Games with Continuous-Time Experimental Protocols

Alexander L. Brown

Daniel G. Stephenson

In many cases, traditional laboratory environments provide the cleanest and clearest tests of economic theory. An experiment can be designed to only match the relevant assumptions of the theory, leaving out all else. If one, for instance, wanted to examine equilibrium selection for a one-shot 2x2 game, one could have individual subjects play strangers in an anonymous environment and earn money proportional to payoffs in the game. Of course, the assumptions of full rationality and induced preferences may not hold. Players might not fully understand how their strategies map into payoffs or they may have preferences that differ from their monetary payoffs in the game. However, a general finding in game theory experiments is that individuals tend to perform more consistently and closer to theoretical predictions after repeated play (Davis and Holt, 1993). The general thought is that subjects gain a better understanding of their strategic environment with repetition, so repetition is often crucial if one wants to test a theory of long run strategic behavior.

Of course, practical concerns limit the number of repetitions that can be observed in the laboratory. It is often infeasible to hold subjects in a laboratory for more than three hours (Friedman and Sunder, 1994). Further, subjects may lose attention well before that time limit is reached. Given that each one-shot period can take a significant number of seconds, there is an upper bound on the number of discrete periods that can be implemented in a laboratory experiment. Continuous-time experimental protocols provide an alternative. They allow subjects to take actions and receive real-time feedback continuously. By allowing subjects the freedom to adjust their actions as frequently as desired and providing them with real-time feedback, continuous-time experiments effectively reduce the length of discrete periods to near instants, accelerating the emergence of long-run behavioral patterns and providing unique insights into the dynamic processes through which such behavioral patterns emerge.

As our short description indicates, it is difficult to succinctly and incontrovertibly define continuous time. In theory, the definition is simple enough: time is defined on a continuous (rather than discrete) interval. However, delays in human perception and reaction time may invariably introduce time-lags to human interaction. That is, there is a minimum amount of time that an individual takes to perceive and react to new information. Recent research indicates this human limitation may have strong implications for some continuous-time games (see Calford and Oprea, 2017, for example). Even in principle, absolutely instantaneous interaction may be infeasible, since even computerized interfaces take some time to perform calculations and provide feedback. Hence no experiment can perfectly implement the theoretical ideal of fully continuous interaction. In this chapter we consider a very broad definition. We consider a continuous-time experiment to be any experiment where interaction occurs with sufficiently high frequency for subjects to perceive it as continuous.

Continuous-time experimental protocols often investigate the long-run patterns of behavior that emerge after individuals become accustomed to a strategic environment through many repeated interactions. By providing subjects with a large quantity of interactive feedback over relatively short periods of time, continuous-time experiments can provide a robust way to examine the emergence of behavioral patterns that would be difficult to observe under conventional discrete-time protocols. The idea that strategic interaction in a population of individuals will ultimately settle down into some consistent pattern of behavior is intuitively appealing and has long been a feature of game theoretic reasoning (see Nash, 1950). Continuous-time experimental protocols provide one of the most powerful ways to test theoretical models of this convergence process.

The findings of continuous-time experiments are varied. Continuous-time experiments that investigate strategic environments with multiple equilibria frequently find that behavior reliably converges to particular equilibria more often than others, providing an experimental test of equilibrium selection models. In some strategic environments, continuous-time interaction is found to encourage the emergence of persistent collusive behavior, providing experimental evidence on the emergence of cooperation. In contrast, in other environments, continuous-time features may promote greater competition, removing collusive opportunities. By collecting real-time data on how subjects choose to adjust their behavior over time, continuous-time experiments often shed new light on the adaptive processes that drive behavioral change over time. Subject behavior in continuous-time experiments is often found to be sensitive to a wide variety of experimental design factors including the matching protocol, the interface through which subjects adjust their strategies, and the way that feedback is provided to subjects.

In this chapter, we will first identify key methodological features common to experiments in continuous-time. We will document and categorize a wide range of continuous-time experiments, noting the key features of the experimental design and results, while also noting the key findings that would have been impossible to observe without continuous time. Finally, in our discussion section, we identify key regularities in results of continuous-time experiments.

1 Methodology

There are certain methodological features that are generally found in many continuous-time experiments. Here, we note four: continuous feedback, continuous payoffs, mean-matching, and restarts between continuous-time periods.

1.1 Continuous feedback

Many continuous-time experiments feature continuous feedback. In this case, all participants observe the current state of the game, which can include the actions of participants and their payoffs. Computer interfaces often utilize dials or slider bars to provide subjects with a default initial position and a way to change their actions as quickly and easily as possible. Randomized initial positions are sometimes utilized so that all subjects have some action selected immediately once the game begins. Figure 1 provides an example of the interface in Stephenson and Brown (2020). Here subjects in an all-pay auction choose bid amounts from the interval $[0, 10]$. Subjects could adjust their bid at any time by clicking on the screen. The interface provides their "instantaneous payoff" (i.e., their payoff at the current moment given their current strategy and all other players' strategies.) The interface also provides the instantaneous payoffs they could currently earn from all other actions (darker line).

1.2 Continuous payoffs

Because the state of the game changes continuously, specific payoff protocols are often employed. Continuous-time games often utilize instantaneous payoffs that are determined by every subject's action and any other relevant state variables at a given moment. In some experimental designs, these payoffs are integrated over the course of the experiment and converted to cash for each subject at the end of the experiment. For instance, a subject who earns \$10 for 90 seconds and \$15 for 30 seconds of a 2-minute period might receive the time-average of \$11.25. In other experimental designs, a subject's earnings are based exclusively on her instantaneous payoff at the end of each period.

1.3 Matching protocols

Experiments that feature the repetition of a one-shot game over several discrete periods often employ random anonymous matching of subjects. This protocol aligns subject incentives with the one-shot game while still allowing a subject to experience

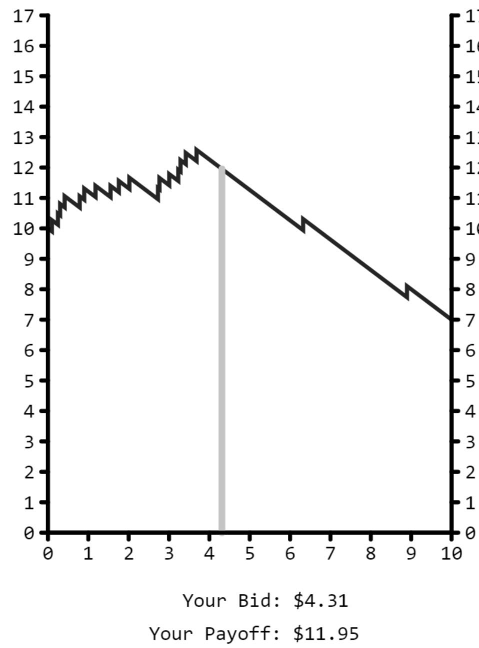

Figure 1: A subject interface in Stephenson and Brown (2020). At any time, subjects could freely select any bid from 0 to 10 by clicking on the graph. Players observe a visual representation of the their current bid, their current instantaneous payoff, and the current instantaneous payoffs to all their feasible strategies (darker line). These values all adjust immediately as soon as any other player in the game adjusts their bid.

multiple periods. Note if there are n subjects in a cohort, a subject's expected payoff from taking a particular action in a single period equals her average payoff from being matched against all other subjects, since the subject is equally likely to be matched against each of the other subjects. As an alternative, *mean-matching* provides this expected payoff directly as earnings rather than randomly selecting a single subject as the opponent. Continuous-time experiments often utilize mean-matching instead of random matching, since random rematching at sufficiently high frequencies closely approximates mean matching by the law of large numbers. Mean-matching also has the added benefit of aiding the convergence process by providing subjects with direct feedback regarding the performance of their strategy against the entire population of opponents.

1.4 Restarts between continuous time periods

Continuous-time experiments often divide periods of continuous interaction into discrete segments by restarting the game at the beginning of each period. These restarts help identify the extent to which convergence occurs within or across games. Across these segments of continuous interaction, conventional random, fixed, or stranger matching procedures may be employed. Since subjects interact at high frequency, continuous-time experimental protocols often achieve convergence to a stable pattern of behavior over relatively short periods of time. Accordingly, continuous-time experiments often implement a relatively small number of restarts between continuous time periods of only a few minutes each.

2 A survey of continuous-time experiments

2.1 The precursor: long-run convergence under discrete-time protocols

Before the emergence of continuous-time experimental protocols, some discrete-time experiments were run with a large number of rounds to test the long-run equilibrium convergence properties of certain games. A good example is Battalio, Samuelson, and Van Huyck (2001) who test equilibrium selection criteria in a variety of stag-hunt games, all of which possess identical best-response correspondences. The payoff matrices for these stag-hunt games are provided by Figure 2. They conduct 75 discrete periods per session. They identify the importance of the optimization premium, the difference between the payoff from employing a best response and the payoff from employing an inferior strategy. They find that subjects respond to large payoff differences more strongly than small payoff differences. That is, over time payoff dominant equilibria are more likely to persist over time in games with a smaller optimization premium.

	Stag	Hare
Stag	45,45	0,35
Hare	35,0	40,40

	Stag	Hare
Stag	45,45	0,40
Hare	40,0	20,20

	Stag	Hare
Stag	45,45	0,42
Hare	42,0	12,12

Figure 2: The stag hunt games investigated by Battalio, Samuelson, and Van Huyck (2001). The expected payoff in the mixed equilibrium of all three games is 36. However, the optimization premiums are larger in the left hand game than in the center game and larger in the center game than in the right hand game.

2.2 Canonical games in continuous time

2.2.1 Coordination games

The aforementioned work of Battalio, Samuelson, and Van Huyck (2001) was the culmination of a decade of work in examining equilibrium selection in coordination games (e.g., Van Huyck et al., 1990, 1991, 1997). This type of work is ripe for analysis in continuous time because many theoretical models, such as evolutionary game theory (see Smith, 1982, Sandholm, 2010) view equilibrium as the outcome of a long-run convergence process. While Battalio, Samuelson, and Van Huyck (2001) have 75 rounds, a continuous-time experiment could feature thousands of very short rounds to examine long-run properties more robustly.

Along those lines, Deck and Nikiforakis (2012) conduct experiments on six-player minimum effort games in continuous time. They build on previous experimental investigation of the minimum effort game in discrete time (see Van Huyck et al., 1990). Throughout each period, subjects choose effort levels from 1 to 7. This game has seven distinct Nash equilibria; in each equilibrium all subjects coordinate on a single effort level. The risk dominant equilibrium has every player exert effort 1, but the Pareto optimal Nash equilibrium has every player exert effort 7. Their experiment involves 10 periods, each lasting for 60 seconds, during which subjects are free to adjust their strategy. The payoffs from each period are calculated based on the effort levels selected by subjects at the end of each period. Their treatments vary the feedback that subjects receive during the 60 second adjustment period.

In their no monitoring treatment, subjects do not observe the effort levels currently selected by any of their neighbors during the adjustment period. In their imperfect monitoring treatment, subjects can observe the effort levels currently selected by their immediate neighbors in a circle network throughout the adjustment period. In their perfect monitoring treatment, subjects can observe the actions currently selected by all of their group members throughout the adjustment period. In the last five periods, the authors observe significantly higher effort levels under the perfect monitoring treatment than under the imperfect monitoring or no monitoring treatments. In contrast, the imperfect monitoring treatment exhibited only slightly higher effort levels than the no monitoring treatment. The authors suggest that strategic uncertainty made coordination more difficult in the imperfect monitoring treatment than the perfect monitoring treatment.

The results observed by Deck and Nikiforakis (2012) in the continuous-time, minimum effort game are strikingly different from those observed in the discrete-time analogue implemented by Van Huyck et al. (1990), where the vast majority of subjects coordinated on the secure strategy of selecting the lowest possible effort level in the discrete-time minimum effort game. Van Huyck et al. (1990) suggested that these relatively low effort levels were the result of strategic uncertainty. In contrast, subjects exhibited significantly higher effort levels in the continuous-time, minimum effort game implemented by Deck and Nikiforakis (2012). As illustrated by Figure 3 this effect is most pronounced in the perfect monitoring treatment. This result suggests that the perfect monitoring may have helped reduce the level of strategic uncertainly experienced by subjects in the minimum effort game by providing them with continuous feedback on the effort levels selected by all other subjects throughout the selection period.

Since earnings are calculated exclusively based on instantaneous payoffs at the end of each period and subjects are free to adjust their strategy at any time during the period, effort levels selected before the end of the period have no effect on payoff and could be described as mere cheap talk. Accordingly, one might hypothesize that providing feedback on behavior during the adjustment period would have no significant effect on behavior. However, the results of Deck and Nikiforakis (2012) suggest otherwise. Subjects exerted significantly higher effort levels and earned

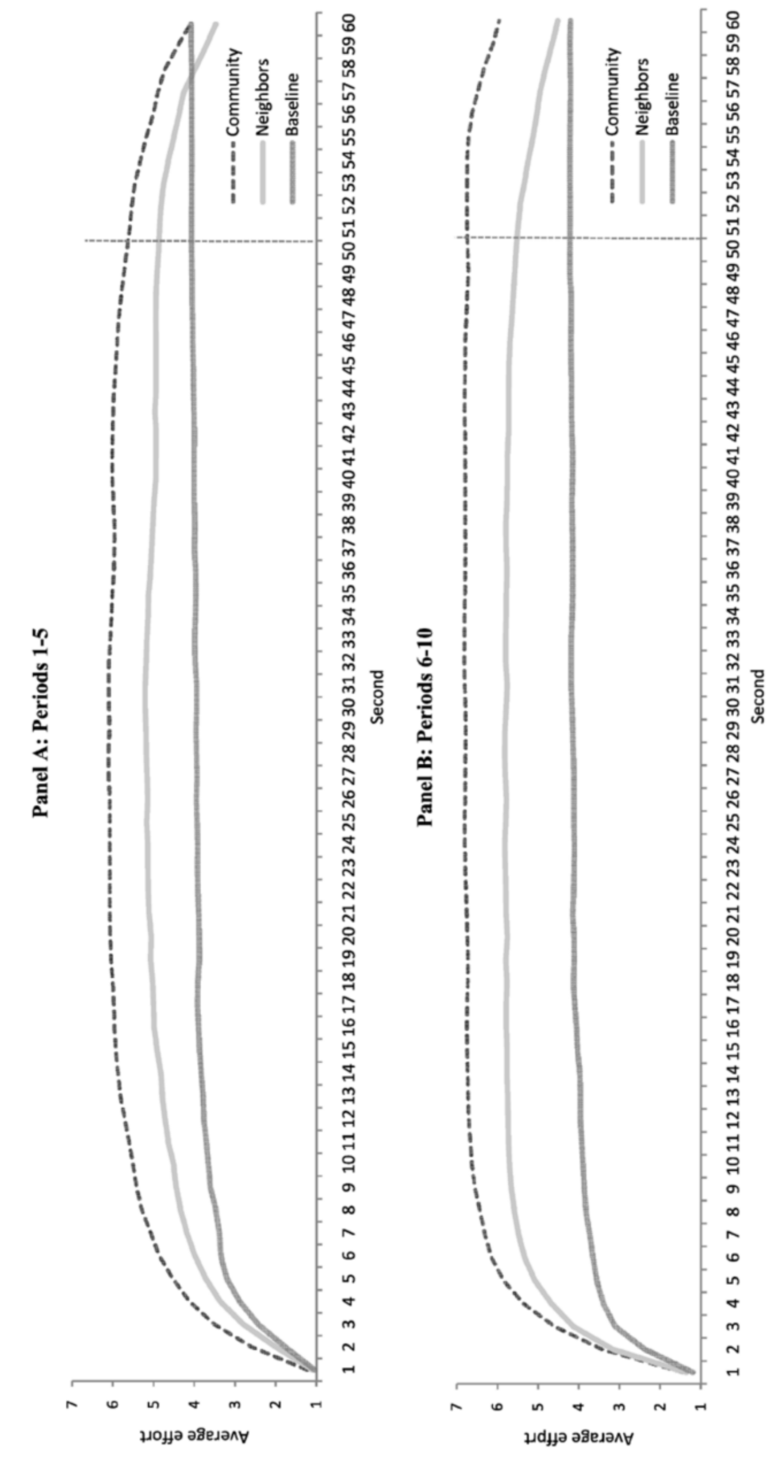

Figure 3: Average effort levels over time within a period of the minimum effort game implemented by Deck and Nikiforakis (2012). The dashed grey line indicates the start of the last 10 seconds within a period.

significantly higher profits in the perfect monitoring treatment, suggesting that they gained valuable experience from interaction and feedback during each adjustment period.

In the currency attack game of Morris and Shin (1998), each player chooses whether or not to participate in a currency attack. Participating in the attack is costly, but if enough players participate in the attack, then it is successful and all attackers receive a large reward. Like the stag-hunt game, this game has multiple equilibria because it is only beneficial to participate in the attack if other players also participate. In one equilibrium, all players participate in the attack. In the other, none of the players participate in the attack.

Cheung and Friedman (2009) conduct laboratory experiments on currency attack games in continuous time. Each subject was assigned a mass and observed a time varying exogenous threshold, which started at 110% of the total player mass. This threshold decreased over time at a constant linear rate to a predetermined value and then remained constant. At any time during a period, subjects could freely switch between passive mode and attack mode. An attack succeeded when the total mass of all attackers exceeded the threshold. Until an attack is successful or the period ends, subjects earned continuous flow payoffs at 100 points per period from passive mode and zero from attack mode. At the first moment an attack is successful, all of the attackers receive an instantaneous bonus of 100 points. For the remainder of a period after a successful attack, both attack mode and passive mode earned continuous flow payoffs of 100 points per period. Figure 4 illustrates the experimental interface that subjects used during the experiment.

The payoff dominant equilibrium for this game has all players simultaneously switching from passive mode to attack mode at the first moment that a successful attack becomes feasible, when the threshold falls under the total player mass. This game also has a wide variety of less efficient equilibria under which all players simultaneously enter at some later point. Each session consisted of 45 periods, each lasting 90 seconds. The experimental design implemented treatments that vary several parameters of the speculative attack game, including the number of players,

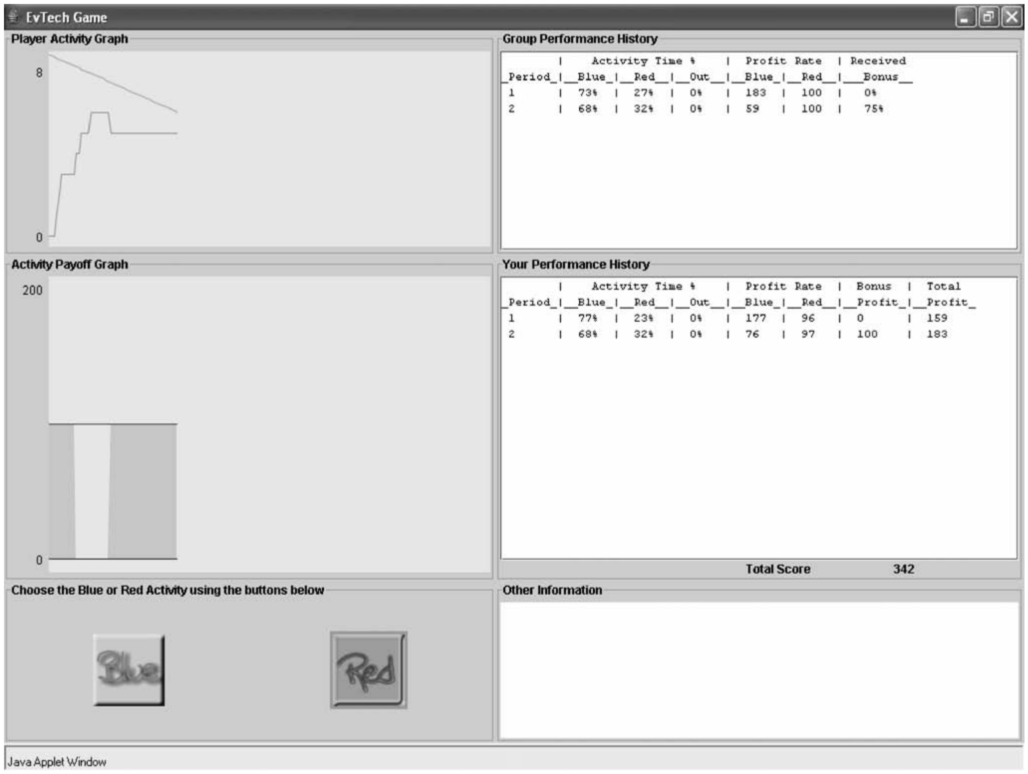

Figure 4: The experimental interface implemented by Cheung and Friedman (2009). Each subject had two buttons, Red and Blue. Red was the passive mode and blue was the attack mode. In the information treatment, the "Player Activity Graph" section of the screen showed the current threshold as the straight line and the total mass of the current attackers as the jagged line. The "Activity Payoff Graph" section of the screen showed the continuous flow payoff rates to each strategy in real time.

the path of the exogenous threshold, the distribution of mass across subjects, and the information players could observe about the actions of others. The authors find that the probability of a successful attack was significantly higher when subjects could continuously observe both the time-varying threshold and the actions selected by others. Attacks were also more likely to be successful against lower thresholds. The presence of a single subject with a larger proportion of the total mass also significantly increased the likelihood of a successful attack over the case where all subjects had equal mass. The authors also found that attacks were significantly more likely to be successful and succeeded significantly earlier when subjects could directly observe the current total mass of players in attack mode throughout each adjustment period.

Each period of Cheung and Friedman (2009) implemented a single dynamic game that proceeds in continuous time. In contrast, each perfect monitoring period of Deck and Nikiforakis (2012) continuously repeated a stage game. The key difference here is a bit like a Markov property. In Deck and Nikiforakis (2012), an individual's instantaneous payoff at time t depends only on the strategy profile at time t, so their subjects effectively played the minimum effort game many times during each period, but were only paid for the last repetition. On the other hand, in Cheung and Friedman (2009), an individual's instantaneous payoff at time t depends on the history of play, in particular whether an attack had succeeded earlier in the period. Hence their subjects effectively play their dynamic currency attack game only once each period. Since the subjects of Deck and Nikiforakis (2012) experienced a larger number of repetitions, they achieved stronger convergence to long-run behavioral patterns. Nonetheless, both studies find that subjects were significantly more successful at coordinating on efficient equilibria when they could observe others in continuous time.

2.2.2 The hawk-dove game

The hawk-dove game, also known as the game of chicken, is one of the simplest symmetric games with both symmetric and asymmetric Nash equilibria. Figure 5 provides the payoff matrix of a hawk-dove game. Evolutionary game theory makes well-known predictions regarding the relative stability of these equilibria depending on the matching protocol employed. The symmetric mixed strategy Nash equilibrium is stable when players are drawn from a single population of agents. Conversely, the asymmetric pure strategy Nash equilibria are stable when the first player is drawn from one population and the second player is drawn from another population. Accordingly, a number of experiments that investigate equilibrium selection have considered variations of the matching procedure in the hawk-dove games.

Early experiments in discrete time confirm many of the predictions of evolutionary game theory. Friedman (1996) investigates several one-population and two-

	Hawk	Dove
Hawk	-4,-4	0, 6
Dove	0, 6	3, 3

Figure 5: Payoff matrix of a hawk-dove game.

population hawk-dove games over 10–16 repeated periods. In single-population hawk-dove sessions, subjects exhibited approximate convergence to the mixed strategy Nash equilibrium. In two-population hawk-dove sessions, subjects exhibited approximate convergence to the pure strategy Nash equilibria. In a foreshadowing to the continuous-time results, somewhat tighter convergence to equilibrium was observed in experimental sessions that implemented mean-matching procedures relative to those that used conventional random matching procedures.

Oprea, Henwood, and Friedman (2011) provide the continuous-time analogue of Friedman (1996). Their treatments were also divided between a one population and a two population game. Specifically, one treatment employed single-population mean matching, where all twelve subjects are in a single mean-matching group, so a subject's earnings rate is proportional to her expected payoff being matched against a randomly selected other subject. The other treatment employed two-population mean matching, where subjects were divided into two equally sized groups and the earnings rate of a subject in one group is proportional to their expected payoff from being randomly matched against a subject from the other group. Under the one-population matching protocol, subjects converged toward the symmetric mixed-strategy Nash equilibrium; behavior converges toward the asymmetric pure strategy Nash equilibrium under the two-population matching protocol.

Thus both discrete-time and continuous-time experimental investigations of the hawk-dove game support the general predictions of evolutionary game theory. The main difference is that Oprea, Henwood, and Friedman (2011) observed stronger convergence towards these evolutionarily stable equilibria, suggesting that the combination of continuous repetition and mean-matching can accelerate convergence to evolutionarily stable equilibria more effectively than either individually.

Using the continuous-time framework, Benndorf, Martinez-Martinez, and Nor-

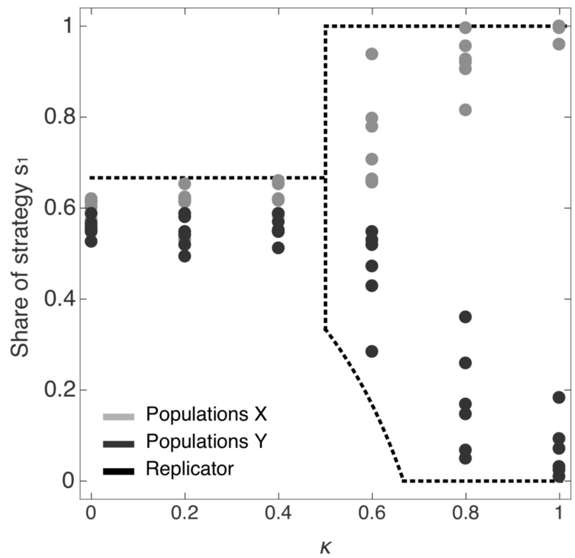

Figure 6: The aggregate results obtained by Benndorf, Martinez-Martinez, and Normann (2016). The vertical axis depicts the share of subjects choosing "hawk" in all sessions. The horizontal axis indicates coupling parameter κ that represents the strength of the interaction between the two populations of subjects. The dotted black line indicates the theoretical predictions from the replicator dynamics.

mann (2016) generalize the previous results, considering a range of matching procedures that lie between one-population and two-population matching in hawk-dove games. The coupling parameter κ denotes the strength of the coupling between the two populations. When $\kappa = 0$, subjects only interact with members of their own population. When $\kappa = 1$, subjects only interact with members of the opposing population. At intermediate values of κ, subjects have some degree of interaction with members of both populations. As illustrated by Figure 6, the authors find that behavior converged to the symmetric mixed strategy equilibrium when subjects primarily interact with members of their own population. Conversely, behavior converged to the asymmetric equilibrium when subjects primarily interacted with members of the opposing population.

2.2.3 Dominant strategy games: the prisoners' dilemma and public goods games

The prisoners' dilemma features a type of social dilemma: individual self-interest dictates that agents play their dominant strategies, but joint payoffs are maximized if all agents play their dominated strategies. Viewed as a conflict between what is best for the individual versus what is best for society, the game is extensively studied throughout the Social Sciences, resulting in hundreds (if not thousands) of experimental studies in discrete time. In general, subjects play their dominated strategies (i.e., cooperate) at levels far greater than the single-shot Nash equilibrium prediction of 0. While standard game theory can explain cooperation in games with indefinite repetition determined by a fixed probability, and has evolved to explain cooperation in finitely repeated games (e.g., Kreps et al., 1982), the evidence is inconclusive whether these factors are predictive of rates of cooperation in discrete-time experiments (c.f., Roth, 1995, Sally, 1995).

With their potential to involve a much larger number of interactions than their discrete-time counterparts, continuous-time experiments may give subjects more opportunity to learn how to utilize and respond to cooperative strategies, which are frequently more complex than the non-cooperative policy of simply playing the stage-game equilibrium strategy in every period. Friedman and Oprea (2012) conduct prisoners' dilemma games with fixed matching. In their discrete-time treatment, subjects played a single prisoners' dilemma with their partner in each period. Each period of their grid-time treatment was divided into n equally long sub-periods. A subject played a single discrete prisoners' dilemma game with their partner in each sub-period. In the continuous-time treatment, subjects interacted continuously with their partner during each period. The duration of each period was exactly 60 seconds in all three treatments. Only the frequency of interaction varied across treatments. They observe significantly higher cooperation rates in continuous-time than in one-shot games. Further, they find that the level of cooperation increased with the number n of interactions per period in the grid-time treatment. Bigoni, Casari, Skrzypacz, and Spagnolo (2015) also investigate the effect of the time-horizon on cooperative behavior in continuous-time prisoners' dilemma games. They implement a 2x2 experimental design varying both the length and the stochasticity of the interaction period. They find that subjects were more successful at achieving cooperation when they faced deterministic time horizons than when they faced stochastic time horizons.

Often considered a multi-player variant of the prisoners' dilemma, the public good game involves n players each of whom select how much of their private resources to invest in a public good. The marginal private return to investment in the public good is generally below the marginal cost of investment, but the marginal social return to the public good generally exceeds the marginal marginal cost of investment. Accordingly, the Nash equilibrium strategy profile of a public goods game generally involves much lower investment levels than the socially optimal strategy profile. As noted by Ledyard (1995), relatively large cooperation rates are often observed in the initial periods of discrete-time public goods experiments, but cooperation rates often decrease considerably in later periods.

Oprea, Charness, and Friedman (2014) investigate four-player public-goods games. They implement a 2x2 experimental design varying both the ability of subjects to communicate and the timing of subject interaction. Their continuous time treatment had subjects make contribution decisions in continuous time over a 10-minute interval. Their discrete-time treatment had subjects make contribution decisions at 10 discrete points of time in this interval. In their communication treatment, subjects can also freely communicate over chat. Their no communication treatments had subjects interact only through their contribution decisions. Individually, both chat and continuous time led a moderate increase in cooperation levels. Moreover, the combination of chat and continuous time produced a large increase in cooperation. When subjects had communication and interacted in continuous time, the median subject contributed completely to the public good with no sign of decay over time. The authors suggest that the need for communication in addition to continuous-time interaction comes from the harder coordination problem in four-player public goods games over two-player prisoners' dilemma games.

2.2.4 Three-person bargaining games

Three-person bargaining games provide a simple but effective glimpse into how individuals bargain. In the game, three players decide how to divide a continuous good. As long as two of the three can agree to an division it is implemented. Thus, there is a conflict between acting strategically and benefitting only two players and acting more equitably and providing significant payoffs for all three parties. Interestingly, any non-wasteful division is a Nash equilibrium.

Battaglini and Palfrey (2012) conduct experiments in discrete time on this game, involving a sequence of repeated periods with a stochastic termination rule. At the beginning of each period, one of the three subjects was randomly selected to propose a distribution of payoffs over all three agents. Then all three subjects voted on whether to approve the proposal. If the proposal received majority approval, it was implemented. Otherwise, the status-quo distribution of payoffs was maintained. Voting behavior is remarkably myopic and self interested: in the vast majority of cases, subjects voted for the option that gave them the highest short-term payoff. Two-way splits that exclude one player and egalitarian three-way spits were the most frequently observed outcomes.

Tremewan and Vanberg (2016) conduct a continuous-time analogue of three-person bargaining games. Subjects were randomly matched into groups of three each period for twenty identical periods. Subjects bargained using an intuitive graphical interface to select a division of payoffs from the two-dimensional simplex of possible divisions (Figure 7). Subjects were free to adjust their selected division at will and could continuously observe the divisions currently selected by other subjects in their group. Each second, if at least two of the three subjects in a group were selecting the same division, then payments were distributed to group members according to this division. This process continued until payoffs had been disbursed thirty times or a total of five minutes had passed. Every group reached thirty seconds of agreement time well before five minutes.

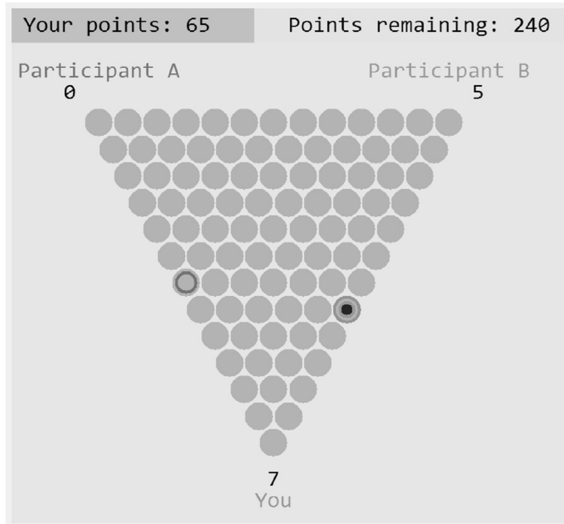

Figure 7: The experimental interface implemented by Tremewan and Vanberg (2016). Each circle represents a feasible allocation of payoffs. At any time, any subject could propose an allocation by clicking on it. This proposal was made visible to all subjects' by highlighting the chosen allocation with the proposer's color. Payments were only disbursed when at least two subjects were proposing the same allocation.

Over time, two-way splits became increasingly more common and three-way splits became increasingly less common, suggesting that subjects were learning to form minimal winning coalitions that exclude one player. By the latter half of periods, over two-thirds of agreements were two-way splits on the boundary of the simplex, meaning one subject received zero payoffs under the agreement. The majority of the remaining agreements were equal three-way splits in the exact center of the simplex. Accordingly, the final distributions of payoffs within a group at the end of a period was often rather unequal, with a least one of the three group members receiving less than one fourth of the available payoffs.

Both the discrete-time, Battaglini and Palfrey (2012), and continuous-time, Tremewan and Vanberg (2016), note two-way splits that exclude one player and egalitarian three-way spits were the most frequently observed outcomes. However, only Tremewan and Vanberg (2016) observe an increase in the frequency of two-way splits over time, suggesting that discrete-time experimental protocols may not have provided subjects with sufficient experience for this trend to emerge.

2.2.5 Trust games

Murphy, Rapoport, and Parco (2006) conduct experiments investigating n-player trust games in continuous time. They consider a real-time trust game with n symmetric players that proceeds in continuous time from 0 to 45 seconds. The value of the winner's prize increases exponentially over this time interval. At time t the value of the winner's prize is given by $5 \times \left(2^{t/5}\right)$. At any time in this period, any player can terminate the game and receive current value of the winner's prize. In this case, the other $n-1$ players receive a fixed fraction δ of the winner's prize. If 45 seconds elapse without any player terminating the game, then all players receive zero payoff. At any time less than 45 seconds, allowing the game to proceed is mutually beneficial to all players since it raises both the winner's payoff and the payoff to other players. However, each player can individually benefit by terminating the game before her opponents, thus securing the winner's payoff. Consequently, in equilibrium, the game unravels and the game will be terminated at time 0, giving each player a payoff of 0.

Each session consisted of 90 rounds. At the beginning of each round, subjects were randomly matched into groups of n players. During each round subjects played the real-time trust game with their group members, during which they were shown both the current clock time and the current value of the winner's prize. Any player in a group could terminate the game at any time. The experimental design implemented three treatment conditions. Treatment 1 had $n = 3$ players per group and losers receiving $\delta = 0.5$ of the winners prize. Treatment 2 had $n = 3$ players per group and losers receiving $\delta = 0.1$ of the winners prize. Treatment 3 had $n = 7$ players per group and losers receiving $\delta = 0.1$ of the winners prize. The authors conducted two sessions with each treatment. Each session was conducted with 21 experimental subjects. Under all three treatment conditions, average stopping times consistently decreased over the 90 rounds of play. Treatment 1 had higher initial stopping times in round 1 and slower rates of decline than the other two treatments.

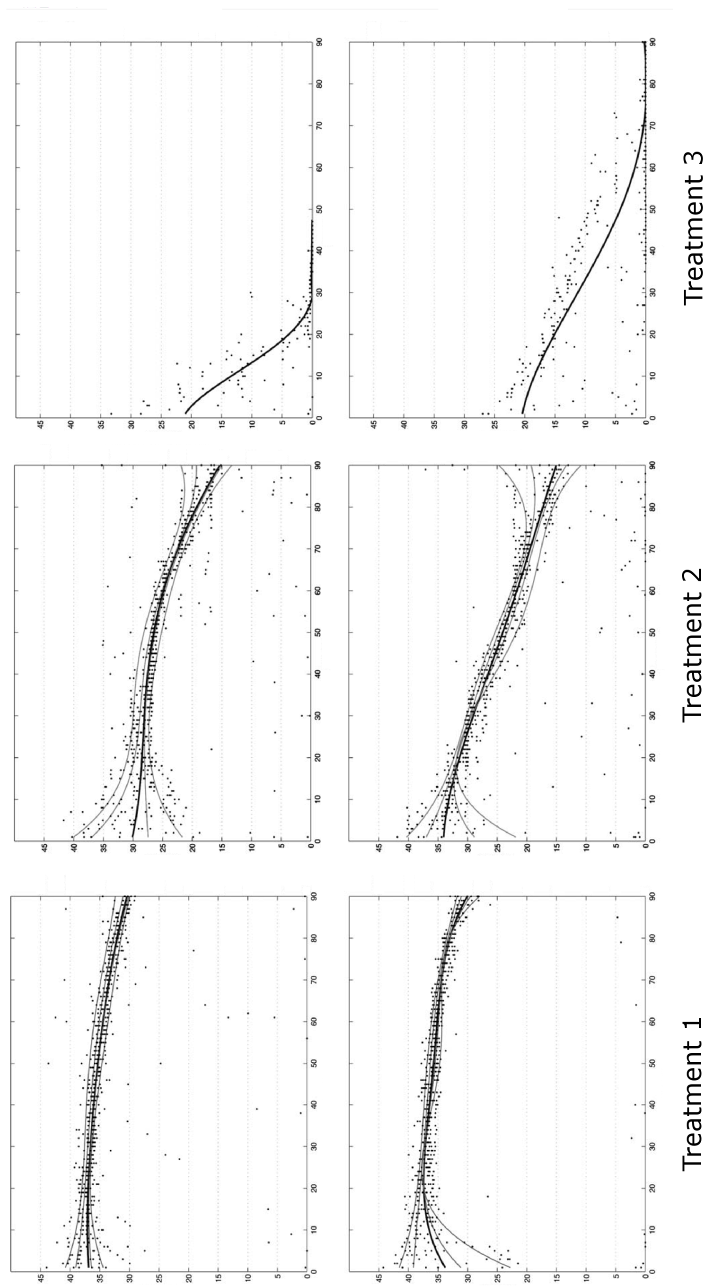

Figure 8: Observed stopping times by round in all three treatments of the real-time trust game implemented by Murphy, Rapoport, and Parco (2006). Each graph illustrates the observed behavior in one experimental session. Stopping times are indicated on the vertical axis of each graph. Rounds from 1 to 90 are indicated on the horizontal axis of each graph. Fourth-order polynomial trend lines are fitted to the data to highlight the dynamics.

Conversely, treatment 3 had lower initial stopping times and higher rates of decline than the other two treatments, suggesting that the breakdown in trust is faster when the group size is larger or the losers receive a smaller share of the winner's payoff. In treatment 3, stopping times converged to nearly zero by the end of all 90 rounds. In the other two treatments, subjects exhibited significantly positive stopping times, even over 90 rounds of repeated play. Figure 8 illustrates the observed stopping times by round in each of the three experimental treatments.

McKelvey and Palfrey (1992) conduct discrete-time experimental investigations of the centipede game, which can be seen as a discrete-time analogue of the continuous-time trust games investigated by Murphy, Rapoport, and Parco (2006). In each non-terminal period of the centipede game, the active player can either terminate the game and take a larger share of payoffs or pass to other player and increase the total size of the payoffs to be divided. In both the centipede game and the continuous-time trust game, total payoffs increase if both players refrain from terminating the game, but each player can individually benefit by terminating the game before her opponent. In both cases, backwards induction implies complete unraveling to the point of immediate termination. McKelvey and Palfrey (1992) find that behavior in the discrete-time centipede game became less noisy as subjects gain experience, but unraveling remained incomplete and the authors do not find support for equilibrium play.

In comparison, the relatively rapid convergence to equilibrium observed in treatment 3 of Murphy, Rapoport, and Parco (2006) can be attributed to the combination of continuous-time interaction and large group sizes. Similar to the role of mean-matching in the continuously repeated hawk-dove games conducted by Oprea, Charness, and Friedman (2014), large group sizes often make continuous interaction more effective at unraveling collusive behavior. Conversely, as observed in the two-player, continuous-time, prisoners' dilemmas conducted by Friedman and Oprea (2012), continuous-time interaction can often facilitate collusion when implemented with sufficiently small groups and conventional matching procedures. In both cases, continuous interaction can accelerate the emergence of long term behavioral patterns,

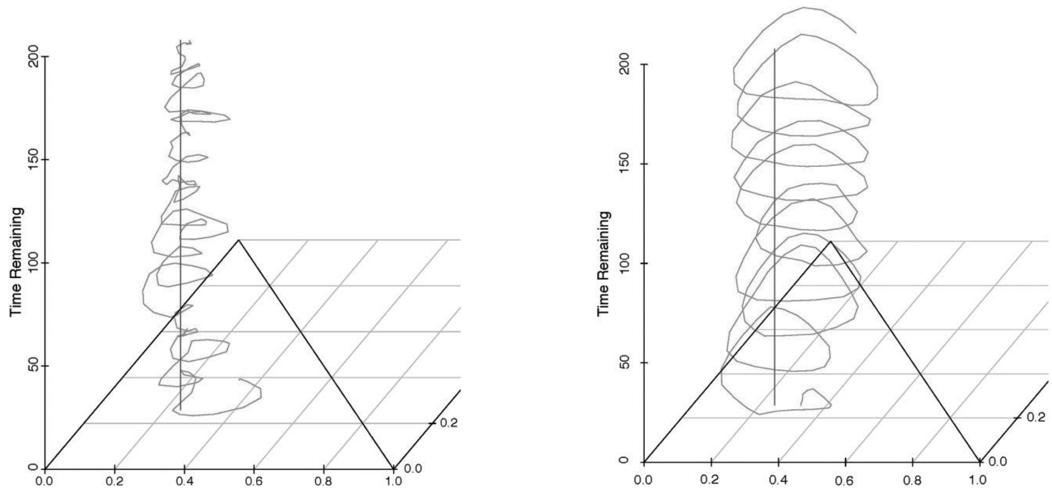

	R	P	S
R	36,36	24,96	66,24
R	96,24	36,36	30,96
S	34,66	96,30	36,36

	R	P	S
R	60,60	72, 0	30,72
R	0 ,72	60,60	66, 0
S	72,30	0 ,66	60,60

Figure 9: Illustration of behavior observed by Cason et al. (2014). The left figure illustrates a period in their stable treatment condition and the right figure illustrates a period in their unstable treatment condition. The vertical axis indicates time within the period. The horizontal axes indicate the percentage of subjects that employed each strategy. The corresponding payoff matrices are shown below each figure.

either collusive or competitive. With a sufficiently large number of repetitions, similar behavioral patterns can eventually emerge in discrete-time settings. But in many cases, the necessary number of repetitions may be prohibitively large for such laboratory experiments.

2.2.6 Mixed-strategy games

Continuous-time experimental methodology provides an excellent framework in which to observe and test the predictions of mixed-strategy equilibria. Under discrete-time protocols, the number of observations per subject is generally too small to determine much about the properties of their strategic play. In contrast, under continuous-time protocols, subjects may make thousands of adjustments in an experiment, more than enough to characterize adherence and deviations from the predictions of mixed-strategy play.

Cason, Friedman, and Hopkins (2014) conduct rock-paper-scissors experiments in continuous-time. They consider three variants of the classic three-strategy rock-paper-scissors game, one with a stable Nash equilibrium and two with unstable Nash equilibria (See Figure 9). The stable condition and one of the unstable conditions have the standard rock-paper-scissors pattern of strategic dominance. The other unstable condition has the reversed rock-scissors-paper pattern of strategic dominance. Each of these conditions yield identical Nash equilibrium predictions, but the best response dynamic only converges to the Nash equilibrium under the stable condition. Under the two unstable conditions, the best response dynamic approaches a stable limit cycle orbiting the Nash equilibrium. Throughout each session, subjects were divided into mean-matching groups of eight, so the payoff to an individual subject equaled the average payoff to her selected strategy against the strategies selected by her other seven group members.

The experimental design included four adjustment conditions: discrete pure, discrete mixed, continuous instant, and continuous slow. In all four adjustment conditions, subjects received graphical payoff feedback using a heatmap of payoffs on the simplex of mixed strategies. Under both of the discrete conditions, each period was divided into 20 subperiods during which the heatmap displayed the payoffs from the previous subperiod. In the discrete pure condition, subjects could could select from the three pure strategies each subperiod. In the discrete mixed condition, subjects could also select any mixture between the three pure strategies each subperiod. Under the continuous conditions, subjects could adjust their strategies continuously, earned continuous flow payoffs, and received real-time heatmap feedback throughout each period. In the continuous instant condition, subjects could adjust their strategy almost instantly with the click of a mouse. Under the continuous-slow condition, actual strategies constantly moved towards the targeted strategies at a fixed rate. Under this continuous-slow condition it would take 10 seconds to move from one vertex of the simplex to another. Each experimental session consisted of five blocks of five periods each.

Each period lasted for exactly three minutes. Each block of periods implemented one of the three payoff matrices and one of the four adjustment conditions. To avoid confounding treatment effects with learning effects, the order of the treatments was balanced across the eleven experimental sessions. Under the continuous slow adjustment condition, the authors observe significant cyclical behavior. As predicted by the best response dynamic, they observe counter-clockwise cycling under payoff matrices with the standard rock-paper-scissors pattern of strategic dominance and they observe clockwise cycling under payoff matrices with the opposite rock-scissors-paper pattern of strategic dominance. Under the unstable payoff matrices, they observe significantly higher cycle amplitude and they find that the time-average of the Shapley polygon (TASP) (see Figure 9) predicts the average mixed strategy better than the Nash equilibrum.

Stephenson (2019) experimentally investigates continuous-time attacker-defender games played by a population of attackers (A) and a population of defenders (D). Each agent faces conflict with members of the opposing population and opportunities for coordination with members of their own population. Each attacker chooses to attack one of two targets and each defender chooses to defend one of two targets. The percentage of population p that selects target t is given by x_t^p. The payoff π_t^p to an agent in population p who selects target t is given by

$$\begin{aligned} \pi_1^A(x) &= Mx_2^D + Cx_1^A & \qquad \pi_2^A(x) &= Mx_1^D + Cx_2^A \\ \pi_1^D(x) &= Mx_1^A + Cx_1^D & \qquad \pi_2^D(x) &= Mx_2^A + Cx_2^D \end{aligned} \tag{1}$$

The experiment implements three treatment conditions: one with zero coordination incentives ($C = 0, M = 5$), one with weak coordination incentives ($C = 0.2, M = 4.8$), and one with strong coordination incentives ($C = 2.4, M = 2.6$). All three experimental treatments exhibit identical Nash equilibrium predictions. Dynamic stability criteria (Hopkins, 1999, Hopkins and Seymour, 2002) classify the Nash equilibrium as neutrally stable under zero coordination incentives but strictly unstable all nonzero coordination incentives. In contrast, explicitly dynamic adaptive mod-

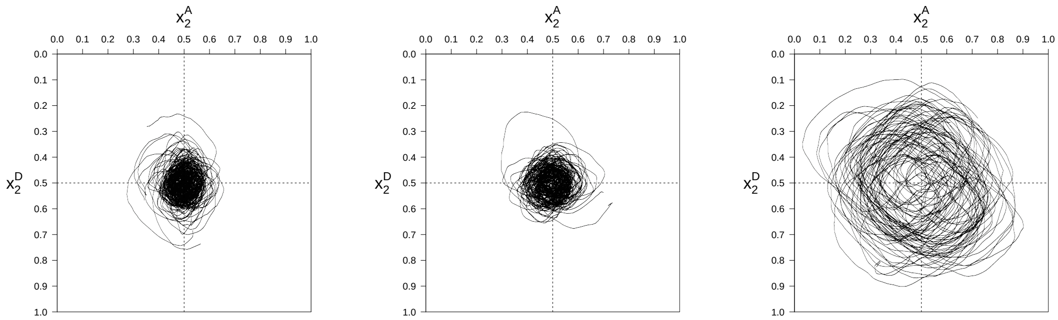

Figure 10: Cyclical behavior observed by Stephenson (2019). Each figure illustrates the path of the social state during every period of one treatment condition. The horizontal axis indicates the percentage of attackers attacking target two and the vertical axis indicates the percentage of defenders defending target two. The strong cordination treatment exhibited significantly larger deviation from equilibrium than the other two treatments.

els predict greater deviation from equilibrium under strong coordination incentives than weak but nonzero coordination incentives.

Each experimental session was conducted with twenty subjects. At the beginning of each session, subjects were randomly divided into two equally sized population groups of ten subjects, one attacker group and one defender group. Each session consisted of eight periods, each lasting for forty seconds, wherein subjects played the coordinated attacker-defender game. Subjects could continuously adjust their mixed strategy and earned continuous flow payoffs throughout each period. Subject behavior was tightly clustered around equilibrium in both the zero coordination treatment and the weak coordination treatment. However, subject behavior was widely dispersed from equilibrium in the strong coordination treatment, suggesting that explicitly dynamic models can capture important aspects of behavior that remain overlooked by dynamic stability criteria. Figure 10 illustrates the observed behavior in every period of this experiment. Subject behavior also violated the widely maintained assumption of sign-preservation, as subjects frequently switched from higher earning strategies to lower earning strategies, suggesting that non-signpreserving evolutionary models may provide a more accurate characterization of human behavior.

Some experimental games like the all-pay auction feature equilibrium in mixed strategies over a continuous strategy space. Despite the increased strategy space, the findings under continuous time feature many of the same properties. Stephenson and Brown (2020) conduct continuous-time experiments investigating long-run behavior in all-pay auction population games played by n agents. In these games, each agent selects a bid and is randomly matched into a group of m bidders. All bidders in this group pay their bid and all but the lowest bidder in this group receives a prize with value v. The authors implement two experimental treatments, one with groups of two bidders ($m = 2$) and one with groups of three bidders ($m = 3$). In both treatments, subjects were endowed with $w = \$10$ and competed for prizes with value $v = \$7$. Each session consisted of four five-minute periods. During each period, subjects could adjust their bids as frequently as desired with a click of the mouse and earned continuous flow mean-matching payoffs. A subject's instantaneous earnings rate was equal to her expected payoff from being matched into a random group of m subjects, consistent with the evolutionary models of "playing the field" (Smith, 1982). As predicted by adaptive but not equilibrium models, they observe significant cyclical bidding behavior. Consistent with dynamic stability criteria (Hopkins, 1999, Hopkins and Seymour, 2002), they find that the empirical distribution of bids exhibits significantly greater stability in the treatment with group size $m = 2$ than in the treatment with group size $m = 3$, suggesting that dynamic stability criteria can help predict when Nash equilibrium will provide a reliable characterization of long-run behavior.

2.3 Applications: industrial organization

Market competition between firms is rarely a static one-shot interaction. Firms often adjust their policies over time and can frequently observe aspects of the policies of their competitors. Further, these dynamic adjustments are rarely constrained to a predetermined sequential order and many firms are free to adjust their prices, hire workers, and adjust output at will. As a result, many industrial organization experiments have employed continuous-time methodology to explore the fundamentally asynchronous nature of market competition between firms.

2.3.1 Market entry

Calford and Oprea (2017) experimentally investigate market entry games where two agents independently decide when to enter a market. In this game, joint delay of market entry can be mutually beneficial to both agents, but each agent can individually benefit by entering the market before her opponent. If agents select from a finite grid of discrete entry times, then the unique subgame perfect Nash equilibrium has both agents entering the market immediately at time zero. However, Simon and Stinchcombe (1989) note that if agents have zero reaction lags and can enter at any time over a continuous range, then the unique equilibrium that survives iterated elimination of strictly dominated strategies has both agents cooperating to maximize their joint payoff by delaying entry until 40% of the time has elapsed. Since agents have zero reaction lags, neither agent can benefit by entering earlier, because her opponent would respond by instantly entering the market before any benefits can accrue to the first entrant. In contrast, if agents have non-zero reaction times, then either agent could benefit by entering slightly earlier, since some some benefits will accrue to the first entrant before her opponent can react. Hence the collusive equilibrium unravels and the unique subgame perfect equilibrium has immediate entry when agents have non-zero reaction times.

To test these theoretical predictions, Calford and Oprea (2017) conduct several experimental treatments. Each 60 second period of their discrete-time treatment is divided into 15 discrete points at which subjects simultaneously decide whether or not to enter the market. The authors induce zero reaction times in their perfectly continuous time treatment by pausing the game time whenever one subject enters the market and giving her opponent a chance to respond before resuming. Their inertial continuous time treatments allowed subjects to enter the market at at time

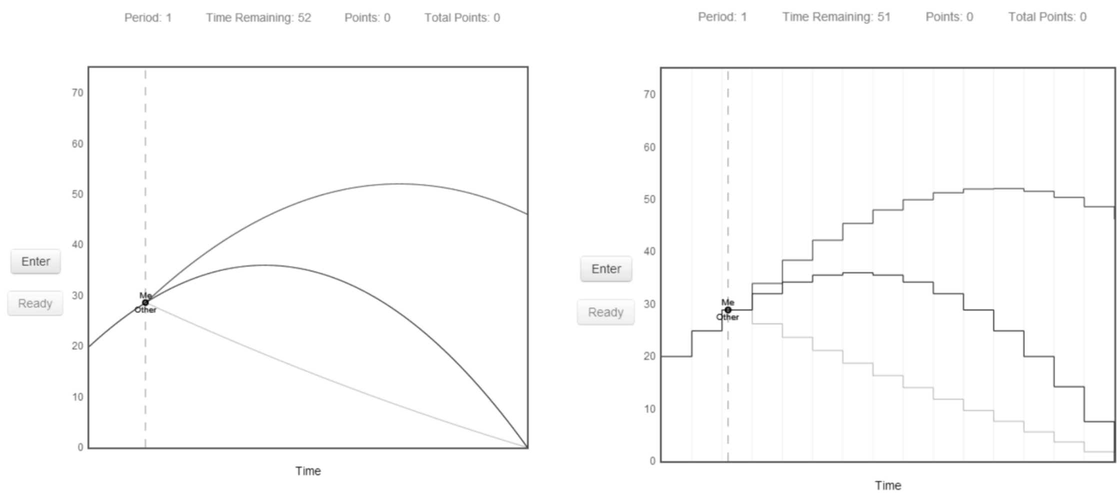

Figure 11: Experimental interfaces employed by Calford and Oprea (2017). The left figure provides a screenshot from the perfectly continuous and initial continuous time treatments. The right figure provides a screenshot from the perfectly discrete time treatment. During each period, the payoff dots labeled "Me" and "Other" move along the joint payoff line (center line) from the left to right. When a subject is the first to enter the market, her payoff dot shifts from the center to the upper line, while her opponent's payoff dot shifts to the lower line. When the second player enters the market, the payoff for both players are determined by the vertical location of each player's dot at the moment of the second entry.

during each period, but are subject to natural human reaction times. They also vary the relative length of human reaction times by adjusting the speed of the game. Figure 11 illustrates the experimental interfaces. After the first 3 periods, they find that most subjects enter immediately under the discrete time treatment and most subjects delay until about 40% of the time has elapsed under their perfectly continuous treatments, in accordance with equilibrium predictions. Under their inertial continuous time treatments they find that entry times vary with the game speed. At faster game speeds, human reaction times are relatively larger and they find that entry times tend to be relatively earlier. At slower game speeds, human reaction times are relatively smaller and they find that entry times tend to be relatively later. This result is inconsistent with the standard predictions of subgame perfect Nash equilibrium. However, it is consistent with ε-equilibrium and is well characterized by the basin of attraction and models of risk dominance. These results suggest that theoretical models of human behavior in continuous time should generally incorporate non-zero reaction times and should not assume the ability to respond instantly with zero delay. None of these conclusions would be possible without an examination of the game under continuous-time protocols.

2.3.2 Oligopolistic competition

Friedman, Huck, Oprea, and Weidenholzer (2015) experimentally investigate Cournot competition in duopolies and triopolies. During each period of their experiment, subjects selected their firm's production quantity. At the end of each period, subjects could observe both their own payoff and the production quantities selected by their competitors. However, subjects were not informed about the exact form of their payoff functions, only that they were symmetric across subjects and remained the same in all periods. Each period lasted for exactly four seconds and each experimental session consisted of exactly 1200 periods. Since four seconds is well above the typical human reaction-time, this experimental design does not fall within the traditional definition of a continuous-time experiment. However, this experimental design provides a high frequency of interaction and a large number of periods relative to conventional discrete-time experiments.

Each session was divided into three 400 period blocks and conducted with 12 subjects. At the beginning of each block, subjects were matched into a new group of competing firms and remained matched against these subjects for the remainder of the block. Under the duopoly treatment, subjects competed in pairs of two firms and in under the triopoly treatment, subjects competed in groups of three firms. The authors observed remarkably competitive behavior during first fifty periods of each session. Over these initial periods, median production quantities increased and then settled between the Cournot-Nash equilibrium and the perfectly competitive Walrasian equilibrium output levels. However, in later periods, production levels persistently decreased below Nash equilibrium output levels. Under the duopoly treatment, median production quantities stabilized near fully collusive levels by the end of the session. Even under the triopoly treatment, median production quantities

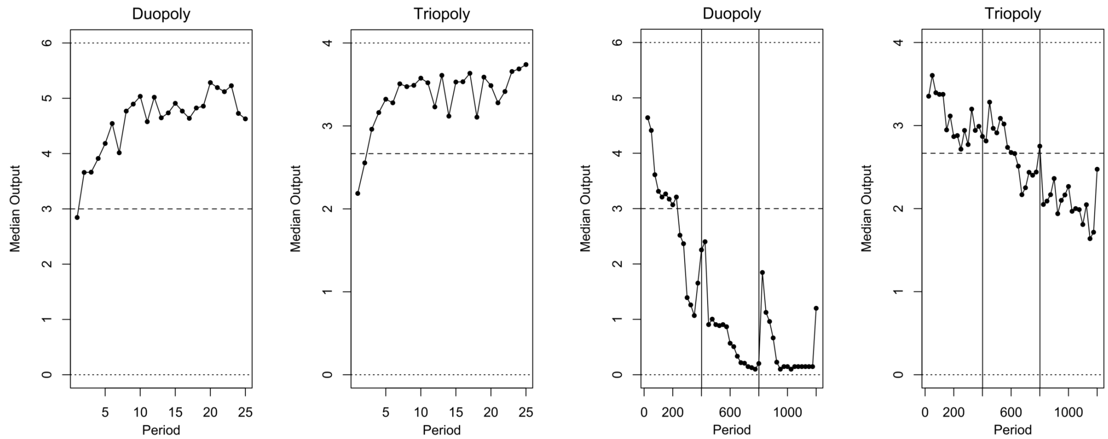

Figure 12: Median output quantities observed by Friedman et al. (2015) in each treatment. The two figures on the left depict the median output quantities in the first 25 periods. The two figures on the right depict the median output quantities in all periods, plotted in 25 period bins. The solid vertical lines indicate restart times at which subjects were rematched into new groups. The horizontal dashed lines indicate the Nash equilibrium output level.

decreased to well below the Cournot-Nash equilibrium levels by the end of the session. Figure 12 illustrates the median production quantities in each period. As noted by the authors, this study demonstrates the importance of providing subjects with sufficient amounts of experience and feedback in order to identify long-run behavioral patterns in a given strategic environment. In contrast to Apesteguia et al. (2010), these results suggest that long run behavior in oligopolistic settings with a small number of competitors is characterized by collusive rather than imitative behavior.

In contrast to several findings that suggest continuous-time protocols may increase collusion, Horstmann, Kraemer, and Schnurr (2016) find the opposite. They investigate oligopolistic competition in an experiment where subjects take the role of symmetric firms producing imperfect substitutes. Each experimental session consisted of sixty periods. The experimental design has three treatment variables, the number of subjects per group, the type of competition, and the timing of interaction. Subjects were randomly matched into groups of two and three, fixed for the entire experimental session, for the duopoly and triopoly conditions, respectively.

The Bertrand (Cournot) condition required subjects to chose prices (quantities) determined by a demand (an inverse demand) function. The discrete-time condition allowed subjects to take as long as desired to select their action each period, but did not receive any feedback on the actions selected by their competitors until the end of the period. In contrast, under the continuous time condition, subjects has exactly thirty seconds each period during which they could observe the observe the actions selected by their group members in real-time and were free to adjust their own actions. Subjects earned a lump-sum payment at the end of each period in discrete time, but earned continuous flow payoffs throughout each period under continuous time. Collusion is present to a greater extent in duopolies than in triopolies, and in Bertrand competition rather than Cournot. Controlling for these factors, the authors also find a significantly higher degree of collusion in the discrete-time treatment than under continuous time.

2.3.3 Hotelling models

In the Hotelling model of locational competition, consumers are uniformly distributed along the unit interval. Each firm chooses a single location on this interval. Firms are regulated to charge a uniform price, so consumers buy from whichever firm has the closest location, breaking ties at random. Hence a firm's market share is determined by the proportion of the unit interval that is closer to their location than the locations of its competitors. In the case of four competing firms, there is a unique pure strategy Nash equilibrium where two firms select neighboring locations at the first quartile of the unit interval and the other two firms select neighboring locations at the third quartile of the unit interval. In an experimental investigation of this model in discrete time, Huck et al. (2002) find limited support for these equilibrium predictions, observing widely dispersed locational choices with clusters around the first quartile, the midpoint, and the third quartile.

Kephart and Friedman (2015) experimentally investigate locational competition in four-player Hotelling games. Their experimental sessions consisted of twelve periods, lasting exactly three minutes. At the beginning of each period, subjects were

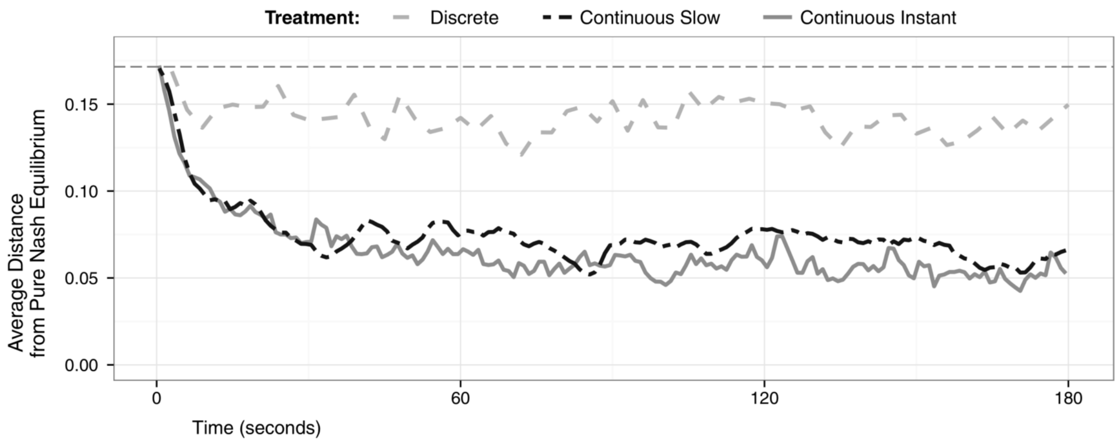

Figure 13: Average distance from the pure strategy Nash equilibrium over time within each period observed by Kephart and Friedman (2015).

randomly matched into groups of four. In one third of these cases, a subject would be matched against automated computer players that implement a predetermined sequence of actions. The experimental design varied the timing of subject interaction. The discrete treatment divided each period into sixty sub-periods, each lasting for exactly three seconds. During each sub-period, subjects were free to adjust their location, but only received feedback regarding locations chosen by other group members at the end of each period. The continuous-time, instant treatment provided subjects with continuous feedback about the locations selected by their group members throughout each period and allowed subjects to instantaneously adjust their location at any time. The continuous-time, slow treatment provided subjects with continuous feedback throughout the period and allowed them to gradually adjust their location at any time, subject to a speed limit at which they could traverse the entire interval in thirty seconds. Figure 13 illustrates the average distance from the pure strategy Nash equilibrium over time within each period. They find that subjects persistently failed to converge to equilibrium in the discrete-time treatment. In contrast, subjects exhibited reliable convergence towards equilibrium under both continuous time treatments.

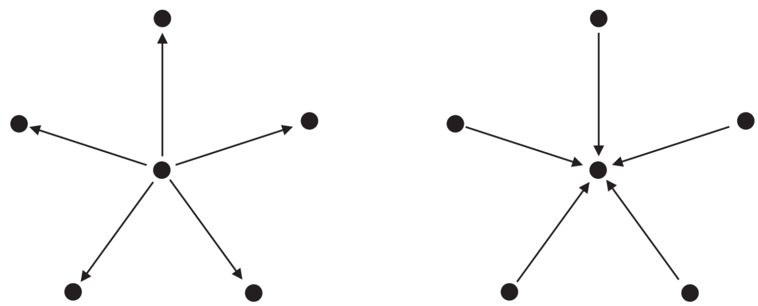

a) center-sponsored star b) periphery-sponsored star

Figure 14: Center-sponsored versus periphery-sponsored stars in the network formation game implemented by Berninghaus et al. (2006). Each arrow points from the players that actively establishes and pays for the link to the other player who is passively connected.

2.4 Applications: network formation

Many economic environments are characterized by networks of voluntary relationships including job search, decentralized markets, and international trade. The formation of these networks is frequently modeled by network formation games where each player makes linking decisions with many others. The outcome of such games is often described as a network of bilateral links between players and the payoff to each individual player is generally a function of the total network structure. Network formation games often have large and complex strategy spaces, so the dynamic process of behavioral adjustment often requires an unusually large amount of feedback to achieve convergence. Accordingly, many network formation experiments allow subjects to adjust their linking decisions freely throughout each period and provide real-time feedback in order to give subjects more opportunity to explore these complex strategy spaces.

Berninghaus, Ehrhart, and Ott (2006) conduct continuous-time network formation games allowing real-time link adjustment. During each experimental session, players could continuously establish or sever their links with others over a period of thirty minutes. Each player only pays for the links she actively establishes. Their treatments vary the cost of establishing a link (c = 2,7,13,16). A player's direct neighbors are the other players to whom they are directly connected. A player's indirect neighbors are the direct neighbors of the players to whom they actively establish a connection. The net revenue earned by a player is increasing her total number of neighbors, both direct and indirect. When c=2, periphery-sponsored stars (see Figure 14) are the only strict Nash networks. When c=7 or c=13, only periphery-sponsored stars and the empty network are strict Nash networks. When c=16, only the empty network is a strict Nash network. In all four treatments, only periphery-sponsored stars are efficient networks. The authors conducted fourteen sessions each with two groups of six subjects. When c=16, they find that subjects spend most of their time in the empty network. In contrast, when c=2,3,13, they find that subjects coordinate on a variety of periphery-sponsored stars. However, they also find that subjects frequently move from one periphery-sponsored star to another. They suggest that this movement between periphery-sponsored stars is driven by inequity aversion.

In later work, Berninghaus, Ehrhart, and Ott (2012) expand the investigation of continuous-time hawk-dove games (see Section 2.2.2) into endogenously constructed matching networks of six players. In these games, subjects play the hawk-dove game continuously with each of their their linked partners. Network links are costly to maintain. The cost of maintaining a link accrues to the subject who established the link. Over the course of a single period, 30-minute session, subjects were free to adjust their strategy, including both their their action in the hawk-dove game and the links they establish with other subjects in their group. Subjects could also observe their current earnings rate, their total accumulated earnings, the hawk-dove actions currently selected by all subjects in their group, and the links currently established by all subjects in their group. Three treatments varied link cost. In all three, it is profitable for a player selecting hawk to establish a link to a player selecting dove. Establishing a link from a dove to another dove is profitable in the middle and low cost treatments, but produces a loss in the high cost treatment. Establishing a link from a dove to a hawk is profitable in the low cost treatment, but produces a loss in the other two treatments. Establishing a link from a hawk to another hawk produces a loss in all three treatments.

In addition to Nash equilibrium, the authors also introduce an alternative stability criterion that they refer to as reaction-anticipating stability. A strategy profile is said to be "reaction-anticipatingly" stable if each player maximizes their expected payoffs under the assumption that other players will respond solely by removing unfavorable links and establishing favorable links. Subjects in this experiment are found to spend the majority of their time near reaction-anticipatingly stable configurations or Nash equilibrium configurations. Subjects spend about twice as much of their time near reaction-anticipatingly stable configuration as near Nash equilibrium configurations but tended to play reaction-anticipatingly stable configurations later in the period. On average, subjects first played Nash equilibrium configurations 803 seconds into each period and first played reaction-anticipatingly stable configurations 1097 seconds into each 1800 second period. Hence Nash equilibrium tended to characterize the initial phase of the convergence process while reaction-anticipatingly stable configurations more reliably characterized the long run behavior. Implementing this laboratory experiment in discrete-time would allow far fewer repetitions, so it might not be possible to observe long-run convergence to reaction-anticipatingly stable configurations without the use of continuous-time experimental methodologies.

Knigge and Buskens (2010) conduct continuous-time experiments where groups of four subjects simultaneously form networks and produce network goods with complementarities. Each experimental session consisted of 30 periods. The length of each period was randomly selected between 90 and 120 seconds, so subjects did not know exactly when a period would end. At the beginning of each period, each subject was randomly matched with three other subjects into a group of four subjects. Throughout each period, subjects were free to adjust both their proposed links and their investment level. Two subjects were considered "neighbors" if each subject proposed a link to the other, so network connections required mutual consent in this experiment.

Subjects could continuously observe both their instantaneous payoff and the links formed by other players. The high information treatment also allowed subjects to continuously observe the investment levels selected by other group members. Even in the low information treatment, subjects could always observe the effect on their own payoffs from linking to other group members. Payoffs were calculated from links and investment levels selected by group members at the end of each period. A subject's payoff depended on her investment level, her number of neighbors, and the investment of her neighbors. Each subject faced an increasing marginal cost of investment, a fixed marginal cost per neighbor, a marginal benefit from investment that increased in the investment level of neighbors, leading to a strategic complementarity of neighboring investment. The two-by-three experimental design has two treatment variables, the link cost and the level of information. They implement three different levels of the link cost and two different informational conditions for a total of six experimental treatments. The authors use a within-subject treatment protocol, so each experimental session implemented all six treatments. Treatment order was blocked to control for potential ordering effects.

The authors define a pairwise equilibrium as a Nash equilibrium profile of investment levels and link proposals under which no pair of subjects could mutually benefit by forming a link. The complete network is the unique pairwise equilibrium network under the low cost treatment. The empty network is the unique pairwise equilibrium network under the high cost treatment. Both are the pairwise equilibrium networks under the intermediate cost treatment. In the low cost treatment, 95% of groups arrived at the complete network. Conversely, in the high cost treatment, 91% arrived at the empty network. In the intermediate cost treatment, 55% of groups arrived at the complete network and 22% arrived at the empty network.

Across all treatments, subject arrived at a pairwise equilibrium network 88% of the 1257 paid periods. Subjects arrived at pairwise equilibrium investment levels in 79% of cases. In some networks is was possible for subjects earn above equilibrium payoffs by cooperating on above equilibrium investment levels. Cooperation was observed in only 13% of cases. However, a conditional logit regression finds that subjects were significantly more likely to cooperate on above-equilibrium investment levels when they could directly observe the investment levels of their group members.

The regression also finds that the likelihood of cooperation exhibited a small but significant upward trend over time.

Teteryatnikova and Tremewan (2016) experimentally investigate network formation games where payoff are derived directly from the structure of the network rather than from additional interactions between networked players. They investigate two three-player network formation games with symmetric payoff functions where each link requires bilateral consent from both parties. The authors consider two distinct classes of pairwise stability concepts: myopic stability concepts and farsighted stability concepts. Myopic stability concepts characterize players whose decision to add or remove links is entirely based on the immediate effects on payoffs. Farsighted stability concepts characterize players who consider the chains of reactions that they anticipate when they evaluate the effect of adding or removing a link. The authors also identify a particular subclass of cautious farsighted stability concepts under which players consistently anticipate the worst possible chain of reactions from adding or deleting a link. In their first game, myopic stability concepts only identify the complete network as stable, while all farsighted stability concepts identify both the complete network and all 1-link networks as stable. In their second game, myopic stability concepts only identify 1-link network as stable, and cautious farsighted stability concepts identify both the complete network and all 1-link networks as stable. Other farsighted stability concepts yield a variety of predictions from their second game.

Each experimental session consisted of 20 rounds, each lasting for exactly 30 seconds. At the beginning of each round, subjects were randomly matched into groups of three. During each round, subjects were able to propose, accept, and delete links with any of their group members at any time. Link formation required bilateral acceptance, but links could be unilaterally deleted by either party. Figure 15 depicts their experimental interface. Throughout each round, payoffs were earned every second based on the current configuration of links within the group. The authors find that the complete network was the most stable in their first games, but that 1-link networks were the most stable. This result is in line with the predictions of myopic

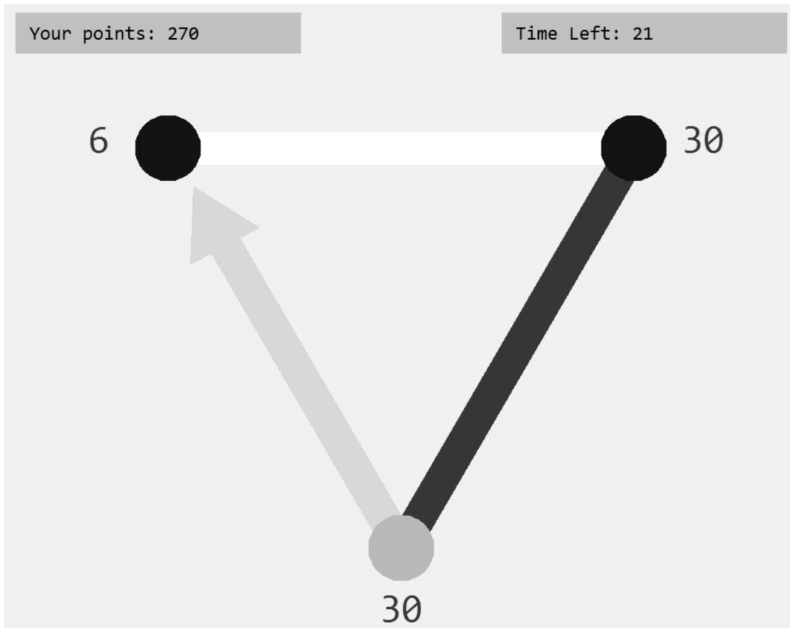

Figure 15: The experimental interface immplemented by Teteryatnikova and Tremewan (2016). The lower dot represents the subject viewing the interface and each upper dot represents another subject. The arrow indicates a proposed link between two subjects. The solid dark line indicates a link that has been confirmed by both subjects. Links could be proposed, cancelled, or confirmed at any time by clicking on the interface.

stability concepts, so the authors conclude that myopically stable networks exhibit the greatest stability in both of their experimental games. However, the authors also find that cautiously farsightedly stable networks exhibited greater stability than networks that were neither farsightedly nor myopically stable, providing evidence for the predictive power of cautiously farsighted stability concepts.

2.5 Applications: mechanism design

Stephenson (2017) conducts experiments investigating the effect of providing continuous feedback in three school choice mechanisms: the deferred acceptance mechanism, the top trading cycles mechanism, and the Boston mechanism. The continuous-feedback treatment condition provided experimental subjects with feedback on their school assignment throughout the preference reporting period. The discrete-feedback treatment condition only provided subjects with feedback on their school assignment at the end of each preference reporting period, as is typical in standard field implementations of school choice mechanisms. The provision of continuous assignment feedback does not effect the Nash equilibrium predictions, but the adaptive best response dynamic predicts that continuous assignment feedback can help subjects achieve equilibrium assignments by providing boundedly rational participants with increased opportunity for learning and adjustment.

Each experimental session consisted of twelve periods, each lasting for exactly 60 seconds. At the beginning of each reporting period, subjects were informed about the earnings that they could receive from being assigned each of three options: a, b, or c. During each period, subjects could freely adjust their preference reports. At the end of each reporting period, all preference reports were finalized and assignments were made. Subjects achieved equilibrium assignments far more often when they received continuous assignment feedback under all three student assignment mechanisms. Figure 16 illustrates the equilibrium assignment percentage observed under each mechanism. These results suggest that that policy makers may be able to improve the effectiveness of school choice mechanisms by providing participants with more feedback during the preference reporting period.

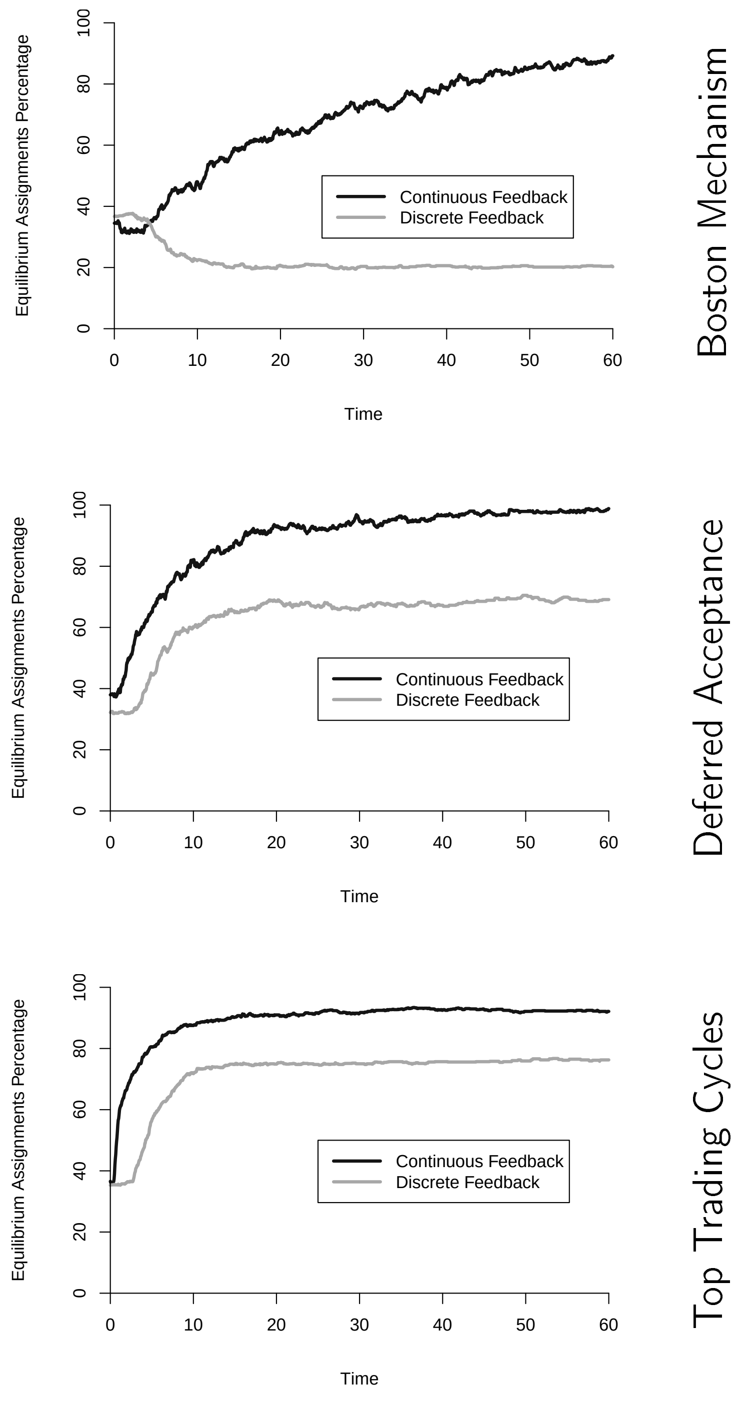

Figure 16: The percentage of suibjects recieving equilibrium assignments over time within each period under each of the student assignment mechanisms implemented by Stephenson (2017). The darker lines illustrate the equilibrium assignment percentage under the continuous feedback. The lighter lines illustrate the equilibrium assignment percentage under discrete feedback.

3 Discussion

Continuous-time protocols often provide new conclusions that significantly differ from those of discrete-time experiments. To broadly classify our survey we note the following.

- When group sizes are small and the gains from cooperation are large, continuous-time can increase cooperation. (Deck and Nikiforakis, 2012, Cheung and Friedman, 2009, Friedman and Oprea, 2012, Bigoni et al., 2015, Oprea et al., 2014)

- Experimental investigations of dynamic games in continuous time can help investigate the unraveling process and identify the role of reaction time. (Calford and Oprea, 2017, Murphy et al., 2006)

- Continuous repetition of a stage game can accelerate the emergence of long-run behavioral patterns by providing many repeated interactions over a relatively short of time. (Friedman et al., 2015, Kephart and Friedman, 2015, Stephenson, 2017)

- Experiments in continuous time can help test equilibrium selection concepts by examining the long run patterns of behavior that emerge as subjects as accumulate experience. (Deck and Nikiforakis, 2012, Oprea et al., 2011, Benndorf et al., 2016, Berninghaus et al., 2012, Tremewan and Vanberg, 2016)

- Continuous-time experiments can give subjects the opportunity to explore complex strategy spaces that are difficult to explore in a small number of discrete periods. (Berninghaus et al., 1999, Knigge and Buskens, 2010, Kephart and Friedman, 2015)

- Experiments in continuous time can help test theoretical models of the convergence process by recording strategy adjustments in real-time. The observed convergence process often conforms to models of evolutionary dynamics in some form. (Cason et al., 2014, Stephenson, 2019, Stephenson and Brown, 2020)

4 Summary

In this chapter, we have described and categorized an extensive collection of laboratory experiments with continuous time protocols. We note that, in general, continuous-time experimental methodology has the advantage of letting subjects freely adjust their actions over a continuous time interval and provides subjects with real-time feedback on the state of the game. By giving subjects more opportunities for learning and adjustment, continuous-time experiments can help subjects explore complex strategy spaces and arrive at stable patterns of behavior over relatively short periods of time. These unique characteristics of continuous-time experiments have provided researchers many valuable insights, and will likely continue to do so in the future as an increasing number of experimental researchers discover continuous-time experimental protocols.

References

Jose Apesteguia, Steffen Huck, Jörg Oechssler, and Simon Weidenholzer. Imitation and the evolution of walrasian behavior: Theoretically fragile but behaviorally robust. *Journal of Economic Theory*, 145(5):1603–1617, 2010.

Marco Battaglini and Thomas R Palfrey. The dynamics of distributive politics. *Economic Theory*, 49(3):739–777, 2012.

Raymond Battalio, Larry Samuelson, and John Van Huyck. Optimization incentives and coordination failure in laboratory stag hunt games. *Econometrica*, 69(3):749–764, 2001.

Volker Benndorf, Ismael Martinez-Martinez, and Hans-Theo Normann. Equilibrium selection with coupled populations in hawk–dove games: Theory and experiment in continuous time. *Journal of Economic Theory*, 165:472–486, September 2016.

Siegfried Berninghaus, Karl-Martin Ehrhart, and Marion Ott. A network experiment in continuous time: The influence of link costs. *Experimental Economics*, 9(3): 237–251, September 2006.

Siegfried K Berninghaus, Karl-Martin Ehrhart, and Claudia Keser. Continuous-time strategy selection in linear population games. *Experimental Economics*, 2 (1):41–57, 1999.

Siegfried K. Berninghaus, Karl-Martin Ehrhart, and Marion Ott. Forward-looking behavior in hawk-dove games in endogenous networks: Experimental evidence. *Games and Economic Behavior*, 75(1):35–52, 2012.

Maria Bigoni, Marco Casari, Andrzej Skrzypacz, and Giancarlo Spagnolo. Time Horizon and Cooperation in Continuous Time. *Econometrica*, 83:587–616, March 2015.

Evan Calford and Ryan Oprea. Continuity, inertia, and strategic uncertainty: A test of the theory of continuous time games. *Econometrica*, 85(3):915–935, 2017.

Timothy N. Cason, Daniel Friedman, and Ed Hopkins. Cycles and Instability in a Rock–Paper–Scissors Population Game: A Continuous Time Experiment. *Review of Economic Studies*, 81(1):112–136, 2014.

Yin-Wong Cheung and Daniel Friedman. Speculative attacks: A laboratory study in continuous time. *Journal of International Money and Finance*, 28(6):1064–1082, October 2009.

Douglas D Davis and Charles A Holt. *Experimental economics*. Princeton university press, 1993.

Cary Deck and Nikos Nikiforakis. Perfect and imperfect real-time monitoring in a minimum-effort game. *Experimental Economics*, 15(1):71–88, March 2012.

Daniel Friedman. Equilibrium in evolutionary games: Some experimental results. *The Economic Journal*, pages 1–25, January 1996.

Daniel Friedman and Ryan Oprea. A Continuous Dilemma. *American Economic Review*, 102(1):337–363, February 2012.

Daniel Friedman and Shyam Sunder. *Experimental methods: A primer for economists*. Cambridge University Press, 1994.

Daniel Friedman, Steffen Huck, Ryan Oprea, and Simon Weidenholzer. From imitation to collusion: Long-run learning in a low-information environment. *Journal of Economic Theory*, 155(C):185–205, 2015.

Ed Hopkins. A note on best response dynamics. *Games and Economic Behavior*, 29(1):138–150, 1999.

Ed Hopkins and Robert M Seymour. The stability of price dispersion under seller and consumer learning. *International Economic Review*, 43(4):1157–1190, 2002.

Niklas Horstmann, Jan Kraemer, and Daniel Schnurr. Oligopoly competition in continuous time. *Available at SSRN 2630664*, 2016.

Steffen Huck, Wieland Müller, and Nicolaas J Vriend. The east end, the west end, and king's cross: On clustering in the four-player hotelling game. *Economic Inquiry*, 40(2):231–240, 2002.

Curtis Kephart and Daniel Friedman. Hotelling revisits the lab: equilibration in continuous and discrete time. *Journal of the Economic Science Association*, 1(2): 132–145, 2015.

Antonie Knigge and Vincent Buskens. Coordination and cooperation problems in network good production. *Games*, 1(4):357–380, 2010.

David M Kreps, Paul Milgrom, John Roberts, and Robert Wilson. Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic theory*, 27(2):245–252, 1982.

John Ledyard. Public goods: A survey of experimental research. In J. Kagel and A. Roth, editors, *Handbook of Experimental Economics*. Princeton University Press, 1995.

Richard D McKelvey and Thomas R Palfrey. An experimental study of the centipede game. *Econometrica*, pages 803–836, July 1992.

Stephen Morris and Hyun Song Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, pages 587–597, June 1998.

Ryan O Murphy, Amnon Rapoport, and James E Parco. The breakdown of cooperation in iterative real-time trust dilemmas. *Experimental Economics*, 9(2): 147–166, 2006.

John Nash. *Equilibrium points in n-person games*. PhD thesis, Princeton University, 1950.

Ryan Oprea, Keith Henwood, and Daniel Friedman. Separating the Hawks from the Doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, 146(6):2206–2225, 2011.

Ryan Oprea, Gary Charness, and Daniel Friedman. Continuous time and communication in a public-goods experiment. *Journal of Economic Behavior & Organization*, 108(C):212–223, 2014.

Alvin E. Roth. Introduction to experimental economics. In J. Kagel and A. Roth, editors, *Handbook of Experimental Economics*. Princeton University Press, 1995.

David Sally. Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and society*, 7(1):58–92, 1995.

William H Sandholm. *Population Games and Evolutionary Dynamics*. MIT press, 2010.

Leo K Simon and Maxwell B Stinchcombe. Extensive form games in continuous time: Pure strategies. *Econometrica: Journal of the Econometric Society*, pages 1171–1214, September 1989.

J.M. Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982. ISBN 9780521288842.

Daniel Stephenson. Continuous feedback in school choice mechanisms. *Working Paper*, 2017. URL http://danielgstephenson.com/items/SchoolChoicePaper.pdf.

Daniel Stephenson. Coordination and evolutionary dynamics: When are evolutionary models reliable? *Games and Economic Behavior*, 113:381–395, 2019.

Daniel Stephenson and Alexander L. Brown. Playing the field in all pay auctions. *Working Paper*, 2020. URL http://danielgstephenson.com/items/AllPay.pdf.

Mariya Teteryatnikova and James Tremewan. Stability in network formation games with streams of payoffs: An experimental study. *Vienna Economics Papers 1508*, July 2016. URL https://homepage.univie.ac.at/mariya.teteryatnikova/Research/NetworkExperiment_160317.pdf.

James Tremewan and Christoph Vanberg. The dynamics of coalition formation: A multilateral bargaining experiment with free timing of moves. *Journal of Economic Behavior & Organization*, 130(C):33–46, 2016.

John Van Huyck, Joseph Cook, and Raymond Battalio. Adaptive behavior and coordination failure. *Journal of Economic Behavior and Organization*, 32:483–503, April 1997.

John B Van Huyck, Raymond C Battalio, and Richard O Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, 80(1):234–248, 1990.

John B Van Huyck, Raymond C Battalio, and Richard O Beil. Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *The Quarterly Journal of Economics*, 106(3):885–910, 1991.