

NBER WORKING PAPER SERIES

MAKING THE MOST OF LIMITED GOVERNMENT CAPACITY:
THEORY AND EXPERIMENT

Sylvain Chassang
Lucia Del Carpio
Samuel Kapon

Working Paper 28042
<http://www.nber.org/papers/w28042>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2020

We are indebted to Vishal Ashvinkumar, Kathleen Ngangoue, Paula Onuchic, and Tianzan Pang for feedback on our experiment design. Early discussions with Dilip Abreu, Abhijit Banerjee, Daniel Diermeier, Guillaume Frechette, Navin Kartik, John Klopfer, Alessandro Lizzeri and Ariel Rubinstein helped refine our thoughts on this topic. We are grateful to seminar participants at PSE, LSE and MIT for very helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Sylvain Chassang, Lucia Del Carpio, and Samuel Kapon. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Making the Most of Limited Government Capacity: Theory and Experiment
Sylvain Chassang, Lucia Del Carpio, and Samuel Kapon
NBER Working Paper No. 28042
October 2020
JEL No. C72,C73,C92,D73,D82,D86,H26

ABSTRACT

Limits on a government's capacity to enforce laws can result in multiple equilibria. If most agents comply, limited enforcement is sufficient to dissuade isolated agents from misbehaving. If most agents do not comply, overstretched enforcement capacity has a minimal impact on behavior. We study the extent to which divide-and-conquer enforcement strategies can help select a high compliance equilibrium in the presence of realistic compliance frictions. We study the role of information about the compliance of others both in theory and in lab experiments. As the number of agents gets large, theory indicates that providing information or not is irrelevant in equilibrium. In contrast, providing individualized information has a first order impact in experimental play by increasing convergence to equilibrium. This illustrates the value of out-of-equilibrium information design.

Sylvain Chassang
Department of Economics
Princeton University
Julis Romo Rabinowitz Building
Princeton, NJ 08544
and NBER
chassang@princeton.edu

Samuel Kapon
New York University
19 West 4th Street,
New York, NY 10003
sdk301@nyu.edu

Lucia Del Carpio
INSEAD
3022 Boulevard de Constance
77300 Fontainebleau
France
Lucia.delcarpio@insead.edu

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 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. The goal of this paper is to better understand the extent to which divide-and-conquer enforcement strategies can help select a high compliance equilibrium in the presence of realistic compliance frictions.

We study the problem of a government entitled to collect an amount of taxes D from each of N agents. The government is able to forcefully collect the amount D , but only after conducting an audit of the targeted agent. The difficulty is that: (i) the government is only able to audit a share $\alpha \in (0, 1)$ of agents; (ii) upon audit, the maximum amount of damages the government can claim is D . Instead of collecting taxes through audits, the government can offer agents to settle their taxes for a given price P . Agents who accept to settle (or comply) are not audited. Audit capacity is spent on auditing non-complying agents. The government’s main policy instrument is to commit to an auditing rule, i.e. an order in which to audit non-compliant agents.

We first establish benchmark results in a static frictionless environment in which agents

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.

are able to settle with probability 1, and the amount of tax they owe is common-knowledge. We contrast uniform auditing rules, in which a number αN of agents are randomly chosen from non-complying agents, with prioritized auditing rules, in which agents are given common-knowledge priorities ahead of settling and non-complying agents are audited in order of priority. We show that uniform random auditing leads to multiple equilibria, including: high collection equilibria in which the government charges a settlement price P slightly below D , and all agents settle; and low collection equilibria in which the government charges a settlement price $P = \alpha D$, and agents all settle at that price, but refuse to settle if the government charges a higher price. The reason for this multiplicity is that incentives to settle depend on how many agents choose to settle or not. If most agents settle, then even limited auditing capacity is sufficient to enforce compliance. If instead few agents are expected to settle, auditing capacity is spread thinly across many non-complying agents so incentives to settle are very weak. In contrast, under prioritized auditing, a unique, high collection equilibrium is selected. The reason for this is that regardless of the overall behavior of agents, incentives are tightly focused on a small group of marginal agents: it is dominant for the αN agents with the highest priority to settle. Given this, 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 audits. Specifically, we assume that with exogenous probability q an agent is simply unable to settle.³ Then, agents mechanically unable to settle risk interrupting 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. In order to get full collection, audit capacity α must be larger than settlement friction q . We show that when taxes owed D , and frictions q are heterogenous among agents, it is optimal to rank agents according to

³This may be interpreted as the impact of psychological payoff shocks along the lines of quantal response equilibrium (McKelvey and Palfrey, 1995). Alternatively, if the taxes owed determined by an audit are uncertain and simply predicted, then q may be the share of agents for whom settlement price P is less than the taxes D they really owe.

score D/q .

Finally, we let settlement decisions take place over time and study the impact of providing agents information about the settlement decisions of others. This analysis is motivated by a recent empirical literature evaluating such policies in the context of taxation (Del Carpio, 2014, Castro and Scartascini, 2015, Dwenger et al., 2016), as well as recent interest in the value of information design in equilibrium (Kamenica and Gentzkow, 2011). We establish an irrelevance benchmark result: as the number of agents N gets large, for any information structure with probability 1, the share α/q of agents with the highest priority settle as soon as they are able to. This implies that any effect of information on play must be assigned to either behavioral forces, or non-equilibrium play.

We complement our theoretical analysis of the auditing strategies for three reasons. First, prioritized auditing need not improve on uniform random auditing if agents somehow coordinate on a high collection equilibrium. Second, the effectiveness of prioritized auditing relies on agents being able to perform many rounds of iterated best-response ahead of choosing their play. Prioritized auditing may be much less effective in environments where players exhibit bounded rationality. Third, in the presence of bounded rationality, information that helps players make better individual choices may improve the effectiveness of prioritized auditing.

Experimental data collected from an implementation of our dynamic settlement game both supports and qualifies the theory. Prioritized auditing does improve over uniform random auditing. However, agents do not systematically coordinate on a low settlement equilibrium under uniform auditing, and settlement under prioritized auditing falls short of our theoretical benchmark. Additionally we find that contrary to predictions of our irrelevance result, information can significantly improve settlement rates, but that this depends on the nature of the information. Information targeted to individual agents clarifying their effective rank given settlement by others is more effective than aggregate information simply reporting the overall share of agents that have settled. This suggests that information design plays an important role out-of-equilibrium by helping players approach equilibrium play.

Our work contributes to the literature on full implementation in mechanism design (Maskin, 1999, Jackson, 1992) that emphasizes implementation of desirable outcomes in all equilibria. In particular, the work of Abreu and Matsushima (1992) on virtual implementation highlights the value of more sophisticated divide-and-conquer schemes to achieve full implementation. Divide-and-conquer schemes also play an important role in the literature on contracting with externalities, including Segal (2003), Winter (2004) and more recently Halac et al. (2019, 2020). Halac et al. (2020) also studies the impact of information design, but allows for schemes in which the rank of agents is not common knowledge. In contrast to our results (which pertain to revealing information about the compliance of others), they find that appropriately designed information about others' rank can increase the principal's surplus. Our experimental findings suggest that it may be worth investigating the validity of such results in realistic non-equilibrium settings, and that the details of how the information is conveyed to agents' may be of considerable importance in practice.

The paper also contributes to the literature on out-of-equilibrium mechanism design. A large body of experimental work reviewed in Chen and Ledyard (2010) has emphasized the learning properties of mechanisms, and in particular dominance solvable mechanisms. Among others, Healy (2006) shows that in repeated public goods mechanisms, contribution behavior is well explained as a best-response to recent play. Mathevet (2010), Healy and Mathevet (2012) emphasize design steps that can be taken to ensure a mechanism is super-modular, ensuring more stable learning dynamics. Crawford and Iriberri (2007) shows that a level- k model can explain overbidding in experimental auctions. More recently De Clippel et al. (2019) studies implementation using k iterations of best reply as a solution concept. We emphasize the value of information design out-of-equilibrium.

Most importantly, the paper hopes to stimulate and guide the evaluation of divide-and-conquer mechanisms in real-life settings. Operation Ceasefire, introduced in Boston in the mid-1990s to reduce gang violence, and applied in several dozen cities, illustrates the value of divide-and-conquer strategies in practice. As described in Braga et al. (2001) and Kennedy (2011), a key aspect of Operation Ceasefire was to commit dedicated police resources to deal

with the first few gang related murders that would happen after a specific call in date during which gang members were invited together and commonly informed of these new policies.⁴ By treating the next few murders differently from pre-existing cases, Operation Ceasefire refocused a police force, overstretched by the crack epidemic, to unravel the equilibrium logic of violence: if no gang wants to initiate violence, then no violence happens. Chassang and Chen (2020) evaluate more than 20 years of evidence across dozens of cities and show that Operation Ceasefire had a significant impact on homicide rates, especially when evaluated as a real option. In addition, recent work has demonstrated the real-life effectiveness of mechanism design steps in improving governance: Duflo et al. (2013) evaluate the value of cross validating messages from multiple environmental inspectors in India, Del Carpio (2014) studies the importance of information about group compliance on own compliance, similarly Pomeranz (2015) highlights the usefulness of the value-added tax in generating actionable information about the profits of other firms permitting cross validation. We hope that by clarifying the impact of realistic frictions and how to address them, this paper serves as a blueprint for the implementation of divide-and-conquer strategies to improve governance capabilities in real-life environments.

The paper is structured as follows. Section 2 sets up a simple model in the context 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 agent non-response on the effectiveness of enforcement priorities. Section 4 casts compliance decisions in a dynamic context, and shows that schemes that provide agents with information about the compliance of others have a negligible impact on collection as the number of agents gets large. Section 5 describes our experimental hypotheses and the corresponding design. Section 6 describes our findings. Section 7 discusses applications besides tax collection. A subset of player instructions is included in Appendix A.

⁴Operation Ceasefire included other important steps, such as involving respected community members, and convincingly demonstrating the police and district attorney’s ability to deliver on their commitments.

2 Static Enforcement without Errors

2.1 Framework

We flesh out our model in the context of tax auditing, and discuss other applications in Section 7. N agents indexed by $i \in I \equiv \{1, \dots, N\}$ each owe a principal a fixed amount D . The principal can potentially collect D from each agent, but can only do so following a *due process* requiring a formal audit. The agents and the principal are all risk-neutral.

The difficulty is that the principal has limited capacity. Specifically, the principal can only audit $\alpha N \geq 1$ agents with $\alpha \in (0, 1)$. If the principal audits an agent, he collects D but can do nothing more. With experimental evaluation in mind, we are interested in a specific class of collection mechanisms. The principal can make settlement offers and commit to an audit schedule according to the following extensive-form game:

- (i) The principal gives each agent the possibility to settle at a common price P . Agents who accept are spared from audits.
- (ii) Agents simultaneously decide whether or not to settle and pay price P or not.
- (iii) The principal audits agents who have not settled according to a complete order \prec over I . An audited agent pays D .

We consider two possible enforcement priorities \prec :

- Random priorities \prec_R : audited 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 non-compliant agents are audited in order of their index $i \in \{1, \dots, N\}$.⁵

⁵Halac et al. (2020) show in a related context that incomplete information over rank can increase the principal's surplus. Since we have experimental implementation in mind, we prefer to focus on a relatively simpler class of mechanism. Even under this simpler class of mechanism, common knowledge of rationality does not appear to hold in our data.

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 directly audited agents. This simplifies computations and reflects the fact that the net benefits of audits may be ambiguous: investigation costs may be well above the amount that can be legally collected from agents. This also clarifies that the value of audits comes from incentive provision rather than actual collection.

We use both Bayes Nash Equilibrium and rationalizability as solution concepts.

Modeling assumptions. In our model, audits are free to the principal but constrained in number. This assumption implies that the principal cannot commit to audit every agent that does not settle. Our result would continue to hold 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 (investigators must be hired, trained, and paid even if there are no crimes to investigate). Alternatively, even if auditors are only paid in the event an actual audit happens, then limited capacity may capture limits on the principal's ability to commit to run expensive audits ex post.

We note that we do not allow for dissuasive punishments in the style of Becker (1968). Our results would be changed if arbitrarily high punishments were available, but not if only intermediate punishments were available. If arbitrarily high punishments are allowed, they can compensate for very limited enforcement capability. However, in practice, there are limits to legitimate levels of punishments. In the US where 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 underreported income is 20%.

In principle, the use of prioritized audits may raise fairness concerns. Indeed, similar

agents can be treated differently when they do not accept settlement. However, because settlement price P applies uniformly to all agents, agents that comply are treated identically. We believe this is a socially acceptable relaxation of fairness constraints. In contrast we think that charging similar agents a different settlement amount would not be socially acceptable because it would treat similar compliant agents differently.⁶ In other terms, unequal treatment of equals strikes us as much more acceptable off-of-the-equilibrium-path, than on-the-equilibrium-path. Ortner and Chassang (2018) make a similar point in the context of counter-corruption measures.

2.2 The Value of Common-Knowledge 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 offer $P \in [\alpha D, D]$, there exists a Nash equilibrium such that the principal makes offer P , all agents accept offers below or equal to P and reject any offer higher than P . There does not exist an equilibrium such that the principal makes an offer $P \notin [\alpha D, D]$.*

Proof. We first establish that for any $P \in [\alpha D, D]$, the profile of strategies such that the principal offers settlement price P , while agents accept any offer less than P and reject all options higher than P is a Nash equilibrium. Offering P is a best response for the principal since lower offers yield lower collections and a higher offer leads to no collection at all. Accepting an offer $P' \leq P$ is a best response for an agent if and only if $P' \leq D$, which holds since $P \leq D$. Rejecting an offer $P' > P$ is a best response provided that $P' \geq \alpha D$, using the

⁶Note that charging different agents a different price may increase the amount of revenue that could be raised. This point is emphasized in a joint production context by Winter (2004), and in a fund-raising context by Halac et al. (2019).

fact that under random order \prec_R , an agent is audited with probability α when all agents do not comply.

There can be no equilibrium such that $P > D$, since in that case, it is dominant for agents not to settle, and the principal collects an amount equal to zero. In addition, there cannot be an equilibrium such that $P < \alpha D$, since when the principal offers any $P' \in (P, \alpha D)$, it is dominant for all agents to accept: a non-compliant agent gets audited with probability at least α , regardless of the behavior of other players. ■

A corollary of Proposition 1, is that for any given settlement amount $P \in (\alpha D, D)$, random enforcement is consistent with two corner equilibria: a high compliance equilibrium in which all agents comply so that even limited enforcement capacity is enough to discipline

Proposition 2 (equilibrium selection via common-knowledge priorities). *Under common-knowledge enforcement order \prec_{CK} , a unique strategy profile survives iterated elimination of dominated strategies. The principal makes an offer P , which all agents accept.*

Proof. We show that for every settlement offer $P < D$, it is iteratively dominant for all agents to settle, so that the principal collects an amount 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 $k = 0$ since the highest priority agent is audited with probability 1 in the event they do not comply. In turn, if the hypothesis holds for $k \geq 0$, then an agent of rank $k + 1$ that does not comply is audited with probability 1. Hence, it is conditionally dominant for an agent of rank $k + 1$ to comply, which establishes the induction step.

In turn, it is dominant for all agents to refuse any offer $P > D$ regardless of how other agents behave. It follows that setting $P = D$ maximizes the principal's payoff. ■

The proof makes clear that the important aspect of \prec_{CK} is that players know their own rank, and this is common knowledge. Since players are symmetric, whether or not they know the

rank of others does not change Proposition 2.

It is also worth noting that in the absence of frictions, prioritized enforcement is extremely effective: an auditing capacity of one is enough to induce any arbitrary number of agents to settle. Sections 3 and 4 study how realistic frictions perturb the effectiveness of prioritized enforcement.

3 Static Enforcement with Errors

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 the principal's ability to essentially "recycle" enforcement capacity, and relies on a high degree of confidence that 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 audits.

Exogenous non-compliance rate q can be interpreted in different ways. In practice, the agent may simply not be aware of the collection problem, or experience liquidity shocks preventing any payment. Alternatively, the agent may know that the amount collected by the principal is erroneous, and that an investigation will prove she does not owe money. Such errors are especially likely in a setting where the principal attempts to predict owed taxes on the basis of informative but imperfect data. In such settings, it may be reasonable to assume that non-compliance rate q is an increase function of price P , or that both non-compliance rate q and taxes due D are heterogeneous across agents. We discuss these extensions below.

Going forward, 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 3. *Consider prioritized enforcement order \prec_{CK} . Set a settlement price $P \in (0, D)$, and fix $\epsilon > 0$. For N large enough, under all rationalizable strategy profiles,*

(i) agents with rank $\rho > \frac{\alpha}{q} + \epsilon$ do not comply;

(ii) agents with rank $\rho < \frac{\alpha}{q} - \epsilon$ comply if they are able to settle.

Proof. We first show that for N large enough, it is dominant for agents with rank $\rho > \frac{\alpha}{q} + \epsilon$ not to comply. 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 profile, an agent with rank $\rho > \frac{\alpha}{q} + \epsilon$ gets audited with probability 0. Hence, for N large enough, it is dominant to refuse any settlement offer $P > 0$.

Now consider the case of agents with rank $\rho < \frac{\alpha}{q} - \epsilon$. We define the sequence $\rho_K \equiv \alpha \sum_{k=0}^K (1-q)^k$. 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, $A(\rho_K - (K+1) \times \nu) \leq q\rho_K$.
- (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 $K = 0$. Since $\rho_0 = \alpha$, it is dominant for all players with rank $\rho \leq \rho_0$ to comply if they can. Since the exogenous non-compliance rate is q , it follows that with probability 1 as N gets large, $A(\rho_0 - \nu) \leq q\rho_0$.

We now show that $H_{K,\nu}$ implies $H_{K+1,\nu}$. Indeed, since $A(\rho_K - (K+1) \times \nu) \leq q\rho_K$ with probability approaching 1, this means that spare audit capacity that can be used on agents with rank greater than $\rho_K - (K+1)\nu$ is greater than $\alpha - q\rho_K$ with probability one. This implies in all strategy profiles surviving $K+1$ iterations of iterated elimination of dominated

strategies, non-complying agents with rank less than $\rho_K - (K+1)\nu + \alpha - q\rho_K = \rho_{K+1} - (K+1)\nu$ get audited with probability 1. Hence, all strategy profiles surviving $K + 1$ rounds are such that agents with rank less than $\rho_{K+1} - (K + 1)\nu$ settle. By the Law of Large Numbers, this implies that $A(\rho_{K+1} - (K + 2)\nu) \leq q\rho_{K+1}$ 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). ■

Corollaries. Proposition 3 admits two corollaries that are relevant for applications. First, in a realistic context, the principal may be able to estimate a “demand curve” for settlement $Q(P)$, mapping settlement offer P to a non-compliance rate $q = Q(P)$. For instance, if there was uncertainty about the actual amount of taxes due, we would have $q = Q(P) = \text{prob}(D > P)$. In this context, Proposition 3 implies the following:

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 expression for profits differs from the usual expression for revenue $P \times (1 - Q(P))$ because of the endogenous externality across agents. Non-compliant agents not only fail to settle, but they reduce the incentives of lower priority agents to settle.

Another realistic possibility is that agents are heterogeneous. Assume that agents belong to a finite number of groups indexed by $g \in \{1, \dots, G\}$, associated with population weights m_g (summing to 1), settlement offers P_g and non-compliance rates q_g . In addition, the auditing difficulty of different groups may vary, so that auditing an agent from group g consumes λ_g from the principal’s auditing capacity.

Corollary 2 (heterogeneous agents). *As N gets large, the maximum profit Π the principal can achieve under any prioritized enforcement mechanism is asymptotically equal to*

$$\max \left\{ \sum_{g=1}^G \delta_g m_g (1 - q_g) P_g \mid (\delta_g)_{g=1}^G \in [0, 1]^G \text{ such that } \sum_{g=1}^G \delta_g m_g q_g \lambda_g \leq \alpha \right\}. \quad (2)$$

This maximum is achieved by prioritizing groups in order of decreasing score $\frac{(1-q_g)P_g}{q_g \lambda_g}$, and setting an arbitrary ordering of agents within each group.

4 Dynamic Settlement

In this section, we embed the decision to comply or not over time. This is realistic: in practice agents need some time to respond. In addition this allows us to explore two relevant policy dimensions: discounts for early settlement of taxes (equivalent to penalties for late payment), as well as revealing information about the compliance of others as in Del Carpio (2014), Castro and Scartascini (2015), Dwenger et al. (2016).

We consider the following variant of the static game introduced in Section 3. Time $t \in [0, 1]$ is continuous. The principal commits to a deterministic settlement schedule $(P_t)_{t \in [0, 1]}$. Each agent $i \in \{1, \dots, N\}$ becomes able to settle according to a Poisson process with intensity $1 - q$. 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. 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 priorities \prec_{CK} .

We allow the principal to 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 $z_{i,t}$ 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 .

The following benchmark is useful. If (i) the path of settlement prices is constant, i.e.

$P_t = P$ for all $t \in [0, 1]$, and (ii) no information is provided to agents over time, then equilibrium settlement in this game will be identical to equilibrium settlement in the static game with frictions of Section 3. The timing of settlement, however, will be indeterminate: agents may settle as early as possible, or right at the end of the game. From this benchmark, providing price discounts would incentivize early settlement. In contrast, providing information about the settlement of others would delay settlement (so that players can capture the option value of information). This suggests that both discounts and information policy may have a non-trivial effect on settlement behavior. We show that this is not the case: to a first order approximation, both information and discount design are irrelevant in equilibrium.

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

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

Proposition 4 (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.*

Proof. The proof is closely related to that of Proposition 3. 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 mass of agents with rank less than ρ who get audited. 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 be audited 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 audits $A(\rho_{K-1})$ converges to $q\rho_{K-1}$ as N grows large. This means that the spare audit 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 audit 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 $-qD$. 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.⁷ Let Settle denote the event that the agent settles, Delay denote the event that the agent

⁷In other terms, payoffs conditional on being able to settle at some point are an affine transformation of unconditional payoffs.

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

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

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. ■

5 An Experiment

Our theoretical analysis makes three potentially important points for practice:

- (i) random enforcement can lead to multiple equilibria, and in particular to a low collection equilibrium;

- (ii) prioritized enforcement selects a high enforcement equilibrium that can improve on random enforcement, even in the presence of frictions;
- (iii) information design (as well as discounts for early settlement) has a negligible impact on collection.

However, properties of prioritized enforcement rely crucially on agents being able to perform many iterated eliminations of dominated strategies. In contrast, experimental work identifying levels of rationality (see for instance Camerer et al., 2004, Costa-Gomes and Crawford, 2006) suggests that in many settings, a majority of players engage in less than two rounds of iterated best-reply.⁸ This section investigates the extent to which points (i), (ii), and (iii) hold in experimental data. We show that prioritized enforcement improves over random enforcement, but that prioritized enforcement falls short of theoretical expectations, likely because of bounded rationality frictions. In this context, we show that information has a role to play by making it easier for players to optimize, and anticipate the behavior of others.

Baseline game. We implement essentially as is the dynamic settlement game of Section 4, with the experimenter playing the role of the principal, and recruited participants playing the role of agents. Parameters were specified as follows: the number N of agents was set to 10. All agents received an initial endowment of 100 points and owed the same amount $D = 100$ pts. The initial and final settlement prices P_0 and P_1 were set to $P_0 = 80$ pts and $P_1 = 92$ pts. Settlement prices evolved linearly over time. Time $t = 1$ corresponded to 30 seconds.

The probability q with which an agent is exogenously unable to settle was set to $q = 20\%$. The principal's audit capacity was also set to $\alpha = 20\%$. This implies that under the benchmark of Proposition 4, in principle, the share of agents able to settle who do so should be close to 1.

⁸Nonetheless, in dominance-solvable games, a non-zero share of players seems to play in equilibrium.

The key departure from Section 4 is that we constrained the realization of times after which agents were able to settle. This would not change our findings for large numbers of players. Our restrictions seek to reduce sampling noise by ensuring that all realizations are typical enough. 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.

Treatments. We experimented with four treatments corresponding to different enforcement policies and different information designs. In a random enforcement treatment, participants were not informed of the order in which they would be audited, 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.

We implemented three prioritized enforcement treatment. In each of these treatments, participants were informed of their rank in the audit line. The treatments differed in the additional information provided to participants about the settlement behavior of others:

- In the **priority only** treatment, players were given no information about the realized settlement of others.
- In the **aggregate information** treatment, players were informed of the aggregate settlement rate in their group. This matches taxpayer information experiments (Del Carpio, 2014, Castro and Scartascini, 2015), as well as experiments seeking to increase pro-social behavior through norms (Allcott, 2011). Additionally, such information may foster learning if a majority of players settle, and some behavioral players choose their play by imitating others.
- In the **targeted information** treatment, players are informed of both the aggregate set-

tlement rate in their group, and of their *real time effective rank*, i.e. their audit rank after taking into account settlement by other players. This targeted information has an immediate strategic interpretation for players: if their effective rank is 1 or 2, it is dominant for players to settle immediately.

Protocol. Experimental sessions took place on MTurk between March 2020, and August 2020. 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 experimental instructions were conveyed to players through their browser. Screenshots of instructions are reproduced in Appendix A.⁹

We ran 20 sessions, with 40 participants in each session. Participants were randomly assigned to the 4 treatment groups, and played the collection game 4 times. The first collection game did not count towards participants' final payoff. Points earned in the last three collection games were averaged across games, and converted to cash at the rate of USD8 for 100pts. Players were *not* reallocated across different treatments over time. To reduce noise due to sampling, the realization of the times at which players can start settling, which we call *wake-up times*, was kept the same across different treatments of a given session.

On average, completing the experiment took 23 minutes. Participants earned a USD3.5 fee for showing up at a pre-announced time. A subset of 40 participants was selected from this group to continue with the experiment. Participants earned between USD0 and USD8 from their play in the collection game, with mean earnings at USD3.2. For reference the mean payoff from settling as soon as possible was equal to USD.992.

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 50 tasks on MTurk.¹⁰ Experiments were all run between 11am and 4pm EST, Monday to Friday.

⁹A preliminary version of the experiment (excluding the targeted information treatment) ran in a physical lab for 12 sessions before the COVID19 pandemic started. Findings from this preliminary round are qualitatively similar.

¹⁰These selection criteria are meant to filter out bots.

6 Findings

Throughout this section, we report findings for payoff relevant games. We begin with the treatment effect of different enforcement mechanisms on settlement behavior.

6.1 Treatment Effects




Figure 1: Mean settlement rate by treatment conditional on waking up.

	mean settlement rate	pct point diff over random	pct diff over random
random	0.477		
priority only	0.544	6.8 (.02)	14.2%
aggregate info	0.596	11.9 (.002)	25.0%
targeted info	0.666	18.9 (.000)	39.7%

Table 1: Settlement rates across treatment conditional on waking up.

Two-sided p-values in parentheses. Treatment effect estimates include session fixed-effects; standard-errors are clustered at the (treatment, session) level.

Mean settlement by treatment. Figure 1 and Table 1 display the mean settlement rate across treatments in the population of agents that are able to settle at some point. The key takeaways are the following:

- Roughly 50% of players settle under random enforcement, halfway between the high compliance equilibrium, and the low compliance equilibrium.
- Prioritized enforcement improves significantly over random assignment, but does not attain the theoretical 100% settlement rate.
- Contrary to the prediction of Proposition 4, information, and especially targeted information, helps improve settlement rates.




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

The distribution of settlement rates. Figure 2 plots the c.d.f. of group-level settlement rates by treatment. The most noticeable fact is that prioritized enforcement doesn't just increase the mean settlement rate, but induces a first-order stochastic dominance (FOSD) increase in settlement rates. This means that even if the principal is risk averse over settlement and collection outcomes, prioritized enforcement improves on random enforcement.

Focusing on prioritized enforcement treatments, there is no FOSD ranking of the priority only treatment and the priority with aggregate information treatment. Specifically, the aggregate information treatment seems to increase the mean settlement rate, but perhaps at the cost of a small increase in the spread of the distribution of settlement rate. In contrast, the targeted information treatment generates a first order increase in settlement shares compared with any other treatment.

Finally, we note that the distribution of settlement under random enforcement is unimodal. This suggests that the data is not well explained by a mixture of agents coordinating on low and high enforcement equilibria.

Collection. Figure 3 shows that the findings of Figure 1 are also reflected in average collection amounts. Collection averages to 40.1pts under random enforcement, and respectively increases to 46.2, 50.9, and 56.8 under prioritized enforcement alone, with aggregate information, and with targeted information.




Figure 3: Mean settlement amount by treatment.

6.2 Delay

In principle, giving players more information may delay settlement: the anticipation of receiving information creates an option value for waiting since it may reveal that other players are in fact not settling.

Information does not delay settlement. We define settlement delay as the difference between the time at which a player wakes up and the time at which they settle. Figure 4 illustrates the distribution of settlement delay for players who do wake up. The settlement delay of players who choose not to settle is set to 30s. Figure 4 shows that prioritized




Figure 4: CDF of settlement delay conditional on waking

enforcement causes a FOSD reduction in delay compared to random enforcement. The impact of aggregate information on delay is more ambiguous: under prioritized enforcement, no information yields a higher mass of settlement times under 5s than providing aggregate information, but the difference is small. Targeted information yields a more sizeable FOSD reduction in delay relative to all other treatments.

Information delays settlement, conditional on settlement. While information does not increase delay, Figure 4 suggests that both targeted information and aggregate information increase the mass of players settling after some delay, say after 10s from waking up. This is confirmed by Figure 5 which plots the CDF of settlement delay conditional on settling.




Figure 5: CDF of settlement delay conditional on settling

Conditional on settlement, prioritized enforcement without information decreases settlement delay compared to random enforcement. In contrast, prioritized enforcement with either aggregate or targeted information increases settlement delay compared to random enforcement. Targeted information causes the largest increase in delay.

6.3 Rationality

The impact of information on settlement and delay suggests that information is effective through its interaction with players' rationality. We are specifically interested in understanding the rationality of players in prioritized enforcement treatments with no information

and targeted information since they correspond to very clearly interpretable information environments.

In principle, several possible forces may be active. First, information can reduce the strategic difficulty of decision problems faced by agents. Keeping rationality the same, this may increase convergence to equilibrium behavior. Second, information may affect rationality. It may increase rationality by making the incentives of other players more transparent. Alternatively it may decrease rationality. For instance, targeted information may lead players to stop thinking about the incentives of others and instead adopt the rule-of-thumb consisting in settling whenever their effective rank (i.e. their rank net of higher ranked agents who have settled) drops to 2.

Correspondingly, we are interested in distinguishing the three following scenarios:

- A. Information decreases the rationality of players; as a result agents settle only when it is dominant to do so.
- B. Information does not affect the rationality of players, but increases the likelihood of information sets at which settling is a best-response given the rationality of players.
- C. Information increases the rationality of players.

As Proposition 4 shows, the game is essentially dominance solvable, with higher rank players requiring more iterated eliminations of strategies to settle. Correspondingly, we can use players' rank at time of settlement as an indicator of rationality.

We are interested in both the players' *effective rank* and *max rank* at the moment they wake up, and at the moment they make a settlement decision, if they do. Effective rank corresponds to a player's initial rank net of the number of players ahead of them who have chosen to settle. Effective rank is the same across both treatment but is only in the information set of players in the targeted information treatment. In contrast, max rank corresponds to the maximum effective rank consistent with the player's information. Under the targeted information treatment, max rank and effective rank are identical. Under the priority only treatment, max rank is equal to a player's initial rank. Settlement at effective rank 1 or 2

indicates a single round of best-reply. Settlement at ranks 3 or higher indicates strategic thinking taking into account the behavior of others.

Rank at settlement is ambiguous. Figure 6 displays the proportions of the overall population choosing to settle at different effective ranks.¹¹ Targeted information results in more agents settling at most ranks. There is a particularly large increase in the number of agents settling at effective rank 2, but also a noticeable increase in the number of agents settling at effective ranks 3 and 4.




Figure 6: Share of population settling by effective rank at settlement.

Considering players' max rank at the time of settlement is a better reflection of rationality since players in the no information treatment do not receive information about their effective rank. The only hard information they have is their initial rank. Figure 7 displays the proportions of the overall population choosing to settle at different max effective ranks. As expected, the max rank of agents settling under priority only is shifted upwards. There is a greater mass of agents settling with max rank equal to 3 under targeted information, but

¹¹The population shares displayed do not sum to 1: players who choose not to settle are the excluded category.

the mass of agents settling with rank 4 are similar across treatments. This suggests that targeted information acts mainly through reducing the difficulty of decision problems, and not by increasing rationality. This hypothesis is confirmed by further analysis.




Figure 7: Share of population settling by max rank at settlement.

Targeted information reduces max rank at wake. Figure 8 illustrates the distribution of agents' max rank at wake (population shares sum to one since we condition on agents' waking). This is not informative of their rationality. Rather, it clarifies that targeted information leads to a different distribution of decision problems than the no information treatment: mass assigned to difficult problems, corresponding to a rank of 8 or higher, is shifted to relatively simpler problems of a rank 4 or less.

We note that for both treatments, the distribution of max rank at wake for players with a max rank below 4 is approximately uniform. This means that focusing on the subset of events where an agent wakes up with a max rank less 4 allows us to control for the impact of information on the difficulty of problems faced by agents, and identify the impact of information on rationality.




Figure 8: Distribution of max rank at wake.

Settlement given max rank at wake, and rationality. Figure 9 displays the settlement rate of agents conditional on their max rank at wake, for max rank at wake less than 4. Conditioning on max ranks such that both no information and targeted information have the same distribution of max rank lets us partial out the effect of information on the strategic difficulty of decision problems agents face.

Figure 9 shows that controlling for the difficulty of decision problems entirely removes the increase in settlement rates associated with targeted information. The no information and targeted information treatment lead to the same ultimate settlement behavior. In other terms, targeted information does not act by increasing the strategic sophistication of players.

Figure 10 confirms this conclusion. It plots the distribution of players' max rank at settlement, conditional on settling, and their max rank at wake being less than 4. If anything, the settlement behavior of players under targeted information reflects lower strategic sophistication than the settlement of players under no information.

This decomposition of the aggregate effect of information suggests that scenario B is in




Figure 9: Settlement rate by max rank at wake (max rank at wake ≤ 4).

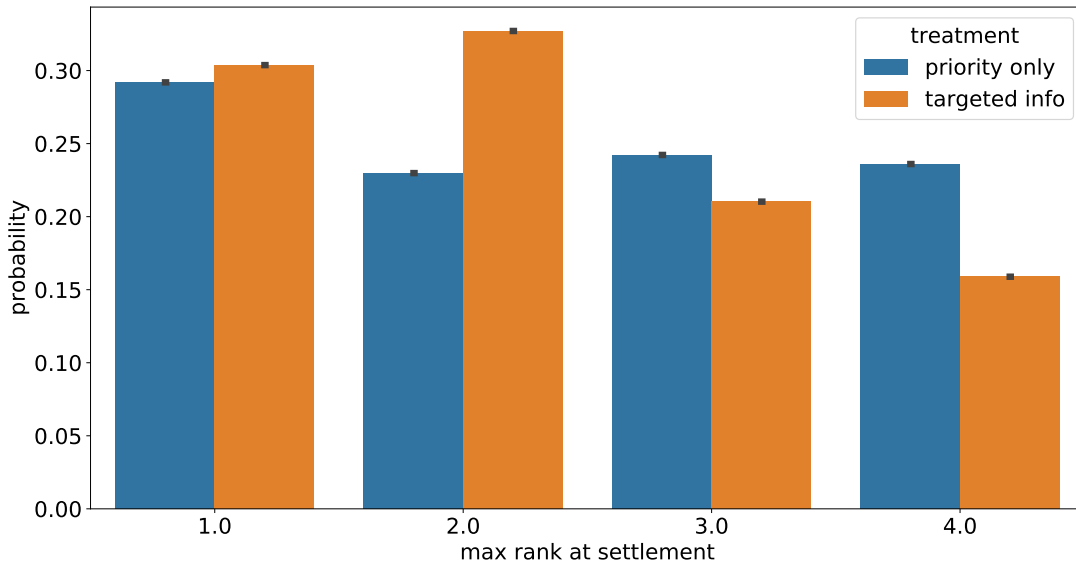


Figure 10: Distribution of max rank at settlement (max rank at wake ≤ 4).

fact the most plausible. Targeted information acts by reducing the strategic difficulty of problems faced by agents. It does not change the strategic sophistication of players (by improving or reducing it) in a large way. If anything, targeted information may slightly reduce the strategic sophistication of players.

7 Discussion

We believe that divide and conquer schemes have wide applicability. Chassang and Chen (2020) provides evidence that divide-and-conquer can be effective in practice based on Operation Ceasefire programs implemented in a large number of US cities over the last 25 years (Braga et al., 2001, Kennedy, 2011, 2012). Operation Ceasefire was initiated in Boston in the mid-90s during a wave of gang related homicides. It first clarified to gangs that the police were in fact fairly good at associating homicides with gang, even though bringing together actionable evidence valid in a court of justice was much more difficult. Second, it shifted some police resources away from uniform enforcement, and instead, promised that a significant share of police effort was going to be directed towards making life difficult for the *first gang* suspected of committing a murder. This is a clever prioritized enforcement strategy that uses a natural common-knowledge rank: the time at which a murder is committed.

Besides debt collection, we believe that prioritized enforcement offers particular promise in fighting 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 an organization: for instance, it may be that cops who take bribes are better than no cops at all. Prioritized enforcement provides a way to initiate organizational change without firing all misbehavers. Proposition 2 provides potential guidance on how to rank agents: prioritized high offenders most likely to comply if given a take-it or leave-it offer.

We conclude with a discussion of limits and possible extensions to our theoretical and experimental analysis.

Dissuasive punishment. The analysis of Sections 3 and 4 assumes that collection in the event of a successful audit is equal to the amount owed D . It turns out that the asymptotic collection rate α/q is not affected if a penalty is applied, so that the amount collected following an audit is $(1 + \gamma)D$ with $\gamma > 0$. The reason for this is that agents with normalized

rank $\rho > \alpha/q$ have a probability 0 of being audited. As a result, raising penalties has a negligible effect on them. However, when the population is finite and fixed, dissuasive penalties certainly raise settlement rates.

Information is relevant in a Level- k model. The experimental results of Section 6 falsify the theoretical prediction of Section 4: information, and especially targeted information, has a large impact, likely because it simplifies the decision problems faced by agents. This observation is amenable to theoretical modeling. Specifically, we can use k -iterated-best-response as our solution concept (with values of k relatively small, say 3 or 4). In that case, even as the sample size N grows large, information design makes a difference. Whenever $\alpha < q$, then under no information, the fraction of agents who settle will be bounded away below the rationalizable threshold α/q , even as the population size N gets large. In contrast, under targeted information, it is sufficient that players be capable of best-replying in order to guarantee that the share who settle approaches α/q as N gets large.

Incomplete information over rank. Halac et al. (2020) show that providing players incomplete information over their rank can improve the effectiveness of divide-and-conquer strategies. Our analysis suggests that information can be quite helpful if it does not require high degrees of sophistication to be interpreted. We believe that there may be ways of providing information that exploit both ideas. For instance, under targeted information, it is likely unimportant that lower ranked players know their rank. It may be more effective to keep such players in some amount of doubt.

Aggregate uncertainty about the settlement rate. In Section 4, providing information about the settlement behavior of other players has little impact because there is essentially no aggregate uncertainty to be uncovered. This would not be the case if the friction level q was uncertain and needed to be learned. In that case, providing information about friction rate q would have an impact. In the context of our experiment, the aggregate

information treatment would provide such information.

Preliminary analysis suggests that providing information about q would raise expected collection. Indeed, consider the extreme case where a small sample of the overall large population is used to estimate q , which is then revealed to remaining agents. The continuation game is the one studied in Section 4, so that a share $\frac{\alpha}{q}$ agents choose to settle. Since this is a convex function of q providing such information raises the expected collection rate.

Small versus large groups. The analysis of Sections 3 and 4 considers games in which agents are put together in a single large group, and follow a single order. An intuitive practical question is whether it is helpful to split agents in many fixed size groups in which a small fixed number of agents can be audited.

Theory suggests that the effect is at best ambiguous. Consider the case of targeted information in the case where $\alpha \geq q$. Proposition 4 implies that all agents who are able to settle will settle. This is not the case in the small group case: due to random variation, some small groups will have their audit capacity exhausted, leading low rank individuals not to settle. As a result, under a many-small-groups design a positive share of agents who can settle will choose not to.

What would happen in an experiment or in the field is less clear. Intuitively, it may be that smaller group help place agents in a context where they think strategically. Also if players react with delay, it may be that large groups cause very large settlement delays. In contrast small groups may allow for parallel processing. We speculate that this force is important in practice.

A Player instructions

This section reproduces instructions given to participants in different treatments.

A.1 Instructions for Random Treatment

You have 10 Minutes to read the instructions. At the end of 10 minutes, the next page will load. These instructions will be available to you in the rest of the experiment as a pull-out tab on the left hand side of your screen. The first round is practice and so results there do not count towards final payment.

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 more significant 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 profit, which is calculated in points. At the end of the experiment your points will be converted into US Dollars (USD) according to the following exchange rate:

$$100 \text{ points} = \text{USD } 8$$

There is a USD 3.5 participation fee for the experiment. At the end of the experiment, we will calculate your **average points across rounds**. We then convert the points average using the above exchange rate.

Summary of the Experiment

In this experiment, an *automated collection authority* wants to extract money from you and other participants. Your goal is to finish the experiment with the most possible amount of money. General details are:

- There are 10 participants in this experiment, including you.
- There are 4 rounds including 1 practice round
- rounds take around 2 minutes and the entire experiment should last roughly 35 minutes.
- each round consists of 2 stages.
 - stage 1: you may receive a settlement offer from the *collection authority*.
 - stage 2: the collection authority will be able to collect directly from two participants according to a rule that will be explained in the *Collection Stage — Details* tab.

Timeouts

A timer may be displayed on some pages. If you do not complete the page (by clicking next, answering questions, or accepting settlement), you will be automatically pushed to the next page. The same is true for all other players.

Stages Overview

Collection Overview

You start each round with 100 points, but the automated collection authority (CA) can claim up to 100 points from you. However, the CA must directly *investigate* you in order to collect payment. If the CA investigates you, then you will certainly pay 100 points; otherwise you keep all 100 points. The CA is only able to investigate **2 players**. Details of the collection procedure are given in the *Collection Stage — Details* tab.

Settlement Overview

Prior to the collection stage, the CA offers you and all other players an identical settlement opportunity to avoid investigation. This stage is described below:

- Each round, you will enter the settlement stage with the other participants
- During settlement stage, you will have 30 seconds to accept the offer made by the CA
- There is *delay* in how quickly you can accept. Furthermore, 20 percent of participants in each round will NOT have an opportunity to settle at all. Whether or not you will have an opportunity to settle is randomly assigned each round.
- If you don't settle, you enter collection stage, where you may be investigated and forced to pay 100 points

Collection Stage — Details

Collection Procedure

The collection authority will RANDOMLY choose 2 players among those who don't settle to pay 100. Other players who do not settle pay 0.

Settlement Stage — Details

Your Decision

You start the round with 100 points. You will be offered a settlement by the collection authority to keep an amount -- this amount decreases over time. Before the round starts, you will be shown a screen with the initial offer, the rate of decrease per second and the final offer. If you do not accept the offer by the deadline and the collector chooses you (according to the procedure in the *Collection Stage — Details* tab), you will pay 100. If you are not chosen, you will pay 0.

Delayed Decision Opportunity

The button to accept a settlement offer will not be immediately available. The button will become available after a random amount of time. Once the acceptance button becomes clickable, it will stay clickable. **For 20% of participants each round, the button never becomes clickable** (in which case the settlement cannot be accepted). Whether or not you are among those whose button becomes clickable is randomly assigned each round.

Other Players

All other players are offered the same settlement. Their buttons become clickable after a random amount of time (this time IS NOT the same for all players). For 20% of all participants in a round, the button will not become clickable. Whether or not a participant's button becomes clickable is randomly assigned each round.

Snapshots

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

The snapshot below is of the discount page that will be shown before the settlement stage. This page describes the settlement offers that will be active that round.

Instructions

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

Time left to complete this page: 0:24

Discounts in this Round (#1)

This screen will be shown each round. Values will not change across rounds.

Amount you keep if you accept at start of round : 20

Rate of decrease per second : 0.4

Amount you keep if you accept at the end of the round: 8

Continue to settlement stage

Indicates the round number

Amount you keep if you settle immediately at the start of the round

Each second, the amount you keep if you settle decreases by this amount

After 30 seconds of decreasing at the rate above, this is the final amount you keep if you settle at the end of the round.

The snapshot below is of the settlement stage.

Instructions

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

Time left to collection stage: 0:27

Settlement Stage

Amount you keep if you accept offer now: 9.8

Settlement cannot be accepted yet

Time until you are sent to the collection stage

Amount you keep if you immediately accept settlement

Indicates you cannot yet accept the settlement offer

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

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

Instructions

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

Time left to collection stage: 0:22

Settlement Stage

Amount you keep if you accept offer now: 10.0

Accept Offer

At the end of the round, you will see a results page.

Instructions

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

Time left to complete this page: 0:16

Collection and Results Stage

You avoided investigation! ← Your (practice) payoff this round: 100.00 points

Continue to next round?

Indicates that you were NOT investigated -- you will also be informed in the event that you are investigated

Your payoff from the current round is shown here. In this round, which is practice, you were not investigated, so you keep your initial endowment of 100 points.

Payment

At the end of the experiment, we will calculate your average points across rounds. We will then convert the points average to US Dollars (USD) using the exchange rate:

100 points = USD 8

To this amount we will then add the USD 3.5 participation fee.

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 0:09

Settlement Stage

Amount you keep if you accept offer now: 11.31

Accept Offer

A.2 Instructions for Priority Treatment

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

Collection Stage — Details

Collection Procedure

The CA will choose according to a pre-specified list. You will be assigned an initial position in line at the start of the settlement stage. The two players with the **lowest** initial positions in line who don't settle (either by choice or because they didn't have opportunity) are investigated and forced to pay 100. The lowest position is 1 and the highest is 10. Others who do not settle pay 0. In the settlement stage, your initial position in line is displayed to you as follows. If your initial position is X (where X can be between 1 and 10), you will be shown the phrase: **Your initial position in line to be investigated is X.**

During the game, players were shown the following screen, reminding them of their *initial rank*.

Time left to the collection stage 0:20

Settlement Stage

Amount you keep if you accept offer now: 15.86

Your **initial position in line to be investigated is **7****

Accept Offer

A.3 Instructions for Priority with Aggregate Information Treatment

The instructions are identical to the priority treatment, except for an added description of the information (and the snapshots page).

Information

You will receive information on the share of players who have decided to accept settlement. For instance, if 1 player has settled, you will be shown **1/10**

During the game, players were shown the following screen, informing them of both their initial rank, and of the aggregate share of players in their group who had settled.

Time left to the collection stage 0:02

Settlement Stage

Amount you keep if you accept offer now: 8.41

Your **initial** position in line to be investigated is **10**

Share Settled
1/10

Accept Offer

A.4 Instructions for Priority with Targeted Information Treatment

The instructions are identical to the priority treatment, except for an added description of the information and an updated version of the collection stage (and the snapshots page).

Information

You will receive information on the share of players who have decided to accept settlement. For instance, if 1 player has settled, you will be shown **1/10**. You will also receive information about your location in line to be investigated, which is updated as other participants choose to settle. For instance, if your initial position is 3 and the participant with initial position 1 chooses to settle, then you will move into 2nd in line to be investigated. 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**.

Collection Stage — Details

Collection Procedure

The CA will choose according to a pre-specified list. You will be assigned an initial position in line at the start of the settlement stage. Your position in line will be updated throughout the round (described in the Information section of the *Settlement Stage — Details* tab). If a player does not settle (either by choice or because they didn't have opportunity) and ends the settlement stage with a position in line less than or equal to **2**, then they are investigated and forced to pay 100. Others who do not settle pay 0. In the settlement stage, your current position in line is displayed to you as follows: If your current position is X (where X can be between 1 and 10), you will be shown the phrase, **Your current position in line to be investigated is X.**

During the game, players were shown the following screen, informing them of their effective rank.

Time left to the collection stage 0:10

Settlement Stage

Amount you keep if you accept offer now: 11.72

Your **current** position in line to be investigated is **4**

Share Settled
1/10

Accept Offer

References

- ABREU, D. AND H. MATSUSHIMA (1992): “Virtual implementation in iteratively undominated strategies: complete information,” *Econometrica: Journal of the Econometric Society*, 993–1008.
- ALLCOTT, H. (2011): “Social norms and energy conservation,” *Journal of public Economics*, 95, 1082–1095.
- 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.
- CAMERER, C. F., T.-H. HO, AND J.-K. CHONG (2004): “A cognitive hierarchy model of games,” *The Quarterly Journal of Economics*, 119, 861–898.
- 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. AND M. CHEN (2020): “Social Change as a Real Option: the Case of Operation Ceasefire,” *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.
- COSTA-GOMES, M. A. AND V. P. CRAWFORD (2006): “Cognition and behavior in two-person guessing games: An experimental study,” *American Economic Review*, 96, 1737–1768.
- 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.

- DE CLIPPEL, G., R. SARAN, AND R. SERRANO (2019): “Level-k mechanism design,” *The Review of Economic Studies*, 86, 1207–1227.
- DEL CARPIO, L. (2014): “Are the neighbors cheating? Evidence from a social norm experiment on property taxes in Peru,” *Unpublished Manuscript, Princeton University*.
- DUFLO, E., M. GREENSTONE, R. PANDE, AND N. RYAN (2013): “Truth-telling by third-party auditors and the response of polluting firms: Experimental evidence from India,” *The Quarterly Journal of Economics*, 128, 1499–1545.
- 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.
- 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. (2006): “Learning dynamics for mechanism design: An experimental comparison of public goods mechanisms,” *Journal of Economic Theory*, 129, 114–149.
- HEALY, P. J. AND L. MATHEVET (2012): “Designing stable mechanisms for economic environments,” *Theoretical economics*, 7, 609–661.
- JACKSON, M. O. (1992): “Implementation in undominated strategies: A look at bounded mechanisms,” *The Review of Economic Studies*, 59, 757–775.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian persuasion,” *American Economic Review*, 101, 2590–2615.
- 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.
- MASKIN, E. (1999): “Nash equilibrium and welfare optimality,” *The Review of Economic Studies*, 66, 23–38.

- 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.
- ORTNER, J. AND S. CHASSANG (2018): “Making Corruption Harder: Asymmetric Information, Collusion, and Crime,” *Journal of Political Economy*.
- POMERANZ, D. (2015): “No taxation without information: Deterrence and self-enforcement in the value added tax,” *American Economic Review*, 105, 2539–69.
- SEGAL, I. (2003): “Coordination and discrimination in contracting with externalities: Divide and conquer?” *Journal of Economic Theory*, 113, 147–181.
- WINTER, E. (2004): “Incentives and discrimination,” *American Economic Review*, 94, 764–773.