



Preordered service in contract enforcement [☆]

Jan U. Auerbach ^{a,1}, Miguel A. Fonseca ^{a,b,*,1}

^a University of Exeter, United Kingdom

^b NIPE, Universidade do Minho, Portugal



ARTICLE INFO

Article history:

Received 2 June 2018

Available online 22 April 2020

JEL classification:

K00

K12

K40

O17

C92

Keywords:

Judicial system

Courts

Judiciary performance

Legal procedure

Civil cases

Caseload

Contract enforcement

Population of investment games

Unraveling

Experiments

ABSTRACT

To address delay and backlog at civil courts, we propose a procedural rule that we refer to as preordered service to replace sequential service of low-profile cases for breach of contract. Courts preannounce a list that uses uniquely identifying information to rank potential low-profile contracts, like a combination of contracting parties' taxpayer numbers. They use this list to schedule initial hearings of filed low-profile contract cases in that order. In theory, unlike sequential service, preordered service ensures efficiency in a population of investment games through unraveling. Results from a laboratory experiment suggest that it may substantially reduce court caseloads.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Many countries' civil courts experience long delay and large case backlog (e.g., Agor et al. 2015; Palumbo et al. 2013). The time and resources it takes to enforce contracts may discourage investment into profitable business and impede economic prosperity. Slower courts and larger case backlog have been linked to more breaches of contract and less investment (Chemin

[☆] We thank Tim Miller for his outstanding help in programming the experimental software. We also thank two anonymous reviewers, Martin Brown, Surajeet Chakravarty, Cecilia Chen, Ramón Cobo-Reyes, Martin Dufwenberg, Nuno Garoupa, Brit Grosskopf, Yoske Igarashi, Christos Kotsogiannis, Simone Meraglia, Quyen Nguyen, Charles Noussair, Kim Peters, Robert Sugden, as well as participants at the Institutions and Cooperation Workshop at U. of Exeter, the 2017 SAET Conference in Faro, the Midwest Economic Theory Conference at Drexel U., the North American Summer Meeting of the Econometric Society at U.C. Davis, the International Conference on Game Theory at Stony Brook U., the Annual Congress of the European Economic Association at U. of Cologne, and the 2019 Utah Experimental Economics Conference; and seminar participants at U. Catolica Porto, U. of East Anglia, U. of Exeter, U. of Vienna, Tsinghua U., Utah State U., and Birkbeck College for their valuable comments and suggestions. Financial support from ESRC grant ES/N00762X/1 and the U. of Exeter Business School BIP cluster is gratefully acknowledged. Errors are ours.

* Corresponding author.

E-mail addresses: j.auerbach@exeter.ac.uk (J.U. Auerbach), m.a.fonseca@exeter.ac.uk (M.A. Fonseca).

¹ Postal address: University of Exeter Business School, Streatham Court, Rennes Drive, Exeter EX4 4ST, United Kingdom.

2012), less lending and tighter credit constraints (Jappelli et al. 2005), higher firm financing costs and smaller firms (Fabbrri 2010). So, what reforms can help speed up the courts?

Among other things, some have suggested that fewer cases for the courts to look at can help reduce the time it takes to resolve civil disputes (see, e.g., Palumbo et al. 2013). To this end, we propose a procedural rule that aims at reducing the number of court case filings associated with a category of contracts of particularly low profile. The implied reduction in the overall caseload should speed up the courts. As the empirical evidence suggests, faster courts in turn are likely to encourage more investment into profitable business of higher profile—business that is more economically significant and legally complex—leading to more lending, more investment, and more firm growth.

To reduce the caseload, we focus on low-profile cases for breach of contract with a sure outcome in court, which make up the bulk of the caseload at civil courts. For example, in some U.S. state courts in 2012–13, almost two-thirds of all non-domestic civil cases were contract disputes; when recorded, damage awards in those disputes averaged at less than \$10,000.² These cases mostly consist of debt collections, landlord-tenant disputes, and foreclosures.³ They are brought to secure a sure ruling, which then allows the claimant to initiate the legal enforcement of a payment (see Agor et al. 2015, p. 35).

With this focus, we suggest to induce unraveling effects of the kind that have been described to be at work in many contexts in law and economics (e.g., Bar-Gill and Ben-Shahar 2009; Dari-Mattiacci and De Geest 2010). Specifically, we propose to preorder the service of such low-profile contract cases at the courts, replacing sequential service, which processes them in order of arrival. If nobody wants to be first in line at the courts, then unraveling reduces the number of such low-profile contract cases being filed. As a consequence, more resources can be allocated to the speedier resolution of higher-profile disputes. The reliance of preordered service on unraveling may undermine its performance in practice (e.g., Nagel 1995; McKelvey and Palfrey 1992). Could it have an effect? We provide proof of concept in the context of a population of low-profile contracts, in theory and in an experiment.

For the sake of exposition, suppose that all contracts are low-profile and have a sure outcome in court if breached. Suppose further that every entity can be party to at most one contract so that every contract can be uniquely identified using any of the parties to it. Preordered service then requires the judiciary to list all entities that may enter contracts in some order, using some uniquely identifying information, like taxpayer numbers. It makes publicly known that it will serve court cases for breach of contract filed within an *ex ante* specified period of time in the order of that list, using the listed entities to identify contracts and thus cases. Courts then collect all cases filed during the specified time period. At the end of it, they order the collected cases according to the list and serve the first ones in line first.⁴

We present preordered service of such low-profile contract cases at the courts in a stylized economy. We interpret this economy as an extreme version of a richer environment with two features that preordered service uses. First, every contract can be associated with a uniquely identifying label. Second, the prospect of a potential court case being served first, rather than somewhere down the line, tips agents over from breaching the contract to honoring it.

Our stylized economy consists of many investors and many entrepreneurs. While entrepreneurs have productive projects, investors are endowed with the capital required to implement them. Investors and entrepreneurs are randomly matched with one another and can enter a contract to use the investor's capital in the entrepreneur's project. After production, however, the entrepreneur can breach the contract and keep all the output, in which case a claim is filed with the court. For simplicity, we assume that the court can serve exactly one claim, in which case the respective breaching party incurs a cost; all other claims go unenforced. This simplification is an extreme version of costs and gains to the respective parties to a breached contract that are associated with long court delays.

In this environment, we represent sequential service by a random draw of one breached contract for service. This assumption captures the idea that civil cases queue randomly when filed, and only the first in line is served. Modeled in this way, sequential service can lead to inefficiencies: some investors do not invest because all contracts are breached, but not all are enforced. In contrast, we model preordered service using an *ad hoc* label to represent some uniquely identifying information relating to the contract or the parties to it. This way we capture the idea that, to do its job the best it can, the judiciary could in principle make use of the identities of the parties to a dispute. The ranking of potential contracts by these labels is announced before they are signed. The judiciary then serves the highest-ranked contract that is breached. In this environment, preordered service achieves efficiency. The intuition is that no entrepreneur wants to be associated with the highest-ranked contract that is breached. Thus, in equilibrium, no entrepreneur breaches the contract and all investors safely enter contracts with entrepreneurs, which maximizes aggregate production and consumption.

This theoretical success of preordered service builds on iterated elimination of dominated strategies and backward induction. Experimental evidence suggests that this fact may undermine its effectiveness in practice (e.g., Nagel 1995, McKelvey and Palfrey 1992, respectively), which we cannot examine in the field as preordered service is counterfactual. Therefore, we test its performance in a laboratory experiment. We implement our environment and undertake *ceteris paribus* comparisons

² According to Agor et al. (2015), 64% of all (925,344—about 5% of U.S. state civil cases) non-domestic civil cases disposed between July 1, 2012 and June 30, 2013 in 152 courts in 10 urban counties were contract cases. Despite fluctuations, this share tends to be above 50% over time (p. 36). The mean of recorded damage awards was \$9,428; award values at the 25th, 50th, and 75th percentiles were \$1,251, \$2,272, and \$4,981, respectively.

³ Agor et al. (2015) report that debt collections, landlord-tenant disputes, and foreclosures make up 37%, 29%, and 17% of the contract cases in their dataset, respectively.

⁴ In Section 5 we explain, among other things, how preordered service deals with many contracts per entity.

of individual and aggregate behavior with preordered and sequential service. We find that, over the course of the experiment, the court caseload is more than 40 percent lower with preordered service than it is with sequential service. In the last round of the experiment, preordered service results in half as many court cases as sequential service. While investment levels do not differ across both rules, with preordered service, as fewer contracts are breached, investors on average secure a higher payoff.

We discuss related literature in Section 2 and present our stylized environment and its predictions in Section 3. In Section 4, we report our experimental design, hypotheses, and empirical findings. Departing from the simplified exposition we adopt for most of the paper, we discuss how preordered service could be taken to practice in Section 5. We conclude in Section 6.

2. Relation to the literature

We focus on a large inflow of civil cases as one source of court delay and case backlog in the judicial system. Our procedural rule, preordered service, reduces the case inflow and thus the workload facing the judiciary. A related source of delay and backlog is the case flow time—the time it takes to conclude a case. Building on insights of Coviello et al. (2014, 2015) into task juggling, Bray et al. (2016) take an operations management perspective to reduce case flow times in a field experiment in an Italian court. They focus on scheduling the sequence of hearings a case requires, when many cases are in progress at the same time, taking the inflow of cases as given. By contrast, our focus is on scheduling initial hearings—i.e., the order in which cases are opened—to affect the inflow of cases, taking case flow times as given.

We model contractual relationships with an option to breach similar to the underlying game in Bohnet et al. (2001), who modify the original investment game introduced by Berg et al. (1995). We use a population of such investment relationships and think of it as a stylized economy with aggregate output. When service is sequential, the probability of any one breached contract to be enforced—i.e., to arrive first in line—is endogenous. Of course, the patterns of economic interactions are more complex than we model them. Basu et al. (2009) let many investors and many entrepreneurs have multiple investment relationships at the same time to study record keeping and trust. Preordered service can incorporate multiple investment relationships in a straightforward fashion. For example, contracts can be ranked lexicographically, where the first and second coordinates of the label are the respective taxpayer numbers of the associated investor and entrepreneur. This approach also works with financial intermediaries.

Our focus is on low-profile civil cases for breach of contract for which there is no uncertainty about the outcome in court. However, preordered service could potentially be adopted whenever accurate reports of violations can be generated at low enough costs. In terms of Mookherjee and Png (1992)'s analysis of costly monitoring versus investigation of violations of law or regulation, we focus on the investigation of violations, which in our setup are reported with probability one. Since reports are accurate and investigation costs are low in our environment, Mookherjee and Png (1992) would in fact suggest to use only investigation. In their terminology, we describe and compare preordered and sequential investigation.

Bar-Gill and Ben-Shahar (2009) and Dari-Mattiacci and De Geest (2010) study similar strategic interactions. Bar-Gill and Ben-Shahar (2009) use an argument based on unraveling to explain the prevalence of plea bargains in simultaneous criminal proceedings that involve a resource-constrained prosecutor handling many cases. The threat by a resource-constrained prosecutor to take to court all those who do not accept the plea offer should not be credible. However, common knowledge of the priorities according to which the prosecutor selects who to prosecute among those who reject the plea offer gives rise to a similar unraveling as in our case, so that all plea offers are accepted in equilibrium. More generally, studying a multiplication effect inherent to the threat of punishment but not to the promise of reward, Dari-Mattiacci and De Geest (2010) show that a similar unraveling argument allows a principal to enforce compliance of many agents in a setting with monitoring. They use their insights to explain a number of observations in law and economics. The nature of our contribution is quite different. Rather than using an unraveling argument to explain an observation, we are proposing a hitherto untested procedural rule based on a similar logic to address delay and backlog at civil courts. While preordered service builds on similar unraveling arguments, our environment focuses on enforcing low-profile contracts, which largely elude monitoring and make up the bulk of civil court caseloads. Assuming that claims are filed when a contract is breached, we show that preordered service has bite across a population of otherwise strategically independent contracting games, both in theory and in the lab.

Finally, in the literature on public goods games, Andreoni and Gee (2012, 2015) and Kamijo et al. (2014) study centralized punishment institutions to sanction non-cooperative behavior that induce a similar strategic interaction. The institutions in these papers punish the least cooperative player. Thus, in equilibrium, all players contribute just enough to the public good in order not to be the least cooperative player. In such strategic games with small self-governing groups, the authors argue, this sort of centralized punishment institution comes quite naturally. However, it requires information about the contributions of all players in order to rank them and calculate the fines. By contrast, while preordered service might not come quite as naturally—and to some may even appear unfair at first—it requires only little information: it relies solely on a report being filed if and when an individual or an entity has violated a convention, a law, a regulation, or a contract. A notification that a violation occurred is enough. In principle, while it can be used, the degree of the violation is irrelevant. Hence, as preordered service can make use of existing uniquely identifying information, like taxpayer numbers, the additional administrative costs of introducing it are likely manageable.

3. The model and predictions

In this section, we discuss a stylized environment with many contracts in which (i) every contract is associated with a uniquely identifying label, and (ii) the prospect of a potential court case being served first, rather than somewhere down the line, tips agents over from breaching the contract to honoring it. We describe the environment in Section 3.1, present its predictions in Sections 3.2–3.4, and provide a discussion of our assumptions, interpretations, and predictions as well as an extension in Section 3.5. All proofs are in Appendix A.

3.1. The environment

There are $2n$ risk neutral agents, n investors and n entrepreneurs. Investors are endowed with one unit of productive capital each that they can either rent out or convert into one unit of the consumption good. Entrepreneurs are endowed with one unit of time each that they can allocate indivisibly to one of two uses. They can either produce one unit of the consumption good or implement a project that requires one unit of capital. The project is risk-free and converts one unit of capital into $2w$ units of the consumption good, where $n > w \geq 2$.

Investors and entrepreneurs are randomly paired. Before any interaction, the n pairs are assigned distinct labels z_1, z_2, \dots, z_n . These labels could be associated with only one party to the pair, like one of their taxpayer numbers. All agents know the label of their pair. The investor in a pair moves first and decides whether or not to enter a contract with the entrepreneur in that pair. All investors move simultaneously. If the investor in a pair declines, then both parties separately produce one unit of the consumption good, implying payoff 1 for both the investor and the entrepreneur. If the investor enters the contract, then the entrepreneur implements the project and produces $2w$ units of the consumption good. The terms of the contract are fixed exogenously to sharing the project's output equally.

Then, all entrepreneurs move simultaneously and decide whether or not to honor the contract, if they have one. An entrepreneur can breach the contract by keeping all the output. In this case, a claim is filed with the judiciary. The order of arrival of these claims at the court is random. The judiciary then selects exactly one such claim for enforcement according to one of two possible procedural rules: either (i) sequential service or (ii) preordered service.

If service is sequential, then the first claim in line for enforcement is served. That is, as the order of arrival is random, one claim is randomly selected for enforcement, with equal probability for each claim. Letting $x \geq 0$ be the number of all claims filed with the judiciary, if $x \geq 1$, then the probability of any individual claim to be enforced is $p = 1/x$.⁵

If service is preordered, then at the outset, the judiciary generates and announces an ordered list of the labels z_1, z_2, \dots, z_n . Every claim is associated with a breached contract, the parties to which belong to a pair identified by one such label. The judiciary then uses the ordered list of labels to order the claims and enforces the first claim in line. The ordered list is a finite sequence generated by some bijection $f : \{1, 2, \dots, n\} \rightarrow \{z_1, z_2, \dots, z_n\}$ associating positions $1, 2, \dots, n$ in the line at the court with the pairs via their labels z_1, z_2, \dots, z_n . Without loss of generality, we assume that $f(i) = z_i$ for all $i \in \{1, 2, \dots, n\}$. We can therefore represent the ordered list by the n -tuple (z_1, z_2, \dots, z_n) , and the label index $i \in I = \{1, 2, \dots, n\}$ directly indicates pair z_i 's position in the line.

If a claim is selected for enforcement, then it is enforced with certainty. In this case, the pair's investor is made whole by a payment w from the pair's entrepreneur, who in addition incurs a cost $c \in (0, w)$. This cost makes breach of contract costly upon enforcement and can be interpreted as capturing many kinds of costs, like, e.g., attorney fees or reputation costs. Therefore, while the entrepreneur's payoff from honoring the contract if they have one is w , their expected payoff from not honoring it is

$$p(w - c) + (1 - p)2w = 2w - p(w + c),$$

where p is the probability of the contract being enforced, which is determined by the procedural rule in interaction with the breach decisions of entrepreneurs who have a contract. The investor's payoff from entering the contract is w if it is honored. Conditional on the contract not being honored, the investor's expected payoff from entering it is

$$pw + (1 - p) \cdot 0 = pw.$$

The game as described is common knowledge among all agents. In autarky, all agents produce separately and aggregate production and consumption equal $2 \cdot n$. In the efficient outcome, all contracts are entered and honored: no agent incurs the cost c and aggregate consumption equals maximum aggregate output $2 \cdot n \cdot w$. Fig. 1 depicts the game tree for the interaction between one investor and one entrepreneur. The triplet (n, w, c) summarizes the economic fundamentals. We focus on pure-strategy Nash equilibria.

⁵ The probability of an individual case to be served (i.e., to arrive first in line) equals the number of all permutations of the labels associated with claims filed that have that case listed first, $(x - 1)!$, divided by the number of all permutations of those labels, $x!$.

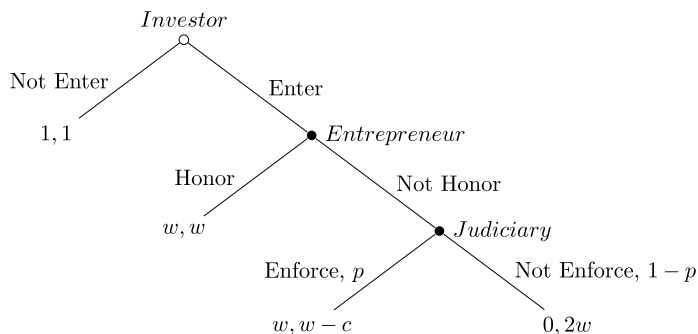


Fig. 1. The game between one investor and one entrepreneur.

3.2. No service

As a benchmark, without enforcement, the entrepreneur can keep all the output, knowing that the contract will not be enforced. Thus, no contract is entered and no project is implemented.

Proposition 1 (NONE). *With no service, the unique equilibrium is autarky.*

However, as soon as one claim is enforced, autarky is not an equilibrium anymore. The contract of an investor who deviates to entering is either honored or enforced.

3.3. Sequential service

With sequential service, a claim's position in the queue for service is random, and only the first claim in line is enforced. This procedural rule can lead to inefficient outcomes.

Proposition 2 (SEQSERVE). *With sequential service, there is an efficient equilibrium as well as inefficient equilibria in which $\lfloor w \rfloor < n$ contracts are entered and not honored.*

The efficient outcome is an equilibrium because enforcement and the associated cost are sure if only one entrepreneur were to decide to breach their contract. However, sequential service also allows for inefficient equilibria. In these equilibria, investment is inefficiently low and all entrepreneurs that have a contract breach it. If enough entrepreneurs breach their contract, then enforcement and the associated cost become unlikely enough for all entrepreneurs that have a contract to find it optimal to breach it. At the same time, enforcement is still likely enough to motivate some investors to enter the contract.

3.4. Preordered service

With preordered service, the order of claims is preannounced, and the first claim in line is enforced. This procedural rule induces unraveling effects that lead to the efficient outcome.

Proposition 3 (PREORDER). *With preordered service, the unique equilibrium is efficient.*

Using a preannounced ordered list to line up breached contracts and enforce the first in line induces full compliance; enforcement is not needed. The intuition is as follows. The entrepreneur whose contract is listed first does not breach it because they face enforcement and the associated cost with certainty. The entrepreneur whose contract is listed second understands this reasoning and concludes that upon breach, they face enforcement and the associated cost with certainty. As this entrepreneur thus does not breach the contract, the same argument holds for the entrepreneur whose contract is listed next—and so on. It follows that all contracts are honored. Therefore, all contracts are entered and aggregate production and consumption are maximized. Compared to some outcomes with sequential service, preordered service reduces the case inflow (and thus the caseload) facing the courts to zero.

3.5. Discussion and an extension

In this section, we provide a discussion of a number of aspects that might be of interest. In particular, we discuss our parameter choices in Section 3.5.1, heterogenous project productivity in Section 3.5.2, our assumptions on the nature of the

contracts in Section 3.5.3, and our assumptions on the judicial process in Section 3.5.4. Finally, we extend the environment to allow for mistakes and exogenous breach of contracts in Section 3.5.5. We then show under what conditions preordered service improves the outcome compared to sequential service.

3.5.1. Parameter choices

Regarding our parameter choices, it can be verified that with $w < 2$, there is a unique equilibrium with sequential service, and it is efficient. Similarly, it can be verified that with $w \geq n \geq 2$, in the inefficient equilibrium with sequential service, while no contract is honored, all contracts are entered and thus aggregate production is maximized. This equilibrium is inefficient because one claim is enforced and the associated cost is incurred by some agent, so that aggregate consumption is not maximized. With a focus on the case of many contracts, our parameterization $n > w \geq 2$ rules out these uninteresting outcomes.

Regarding $c \in (0, w)$, the assumption that $c > 0$ generates a meaningful trade-off between honoring and not honoring the contract for entrepreneurs who have one. If there was no cost to the entrepreneur associated with the enforcement of a breached contract, then not honoring the contract always guarantees at least as high a payoff as honoring the contract. As to the assumption that $c < w$, preordered service induces the unraveling effects we describe irrespective of how high c is as long as $c > 0$. When service is sequential, a high enough c could guarantee that no contract is ever breached (e.g., $c > (n - 1)w$ would suffice). However, the observation that low-profile contract cases clog up the civil courts suggests that c is not high enough in this sense. To the extent that c cannot be arbitrarily high for there to be contracts that are entered and breached when service is sequential, $c < w$ is a convenient restriction. In certain situations, this restriction can be interpreted as limited liability with respect to certain costs associated with the court case, but not related to the investor's compensation. Of course, the cost c could be another lever policy makers could use to reduce the caseload at courts, one that we do not focus on here. An example would be to impose extraordinarily high fines or even time in prison on the breaching party upon enforcement, in addition to compensation to the investor. However, with such a policy, concerns about the punishment fitting the offense would almost certainly arise.

3.5.2. Heterogeneous projects

The nature of our predictions is largely unchanged when allowing for heterogeneous projects. All else the same, let there be two types of projects, $n_L \geq 1$ projects with low productivity $w_L > 1$ and $n_H \geq 2$ projects with high productivity $w_H \geq 2$, $\lfloor w_H \rfloor > \lfloor w_L \rfloor$. Assume that $n_L + n_H > w_L > c > 0$. With one claim being enforced, the efficient outcome is always an equilibrium. The breached contract of an entrepreneur who deviates to not honoring their contract is the only one in line for service and thus enforced with certainty. With preordered service on the one hand, this equilibrium is the unique equilibrium because a project's productivity does not interfere with the unraveling effect at work.

With sequential service on the other hand, an inefficient equilibrium exists. If $n_H > w_H$, then any profile in which exactly $\lfloor w_H \rfloor$ investors enter a contract with an entrepreneur who has a project with productivity w_H , no other investor enters a contract, and all entrepreneurs breach their contract if they have one is an equilibrium. If $n_H \leq w_H$, then there are two cases. If $w_L < n_H + 1$, then the profile in which exactly n_H investors enter a contract with an entrepreneur who has a project with productivity w_H , no other investor enters a contract, and all entrepreneurs breach their contract if they have one is an equilibrium. If $w_L \geq n_H + 1$, then any profile in which n_H investors enter a contract with an entrepreneur who has a project with productivity w_H , $\lfloor w_L \rfloor - n_H$ investors enter a contract with an entrepreneur who has a project with productivity w_L , no other investor enters a contract, and all entrepreneurs breach their contract if they have one is an equilibrium. In all cases, production is inefficiently low and all contracts are breached so that one contract is enforced.

3.5.3. Contracts, sure outcomes in court, and settlement

We focus on low-profile contract cases with a sure outcome, which constitute the bulk of the caseload at civil courts. In the interest of simplicity, we assume that claims are either enforced in this period, or not all. To the extent that preordered and sequential service determine the order in which arriving court cases are served, this all-or-nothing assumption can be interpreted as an extreme version of discounting the benefits and costs of a breached contract being enforced at some point in the future. The exact terms of the contract, which we fix exogenously, are not important, as long as there is a meaningful trade-off between breaching the contract and honoring it. Allowing agents to enter multiple contracts would not be a problem as long as every potential contract is associated with a uniquely identifying label. With the same qualification, allowing multiple parties in the same contract would not be a problem either.

In line with our focus on low-profile contract cases with a sure outcome, we assume that there is no uncertainty over the court's ruling. Risk aversion thus plays no role when service is preordered. If a contract tops the preordered list, then enforcement and the associated cost are sure upon breach. If sequential service allows for an inefficient equilibrium, preordered service compares favorably to it. If the court's ruling were uncertain, then there might be a role for risk aversion. However, with sequential service, a probability of enforcement upon service combines with the probability of arriving first in line, and thus of service. By contrast, for the first contract on the preordered list, the probability of a potential court case being served is one. Therefore, the probability of facing enforcement and the associated cost is higher when service is preordered than when it is sequential. Given a probability of enforcement upon service and an associated cost, the theoretical predictions then remain unaffected for certain levels of risk aversion. However, recall that Agor et al. (2015,

p. 35) suggest that most contract cases in their dataset are brought only to secure a sure ruling which then allows the claimant to initiate the legal enforcement of a payment.

For these low-profile contract cases with a sure outcome we focus on, settlement after filing but before a ruling is issued is not very important.⁶ However, the probability p of enforcement in our environment can be interpreted as the probability of the case being disposed by the court (in this period), be it by ruling or by settlement. The payoffs upon enforcement (settlement) could easily be altered to account for a potential haircut associated with settlement, or disposal of the case more generally, without affecting the strategic interaction or the predictions. Settlements of claims filed with the courts do suffer from court delay when they come about only after initial or even multiple hearings. With regards to the implementation of preordered service, when more than one case is served in a period, settled cases simply vacate a spot in the queue. If settlement currently takes place with a large haircut because the delay until court service is too long, then preordered service should lead to fewer settlements of filed claims as court delay decreases.

3.5.4. The judiciary, compensation, and false claims

We assume that the intent of courts and their agents to enforce contracts does not compete with other objectives, such as political values (see Landes and Posner 2009). More generally, of course, judges are subject to the incentives posed by the environment they work in (see Posner 2005). The possibility of convergence to efficient law has been discussed by Gennaioli and Shleifer (2007) and studied empirically by Niblett et al. (2010).

We assume that upon enforcement of the contract, the investor is fully compensated. We do not take a stand on what form this compensation takes. Schwartz (1979) and Shavell (1980, 1984) discuss damages and specific performance as remedies for breach of contract.

Finally, we abstract from intentionally false claims by investors. Suppose that we allowed for such fraudulent claims in our context without uncertainty over the outcome in court. If investors whose claims are found to be unsubstantiated face some sort of cost that they would not otherwise incur, e.g., reputation cost or attorney fees, then preordered service would prevent these claims from being filed: either the entrepreneur or the investor associated with the first claim on the list faces a sure cost and would thus prefer that the claim did not exist.

3.5.5. An extension: mistakes and exogenous breach

In this section, we allow for contracts to be breached by an exogenous shock. Possible interpretations are a liquidity shock or a randomly occurring mistake by the entrepreneur to the effect that they breach the contract irrespective of whether or not it is optimal to do so. Suppose that after production and collection of their projects' proceeds, entrepreneurs who have a contract are subject to an exogenous shock that is independent and identically distributed across entrepreneurs. This shock breaches the entrepreneur's contract exogenously with probability $1 - \alpha \in (0, 1)$. In this case, a claim is filed with the courts. Each entrepreneur has a choice whether or not to honor their contract with probability $\alpha \in (0, 1)$.⁷ All claims arrive at the court at the same time and in random order, irrespective of whether the contract was breached exogenously or deliberately. The judiciary cannot distinguish between the nature of the breach and thus treats all claims equally. The possibility of exogenous breach due to this random shock as well as the otherwise unchanged structure of the game more generally are common knowledge among all players. The modified game between one investor and one entrepreneur is depicted in Fig. 2. As before, we focus on pure-strategy Nash equilibria.

For the case of sequential service, we first note that in equilibrium, either all entered contracts are breached with certainty, or all breached contracts are breached exogenously.

Lemma 1. *In equilibrium with sequential service, either all entrepreneurs in a contract breach it if they have a choice, or all entrepreneurs in a contract honor it if they have a choice.*

Lemma 1 states that every equilibrium has to be of one of two types: either (1) all entrepreneurs in a contract honor it if they have a choice, or (2) all entrepreneurs in a contract breach it if they have a choice. In the former type of equilibrium, in which all entrepreneurs in a contract honor it if they have a choice, all breached contracts are necessarily breached exogenously. In this sense, the caseload at the courts is determined exogenously. There is no point in attempting to reduce the caseload at the courts using policy interventions that aim at affecting entrepreneurs' breaching decisions. Entrepreneurs already honor their contracts if they have a choice. Thus, this type of equilibrium is of limited interest for our purposes. For completeness, we report conditions an equilibrium of this type has to satisfy, and that for a low enough probability of exogenous breach, an equilibrium exists in which all contracts are entered and honored if there is a choice.

Proposition 4. *With sequential service, a strategy profile in which exactly $\hat{k} > 1$, $\hat{k} < n$, investors enter the contract and all entrepreneurs honor the contract if they have one and have a choice is an equilibrium if and only if \hat{k} satisfies*

⁶ Agor et al. (2015) report that about 7% of the contract cases in their data were settled; some settled contract cases might have been reported as dismissed.

⁷ While entrepreneurs do not have a choice whether or not to enter the contract once the investor enters it, the probability of exogenous breach should be consistent with them choosing to enter the contract when they would honor it if they had a choice. A sufficient condition is $\alpha w + (1 - \alpha)(w - c) \geq 1$ or $(1 - \alpha) \leq (w - 1)/c$.

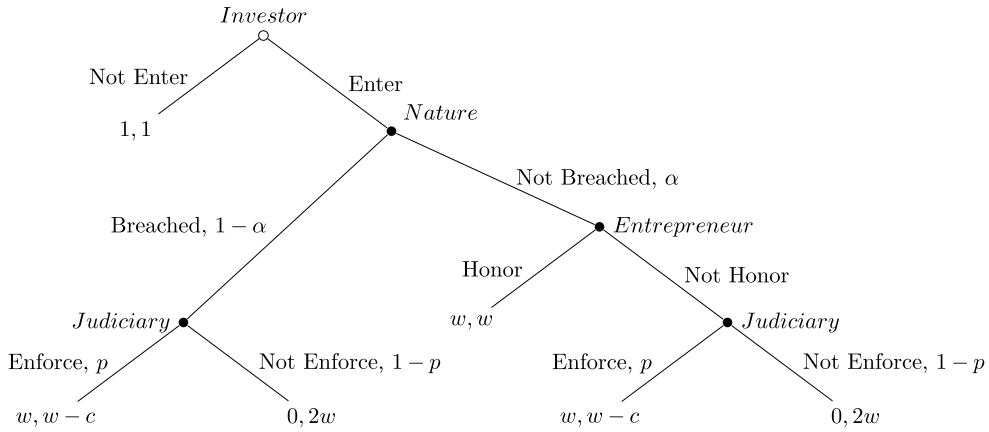


Fig. 2. The game between one investor and one entrepreneur with exogenous breach.

$$\frac{1 - \alpha^{\hat{k}}}{(1 - \alpha)\hat{k}} \geq \frac{w}{w + c} \quad \text{and} \quad \frac{1}{w} \geq \alpha + \frac{1 - \alpha^{\hat{k}+1}}{\hat{k} + 1}.$$

For high enough α , the strategy profile in which all investors enter the contract and all entrepreneurs honor the contract if they have one and have a choice is an equilibrium.

By contrast, in the second type of equilibrium, in which all entrepreneurs in a contract breach it if they have a choice, all contracts that are not breached exogenously are breached endogenously. In this case, potentially, there is a role for policy interventions that aim at affecting entrepreneurs' breaching decisions to reduce the caseload at courts. In this sense, this type of equilibrium is of more interest for our purposes. We have the following result.

Proposition 5. *With sequential service, every strategy profile in which exactly $\lfloor w \rfloor$ investors enter the contract and all entrepreneurs breach the contract if they have one and have a choice is an equilibrium. There is no such equilibrium in which more than $\lfloor w \rfloor$ contracts are entered.*

In the relevant equilibrium with sequential service, as in the case without the possibility of exogenous breach, all entered contracts are breached with certainty. Therefore, the possibility of exogenous breach does not affect investors' entering decisions and thus the overall equilibrium outcome or its intuition. When service is preordered, we have the following result.

Proposition 6. *With preordered service, an equilibrium exists. Define*

$$k^* := \left\lceil \frac{\log\left(\frac{w}{w+c}\right)}{\log \alpha} + 1 \right\rceil. \tag{1}$$

If $k^ \geq n$, then in every equilibrium, the investors in all n pairs enter the contract and all entrepreneurs honor it if they have a choice, except possibly the one in pair z_n .*

If $k^ < n$, then in every equilibrium, either k^* or $k^* + 1$ contracts are entered and at most one entrepreneur with a contract breaches it if they have a choice.*

The intuition is similar to before and again derives from unraveling. However, due to the possibility of exogenous breach, the unraveling effect may extend only to some number k^* or $k^* + 1$ of the highest-ranked pairs. The exact number of entered contracts depends on parameters, but for large enough n it is at least k^* . In particular, if the probability of exogenous breach is small enough (for a large enough α), then this economy enters more contracts, hence produces more output, and thus experiences higher welfare with preordered service than with sequential service. Proposition 7 provides a sufficient condition.

Proposition 7. *If $\alpha \geq (w/(w + c))^{1/(w-1)}$, then more contracts are entered with preordered service than with sequential service.*

The condition in Proposition 7 is sufficient but not necessary. Fig. 3 depicts the number of contracts entered (vertical axis) as a function of α (horizontal axis) for both preordered service, which is assumed to be k^* (rather than $k^* + 1$, assuming that n is large), and sequential service, which is $\lfloor w \rfloor$. The top-left panel uses the parameters we use in our experiment. The

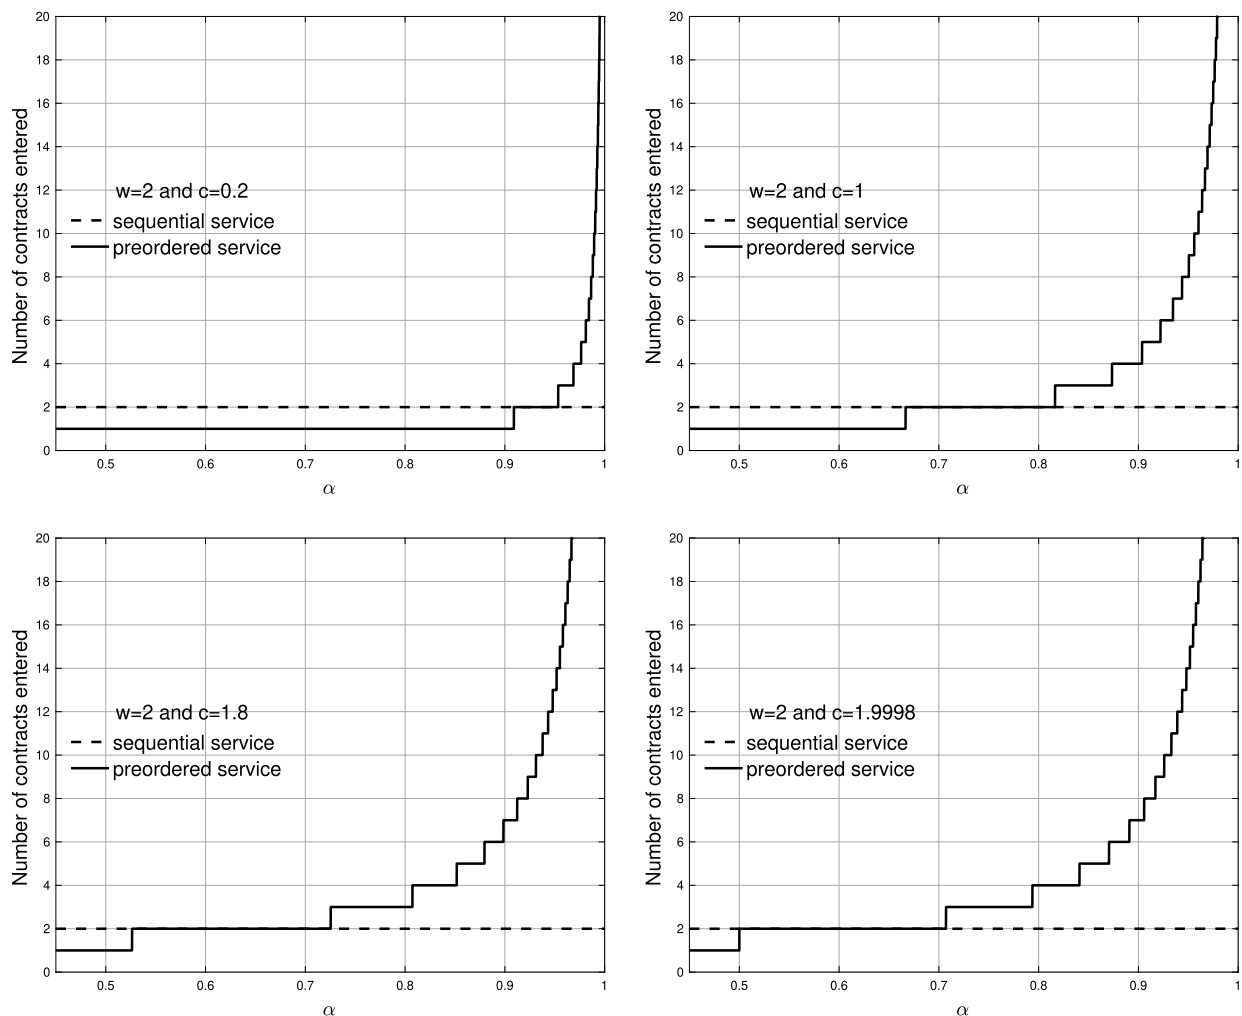


Fig. 3. Number of contracts entered for $w = 2$ and various costs c .

top-right and bottom-left panels vary the cost c incurred by the entrepreneur upon enforcement of a breached contract in a symmetric fashion, from 10 percent of w in the top-left panel to 50 percent of w in the top-right panel to 90 percent of w in the bottom-left panel. The bottom-right panel considers the case of very high costs c of 99.99 percent of w . As can be seen in all panels, for high enough $\alpha < 1$, more contracts are entered with preordered service than with sequential service. How high α has to be depends on the parameters.

To summarize, more contracts are entered with preordered service than with sequential service for low enough but strictly positive probabilities $1 - \alpha$ of exogenous breach. In fact, depending on the parameters, the probability of exogenous breach can be fairly high and still more contracts are being entered with preordered service than with sequential service. In this sense, the predictions concerning overall investment for the environment described in Section 3.1 appear to be robust to the addition of a small chance of exogenous breach via an independent and identically distributed shock.

What can we say about the caseload at the courts? With sequential service, by Proposition 5, the caseload is $\lfloor w \rfloor$, as all entered contracts are breached. By Proposition 6, with preordered service, either k^* , $k^* + 1$, or n contracts are entered, and at most one contract is breached if there is a choice. It follows that the expected caseload at the courts is lower with preordered service than with sequential service as long as not too many more contracts are entered with preordered service than with sequential service. Proposition 8 provides a sufficient condition.

Proposition 8. Letting \hat{k} be the number of contracts entered with preordered service, the expected caseload at courts is lower than with sequential service if $\hat{k} \leq (\lfloor w \rfloor - \alpha)/(1 - \alpha)$.

The condition in Proposition 8 is sufficient but not necessary. If at most as many contracts are entered with preordered service as with sequential service, then the expected caseload is smaller with preordered service because most or all con-

Table 1
Experimental design.

	NONE	SEQSERVE	PREORDER
# subjects	96	96	96
# economies	6	6	6

tracts are honored if there is a choice while with sequential service, all contracts are breached with certainty. For the same reason, as $\lfloor w \rfloor \geq 2$ implies that $(\lfloor w \rfloor - \alpha)/(1 - \alpha) > \lfloor w \rfloor$, if more contracts are entered with preordered service than with sequential service, but not too many more, then the expected caseload is still lower with preordered service than with sequential service. If very many more contracts are entered with preordered service than with sequential service, then the expected caseload may be higher with preordered service than with sequential service—for the purely mechanical reason that so many more contracts are entered and each one of them faces an independent chance of being breached exogenously.

To summarize, the expected caseload is lower with preordered service than with sequential service unless there are too many more contracts entered with preordered service than with sequential service. In this sense, the predictions concerning the caseload at the courts for the environment described in Section 3.1 appear to be robust to the addition of a small chance of exogenous breach via an independent and identically distributed shock as well.

4. Evidence

For preordered service to be effective, it requires individuals to iteratively eliminate dominated strategies and backward induct. As we do not observe preordered service in practice, we ran a laboratory experiment to verify whether or not, compared to sequential service, preordered service may reduce the number of contract cases filed with the courts. We describe the experimental design in Section 4.1, state our hypotheses in Section 4.2, and report our findings in Section 4.3. We discuss and interpret our empirical findings in Section 4.4.

4.1. Experimental design

Our experiment has three treatments that implement the model studied in Section 3, comprising a population of the game described in Fig. 1. The exogenous contract terms allow us to focus on the strategic interaction the different procedural rules give rise to. The payoffs to both players are equal when there is no investment, as well as when the investor enters the contract and the entrepreneur honors it. Thus, aversion to inequality in payoffs cannot be a potential motive for investors' actions.⁸ The binary action set of both players lends conceptual simplicity. If the investor had a continuous investment decision as in the standard trust game, then the entrepreneur would have to have a choice to decline the contract: the entrepreneur's payoff in the case of no investment might be greater than the payoff implied by a very small investment. Such additional complications would be a distraction.

When describing the game in the experimental materials we avoided the terminology used in the paper. Instead, we used generic labels for both players, as well as their actions. In all experimental sessions, we randomly assigned subjects to one of two roles: Player *A* and Player *B*. Each Player *A* was given a label, A1 to A8, which he or she kept for the duration of the experiment; Player *B*s had no labels. At the start of each experimental round, Player *B* was told the label of the Player *A* he or she was matched with. Player *A* had the option to *Enter* or to *Stay*; in the latter case, both players got a payoff of 1 Experimental Currency Unit (ECU). If Player *A* chose *Enter*, then Player *B* had the option of either *Send* or *Keep*. The former option gave each player 2 ECU; the latter option gave 0 ECU to Player *A* and 4 ECU to Player *B* if the contract was not enforced, and 2 ECU to Player *A* and 1.8 ECU to Player *B* if the contract was enforced. In terms of our model, this corresponds to $w = 2$ and $c = 0.2$, or 10% of Player *B*'s share of the surplus.

The three treatments differ in the way we operationalize the probability of enforcing a claim in the event of Player *B* choosing *Keep*. Our baseline condition is the NONE treatment, in which the probability of enforcement, p , is set to zero. This treatment corresponds to the case where there is no contract enforcement. Its purpose is to establish the baseline level of investment and compliance—the extent to which entered contracts are being honored—in our sample. The SEQSERVE treatment introduces enforcement of exactly one claim, which is chosen at random from the pool of claims in a given experimental round. This treatment captures sequential service at courts: claims are processed on a first-come-first-served basis and a claim's spot in the queue is random. Finally, the PREORDER treatment implements preordered service. All pairs in which Player *B* chose *Keep* in a given round are ranked in descending order of Player *A* label. The pair whose Player *A* has the highest label is selected for enforcement with certainty, and all others are not selected. Table 1 outlines the experimental design, and the sample size in each treatment. The unit of observation is an economy, which consists of eight investor-entrepreneur pairs. We collected six economies per treatment.

⁸ Bohnet and Huck (2004) and Bohnet et al. (2005) study binary trust games in which the no-investment outcome has first-movers earning more than second-movers. In contrast, the no-investment outcome in Bohnet et al. (2001) is zero for both players.

4.1.1. Information and feedback

We assigned labels to Player *As* because they are essential to operationalize the *PREORDER* treatment; Player *As* in *NONE* and *SEQSERVE* also had labels for consistency. Individual Player *As* kept their label throughout the experiment for two reasons. First, had we randomly assigned Player *A* labels every round, then the positions of Player *As* on the list in *PREORDER* would have changed every round. That is, while the strategic environment for Player *As* is unchanged from round to round in *SEQSERVE* and *NONE*, it would have changed every round for Player *As* in *PREORDER*. By fixing labels, we made the strategic environment stationary for Player *As* in all three treatments. Second, fixed labels are also realistic, as most identifying information for individuals and firms (e.g., taxpayer numbers) are unique to those individuals and do not vary. We did not assign labels to Player *Bs* because we did not want to introduce reputation concerns: if Player *As* could keep a history of past interactions for each Player *B* in the economy, they would be able to discriminate between trustworthy and untrustworthy Player *Bs*. This discrimination in turn can work as an extremely effective disciplining device, even in finitely repeated trust games (Bohnet and Huck 2004; Bohnet et al. 2005).

In all treatments, the end-of-round feedback screen subjects saw had two separate sections: one section had information about the round for all pairs in the economy, and the other had information about the outcome for that subject's pair (see Appendix C.4). The former section provided information about the 'Number of Player *Bs* that chose *Keep*' and the 'Number of Pairs *Selected*'.⁹ This feedback levels the playing field from an informational point of view for all subjects within a treatment, as well as across treatments. Even though our experiment simulates a one-shot environment, subjects can still learn about the population of players with whom they interact. Therefore, the informational environment should provide the same conditions for learning in the three treatments. This is particularly so in cases where the investor enters the contract, the entrepreneur breaches it, and their contract is enforced. In *PREORDER*, both players can immediately infer that either there were no instances of contract breach for any pair whose Player *A* had a higher label and invested, or that no Player *A* with a higher label invested. That is, those two players can infer an upper bound on the number of breaches in a round. This inference is not possible in *SEQSERVE*. Providing the total number of breaches in the economy in that round levels the information subjects can use to learn from round to round in both treatments.

4.1.2. Procedures

We implemented economies with eight pairs. We ran 10 experimental sessions, eight of which had 32 participants while two sessions had 16 participants for a total of 288 participants. In each of the 32-participant sessions, we ran two separate economies in parallel.¹⁰ In each session, subjects sat at a booth which did not allow for visual or verbal communication. The experimenter handed written copies of the instructions, which included a quiz to check for understanding.¹¹ Subjects read the instructions and completed the quiz in their own time; the experimenter checked their answers and clarified any questions individually if necessary. Once all subjects had their questions checked, the experiment started. There were two practice rounds so that subjects could get familiarized with the software interface. After the second practice round ended, the software informed subjects that the incentivized part would begin. The incentivized part of the experiment consisted of 20 rounds, three of which were randomly selected by the computer for payment. In all rounds, the computer randomly matched a Player *A* with a Player *B*. We programmed the experimental software in Z-Tree (Fischbacher 2007), and we recruited our subjects from a pool of undergraduate students from the University of Exeter from a variety of majors using ORSEE (Greiner 2015). No subject in the sample took part in more than one session, and nobody in the sample had registered for a trust game experiment before taking part in our experiment. The sessions took place in the FEELE lab at the University of Exeter in October and November 2016. Payoffs in the experiment were denominated in Experimental Currency Units (ECU). One ECU was worth £2; subjects also received £3 for participating. Sessions lasted on average 50 minutes and the average payment was £12.26 (\$15.30).

4.2. Hypotheses

The objective of the experiment is to check if preordered service decreases the number of contract cases filed with courts relative to sequential service in an economy consisting of many two-party contracts. As shown in Section 3, the unique equilibrium with preordered service is efficient, while sequential service allows for an efficient equilibrium as well as inefficient equilibria; without enforcement, autarky prevails. Our first hypothesis concerns the caseload faced by courts, and the impact the different procedural rules have on cases filed. We therefore focus on the two treatments where there is a judiciary to enforce claims.

Hypothesis 1 (Caseload). *The caseload faced by courts in PREORDER is smaller than or equal to that in SEQSERVE.*

Our next two hypotheses decompose caseload by looking at its two behavioral determinants: entrepreneur compliance, and investors' decision whether or not to invest. The *NONE* treatment serves as a no-enforcement benchmark condition.

⁹ The second sentence was not included in the feedback screen in the *NONE* treatment.

¹⁰ A large number of no-shows in one of the scheduled sessions forced us to run the two 16-person sessions.

¹¹ All instructions as well as screenshots of the experimental interface are collected in Appendix C.

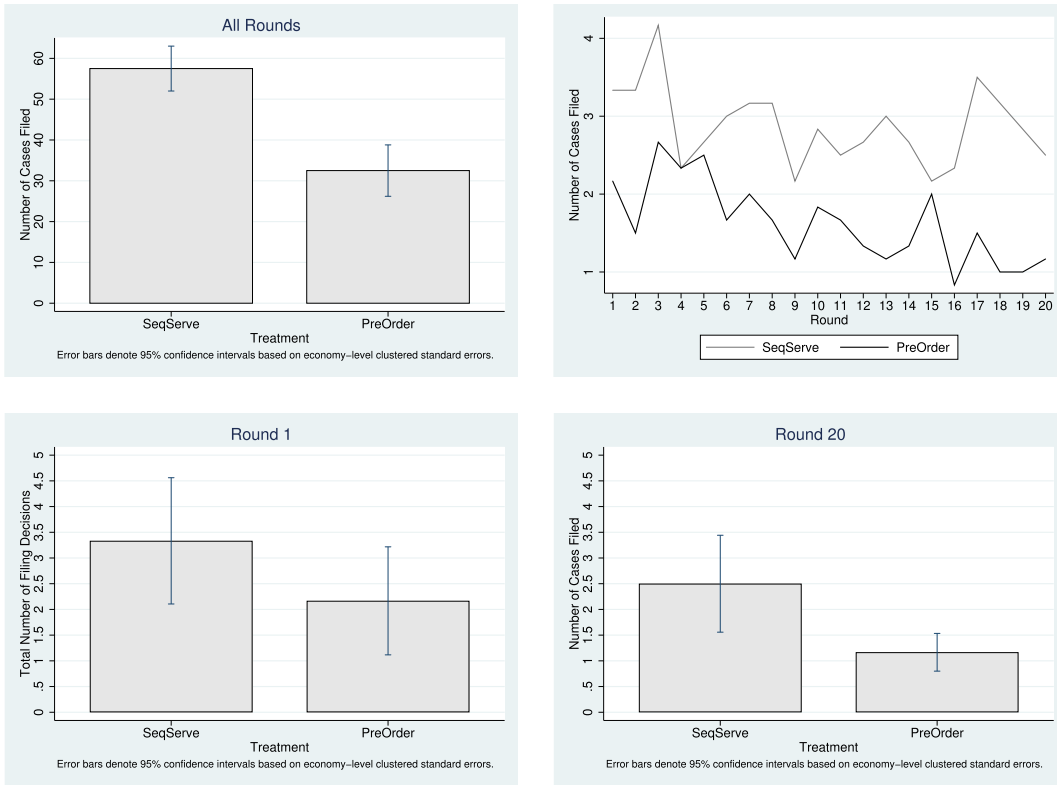


Fig. 4. Caseload. **Top Left:** All rounds. **Top Right:** Time series. **Bottom Left:** First round of the experiment. **Bottom Right:** Last round of the experiment.

Hypothesis 2 (Compliance). Treatments are ranked on the basis of the proportion of contracts being honored: $PREORDER \geq SEQSERVE > NONE$.

Hypothesis 3 (Investment). Treatments are ranked on the basis of the number of contracts being entered: $PREORDER \geq SEQSERVE > NONE$.

4.3. Empirical findings

In our analysis of the experimental data we use two complementary approaches. We test average treatment effects using non-parametric statistical tests, using economy averages as independent observations.¹² Panel data econometric techniques, which take advantage of the repeated nature of the experiment, merely confirm the results; as such, this analysis is relegated to Appendix B. Given that we have directional hypotheses, *p*-values relating to hypothesis tests refer to one-sided tests. To facilitate the discussion, and to be consistent with the terminology of our model, we refer to the first-mover player (Player A in the experimental instructions) as the investor. Likewise, we refer to the second-mover player (Player B in the experimental instructions) as the entrepreneur. In Sections 4.3.1–4.3.5, we present our results with support. In Section 4.4, we offer an integrated discussion and interpretations.

4.3.1. Caseload

In this section, we compare the caseload generated by economies under the two procedural rules. We compare the average of six observations of the total caseload in an economy, over 20 rounds, as well as the first and final rounds to ascertain for learning effects. The top left panel in Fig. 4 displays the total caseload in SEQSERVE and PREORDER over the course of the twenty rounds; the top right panel shows average caseload round-by-round. The bottom panels show average caseload in the first and final rounds of the experiment.

Result 1 (Caseload). Caseload is substantially lower in PREORDER than in SEQSERVE.

¹² We use MW to denote the Mann-Whitney test, WSR to denote the Wilcoxon Signed Ranks test and JT to denote the Jonkheere-Thepstra test.

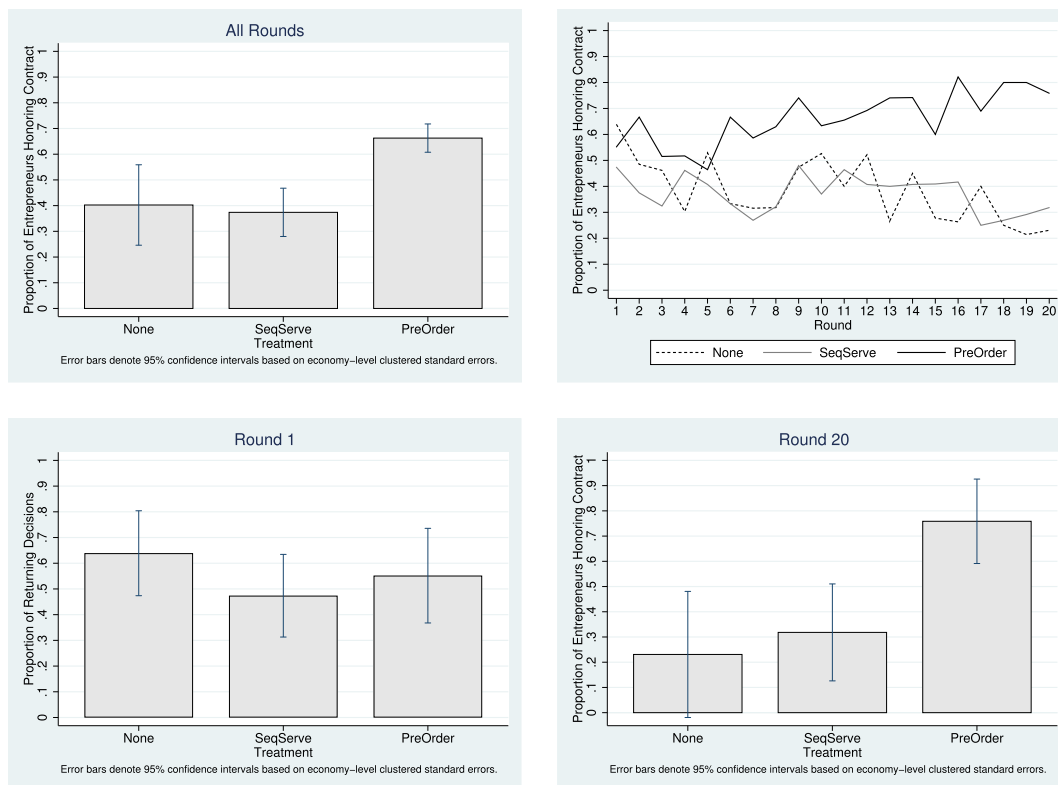


Fig. 5. Compliance. **Top Left:** All rounds. **Top Right:** Time series. **Bottom Left:** First round of the experiment. **Bottom Right:** Last round of the experiment.

Support. In Fig. 4, we observe a significantly lower (average) caseload in PREORDER than in SEQSERVE. Over the whole sample, there were just over 32 cases filed on average in PREORDER and just over 57 in SEQSERVE ($z = 2.882$, $p = 0.002$, MW test). In the first round, 3.33 cases were filed on average in SEQSERVE and 2.17 in PREORDER ($z = 1.187$, $p = 0.235$, MW test). However, by the final round of the experiment, there were 1.17 cases filed on average in PREORDER and 2.5 in SEQSERVE ($z = 2.331$, $p = 0.010$, MW test). The analysis of caseload over time corroborates this result. Average caseload is always lower in PREORDER than in SEQSERVE; we also find a negative and significant linear time trend on caseload in both PREORDER (-0.064 , $p < 0.001$), and SEQSERVE (-0.029 , $p = 0.003$). However, the linear trend in PREORDER is significantly different from that in SEQSERVE ($F(1, 11) = 7.48$, $p = 0.019$), indicating a faster rate of decline in caseload in PREORDER.

In the next two sections, we decompose this result; first by looking at average compliance conditional on investment, and then by looking at average investment. We add to the analysis the data from the NONE treatment, which gives us benchmark compliance and investment in the absence of courts.

4.3.2. Compliance

In this section, we compare entrepreneurs' behavior in the three treatments. We take the average of compliance decisions for all eight entrepreneurs in an economy (conditional on investment taking place) over 20 rounds ($N \leq 160$ per economy, since investment did not always happen) as one observation. We have six such observations. Fig. 5 depicts the relative frequency with which entrepreneurs complied with the contract, that is, honored it. The top left panel considers all twenty experimental rounds; the top right panel looks at average compliance round-by-round. The bottom panels depict compliance in the first and last round.

Result 2 (Compliance). Average compliance is higher in PREORDER than in both SEQSERVE and NONE. There is no difference in average compliance between SEQSERVE and NONE.

Support. In Fig. 5, we observe a significantly higher average compliance level in PREORDER than in SEQSERVE or NONE, when considering the whole sample (PREORDER vs SEQSERVE: $z = 2.882$, $p = 0.002$; PREORDER vs NONE: $z = 2.402$, $p = 0.008$; MW test). In the first round of the experiment, all three treatments perform rather similarly (PREORDER vs SEQSERVE: $z = 0.734$, $p = 0.463$; PREORDER vs NONE: $z = 0.567$, $p = 0.570$; SEQSERVE vs NONE: $z = 1.043$, $p = 0.297$; MW test). However, as the bottom left panel of Fig. 5 shows, by the final round of the experiment the difference in compliance between PREORDER and the other two treatments is large and highly significant (PREORDER vs SEQSERVE: $z = 2.887$, $p = 0.002$; PREORDER vs

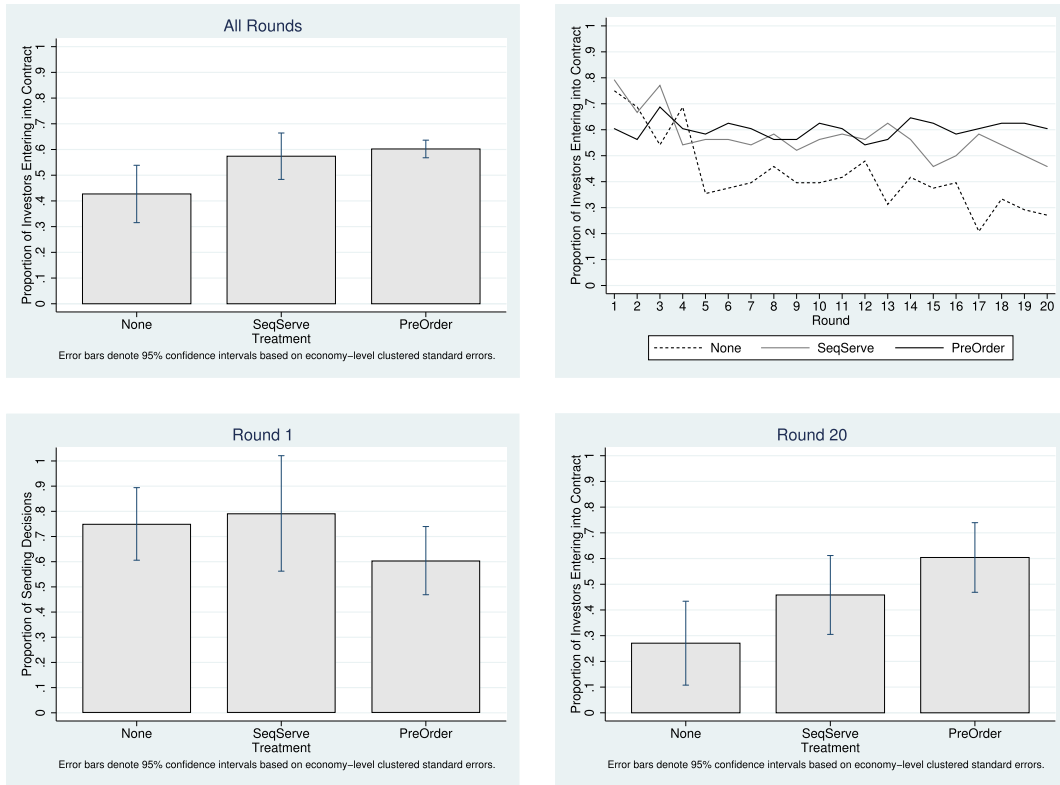


Fig. 6. Investment. **Top Left:** All rounds. **Top Right:** Time series. **Bottom Left:** First round of the experiment. **Bottom Right:** Last round of the experiment.

NONE: $z = 2.917, p = 0.002$; MW test). In contrast, we do not find a significant difference in average compliance between SEQSERVE and NONE over all rounds ($z = 0.000, p = 0.500$; MW test) but a marginal significance in the final round ($z = 1.624, p = 0.052$; MW test).

The top right panel of Fig. 5 corroborates this result by displaying the evolution of compliance over time. Fitting a linear time trend on the average per round compliance in each treatment, we find a negative and significant trend in NONE ($-0.014, F(1, 17) = 11.11, p = 0.004$), a negative, though not significantly different from zero, trend in SEQSERVE ($-0.002, F(1, 17) = 1.85, p = 0.191$), and a positive and highly significant time trend in PREORDER ($0.012, F(1, 17) = 24.76, p < 0.001$).¹³ That is, while compliance declined over time in the NONE treatment—which is consistent with the existing evidence on trust games (e.g., Bohnet and Huck 2004)—compliance remained roughly constant in the SEQSERVE treatment and actually increased over time in the PREORDER treatment.

4.3.3. Investment

In this section, we compare investors' behavior in the three treatments. We take the average of investment decisions for all eight investors in an economy over 20 rounds as one observation ($N = 160$ per economy). We have six such observations. Fig. 6 depicts the relative frequency with which investors entered into a contract with the entrepreneur. As in previous figures, we display average investment over 20 rounds, the first and last round, as well as a round-by-round plot.

Result 3 (Investment). Average investment is higher in both PREORDER and SEQSERVE than in NONE. There is no difference in average investment between PREORDER and SEQSERVE.

Support. In Fig. 6, we observe a significantly higher level of investment in PREORDER than in NONE over all rounds ($z = 2.406, p = 0.008$, MW test). This difference belies the evolution of investment over the course of the experiment. In the first round, there is no difference in average investment ($z = 1.338, p = 0.181$, MW test); by the final round, average investment in PREORDER is significantly higher than NONE ($z = 2.303, p = 0.011$, MW test). While SEQSERVE significantly outperforms NONE over all rounds ($z = 2.330, p = 0.010$, MW test), there is only a marginally significant difference in average investment in both the first ($z = 0.818, p = 0.413$, MW test) and the final round ($z = 1.483, p = 0.069$, MW test). When comparing

¹³ All time trends are significantly different from each other at the 5% level or lower. See Table B1 in the Appendix for details.

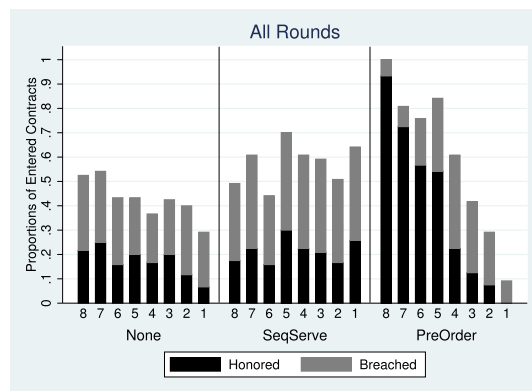


Fig. 7. Breakdown of average investment and compliance by investor label.

SEQSERVE to PREORDER, the difference in average investment is not significant over the course of the experiment ($z = 1.000$, $p = 0.500$, MW test). Investment in SEQSERVE is marginally significantly higher than in PREORDER in the first round ($z = 1.721$, $p = 0.085$, MW test); however, by the final round average investment is nominally (though not significantly) higher in PREORDER: $z = 1.257$, $p = 0.105$; MW test).

This again illustrates the role of the evolution of average investment over time in the three treatments as depicted in the top right panel of Fig. 6: in NONE, there is a negative and significant (-0.019 , $F(1, 17) = 25.35$, $p = 0.002$) time trend on investment. The time trend on investment in SEQSERVE is also negative and significantly different from zero (-0.010 , $F(1, 17) = 12.15$, $p = 0.003$). In contrast, the linear trend on investment in PREORDER is flat and not significantly different from zero (0.0003 , $F(1, 17) = 0.00$, $p = 0.948$).

4.3.4. Compliance and investment conditional on labels

In this section we look at individual behavior in more detail. A closer look at Figs. 5 and 6, in particular the 95% confidence intervals around the sample means, suggests that the variance in both compliance and investment is higher in SEQSERVE than in PREORDER. Fig. 7 breaks down investment and compliance by investor label. Column heights represent average investment for each label. The dark portion in each column represents the proportion of cases where the entrepreneurs honored the contract; the light portion represents the proportion of cases where entrepreneurs breached the contract.

Result 4 (Labels).

- There is a positive relationship between investment/compliance and investor labels in PREORDER but not in NONE or SEQSERVE.
- The levels of investment and compliance for investors with the six highest labels in PREORDER are greater than or equal to those levels for the average investor in SEQSERVE and NONE.
- There is less variance in individual investment decisions in PREORDER than either in NONE or SEQSERVE.

Support. *a.* Visual inspection of Fig. 7 suggests that average investment and compliance vary with label in PREORDER but not in NONE or SEQSERVE. To test for a relationship between labels and investment/compliance, we compute the average investment by individual investor, as well as the average compliance experienced by individual investors over twenty rounds in an economy and take that to be our independent unit of observation. We then compute Spearman correlation coefficients (ρ), as well as the (non-parametric) Jonkheere-Thepstra test (JT) for ordered alternatives.¹⁴

In PREORDER we find evidence of a strong relationship between investor labels and investment. There is a positive correlation between investment and investor labels ($\rho = 0.76$, $p < 0.001$); the JT test rejects the null of joint equality against the ordered alternative ($J^* = 5.950$, $p < 0.001$). We also find evidence of a strong relationship between investor labels and compliance. The correlation between compliance and labels is positive and highly significant ($\rho = 0.92$, $p < 0.001$); the JT test rejects the null of joint equality against the ordered alternative ($J^* = 7.213$, $p < 0.001$).

In SEQSERVE we find no evidence of a relationship between investor labels and investment. There is a negative but non-significant correlation between investment and investor labels ($\rho = -0.14$, $p = 0.350$); the JT test does not reject the null of joint equality ($J^* = -0.945$, $p = 0.345$). We also find no evidence of any relationship between investor labels and compliance. The correlation between compliance and labels is essentially zero and non-significant ($\rho = 0.03$, $p = 0.844$); the JT test does not reject the null of joint equality ($J^* = 0.296$, $p = 0.767$).

¹⁴ The JT test tests for the null hypothesis of joint equality of medians against an ordered alternative hypothesis. The alternative is that the median investment level is increasing or decreasing in the investor labels. We report J^* -statistics that correct for ties.

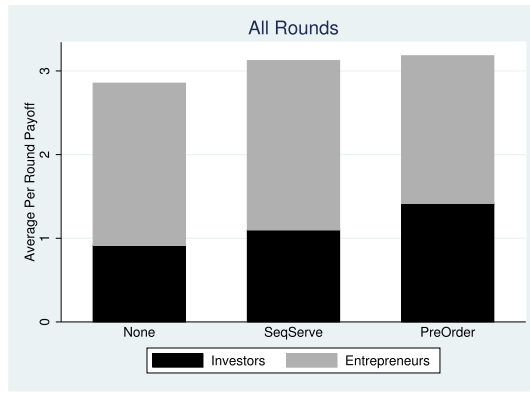


Fig. 8. Average payoffs to investors and entrepreneurs.

In NONE we find very weak evidence of a positive relationship between investor labels and investment. There is a positive, weakly significant correlation between investment and investor labels ($\rho = 0.24, p = 0.098$), but the JT test does not reject the null of joint equality ($J^* = 1.637, p = 0.102$). We find no evidence of any relationship between investor labels and compliance. The correlation between compliance and labels is positive but non-significant ($\rho = 0.16, p = 0.283$); the JT test does not reject the null of joint equality ($J^* = 1.037, p = 0.300$).

b. We compare the average investment by an investor in PREORDER conditional on a label to the average investment by all investors in NONE and SEQSERVE. To do this, we take the average investment over the twenty rounds of the experiment for each investor as the independent observation.¹⁵

The level of investment by investors with labels A8–A5 in PREORDER is significantly higher than that by investors in NONE (all comparisons, $z \geq 2.193, p \leq 0.028$, MW test) and SEQSERVE (A6: $z = 1.740, p = 0.082$; all other comparisons, $z \geq 2.168, p \leq 0.030$, MW test). Also, the level of investment by investors with labels A4–A2 in PREORDER is not significantly different to that by investors in NONE (all comparisons, $z \leq 1.448, p \geq 0.148$, MW test); the average level of investment by investors with labels A4–A3 in PREORDER is not significantly different to that by investors in SEQSERVE (all comparisons, $z \leq 1.269, p \geq 0.204$, MW test). In contrast, the level of investment by investors with label A1 in PREORDER is significantly lower than that by investors in NONE ($z = 3.146, p = 0.002$, MW test). Likewise, the level of investment by investors with labels A2 and A1 in PREORDER is significantly lower than that by investors in SEQSERVE (all comparisons, $z \geq 2.168, p \leq 0.030$, MW test).

We observe the same pattern when looking at compliance conditional on investment. The level of compliance experienced by investors with labels A8–A5 in PREORDER is significantly higher than that experienced by investors in NONE (all comparisons, $z \geq 3.175, p \leq 0.001$, MW test) and SEQSERVE (all comparisons, $z \geq 3.622, p \leq 0.001$, MW test). The level of compliance experienced by investors with labels A4 and A3 in PREORDER is not significantly different to that experienced by investors in NONE (all comparisons, $z \leq 0.734, p \geq 0.463$, MW test) or SEQSERVE (all comparisons, $z \leq 0.868, p \geq 0.385$, MW test). In contrast, average compliance experienced by investors with labels A2 and A1 in PREORDER is significantly lower than that by investors in NONE (all comparisons, $z \geq 1.658, p \leq 0.097$, MW test) or SEQSERVE (all comparisons, $z = 1.936, p \leq 0.053$, MW test).

c. We take advantage of the fact that the investors’ decision is a binary variable. Therefore, the sample standard deviation of an individual’s investment decisions over the twenty rounds in the experiment is maximized when a given investor invests half of the time (i.e., in ten out of twenty rounds), and it is minimized when an investor either always invests or never invests. We calculate the standard deviation of investment decisions for each investor in our sample. The mean standard deviation in investment decisions in PREORDER is 0.248, and it is significantly lower than that in SEQSERVE (0.401; $z = 4.260, p < 0.001$, MW test) and NONE (0.403, $z = 3.945, p < 0.001$, MW test). In contrast, there is no difference in mean standard deviation between SEQSERVE and NONE ($z = 0.774, p = 0.439$; MW test).

4.3.5. Payoffs

We complete our analysis by examining players’ payoffs across the three treatments. Fig. 8 depicts average payoffs for both types of players over the 20 rounds. The height of the bar represents average total payoffs. The dark portion of the bar denotes average investor payoffs, and the light portion denotes average entrepreneur payoffs.

¹⁵ Note that since we are making comparisons across treatments and labels were fixed within an economy, these are statistically independent observations. Also note that the non-parametric tests we employ are robust to unequal sample sizes.

Result 5 (Payoffs).

- a. In all treatments, the average payoff to investors from investing is at least as high as that from not investing.
- b. Average investor payoff is higher in both PREORDER and SEQSERVE than in NONE. Average investor payoff is higher in PREORDER than in SEQSERVE.
- c. Average entrepreneur payoff is lower in PREORDER than in SEQSERVE.

Support. a. Using the average payoff of all investors in an economy as an independent observation, we test against the null that average payoff is equal to one. We reject the null in the cases of PREORDER ($z = 3.064$, $p < 0.01$, WSR test) and SEQSERVE ($z = 2.680$, $p < 0.01$, WSR test), but we cannot reject it in the case of NONE ($z = 1.503$, $p = 0.133$, WSR test).

b. Using the average investor payoff in an economy as an independent observation, we perform a series of pairwise tests of equality of payoffs. The average investor payoff in PREORDER is 1.42, which is significantly higher than the average investor payoff in SEQSERVE (1.10, $z = 2.882$, $p < 0.01$, MW test), and significantly higher than the average investor payoff in NONE (0.93; $z = 2.882$, $p < 0.01$, MW test). In turn, average investor payoff in SEQSERVE is also significantly higher than that in NONE ($z = 2.082$, $p = 0.04$, MW test).

c. We perform the same exercise using average entrepreneur payoffs. The average entrepreneur payoff in PREORDER is 1.77, which is not significantly different from the average entrepreneur payoff in NONE (1.94, $z = 1.601$, $p = 0.11$, MW test); it is significantly different from the average entrepreneur payoff in SEQSERVE (2.02, $z = 2.817$, $p < 0.01$, MW test). Finally, there is no statistically significant difference in average entrepreneur payoffs between SEQSERVE and NONE ($z = 1.506$, $p = 0.132$, MW test).

4.4. Discussion of the experiment and our findings

The objective of our experiment is to compare the performance of preordered service to that of sequential service. We find that preordered service indeed substantially reduces the case inflow facing the judiciary. It does so throughout the experiment, up to the very last round. We find this is primarily driven by higher entrepreneur compliance.

When examining the performance of the procedural rules, it is useful to compare them with NONE, which is a binary version of the trust game (Berg et al. 1995). NONE provides the baseline level of investment and compliance in the absence of any procedure for contract enforcement. Consistent with the existing literature on trust games,¹⁶ we observe positive, albeit declining levels of trust and trustworthiness in our sample. How do SEQSERVE and PREORDER compare? Investment levels are significantly higher in SEQSERVE than in NONE, but compliance levels are not. The increase in investment is easily rationalizable, as investors have a positive probability of at least 1/8 of being made whole, as opposed to zero in NONE. Investors in SEQSERVE and NONE earn on average the same payoff: that which they would have earned had they not invested at all. However, average aggregate payoff in both treatments is 50% higher than that if investment had not taken place. Therefore, investors could be motivated by efficiency concerns (Engelmann and Strobel 2004). In contrast, the PREORDER treatment improves both investment and compliance relative to NONE. In other words, effective enforcement is not just about the availability of arbitration per se, but also about what procedures govern arbitration itself.

Comparing the two procedural rules directly, we find that preordered service outperforms sequential service in terms of entrepreneur compliance, but not in terms of investment. We did not consider a longer time horizon in the lab, as this could lead to subject fatigue and compromise decision quality. However, the negative time trend in investment in SEQSERVE and the flat time trend in PREORDER suggest that with a sufficiently long time horizon, PREORDER might have outperformed SEQSERVE in this dimension as well.

While we cannot demonstrate in our experiment that preordered service leads to more low-profile contracts being entered than sequential service, it does reduce the inflow of court cases for breach of such low-profile contracts. Informed by the literature mentioned above (Palumbo et al. 2013; Chemin 2012; Jappelli et al. 2005; Fabbri 2010), we would expect the resulting reduced overall caseload to speed up the courts, leading to a more efficient allocation of resources—in particular, to more lending and investment with higher profile, and more firm growth—and thus higher welfare. Moreover, it is conceivable that breaching a contract is associated with additional efficiency costs that we have not accounted for. For example, upon breach, the entrepreneur might not be able to enjoy the full output of the project because, e.g., diverting resources through unofficial channels and disguising the illegitimate nature of the proceeds might incur a deadweight loss. That is, every breached contract might be associated with an efficiency loss. It might thus be the case that preordered service increases welfare due the smaller number of breached contracts, even if overall investment does not increase.

The preordered service rule implemented in PREORDER exhibits sorting characteristics: investors whose labels put them on the top half of the list almost always invested, and on average their contracts were almost always honored. In turn, investors on the very bottom of the list rarely invested. Remarkably, this sorting occurred with only very little information compared to reputation transmission environments that successfully promote trust (Bohnet and Huck 2004; Bohnet et al. 2005). In addition to their own pair's label, players only knew the total number of breaches by entrepreneurs in previous rounds, which at best puts a lower bound on the total level of investment. This lack of detailed information is particularly

¹⁶ For a review of evidence from the lab, see Camerer (2003); for field data, see Henrich et al. (2001).

important when service is preordered. The equilibrium in this case requires players to realize that entrepreneurs should choose to comply, and that therefore investment is safe, which might be easier with more detailed information.

Had we provided full disclosure about the outcome in each pair in the economy, the performance in PREORDER might have improved (faster), as subjects would have observed almost perfect compliance for the pairs with label A8 and A7. Had we not provided any feedback to subjects on breaches by other entrepreneurs in the economy, PREORDER would have had an inherent advantage over SEQSERVE. In PREORDER, if the entrepreneur in a pair breaches the contract, then the players in that pair can infer from their enforcement outcome the behavior of players in pairs in line before them as well as, potentially, an upper bound on the number of breaches in that round. This sort of inference is not possible in SEQSERVE. As such, our choice of end-of-round aggregate feedback removed an informational edge inherent to PREORDER.

It is clear that the behavior in the experiment did not match the equilibrium predictions of our model, particularly in the case of PREORDER. That being said, we are not interested in testing theoretical point predictions, but rather treatment effects. While efficiency and reciprocity preferences likely played a role in determining investment and compliance, they do not explain the decline in compliance and investment along the queue. A natural candidate explanation is that subjects in our experiment were boundedly rational, and that they may have had difficulty performing the reasoning inherent to the equilibrium of the model.

However, popular models of boundedly rational agents struggle to explain behavior in this treatment. Take for instance the level- k model proposed by Nagel (1995) and Stahl and Wilson (1995). Assuming that a level-0 player selects a strategy at random, it can be verified that the optimal action by a Player B with sophistication level 1 (and a reasonable degree of risk aversion) in any position below A8 is to breach. By breaching the contract, such a player can either achieve a payoff of 4 ECU with probability 0.5, or a payoff of 1.8 ECU with probability 0.5; by honoring the contract, that player gets a sure payoff of 2 ECU.¹⁷ In turn, a level-2 Player B will find it dominant to breach. As such, any Player B that has a sophistication level greater than or equal to 1 will optimally breach. As a consequence, the cognitive hierarchy model of Camerer et al. (2004) makes the same prediction.

As the success of preordered service builds on iterated elimination of dominated strategies and backward induction, it is conceivable that it will fail to outperform sequential service for a large enough number of pairs in our experimental economy, in which only one breached contract is enforced. We think that the robustness of our findings with respect to the number of pairs in the economy is of some, but limited interest. As long as preordered service is effective for some nontrivial list length, the courts can use many lists of at most that length. What that critical length is and how the many lists are best generated and employed is an empirical question that a laboratory experiment is not well-suited to answer. We discuss these and other aspects of preordered service in practice in Section 5.

5. Taking preordered service to practice

In this section, we depart from the simplified exposition we adopt throughout to discuss how preordered service could be taken to practice. To start, note that throughout we use “enforcement” to mean that the courts hold hearings and issue a ruling that can then be used in enforcement proceedings; and preordered service determines the order of initial hearings of court cases. Also recall that the focus of this paper as well as the bulk of the caseload at civil courts are low-profile contract cases with a sure outcome in court.

Regarding concerns that some might have about equal access to justice when service at courts is preordered—which may explain why this sort of procedural rule has not emerged—we submit two observations. First, in theory, while there is a hierarchy of potential contract cases *ex ante*, in equilibrium, access to justice is equal. Second, to some extent, many legal systems already distinguish between various types of court cases at a general level. Often civil cases are differentiated by the amount claimed. For instance, small claims are typically dealt with by different courts than larger commercial disputes. That is, specialized courts already exist, and they give priority to certain types of cases by definition.

The additional administrative costs of introducing preordered service of low-profile contract cases with a sure outcome in practice are likely to be manageable. One would have to identify the categories of contract cases that preordered service should be applied to. For each of these categories, an appropriate period of time to be used, which might vary across categories, has to be determined. These time frames would likely relate to payment patterns commonly associated with the respective contract category, account for a grace period, and accommodate other timing-related aspects inherent to both the contract category and the legal system. Recall that many of the low-profile contract cases we focus on are debt collections and landlord-tenant disputes. As examples that could lead to disputes fitting into these categories, cellphone and housing rental contracts tend to specify monthly payments.

Given these categories and the associated specified periods of time, with a preordered list in place and announced, when new court cases are filed, those cases belonging to these categories have to be identified. One possibility is to use a form on which the type of contract the case concerns is to be indicated. Court personnel would have to collect and store the information relating to filed contract cases belonging to these categories. At the end of the specified period, they would

¹⁷ A player would require a coefficient of relative risk aversion of at least 7.53 to prefer compliance; this degree of risk aversion is not only implausible, as argued by Rabin (2000), but also not observed in most laboratory or field experiments measuring individual risk aversion (Biswanger 1980; Holt and Laury 2002).

have to order the filed cases according to the preordered list. Then, one might have to hold special hours at, e.g., small claims courts to deal with these cases. To the extent that many of the types of cases in question are already heard at small claims courts, a reduced caseload due to preordered service of these cases reduces delay at these courts, if any, and allows to reallocate resources to other courts and activities.

As to the list, a prerequisite for preordered service is that collating and ordering uniquely identifying information to the extent required is not unlawful, e.g., regarding privacy concerns. Once such concerns are addressed, preordered service can make use of existing uniquely identifying information, such as taxpayer numbers. It can equally well use newly randomly generated and assigned numbers. In principle, it does not matter whether the identifying label associated with a claim refers to the plaintiff, the defendant, or is a combination of information relating to both. Courts can use any randomly generated ordered list of these labels. There can be many different lists for different contract types—as well as many different lists for the same contract type, in which case each potential contract can appear on only one of the lists. All of these lists can order the labels in the same way, or in different ways. Notice that our focus on serving only one court case is a stark simplification. Agor et al. (2015) suggest that the U.S. court system has the capacity to dispose more than 18 million civil cases per year.¹⁸ That is, the court system can promise that, for each of these lists, very many filed court cases can be served so that, potentially, very many breached contracts can be enforced.

These last couple of points also suggest a way to allay the concern that very long lists may make it more difficult for preordered service to work as theory suggests: courts could use many shorter lists. These lists could be based on subject matter categories, contract type, industry, geography, or any other relevant characteristic. In fact, as mentioned in Section 3.5, Agor et al. (2015) report that the contract cases in their dataset are mostly debt collections, landlord-tenant disputes, and foreclosures, suggesting rather broad natural categories to start from. As argued above, on each list, many cases can be served. How many (or how few) lists are necessary to successfully implement preordered service of low-profile contract cases will depend on legal jurisdiction and legal custom. Answering this question—which is inevitably an empirical one—is beyond the scope of this paper and straddles economics, public policy, and law.

In order to allow for multiple contractual relationships per entity, contracts can be ranked lexicographically, where the first and second coordinates of the label could be the respective taxpayer numbers of the associated parties. To allow for multiple contracts between the same parties the label can be extended by an index in any arbitrary way. At the outset, a conceivable number of contracts between the same parties within the reference period, say 100, or 1000, could be fixed, and an index running from 1 to 100, or 1000, can be attached to the identifying label of the contracting parties to thereby uniquely identify any one of their contracts. This approach also works with financial intermediaries. Similarly, contracts with more than two parties can be accommodated in much the same way, as long as every contract is associated with a uniquely identifying label. There could be separate lists for contracts involving two parties, three parties, four parties, and so on.

Alternatively, for each entity that can enter contracts, one could simply use its taxpayer number, combine it with an index indicating the count of contracts the entity is a party to, and identify contracts with multiple labels, each one associated with one of the involved parties. That is, a contract with three parties could have three different labels, each one made up of, for example, one party's taxpayer number and an index referring to this specific contract among all the hundreds, or thousands, of contracts this entity signed during the reference time period. Of course, uniquely identifying a contract with multiple labels is not a problem. A filed court case associated with two different uniquely identifying labels is served at the position of the higher ranked label of the two. The position at the lower ranked label of the two is simply vacant as the case has already had its initial hearing. These approaches also work when contracts are made through intermediaries.

All preordered service of low-profile contract cases with a sure outcome requires is that all potential contracts are associated with at least one uniquely identifying label, according to which they can be ordered at the outset. There could be multiple such labels for each contract, or only for some, and many ordered lists. Identifying the best way to implement it requires further research and, again, straddles economics, public policy, and law.

6. Concluding remarks

Delay and case backlog are evident at civil courts in many countries and have been shown to be linked to more breaches of contract and less investment, less lending and tighter credit constraints, higher firm financing costs and smaller firms. We offer a conceptual demonstration of a procedural rule that may help mitigate an important source of delay and backlog: the large number of civil cases filed. Our procedural rule stipulates that courts line up low-profile contract cases for service using a preannounced ranking by some uniquely identifying information of at least one party to the contract. We call this procedure preordered service, as compared to sequential service, which lines up civil cases in order of arrival.

We offer proof of concept by presenting preordered service in a stylized environment. We demonstrate that replacing sequential service of low-profile contract cases with preordered service has the potential to reduce the caseload at courts. In an experimental investigation, preordered service reduces the number of contract cases filed with the courts by more than 40 percent compared to sequential service. As contracts are honored more frequently, investors can expect a higher payoff. While our analysis is not suited to discuss the economy's progression when preordered service is first adopted to replace

¹⁸ Agor et al. (2015, p. 16) state that the 925,344 cases in their dataset, which were disposed between July 1, 2012 and June 30, 2013, amount to about five percent of the U.S. state courts civil caseload.

sequential service, to some extent the sequence of one-shot games our subjects played in the experiment can be interpreted as one-shot interactions in a stylized economy. In that case, the time series of the investment and breach decisions seems to suggest that the positive effects of preordered service compared to sequential service become more pronounced over time.

While we cannot demonstrate in our experiment that preordered service leads to more low-profile contracts being entered than sequential service, it does reduce the inflow of court cases for breach of such low-profile contracts. Informed by the insights of a growing literature (e.g., Palumbo et al. 2013; Chemin 2012; Jappelli et al. 2005; Fabbri 2010), we would expect the resulting reduced overall caseload to speed up the courts, leading to a more efficient allocation of resources—in particular, to more lending and investment with higher profile, and more firm growth—and thus higher welfare.

There are several avenues for future research. One line of inquiry is to explore whether the provision of more information and a role for reputation can make preordered service even more effective. As the cost upon enforcement was very small in our experiment, one may ask whether higher costs, maybe due to large fines, can further improve the performance. Similarly, one may ask what the effect of a cost associated with filing a court case is. Beyond the laboratory, we think of our findings as suggestive that preordered service should be assessed in a field experiment, which may also be informative as to how to best implement it.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.geb.2020.04.007>.

References

- Agor, P.H., Graves, S.E., Miller, S.S., 2015. The Landscape of Civil Litigation in State Courts. National Center for State Courts (NCSC) Report.
- Andreoni, J., Gee, L.K., 2012. Gun for hire: delegated enforcement and peer punishment in public goods provision. *J. Public Econ.* 96, 1036–1046.
- Andreoni, J., Gee, L.K., 2015. Gunning for efficiency with third party enforcement in threshold public goods. *Exp. Econ.* 18, 154–171.
- Bar-Gill, O., Ben-Shahar, O., 2009. The Prisoners' (plea bargaining) dilemma. *J. Leg. Anal.* 1 (2), 737–773.
- Basu, S., Dickhaut, J., Hecht, G., Towry, K., Waymire, G., 2009. Recordkeeping alters economic history by promoting reciprocity. *Proc. Natl. Acad. Sci.* 106 (4), 1009–1014.
- Berg, J., Dickhaut, J., McCabe, K., 1995. Trust, reciprocity, and social history. *Games Econ. Behav.* 10, 122–142.
- Biswanger, H.P., 1980. Attitude toward risk: experimental measurement in rural India. *Am. J. Agric. Econ.* 62, 395–407.
- Bohnet, I., Frey, B.S., Huck, S., 2001. More order with less law: on contract enforcement, trust, and crowding. *Am. Polit. Sci. Rev.* 95, 131–144.
- Bohnet, I., Hamgart, H., Huck, S., Tyran, J.-R., 2005. Learning trust. *J. Eur. Econ. Assoc.* 3 (2–3), 322–329.
- Bohnet, I., Huck, S., 2004. Repetition and reputation: implications for trust and trustworthiness when institutions change. *Am. Econ. Rev.* 94 (2), 362–366.
- Bray, R.L., Coviello, D., Ichino, A., Persico, N., 2016. Multitasking, multiarmed bandits, and the Italian judiciary. *Manuf. Serv. Oper. Manag.* 18 (4), 545–558.
- Camerer, C.F., 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Russell Sage Foundation, New York.
- Camerer, C.F., Ho, T.-H., Chong, J.-K., 2004. A cognitive hierarchy model of games. *Q. J. Econ.* 119 (3), 861–898.
- Chemin, M., 2012. Does court speed shape economic activity? Evidence from a court reform in India. *J. Law Econ. Organ.* 28 (3), 460–485.
- Coviello, D., Ichino, A., Persico, N., 2014. Time allocation and task juggling. *Am. Econ. Rev.* 104 (2), 609–623.
- Coviello, D., Ichino, A., Persico, N., 2015. The inefficiency of worker time use. *J. Eur. Econ. Assoc.* 13 (5), 906–947.
- Dari-Mattiacci, G., De Geest, G., 2010. Carrots, sticks, and the multiplication effect. *J. Law Econ. Organ.* 26 (2), 365–384.
- Engelmann, D., Strobel, M., 2004. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *Am. Econ. Rev.* 94 (4), 857–869.
- Fabbri, D., 2010. Law enforcement and firm financing: theory and evidence. *J. Eur. Econ. Assoc.* 8 (4), 776–816.
- Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10, 171–178.
- Gennaioli, N., Shleifer, A., 2007. The evolution of common law. *J. Polit. Econ.* 115 (1), 43–68.
- Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* 1, 114–125.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., McElreath, R., 2001. In search of homo economicus: behavioral experiments in 15 small-scale societies. *Am. Econ. Rev.* 91 (2), 73–78.
- Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. *Am. Econ. Rev.* 92 (5), 1644–1655.
- Jappelli, T., Pagano, M., Bianco, M., 2005. Courts and banks: effects of judicial enforcement on credit markets. *J. Money Credit Bank.* 37 (2), 223–244.
- Kamijo, Y., Nihonsugi, T., Takeuchi, A., Funaki, Y., 2014. Sustaining cooperation in social dilemmas: comparison of centralized punishment institutions. *Games Econ. Behav.* 84, 180–195.
- Landes, W.M., Posner, R.A., 2009. Rational judicial behavior: a statistical study. *J. Leg. Anal.* 1 (2), 775–831.
- McKelvey, R.D., Palfrey, T.R., 1992. An experimental study of the centipede game. *Econometrica* 60 (4), 803–836.
- Mookherjee, D., Png, I.P.L., 1992. Monitoring vis-à-vis investigation in enforcement of law. *Am. Econ. Rev.* 82, 556–565.
- Nagel, R., 1995. Unraveling in guessing games: an experimental study. *Am. Econ. Rev.* 85 (5), 1313–1326.
- Niblett, A., Posner, R.A., Shleifer, A., 2010. The evolution of a legal rule. *J. Leg. Stud.* 39 (2), 325–358.
- Palumbo, G., Giupponi, G., Nunziata, L., Mora-Sanguinetti, J., 2013. *Judicial Performance and its Determinants: A Cross-Country Perspective*. OECD Economic Policy Paper No. 05.
- Posner, R.A., 2005. Judicial behavior and performance: an economic approach. *Florida State Univ. Law Rev.* 32, 1259–1279.
- Rabin, M., 2000. Risk aversion and expected-utility theory: a calibration theorem. *Econometrica* 68 (5), 1281–1292.
- Schwartz, A., 1979. The case for specific performance. *Yale Law Policy Rev.* 89, 271–306.
- Shavell, S., 1980. Damage measures for breach of contract. *Bell J. Econ.* 11 (2), 466–490.
- Shavell, S., 1984. The design of contracts and remedies for breach. *Q. J. Econ.* 99 (1), 121–148.
- Stahl, D.O., Wilson, P., 1995. On players' models of other players: theory and experimental evidence. *Games Econ. Behav.* 10, 213–254.