Coordination with Differential Time Preferences: Experimental Evidence*

Marina Agranov† Jeongbin Kim‡ Leeat Yariv§

September 5, 2023

Abstract

The experimental literature on repeated games has largely focused on settings where players discount the future identically. In applications, however, interactions often occur between players whose time preferences differ. We study experimentally the effects of discounting differentials in infinitely repeated coordination games. In our data, differential discount factors play two roles. First, they provide a coordination anchor: more impatient players get higher payoffs first. Introducing even small discounting differentials reduces coordination failures significantly. Second, with pronounced discounting differentials, intertemporal trades are prevalent: impatient players get higher payoffs for an initial phase and patient players get higher payoffs in perpetuity afterward.

Keywords: Repeated Games, Discounting, Intertemporal Trade, Experiments.

JEL: C73, C92, D15, D25.

*We thank Guillaume Fréchette, Drew Fudenberg, and Pëllumb Reshidi for many helpful comments. We also grateful to the Coeditor Dirk Bergemann and 5 anonymous reviewers for numerous helpful suggestions.

†California Institute of Technology and NBER; magranov@hss.caltech.edu

‡Florida State University; jkim33@fsu.edu

§Princeton University, CEPR, and NBER; lyariv@princeton.edu

1 Introduction

Societies are often capable of implementing efficient outcomes that are fragile to opportunistic behavior, even absent contractual instruments, see Ostrom (1990) and Ellickson (1991). The theory of repeated games has provided a foundation for understanding the informal emergence of cooperation. A burgeoning experimental literature studies the determinants of efficient outcomes when interactions are repeated. Most of this literature, however, presumes that agents exhibit identical time preferences and trade off current and future payoffs identically; see Dal Bó and Fréchette (2018).

In a variety of settings, agents value time differently: household members may face different life expectancies; legislators can represent diverse districts, subject to disparate time pressures; board members of a company may differ in age and investment portfolios. Different time preferences allow for new behaviors and outcomes to emerge through repeated interactions. Importantly, a patience differential across players can give rise to *intertemporal trades* in efficient equilibria, whereby more patient agents forgo payoffs early on in favor of their impatient partners and, in return, gain high payoffs in later periods. While some theoretical work has analyzed the implications of differential discount factors, its predictions have, by and large, remained untested (see our literature review).

We report results from a set of laboratory experiments designed to assess the impacts of heterogeneous time preferences. We show that differential discount factors play two roles: they ease coordination on particular equilibria and allow for efficient intertemporal trades. Our results suggest that behavior in the knife-edge case of equal discount factors may not represent behavior in settings in which players differ in their time preferences. The emergence of cooperation through repeated interactions takes a different form. Nonetheless, even with a discounting differential, efficient equilibria still have strong drawing power.

Our design is based on a simple two-player 2×2 coordination game, with each of the two coordination outcomes favoring (symmetrically) a different player. Four treatments correspond to infinite repetitions of the same coordination game with different profiles of discount factors. In two treatments—our Equal Low and Equal High treatments—players' discount factors coincide, but are either low (30%) or high (90%). In two other treatments—our Unequal Low and Unequal Mixed treatments—players' discount factors differ. In the Unequal Low treatment, both players are relatively impatient (discount factors of 30% and 40%), while in the Unequal Mixed treatment, players' discount factors are more disparate (30% and 90%). Since the resulting repeated games in our treatments exhibit many equilibria, we allow participants to communicate prior to each of the 10 repeated games played in every experimental session.

Theoretically, when discount factors coincide, equilibrium outcomes are always within the convex hull of the stage-game payoffs. This is no longer the case when discount factors differ. In our setting, with identical discount factors, any strategy profile that implements a coordination outcome in each round constitutes a utilitarian efficient equilibrium. In contrast, when discount factors differ, as in our Unequal Low and Unequal Mixed treatments, the most efficient equilibria entail intertemporal trades and generate payoffs outside the convex hull of the stage-game payoffs. Specifically, for our Unequal Low or Unequal Mixed treatments, in the most efficient equilibrium, the more impatient agent receives the high coordination payoff only in the first round or first two rounds, respectively. Afterward, the patient agent receives the high coordination payoff forever.

The empirical relevance of the theoretically efficient equilibria is not obvious. Implementing these equilibria requires non-trivial sophistication and trust. Players need to grasp the heterogeneous intertemporal tradeoffs in the game. Furthermore, patient agents need to trust that early concessions to their impatient counterparts will yield the promised rewards in later periods. Data on the implications of different discount factor distributions on behavior and outcomes are scarce. Our study starts filling this gap.

In our data, participants with identical discount factors largely alternate between the two pure stage-game equilibria. However, we see significantly more miscoordination when players' discount factors are low, when both players are impatient. Both players want to receive the high coordination payoff in the first round and the resulting tension creates coordination failures.

A small difference in players' discount factors acts as a coordination instrument. In our Unequal Low treatment, the more impatient of the two players receives the high coordination payoff first in nearly all games and coordination failures are significantly less frequent relative to our Equal Low treatment. The modal strategy profile is still alternation. However, a substantial fraction of player pairs intertemporally trade, with the more patient player receiving the high coordination payoff from the second round on. These pairs implement the most efficient equilibrium.

Players exploit heavily the possibility of trading payoffs over time when their discount factors differ substantially. Indeed, in our Unequal Mixed treatment, players predominantly implement efficient intertemporal trades. In nearly all repeated games in our data, the impatient player receives the high coordination payoff for two or three rounds at the start. The patient player receives the high coordination payoff in all later rounds.

Taken together, our results suggest that, in repeated interactions, both the magnitude and the distribution of discount factors can have important effects on behavior and coordination levels. In particular, with a discounting differential, agents are still drawn to efficient equilibrium outcomes and are able to implement the necessary efficient intertemporal trades.

2 Related Literature

The theoretical literature has made important advances in our understanding of repeated interactions in the presence of non-homogeneous discount factors. The seminal work of Lehrer and Pauzner (1999) analyzes two-player (complete information) infinitely repeated games in which discount factors are unequal. They characterize the sets of Nash and subgame perfect equilibrium payoffs. They show that players with unequal discount factors can make intertemporal payoff trades that expand the set of feasible payoffs of a repeated game beyond the convex hull of the stage-game payoffs. Guéron et al. (2011) and Chen and Takahashi (2012) extend the findings of Lehrer and Pauzner (1999) to repeated games with more than two players.

Players with differential discounting have also been dealt with in the reputation literature, going back to Kreps and Wilson (1982) and Milgrom and Roberts (1982). In that literature, incomplete information about the "type" of one player, typically the more patient player, is key. That literature often identifies conditions under which a "reputation" for being of a particular type can emerge in equilibrium. In particular, intertemporal trades across players are not central.[1]

A growing experimental literature tests various predictions of the theory of repeated games; see the survey of Dal Bó and Fréchette (2018). Experimental studies of complete-information repeated games have focused on players with identical discount factors. In fact, our benchmark treatments with homogeneous discount factors resemble several recent studies. For instance, Cason et al. (2013) study coordination behavior in a repeated assignment game and find that participants are more likely to coordinate on turn-taking strategies with a low degree of conflict. In repeated and symmetric battle of the sexes games, Kuzmics et al. (2014) and Romero and Zhang (2018) document frequent turn-taking strategies. These behaviors are consistent with the alternation profiles we observe.

Some work inspects reputation motives á la Milgrom and Roberts (1982). For example, Tingley and Walter (2011) emulate a reputation-building setting, where a long-run "defender" of an uncertain and private type faces a sequence of short-lived "entrants." To our knowledge, there is no experimental work on complete-information repeated games played by agents with differing time preferences. In particular, whether equilibria exhibiting intertemporal trades have empirical relevance has been an open question.[2] Our study provides an answer within the coordination game setting.

[1] Lehrer and Yariv (1999) study two-person repeated games where only one player is informed about a realized state. They show that intertemporal trades can occur in equilibrium with unequal discount factors even when players are arbitrarily patient and the stage game is zero-sum. See also Cripps and Thomas (2003).

[2] There is some work on the composition of households in terms of time preferences; see, for example, Schaner (2015) and references therein. Jackson and Yariv (2014) use experiments to assess agents' social rankings of consumption streams when society's individuals vary in their time preferences, speaking to theoretical results in Jackson and Yariv (2015).

3 Experimental Design

Our experimental treatments all involved repeated play of a two-player, symmetric coordination game with two pure equilibria, each preferred by one of the players. The corresponding stage game, displayed in terms of experimental points, was given by:

	R	B
G	1, 6	0, 0
P	0, 0	6, 1

where $\{R, B\}$ are the column player's actions and $\{G, P\}$ are the row player's actions. For our interface's transparency, each action was associated with a color—R with red, B with blue, G with green, and P with purple.[3] Furthermore, in order to avoid effects driven by the payoff matrix' orientation, our interface displayed the payoff matrix so that each participant was a row player. In all our sessions, participants played a repeated version of this game.

Theoretically, expected payments of an exponential discounter with a discount factor of δ coincide with those of an agent who, after each round, has a constant probability of δ that her participation continues, and a probability of $1-\delta$ that her participation terminates.[4] As we soon describe in detail, we used this equivalence and varied the continuation probability with which the repeated game proceeded from one round to the next for each player, thereby inducing different discount factors (see Roth and Murnighan, 1978). Specifically, we had four treatments. In the **Equal Low** treatment, both players had a continuation probability of 30%, while in the **Equal High** treatment, both players had a continuation probability of 90%. In the **Unequal Low** treatment, one player had a continuation probability of 40% and the other of 30%, while in the **Unequal Mixed** treatment, one player had a continuation probability of 90% and the other of 30%.

In the main part of the experiment, participants played 10 repeated games that, following the literature, we term *super-games*. Between super-games, the role for each player—relevant in the treatments with unequal discount factors—was randomly determined and players were rematched at random to someone else in the session.

In order to implement different time preferences across players, and achieve comparable data sets across players' types and across sessions, we used the *block random termination method* (see, for example, Fréchette and Yuksel, 2017). Specifically, for any profile (δ_1, δ_2) of two players' discount factors, we used the equivalent presentation of payoffs via different continuation probabilities. Namely, at the end of each round, player i's involvement in a super-game continued with probability δ_i and terminated with probability $1-\delta_i$, $i = 1, 2$.

[3] Detailed experimental instructions are available in the Online Appendix.

[4] See, for example, Section 12 in Mas-Colell et al. (1995) or Section 4 in Mailath and Samuelson (2006).

In all treatments, continuation between one round to the next was determined independently for the two players.[5] Each participant was exposed to one treatment only. If the super-game continued for a player up to a certain round t, we say the player was *active* up to time t, and her super-game payoff was the sum of payoffs up to and including round t. In what follows, we use discount factors and continuation probabilities interchangeably.

In our block design, participants played every super-game in blocks of 12 rounds. Throughout each block, participants were not informed whether their participation had ended or not. At the end of each block of 12 rounds, if the super-game did not terminate for at least one player in the pair, another block of 12 rounds was played. If the super-game terminated for both players in a block, players were informed of the round at which the super-game had terminated for each, and their respective payoffs: the sum of round payoffs across rounds in which each was active.[6] This block design guarantees we have action choices for at least 12 rounds for both players in each super-game across our treatments.[7]

Participants could monitor one another perfectly throughout each super-game: the interface allowed participants to observe and recall the actions played in all previous rounds.

Coordination games are known to be challenging (Crawford, 2016), and the repeated games we study allow for a vast set of equilibrium predictions, thereby compounding the coordination problem. In practice, individuals trying to coordinate, in a household or the workplace, often communicate. Since our goal was to identify which repeated-game equilibrium had strong drawing power, our design enabled participants to discuss a joint strategy beforehand. In particular, at the start of every super-game, each matched pair could communicate freely through a chat interface. When either decided to terminate the communication phase, the game started. Participants did not communicate further during the super-game.

At the start of each session, participants had a five-minute training period. In the training period, participants could specify arbitrary 12-round strategies for themselves and their paired player and observe the resulting expected payoffs were the super-game to end within that block, as well as the realized super-game length for each of the players.

[5] In order to ease participants' understanding of the implications of various continuation probabilities, "cheat sheets" were distributed throughout experimental stations with a table detailing the resulting probability of each game length for both discount factors utilized in the session.

[6] For example, consider the Unequal Mixed treatment and suppose that the player with a continuation probability of 30% was active for 2 rounds, while the player with a continuation probability of 90% was active for 7 rounds. Then, at the end of the first block of 12 rounds, both players would be informed of their active rounds and their resulting payoffs, and proceed to the next super-game. Similarly, suppose the player with a continuation probability of 30% was active for 2 rounds, while the player with a continuation probability of 90% was active for 15 rounds. Then, at the end of the first block of 12 rounds, both players would be informed that at least one of them was still active and proceed to play the second block of 12 rounds, at the end of which they would proceed to the next super-game.

[7] In order for all pairs in a session to progress through the experiment at the same pace, we randomly determined the active rounds for each player in each super-game, using their specific continuation probability, at the session level.

Table 1: Experimental Design

Treatment	Discounting		# Sessions	# Participants
Equal Low	equal	low(30)-low(30)	4	42
Equal High	equal	high(90)-high(90)	4	46
Unequal Low	unequal	low(40)-low(30)	5	56
Unequal Mixed	unequal	high(90)-low(30)	5	60
Total			18	204

After 10 super-games were played, we also elicited risk attitudes using the design of Gneezy and Potters (1997) with two parametrizations, as well as altruism elicitations using dictator games with two specifications.[8] At the end of each session, we asked participants to write free-form descriptions of what guided their behavior. In the Online Appendix, we report the analysis of these reports that, in particular, suggests participants' comprehension of the underlying tasks.

The experiment was run with 204 participants at the Princeton Experimental Laboratory for the Social Sciences (PExL). Participants were paid for the sum of their payoffs in all the super-games, with a conversion rate of \$7.50 for each 100 experimental points. They were also paid for the auxiliary risk and altruism elicitation: one of the two tasks corresponding to each elicitation was randomly selected for payment. Overall, the average payment was \$23.68. In addition, there was a \$10 show-up fee. Table 1 summarizes our treatments, as well as the number of separate sessions and participants in each.

Efficient Equilibria Theoretically, when players' discount factors coincide, any strategy profile generating coordination in each period is utilitarian efficient. When player's discount factors differ, there is beneficial intertemporal trade, whereby the impatient agent receives the high coordination payoff in the first k periods, and the patient agent receives the high coordination payoff later on. With our parameters, to achieve utilitarian efficiency, $k = 1$ in the Unequal Low treatment and $k = 2$ in the Unequal Mixed treatment.[9]

[8]Specifically, in our risk elicitations, participants had to allocate \$2 to a safe project that paid the amount invested with certainty, and a risky project that paid 2.5 times (or 3 times) the amount invested with probability 0.5 (or 0.4). In our altruism elicitations, participants had to allocate \$1 or \$2 between themselves and a random other. At the end of the experiment, we split each session into two distinct pairings, one for each allocation amount. In each allocation task, we randomly implemented the choice of one of each pair members.

[9]Formally, to maximize utilitarian efficiency, players solve the following:

$$max_{\pi_t \in \{1,6\}} \quad 7 + \sum_{t=1}^{\infty} [(1-\delta_P)\delta_P^{t-1} - (1-\delta_I)\delta_I^{t-1}]\pi_t,$$

where δ_P and δ_I are the patient and impatient players' discount factors, respectively, and π_t (or $7 - \pi_t$) is the patient (or impatient) agent's coordination payoff. Optimally, when $[(1-\delta_P)\delta_P^{t-1} - (1-\delta_I)\delta_I^{t-1}] > 0$, players would set $\pi_t = 6$ and otherwise, they would set $\pi_t = 1$. Since the term in the square parentheses crosses 0 only once, the optimal window size k can be readily calculated.

Approach to Data Analysis In what follows, we describe our results for the last 5 super-games of each session. Considering all 10 super-games yields qualitatively similar results, reported in the Online Appendix. We do, however, see more noise and some evidence of learning in the first 5 super-games of our sessions. As we are interested in strategy profiles that pairs pursue, the unit of observation for most of our analysis is a pair in a super-game.[10] We focus our analysis on the first block of 12 rounds within each super-game, when at least some rounds have substantial probability (> 1%) of counting towards either player's payoffs and when we have balanced data for all treatments.[11] In the Online Appendix, we include further analysis of later blocks.

4 Discount Profiles and Coordination Patterns

Small differences in discount factors assist impatient players in coordination, while large differences in discount factors lead to intertemporal trades, as shown in Figure 1. The figure illustrates coordination rates for each of the 12 rounds within super-games' first block. For each combination of players' discount factors, we display the observed frequencies at which players coordinated on an action profile generating the highest payoff to the high-discount factor player, to the low-discount factor player, or miscoordinated, generating a 0 payoff for both. To display analogous figures when players' discount factors coincide, we randomly determine who is labeled the high-discount factor player among the two players in each super-game so that the surplus is split equally.

When players' discount factors coincide, both players' patience affects the rate of miscoordination, as seen in the top two panels of Figure 1. In the Equal Low treatment, miscoordination occurs at a significantly higher frequency than in the Equal High treatment.[12]

When players' discount factors differ, coordination patterns take a different form, as illustrated in the bottom two panels of Figure 1. The modal action profile in the first round of each super-game is the pure equilibrium preferred by the low-discount factor player, regardless of the wedge between the players' time preferences. In the Unequal Low treatment, when both agents are impatient, we see alternation between the pure equilibria throughout the super-game. The difference in discount factors affects the precise sequencing: the more impatient players are far more likely to generate their most preferred pure equilibrium profile in round 1 relative to what we see in the Equal Low treatment.[13] Play-

[10] In the Online Appendix, we consider individual behavior and its limited dependence on risk attitudes and altruism. At the other extreme, we also show that session-level averages exhibit similar comparisons to those we report.

[11] In treatments involving two impatient players, either Equal Low or Unequal Low, all super-games terminated within the first block.

[12] The p-values corresponding to miscoordination rates across the two treatments in the first round, or at least once within the first block, are 0.021 and 0.043, respectively. See details in our Online Appendix.

[13] The resulting coordination rates in the first round are, indeed, statistically different across the two treat-

Figure 1: Coordination Patterns across Treatments

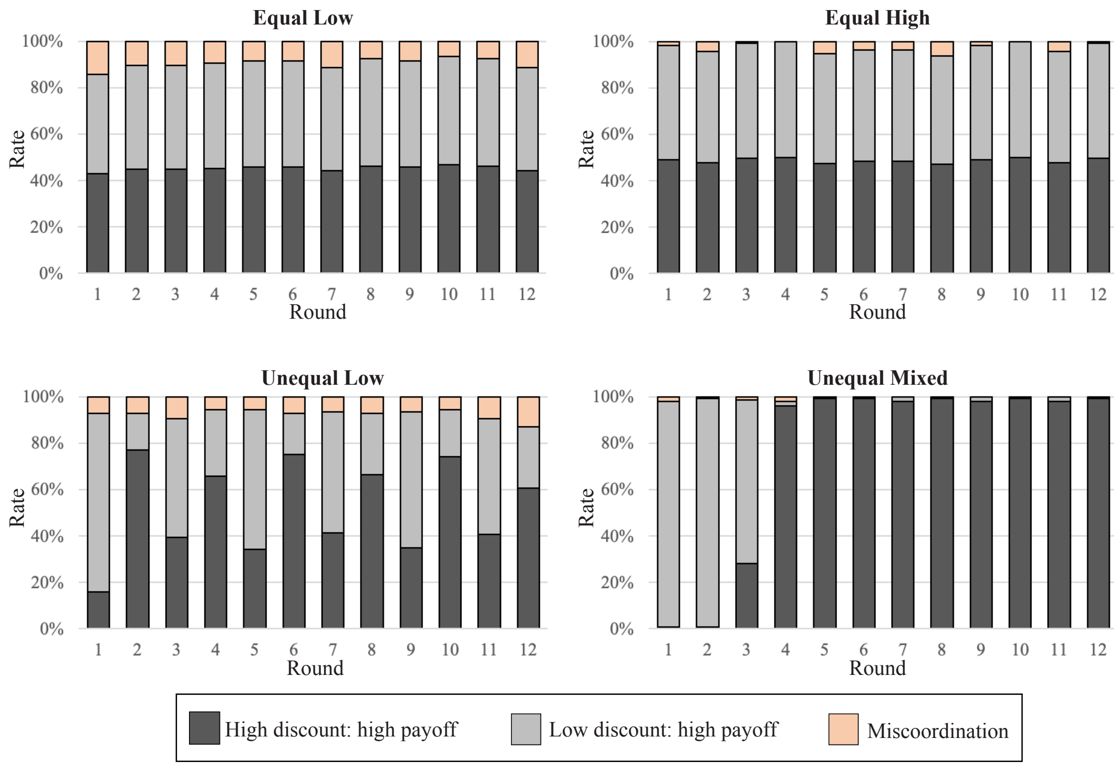

ers' discount factors seem to provide a coordination instrument in and of themselves, as they have a strong impact on which of the alternation equilibria in the super-game gets selected. The alternation profile that provides the impatient player a high payoff initially is the more efficient alternation profile. As we will soon show, the most efficient equilibrium in which the impatient player receives a high payoff only in the first period also has some drawing power. These two profiles—efficient alternation and efficient intertemporal trade—differ only in actions taken starting in round 3. In the Unequal Low treatment, these differences have a small effect on overall expected payoffs. In other words, players get close to achieving efficiency.

In the Unequal Mixed treatment, when the wedge between discount factors is large, we see far less alternation between the pure equilibria throughout the super-game. Less patient agents receive the highest possible payoff for $2-3$ rounds. Afterward, patient players reap the benefits of the pure equilibrium that is more beneficial for them. This coordination pattern is in line with the theoretically-predicted efficient intertemporal trade.[14]

ments, with a p-value of 0.043.

[14]The rates of coordination in the Unequal Mixed treatment are not statistically different from those seen in

Figure 2: Efficiency and Surplus Distribution

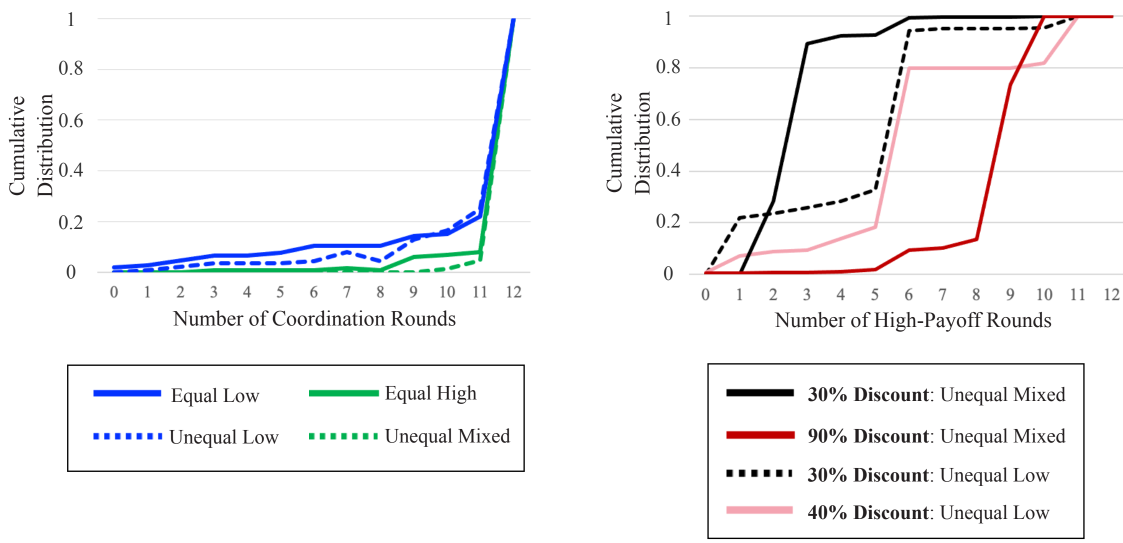

The patterns seen in Figure 1 are echoed when considering outcomes across supergames, as shown in Figure 2. The left panel of Figure 2 presents the cumulative distributions of the number of times each pair successfully coordinated within the first block of the game. We see the least coordination when both agents are impatient, and coordination rates increase with the wedge between agents' patience. Treatments in which both players are fairly impatient—the Equal Low and Unequal Low treatments—yield the least coordination. Their corresponding distributions are first-order stochastically dominated by those corresponding to treatments in which at least one of the players was fairly patient. Furthermore, the distribution corresponding to the Unequal Mixed treatment first-order stochastically dominates all others.[15]

The right panel of Figure 2 illustrates that when agents' discount factors differ, the more patient agents receive the high coordination payoff for longer durations within super-

the Equal High treatment. However, they are significantly higher than those seen in the Equal Low treatment: the p-values corresponding to miscoordination rates across the two treatments in the first round, or at least once within the first block, are 0.002 and 0.004, respectively.

[15]To determine the statistical significance of these comparisons, we use the test proposed by Barrett and Donald (2003). We cannot reject the null that the distributions generated by the Unequal Mixed or Equal High treatments dominate that of the Equal Low treatment (p-value > 0.845), while we strongly reject the reverse null (p-value = 0.001). We also cannot reject the null that the distributions generated by the Unequal Mixed or Equal High treatments dominate that generated by the Unequal Low treatment (p-value > 0.851), while we strongly reject the reverse null (p-value = 0.001). We cannot reject the null that the distribution in the Unequal Mixed treatment dominates that in the Equal High treatment (p-value = 0.849), while we reject the reverse null (p-value = 0.013). Kolmogorov-Smirnov tests comparing distributions yield qualitatively similar results.

games. The panel displays the cumulative distributions corresponding to the number of rounds each player receives the high coordination payoff within the first block of super-games. In both the Unequal Low and the Unequal Mixed treatments, the distributions associated with the more patient agents first-order stochastically dominate the distributions associated with the impatient agents. The distance is substantially more pronounced when the wedge in discount factors is large. In particular, in the Unequal Mixed treatment, a large majority of the low, 30% discount factor players receive the high coordination payoff for fewer than 3 rounds and, accordingly, a large majority of the 90% discount factor players receive the high coordination payoff for more than 9 rounds. In the Unequal Low treatment, the modal distance between the number of rounds players receive the high coordination payoff is lower than 1.

5 Strategies Utilized

The previous section illustrated the outcome patterns emerging from different players' patience profiles. We now consider the strategies that underlie these outcomes. We show that players both communicate suggestions for and play predominantly two classes of strategy profiles, ones involving turn-taking and ones involving intertemporal trades. Players' patience profile governs which of the two prevails and how it gets implemented.

Formally, a general *turn-taking* strategy profile corresponds to a profile in which players alternate between blocks of k consecutive rounds in which one of the two coordination outcomes is implemented. One instance of such a strategy profile is a simple turn-taking strategy, where $k = 1$, and players alternate between the two coordination outcomes. An *intertemporal-trade* strategy profile is one in which one player receives the high coordination payoff for the first k rounds, and the other player receives the high coordination payoff for the remaining rounds.[16] As in the previous section, we inspect strategies in the first block of super-games, where we have comparable volumes of data across treatments. Similar patterns emerge for treatments in which further blocks occurred.

Communication Protocols In each of our treatments, non-trivial communication occurs in more than 98% of the last 5 super-games.

In the Equal Low treatment, getting a high payment in the first round is extremely valuable for both players: there is a 70% chance that the super-game ends for either player after that round. This generates a tension between the players, and 10% of conversations culminate in no agreement. Of the pairs who reach an agreement through conversation,

[16] In general, intertemporal trade can take different forms; see, for example, Lehrer and Yariv (1999). We focus on these forms of intertemporal trade since they are prevalent in the data and yield the utilitarian efficient outcomes in our Unequal Low and Unequal Mixed treatments.

72% plan to play a simple turn-taking strategy profile (with blocks of length $k = 1$) and 21% plan on a simple intertemporal-trade strategy profile (with $k = 1$), whereby one of the players receives the high coordination payoff only in the first round.[17]

In the Equal High treatment, there is less tension between players, and 98% of conversations yield agreement. Of those, 95% specify a turn-taking strategy profile, and 89% specify a simple turn-taking strategy profile, alternating between coordination outcomes.

The wedge in discount factors affects the communication protocols we observe. In the Unequal Low treatment, over 99% of discussions reach an agreement. Of those, 53% focus on simple turn-taking strategy profiles. In 85% of these, players agree to grant the more impatient player the high coordination payoff in the first round. In 27% of conversations, players agree on a simple intertemporal-trade strategy profile, and in 81% of these, the lower discount-factor player receives the high coordination payoff in the first round only. Thus, relative to the Equal Low treatment, in the Unequal Low treatment, the difference between continuation probabilities serves as a coordination device: players more frequently agree on who receives the high coordination payoff first and, throughout the super-game, coordination improves.

Virtually all conversations in the Unequal Mixed treatment reach an agreement, and 96% of discussions spell out a plan for an intertemporal-trade strategy profile. There is some variation in the precise intertemporal-trade strategy profile players agree on. In 30% of cases, players agree on the more impatient player receiving the high coordination payoff for the first $k = 2$ rounds; in 69% of cases, the agreed-upon block is of $k = 3$ rounds; while in the remaining 1% of conversations, the initial block of $k = 4$ rounds is agreed upon.

Observed Action Profiles Consistent with the strategy profiles discussed during the communication phase, observed action profiles, by and large, correspond to either turn-taking or intertemporal-trade strategies, as shown in Table 2. The table displays the frequency of observed action profiles in sessions' first block of 12 rounds.

Turn-taking profiles are the modal outcome when players' discount factors coincide.[18] They are particularly high when agents are both patient (91% relative to 53% when players are both impatient). We see cycles of length $k = 2$ or $k = 3$ only in 6% of turn-taking strategies in the Equal High treatment. Otherwise, players tend to use simple turn-taking strategy profiles (with $k = 1$).[19] Intertemporal-trade profiles are the modal outcome in

[17] The remaining conversations are spread between explicit specifications of play across rounds and turn-taking with longer blocks.

[18] In the Online Appendix, we show that there is limited evidence for individual "types"—innate features affecting behavior—by looking at the distributions of individuals' tendencies to concede in the first round of each super-game.

[19] Our classification of turn-taking strategies is demanding since we require players to alternate as prescribed in all 12 rounds. In the Equal Low treatment, in 82%, 56%, and 55% of games, players alternated between the two coordination outcomes across the first 2, 4, and 6 rounds, respectively.

Table 2: Distribution of Strategy Profiles

Treatment	Turn-Taking frac	Turn-Taking chat if play	Intertemporal-Trade frac	Intertemporal-Trade chat if play	length | play 1	2	3	Other Strategies frac	round 1 miscoor	round 1 high-low
Equal Low	0.53	0.95	0.19	0.90	1.00			0.28	0.52	
Equal High	0.91	1.00	0.00					0.09	0.20	
Unequal Low	0.39	1.00	0.22	0.97	1.00			0.39	0.18	0.67
Unequal Mixed	0.01	1.00	0.93	0.99		0.29	0.70	0.06	0.33	0.67

Notes: "Chat if play" corresponds to how often pair members discussed playing the relevant strategy conditional on playing it. In the last two columns, we report two characteristics of other types of strategies played: the rate of miscoordination and the frequency of allocating a high payoff to a low-discount factor player in the first round.

the Unequal Mixed treatment, appearing in 93% of super-games, with the impatient agent receiving the high coordination payoff in the first $k = 2$ or $k = 3$ rounds, approximating the efficient equilibrium profile. Intertemporal-trade profiles also appear at high rates in treatments with at least one impatient agent, although they take a different form: in the Equal Low and Unequal Low treatments, impatient agents receive the high coordination payoff for only one round ($k = 1$), again approximating the efficient equilibrium profile. In the Online Appendix, we show that similar profiles emerge even in the first 5 super-games of sessions, suggesting they have particularly strong drawing power.

Players' communication is, by and large, in line with their behavior. Whenever players play either a turn-taking or an intertemporal-trade strategy profile, they agree upon it during the communication phase in at least 90% of cases. The greatest discrepancies between behavior and communication agreements occur in the Equal Low treatment, where there is no focal player and both players have a strong incentive to receive high payoffs early on.

When players follow other strategies, miscoordination rates are substantial. Interestingly, even with those strategies, in treatments with asymmetric patience across the players, there is a strong draw for the more impatient agent to receive the high coordination payoff first (67% of cases).

From Table 2, in the Equal High and Unequal Mixed treatments, there is prevalent use of turn-taking and intertemporal-trade strategies, respectively. Furthermore, we see no deviations when players agree to play one of these profiles. Therefore, we cannot inspect off-path behavior. In the Equal and Unequal Low treatments, we do see some deviations from agreements. Specifically, of those pairs that agreed to play the modal strategy profile of simple turn-taking, 7 pairs in each of the treatments had one player deviate within the first two rounds of a super-game.[20] In 11 of these—6 out of 7 in the Equal Low treatment and 5 out of 7 in Unequal Low treatment—the partner of the deviating agent uses the

[20] We focus on super-games where players agree on a strategy profile since, otherwise, we cannot determine whether seeming deviations reflect true deviations, mistakes, or attempts to coordinate on a different profile.

Figure 3: Current Round Payoffs Conditional on Previous Round Payoffs

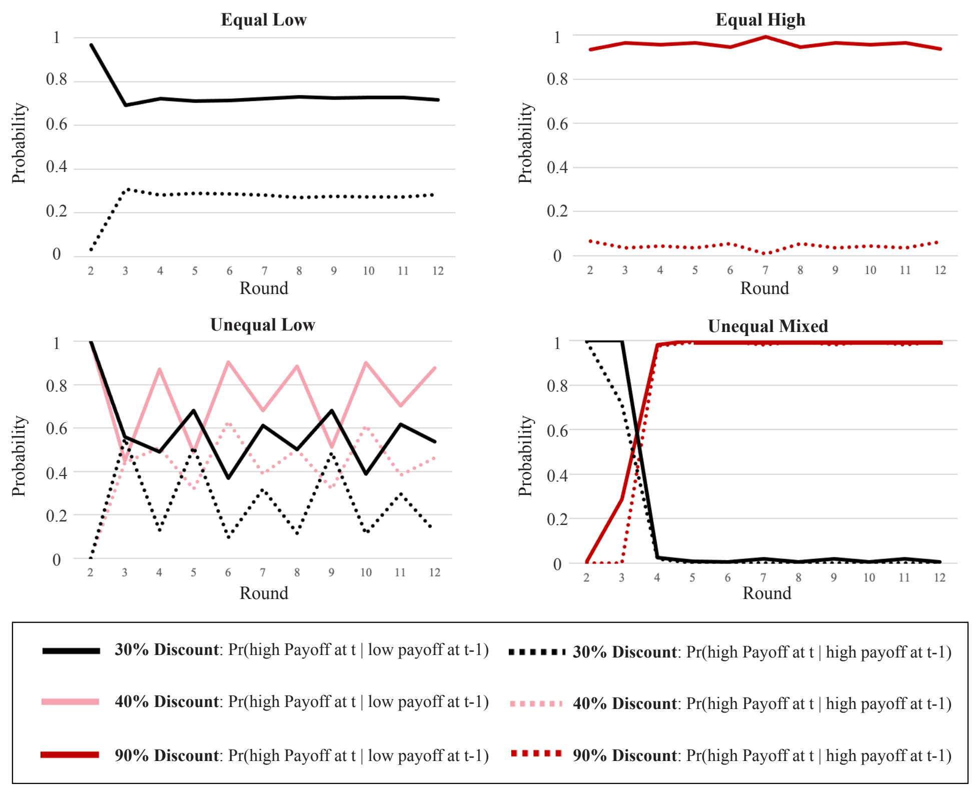

action that would generate the high-coordination payoff for herself, the minmax strategy against the deviating agent, in the following round.

Figure 3 depicts how the profile of discount factors affects the round-to-round choices players make. The figure displays the frequency with which each player receives the high coordination payoff at any round t conditional on receiving either a low or a high coordination payoff at the preceding round $t-1$, when players coordinate in round t as well.[21]

In the Equal Low treatment, alternation is nearly perfect in the first two rounds. Starting in round 3, alternation is partial, and responses are constant from that round on. In the Equal High treatment, alternation is nearly perfect throughout the first block of play.

In the Unequal Low treatment, the rates with which players receive the high coordination payoff are always higher when players' previous payoffs were low. In addition, the higher, 40% discount factor players consistently receive the higher coordination payoffs after round 3. Nonetheless, play patterns are more haphazard, reflecting the substantial diversity of strategy profiles players use.

In the Unequal Mixed treatment, we see a crossover in the fraction of high coordination

[21] We use the shorthand "Pr" for probability to reflect the observed frequency.

payoffs received by patient and impatient agents. In rounds 2 and 3, it is the impatient agents who receive the high payoffs, regardless of what transpired in previous rounds. This reverses starting from round 4.

Figure 3 suggests that, in our treatments, much of the variability in play occurs in the first few rounds. In fact, coordination in the first two rounds is highly indicative of successful coordination in later rounds.[22]

Fairness and Equity Considerations An abundance of work documents lab participants' draw to egalitarian outcomes; see, for instance, Gächter and Fehr (2002). In principle, social concerns could play a role in our experiment. However, they do not appear to organize the data in a meaningful fashion. First, as noted, we show in the Online Appendix that elicited altruistic tendencies have limited power in explaining behavior. Second, while miscoordination produces perfectly egalitarian outcomes round by round—both players receive a payoff of 0—it is rarely utilized. Third, straightforward calculations of payoffs indicate that in the Unequal Low treatment, both efficiency and equity concerns should drive participants to choose the intertemporal-trade strategy, but participants play the alternation strategy more frequently. In the Unequal High treatment, intertemporal trade yields arguably less equitable outcomes, but is the predominant choice. Last, participants' free-form descriptions of their choices mention such considerations at low rates.[23]

6 Conclusions

We provide a first look at the consequences of differential discount factors in a complete-information repeated coordination game using a sequence of lab experiments.

When players' discount factors coincide, they tend to alternate across the stage-game coordination outcomes. However, with lower coincident discount factors, both players have a strong incentive to receive a high payment first. Perhaps as a consequence, we see more coordination failures when both players are identically and severely impatient.

A discount factor differential has two effects. First, it provides a form of coordination instrument: the more impatient agent frequently receives the highest possible payoff first, even when discount factors are close to one another. Second, players use intertemporal

[22] In the Equal Low, Equal High, Unequal Low, and Unequal Mixed treatments, in the first super-game block, the number of miscoordination rounds conditional on first-round coordination is 0.32(0.20), 0.24(0.14), 0.68(0.36), and 0.03(0.02), respectively (robust standard errors reported in parentheses, where clustering is at the session level). Conditional on miscoordination in the first round, these numbers are 6.13(1.56), 5.00 (with too few observations to calculate standard errors), 3.70(0.54), and 1.67(0.22), respectively.

[23] We provide detailed analysis in the Online Appendix. As we show, fairness and equity considerations were mentioned in fewer than 15% of participant reports in all treatments other than the Equal Low treatment, when they were noted in roughly 30% of reports. In fact, in the Unequal Mixed treatment, they were noted by fewer than 5% of participants.

trades, whereby impatient agents receive high payoffs early on and patient agents receive high payoffs in all later rounds.

Our results indicate that the knife-edge case of equal discount factors, which has been the main focus of study in the literature, may miss important aspects of behavior. Nonetheless, even with a discounting differential, efficient cooperation has a strong draw.

Our study opens the door to further investigations of the consequences of differential time preferences. For example, a larger-scale study considering a variety of games, covering a richer set of discount factor profiles could be useful for generalizing and refining the insights we offer. It would also be interesting to consider the effects of various communication possibilities, ranging from no communication at all to more frequent communication prior to every round in the super-game.

References

Barrett, G. F. and S. G. Donald (2003). Consistent tests for stochastic dominance. *Econometrica 71*(1), 71–104.

Cason, T. N., S.-H. P. Lau, and V.-L. Mui (2013). Learning, teaching, and turn taking in the repeated assignment game. *Economic Theory 54*(2), 335–357.

Chen, B. and S. Takahashi (2012). A folk theorem for repeated games with unequal discounting. *Games and Economic Behavior 76*(2), 571–581.

Crawford, V. P. (2016). New directions for modelling strategic behavior: game-theoretic models of communication, coordination, and cooperation in economic relationships. *Journal of Economic Perspectives 30*(4), 131–50.

Cripps, M. W. and J. P. Thomas (2003). Some asymptotic results in discounted repeated games of one-sided incomplete information. *Mathematics of Operations Research 28*(3), 433–462.

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

Ellickson, R. C. (1991). *Order without law: How neighbors settle disputes*. Harvard University Press.

Fréchette, G. R. and S. Yuksel (2017). Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics 20*(2), 279–308.

Gächter, S. and E. Fehr (2002). Fairness in the labour market: A survey of experimental results. In *Surveys in experimental economics: Bargaining, cooperation and election stock markets*, pp. 95–132. Springer.

Gneezy, U. and J. Potters (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics 112*(2), 631–645.

Guéron, Y., T. Lamadon, and C. D. Thomas (2011). On the folk theorem with one-dimensional payoffs and different discount factors. *Games and Economic Behavior 73*(1), 287–295.

Jackson, M. O. and L. Yariv (2014). Present bias and collective dynamic choice in the lab. *American Economic Review 104*(12), 4184–4204.

Jackson, M. O. and L. Yariv (2015). Collective dynamic choice: the necessity of time inconsistency. *American Economic Journal: Microeconomics 7*(4), 150–178.

Kreps, D. M. and R. Wilson (1982). Reputation and imperfect information. *Journal of Economic Theory 27*(2), 253–279.

Kuzmics, C., T. Palfrey, and B. W. Rogers (2014). Symmetric play in repeated allocation games. *Journal of Economic Theory 154*, 25–67.

Lehrer, E. and A. Pauzner (1999). Repeated games with differential time preferences. *Econometrica 67*(2), 393–412.

Lehrer, E. and L. Yariv (1999). Repeated games with incomplete information on one side: The case of different discount factors. *Mathematics of operations research 24*(1), 204–218.

Mailath, G. J. and L. Samuelson (2006). *Repeated games and reputations: long-run relationships*. Oxford University Press.

Mas-Colell, A., M. D. Whinston, J. R. Green, et al. (1995). *Microeconomic theory*, Volume 1. Oxford University Press.

Milgrom, P. and J. Roberts (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory 27*(2), 280–312.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.

Romero, J. and H. Zhang (2018). Normative conflict and history dependence in repeated coordination games.

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

Schaner, S. (2015). Do opposites detract? intrahousehold preference heterogeneity and inefficient strategic savings. *American Economic Journal: Applied Economics 7*(2), 135–174.

Tingley, D. H. and B. F. Walter (2011). The effect of repeated play on reputation building: an experimental approach. *International Organization 65*(2), 343–365.