Repeated Games

Guillaume R. Fréchette
NYU

Emanuel Vespa
UCSD

Sevgi Yuksel
NYU

December 3, 2024

1 Introduction

An important contribution of game theory has been to elucidate how repeated interactions can affect outcomes by facilitating credible punishments and rewards. For instance, repeated interaction can enable cooperation to arise in settings where this is ruled out by individual rationality in a one-shot interaction. Many problems that are of interest in economics, other social sciences, and in the natural sciences involve repeated interactions. As such, repeated games have been useful in providing insights on a wide range of applications, such as collusion, the transition between democracy and dictatorship, bargaining and public debt. Empirically studying behavior in repeated games can serve multiple purposes: (1) To test predictions of game theory—do people react to dynamic incentives in ways consistent with theory? (2) To provide guidance for equilibrium selection, i.e. to better understand if behavior is systematic and whether certain factors can organize and predict outcomes. (3) To understand the strategies subjects use to support those outcomes. These goals are particularly relevant for infinitely repeated games, where standard theory often fails to make sharp predictions. Experiments on repeated games can be particularly useful as the laboratory allows control over important features of the environment that are difficult to observe or measure in naturally occurring settings, such as the discount factor and the type of information that is revealed between interactions.

In this chapter, we discuss the methodological challenges associated with designing experiments to study behavior in repeated games. Our goal is to summarize the methods developed in the literature over the last few decades to overcome these challenges. We will focus heavily on experiments in repeated prisoner's dilemma games, where much of the literature has concentrated. But, throughout this chapter, we will highlight how these methods are applicable to the study of other repeated games and can be useful in settings beyond repeated games.[1] We motivate and describe the challenges by revisiting the first experiment on the prisoner's dilemma, but a brief summary of the challenges is provided in the last paragraph of the introduction.

In January of 1950, Merril M. Flood and Melvin Dresher, two mathematicians working at the RAND corporation, invited UCLA economist Armen Alchian and RAND mathematician John D. Williams to play a game that has since become known as the prisoner's dilemma (PD). Each player simultaneously selects either to cooperate (C) or defect (D); the joint-revenue maximizing payoffs occur when both select C, but the unique Nash equilibrium of the one-shot play involves both selecting the dominant action D. For our purposes here we do not need more details on the actual payoffs, but this information is provided in Flood (1958).

Flood and Dresher were interested in testing the predictions of Nash equilibrium in the one-shot game. Even though their participants were familiar with developments in game theory, they thought that providing participants with experience was still important. Hence, play was repeated for 100 rounds. They kept track of choices, provided feedback at the end of each round, and recorded each participant's comments after their choices were entered. Table 1 reproduces the record for the first nine rounds of play.

Both players defecting took place in only 14 of the 100 rounds. Observing this behavior, Flood and Dresher reached out to Nash, whose reply is particularly useful to introduce the challenges that repeated-game experiments face. The first comment (footnote 11 of Flood (1958)) articulates the idea that the Flood-Dresher experiment is really a test of a repeated game, not a one-shot game:

> The flaw in this experiment as a test of equilibrium point theory is that the experiment really amounts to having the players play one large multimove game. One cannot just as well think of the thing as a sequence of independent games as one can in zero-sum cases. There is too much interaction, which is obvious in the results of the experiment.

Tests of one-shot games now typically deal with this concern by having multiple repetitions of the task, but with random re-matching of participants. Repetition is necessary because the stage games are an abstraction that participants are usually not familiar with. Facing the same task many times helps participants gain experience with the environment.

[1] Most of our examples are drawn from our own work simply because we are more familiar with our papers. However, there are many other authors that have contributed to the experimental literature on repeated games. Although we will mention some of them, this is in no way an exhaustive review of the work on this topic and many excellent papers are not mentioned.

Table 1: First nine rounds (R) of play between Armen Alchien (AA) and John D Williams (JW) in the Flood-Dresher experiment (Flood 1958)

R	AA	JW	AA Comments	JW Comments
1	D	C	JW will play D-sure win.	Hope he's bright
			Hence if I play C-I lose.	
2	D	C	What is he doing?!!	He isn't, but maybe he'll wise up.
3	D	D	Trying mixed?	Ok, dope.
4	D	D	Has he settled on D?	Ok, dope.
5	C	D	Perverse!	It isn't the best of all .
				possible worlds.
6	D	C	I'm sticking to D since he will	Oh no! Guess I'll have to
			mix for at least 4 more times.	give him another chance.
7	D	C		Cagey, isn't he? Well...
8	D	D		In time he could learn,
				but not in ten moves or so.
9	D	D	If I mix occasionally, he will switch-	
			but why will he ever switch from D?	

How should the experimental design allow for experience when the aim is to study behavior in a repeated game? Note that if we consider the Dresher-Flood experiment as an experimental study of a repeated game (as Nash did), it fails to provide participants with appropriate experience, because, despite the fact that the stage game was repeated for 100 rounds, the participants only played *one* repeated game. Not providing experience in experimental studies of repeated games had an important effect on what early researchers concluded about the relevance of dynamic incentives for behavior. In this chapter, we will review how the literature has overcome this challenge.

The subsequent comment from Nash anticipates the challenges associated with the *discovery of strategies* and, in addition, with the implementation of repeated games more broadly:

> Viewing it as a multimove game a strategy is a complete program of action, including reactions to what the other player has done. In this view it is still true the only real absolute equilibrium point is for Armen Alchien (AA) always to play D and John Williams (JW) always D. However, the strategies: AA plays C 'til JW plays D, then D everafter, JW plays C 'til AA plays D, then D everafter are very nearly at equilibrium and in a game with an indeterminate stop point or an infinite game with interest on utility it *is* an equilibrium point.

In other words, Nash describes what we now know as a grim-trigger strategy. His description highlights a key methodological difference between laboratory tests of one-shot and repeated games. In one-shot interactions, instructions describing the action space, by definition, also specify to the participants the set of pure strategies. In a typical experiment on a repeated game, the instructions would similarly describe the action space for the stage game, but the set of pure strategies (which can be infinite) would not be specified. That is, by design, it remains up to the participants to discover what strategies they may want to consider and use. In other words, in contrast to experiments on one-shot games, a key interest of experiments on repeated games is to uncover the strategies participants are naturally drawn to and to study how this is a function of the parameters of the game. In this chapter, we will review methods the literature has developed to elicit strategies directly while minimizing interference with the natural evolution of play.

Finally, the comment also presents the idea of an indeterminate stop point and discounting. The literature did not develop methods to implement infinitely repeated games in the laboratory until the late 1970s. There are also associated considerations with providing proper incentives that have been formally considered only recently. In this chapter, we will review different approaches to implementing infinitely repeated games in the laboratory and discuss incentive schemes that can be paired with such methods.

The chapter is organized around three main methodological challenges. We start with the last challenge we introduced, the implementation of infinitely repeated games, because it eases the discussion of the remaining topics. Subsequently, we document the importance of experience and related considerations. The last portion of the chapter focuses on methods developed to elicit strategies. We provide examples on how methodological advancements in these areas were consequential in shaping our understanding of behavior in repeated games. We also highlight how these methods can be useful beyond the study of repeated games.

2 Termination Methods and Payments

Theoretical predictions for repeated games can vary dramatically depending on the time horizon. If the number of interactions T is finite, the set of subgame perfect Nash equilibria can be derived by applying backward induction. An implication of this is that cooperation cannot be supported in a finitely repeated PD regardless of the time horizon.[2] In an infinitely repeated PD, by contrast, folk theorems describe a large set of outcomes (including full cooperation) that can be supported in equilibrium for sufficiently patient players. The stark theoretical contrast between these environments is one of the motivations for empirically studying behavior with both finite and infinite T.[3]

Implementing a repeated game with finite T is relatively straightforward. By contrast, implementing in the laboratory a theoretical environment where $T = \infty$ is not as obvious. The first approach used in the literature was to implement a very large T. The intuition for why this might work is, in fact, illustrated in Nash's private communication to Dresher and Flood (see Introduction). There, Nash introduces the grim-trigger strategy as behavior that can support cooperation in an infinitely repeated version of the PD. He follows that paragraph with the following.

> Since 100 trials are so long that the Hangman's Paradox cannot possibly be well reasoned through on it, it's fairly clear that one should expect an approximation to this behavior which is most appropriate for indeterminate end games with a little flurry of aggressiveness at the end and perhaps a few sallies, to test the opponent's mettle during

[2] Beyond the standard model we have in mind here, there is important theoretical work that has shown how cooperation can be rationalized in this case, most famously Kreps et al. (1982).

[3] We will revisit the contrast between these two environments in more detail at the beginning of our section on Experience.

the game.[4]

However, as Nash points out, it is reasonable to expect backward induction forces to play a role even for large T, particularly in rounds close to T. As discussed further in section 3, Embrey et al. (2018) and Aoyagi et al. (2024) provide clear evidence that, independent of the specific game, cooperation breaks down in later rounds of a finitely repeated PD and that this breakdown is anticipated in earlier rounds, even among pairs who are on a cooperative path.[5] This is important because it shows that an experimental design involving finite T fails to provide an adequate setting to test theory with infinite T.[6]

Below we describe three alternative approaches to implement an infinite time horizon in the laboratory. Most commonly used methods exploit the fact that infinitely repeated games with payoff discounting are theoretically isomorphic to randomly terminated repeated games without payoff discounting. Such methods, thus, introduce uncertainty with respect to when a repeated interaction ends. The presence of uncertainty opens up the possibility for risk attitudes to play a role. In our last portion of this section we discuss how this can be taken into account when designing payment schemes.

2.1 Termination Methods

Observational studies of time preferences involve multiple choices over time that are separated from each other in a clear and meaningful way (e.g. one day or week apart), allowing for questions to be asked about how participants trade off consumption today with consumption in the future. Meanwhile, most of the experimental literature on repeated games takes a different approach.[7]

In this literature, time preferences are exogenously induced by the experimenter.[8] Note that this is another reason why using a large but finite T is problematic for implementing an infinitely

[4] With the Hangman's Paradox Nash alludes to the idea we now know as backward induction.

[5] This comparison is first made directly in Dal Bó (2005). In the games he studies, cooperation rates are lower in all rounds of the finite game as compared to the indefinite game.

[6] Aoyagi et al. (2024) further show that even at the beginning of an interaction subjects consider finite and indefinite PD differently.

[7] An exception is Kim (2023). Like other designs, participants make all their choices during a laboratory session. However, unlike other designs, the earnings associated with different rounds of the repeated game are paid with time delays. For example, round 1 might be paid a week later (via Venmo), with each subsequent round paid one week apart. This creates a meaningful delay in payments, allowing researchers to study how time preferences affect decision-making.

[8] The time between choices in an experiment is short enough to ignore the subject's own discounting.

repeated game: time preferences, a crucial element from the point of view of theory, has no role or parallel in such an implementation.

To fix ideas, let x_t be the payment that a participant obtains in round t of an infinitely repeated game in which one unit of payment received in period $t+1$ is worth $\delta \in [0,1)$ units in period t. The parameter δ is known as the discount factor (sometimes described as the *shadow of the future*) and discounted payments in an infinitely repeated game are equal to:

$$x_1 + \delta x_2 + \delta^2 x_3 + ... = \sum_{t=1}^{\infty} \delta^{t-1} x_t \tag{1}$$

That is, participants would make a choice in a similar fashion as in the Flood-Dresher experiment, one after the other. The difference between rounds t and $t+1$ is that payments in $t+1$ are multiplied by a pre-determined parameter δ. In this way, it is possible to study choices in a setting where the key parameters are induced and the focus of attention is on what participants choose to do given such primitives.

The first method that implemented such an infinite time horizon in the laboratory used random termination (RT) and is due to Roth & Murnighan (1978). RT works by changing the meaning of δ. Instead of δ capturing the value in t of receiving one unit of payment in $t+1$, it captures the probability that $t+1$ will take place. For instance, assume that the aim is to implement $\delta = 50\%$. In the standard interpretation, this means that a payment in period $t+1$ is worth 50% of that in t. In the interpretation with indefinite time horizon, it captures the idea that period $t+1$ takes place with probability 50% conditional on playing period t. With the complementary probability, the game is over and there is no future. It is straightforward to see that a game where at the end of each period t, there is a probability δ that players face the same stage game in period $t+1$ implements (under risk neutrality) the same discounted payoffs as described above.

The following passage from Dal Bó & Fréchette (2019)'s instructions (describing $\delta = 50\%$) provides an example of how random termination of a repeated game, referred to as *match* below, is commonly explained to participants:

> After each round, there is a 50% probability that the match will continue for at least another round. This is as if we would roll a four sided die and end if 1, or 2 come up and continue if 3, or 4 come up. So, for instance, if you are in round 2, the probability there will be a third round is 50% and if you are in round 9, the probability there will be another round is also 50%.

From now on we will use the term *supergame* to refer to an implementation of a repeated game, i.e. the series of rounds that participants play in a fixed group until the game randomly terminates; what is referred to as a *match* in the instructions quoted above. When δ is implemented with random termination, the expected number of rounds in a supergame is $\frac{1}{1-\delta}$.

By definition, RT creates uncertainty over supergame lengths. In an experiment that involves a few sessions per treatment, it is in principle possible that the random realization of supergame lengths in each treatment impact results.[9] For instance, suppose that the length of the first supergame is two, one, and three rounds respectively in the three sessions conducted for a control treatment. Assume the first supergame length changes to five, three, and seven rounds respectively in the three sessions conducted for the treatment. Some metrics of comparison between control and treatment may be affected by the differences in the realization of supergame lengths. One approach in the literature to deal with confounds (e.g., Fréchette & Yuksel (2017)) is to fix seeds (that determine realization of supergame lengths) in sessions across treatments. Specifically, suppose that a design involves Y sessions per treatment. A seed is randomly selected for each session y, so that, for instance, session 3 of every treatment will have supergames of exactly the same length.[10]

With RT, the distribution of supergame lengths is geometric, and the median will be lower than the mean. This can create limitations, particularly in games where observing a large number of rounds is particularly useful. The literature has relied on two alternative methods to RT to collect more observations per supergame without using high δ.

Discount followed by random termination (D+RT), introduced by Sabater-Grande & Georgantzis (2002) and also used in Cabral, Ozbay & Schotter (2014) and Vespa (2020), divides a supergame into two parts. The first part consists of T^{D+RT} rounds that are played with certainty, but with actual discounting of payoffs given by δ. From round $T^{D+RT}+1$ onwards, there is no longer payoff discounting but the game ends with probability δ. In other words, in the second part, there is random termination. This method implements supergames of length that are T^{D+RT} rounds larger than RT.

The second method is known as block random termination (BRT) and was introduced by Fréchette & Yuksel (2017), and has since been used in Wilson & Wu (2017), Vespa and Wilson (2017, 2019), and Aoyagi et al. (2024) among others. Participants play as in RT, but in blocks of

[9]It is common to have everyone in a session share the same random termination to facilitate re-matching of participants between supergames.

[10]The seeds associated with each session should be randomly selected (or at least not selected after looking at the random number it generates). One way to do this is for the seed of the sessions for one treatment to be set by the internal clock of the computer when the sessions start. This seed is saved and then re-used in the sessions of the other treatments.

Table 2: Stage-game payoffs for the row-player in a symmetric PD, where each player decides whether to cooperate (C) or defect (D). (a) Payoffs are such that $T > R > P > S$. (b) Normalization of stage-game payoffs: Subtracting P from all payoffs and dividing by $R - P$ to set payoff associated with (C,C) and (D, D) to 1 and 0, respectively. (c) Summary of normalized payoffs.

	C	D
C	R	S
D	T	P

(a) Payoffs

	C	D
C	$\frac{R-P}{R-P} = 1$	$\frac{S-P}{R-P} = -l$
D	$\frac{T-P}{R-P} = 1 + g$	$\frac{P-P}{R-P} = 0$

(b) Normalization

	C	D
C	1	$-l$
D	$1 + g$	0

(c) Normalized Payoffs

a pre-announced fixed number of rounds (T^{BRT}). Within a block, participants receive no feedback about whether or not the supergame has continued until that round, but they make choices that will be payoff-relevant contingent on the supergame actually having reached that point. Once the end of the block is reached, participants are told whether the supergame ended within that block and, if so, in what round; otherwise they are told that the match has not ended yet, and start a new block. Subjects are only paid for *valid* rounds, and all decisions for subsequent rounds within the last block are void. In most recent implementations of this method (including in the papers citated above), the experimental design involves only one block followed by random termination on a round-by-round basis thereafter.[11]

A main focus of the experimental literature on the infinitely repeated PD has been to study the extent to which people use dynamic incentives to sustain cooperation. Cooper & Kühn (2014) proposes an alternative method for implementing an infinitely repeated PD in the laboratory. The method, which we will refer to as discounting followed by coordination (D+C), uses a coordination game, instead of random termination, to allow for dynamic incentives. First, participants play a PD for T^{D+C} rounds with discounting. Subsequently, participants play a one-shot coordination game. The coordination game has two equilibria in pure strategies. The payoffs associated with these equilibria correspond to the largest and smallest achievable symmetric payoffs in the continuation game of the infinitely repeated game. Namely, the coordination game payoffs capture the reduced continuation game (in the infinitely repeated PD) where players can only choose between two strategies: the grim-trigger strategy that starts by selecting C and moves to always defect after any deviation and the always-defect strategy.[12]

[11]Note that unlike in the strategy method (Selten 1967), there are realizations where all choices made using the block random method are payoff relevant.

[12]Payoffs in the coordination game are discounted by δ.

An advantage of the D+C method is that it does not need to introduce uncertainty to implement an infinite time horizon. A limitation is that it constrains the coordination portion of the game to two strategies of the infinitely repeated game. A second shortcoming is that it is not straightforward to extend the method beyond the repeated PD.

Different methods may be more appropriate for different games, but it is useful to know whether there are systematic differences introduced by a given method for a fixed environment. Fréchette & Yuksel (2017) conduct a series of lab experiments using the infinitely repeated PD. Their design involves a between-subjects comparison across the four methods and a within-subjects manipulation of the stage-game payoffs. As a reference, Table 2a presents generic stage-game PD payoffs from the perspective of the row player. The normalization of payoffs in panel (b) is a common way to reduce the number of parameters used to describe and compare PD's. In the first part of each session, the parameter for R is set to be relatively high, while it is lower in the second part. Results can be summarized as follows: (1) All four methods generate behavior consistent with the comparative-static prediction of more cooperation with high R; (2) With D+C, as participants gain experience, their response to the coordination game becomes independent of the history of play. Hence, the way subjects perceive this implementation fails to generate dynamic incentives. (3) Within the three methods that implement an infinite time horizon as an indefinite time horizon, the highest levels of cooperation are observed with R+T. However, D+RT is the method that generates the most stable cooperation rates. Furthermore, D+RT and BRT are significantly less affected by past experiences within a session. Taken together, the evidence in Fréchette & Yuksel (2017) indicate that all three methods successfully induce dynamics incentives and suggest that D+RT and BRT can be more desirable than RT when important variations in the realized length of the supergames are expected and the samples are small.

Being able to observe more rounds of play—by using D+RT or BRT instead of RT—can be extremely useful in some settings. To illustrate this, we revisit the data collected using BRT in Fréchette & Yuksel (2017), which guaranteed at least four rounds of play to be observed in an indefinitely repeated game with a discount factor of $\delta = 0.75$. If the implementation had involved RT, only a subset of these rounds would have been observed. We ask whether observing only this subset would have changed certain conclusions that can be drawn from this data. To provide an example of this, we compute the share of subjects whose behavior is consistent with Win-Stay-Lose-Shift, a strategy that starts out cooperating, switches to defection after observing a defection, but reverts back to cooperation after a mutual defection. When all observed rounds are considered: 27 and 26 percent of choices are consistent with Win-Stay-Lose-Shift in the first and last three supergames, respectively.[13] However, if we consider only the subset of rounds that would have been observed under RT, this share increases to 54 and 42 percent for the first and last three supergames, respectively. This is partly because observing a response to defection is less likely when fewer rounds of data are observed.

As an alternative example, consider a simple version of the dynamic common-pool problem. In each period t, each of two players decides a share to extract from a common resource (e.g. fish in a lake). At the end of the period, the remaining stock reproduces at some known rate, which means that the total stock s_t changes over time. This is known in the literature as a *dynamic game* with an underlying state variable s_t. Once s_t is known, it pins down a specific stage game, but the payoffs in the stage came can change over time. A version of this game is studied in Vespa (2020). Strategies in dynamic games can condition on past history (as in repeated games) *and* on the state variable s_t. To understand what participants condition their behavior on, it is particularly useful to observe their choices for many values of s_t, which is feasible only if there are many rounds of play in each supergame. Implementing RT would require relatively high discount rates to achieve this. D+RT and BRT provide alternative ways of achieving this for a fixed discount rate.

Visualizing the repeated game

A repeated game can be theoretically represented in extensive form with decision nodes and information sets. In experiments, however, participants are typically presented with the stage game on the screen, where they indicate their choice for the round. As we mentioned in the Introduction, and as we will go into further detail in Section 4, it is up to participants to visualize the game beyond the current round. In the case of repeated games this task is facilitated by the fact that rounds in the future look the same as in the current one: the stage game is fixed. In games where the stage game can change (as in dynamic or stochastic games), however, picturing the full game is more challenging.

For instance, consider the dynamic common-pool problem introduced earlier. To make the example more concrete, suppose that there are only two actions available: extract 10% or 25% of the resource at round t. In this environment, stage-game payoffs change endogenously. That is, once the stock of the resource is established in a given round, players can compute the payoff they would get for each action. In other words, depending on the outcome in period t participants would

[13] Note that mutual cooperation by both players in all rounds is consistent with Win-Stay-Lose-Shift as well as a numerous other strategies (if both players is playing one of them) such as Always Cooperate, Tit-for-Tat and Grim Trigger.

face a different stage game in period $t+1$. If joint extraction is low in t, s_{t+1} will be lower than if joint extraction is high in t. Thus, visualizing the possible future consequences of a choice in t in an environment with a large set of possible states is much more challenging relative to a standard repeated game where the stage game remains unchanged.

To deal with this difficulty, Vespa (2020) uses a visual tool to anticipate hypothetical futures. Specifically, in any round, before participants make a choice, the interface allows them to click on any of the possible stage-game outcomes and shows them the stage game that would subsequently follow from this outcome. Participants can repeat this exercise to build further a possible hypothetical future path, where the interface shows them sequence of stage-games that would be visited along this path. Participants can also see on their screens the payoffs that would result from such a path. They can re-start and repeat the exercise to build alternative future paths before deciding on a choice.[14]

2.2 Payment Methods

We have described three different alternatives that can be used to successfully implement infinite time horizon in the laboratory. However, all three methods introduce uncertainty over supergame lengths, and risk preferences may affect choices. Here, in light of these considerations, we discuss different ways to pay participants and these different implementations can help assess the extent to which risk preferences play a major role.

In a typical laboratory session that implements a repeated game, participants can easily make more than 100 decisions. Decisions are broken up in several supergames. Once a supergame terminates, the participant is randomly matched to another and starts a new supergame. How should participants be incentivized? Pay every choice? A subset of choices?

Azrieli et al. (2018) provide a theoretical framework to consider the implications of different incentive schemes in experiments where participants make multiple decisions. They discuss conditions under which it is optimal for the experimenter to randomly select one choice for payment rather than pay for every decision.

To simplify the presentation, let us discuss these options in the context of repeated games for a hypothetical session that consists of one supergame. Paying for every round is the most standard approach in the literature and it amounts to the payoff in Equation (1). Suppose, alternatively,

[14]This visualisation tool also allows the experimenter to record the hypothetical futures that participants considered before making a choice.

that we pay by randomly selecting one round. For each possible x_t in Equation (1), it is possible to compute the probability that such x_t is selected.

Sherstyuk, Tarui & Saijo (2013) show that in repeated games, paying for one randomly selected round induces players to discount the future more heavily than paying all rounds. Intuitively, this results from the fact that payoffs for high t can only be selected if the game lasts long enough, but payoffs for early rounds, say for round 2, are potentially considered for payment for a larger set of possible supergame lenghts. In other words, they predict that randomly paying for one round will tend to result in more myopic choices relative to paying all rounds, where myopia translates in the context of the PD to less cooperation. Indeed, their results support this prediction. For instance, the cooperation rate in their experiment when all rounds are paid is 55 percent, but it drops to a significantly lower 36 percent in a treatment that pays one randomly selected round. An additional treatment further reinforces the conclusion that randomly selecting one round distorts behavior in a myopic direction. In this treatment, subjects are paid for the outcome in the last round only. That is, if the game randomly terminates in period $T^\star$, then subjects are paid $x_{T^\star}$. Sherstyuk et al. (2013) show that the last-round scheme uses only one round but induces the same incentives as paying all rounds (for risk neutral agents). In their data, overall cooperation if participants are paid based on the last round (53%) is not different relative to paying all rounds (55%), but is significantly higher than paying one random round (36%).

When there is uncertainty about the supergame length, risk attitudes can also impact how participants trade off a sure payment in the current round with an uncertain payment in a future round. That is, participants may make choices in the risky *indefinitely* repeated game that differ from their choices in the riskless *infinitely* repeated game. For risk-averse participants, for instance, cooperation may not be part of an equilibrium in the indefinitely repeated game, even when it can be supported in the corresponding infinitely repeated game. Chandrasekhar & Xandri (2023) show that, under some assumptions on the utility function such as time separability, paying the last (randomly occurring) round is the only payment method that does not distort incentives regardless of the risk preferences of the agents. Nonetheless, there isn't so far much evidence to suggests that risk preferences play a major role in this environment (at least for the payoff configurations that have been commonly studied). The consistency of behavior between the last-round and the all-rounds treatments in Sherstyuk et al. (2013) is also supportive of this. However, paying all rounds of a supergame is simple to explain and remains the most popular alternative.

Our discussion so far has focused on payment for a given supergame but has not addressed how to deal with the fact that sessions typically have multiple supergames. The Sherstyuk et al. (2013)

design indeed involves several supergames and, while the payment for each supergame changes across treatments, session payments are determined by adding payoffs across supergames. To our knowledge there is no direct test of whether paying for one or for several supergames impacts behavior, even though theoretically the assumptions needed to pay for more choices appear more demanding than paying for a single choice (Azrieli, Chambers & Healy (2018)). Boczoń et al. (2024), however, provides some evidence suggesting that this choice is of no important consequence. In that paper, payments are decided by the last round of one randomly selected supergame. By contrast, in most repeated PD experiments, including those from the meta-study of Dal Bó & Fréchette (2018), payments are determined by the sum of payoffs from all rounds of all supergames. Using those data, one can compute counterfactual cooperation rates for the payoff and continuation probability parameters of Boczoń et al. (2024). Comparing the two reveals no visible difference in cooperation rates.

2.3 Choosing between termination and payment methods

When considering a design that involves an unknown time horizon there are several choices to make. Often, the choice of a termination method and the choice of a payment method are made in combination. For many of these options, there is no evidence that a specific implementation dominates others. The choice thus depends on the question and the environment under study. Before exploring these trade-offs, we also note that the methods described here in the context of repeated games also apply to other types of games, in particular environments with an infinite time horizon with discounting, such as bargaining or network games.

As background for the discussion, imagine a session where participants play several supergames. We consider two possibilities for payoffs: either subjects are paid for all rounds in all supergames or for the last choice of one randomly selected supergame. The most common implementation is the former, in which every choice has a payoff consequence. In practice, it often involves introducing a conversion between *points* subjects see in the stage-game and the unit of currency in which subjects are paid. Since paying every round of every supergame involves adding up several payoffs, in order for the final amount in currency to be reasonable, either the actual payoffs in each stage game are quite small, or stage-game payoffs have larger numbers that are eventually converted into currency at a certain conversion rate. The most common approach in the literature uses the latter.[15] A conversion rate needs to be determined before treatments are conducted and setting it in a manner

[15] Note that Kagel & Schley (2013) show that behavior is not affected by changing the units in which the stage game is presented (and the exchange rate such that actual dollar payoffs remain unchanged).

that is consistent with the experimenter's budget often involves a prediction about how participants will behave. Since this prediction carries some uncertainty, total expenditures can be challenging to predict.

At the other extreme, consider the case of paying the last round of one randomly selected supergame. On the one hand, describing this payment method is less natural as only one of the several choices that participants make will count. Yet an advantage of this method is that stage-game payoffs can be directly expressed in currency and that the total experimenter's expenditures (or its range) is easier to predict. In addition, paying the last round is a method that seems particularly useful in environments where risk attitudes can be expected to play an important role.

The choice of termination method also involves evaluating trade-offs. Payoff discounting followed by random termination (D+RT) and block random termination (BRT) both require instructions beyond those for the standard random termination (RT). But for a more detailed comparison of the methods, notice first that D+RT and BRT essentially disconnect δ from the number of choices that participants make. Specifically, in RT the expected number of rounds $E(r) = \frac{1}{1-\delta}$ and the number of choices that participants are expected to make in a supergame coincides. Fixing δ, D+RT will have longer supergames with the expected number of choices equal to $E(r)$ plus the rounds that are played with discounting. In BRT subjects will also make more choices per supergame, as they will likely make choices in rounds that will not be counted for payment. In other words, for a fixed δ, the trade-off is clear: D+RT and BRT provide longer supergames than RT, but if the amount of time devoted to a session is fixed, this means that in RT participants will experience more supergames.

For a different comparison across termination methods, suppose that the experimenter needs relatively long supergames. For instance, assume they need supergames with about 10 rounds of play. Using D+RT, this can be done by playing six rounds for certain followed by a randomly determined number of rounds using a discount factor of 0.75. Alternatively, one could use RT using a discount factor of 0.9, which also implements supergames that will last for 10 rounds in expectation. The trade-off is that using random termination with a $\delta = 0.90$ implies much more variation in supergame length, and some potentially very long supergames, while using D+RT with a smaller δ will have more supergames with similar length close to the target average.

Finally, there is a related choice that we have only tangentially discussed: the number of supergames in a session. There are essentially two approaches: either the number of supergames is set or the minimum time over which subjects will play is fixed (the first supergame that terminates after x minutes of play marks the end of the session).[16] Both of these methods have been used in multiple papers and they each have advantages and disadvantages. Fixing the minimum amount of time reduces variation in the amount of time spent in the laboratory, but introduces complications when analyzing data since different sessions will have different number of supergames. Fixing the number of supergames eliminates this aspect, but some sessions may last substantially longer than others (which can make setting appropriate remuneration difficult).[17] Past experience (which of course depends on details of implementation and instructions) suggests that for a standard repeated PD experiment, considering time for instructions and payments, participants make somewhere between 100 and 150 choices in 90 minutes. From this and the choice of δ, one can work out the number of supergames the session can be expected to have. For environments that are less explored, the number of supergames that can be conducted in a fixed amount of time can be more difficult to predict.

3 Experience

3.1 Experience as Rounds Within a Supergame

Roth & Murnighan (1978) offers the first controlled indefinite repeated game in the laboratory. Between Flood (1958) and Roth & Murnighan (1978), experimental studies of the indefinitely repeated PD were predominantly conducted by psychologists and use long (sometimes hundreds) repetitions of the stage game without announcing the number of repetitions (see for instance Rapoport & Chammah (1965)). Given the crucial role of having a known number of repetitions versus an infinite horizon (as discussed in the previous section), there are substantial challenges to interpreting results from such implementations as providing tests of economic theory. Following Roth & Murnighan (1978), in the 80's and 90's, more experimental papers on finitely and indefinitely repeated games (mostly PD's) were published in economic journals (Murnighan & Roth (1983), Selten & Stoecker (1986), Andreoni & Miller (1993), Feinberg & Husted (1993), Kahn & Murnighan (1993), Palfrey & Rosenthal (1994), Cooper et al. (1996)).

Looking back, experiments on finitely repeated games during this time were primarily concerned

[16] A third method would be a minimum number of total rounds (summing over all supergames), but as far as we know this has not been used.

[17] When this method is combined with all-rounds payment, sessions in which there is a participant who takes more time to make choices will be shorter and average payoffs will be lower. If this is anticipated as a potential complication, it suggests that when using a fixed total time of play it may be useful to also only pay the last round of a single supergame.

with shedding light on why participants cooperated in this environment despite the standard theoretical predictions of joint defection in all rounds. Hence, many of the early papers on the finitely repeated PD were designed to study whether observed cooperation in this game could be explained via alternative models that allow for incomplete information about players' types (Kreps et al. 1982). These experiments typically involved multiple supergames per session. One common finding is a pattern of early cooperation followed by a collapse of cooperation by the end of the game.

The few early papers on the indefinitely repeated PD, by contrast, all involved a single supergame. Those report positive, but very modest effects of increasing δ on cooperation. Thus, results at the time suggested cooperation to be behaviorally difficult to achieve in this game, even for sufficiently patient agents for whom full cooperation could easily be theoretically sustained.

As a whole, these early results seem difficult to reconcile: too much cooperation in the finitely repeated PD and comparatively little cooperation in the indefinitely repeated PD with high δ. Add to this the collapse of cooperation in the last rounds of the finitely repeated PD and the non-trivial amount of cooperation in the indefinitely repeated PD with low δ, and the results are not only difficult to make sense of together, they are also difficult to reconcile with standard theoretical predictions. These puzzling observations point toward an important challenge in conducting experiments on repeated games: evolution of behavior with experience, which is the focus of this section.

Because we want our experiments to isolate specific forces, we aim to abstract away from context that is highly relevant in naturally occurring settings. Thus, by design, our experiments present subjects with environments they do not recognize. A carefully designed experiment will have taken care to write instructions as clearly as possible. Nonetheless, reading such an abstract description of a decision environment makes it challenging to fully grasp the strategic tensions. In a naturally-occurring setting, people may have cues that they have learned to react to that are now absent in the synthetic version created in the laboratory. This creates an important potential role for experience (and thus feedback).[18]

[18] Note that we do not claim experience to be a common feature in most applications of interest outside of the laboratory. Instead, we argue that allowing for subjects to gain experience in the laboratory substitutes for other contextual cues that help players to understand the strategic interaction better.

3.2 Experience as Repetitions of the Supergame

What do players in repeated games learn with experience? One important feature of repeated games is that the set of strategies is distinct from the set of actions. Even a simple indefinitely repeated PD with only two choices and four possible outcomes in a round has an infinite number of strategies. Hence, even with the most carefully crafted instructions that describe the stage game and the continuation probability in detail, participants need to discover the strategies available to them on their own.[19] At the same time, participants in a repeated game experiment need to form expectations about the strategies of others: Even if they do not think in terms of strategies, they might form beliefs about the reactions of others to specific outcomes. In naturally occurring situations involving repeated interactions, people are often familiar with the environment and the context helps them form reasonable expectations about the behavior of others and how to best respond. In the laboratory, in the absence of context, experience plays a key role helping participants form a sophisticated understanding of their strategic environment. From this perspective, the appropriate notion of experience in a repeated game is not the play of multiple rounds in a supergame, but rather the play of multiple supergames, since one needs to discover strategies.

As we mentioned already, the first experiments on the indefinitely repeated PD implemented a single supergame. This first changed with Dal Bó (2005), who conducted an experiment where subjects played an indefinitely repeated PD with the same parameters six to ten times. Subsequently, Dreber et al. (2008), Aoyagi & Fréchette (2009), Duffy & Ochs (2009), and Blonski, Ockenfels & Spagnolo (2011) each implemented indefinitely repeated PD games where subjects experienced multiple supergames, between eight and 27 depending on the paper and treatment. Then, Dal Bó & Fréchette (2011) conducted an experiment where subjects played from 33 to 77 supergames across different parameter combinations, which helped establish how behavior changes with experience. For instance, Figure 1 displays the evolution of round-one cooperation over supergames in two of their treatments, where treatments are characterized in terms of payoff parameters g and l, described in the normalization presented in Table 2c, and the discount rate, δ.[20] As the figure makes it clear, initially round-one cooperation rates are almost the same in those two treatments (and there is overlap across sessions). However, by supergame 35, most subjects cooperate in the

[19] By *discovering strategies* here we do not mean that the participants explicitly think in terms of a complete plan of action that indicates a choice for every information set. We simply mean that because it is not described in the instructions, it is up to participant to consider the different ways in which they can condition their choices on the observed history of play.

[20] In terms of analysis, round-one cooperation rates are useful, as they provide a measure of participant's initial intentions that are not conflated with the choices of the person they are matched with.

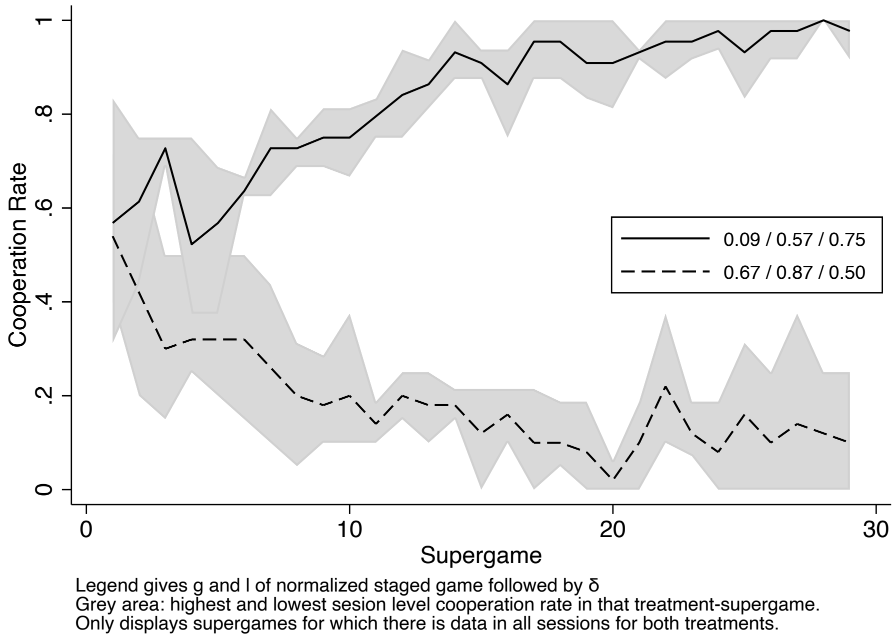

Figure 1: Round One Cooperation

treatment with $\delta = \frac{3}{4}$, while less than 20% do in the treatment with $\delta = \frac{1}{2}$. This observation, that experience is important and affects behavior in different ways depending on the treatment parameters, is carefully demonstrated in Dal Bó & Fréchette (2018). That paper puts together a data set of all experiments involving one-shot and indefinite PD experiments with perfect monitoring available up to that point. This consists of 44 paper-treatment, 141 sessions, 2,415 subjects, making a total of 157,170 choices. Using that data set, they clearly show that the pattern represented in Figure 1 is not specific to Dal Bó & Fréchette (2011), but rather a general finding (and that it extends beyond round one).

Beyond changes in whether subjects become more or less cooperative, behavior also evolves in other important ways with experience. In particular, for some parameter combinations, subjects clearly become more *responsive* with experience. We define responsiveness as the measure of the difference in probability of cooperation following a cooperate and defect choice by the other player. Figure 2 displays cooperation rates following both cooperation or defection by the other and the difference between these two. The left panel reproduces a figure from the first paper to have discussed this evolution (Fréchette & Yuksel 2017). As can be seen, responsiveness increases from slightly below 50% to about 75% over the course of 12 supergames. The second paper to note such an evolution is Aoyagi et al. (2019). In that case, as represented in the middle panel of Figure 1, responsiveness increases by about 40 percentage points over nine supergames. The last panel of

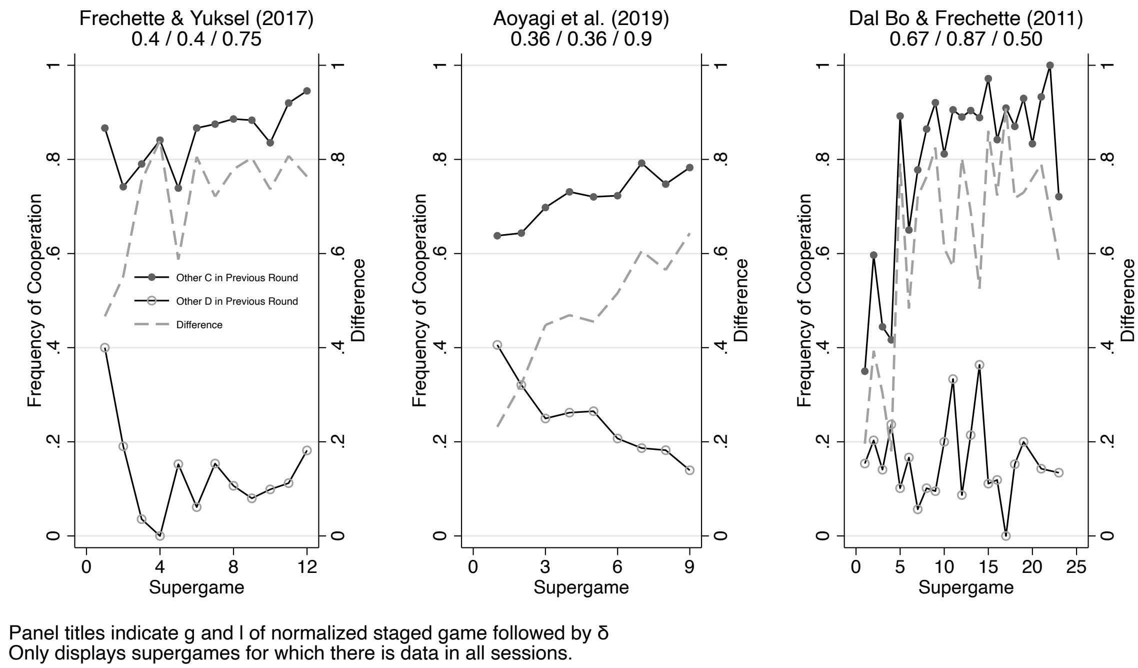

Figure 2: Responsiveness

Figure 1 presents data from one of the treatments of Dal Bó & Fréchette (2011). That treatment also shows an increase of about 40 percentage points in the responsiveness measures. Such large increases are not always observed, but they certainly occur with some parameter configurations. Although responsiveness abstracts from many potentially relevant features of a supergame history (one's own past choices or choices of the other beyond the immediate previous round), it is directly informative about a key aspect of the strategies participants are using. Note, for example, that strategies such as Tit for Tat (TFT) or grim are consistent with a responsiveness measure equal to one.[21] By contrast, other strategies, such as always defects or always cooperate are consistent with a responsiveness measure of zero.

In their paper, Proto et al. (2019) treat a defection following joint cooperation as an error. They show (see also Proto et al. (2022)) such errors to be less likely among participants with high intelligence scores. The clearest case is a defection in round two, when both subjects cooperated in round one. The left panel of Figure 3, using the data put together in Dal Bó & Fréchette (2018), shows such errors to decline with experience. In all three cases, corresponding to treatments with $\delta = 0.9$, errors in supergame one are between 14% and 23%, but fall to below 5% with experience.

[21] The grim-trigger strategy cooperates until there is a defection and defects from then on. Tit for tat starts cooperating and subsequently plays the action selected by the other in the previous period.

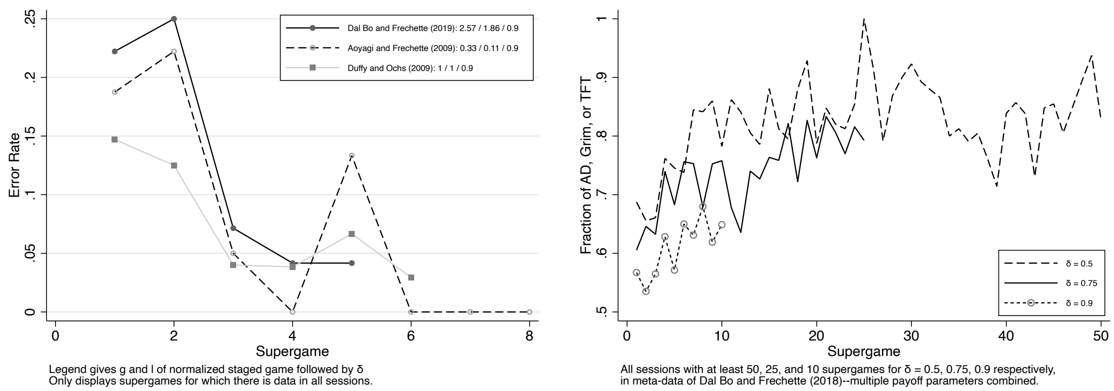

Figure 3: [Left Panel] Errors in Round Two; [Right Panel] Choices Consistent with AD, Grim, or TFT.

In other treatments, we sometimes see a decrease that is slower, while in others these errors are rare to begin with.

The changes in choices that occur with experience can be considered from the lens of strategies. Consider the three following strategies: always defect (AD), the grim-trigger strategy, and TFT. In the next section, we will discuss strategies and how to identify them in more detail, but here we will simply report whether *all* choices of a subject in a supergame are consistent with those strategies. The right panel of Figure 3 reports the fraction of subject-supergame observations consistent with one of these three strategies using the relevant treatments of the meta-data of Dal Bó & Fréchette (2018). To provide a meaningful test, the focus is on a subset of supergames where at least two rounds of play are observed. Note that 10 supergames with $\delta = 0.9$ involves about the same number of total rounds as 25 with $\delta = 0.75$ and 50 with $\delta = 0.5$. As the Figure makes clear, there are important increases in the fraction of play consistent with these strategies. In the $\delta = 0.9$ sessions, it increases by 8 percentage points between the first and tenth supergame. With $\delta = 0.75$, the increase is 19 percentage points between supergame one and 25. Finally, the increase is 14 percentage points when $\delta = 0.5$. One interesting trend not depicted in the Figure is a substantial increase in play of AD in treatments where cooperation cannot be supported in equilibrium (a subset of the $\delta = 0.5$ sessions with stage game parameters that are not conducive to cooperation). In these treatments, the fraction of subjects whose choices are consistent with AD is 55% in supergame one, but it climbs to 92% in supergame 50.

To summarize, experience matters for multiple facets of play. Subjects learn whether they want to adopt cooperative or defective strategies, as well as how to use dynamic incentives in their play. The observed changes are always in a direction that is sensible: play becomes more predictable, it is more reactive both to what is happening within a supergame and also to the parameters of the environment. But is this peculiar to indefinitely repeated games? Does it have to do with participants understanding the complexities of random termination? To explore these question, we return to the finitely repeated game.

As we mentioned earlier, unlike the indefinitely repeated PD, early finitely repeated PD experiments often involved multiple supergames. The typical (and surprising) result was an increase in average cooperation rate as subjects gained experience. But in the 2000's, both Dal Bó (2005) and Friedman & Oprea (2012) conducted experiments where they included finitely repeated PD treatments that served as control. Results from Dal Bó (2005) were that with experience, average cooperation *decreased*; while Friedman & Oprea (2012) found it to increase. Many aspects of these experiments—such as differences in procedures, instructions, interface, or subject pool—could potentially play a role in generating these conflicting results. The other possibility is that differences in behavior observed in these experiments are primarily explained by the specific parameters used in these studies.

To explore this further, Embrey et al. (2018) conducts four treatments of finitely repeated PD that vary payoffs and the number of rounds in a supergame, all within the same experimental paradigm and subject pool. They show that they can reproduce both the increase and decrease in cooperation over supergames previously observed in different papers. How can we make sense of these seemingly conflicting results? Partly to answer this question, Embrey et al. (2018) also assemble a data set of the previous finitely repeated PD experiments. Using this meta data (the previous experiments and their new treatments), they identify an aspect of behavior that changes with experience consistently in the same way across different environments covering a range of parameters. Early in a session, many subjects display a tendency to revert back to cooperation after observing a defection. Such behavior is inconsistent with the use of a threshold strategy. A threshold strategy is conditionally cooperative (following grim trigger) up to a certain *threshold* round and then switches to defection. Between 29% and 68% of subjects, depending on the treatment, follow a threshold strategy in the first supergame. Numbers in the sixty percent are never observed in supergames of more than four rounds. In the three experiments with supergames of ten rounds, the numbers are 29%, 30%, and 42%. However, in every treatment for which the authors could obtain the original data (thirteen treatments), that number is substantially higher in the last supergame of the experiment. In four of them, 90% or more of the subjects play a threshold strategy in the last supergame. In summary, the evidence shows that, with experience, behavior changes to become more consistent with the use of threshold strategies. The paper further demonstrates that the increase in the popularity of threshold strategies has two opposing effects on cooperation: (i) increase in cooperation rates in the early rounds of a supergame, and (ii) unravelling of cooperation in later rounds of a supergame. How experience impacts aggregate cooperation in the game depends on which force dominates. These opposing forces also explain how we can observe conflicting results on the impact of experience within the same experimental framework (as in Embrey et al. (2019)) as we vary the parameters.[22]

Taken together, the evidence suggests that subjects discover many of the relevant features of a repeated game with experience. For instance, participants in a repeated PD appear to learn the *value of cooperation*, which combines returns to cooperating in the current round with returns in the continuation game, with experience. In indefinite games, this is reflected in the fact that, the length of the supergame (which varies because of random termination) affects subsequent interactions. This has been observed first in the trust game (Engle-Warnick & Slonim 2006*b*), where longer supergames generate more trust and more reciprocity than shorter ones, and then in the PD game (Dal Bó & Fréchette 2011), where round one cooperation increases (decreases) after long (short) supergames. Participant's beliefs about the cooperativeness of others also appear to change depending on their past experiences. For instance, subjects are more likely to start cooperating if they interacted with someone who started by cooperating in the previous supergame. This is observed in both indefinitely repeated and finitely repeated PDs (see Dal Bó & Fréchette (2011) and Embrey et al. (2018), respectively). Note that although these forces create some path dependence in how behavior evolves over a session, these effects are small in magnitude relative to the effect of the underlying game parameters. Finally, participants also appear to learn how they can condition their actions on the history of play. Behavior evolves with experience to exhibit key features consistent with specific strategies: responsiveness increases and errors decrease in indefinitely repeated PDs and consistency with threshold strategies increases in finitely repeated PDs.

We also note that the need to provide participants with sufficient experience in repeated games limits the feasibility of conducting such experiments with online populations, as there are challenges to keeping such participants engaged in the long experiments necessary to provide this experience.

[22] One key feature of the new experiment in Embrey et al. (2019) is that it includes many more supergames (30) than previous experiments. As a consequence of that, the experimental results show how cooperation can first increase and then decrease with sufficient experience for a fixed set of parameters as the force described as (i) dominates in early supergames while (ii) dominates in later supergames.

3.3 Practical comments on conveying feedback

Because experience is important, one needs to carefully think about how feedback is provided. Unfortunately, we do not have much in the way of evidence with respect to what is useful feedback and how it should be presented. Given limited research on the topic, our recommendations on this matter are based on what we typically now do in our own repeated game experiments. One general comment is that more is not necessarily better, since information competes for attention. The two central aspects of any interface we use involve immediate feedback and history. Let us start with the simplest case: games with perfect monitoring, which capture the case where at the end of the round a participant can perfectly observe the choice of the other. First, the outcome and choices of both players in the stage game must be clear. This can be achieved by highlighting with distinct colors the choices and outcome associated with that round in the stage game matrix presented to subjects.[23] Second, we provide a history table on the screen that allows the subject to see what has happened in the previous rounds of the current supergame. The information provided in the table is: supergame, round, the subject's choice, and their partner's choice. The table lists the latest round first, and a scroll bar allows to see the complete history when the supergame is long. In addition, the subject can select to see in that same window the history of past supergames. The interface must also make clear when a supergame ends (and indicate that a new one starts).

In games with imperfect monitoring, more choices have to be made. These are games in which subjects' receive only noisy signals about their opponent's choice. During a supergame, a subject would be informed and reminded of the signals they received about their opponent's choice, but would not see their actual choices. Once a supergame is over, the experimenter can then inform the subject of their opponent's choices. Opinions on the desirability of that may differ. One position is that doing so would help subjects understand what is actually happening in this environment. Maybe, for instance, they had missed that signals are noisy, and this makes it clear. A different perspective is that such feedback is not part of the model or representative of the typical situation of interest outside of the laboratory, and thus providing it is artificial in a way that interferes with what we want to learn.

[23] While there is some variation on this in the literature, we always opt for using a stage game matrix to present actions and payoffs associated with a round.

3.4 Lessons for other environments

As we have already pointed out, in repeated game experiments, participants need to discover the strategies on their own. As we have shown, in such environments, experience is often important. In fact, these considerations apply to a broader range of games. Consider, for instance, bargaining. Many interactions involving bargaining cannot be modeled as a repeated game, but, as in repeated games, dynamic incentives play a central role in our understanding of these strategic interactions. In a standard bargaining game, the proposer must anticipate how others will respond to different types of offers in the current round, and how the game will evolve in future rounds conditional on the current offer being rejected. An example of this can be seen in Fréchette et al. (2003), which studies the bargaining game of Baron & Ferejohn (1989). Under the closed amendment rule, a proposer suggests a division of money to a group, which is implemented if it receives a majority of votes; otherwise payoffs are discounted and the game starts over. In their experiment, the theoretical prediction is that the proposer should take slightly more than two-thirds of the money and give the rest to half of the group members to ensure a winning coalition. What they observe is that between the first supergame and the last one (corresponding to the 15th), the share the proposer takes for themselves increases by more than 40%. The main force behind this increase is that proposers discover that they can offer money to only half of the subjects without risking the offer being rejected. This could be either because they need to realize the role of majority voting or to better understand the preferences of other subjects. In another treatment, the experiment implements an open amendment rule, where the theory predicts a lower share to the proposer. The results are such that, if they had only studied a single supergame, they would have observed no treatment differences. By contrast, experienced behavior revealed important treatment effects in the direction suggested by theory: proposer take a much larger share of the pie under the closed rule.

There are many games where, in equilibrium, one player's behavior is a function of the other's type, or of a state, for instance. In such games, strategies and actions are again distinct, and designing experiments where subjects gain experience may be particularly desirable.

This being said, even in extremely simple one shot games where actions and strategies are equivalent, experience can be important. Subjects do not necessarily appreciate the strategic tensions at the heart of a game from the description provided in the instructions. Playing the game can help them realize the relevant features, such as an action being dominant (or dominated). It also allows them to form more accurate beliefs. As we pointed out earlier, the description of the games used in experiments are, by design, disconnected from the subject's usual point of reference, and thus forming expectations may be at first challenging. This can explain why, for instance, cooperation rates decrease as subjects gain experience in a simple game such as the one-shot PD.

3.5 Data Analysis: A Focus on Experienced Behavior

Some experimenters view the use of repetitions in experiments as a cost-effective way to increase sample size. Others think of it as a way to allow the analyst to control for subject level heterogeneity. Here, we argue that in some settings experience may indeed be a necessary tool to obtain the right answer to the research question studied in an experiment, as inexperienced data may not properly reflect subjects' understanding of a strategic environment, especially in repeated games where the strategies are distinct from the choices. From this perspective, data analysis should be performed on data after subjects have gained experience.

However, it is most often impractical to focus solely on the last decision of a subject. Using a larger set of decisions from each participant can help the researcher by increasing the sample size and allow to control for heterogeneity at the subject or session level by taking advantage of the panel structure of the data. Note that this is not to say that early behavior is not useful or informative. But the information it provides is often primarily informative of how participants learn about the strategic interaction, which may not align with answers to the main research question guiding the experiment on how behavior differs between some treatment and control conditions. We think that in many cases, it makes sense to present the main data analysis based on somewhat experienced behavior in combination with some analysis on the entire data and key figures displaying evolution of behavior to help readers understand how the specific choice of what part of the data the main analysis is performed on may have mattered for the key results.

4 Strategies

The distinction between actions and strategies in repeated games poses many challenges. In the previous section, we discussed how experiments need to be designed to account for learning through experience that allow subjects to properly understand the strategic environment. In this section, we discuss one of the challenges faced by researchers trying to understand and organize the subjects' behavior. Namely, the experimenter does not observe strategies, they only see choices, but the object of interest is, at least for some research questions, strategies.

In indefinitely repeated PDs, for instance, sufficiently patient players can support equilibrium outcomes that are more socially desirable than what would be possible in a one-shot or a finitely repeated PD, or among less patient players in the same game. In such instances, there is typically a plethora of equilibria; and those cooperative outcomes require the use of dynamic incentives to be supported. This raises two types of questions: 1) What guides equilibrium selection? In other words, are equilibrium outcomes predictable and do they depend in some systematic way on the parameters of the game? 2) What strategies support these outcomes? And can we make sense of the strategies that are used—are they equilibrium strategies, are they best responses, etc?

Importantly, let us state that, from an empirical perspective, we think of strategies as an "as if" description. We do not posit that people necessarily make complete contingent plans of action for every potential history. Instead, we think of strategies as providing a description of how history maps into choices for subjects in our experiments. Immediately, this underscores a challenge: how can we discuss something like a *contingent plan of action* given that in a supergame, we only observe choices along a specific history.

The right panel of Figure 3 can help us illustrate this challenge more concretely. Taking the data set used in that figure as a whole, in 75% of the supergames, choices are consistent with one of the three strategies: AD, grim, or TFT. On the one hand, this seems impressive, as it suggests that choices can easily be rationalized. However, even though this is a very small set of strategies (with only two cooperative strategies), 30% of supergames involve choices that are equally consistent with either grim or TFT. Hence, we have learned much less than it first appears as we cannot easily distinguish between strategies that are the same along common histories. The resulting question is how to use the data or what additional data to bring to bear to better understand subjects' strategies.

Different approaches have been used to tackle this challenge. On the estimation side, they often rely on some degree of aggregation. On the experimental side, they have to do with enriching the data: through experimental devices to elicit additional choices that would not occur outside of the laboratory, via methods that make subjects explain what they are thinking, or through mechanisms to directly elicit strategies.

Some of the first empirical work on strategies in repeated games is done by Axelrod (1980*a*, 1980*b*), who organizes contests where participants propose strategies. He then studies which strategies perform better in his tournaments of simulated play. A method to recover strategies from choices in repeated trust games is proposed in Engle-Warnick & Slonim (2006*a*) and its application first appears in print in Engle-Warnick & Slonim (2004). This approach, in essence, provides a way to trade-off goodness of fit and number of strategies from a pre-specified set. A different approach applied to indefinite PDs with imperfect monitoring is proposed in Aoyagi & Fréchette (2009). It recovers a representative strategy and allows for heterogeneity according to some distribution away from the central strategy. Dal Bó & Fréchette (2011) recover strategies in indefinite PDs using a method they refer to as the Strategy Frequency Estimation Method (SFEM). Recognizing that past outcomes in a supergame can be treated as data: a simple mixture model over a pre-specified set of strategies can be estimated. Recently, using data from indefinite PD games, Heller & Tubul (2023) suggests a method to adapt clustering techniques to recover strategies. Many variations on these approaches exists, but those cover the main types. However, this being a chapter on experimental methods, instead of diving further into the estimation of strategies, we will explore experimental methods designed to help us understand strategies.

Elicitation of Strategies

Experimental economists are familiar with the strategy method (Selten 1967): the subject is asked for choices conditional on some condition, state, outcome, etc. In an extensive form game, the experimenter can ask the subject for choices at each node, after the choices have been recorded, the subjects are informed of what specific outcome was realized. Simply importing the strategy method to repeated games is problematic for a few reasons.

First, unlike in a simple extensive form game, the strategy space in a repeated game can be infinite, making the selection of available options for participants a key decision. On one hand, as we will see, computers offer flexibility, enabling us to provide participants with an extremely rich strategy space. On the other hand, this richness can overwhelm subjects, potentially causing them to behave in ways they might not in natural settings. This trade-off must be carefully considered, but conducting experiments with different implementations can help mitigateor highlightthese concerns

Second, subjects may not naturally think in terms of strategies, at least not in the way a game theorist would. As mentioned, we use the term in an *as if* sense. In this context, the *as if* becomes significant because we are directly asking subjects to describe how they are playing the game. To make this accessible, we aim to use language that participants understand. For example, one of the methods we discuss below uses the term *plan of action*simple, clear language that effectively conveys the meaning of a strategy. Additionally, it's important to provide opportunities for subjects to learn how to use technology to express how they intend to play the game. This can take the form of feedback meant to facilitate that aspect of understanding.

Third, and related to the first two points, how can this be done without affecting what subjects do? Since this method is an artificial construct designed for the experimenter to see what is not visible in the naturally occurring setting, it's crucial that the behavior of interest remains unaffected. There has been a long standing debate about whether the strategy method affects choices (see for instance Brandts & Charness (2011) who report that no treatment effect found with the strategy method is not found with the direct method). That debate often focuses on the differing emotional state at the point of decision: hot versus cold. In the case of repeated game, we think there is a more important challenge. As we have argued in the previous section, players of repeated games must discover strategies, something they do through repeatedly playing supergames. Introducing subjects to strategies may very well interfere with their normal process of learning and change their behavior. Hence, to properly elicit strategies, one must find ways to mitigate the interference of the method on the learning process.

A few early experiments have elicited strategies. Selten et al. (1997) experimentally investigated strategies in a finitely repeated Cournot duopoly where subjects, participants at a workshop, were taught a computer language to program strategies. However, their design limits what we can learn: subjects were not paid, they were instructed that their objective should be to attain a sum of profits as high as possible, after every play of their program—which they could modify in between play—they had group discussions, etc. Later, Bruttel & Kamecke (2012) designed an experiment where subjects play the three first rounds of an indefinite PD, followed by what choice they want to make for each possible outcome in the previous round: a memory-one strategy (the paper also includes other treatments). However, they are not interested in strategies per se, but rather they view this as an alternative way to induce indefinite games in the laboratory. Before pointing out two interesting observations they make, we note that they report differences between their random termination treatment and their strategy treatments (cooperation drops over rounds of the random termination treatment whereas it is more stable in their strategy treatment). This being said, they conclude that a large fraction of choices in this game is consistent with the memory-one restriction in their strategy treatment. They also note that some subjects attempt to re-establish cooperation following joint defection, but that such behavior disappears with experience. This is reminiscent of the increased responsiveness of Figure 2 and of the increase in threshold strategies from finitely repeated PD experiments discussed in the previous section.

Dal Bó & Fréchette (2019) propose a novel design to elicit strategies motivated by the three concerns highlighted above: rich strategy space, helping subjects discover how to express their strategies, and doing so with minimal impact on the discovery process. Their design is in three phases. Phase 1: subjects play randomly terminated PDs, just as they would in a standard implementation (in fact this is as in Dal Bó & Fréchette (2011)). Phase 2: Subjects are given additional instructions where they are presented with a computer interface allowing them to enter their *plan of action* in between each supergame (we will come back to what plans of action could be selected later). This plan of action is passive, as subjects are still required to enter choices in each round. However, in between supergames, they are shown the choices they and their partner made, but also the choices their plan of action would have selected. They are informed that after a certain amount of time as elapsed, Phase 2 will end. Because the length of supergames is random, any supergame may be the one in which Phase 2 terminates. Phase 3: the plan of action specified at the beginning of the last supergame of Phase 2 takes over and plays for the subject the remainder of that supergame and a number of additional supergames (at which point the subject can no longer alter their plan of action).

The purpose of Phase 1 is to allow people to gain an understanding of the strategic situation as they would absent strategy elicitation. Phase 2 is a didactic tool to help subjects learn to express what choices they want to make in the form of a strategy. Phase 3 serves to incentivize the selection of strategies in Phase 2.

The richness of the strategy space is a function of what plans of action are available to the subjects. The paper explores different treatments: some offer a simple interface of memory-one strategies that are easy to understand (if-then statements), but restrict the strategies available; while other offer a very large set of strategies from a combination of conditional strategies (up to memory-two) and lists of strategies with free parameters (for example, cooperate for X rounds and defect forever thereafter).

Figure 4 shows the evolution of round one choices in two treatments with the standard, direct choice, method (Dal Bó & Fréchette 2011); and in the strategy elicitation treatments of Dal Bó & Fréchette (2019).[24] For the data with strategy elicitation, the part of the line that is solid only involves direct elicitation. We highlight these two treatments because they feature opposite evolution: in one, cooperation increases while in the other it decreases. Importantly, these comparative statics are preserved in the two implementations, suggesting that the strategy elicitation paradigm implemented did not significantly affect the subjects' ability to learn from experience.[25]

[24]The figure displays supergames for which there is data in all sessions, except for supergames 24 through 29, because in Dal Bó & Fréchette (2019) with $\delta = 0.75$ one session has only 23 supergames. The grey area and the vertical bars provide session clustered 95% confidence intervals.

[25]We note that in one of their treatment, the data from Dal Bó & Fréchette (2011) and Dal Bó & Fréchette (2019) are significantly different, but this is true from the start: before any strategy elicitation has started. Hence, they

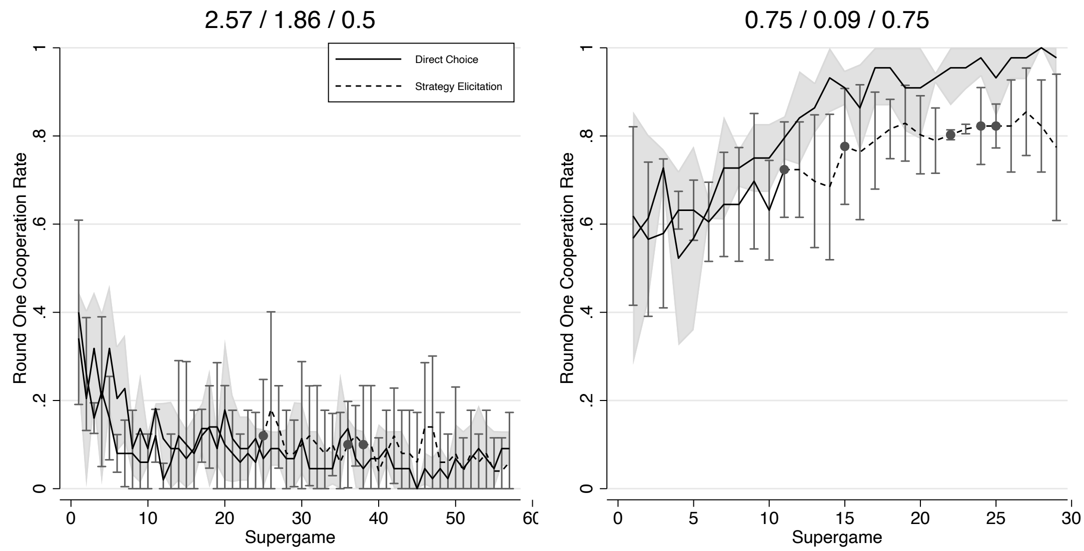

Panel titles indicate g and l of normalized staged game followed by δ.
The start of the dashed line indicates the first session that moves into Phase 2. Each circle indicates the beginning of Phase 2 for one session. After the last circle, all 2019 sessions are in Phase 2.

Figure 4: Evolution of Round One Choices in Dal Bó and Fréchette (2011, 2019)

We will revisit the findings of this paper at the end of the section, but we want to emphasize that, in our view, Phase 1 plays a pivotal role in this paradigm. Further evidence supporting this can be found in Romero & Rosokha (2018, 2023). The first paper introduces an alternative strategy elicitation method. While the primary difference relative to Dal Bó & Fréchette (2019) lies in how strategies are specified, they also use a paradigm that includes an initial Phase 1, where subjects make decisions one round at a time before strategy elicitation is introduced. The second paper expands the interface to include strategies that randomize between C and D. In both this, and the 2018 paper, one of their treatments overlaps with one of the treatments in Dal Bó & Fréchette (2019)—the one with $\delta = 0.95$. They show (see Figure OA5 in their appendix) that the evolution of cooperation (all rounds and round one) are very similar in all three papers. Hence, using a gradual, multi phase, design; it seems possible to elicit more than just choices without significantly impacting the evolution of behavior. However, it is important to note that this consistency is absent in earlier work where strategies were elicited from the outset without the use of an initial Phase 1 (Romero & Rosokha 2016).

The strategy elicitation method introduced in Romero & Rosokha (2018) has subjects define *rules* that map histories to a choice. Subjects can define as many rules as they want. The approach is guaranteed to yield a choice because: 1) Subjects are forced to define a rule (choice) for round

conclude that this must be attributed to variation that is unrelated to the elicitation method.

one and they must define a default rule, a rule that applies if no other rule indicates what to do. 2) If two rules disagree, the one with the longest memory is prioritized. The rule set determines the strategy.

Collectively, we now have several papers that employ a multi-phase, gradual paradigm, though they differ in the constraints placed on the strategies that can be formulated and in the way participants are asked to articulate those strategies. Interestingly, in their common treatment ($g = 2.57, \ell = 1.86, \delta = 0.95$), the qualitative results are very consistent: 1) The majority of strategies have memory-one (75% in both Dal Bó & Fréchette (2019) and Romero & Rosokha (2018)). 2) That fraction increases as δ decreases (Dal Bó & Fréchette (2019) and Romero & Rosokha (2018)). 3) A large fraction of strategies come from a small set. Again, in the treatment that overlap, Dal Bó & Fréchette (2019) report that 53% of strategies are AD, grim, TFT, or Suspicious Tit-For-Tat (STFT).[26] Romero & Rosokha (2018) finds that 40% exactly correspond to that set, while 72% are very close. Romero & Rosokha (2023) have instead that 38% are exactly AD, grim, or TFT; while the fraction that is very close to one of those is 54%.

None of these designs are immune to criticism—some are more complex to understand, others are more restrictive. On the other hand, estimation is also no silver bullet, as it imposes many assumptions. Dal Bó & Fréchette (2019) summarizes the results of papers that have estimated strategies using SFEM with a similar set of strategies. They find that, in almost all treatments (across a range of treatment parameters), the majority of strategies correspond to one from the same set as above (plus AC in some treatments): i.e. AD, grim, TFT, STFT, or AC. A more recent paper, Heller & Tubul (2023), using k-means clustering (based on the meta-data from Dal Bó & Fréchette (2018)), arrives at a similar conclusion: most behavior can be explained by one of these five strategies (with some deviations from the pure strategies). Thus, while each approach has its limitations, together they present a rather consistent picture.

Outside of the indefinitely repeated PD with perfect monitoring, the analysis of strategies in games with imperfect monitoring suggests a greater reliance on strategies with longer memory (Fudenberg et al. (2012) and summarized in Dal Bó & Fréchette (2018)). In contrast, strategies in finitely repeated games are of a different nature. Embrey et al. (2019) shows that subjects often use threshold strategies, which are also estimated using SFEM in Aoyagi et al. (2024). Beyond these two games, strategies have been studied in the trust game (Engle-Warnick & Slonim 2004), but relatively few papers have systematically explored strategies outside of the repeated PD.

[26] STFT behaves like TFT after round one but starts with defection.

Experimental designs involving strategy elicitation seemed more promising when they were initially explored. Although those papers have yielded very useful information that has clarified our understanding, applying this approach even in very simple environments seems challenging enough to limit what can be done in more involved repeated games. In settings where more needs to be learned, if the feedback is more complex or the choice space is richer, achieving the goals of including relevant strategies, ensuring subjects understand how to express their strategies, and doing so without interfering with their learning now seems to offer limited possibilities.

Procuring data beyond the path of play

One-Period Ahead Strategy Method

Besides eliciting complete strategies, a more subtle intervention can involve obtaining more information on choices that are not on the path of play. Earlier, we mentioned that in an indefinitely repeated PD, it is not possible to distinguish between, say, TFT and the Grim trigger if along the observed path of play subjects always cooperate. The problem of strategy identification based on observed choices is exacerbated in dynamic games, where strategies can also depend on an underlying state variable that changes over time. To illustrate this, let us revisit the dynamic common pool problem. In applications of dynamic games, it is common to focus on Markov strategies. A Markov strategy in the dynamic common pool problem specifies how much of the resource to extract for each possible state (stock of fish in the pond). These are strategies that do not condition on past play. Suppose that there is a relatively low extraction rate that lets the stock grow over time and captures the efficient outcome. Observing participants choosing low extraction rates is consistent with the Markov strategy 'select the low extraction rate for all levels of the state,' but it is also consistent with any Non-Markov strategy (i.e. a strategy that conditions on past play) that selects the low extraction rate in the 'cooperative' phase.

To procure data outside of the path of play in an incentivized manner, Vespa (2020) introduces the one-period ahead strategy method. In round one of a supergame, participants make a choice as in the standard implementation, but at the end of round one participants are not provided with feedback on what the other agents they are facing did. In round two, the subject is asked to specify what she would choose in that round for *any possible* choice that other participants could have made in round one. After the subject specifies choices for all such contingencies, the actual round one choices are revealed and the round two choice that corresponds to the actual contingency is implemented. The procedure is repeated in round two and so on. This procedure incentivizes subjects to report choices that, ex-post, are not on the path of play.

Using the strategy method can in principle change the way that people behave in a game. In the first supergames of Vespa (2020), participants gain experience in a standard setting without the one-period ahead strategy method. Once participants have had experience, the one-period ahead strategy method is introduced and in that paper the data does not point towards a change of behavior due to the introduction of the method. However, we are not aware of a more systematic study of the potential impact of this elicitation.

Another aspect worth noting is that asking subjects to think about several hypothetical contingencies in each round is only feasible if the action set is discrete and small and the set of players in the game is small as well. Asking subjects to make choices using the one-period-ahead strategy method requires more time, creating a trade-off with the number of supergames participants can experience within a typical session.

Teams

Instead of uncovering strategies by explicitly eliciting choices at nodes beyond the ones encountered in the normal path of play, Kagel & McGee (2016) and Cooper & Kagel (2023) propose an alternative approach. The role of each agent in the game (finite PD and indefinite PD respectively), is played by a pair of subjects. The subjects must agree on a decision and have access to an online chat for communication. This chat offers valuable insight into the strategic considerations made by the participants.

Some interesting results that come out of this approach for finite PD games are the following: 1) Cooperation in the last round comes from "mistakes, confusion, or naïveté." 2) When deciding to defect one round earlier from one supergame to the next, subjects do not anticipate that others may do the same. However, this is not inconsistent with profit-maximizing behavior, given the average behavior of others. In the indefinite PD, teams tend to start by playing AD. This is driven by distrust of the opponent. Over supergames, many evolve to attempt cooperative strategies that are of the grim type (grim and lenient versions of grim). When attempting to cooperate, they often do so for a few rounds in the hope that the others will change strategies.

This teams paradigm has some clear advantages. Foremost, it imposes very little structure. Hence, it can allow the discovery of features of strategic thinking that we are not expecting. It also reduces the concerns having to do with demand induced effects present in other strategy elicitation techniques. However, the main disadvantages are twofold: First, the content requires interpretation. Second, and maybe more importantly, it provides insight into a different environment. Teams behave differently from individuals (as these papers indeed show). Hence, the considerations expressed may be the relevant ones, but it is impossible to know to what extent the specific findings are only relevant for team play. As with most of these methods, we find their complementary features to be the ultimate value. Taken together, they help paint a more complete understanding.

Pre-play chat

An alternative approach to understanding what type of strategies participants consider in a repeated game is to provide them with the opportunity to exchange free messages before the supergame starts. Notice that this is different from teams, as described earlier. In a setting with teams, two or more participants make choices acting as a single agent, facing other teams. Chat in that case can provide insight into how that single agent makes a decision when there is no opportunity to communicate with the other agent (represented by the other team). The type of pre-play chat we consider here does not involve teams. Instead, each player is allowed to talk to their opponent before the supergame starts.

Since there is no commitment to any messages, chat is cheap talk and the set of equilibria remains unchanged. But pre-play communication can be seen as a tool for revealing strategies that subjects were considering but failed to implement due to coordination problems For instance, consider a participant in an indefinitely repeated PD who is not cooperating. It could be that she never considered cooperation as an option, or that she did consider it but ruled it out, for instance, due to strategic uncertainty, in the sense that she does not know whether others would be willing to coordinate. Pre-play chat can provide information to distinguish between these two cases.

Figure 5 illustrates the effects of pre-play communication. It reports initial cooperation rates in the four-player indefinitely repeated prisoner's dilemma studied in Boczoń et al. (2024). The standard baseline treatment with no pre-play communication (NoChat) is displayed as a dashed black line. Initially, cooperation rates are very low and move toward the lower bound, remaining low, even though efficient outcomes could theoretically be supported with the implemented discount factor of $\frac{3}{4}$. The solid black line captures a treatment with the same discount factor and no communication in the first ten supergames, but where pre-play chat is allowed in the second half. Immediately after chat is allowed, there is a large effect and cooperation rates jump from close to 0 to close to 100 percent.[27] This suggests that prior to the introduction of chat many subjects were

[27]Relatedly, Agranov & Yariv (2018) argue that repeated auctions yield lower collusion than one-shot auctions with communication.

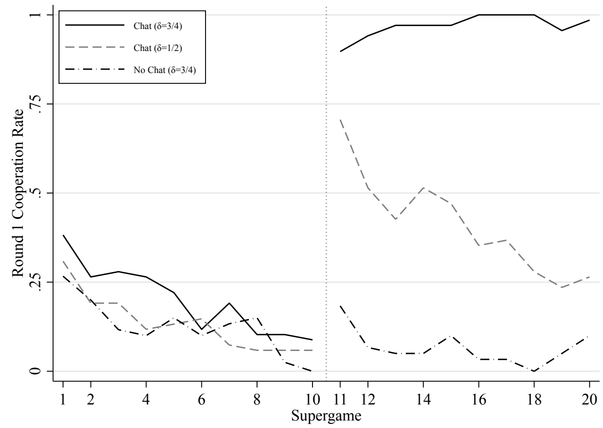

Figure 5: Initial cooperation rates in the four-player prisoner's dilemma studied in Boczoń et al. (2024)

Notes: In the three between-subject treatments the first ten supergames do not allow for pre-play communication (chat). Treatments differ depending on whether the discount factor equals $\frac{3}{4}$ or $\frac{1}{2}$, and on whether there is Chat allowed or not in the last 10 supergames, as indicated in the Figure.

considering strategies that could implement cooperation but, because of strategic uncertainty, they defected instead.[28]

The large effect of chat, however, raises the question: can pre-play communication sustain cooperation regardless of the actual incentives to support cooperation? To evaluate this option, Boczoń et al. (2024) conduct a third treatment with chat in the second half of the session, but where the discount factor is reduced to $\frac{1}{2}$. The minimum discount factor at which cooperation can be theoretically supported in this game is exactly $\frac{1}{2}$, so this treatment tests whether communication supports cooperation when we are exactly at the boundary. Initially, subjects aim to leverage communication (dashed gray line), but with experience the effect dissipates and the cooperation rate strongly trends downwards. This suggests that pre-play communication can help with coordination provided that the underlying incentives for cooperation are strong enough.[29]

[28] An alternative way to assess the role of strategic uncertainty involves comparing a simultaneous-move game to a sequential version in which one player knows the choice of the other. For designs with this feature, see Kartal & Müller (2022) and Ghidoni & Suetens (2022).

[29] Other recent experiments that study communication in repeated settings include Harrington et al. (2016), Cooper & Kühn (2014), and Wilson & Vespa (2020).

Lessons For Other Environments

In our opinion, the multi-phase design introduced to allow for strategy elicitation without disrupting learning can be a useful design for other purposes as well. For instance, Aoyagi et al. (2024) use a two phase design to add belief elicitation to a repeated PD. They study both finite and indefinite treatments. The parameters of the finite treatment correspond to those of Embrey et al. (2019). They show that: 1) cooperation rates across the two papers (with and without belief elicitation) are not different, and 2) the cooperation rates after belief elicitation starts follow the same trends as they do before elicitation.

Clearly, using a multi-phase design comes at a cost. For instance, in Aoyagi et al. (2024) they have no record of beliefs for early supergames. Nonetheless, in cases where experience is important, we think that the benefits outweigh the potential costs.

In a completely different setting, Nagel et al. (2024) adapt the procedure to elicit strategies in first-price common-value auctions. Initially participants plays in a series of common-value auctions, using the canonical environment of Kagel & Levin (1986). Essentially, there is an object for sale and its value is unknown to participants, but each participant receives an informative signal. Participants bid simultaneously, the winner gets the object and pays the submitted bid. Subjects play several rounds of this game with random matching, which aims to provide them with experience in this environment, as in Phase 1 of the elicitations described earlier. Subsequently, the interface estimates a bidding function for each participant using the bids that they submitted in previous auctions for each signal they observed. The estimated bid function is presented to subjects in a figure that has the signal space on the horizontal axis and bids on the vertical axis. The figure also displays the bids that subjects actually placed for the signals that they observed (i.e. the information that was used to estimate the bid function). Participants then face a few rounds of practice in which they can change the bid function if they want to. This part of the implementation is inspired by Phase 2 of the procedure in Dal Bó & Fréchette (2019): participants are presented with a strategy and a way to select other strategies and they can make changes. After a few auctions in which participants can change their bid function, the bid function takes over and determines bids for them in a subsequent set of rounds, matching the ideas in Phase 3 of Dal Bó & Fréchette (2019).

The strategies that can be recovered from subjects' choices in what corresponds to the first phase are not different from the strategies that they eventually specify, suggesting that the procedure does not systematically affect the choices that subjects make in the auction environment either.

5 Conclusion

Writing on field experiments in ecology, Hurlbert (1984) identifies threats to the validity of an experiment ("sources of confusion") and how to minimize them. For instance, there could be a temporal change (in a field experiment, measurements are taken over an extended period of time), and the solution is to have a control. Another is experimenter bias, and this can be addressed with a blind design. He identifies seven such potential problems, the last being *demonic intrusion*, which he describes in the following way: "If you worked in areas inhabited by demons, you would be in trouble regardless of the perfection of your experimental designs. If a demon chose to do something to each experimental unit in treatment A but to no experimental unit in treatment B, and if his/her/its visit went undetected, the results would be misleading." He gives a hypothetical example of studying the effect of fencing on fox predation. If hawks use fence posts as perches and that helps them attack animals (and this is not known to the experimenter), even though foxes are blocked by the fence, the experimenter would incorrectly conclude that fencing increases fox predation. We mention this as it highlights the potential pitfalls of writing about experimental design: a good experimental design is only good or bad with respect to what researchers understand or can conceive of at the point in time where it is elaborated. In this chapter, we described what we think are some important considerations when designing repeated-game experiments, however, the reader should keep in mind that this is from the perspective of our current state of knowledge.

We believe that many of the aspects of experimental design and data that we have highlighted are relevant beyond repeated game experiments. In particular, the effect of experience on behavior is relevant in multiple settings. Because of that, special care must be given to the computer interface and the feedback it gives. This also translates into experienced data as being relevant for key hypothesis tests.

Note also that the special role of experience and the importance of attention to feedback in repeated games experiments may suggests that online platform do not offer the ideal setting. For instance, **?** give an example of an experiment where learning is obvious and important when the experiment is conducted in a traditional laboratory, but is non-existent in an equivalent online experiment. [30] Hence, the results from the laboratory experiment are much more in line with theory than those from the online experiment.

Finally, we think that some of the experimental tools we develop to obtain additional data beyond the basic choice can sometime benefit from being introduced using a multi-phase approach.

[30] The experiment explores a lemons problem where buyers must learn not to overpay for the item.

This is not always true, but probably more important when the elicitation may trigger different consideration, or an alternative perspective, on the choice at hand.

As we hope this chapter has made clear, very few design choices in repeated game experiments have clear cut answers. There are advantages and disadvantages to most options available and reasonable people could disagree on the best design. But along some dimensions, such as the need for experience, our opinion is more firmly established. In the end, demonic intrusions are difficult to completely prevent, only vigilance and care can be recommended.

References

Agranov, M. & Yariv, L. (2018), 'Collusion through communication in auctions', *Games and Economic Behavior* **107**, 93–108.

Andreoni, J. & Miller, J. H. (1993), 'Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence', *The Economic Journal* **103**(418).

Aoyagi, M., Bhaskar, V. & Fréchette, G. R. (2019), 'The impact of monitoring in infinitely repeated games: Perfect, public, and private', *American Economic Journal: Microeconomics* **11**(1), 1–43.

Aoyagi, M. & Fréchette, G. R. (2009), 'Collusion as public monitoring becomes noisy: Experimental evidence', *Journal of Economic theory* **144**(3), 1135–1165.

Aoyagi, M., Fréchette, G. R. & Yuksel, S. (2024), 'Beliefs in repeated games: An experiment', *Available at SSRN 3784841* .

Axelrod, R. (1980*a*), 'Effective choice in the prisoner's dilemma', *Journal of Conflict Resolution* **24**(1), 3–25.

Axelrod, R. (1980*b*), 'More effective choice in the prisoner's dilemma', *Journal of Conflict Resolution* **24**(3), 379–403.

Azrieli, Y., Chambers, C. P. & Healy, P. J. (2018), 'Incentives in experiments: A theoretical analysis', *Journal of Political Economy* **126**(4), 1472–1503.

Baron, D. P. & Ferejohn, J. A. (1989), 'Bargaining in legislatures', *American political science review* **83**(4), 1181–1206.

Blonski, M., Ockenfels, P. & Spagnolo, G. (2011), 'Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence', *American Economic Journal: Microeconomics* **3**(3), 164–192.

Boczoń, M., Vespa, E., Weidman, T. & Wilson, A. J. (2024), 'Testing models of strategic uncertainty: Equilibrium selection in repeated games', *Journal of the European Economic Association (forthcoming)* .

Brandts, J. & Charness, G. (2011), 'The strategy versus the direct-response method: a first survey of experimental comparisons', *Experimental Economics* **14**, 375–398.

Bruttel, L. V. & Kamecke, U. (2012), 'Infinity in the lab: How do people play repeated games?', *Theory and Decision* **72**(2), 205–219.

Cabral, L., Ozbay, E. Y. & Schotter, A. (2014), 'Intrinsic and instrumental reciprocity: An experimental study', *Games and Economic Behavior* **87**, 100–121.

Chandrasekhar, A. G. & Xandri, J. P. (2023), 'A note on payments in the lab for infinite horizon dynamic games with discounting', *Economic Theory* **75**(2), 389–426.

Cooper, D. J. & Kagel, J. H. (2023), 'Using team discussions to understand behavior in indefinitely repeated prisoners dilemma games', *American Economic Journal: Microeconomics* **15**(4), 114–145.

Cooper, D. J. & Kühn, K.-U. (2014), 'Communication, renegotiation, and the scope for collusion', *American Economic Journal: Microeconomics* **6**(2), 247–278.

Cooper, R., DeJong, D. V., Fosythe, R. & Ross, T. W. (1996), 'Cooperation without reputation: experimental evidence from prisoner's dilemma games', *Games and Economic Behavior* **12**(2), 187–218.

Dal Bó, P. (2005), 'Cooperation under the shadow of the future: experimental evidence from infinitely repeated games', *The American Economic Review* pp. 1591–1604.

Dal Bó, P. & Fréchette, G. R. (2011), 'The evolution of cooperation in infinitely repeated games: Experimental evidence', *The American Economic Review* **101**(1), 411–429.

Dal Bó, P. & Fréchette, G. R. (2018), 'On the determinants of cooperation in infinitely repeated games: A survey', *Journal of Economic Literature* **56**(1), 60–114.

Dal Bó, P. & Fréchette, G. R. (2019), 'Strategy choice in the infinitely repeated prisoners dilemma', *American Economic Review* **109**(11), 3929–3952.

Dreber, A., Rand, D. G., Fudenberg, D. & Nowak, M. A. (2008), 'Winners don't punish', *Nature* **452**, 348–351.

Duffy, J. & Ochs, J. (2009), 'Cooperative behavior and the frequency of social interaction', *Games and Economic Behavior* **66**(2), 785–812.

Embrey, M., Fréchette, G. R. & Yuksel, S. (2018), 'Cooperation in the finitely repeated prisoners dilemma', *The Quarterly Journal of Economics* **133**(1), 509–551.

Embrey, M., Mengel, F. & Peeters, R. (2019), 'Strategy revision opportunities and collusion', *Experimental Economics* **22**, 834–856.

Engle-Warnick, J. & Slonim, R. L. (2004), 'The evolution of strategies in a repeated trust game', *Journal of Economic Behavior & Organization* **55**(4), 553–573.

Engle-Warnick, J. & Slonim, R. L. (2006*a*), 'Inferring repeated-game strategies from actions: Evidence from repeated trust game experiments', *Economic Theory* **28**(3), 603–632.

Engle-Warnick, J. & Slonim, R. L. (2006*b*), 'Learning to trust in indefinitely repeated games', *Games and Economic Behavior* **54**(1), 95–114.

Feinberg, R. M. & Husted, T. A. (1993), 'An experimental test of discount-rate effects on collusive behaviour in duopoly markets', *The Journal of Industrial Economics* pp. 153–160.

Flood, M. M. (1958), 'Some experimental games', *Management Science* **5**(1), 5–26.

Fréchette, G. R., Kagel, J. H. & Lehrer, S. F. (2003), 'Bargaining in legislatures: An experimental investigation of open versus closed amendment rules', *American Political science review* **97**(2), 221–232.

Fréchette, G. R. & Yuksel, S. (2017), 'Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination', *Experimental Economics* **20**, 279–308.

Friedman, D. & Oprea, R. (2012), 'A continuous dilemma', *American Economic Review* **102**(1), 337–363.

Fudenberg, D., Rand, D. G. & Dreber, A. (2012), 'Slow to anger and fast to forgive: Cooperation in an uncertain world', *American Economic Review* **102**(2), 720–749.

Ghidoni, R. & Suetens, S. (2022), 'The effect of sequentiality on cooperation in repeated games', *American Economic Journal: Microeconomics* **14**(4), 58–77.

Harrington, J. E., Gonzalez, R. H. & Kujal, P. (2016), 'The relative efficacy of price announcements and express communication for collusion: Experimental findings', *Journal of Economic Behavior & Organization* **128**, 251–264.

Heller, Y. & Tubul, I. (2023), 'Strategies in the repeated prisoners dilemma: A cluster analysis'. Working Paper.

Hurlbert, S. H. (1984), 'Pseudoreplication and the design of ecological field experiments', *Ecological monographs* **54**(2), 187–211.

Kagel, J. H. & McGee, P. (2016), 'Team versus individual play in finitely repeated prisoner dilemma games', *American economic Journal: Microeconomics* **8**(2), 253–76.

Kagel, J. H. & Schley, D. R. (2013), 'How economic rewards affect cooperation reconsidered', *Economics Letters* **121**(1), 124–127.

Kagel, J. & Levin, D. (1986), 'The winner's curse and public information in common value auctions', *The American Economic Review* pp. 894–920.

Kahn, L. M. & Murnighan, J. K. (1993), 'A general experiment on bargaining in demand games with outside options', *The American Economic Review* **88**(5), 1260–1280.

Kartal, M. & Müller, W. (2022), 'A new approach to the analysis of cooperation under the shadow of the future: Theory and experimental evidence', *Working Paper* .

Kim, J. (2023), 'The effects of time preferences on cooperation: Experimental evidence from infinitely repeated games', *American Economic Journal: Microeconomics* **15**(1), 618–637.

Kreps, D. M., Milgrom, P., Roberts, J. & Wilson, R. (1982), 'Rational cooperation in the finitely repeated prisoners' dilemma', *Journal of Economic Theory* **27**(2), 245–252.

Murnighan, J. K. & Roth, A. E. (1983), 'Expecting continued play in prisoner's dilemma games: A test of several models', *Journal of conflict resolution* **27**(2), 279–300.

Nagel, L., Niederle, M. & Vespa, E. (2024), 'Decomposing the winner's curse', *Working Paper* .

Palfrey, T. R. & Rosenthal, H. (1994), 'Repeated play, cooperation and coordination: An experimental study', *The Review of Economic Studies* **61**(3), 545–565.

Proto, E., Rustichini, A. & Sofianos, A. (2019), 'Intelligence, personality, and gains from cooperation in repeated interactions', *Journal of Political Economy* **127**(3), 1351–1390.

Proto, E., Rustichini, A. & Sofianos, A. (2022), 'Intelligence, errors, and cooperation in repeated interactions', *The Review of Economic Studies* **89**(5), 2723–2767.

Rapoport, A. & Chammah, A. M. (1965), *Prisoner's dilemma: A study in conflict and cooperation*, Vol. 165, University of Michigan press.

Romero, J. & Rosokha, Y. (2016), 'Constructing strategies in the indefinitely repeated prisoners dilemma game'. Working Paper.

Romero, J. & Rosokha, Y. (2018), 'Constructing strategies in the indefinitely repeated prisoners dilemma game', *European Economic Review* **104**, 185–219.

Romero, J. & Rosokha, Y. (2023), 'Mixed strategies in the indefinitely repeated prisoner's dilemma', *Econometrica* **91**(6), 2295–2331.

Roth, A. & Murnighan, J. (1978), 'Equilibrium behavior and repeated play of the prisoner's dilemma', *Journal of Mathematical Psychology* **17**(2), 189–198.

Sabater-Grande, G. & Georgantzis, N. (2002), 'Accounting for risk aversion in repeated prisoners dilemma games: An experimental test', *Journal of economic behavior & organization* **48**(1), 37–50.

Selten, R. (1967). "Die Strategiemethode zur Erforschung des eingeschraenkt rationalen Verhaltens im Rahmen eines Oligopolexperiments" in *Beitrge zur experimentellen Wirtschatsforschung*, ed. H. Sauermann, Tbingen, J.C.B. Mohr, pp.136-168.

Selten, R., Mitzkewitz, M. & Uhlich, G. R. (1997), 'Duopoly strategies programmed by experienced players', *Econometrica* **65**(3), 517–556.

Selten, R. & Stoecker, R. (1986), 'End behavior in sequences of finite prisoner's dilemma supergames a learning theory approach', *Journal of Economic Behavior & Organization* **7**(1), 47–70.

Sherstyuk, K., Tarui, N. & Saijo, T. (2013), 'Payment schemes in infinite-horizon experimental games', *Experimental Economics* **16**, 125–153.

Vespa, E. (2020), 'An experimental investigation of cooperation in the dynamic common pool game', *International Economic Review* **61**(1), 417–440.

Vespa, E. & Wilson, A. J. (2017), 'Experimenting with equilibrium selection in dynamic games', *Working paper* .

Vespa, E. & Wilson, A. J. (2019), 'Experimenting with the transition rule in dynamic games', *Quantitative Economics* **10**(4), 1825–1849.

Wilson, A. J. & Vespa, E. (2020), 'Information transmission under the shadow of the future: An experiment', *American Economic Journal: Microeconomics* **12**(4), 75–98.

Wilson, A. J. & Wu, H. (2017), 'At-will relationships: How an option to walk away affects cooperation and efficiency', *Games and Economic Behavior* **102**, 487–507.