Making the Most of Limited Government Capacity: Theory and Experiments

Sylvain Chassang*,† — NYU and NBER
Lucia Del Carpio — INSEAD
Samuel Kapon — NYU

September 25, 2019

Abstract

Limited government capacity frequently results in multiple equilibria. If most agents comply with government policy, then limited enforcement is sufficient to dissuade any given agent from misbehaving. If most agents do not comply, then limited enforcement capacity has a negligible impact on incentives. We study the extent to which properly designed enforcement priorities can help select a high compliance equilibrium. Our analysis emphasizes the impact of design features such as discounts and information provision both in theory, and in experimental play. While large discounts and information provision have a negligible impact on collection in theory, we believe that the two are effective complements in practice.

KEYWORDS: government capacity, limited enforcement, common knowledge enforcement priorities, bounded rationality, mechanism design.

1 Introduction

Limited government capacity frequently results in multiple equilibria. If most agents comply with government policy, then limited enforcement is sufficient to dissuade any given agent from misbehaving. If most agents do not comply, then limited enforcement capacity has a

*Chassang: chassang@nyu.edu, Del Carpio: lucia.delcarpio@insead.edu, Kapon: sdk301@nyu.edu.

†We are especially grateful to Vishal Ashvinkumar, Kathleen Ngangoue, Paula Onuchic, and Tianzan Pang for feedback on our experiment design. Early discussions with Dilip Abreu, Abhijit Banerjee, Navin Kartik, John Klopfer and Alessandro Lizzeri helped refine our thoughts on this topic.

negligible impact on incentives. We study the extent to which properly designed enforcement priorities can help select a high compliance equilibrium.

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 shows that in theory, discounts and information provision 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 provides some pre-analysis. Appendix A provides a sample of our instructions.

2 Static Enforcement without Errors

2.1 Framework

We flesh out our model in the context of tax auditing. N agents $i \in I \equiv \{1, \cdots, N\}$ 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 and can only audit αN agents with $\alpha \in (0, 1)$. The objective of this paper is to better understand how different carrot-and-stick approaches can help the principal collect efficiently in spite of limited enforcement capacities.

The principal can make settlement offers and commit to an audit schedule according to the following extensive-form game:

1. The principal gives each agent the possibility to settle at a common price P.

2. Agents simultaneously decide whether or not to settle and pay price P or not.

3. The principal audits agents who have not settled according to a complete order $\prec$ over I.

We consider two possible enforcement priorities $\prec$:

- Random priorities $\prec_R$: audited agents are drawn sequentially ex post (i.e. period 3), with uniform probability and without replacement;
- Common knowledge priorities $\prec_{CK}$: the ordering is specified ex ante (i.e. period 1) and common knowledge among players.[1] We index agents by their rank i under $\prec_{CK}$.

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 through their strategic use.

Our solution concept throughout the paper is Bayes Nash Equilibrium.

Modeling assumptions. Audits are free to the principal but constrained in number. A key implication for our results is that the principal cannot commit to arbitrary enforcement off of the equilibrium path. A model with unlimited but costly audits would change our results if auditors are paid only in the event of an audit. In that case, the government could commit to audit all non-compliers. Since there would be no audits on the equilibrium path, we would obtain full enforcement at no cost. If instead maintaining the capacity for off-path response is costly on the equilibrium path — investigators must be hired, trained, and paid

[1] In principle it may be useful to induce incomplete information over enforcement priorities. We choose to focus on an intuitive focal design which we study in depth, both theoretically and experimentally.

even if there are no crimes to investigate — a model with costly audits would make the same predictions as a model with limited enforcement. The cost of audits can be interpreted as a Lagrangian on the principal's aggregate enforcement constraint.

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 late payment or misreports is 25%.

Finally, we note the mechanisms we consider can treat different agents differently, but within tight limits. Indeed, when agents are investigated according to a common-knowledge order, some agents know they have a high audit priority. We believe this is socially acceptable because compliant agents are all treated the same. Only non-compliers are potentially treated differently. In contrast we think that charging different agents a different amount would not be socially acceptable because it would treat compliant agents differently.[2]

2.2 The Value of Common-Knowledge Enforcement Priorities

Proposition 1 (multiple equilibria under random enforcement)**.** *Under random enforcement order $\prec_R$, for any settlement offer $P \in [\alpha D, D]$, there exists an 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 $P \notin [\alpha D, D]$.*

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.*

[2] Note that charging different agents a different price may in fact increase the amount of revenue that could be raised.

In Sections 3 and 4 we study how realistic frictions on the side of the agents perturb the effectiveness of common-knowledge priorities.

3 Static Enforcement with Errors

We now consider a variant of the game introduced in Section 2 in which agents are exogenously and independently unable to comply with probability q. As a result the largest amount the principal can collect is $\Pi = (1 - q)D$. We denote by $\mathcal{S}_q$ this game of static enforcement with errors.

As a result, some agents may fail to settle even though there is common knowledge that they will be audited with probability 1. Exogenous non-compliance can be interpreted in different ways. The agent may simply not be aware of the collection problem. The agent may experience a liquidity or preference shock stopping her from making a payment. Finally, the agent may know that the amount collected by the principal is erroneous, and that an investigation will prove she does not owe money.[3] In some of these settings, it may be reasonable to assume that q is a function of price P.

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 of agents N gets large.

Proposition 3. *Under common-knowledge ordering $\prec_{CK}$, for any price $P \in (0, D)$:*

(i) all equilibria take a threshold form: there exists ρ^* such that all agents with rank $\rho < \rho^*$ comply, and agents with rank $\rho > \rho^*$ do note comply;

(ii) for any $\epsilon > 0$, with probability 1 as N grows large

$$\min\left(1, \frac{\alpha}{q} - \epsilon\right) \leq \rho^* \leq \min\left(1, \frac{\alpha}{q} + \epsilon\right). \tag{1}$$

[3] Such errors are likely in a setting where the principal attempts to predict owed taxes on the basis of informative but imperfect data.

When q is truly exogenous, Proposition 3 implies that the principal will maximize collection by making a settlement offer P slightly below D. When error rate q depends on the offer P, taking the form $q(P)$, then the principal's profits will be asymptotically equal to

$$\Pi(P) \sim \min\left(1, \frac{\alpha}{q(P)}\right) P. \tag{2}$$

It may then be optimal to set a price strictly below D even as the number of agents N grows large.

4 Dynamic Settlement

4.1 Framework

In many practical settings, the decision to comply or not happens over time: agents need some time to respond. We are interested in how discounts and potential information provision affect the timing and effectiveness of collection.

We consider the following variant of static game $\mathcal{S}_q$ introduced Section 2. Time $t \in [0, 1]$ is continuous. The principal commits to a deterministic settlement schedule $(P_t)_{t \in [0,1]}$. Each agent $i \in \{1, \cdots, N\}$ becomes able to comply according to a Poisson process with intensity $1-q$. If an agent is able to comply at date t, she is able to comply at all further dates $t' \in (t, 1]$. Compliance 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 $\prec_{CK}$. Potential non-response implies that agents cannot settle at any point with probability q, as under game $\mathcal{S}_q$. We denote this dynamic collection game with errors by $\mathcal{D}_q$.

We first study this game in the case where price P_t is constant over time and equal to P and no information is provided to agents. We then study the impact of providing discounts for early settlement, so that $P_t < P_{t'}$ whenever $t < t'$. Finally we study the impact of information on collection.

4.2 Constant Prices and No Information

Assume that the path of prices is constant, i.e. $P_t = P$ for $t \in [0, 1]$, and that no information is provided to agents over time.

Proposition 4 (static equivalence, indeterminate timing)**.** *The following hold:*

(i) Equilibria take a threshold form: there exists ρ^* such that all agents with rank $\rho < \rho^*$ ultimately comply if they are able to; all agents with rank $\rho > \rho^*$ do not comply.

(ii) A threshold ρ^* is an equilibrium threshold in dynamic game $\mathcal{D}_q$ if and only if it is an equilibrium threshold of static game $\mathcal{S}_q$.

(iii) Given an equilibrium threshold ρ^*, any feasible timing of settlement (consistent with the threshold rule) is in equilibrium. In the fastest settlement equilibrium, all agents settle as fast as they are able to. In the slowest settlement equilibrium, all agents settle at date $t = 1$ (if they can).

Of course indeterminacy of timing relies strongly on the fact that settlement prices P_t are constant. As we show below, it is possible to break indifferences towards fast settlement without sacrificing much payoff.

4.3 Discounts

We now show that by providing small discounts for early payments, it is possible to speed the timing of settlement, while also approaching the highest possible collection amount. In contrast providing large discounts results in collection amounts that are bounded away from the highest collection amount.

For the sake of concision we focus on weakly increasing price schedules such that $\forall t < t' \leq 1$, $P_t \leq P_{t'}$, and such that $P_0 > 0$ and $P_1 < D$. We refer to such price schedules as discounted. We say that a price schedule is *strictly discounted* if it is discounted, and for all $\epsilon > 0$, there exists $\eta > 0$ such that for all $t \in [0, 1]$,

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

Note that in this model, the decision to settle will not depend on rank ρ alone, but also on the time at which an agent becomes able to settle. This means that equilibria no longer exhibit a one-dimensional threshold structure, and the share of agents who do settle is a random variable. However, we show that when N grows large, the share of agents who settle becomes approximately deterministic. Let $\overline{s} \equiv \frac{1}{(1-q)N} \sum_{i=1}^{N} s_i$ denote the realized share of agents who settle among those who can in expectation (note that because we scale by $\frac{1}{1-q}$, we may have $\overline{s} > 1$ for some realizations).

Proposition 5. *(i) For any strictly discounted price schedule, agents settle as soon as possible, or do not settle at all.*

(ii) For any discounted price schedule, for any $\epsilon > 0$, with probability approaching 1 as N gets large,

$$\min\left(\frac{\alpha}{q} - \epsilon, 1\right) \leq \overline{s} \leq \min\left(\frac{\alpha}{q} + \epsilon, 1\right).$$

Together, points (i) and (ii) imply that optimal pricing schedules: should offer small price incentives for early settlement; but should *not* offer large incentives for early settlement. Indeed, Proposition 5 shows that the magnitude of discounts has at most a negligible impact on the share of agents who settle. However, large discounts will reduce the amount that agents settle for.

4.4 Information Provision

We now consider the possibility of information provision. In each period t, given a history of settlement decisions $h_t = (s_{i,t'})_{i\in\{1,\cdots,N\},t'<t}$ 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.

We emphasize that we do not introduce incomplete information regarding all aspects of the game. For instance we do not consider the possibility of providing the agents incomplete information about their enforcement priority. We believe this is an interesting question, but it is beyond the intended scope of this paper.

Proposition 6 (information provision is irrelevant)**.** *For any strictly discounted price schedule, and for any $\epsilon > 0$, as N becomes large,*

(i) with probability approaching 1, an agent with rank $\rho < \min\left(\frac{\alpha}{q} - \epsilon, 1\right)$ 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.

5 Empirical Predictions and Experiment Design

Table 5 summarizes the insights we wish to test, and what we expect the data to show.

	THEORETICAL PREDICTION	SUBJECTIVE EMPIRICAL PREDICTION
(1)	Common knowledge enforcement priorities improve collection over random priorities	True
(2)	Discounting affects the timing of settlement but does not impact ultimate collection	Maybe
(3)	A small discount achieves fast settlement while maximizing collection	False; a significant discount is needed to speed up settlement
(4)	Information has no impact on collection	False; information significantly improves the speed (and maybe the amount) of collection, but only in the presence of meaningful discounts

Table 1: theoretical and empirical predictions

5.1 Design

The game structure, treatments and design are described below.

Baseline game structure.

- Time interval $[0, 1]$ is 30 seconds long
- $\alpha = 0.2$, i.e. capacity for audit is 20%
- $q = 0.2$, i.e. the probability that an agent is unable to settle is 20%

Treatments:

- **A**: $\prec_R$ without information on the settlement decision of others
- **B**: $\prec_{CK}$ without information on the settlement decision of others
- **C**: $\prec_{CK}$ informed at all times of the *total share of others who decided to settle*

In each treatment, the discount level and rate of increase will be randomly pre-selected.

Design:

- students to be recruited at NYU Center for Experimental Social Science (CESS)
- 12 sessions with 15-25 players
- players split into two groups each round with random re-matching between rounds
- treatments in a session consist of a randomly pre-selected order of all three treatments A, B and C (with multiple rounds in each treatment)

6 Empirical Pre-Analysis

The key statistics that we are interested in are:

- Total amount settled
- Share of players settled
- Speed of settlement

Let X be a vector consisting of treatment characteristics: an indicator for whether or not information about others' settlement decisions was shown, discount level and slope, an indicator for random or priority audit, and indicators for the order in which the treatment was presented.

An observation is a single round of one of the treatments. For the total amount settled and the share settled, we'll run a regression on a second-order polynomial of X (to capture interaction effects), with clustering at the session level.

For the speed of settlement, we estimate a parametric model. Let $M(T, X)$ be the probability that a player settles by time T in a treatment with characteristics X. We parameterize $M(T, X)$:

$$M(T) := 1 - exp(-f(T, X))$$

where $f(T, X)$ is a linear function of T plus a second-order polynomial in X (again to capture interaction effects). We plan to estimate this model and then test for differences across treatments.

Appendix

A Instructions

We produce here the instructions that will be shown to players for treatment C (the others are similar).

A.1 Introductory Instructions

The first set of instructions are shown to participants at the beginning of the experiment.

Instructions

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. The decisions you make will have an impact on your profit. Your decisions will also influence the profit of other participants, just as the decisions of other participants may influence your profit. Your profit is calculated in points (plus the baseline participation fee). At the end of the experiment your points will be converted into US Dollars (USD) according to the following exchange rate:

100 points = USD 2

All payments are made based on point totals at the end of the experiment, plus a base pay of USD 10.

Summary of the Experiment: In this experiment, an automated collection agency 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:

- 10 participants in this experiment, including you.
- 3 treatments, with 5 rounds including 1 practice round in each treatment
- rounds take 30-40 seconds 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 agency will be able to collect directly from one participant according to a predetermined protocol.

Specific details will be given before each treatment to clarify the settlement and collection protocol used.

Collection: You start the round with 100.0 units, but the automated collection authority (CA) can claim up to 100.0 units from you. However, the CA must directly investigate you in order to collect payment. If the CA investigates you, then you will certainly pay 100.0 units; otherwise you keep all 100.0 units. The CA is only able to investigate 2 people.

Settlement: 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 have 30 seconds to accept the offer made by the CA
- There is delay in how quickly you can accept; with a 20% chance, you will not have an opportunity to settle at all.
- If you choose not to settle, you enter collection stage, where you may be investigated and are forced to pay 100.0 units
- There is a wikipedia link available to pass time; this is not relevant for payment

Timeouts: On each page will be displayed a timer. 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.

A.2 Comprehension Quiz

The players will next be given a comprehension quiz. The screenshot of this quiz is given below:

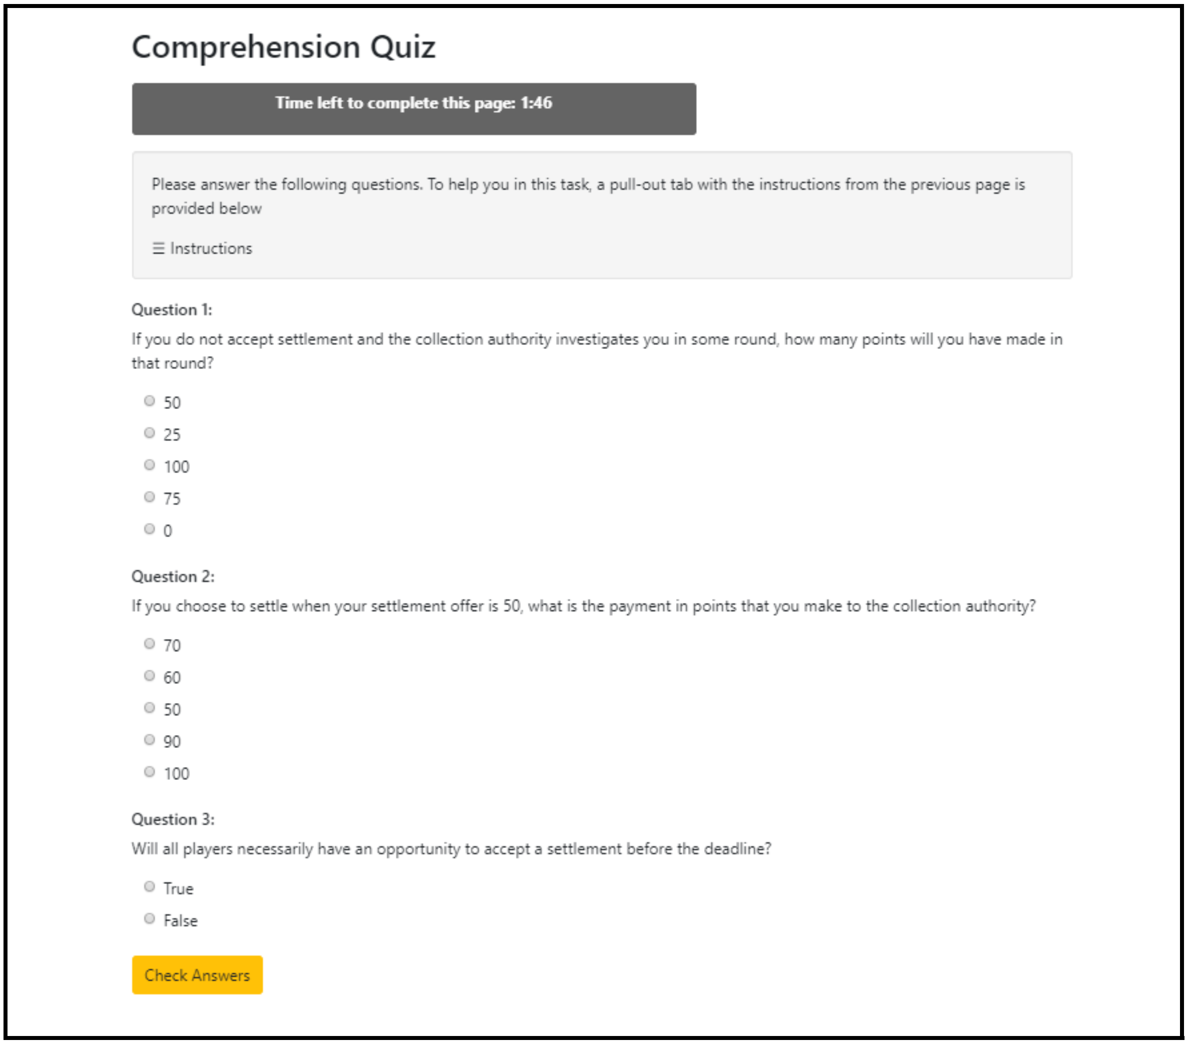

Comprehension Quiz

Time left to complete this page: 1:46

Please answer the following questions. To help you in this task, a pull-out tab with the instructions from the previous page is provided below

☰ Instructions

Question 1:

If you do not accept settlement and the collection authority investigates you in some round, how many points will you have made in that round?

- 50
- 25
- 100
- 75
- 0

Question 2:

If you choose to settle when your settlement offer is 50, what is the payment in points that you make to the collection authority?

- 70
- 60
- 50
- 90
- 100

Question 3:

Will all players necessarily have an opportunity to accept a settlement before the deadline?

- True
- False

Check Answers

The players will be given 2 minutes to complete the quiz. No penalty is given for incorrect answers, but they can progress faster if they answer correctly.

A.3 Treatment C

The treatment instructions for treatment C (priority audit with information) are presented below.

New Treatment

This is the beginning of a new treatment. The first round is practice and payoffs will not accrue in the total payoffs row of the results page. Starting in the second round, payoffs will count towards final payment and will begin to accrue.

Treatment Details

Your Decision: You start the round with 100.0 units. You will be offered a settlement by the collection authority to pay an amount – this amount increases over time by 0.01 each second.If you do not accept the offer by the deadline and the collector chooses you (according to the procedure below), you will pay 100.0. If you are not chosen, you will pay 0.

Collection Procedure: The CA will choose according to a pre-specified list. You will be assigned a priority number on the next page. The 2 players with the **lowest** priority numbers (1 is the lowest) who chose not to settle are audited and forced to pay 100.0. Others who chose not to settle 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. **With 20% chance, the button never becomes clickable**.

Information: You will receive information on the share of other players who have decided to accept settlement.

Other Players: All other players are offered the same settlement after a random amount of time (this time IS NOT the same for all players). Others' buttons have a 20% chance to remain un-clickable. If your button becomes clickable, this IS NOT informative about whether or not others' buttons have become clickable.

Treatment Snapshots:

Below, we produce two snapshots of the screens you will see during the round. Red text is commentary that won't appear in the round.

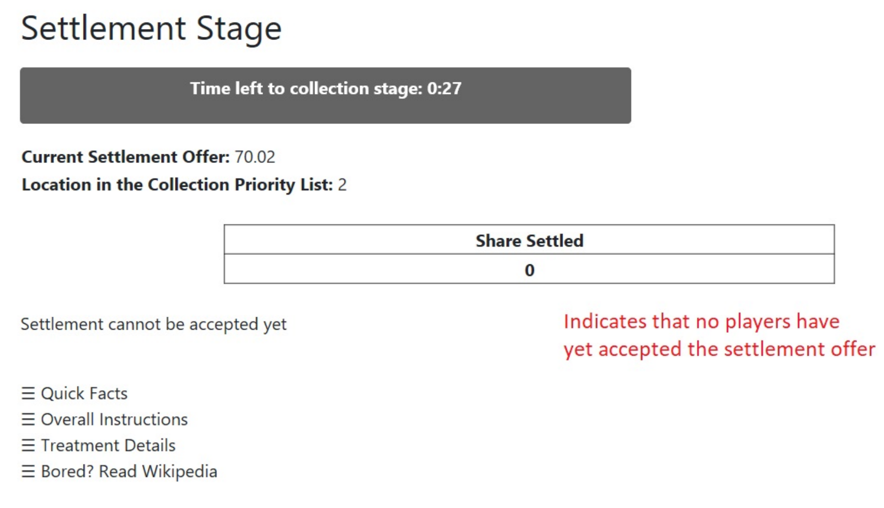

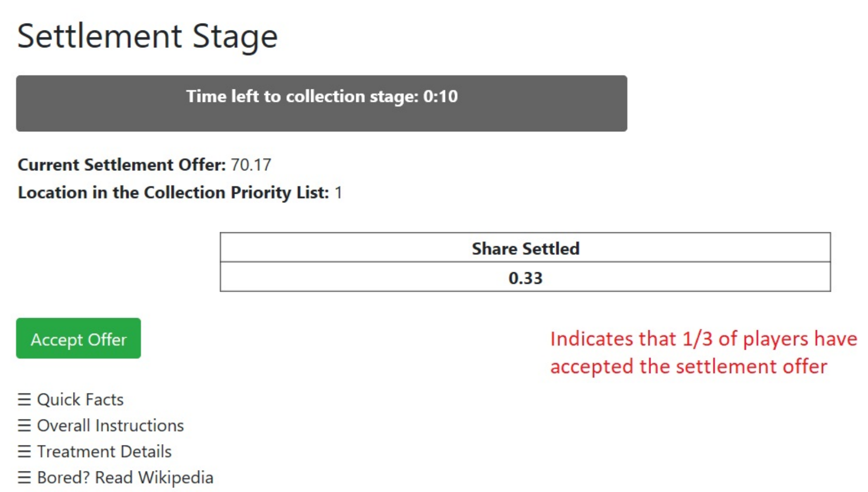

References

Abreu, D. and H. Matsushima (1992): "Virtual implementation in iteratively undominated strategies: complete information," *Econometrica: Journal of the Econometric Society*, 993–1008.

Abreu, D. and A. Sen (1991): "Virtual implementation in Nash equilibrium," *Econometrica: Journal of the Econometric Society*, 997–1021.

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.

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.

Halac, M., I. Kremer, and E. Winter (2018): "Raising capital from heterogeneous investors," .

Jackson, M. O. (1992): "Implementation in undominated strategies: A look at bounded mechanisms," *The Review of Economic Studies*, 59, 757–775.

Kennedy, D. M. (2011): *Don't shoot: one man, a street fellowship, and the end of violence in inner-city America*, Bloomsbury Publishing USA.

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.