alpha$, at any time $t' > t$, agent i must have settled their taxes if they can.

For $\nu > 0$ but small, consider $t = \frac{\log(q+\nu)}{\log(q)}$. With probability approaching 1 as N gets large, a share at least $1 - q - \nu$ of agents is able to settle, so that with probability approaching 1, a mass at least $(1 - q - \nu)\alpha$ consisting of agents with rank $\rho \leq \alpha$ must have settled.

The following holds: for any $\bar{\rho} < 1 \wedge \frac{\alpha}{q+\nu}$ and any $\eta > 0$, if by time $t + \eta$, with probability approaching 1 as N gets large, a share at least $1 - q - \nu$ of agents with rank $\rho \leq \bar{\rho}$ has settled, then with probability approaching 1 as N gets large, a share at least $1 - q - \nu$ of agents with rank $\rho \leq \bar{\rho} + \alpha - (q + \nu)\bar{\rho}$ must have settled by time $t + 2\eta$ if they are able to. Indeed, with probability approaching 1, agents with rank $\rho \leq \bar{\rho}$ use an amount less than $(q + \nu)\bar{\rho}$ of the government's enforcement capacity. The remainder of the enforcement capacity can thus be allocated to agents with rank $\rho > \bar{\rho}$. Hence by time $t + \eta$, with probability approaching 1, agents with rank $\rho \leq \bar{\rho} + \alpha - (q + \nu)\bar{\rho}$ will receive information that their updated rank is below α . Hence all such agents who are able to must settle before time $t + 2\eta$.

Since η can be taken to be arbitrarily small, this implies that as N grows large, with probability approaching 1, all agents with rank $\rho \leq 1 \wedge \frac{\alpha}{q+\nu} - \epsilon$ must settle if they can. Since ν can be made arbitrarily small, this implies point (i).

We now turn to point (ii). Assume that agents receive no information. Then the strategy profile in which agents with rank $\rho > \alpha$ never settle, and agents with rank $\rho \leq \alpha$ settle

whenever they can, is not obviously dominated. If instead settlement prices are constant, then the strategy profile in which agents with rank $\rho \leq \alpha$ settle at the last instant $t = 1$ and other agents do not settle unless they know that their updated rank is less than α is not obviously dominated. ■

C Player instructions

This section reproduces instructions given to participants in different treatments.

C.1 Instructions for Priority - Info

Introduction

You are about to participate in an experiment. During this experiment you have the opportunity to earn a sum of money that will be paid to you at the end of the experiment. The amount of money you earn may be larger if

- you read the instructions carefully.
- you think carefully about the decisions you make.

In today's experiment, you will interact with other participants via your computer. Your decision as well as others' will affect your payoff, which is calculated in points. The experiment consists of a number of rounds, and at the end of the experiment we will calculate your **average payoff (in points) across rounds**. We then convert this average into US Dollars (USD) according to the following exchange rate:

100 points = USD 8

To compute your final payment, we add to this a USD 3.5 participation fee for the experiment.

Summary of the Experiment

In this experiment, you and other participants interact with an automated *collection authority*. General details are:

- there are 10 participants in this experiment, including you
- all participants read the same set of instructions
- there are 5 rounds including 1 practice round
- each round consists of 2 stages
 - stage 1: *settlement stage*
 - stage 2: *collection stage*

Stages Overview

Stage 1: Settlement Stage

You start each round with **100** points. At the beginning of the round, you will enter the settlement stage with the other participants. The collection authority (CA) offers you and all other participants an identical settlement opportunity to keep a number of points. During the settlement stage, you will have **45** seconds to accept the offer made by the CA. Further details of the settlement stage are given in the *Settlement Stage — Details* tab.

Stage 2: Collection Stage

The CA is able to *investigate 1 participant*. If the CA investigates you, then you will certainly pay 100 points. If you accepted a settlement offer in the settlement stage, the CA will certainly not investigate you. Details of the investigation and collection procedure are given in the *Collection Stage — Details* tab.

Collection Stage — Details

The collection authority (CA) will choose to investigate according to a pre-specified line. You will be assigned an initial position in line at the start of the settlement stage, with no two participants assigned the same position. The **one** participant with the **lowest** initial position in line **among those who do not accept a settlement offer** is investigated and forced to pay 100 points, leaving that participant with a payoff of 0 points in the round. The lowest position is 1 and the highest is 10. Participants who do not accept a settlement offer and are not investigated pay 0 points, leaving each of them with a payoff of 100 points in the round. If all participants accept a settlement offer, the CA does not investigate anybody.

Settlement Stage — Details

Your Decision

You start the round with 100 points. You will be offered a settlement by the collection authority to keep a number of points — this number decreases over time. The initial settlement offer is to keep **11** points. This offer decreases by **0.045** per second, and the final settlement offer is to keep **9** points. If you accept the offer in the settlement stage, the number of points you accept is your payoff in the round. If you do not accept the offer by the deadline, your payoff in the round depends on the outcome of the collection stage described in the *Collection Stage — Details* tab.

Delayed Decision Opportunity

The button to accept a settlement offer may not be immediately available. The button will become available after a random amount of time, before the end of the settlement stage. Once the acceptance button becomes available, it will stay available until the end of the settlement stage.

Information

You will receive information about your *current position in line* to be investigated, which is a value that is updated continuously throughout the settlement stage. At the start of the round, your current position in line is equal to your initial position in line. Afterwards, any time a participant with an initial position in line lower than yours accepts a settlement offer, your current position in line decreases by 1. In general, if your current position in line to be investigated is X, you will be shown the phrase, **Your current position in line to be investigated is X.**

Other Participants

All other participants are offered the same settlement. Their buttons become available after a random amount of time, before the end of the settlement stage.

Snapshots

Below we produce example snapshots of the screens you will see in each round. Text in red is commentary describing the page — please read these comments so you understand the screen.

The snapshot below shows an example of the settlement stage screen.

The screenshot shows a user interface for the settlement stage. On the left is a dark sidebar with the title "Instructions" and a list of menu items: Introduction, Stages Overview, Collection Stage — Details, Settlement Stage — Details, Snapshots, Payment, and Quick Facts. The main content area has a dark header with "Time left to collection stage: 0:40". Below this is the title "Settlement Stage (#1)". The text "Points you keep if you accept offer now : 10.77" is displayed. Below that, "Your current position in line to be investigated is 1" is shown. At the bottom, it says "Settlement cannot be accepted yet". Red arrows point from text annotations to these elements: "Round number" points to "#1"; "Time until you are sent to the collection stage" points to "0:40"; "Amount you keep if you immediately accept settlement" points to "10.77"; "This number indicates your current position in the investigation line" points to "1"; and "Indicates that you cannot yet accept the settlement offer" points to "Settlement cannot be accepted yet".

The instructions remain available to you. There is also a *Quick Facts* tab that lists some of the numbers found in the instructions.

Indicates that you cannot yet accept the settlement offer

Next we produced a snapshot of an example of the settlement stage once the "accept offer" button becomes available.

This screenshot is similar to the previous one, showing the same sidebar and header. The main content area now includes a green button labeled "Accept Offer" below the text "Your current position in line to be investigated is 1".

At the end of the round, you will see a results page. Below is an example of the results page for a case in which the participant did not accept a settlement offer and was not investigated.

Instructions

- Introduction
- Stages Overview
- Collection Stage — Details
- Settlement Stage — Details
- Snapshots
- Payment

Time left to complete this page: 0:16

Round number

Collection Stage and Results (#2)

You avoided investigation! ← Indicates that you were not investigated

Your payoff this round: **100.00 points** ← Your payoff from the current round is shown here. In this round, you did not settle and were not investigated, so you keep your initial endowment of 100 points.

Continue to next round

During the game, players were shown the following screen. Whenever a player was unable to settle, the “Accept Offer” button was deactivated.

Time left to the collection stage 0:10

Settlement Stage

Points you keep if you accept offer now: 9.49

Your current position in line to be investigated is 4

Accept Offer

C.2 Instructions for Priority - No Info Treatment

The instructions are identical to the priority - info treatment, except for the description of the collection stage (and the snapshots page).

Settlement Stage — Details

Your Decision

You start the round with 100 points. You will be offered a settlement by the collection authority to keep a number of points — this number decreases over time. The initial settlement offer is to keep **11** points. This offer decreases by **0.045** per second, and the final settlement offer is to keep **9** points. If you accept the offer in the settlement stage, the number of points you accept is your payoff in the round. If you do not accept the offer by the deadline, your payoff in the round depends on the outcome of the collection stage described in the *Collection Stage — Details* tab.

Delayed Decision Opportunity

The button to accept a settlement offer may not be immediately available. The button will become available after a random amount of time, before the end of the settlement stage. Once the acceptance button becomes available, it will stay available until the end of the settlement stage.

Other Participants

All other participants are offered the same settlement. Their buttons become available after a random amount of time, before the end of the settlement stage.

During the game, players were shown the following screen with their *initial rank*.

Time left to the collection stage 0:20

Settlement Stage

Points you keep if you accept offer now: 9.86
Your initial position in line to be investigated is **7**

Accept Offer

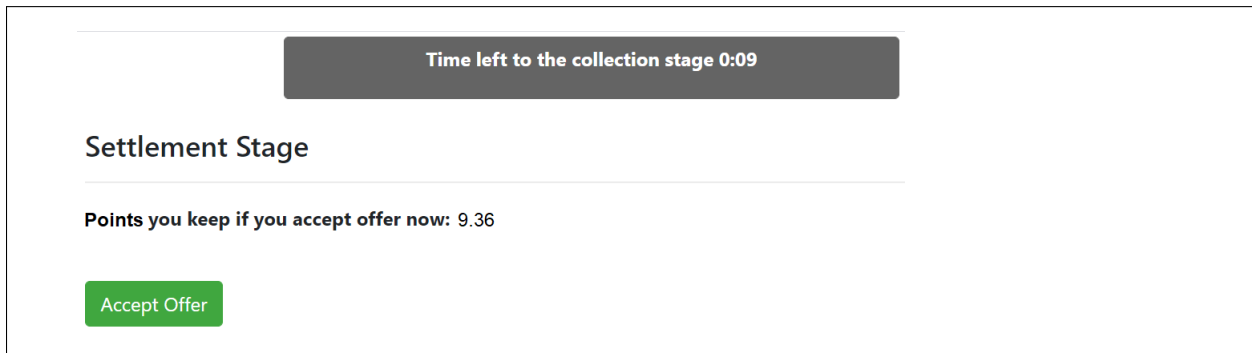
C.3 Instructions for Random Treatment

The instructions are identical to the priority-no info treatment, except for the description of collection (and the snapshots page).

Collection Stage — Details

The collection authority (CA) will RANDOMLY choose **one participant among those who do not accept a settlement offer** to investigate and force to pay 100 points, leaving that participant with a payoff of 0 points in the round. Participants who do not accept a settlement offer and are not investigated pay 0 points, leaving each of them with a payoff of 100 points in the round. If all participants accept a settlement offer, the CA does not investigate anybody.

During the game, players were shown the following screen.



Time left to the collection stage 0:09

Settlement Stage

Points you keep if you accept offer now: 9.36

Accept Offer

D Further empirics

D.1 Results for positive q

In this section, we detail results for the case in which $q = 0.2$. The parameters were the same as in Section 5, except the settlement stage lasted 30 seconds and $\alpha = 0.2$. Between March 2020 and August 2020, we ran 20 sessions, each with 40 participants, assigned to one of four treatments: **random**, **priority + no info**, **priority + aggregate information** and **priority + info**. The first two are identical to their counterparts in Section 5. The third is a version of prioritized enforcement in which participants are told the *total number* of other participants who have settled, but nothing else. The fourth is identical to its counterpart in Section 5, except subjects are also told the total number of other participants who have settled.

To reduce sampling variability, we constrained realization of times after which agents

were able to settle. This would not change our findings for large numbers of players. Specifically, we took the following steps. In each realization of the game, exactly 2 players were exogenously unable to settle: one uniformly selected player with rank less than 5, and one uniformly selected player with rank strictly greater than 5. Among players able to settle, 3/4 were able to settle (at a uniformly drawn date) within the first 15 seconds of the game, and 1/4 were able to settle (at a uniformly drawn date) within the last 15 seconds.

Subjects played 4 rounds, and the first did not count towards final payoff.

D.1.1 Treatment Effects

Similar to its counterpart in Section 5, Table D.1 displays results from regressing an indicator for settlement and settlement amount on treatment status. The table demonstrates that

- Random enforcement generates approximately 47% settlement
- Prioritized enforcement improves over random enforcement, but only modestly
- Information, and especially targeted information, leads to a large improvement over random enforcement

Table D.1: Settlement rates and revenues across treatments.

	settlement rate	tax revenue (per person)	
constant	0.467	40.11	
priority+no-info	0.071 (0.130)	6.067	(0.134)
priority+aggregate-info	0.123 (0.028)	10.78	(0.027)
priority+info	0.190 (0.000)	16.66	(0.000)
Observations	1920	1920	

Two-sided p-values in parentheses. Standard-errors are clustered at the (treatment, session) level.

Similar to its counterpart in Section 5, Figure D.1 shows the CDF of settlement rates by treatment, conditional on waking up. As in Section 5, **priority+info** first-order stochastically dominates all other treatments. For the most part, the CDF of **random** lies above the

other treatments, though it crosses **priority+no-info** at high values. The **priority+no-info** and **priority+aggregate-info** treatments are not clearly ranked.

Figure D.2 shows the CDF of settlement delay, where non-settlers are assigned a settlement time of 30 seconds. As with its $q = 0$ counterpart in Section 5, the figure shows that **priority+info** induces a FOSD decrease in delay over all other treatments.

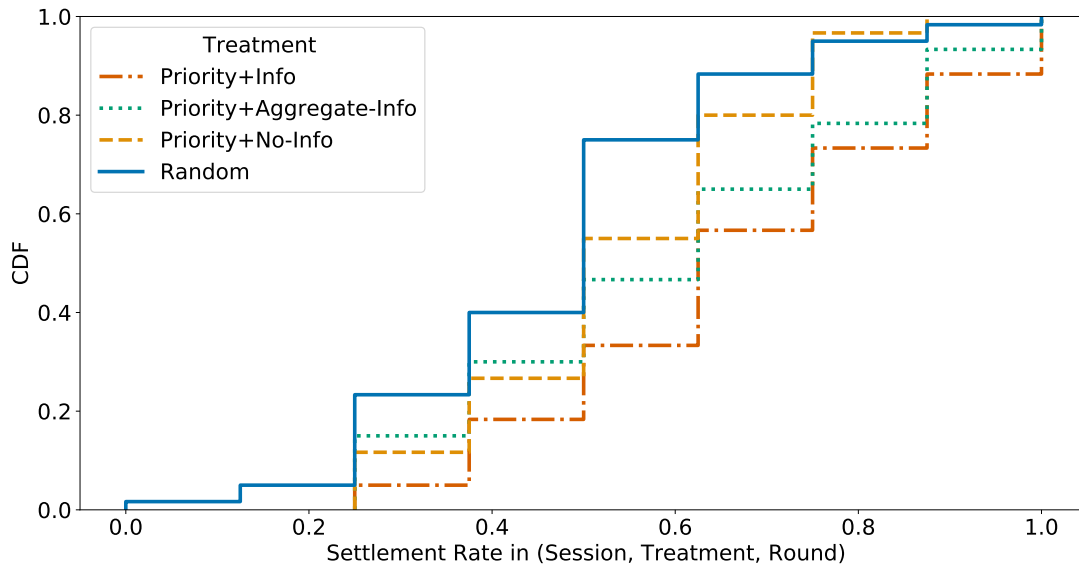


Figure D.1: Cumulative distribution function of settlement rate by treatment.

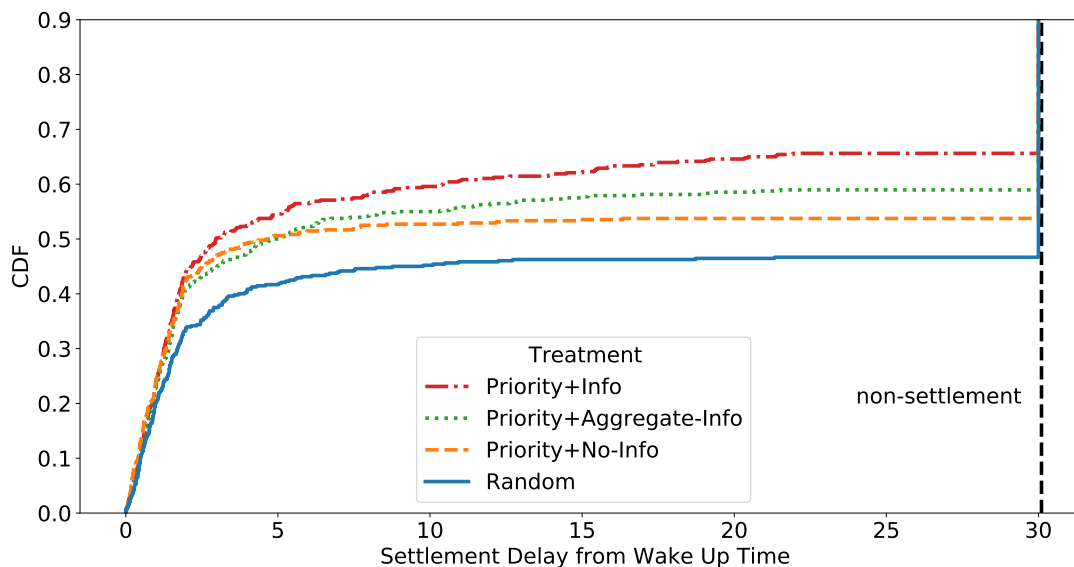


Figure D.2: CDF of settlement delay

References

- ABREU, D. AND H. MATSUSHIMA (1992): “Virtual implementation in iteratively undominated strategies: complete information,” *Econometrica: Journal of the Econometric Society*, 993–1008.
- AKBARPOUR, M. AND S. LI (2020): “Credible auctions: A trilemma,” *Econometrica*, 88, 425–467.
- ALLINGHAM, M. G. AND A. SANDMO (1972): “Income tax evasion: A,” .
- BECKER, G. S. (1968): “Crime and punishment: An economic approach,” in *The economic dimensions of crime*, Springer, 13–68.
- BRAGA, A. A., D. M. KENNEDY, E. J. WARING, AND A. M. PIEHL (2001): “Problem-oriented policing, deterrence, and youth violence: An evaluation of Boston’s Operation Ceasefire,” *Journal of research in crime and delinquency*, 38, 195–225.

- 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.
- CHASSANG, S., M. CHEN, AND M. KOLESÁR (2022): “Measuring the Option Value of Change: Theory and Application to a Homicide Reduction Program,” *Princeton University Working Paper*.
- CHEN, D. L., M. SCHONGER, AND C. WICKENS (2016): “oTree—An open-source platform for laboratory, online, and field experiments,” *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- CHEN, Y. AND J. O. LEDYARD (2010): “Mechanism design experiments,” in *Behavioural and Experimental Economics*, Springer, 191–205.
- CRAWFORD, V. P. AND N. IRIBERRI (2007): “Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions?” *Econometrica*, 75, 1721–1770.
- DAL BÓ, E. (2007): “Bribing voters,” *American Journal of Political Science*, 51, 789–803.
- DE CLIPPEL, G., R. SARAN, AND R. SERRANO (2019): “Level-k mechanism design,” *The Review of Economic Studies*, 86, 1207–1227.
- DE NEVE, J.-E., C. IMBERT, J. SPINNEWIJN, T. TSANKOVA, AND M. LUTS (2021): “How to improve tax compliance? Evidence from population-wide experiments in Belgium,” *Journal of Political Economy*, 129, 1425–1463.
- DEL CARPIO, L. (2014): “Are the neighbors cheating? Evidence from a social norm experiment on property taxes in Peru,” *Unpublished Manuscript, Princeton University*.
- DEL CARPIO, L., S. KAPON, AND S. CHASSANG (2022): “Using Divide and Conquer

- to Improve Tax Collection: Evidence from the Field,” Tech. rep., Princeton University Working Paper.
- 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, 203–32.
- ECKHOUT, J., N. PERSICO, AND P. E. TODD (2010): “A theory of optimal random crackdowns,” *American Economic Review*, 100, 1104–35.
- ELIAZ, K. AND R. SPIEGLER (2015): “X-games,” *Games and Economic Behavior*, 89, 93–100.
- GLAZER, J. AND R. W. ROSENTHAL (1992): “A note on Abreu-Matsushima mechanisms,” *Econometrica: Journal of the Econometric Society*, 1435–1438.
- GLAZER, J. AND A. RUBINSTEIN (1996): “An extensive game as a guide for solving a normal game,” *journal of economic theory*, 70, 32–42.
- HALAC, M., I. KREMER, AND E. WINTER (2019): “Raising capital from heterogeneous investors,” *American Economic Review*.
- HALAC, M., E. LIPNOWSKI, AND D. RAPPOPORT (2020): “Rank Uncertainty in Organizations,” *Available at SSRN 3553935*.
- HEALY, P. J. AND L. MATHEVET (2012): “Designing stable mechanisms for economic environments,” *Theoretical economics*, 7, 609–661.
- KAGEL, J. H., R. M. HARSTAD, AND D. LEVIN (1987): “Information impact and allocation rules in auctions with affiliated private values: A laboratory study,” *Econometrica: Journal of the Econometric Society*, 1275–1304.
- KAGEL, J. H. AND D. LEVIN (2001): “Behavior in multi-unit demand auctions: experiments with uniform price and dynamic Vickrey auctions,” *Econometrica*, 69, 413–454.

- KAPON, S. (2022): “Interactive Experiments on Amazon Mechanical Turk,” Tech. rep., Princeton University Working Paper.
- KATOK, E., M. SEFTON, AND A. YAVAS (2002): “Implementation by iterative dominance and backward induction: An experimental comparison,” *Journal of Economic Theory*, 104, 89–103.
- KENNEDY, D. M. (2011): *Don’t shoot: one man, a street fellowship, and the end of violence in inner-city America*, Bloomsbury Publishing USA.
- (2012): *Deterrence and crime prevention: Reconsidering the prospect of sanction*, Routledge.
- KLEVEN, H. J., M. B. KNUDSEN, C. T. KREINER, S. PEDERSEN, AND E. SAEZ (2011): “Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark,” *Econometrica*, 79, 651–692.
- LI, S. (2017): “Obviously strategy-proof mechanisms,” *American Economic Review*, 107, 3257–87.
- LUTTMER, E. F. AND M. SINGHAL (2014): “Tax morale,” *Journal of economic perspectives*, 28, 149–68.
- MATHEVET, L. (2010): “Supermodular mechanism design,” *Theoretical Economics*, 5, 403–443.
- MCKELVEY, R. D. AND T. R. PALFREY (1995): “Quantal response equilibria for normal form games,” *Games and economic behavior*, 10, 6–38.
- MILGROM, P. (2000): “Putting auction theory to work: The simultaneous ascending auction,” *Journal of political economy*, 108, 245–272.
- MORRIS, S., R. ROB, AND H. S. SHIN (1995): “p-Dominance and belief potential,” *Econometrica: Journal of the Econometric Society*, 145–157.

- PYCIA, M. AND P. TROYAN (2021): “A theory of simplicity in games and mechanism design,” *University of Zurich, Department of Economics, Working Paper*.
- SEFTON, M. AND A. YAVAŞ (1996): “Abreu–Matsushima mechanisms: experimental evidence,” *Games and Economic Behavior*, 16, 280–302.
- SEGAL, I. (2003): “Coordination and discrimination in contracting with externalities: Divide and conquer?” *Journal of Economic Theory*, 113, 147–181.
- SEGAL, I. R. AND M. D. WHINSTON (2000): “Naked exclusion: comment,” *American Economic Review*, 90, 296–309.
- SLEMROD, J. (2019): “Tax compliance and enforcement,” *Journal of Economic Literature*, 57, 904–54.
- SPIEGLER, R. (2000): “Extracting interaction-created surplus,” *Games and Economic Behavior*, 30, 142–162.
- WINTER, E. (2004): “Incentives and discrimination,” *American Economic Review*, 94, 764–773.