

CONTINUITY, INERTIA, AND STRATEGIC UNCERTAINTY: A TEST OF THE THEORY OF CONTINUOUS TIME GAMES

BY EVAN CALFORD AND RYAN OPREA¹

The theory of continuous time games (Simon and Stinchcombe (1989), Bergin and MacLeod (1993)) shows that continuous time interactions can generate very different equilibrium behavior than conventional discrete time interactions. We introduce new laboratory methods that allow us to eliminate natural inertia in subjects' decisions in continuous time experiments, thereby satisfying critical premises of the theory and enabling a first-time direct test. Applying these new methods to a simple timing game, we find strikingly large gaps in behavior between discrete and continuous time as the theory suggests. Reintroducing natural inertia into these games causes continuous time behavior to collapse to discrete time-like levels in some settings as predicted by subgame perfect Nash equilibrium. However, contra this prediction, the strength of this effect is fundamentally shaped by the severity of inertia: behavior tends towards discrete time benchmarks as inertia grows large and perfectly continuous time benchmarks as it falls towards zero. We provide evidence that these results are due to changes in the nature of strategic uncertainty as inertia approaches the continuous limit.

KEYWORDS: Dynamic games, continuous time, laboratory experiments, game theory, strategic uncertainty, epsilon equilibrium.

1. INTRODUCTION

IN GAME THEORETIC MODELS, players usually make decisions in lock-step at a predetermined set of dates—a timing protocol we will call “Perfectly Discrete time.” Most human interaction, by contrast, unfolds asynchronously in unstructured continuous time, perhaps with some inertia delaying mutual responses. Does this difference between modeling conventions and the typical settings in which human interactions occur matter? Theoretical work on the effects of continuous time environments on behavior (developed especially in Simon and Stinchcombe (1989) and Bergin and MacLeod (1993)) focuses on what we will call “Perfectly Continuous time,” a limiting case in which players can respond instantly (i.e., with zero *inertia*) to one another, and arrives at a surprising answer: Perfectly Discrete time and Perfectly Continuous time can often support fundamentally different equilibria, resulting in wide potential gaps in behavior between the two settings.

In this paper, we introduce new techniques that allow us to evaluate these theorized gaps in the laboratory directly and assess their relevance for understanding behavior in naturally occurring environments. We pose two main questions. First, does the gulf between Perfectly Discrete and Perfectly Continuous time suggested by the theory describe real human behavior? Though equilibria exist that produce large differences in behavior (and authors such as Simon and Stinchcombe (1989) argued that these equilibria

¹We are especially grateful to Chad Kendall for invaluable advice and technical assistance and to James Peck and David Cooper whose insightful early comments informed the evolution of the experimental design. We thank James Bergin, Guillaume Frechette, Daniel Friedman, Yoram Halevy, P.J. Healy, Charles Sprenger, Emanuel Vespa, and Sevgi Yuksel for valuable comments and discussion. We are, finally, grateful to the Faculty of Arts at the University of British Columbia and the National Science Foundation under Grant SES-1357867 for supporting this research; to participants in the 2014 Stanford Institute for Theoretical Economics, Section 7, 2014 Caltech Political Economy Experiments Workshop, 2014 LABEL Conference at USC, the 2015 World Meetings of the Economic Science Association in Sydney; and to seminar audiences at Oxford University, Texas A&M, Washington University St. Louis, the University of Pittsburgh, Florida State University, the University of Missouri, and Chapman University.

should be considered highly focal), multiplicity of equilibrium in Perfectly Continuous time means that the effect of continuous time is, ultimately, theoretically indeterminate. Second, we ask how empirically relevant these gaps are: can more realistic, imperfectly continuous time games (games with natural response delays that we call “Inertial Continuous time” games) generate Perfectly Continuous-like outcomes? Subgame perfect Nash equilibrium (hereafter, SPE) suggests not but, as Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) emphasized, even slight deviations from SPE assumptions (à la ε -equilibrium) allow Perfectly Continuous-like behavior to survive as equilibria in the face of inertia, provided inertia is sufficiently small.

Recent experiments have begun to investigate the relationship between continuous and discrete time behavior in the lab (e.g., Friedman and Oprea (2012) and Bigoni, Casari, Skrzypacz, and Spagnolo (2015)) but have not yet directly tested the theory motivating these questions for a simple reason: natural human reaction lags in continuous time settings generate inertia that prevents a direct implementation of the premises of the theory.² These Inertial Continuous time settings—in which players can move at any time but with inertia preventing instantaneous reaction—are empirically important (and of independent interest) but are insufficient for a direct theory test because they generate very different equilibrium behavior from the Perfectly Continuous time environments that anchor the theory (a prediction we test and find strong though highly qualified support for in our data). In our experimental design, we introduce a new protocol (“freeze time”) that eliminates inertia by pausing the game for several seconds after subjects make decisions, allowing them to respond “instantly” to actions of others (i.e., with no lag in game time) and thus allowing us to test Perfectly Continuous predictions. By systematically comparing behavior in this Perfectly Continuous setting to both Perfectly Discrete time and Inertial Continuous time settings, we are able to pose and answer our motivating questions.

We apply this new methodology to a simple timing game similar to one discussed in Simon and Stinchcombe (1989) that is ideally suited for a careful test of the theory.³ In this game, each of two agents decides independently when to enter a market. Joint delay is mutually beneficial up to a point, but agents benefit from preempting their counterparts (and suffer by being preempted). In Perfectly Discrete time, agents will enter the market at the very first opportunity, sacrificing significant potential profits in subgame perfect equilibrium. By contrast, in Perfectly Continuous time, agents can, in equilibrium, delay entry until 40% of the game has elapsed, thereby maximizing (symmetric) joint profits. (Simon and Stinchcombe (1989) emphasized this equilibrium and pointed

²Both Friedman and Oprea (2012) and Bigoni et al. (2015) reported evidence from prisoner’s dilemmas played with flow payoffs in Inertial Continuous time (i.e., subjects in these experiments suffer natural reaction lags that prevent instant response to the actions of others). While the Friedman and Oprea (2012) design varies the continuity of the environment (discrete vs. continuous time interaction) in deterministic horizon games, the Bigoni et al. (2015) design centers on varying the stochasticity of the horizon (deterministic vs. stochastic horizon) in continuous time games. Other more distantly related continuous time papers include experimental work on multi-player centipede games (Murphy, Rapoport, and Parco (2006)), public-goods games (Oprea, Charness, and Friedman (2014)), network games (Berninghaus, Ehrhart, and Ott (2006)), minimum-effort games (Deck and Nikiforakis (2012)), hawk-dove games (Oprea, Henwood, and Friedman (2011)), bargaining games (Embrey, Frechette, and Lehrer (2015)), and the effects of public signals (Evdokimov and Rahman (2014)).

³Compared to, for instance, the continuously repeated prisoner’s dilemma, our game has the advantage of (i) featuring interior symmetric joint profit maximizing strategies that produce a much sharper test of the Simon and Stinchcombe (1989) prediction and especially (ii) a significantly simpler strategy space that both generates easier to interpret results and allows for a much simpler implementation of the “freeze time” protocol at the heart of our design.

out that it uniquely survives iterated elimination of weakly dominated strategies, but many other equilibria—including the inefficient immediate-entry equilibrium—exist in Perfectly Continuous time.) Importantly, Inertial Continuous time of the sort studied in previous experiments leads not to Perfectly Continuous time-like multiplicity in equilibrium but only to the inefficient instant entry predicted for Perfectly Discrete time—as Bergin and MacLeod (1993) pointed out, even a small amount of inertia theoretically erases all of the efficiency enhancing potential of continuous time in SPE.

In the first part of our experimental design, we pose our main question by comparing Perfectly Discrete and Perfectly Continuous time using a baseline set of parameters and 60 second runs of the game. In the Perfectly Discrete time protocol, we divide the 60 second game into 15 discrete grid points and allow subjects to simultaneously choose at each grid point whether to enter the market. In the Perfectly Continuous time protocol, we instead allow subjects to enter at any moment but, crucially, eliminate natural human inertia by freezing the game after any player enters, allowing her counterpart to enter “immediately” from a game-time perspective if she enters during the ample window of the freeze. We find evidence of large and extremely consistent differences in behavior across these two protocols. Virtually all subjects in the Perfectly Discrete time treatment suboptimally enter at the first possible moment while virtually all subjects in Perfectly Continuous time enter 40% of the way into the period, forming a tight mode around the joint profit maximizing entry time. The results thus support the conjecture of a large—indeed, from a payoff perspective, maximally large—gap between Perfectly Continuous and Discrete time behaviors.

In the second part of the design, we study how introducing realistic inertia into continuous time interaction changes the nature of the results observed in our Perfectly Continuous time treatment. Though SPE predicts that even a tiny amount of inertia will force behavior back to Perfectly Discrete-like immediate entry times, alternatives such as ε -equilibrium suggest that Perfectly Continuous-like results may survive as equilibria at low levels of inertia. In Inertial Continuous time treatments, we replicate our Perfectly Continuous time treatment but remove the freeze time protocol, thereby allowing natural human reaction lags to produce a natural source of inertia. We systematically vary the severity of this inertia by varying the speed of the game relative to subjects’ natural reaction lags and find that when inertia is highest, entry times collapse nearly to zero as predicted by SPE. However, when we lower inertia to sufficiently small levels, we observe large entry delays that are nearly as efficient as those observed in Perfectly Continuous time. Thus, realistic Inertial Continuous time behavior is well approximated by the extreme of Perfectly Discrete time when inertia is large and better approximated by the extreme of Perfectly Continuous time when inertia is small. While these patterns are inconsistent with SPE, they are, as both Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) stressed, broadly consistent with ε -equilibrium.

Though ε -equilibrium predicts the positive entry times we observe at smaller levels of inertia, it also generates multiplicity and therefore imprecise predictions. In the final part of the paper, we provide evidence that selection among these ε -equilibria is heavily influenced by strategic uncertainty. As inertia falls towards the continuous limit, the strategic risk of playing cooperative strategies falls with it, pushing the time at which immediate entry becomes a risk dominant strategy (relative to the most cooperative alternative) later and later into the game. We show that the time at which immediate entry first becomes risk dominant in this sense nearly perfectly organizes average entry times across treatments. We validate these results with a pair of diagnostic treatments and data from previous continuous prisoner’s dilemma experiments, where, in both cases, the time at which entry

becomes risk dominant also nearly perfectly organizes behavior. Our results join those from a number of recent papers showing that behavior in dynamic strategic settings is fundamentally shaped by strategic uncertainty (e.g., Dal Bó and Frechette (2011, 2016), Embrey, Frechette, and Yuksel (2016), Vespa and Wilson (2016)).

The results of our experiment suggest a role for Perfectly Continuous time theoretical benchmarks in predicting and interpreting real world behavior, even if the world is never perfectly continuous. Changes in technology have recently narrowed—and continue to narrow—the gap between many types of human interactions and the Perfectly Continuous setting described in the theory. Constant mobile access to markets and social networks, the proliferation of applications that speed up search, and the advent of automated agents deployed for trade and search have the effect of reducing inertia in human interactions. Our results suggest, contra subgame perfect Nash equilibrium, that such movements towards continuity may generate some of the dramatic effects on behavior predicted for (and observed in) Perfectly Continuous time even if inertia never falls quite to zero. Guided by these results, we conjecture that the share of interactions that are better understood through the theoretical lens of Perfectly Continuous time than that of Perfectly Discrete time will grow as social and economic activity continues to be transformed by this sort of technological change.

The remainder of the paper is organized as follows. Section 2 gives an overview of the main relevant theoretical results that form hypotheses for our experiment and Section 3 describes the experimental design. Section 4 presents our results, and Section 5 concludes the paper. Appendix A contains theoretical proofs and Supplemental Material Appendix B (Calford and Oprea (2017)) reproduces instructions to subjects.

2. THEORETICAL BACKGROUND AND HYPOTHESES

In Section 2.1, we introduce our timing game and in Section 2.2, we state and discuss a set of propositions characterizing subgame perfect Nash equilibrium (SPE) and providing us with our main hypotheses. In Section 2.3, we consider alternative hypotheses motivated by ε -equilibrium.

2.1. *A Diagnostic Timing Game*

Consider the following timing game, adapted from one described in Simon and Stinchcombe (1989). Two firms, $i \in \{a, b\}$, each choose a time $t_i \in [0, 1]$ at which to enter a market, perhaps conditioning this choice on the history of the game. Payoffs depend on the order of entry according to the following symmetric function:

$$U_a(t_a, t_b) = \begin{cases} \frac{1-t_b}{2} \left[\Pi_D + (t_b - t_a) \left(1 + \frac{2}{1-t_b} \right) \Pi_F \right] - c(1-t_a)^2 & \text{if } t_a < t_b, \\ \frac{1-t_a}{2} \Pi_D - c(1-t_a)^2 & \text{if } t_a = t_b, \\ \frac{1-t_a}{2} [\Pi_D - (t_a - t_b) \Pi_S] - c(1-t_a)^2 & \text{if } t_b < t_a, \end{cases} \quad (1)$$

with parameters assumed to satisfy $0 < 2c < \Pi_S \leq \Pi_D < 4c$ and $\frac{4c}{3} \leq \Pi_F \leq 4c$. Though the applied setting modeled by this sort of game matters little for our relatively abstract experiment, we can interpret the model as one in which firms face quadratic costs for time spent in the market (parameterized by c), earn a duopoly flow profit rate of Π_D while sharing the market, earn a greater flow profit Π_F while a monopolist, and suffer a

permanent reduced earnings rate (parameterized by Π_s) proportional to the time one's counterpart has spent as a monopolist.

Several characteristics of this game are particularly important for what follows. First, firms earn identical profits if they enter at the same time and this simultaneous entry payoff is strictly concave in entry time, reaching a maximum at a time $t^* = 1 - \frac{\Pi_D}{4c} \in (0, \frac{1}{2})$. Second, if one of the firms instead enters earlier than the other (at time t'), she earns a higher payoff and her counterpart a lower payoff than had they entered simultaneously at time t' . The firms thus maximize joint earnings by delaying entry until an interior time t^* , but at each moment each firm has a motivation to preempt its counterpart and to avoid being preempted.

2.2. Discrete, Inertial, and Perfectly Continuous Time Predictions

What entry times can be supported as equilibria in this game? The key observation motivating both the theory and our experiment is that the answer depends on how time operates in the game. In this subsection, we characterize equilibrium under three distinct protocols: Perfectly Discrete time, Perfectly Continuous time, and Inertial Continuous time.

We begin with Perfectly Discrete time, the simplest and most familiar case. Here, time is divided into $G + 1$ evenly spaced grid points (starting always at $t = 0$) on $[0, 1]$ and players make simultaneous decisions at each of these points. More precisely, each player chooses a time $t \in \{0, 1/G, \dots, (G - 1)/G, 1\}$ at which to enter, possibly conditioning this choice on the history of the game, H_t , at each grid point. Earnings are given by expression (1) applied to the dates on the grid at which entry occurred. As in familiar dynamic discrete time games like the centipede game and the finitely repeated prisoner's dilemma, there is a tension here between efficiency (which requires mutual delay until at least the grid point immediately prior to t^*) and individual sequential rationality (which encourages a player to preempt her counterpart). Applying the logic of backwards induction, strategies that delay entry past the first grid point unravel, leaving immediate entry at the first grid point, $t = 0$, as the unique subgame perfect equilibrium, regardless of G .

PROPOSITION 1: *In Perfectly Discrete time, both firms enter at time 0, regardless of the fineness of the grid, G , in the unique subgame perfect Nash equilibrium.*

In the opposite extreme protocol, Perfectly Continuous time, players are not confined to a grid of entry times but can instead enter at any moment $t_i \in [0, 1]$ (again, possibly conditioning on the history of the game at each t , H_t). Simon and Stinchcombe (1989) emphasized the relationship between the two extremes, modeling Perfectly Continuous Time as the limit of a Perfectly Discrete time game as G approaches infinity. In this limit, players can respond instantly to entry choices made by others: if an agent enters the market at time t , her counterpart can respond by also entering at t , moving in response to her counterpart but at identical dates. Since, in our game, delaying entry after a counterpart enters is strictly payoff decreasing, no player can expect to succeed in preempting her counterpart (or have reason to fear being preempted). This elimination of preemption motives also protects efficient delayed entry from unravelling and thus makes it possible to support any entry time $t \in [0, t^*]$ as an equilibrium.

PROPOSITION 2: *In Perfectly Continuous time, any entry time $t \in [0, t^*]$ can be supported as a subgame perfect Nash equilibrium outcome.*

Though it is possible for Perfectly Discrete and Perfectly Continuous behaviors to radically differ in equilibrium, this is hardly guaranteed. Because of multiplicity, Perfectly Continuous behavior may be quite different or quite similar to Perfectly Discrete behavior in equilibrium ($t = 0$ and t^* are both supportable in equilibrium in Perfectly Continuous time) depending on the principle of equilibrium selection at work—a central motivation for studying these environments in the laboratory. Simon and Stinchcombe (1989) emphasized that t^* is the unique entry time to survive iterated elimination of weakly dominated strategies in our game and argued that this refinement is natural in the context of Perfectly Continuous time games.

REMARK 1: In Perfectly Continuous time, joint entry at $t^* = 1 - \frac{\Pi_D}{4c}$ is the only outcome that survives iterated elimination of weakly dominated strategies.

Finally, Inertial Continuous time lies between the extremes of Perfectly Discrete and Perfectly Continuous time, featuring characteristics of each. Here, as in Perfectly Continuous time, players can make asynchronous decisions and are not confined to entering at a predetermined grid of times. However, as in Perfectly Discrete time, players are unable to respond instantly to entry decisions by their counterparts. In Inertial Continuous time, inability to instantly respond is due to what Bergin and MacLeod (1993) called *inertia* (here, simply response lags of exogenous size δ).⁴ With inertial reaction lags, the logic of unravelling returns as players once again have motives to preempt one another. As a result, the efficient delayed entry supported in equilibrium in Perfectly Continuous time evaporates with even an arbitrarily small amount of inertia, pushing continuous time behavior to that predicted for Perfectly Discrete time.

PROPOSITION 3: *In Inertial Continuous time, both firms enter at time 0, regardless of the size of inertia, $\delta > 0$, in the unique subgame perfect Nash equilibrium.*

Our game is a variation on a timing game studied in Simon and Stinchcombe (1989) and we refer the reader there for equilibrium derivations for Perfectly Continuous and Discrete time settings (Propositions 1 and 2). The Appendix provides a proof of Proposition 3.

2.3. *Alternative Hypothesis: Inertia and ε -Equilibrium*

Continuous time can fundamentally change SPE behavior but this effect is extremely fragile: even a slight amount of inertia will theoretically eliminate any pro-cooperative effects of continuous time interaction. Since inertia is realistic, this frailty in turn calls into

⁴We define inertia, δ , as the fixed fraction of the game that must elapse before an agent can respond to her counterpart's entry decision due to natural human reaction lags (i.e., if reaction lags last δ_0 seconds and the game lasts T seconds, $\delta \equiv \delta_0/T$). This notion of inertia is an instance of a more general notion described in Bergin and MacLeod (1993). There, a strategy satisfies inertia if, at each time t and history h , there is an action a and $\nu > 0$ (which may vary with t , h and a) such that a is played on the interval $[t, t + \nu)$ following every history h' that is the same as h on the interval $[0, t)$ —that is, a is constrained to be independent of the history h at time t over the interval $[t, t + \nu)$. Our application of inertia differs only in that it fixes the reaction lag, exogenously, to be δ . Formally, if a strategy enters at time t' in the history where no one has entered at any time prior to t' , then, we fix $\nu = \delta$ at all $t < t' - \delta$ and $\nu = t' - t$ at all $t' - \delta \leq t < t'$ across all h and a . Thus, an agent cannot respond to her opponent's entry decision at time t before her fixed reaction lag, δ , has elapsed (or, more precisely, an agent cannot condition an entry decision in $[t, t + \delta)$ on her opponent's entry decision at t).

question the usefulness of the theory for predicting and interpreting behavior. Perhaps for this reason, both Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) motivated the theory explicitly with reference to the more forgiving alternative of ε -equilibrium, emphasizing that any Perfectly Continuous time SPE is arbitrarily close to some ε -equilibrium of a continuous time game with inertia (or discrete time game), and vice versa. An ε -equilibrium allows agents to play strategies whose payoffs fall short of best response payoffs by up to some $\varepsilon > 0$.⁵

If agents are willing to tolerate small deviations from best response, they can support Perfectly Continuous-like outcomes as equilibria even in the face of inertia, provided inertia is sufficiently small. Indeed, in our game, when inertia is large, ε -equilibrium is similar to the Perfectly Discrete SPE, with entry occurring within less than a reaction lag of the beginning of the game. However, when inertia falls below a threshold level $\bar{\delta}$ (determined by ε), the equilibrium set expands to support any entry time $t \in [0, t^*]$, instead mirroring the Perfectly Continuous time SPE. Defining player i 's ($i \in \{a, b\}$) *target entry time*, $\hat{t}_i$, to be the time at which i enters if $-i$ has not entered at or before $\hat{t}_i - \delta$, and the *first entry time* in the game as $t_f \equiv \min\{\hat{t}_a, \hat{t}_b\}$, we formalize this in the following proposition.

PROPOSITION 4: *Suppose $\varepsilon < 3c$ (i.e., agents are not perfectly indifferent over all outcomes of the game), $0 < \delta < 1$ (i.e., reaction lags are less than the length of the game), and $\Pi_F = 4c$ (as in our experiment). Then there exists a cutoff value of inertia, $\bar{\delta} = 2 - \sqrt{4 - \frac{\varepsilon}{c}}$, such that (i) for $0 < \delta \leq \bar{\delta}$, joint entry at any time in $[0, t^*]$ can be supported in a subgame perfect ε -equilibrium and (ii) for $\bar{\delta} < \delta < 1$, the range of first entry times that can be supported in a subgame perfect ε -equilibrium is $[0, \bar{\delta}]$.*

See the Appendix for a proof. This result emphasizes that even very small deviations from the assumptions underlying the SPE—for instance, subjects holding incorrect beliefs regarding the size of reaction lags and the timing structure of the game or imprecision in payoffs specified in the game—can make either Perfectly Discrete-like or Perfectly Continuous-like benchmarks more predictive, depending on the severity of inertia. For this reason (and because of the important role ε -equilibrium plays in the theory), we built our experimental design in part with this alternative prediction in mind as an *ex ante* alternative hypothesis to SPE.

3. DESIGN AND IMPLEMENTATION

In Section 3.1, we discuss our strategy for implementing our three timing protocols in the lab and present the experimental software we built to carry out this strategy. In Section 3.2, we present our treatment design.

⁵More formally, suppose agent i has utility u_i and chooses a strategy $s_i \in S_i$. Then a strategy profile (s_i, s_{-i}) is an ε -equilibrium if, for all i , $u_i(s_i, s_{-i}) \geq \max_{s'_i \in S_i} u(s'_i, s_{-i}) - \varepsilon$, for some $\varepsilon > 0$, where ε is fixed across all i . To extend this to subgame perfection, write $S_{H,t}^{s_i}$ to denote the set of strategies for player i that agree with strategy s_i on the history H until time t (i.e., all strategies in $S_{H,t}^{s_i}$ prescribe the same action at all $t' < t$ along the history H , with s_i clearly being a member of $S_{H,t}^{s_i}$). Then a strategy profile (s_i, s_{-i}) is a subgame perfect ε -equilibrium if, for all i , H and t , $u_i(s_i, s_{-i}) \geq \max_{s'_i \in S_{H,t}^{s_i}} u(s'_i, s_{-i}) - \varepsilon$, where ε is fixed across all i , H , and t .




FIGURE 1.—Screen shot from Perfectly Continuous and Inertial Continuous time treatments (L-PC parameters).




FIGURE 2.—Screen shot from Perfectly Discrete time treatments (L-PD parameters).

3.1. *Timing Protocols and Experimental Software*

We ran our experiment using a custom piece of software programmed in Javascript. Figures 1 and 2 show screenshots. Using this software, we implemented the three timing protocols described in Section 2 as follows:

Inertial Continuous time. Figure 1 shows an Inertial Continuous time screenshot. As time elapses during the period, the payoff dots (labeled “Me” and “Other”) move along the joint payoff line (black center line) from the left to the right of the screen. (In most treatments, periods last 60 seconds, meaning it takes 60 seconds for the payoff dot to reach

the right-hand side of the screen.) When a subject is the first player to enter the market, her payoff dot shifts from the black to the green line (the top line), while her counterpart's payoff dot (the dot of the second mover) shifts to the red line (the bottom line).⁶ When the second player enters the market, period payoffs for both players are determined by the vertical location of each player's dot at the moment of second entry. Once both players have entered, they wait until the remaining time in the period has elapsed before the next period begins (see the instructions in Supplemental Material Appendix B for more detail). Because subjects, on average, take roughly 0.5 seconds to respond to actions by others, subjects have natural inertia in their decision making that should theoretically generate Inertial Continuous time equilibrium behavior.

Perfectly Continuous time. The Perfectly Continuous time implementation is identical to the Inertial Continuous time implementation (as shown in Figure 1) except that when either subject presses the space bar to enter, the game freezes (we call this the "freeze time" protocol) and the payoff dots stop moving from left to right across the screen for five seconds. Subjects observe a countdown on the screen and the first mover's counterpart is allowed to enter during this time. If the counterpart enters during this window, her response is treated as simultaneous to her counterpart's entry time and both players earn the amount given by the current vertical location of their payoff dot. Otherwise, the game continues as in Inertial Continuous time once the window has expired. Regardless, subjects must wait until the remaining time in the period has elapsed before the next period begins. The length of the pause was calibrated to be roughly 10 times longer than the median reaction lag measured in Inertial Continuous time, giving subjects ample time to respond, driving inertia to 0 and thus satisfying the premises of Perfectly Continuous time models.

Perfectly Discrete time. Figure 2 shows a screen shot for the Perfectly Discrete treatments, which is very similar to the continuous time screen but for a few changes. First, periods are divided into $G = 15$ subperiods, which begin at gridpoints $t = \{0, 4, 8, \dots, 56\}$ (measured in seconds), each marked by a vertical gray line on the subject's screen. Instead of moving smoothly through time, as in the continuous time treatments, the payoff dots follow step functions and "jump" to the next step on the payoff functions at the end of each subperiod. Actions are shrouded during a subperiod, so payoff dots will only move from the black to the green (or red) payoff lines after the subperiod in which a subject chose to enter has ended. Payoffs are determined according to equation (1), calculated at the grid point that began the subperiod in which the subject entered.⁷

3.2. Treatment Design and Implementation

Our experimental design has three parts. In the first, we implement 60 second timing games using the parameter vector $(c, \Pi_D, \Pi_F, \Pi_S) = (1, 2.4, 4, 2.16)$ under the Perfectly Discrete and Perfectly Continuous time protocols.⁸ We call these baseline treatments PD (Perfectly Discrete) and PC (Perfectly Continuous).

⁶Because of the nature of the payoff function, the green and red line change throughout the period prior to entry and stabilize once one player has entered.

⁷For example, if in our PD treatment a subject entered in the first subperiod and her counterpart entered in the third subperiod, payoffs would be given by $U(0, \frac{2}{15})$ for the subject and $U(\frac{2}{15}, 0)$ for her counterpart.

⁸To display the entire payoff space, we truncated the maximum period payment to 75 points (for context, $U(t^*, t^*)$ is 36 points), a payment that occurs only in the unusual event that a second mover delays entry significantly. This truncation has no effect on the equilibrium sets discussed in the paper.

Second, we examine the effects of inertia on continuous time decisions by running a series of Inertial Continuous time treatments using the same parameters. In the IC_{60} treatment, we run periods lasting 60 seconds each (just as in the PC and PD treatments). In the IC_{10} and IC_{280} treatments, we repeat the IC_{60} treatment but speed up or slow down the clock so that periods last 10 or 280 seconds (respectively). By speeding up the game clock so that the game lasts only 10 seconds (the IC_{10} treatment), we dramatically increase the severity of inertia; by slowing down the game so that it takes 280 seconds to finish (the IC_{280} treatment), we substantially reduce the severity of inertia (recall that inertia is the fraction of the game a subject is delayed by her reaction lag).

Finally, in the Low Temptation treatments (discussed in Section 4.4), we examine the robustness of explanations for our main results by changing the payoff functions in Perfectly Continuous and Discrete time. In the L-PD (Low temptation-Perfectly Discrete) and L-PC (Low temptation-Perfectly Continuous) treatments, we replicate the PD and PC treatments but lower the premium from preempting one's counterpart, Π_F , from 4 to 1.4. Changing Π_F has no effect on SPE in either case but can have substantial effects on strategic risk in Perfectly Discrete time.

All treatments are parameterized such that t^* occurs 40% of the way into the period (the seventh subperiod in Perfectly Discrete time treatments, and 24, 4, and 112 seconds into the period in the PC/L-PC/ IC_{60} , IC_{10} , and IC_{280} treatments, respectively). We ran the PD, IC_{60} , PC, L-PD, and L-PC treatments using a completely between-subjects design. In each case, we ran 4 sessions with between 8 and 12 subjects participating. Each session was divided into 30 periods, each a complete run of the 60 second game, and subjects were randomly and anonymously matched and rematched into new pairs at the beginning of each period. We ran the IC_{10} and IC_{280} treatments using a within-subject design consisting of three blocks, each composed of three IC_{280} periods followed by seven IC_{10} periods, for a total of 30 periods.⁹ Once again, subjects were randomly and anonymously rematched into new pairs each period. We conducted all sessions at the University of British Columbia in the Vancouver School of Economics' ELVSE lab between March and May 2014. We randomly invited undergraduate subjects to the lab via ORSEE (Greiner (2015)), assigned them to seats, read instructions (reproduced in Supplemental Material Appendix B) out loud, and gave a brief demonstration of the software. In total, 274 subjects participated, were paid based on their accumulated earnings and, on average, earned \$ 26.68 (including a \$5 show up payment), with funds provided by the Faculty of Arts at UBC. Sessions (including instructions, demonstrations, and payments) lasted between 60 and 90 minutes.

4. RESULTS

In Section 4.1, we report results from the Perfectly Continuous and Perfectly Discrete time treatments, while in Section 4.2, we study the relationship between the relatively realistic setting of Inertial Continuous time and the extremes of Perfectly Continuous and Discrete time. In Sections 4.3–4.5, we document how strategic uncertainty organizes our results. Unless otherwise noted, references to “entry times” will refer to the timing of first entry throughout. Second movers in the experiment almost universally behave in a manner suggestive of trigger strategies, responding to counterparts' first entry by entering as soon as possible after (this is unsurprising as delay is strictly dominated in our game).

⁹We used a within design, interspersing the fast paced IC_{10} treatment with the extremely slow paced IC_{280} because of concerns for subject boredom in IC_{280} .

After the first 10% of periods (after subjects have had a few periods to become comfortable with the interface), over 95% of second movers in both Perfectly Continuous and Perfectly Discrete time behave in this fashion (entering immediately in PC and no later than the very next subperiod in PD), while in Inertial treatments, we measure the median subject's reaction lag, δ_0 , at 0.5 seconds, closely matching reaction lags documented in previous experimental economics research (e.g., Friedman and Oprea (2012)).^{10,11}

4.1. Perfectly Continuous and Discrete Time

Figure 3(a) plots kernel density estimates of observed entry times for our PD (in gray, centered around 0) and PC treatments (in black). Figure 3(b) complements the kernel density estimates by plotting CDFs of subject-wise median entry times using product limit estimation intended to minimize the potential downward bias introduced by first movers preempting—and therefore censoring—the intended entry times of second movers.¹²




FIGURE 3.—(a) The left-hand panel shows kernel density estimates of entry times (normalized as percentage of the period elapsed) in the PD and PC treatments. For both treatments, t^* , which generates the maximal symmetric payoff, lies 40% of the way into the period. (b) The right-hand panel shows CDFs of subject-wise medians calculated using product limit estimates (Kaplan and Meier (1958)) of intended entry times.

¹⁰About 5% of subjects in the PC protocol entered with a delay of exactly 0.1 seconds, which we believe is due to a rounding error by the software and which we treat as a zero second lag in this calculation.

¹¹Neuroscientists find similar reaction times ranging from 450 ms to 600 ms in visual stimulus response tasks (compare to the 500 ms reaction time in our data), with brain imaging results suggesting that visual recognition of the stimuli occurs 150 to 200 ms before physical reaction (see Amano et al. (2006)).

¹²To form this estimate, we order each subject's observed *first* entry times t_i and for each calculate (i) n_i , the number of observations (periods of the experiment) in which the subject did not enter at a time prior to t_i and (ii) d_i , the number of observations in which the subjects entered at precisely t_i . The product limit estimate of the CDF is $F(t) = 1 - \prod_{t_i \leq t} \frac{n_i - d_i}{n_i}$, which Kaplan and Meier (1958) showed is the nonparametric, maximum likelihood estimate (here, of intended first entry times) in the face of right censoring (such as is caused by preemption in our game). We calculate the full CDF over observed entry times for each subject and use it to calculate a subject-wise median entry time for each subject. For the handful of subjects for whom the product limit median could not be calculated, we substitute the subject's raw median entry time although our estimates are robust to alternative treatment of those subjects (including simply dropping them from the sample).

The results are striking. In the PD treatment, virtually all subjects choose to enter immediately as the theory predicts, generating highly inefficient outcomes. The PC treatment, by contrast, induces radically different behavior: entry times are tightly clustered near t^* , with subjects maximizing joint earnings by delaying entry until about 40% of the period has elapsed.¹³ Recall that though t^* is only one of a continuum of equilibria in PC, it is the outcome uniquely selected by elimination of weakly dominated strategy and was advanced as a focal prediction by Simon and Stinchcombe (1989). The tightly clustered behavior in the PC treatments supports this conjectured focality and suggests that equilibrium selection is very uniform in Perfectly Continuous time. This pattern of behavior thus strongly supports the conjecture that Perfectly Discrete and Perfectly Continuous time induce fundamentally different behaviors in otherwise identical games.

RESULT 1: Perfectly Continuous interaction induces fundamentally different behavior from Perfectly Discrete interaction. While subjects virtually always enter immediately in the PD treatment, they virtually always delay entry until t^* in the PC treatment.

4.2. *Inertia and Continuous Time*

Perfectly Continuous time generates a dramatic change in behavior, but environments with zero inertia are probably rare. How robust are these extreme results to a re-introduction of inertia into the game? In order to study this question, we ran a series of Inertial Continuous time (IC) treatments, varying the severity of inertia from very high to very low. In the IC₆₀ treatment, we duplicated the PC treatment but eliminated the freeze time protocol, allowing subjects' reaction lags to generate natural inertia in the game. In the IC₁₀ and IC₂₈₀ treatments, run within-subject, we sped up (IC₁₀) or slowed down (IC₂₈₀) the game clock relative to the 60 second IC₆₀ periods, generating periods that lasted 10 or 280 seconds, respectively. Speeding up the game dramatically increases the magnitude of inertia (defined, recall, as the ratio of reaction lag length to game length) while slowing down the game reduces inertia substantially.

Figure 4 shows the results, plotting CDFs of subject-wise median product limit estimates of entry times for the IC₁₀ (high inertia), IC₆₀ (moderate inertia), IC₂₈₀ (low inertia), and the PC (zero inertia) treatments (for reference, we also plot the PD treatment in gray). The results reveal large, systematic effects of inertia on continuous time behavior as inertia drops towards zero. First, the tight optimal entry delays observed in the PC treatment almost completely collapse in the high inertia case, generating Perfectly Discrete-like near-immediate entry as predicted by SPE. However, when we reduce the severity of inertia, CDFs shift progressively to the right, with median entry times rising to $t = 0.184$ at medium inertia, $t = 0.323$ (where subjects earn 95% of earnings available at t^*) at low inertia, and finally $t = 0.393 \approx t^*$ when inertia reaches zero.¹⁴ The results thus show that entry times rise monotonically towards Perfectly Continuous levels as inertia falls towards zero, providing us with a next result:

RESULT 2: High levels of inertia cause entry delay to almost completely collapse as subgame perfect Nash equilibrium predicts. However as inertia falls towards zero, entry times progressively approach Perfectly Continuous levels.

¹³Mann–Whitney tests on session-wise median product limit estimates of entry times allow us to reject the hypothesis that PC and PD distributions are the same at the five percent level.

¹⁴An exact Jonckheere–Terpstra test allows us to reject the hypothesis that distributions of session-wise median product limit estimates of entry times are invariant to inertia against the alternative hypothesis that they are (weakly) monotonically ordered by inertia ($p < 0.001$).




FIGURE 4.—CDFs of subject-wise product limit estimates of entry times in each of the main treatments of the experiment.

The survival of high levels of cooperative delay in the face of small amounts of inertia is starkly inconsistent with SPE but broadly consistent with ε -equilibrium. However, ε -equilibrium supports delay by generating severe multiplicity: when δ is small enough to generate any significant delay in ε -equilibrium, it can support *any* $t \in [0, t^*]$. Thus while ε -equilibrium explains why delayed entry can occur in Inertial Continuous time, it cannot itself explain why progressively later ε -equilibrium entry times are selected as inertia falls towards zero. In the next subsection, we consider how inertia drives selection within the ε -equilibrium set.

4.3. Strategic Uncertainty and Inertia

Why do subjects select progressively later entry times within the ε -equilibrium set as inertia drops to zero? One appealing explanation is that subjects are sensitive to strategic risk, cooperating more readily (entering later) when it is less strategically risky to do so. Indeed, strategic risk has arisen as a centrally important variable in explaining variation in cooperation rates in recent experimental work on infinitely repeated prisoner's dilemmas (e.g., Dal Bó and Frechette (2011, 2016), Blonski, Ockenfels, and Spagnolo (2011)), finitely repeated prisoner's dilemmas (Embrey, Frechette, and Yuxsel (2016)), and dynamic games (Vespa and Wilson (2016)). To measure the strategic risk associated with attempting to cooperate, the prior literature has focused on a reduced form comparison between the *least cooperative* strategy (in our game, entering immediately) and the *most cooperative* trigger strategy (in our game, implementing a trigger strategy aimed at entry

at t^*), and examined the *basin of attraction* of choosing the least cooperative option.¹⁵ The basin of attraction of immediate entry (hereafter, BOA) is the probability that a player must assign to her counterpart playing the cooperative trigger so that she is indifferent between playing the cooperative trigger and entering immediately. Defining Δt as a subject's reaction-lagged entry time if her counterpart enters at t while she continues to play the cooperative trigger strategy, the $BOA(t)$ in our game is

$$\frac{U(t, t) - U(\Delta t, t)}{U(t, t) - U(\Delta t, t) + U(t^*, t^*) - U(t, \Delta t)}, \quad (2)$$

when (2) lies in $[0, 1]$ and is 1 otherwise.¹⁶

Intuitively, as the $BOA(t)$ grows large, it becomes increasingly risky for an agent to delay entry and continue with a cooperative trigger strategy since continued cooperation requires her to be increasingly certain her counterpart will continue cooperating as well; when BOA reaches 0.5, it becomes risk dominant for the player to abandon the cooperative trigger and immediately enter. Though the BOA is a reduced form measure (focused as it is on only two among many possible strategies), it has had a great deal of success in organizing prior data on cooperation in games, with higher BOA measures being associated with lower cooperation rates across a variety of settings (see meta-analyses in Dal Bó and Frechette (2016), Blonski, Ockenfels, and Spagnalo (2011), and Embrey, Frechette, and Yuksel (2016)). It also perfectly qualitatively organizes median entry times in our experiment: following the prior literature, and measuring at $t = 0$, the BOA—and with it the risk associated with cooperation—rises as inertia rises, from 0 (in PC) to 1 (in PD and IC₁₀), with IC₂₈₀ and IC₆₀ generating intermediate values. Median entry times thus fall monotonically as the strategic risk associated with cooperation (as measured by the BOA) rises.¹⁷

Corresponding classifications of risk dominance are especially successful at organizing the data: it is *never* risk dominant (relative to the cooperative trigger strategy) to enter prior to t^* in the PC treatment ($BOA(t)$ is 0 for all $t < t^*$) and subjects virtually always delay entry to t^* , while it is *always* risk dominant to enter immediately in the PD and IC₁₀ treatments ($BOA(t)$ is 1 for all t) and subjects almost universally enter immediately, near $t = 0$. In the IC₆₀ and IC₂₈₀ treatments, $BOA(t)$ rises throughout the game, with immediate entry becoming risk dominant at interior times: $t = 0.199$ in IC₆₀ and

¹⁵ An agent playing the cooperative trigger delays entry until t^* unless her counterpart enters at $t < t^* - \delta$ (or, in discrete time, $t < t^*$), in which case she enters δ (or, in discrete time, one gridpoint) after her counterpart's entry time.

¹⁶ The reaction-lagged entry time is $\Delta t = t + \min\{\delta, t^* - t\}$ in continuous time and $\Delta t = t + 1/G$ with $t < t^*$ in discrete time. $U(t, t)$ and $U(t^*, t^*)$ are the returns from mutual play of the cooperative trigger strategy and mutual immediate entry, respectively, while $U(t, \Delta t)$ is the earnings from immediate entry against a trigger playing counterpart and $U(\Delta t, t)$ is the earnings from playing the trigger against an immediately entering counterpart. The numerator of the BOA is an agent's potential loss from failing to enter immediately (due to preemption), while the denominator adds to this the potential loss from choosing to enter immediately (due to foregone joint cooperation payoffs). The BOA at time 0 is equivalent to the sizeBAD statistic Embrey, Frechette, and Yuksel (2016) used to measure strategic risk in the finitely repeated prisoner's dilemma. It is also the measure of strategic risk used by Harsanyi and Selten (1988) to motivate the notion of risk dominance (p. 83).

¹⁷ In the IC₆₀ and IC₂₈₀ treatments, the BOA rises over time, eventually reaching 1 in both cases. However, before either treatment's BOA reaches this boundary, there is a strict ordering of the BOA across inertia levels and this ordering is perfectly negatively associated with median entry times.




FIGURE 5.—Panels (a) and (b) show the time at which immediate entry first becomes risk dominant (in dashed gray lines) and median product limit estimates of entry times (black lines) from our data. Panel (c) plots corresponding measures for timing of defection in continuous prisoner's dilemmas from Friedman and Oprea (2012).

$t = 0.308$ in IC_{280} . Importantly, these dates at which entry first becomes risk dominant are almost precisely equal to the median entry times of 0.184 and 0.323 observed in IC_{60} and IC_{280} , respectively. Figure 5(a) summarizes by plotting the *earliest* time at which it is risk dominant to enter immediately for each treatment and overlays median product limit estimates of entry times (medians of the distributions plotted in Figure 4). The plot reveals that the earliest time of risk dominance almost perfectly organizes the timing of entry across our experimental design, predicting the monotonic convergence to Perfectly Competitive entry times observed in the data and strongly suggesting that strategic risk is driving ε -equilibrium selection in our game.¹⁸

RESULT 3: Entry times are qualitatively well organized by the basin of attraction of immediate entry, with entry times rising as the basin of attraction falls to zero. Corresponding notions of risk dominance quantitatively organize the data nearly perfectly, with the median subject entering almost exactly when immediate entry first becomes risk dominant.

¹⁸Immediate entry becomes trivially risk dominant relative to the cooperative trigger strategy in the PC treatment at $t = 0.4$ because the two strategies become identical.

4.4. Validation

We designed and ran two additional treatments to study whether our explanation for the comparative static effect of inertia can also explain other, distinct comparative statics. In the L-PD and L-PC treatments, we replicate the Perfectly Discrete and Perfectly Continuous treatments but dramatically lower the preemption temptation parameter Π_F from 4 to 1.4. In the PC treatment, lowering this parameter has no effect on strategic risk as measured by the basin of attraction. By contrast, in Perfectly Discrete time protocols, strategic uncertainty changes a great deal when we lower Π_F in the L-PD treatment: while, relative to the cooperative trigger strategy, immediate entry is always risk dominant in the PD treatment, it becomes risk dominant only at an interior time in the L-PD treatment, suggesting that lowering Π_F may generate a later entry time in discrete time via this strategic risk channel. Figure 5(b) plots the time at which entry first becomes risk dominant (relative to the cooperative trigger strategy) for the PD, L-PD, PC, and L-PC treatments and median subject-wise product limit estimates. The results show that the time of risk dominance tracks behavior quite well, with the change in Π_F leading to later entry times in Perfectly Discrete time (where the variable change alters strategic risk) but not in Perfectly Continuous time (where the variable change has no effect on strategic risk).¹⁹

RESULT 4: Changing the temptation to preempt changes entry behavior in Perfectly Discrete time (where it changes the time at which entry becomes risk dominant) but not in Perfectly Continuous time (where it has no effect on risk dominance).

How relevant are these sorts of results for understanding behavior in other continuous time games? To find out, we apply our risk dominance measure to data from the continuous prisoner's dilemma, the simplest game in a broad and empirically important class of games in which efficient outcomes are in tension with individual incentives. The Grid-n treatment in Friedman and Oprea (2012) studied 60 second prisoner's dilemmas that are divided up into 4, 8, 16, 32, and 60 Perfectly Discrete time subperiods, within-subject. This time protocol creates the equivalent of exogenous reaction lags in continuous time lasting 50%, 25%, 12.5%, 6.6%, 3.3%, and 1.6% of the game, respectively, generating a similar effect to inertia in our Inertial Continuous time games. In Figure 5(c), we plot median final mutual cooperation times (measured as a fraction of the period) as a function of the number of grid points.²⁰ Over this we overlay the earliest time at which "always defect" becomes risk dominant (relative to the grim trigger strategy). Again, these risk dominance benchmarks nearly perfectly match median final cooperation times, converging towards the Perfectly Continuous time limit of 1 (cooperation until the very end of the period) as the number of grid points grows large and the forced reaction lag grows small. The results thus provide strong out-of-sample confirmation that the time at which defection (or, in our game, immediate entry) becomes risk dominant organizes convergence paths to Perfectly Continuous time benchmarks.²¹

¹⁹A Mann-Whitney test allows us to reject the hypothesis that session-wise median product limit estimates of entry times in the PD and L-PD treatments are from the same distribution ($p = 0.02$); the same test does not allow us to reject the same hypothesis regarding the PC and L-PC treatments ($p = 0.183$).

²⁰As in the other analyses in this paper, we use the full data set in making these measurements. Restricting attention to the final 2/3 of the session as Friedman and Oprea (2012) did generates similar results.

²¹The earliest time of risk dominance also predicts the high rates of cooperation and low variation over parameters Friedman and Oprea (2012) observed in their (Inertial) Continuous time treatments. These treat-

RESULT 5: The earliest time of risk dominance (for the “always defect” strategy) predicts the progressively later timing of cooperation collapse observed in continuous time prisoner’s dilemmas as reaction lags fall to zero and the setting approaches Perfectly Continuous time.

4.5. Discussion: Strategic Uncertainty

We have shown that a simple measure of strategic uncertainty and corresponding benchmarks of risk dominance organize our data extremely well and do an equally impressive job of organizing behavior from other treatments and games. These results suggest that strategic uncertainty has a crucial role in determining behavior in continuous time games and that, behaviorally, one of the most important things that changes about a strategic environment as it grows more continuous is that cooperative behavior becomes less strategically risky. The model we use to measure strategic risk and benchmark behavior is built on fundamental simplifications, including that subjects select entry times by continually comparing the strategic risk of attempting to cooperate fully (until t^*) against the alternative of entering immediately, notably with little forward-looking consideration of other possible entry times. While it is possible that this, taken literally, is an accurate model of how subjects make dynamic strategic decisions (this sort of myopia and narrow bracketing, after all, has a long precedent in behavioral and experimental economics research), we emphasize that the measure can (and perhaps should) be interpreted instead as a convenient, reduced form shorthand for a richer model of strategic uncertainty. For instance, models of decision making under high degrees of uncertainty and in particular the minimax regret avoidance model (Savage (1951), Milnor (1954), and Stoye (2011)), when applied directly to the set of trigger strategies available in our game, generates benchmarks that also match our data (and data from the continuous prisoner’s dilemma) nearly perfectly. We conclude that careful research designed *ex ante* to distinguish between models of how subjects process and respond to strategic uncertainty is an important next frontier in the experimental study of dynamic strategic behavior.

5. CONCLUSION

Perfectly Continuous and Perfectly Discrete time are both idealizations, but they are illuminating ones, functioning as strategic analogues to vacuums in the physical sciences. Like vacuums, they are environments in which theoretical forces are cast in particularly sharp relief and results can be crisply interpreted in the light of theory. Although Perfectly Discrete time behavior has been exhaustively studied in thousands of experimental investigations, Perfectly Continuous time has never been studied before for a very simple reason: natural frictions in human interaction that loom especially large in the relatively fast paced setting of a laboratory experiment push strategic environments meaningfully

ments study 60 second continuous time prisoner’s dilemmas—prisoner’s dilemmas with flow payoffs realized in continuous time (though with subject-generated inertia). Friedman and Oprea’s 2012 design sets mutual cooperation payoffs of 10 and “suckers” payoffs of 0 and varies the temptation payoff (x) and defection payoff (y) cyclically over 32 periods over four parameterizations: Hard ($x = 18$, $y = 8$), Easy ($x = 14$, $y = 4$), Mix-a ($x = 18$, $y = 4$), Mix-b ($x = 14$, $y = 8$). The median final time of mutual cooperation (in 60 second periods) are 59.6, 58.4, 58.4, 57.5 in the Easy, Mix-a, Mix-b, and Hard treatments, which are very tightly clumped near the Perfectly Continuous benchmark time of 60. This nearly perfect cooperation and minimal variation over parameters is explained by the earliest time of risk dominance, which predicts collapse of cooperation at 58.8, 58.5, 56.5, and 55.5 seconds in these four treatments.

away from the Perfectly Continuous time setting described in the theory. Our paper introduces a methodological innovation that eliminates these frictions, allowing us to observe, for the first time, Perfectly Continuous behavior. By observing and comparing behavior across these two “pure” environments and by comparing both to more naturalistic protocols in-between, we learn some fundamental things about dynamic strategic behavior.

In our baseline treatments, we observe a large and economically significant gulf between Perfectly Discrete and Perfectly Continuous time behaviors, with entry times tightly clustered at the interior joint profit maximizing entry time in the latter case and at zero (immediate entry) in the former. Adding even a small amount of realistic inertia in mutual response to Perfectly Continuous time should, theoretically, erase this gulf, driving entry times to Perfectly Discrete-like levels. Instead, in a series of Inertial Continuous time treatments, we find that the effect of inertia depends fundamentally on its magnitude: entry times collapse to Perfectly Discrete-like early entry when inertia is large, but rise monotonically towards Perfectly Continuous-like efficient entry times as inertia falls towards zero. These patterns are inconsistent with SPE but broadly consistent with ε -equilibrium, where slight deviations from best responses can support equilibrium entry times much greater than zero even in the presence of inertia. We close the paper by showing that inertia also increases the strategic risk subjects face when attempting to cooperate and that aversion to this risk crisply explains the progressively later entry times we observe as inertia falls in our data. We validate these results first by using additional treatments that vary strategic risks by making changes to the payoff function, and then by using data from other continuous time games; in both cases, measures of strategic risk continue to sharply predict average behavior.

The results from our experiment suggest an appealing framework for understanding the relationship between the abstractions of Perfectly Discrete and Perfectly Continuous time and behavior in more realistic settings. Perfectly Discrete and Perfectly Continuous time predictions can be thought of as polar outcomes that each approximate realistic (Inertial Continuous time) behavior when inertia is either very high or very low, respectively. Indeed, we can easily push real time (Inertial) behavior close to either Perfectly Discrete or Perfectly Continuous time benchmarks simply by varying the severity of inertia and with it the strategic risk associated with Perfectly Continuous-like behavior. Concretely, these sorts of results suggest that Perfectly Continuous time benchmarks can, in some cases, be more empirically relevant than Discrete time benchmarks, even if agents face frictions that should be sufficient to eliminate Perfectly Continuous time equilibria under standard theory. The rise of thick online global markets, always-accessible mobile technology, friction-reducing applications, and automated online agents have made strategic interactions more asynchronous and lags in response less severe. These trends, which seem likely to intensify in coming years, have the effect of pushing many interactions closer to the setting of Perfectly Continuous time. Though these technological changes may never drive inertia entirely to the Perfectly Continuous limit of zero, our results suggest that behavior can nonetheless come close to Perfectly Continuous levels as inertia falls. This deviation from standard theory in turn suggests that we might expect Perfectly Continuous time predictions to become an increasingly relevant way of understanding economic behavior relative to the Perfectly Discrete predictions most often used in economic models.

APPENDIX A

As discussed in Section 2.2, Simon and Stinchcombe (1989) studied a game very close to ours and characterized results equivalent to Propositions 1 and 2. We prove Propositions 3 and 4 below.

PROOF OF PROPOSITION 3: SPE is a special case of subgame perfect ε -equilibrium, characterized in Proposition 4, in which $\varepsilon = 0$. When $\varepsilon = 0$, it is easy to see from expression (A.2) in the proof of Proposition 4 that $\bar{\delta} = 0$ so that equilibrium exists *iff* first entry is at $t = 0$. Because payoffs strictly decrease as an agent delays entry in response to her counterpart entering, the best response to an entry time of 0 is to enter at 0 as well so that joint entry at $t = 0$ is the unique SPE. Q.E.D.

PROOF OF PROPOSITION 4: For an arbitrary pair of strategies, to support target entry times of $(\hat{t}_i, \hat{t}_{-i})$ in an ε -equilibrium, the payoff from deviating to $t^{\text{BR}}(\hat{t}_{-i})$ given the strategies, must be weakly less than ε , $U(t^{\text{BR}}(\hat{t}_{-i}), \hat{t}_{-i}) - U(\hat{t}_i, \hat{t}_{-i}) \leq \varepsilon$ for $i \in \{a, b\}$. We first show that all possible first entry times, t_f , can be characterized by restricting attention to symmetric equilibria: if an asymmetric ε -equilibrium exists with $\hat{t}_a \neq \hat{t}_b$, then a symmetric equilibrium exists with joint entry at $\hat{t} = \min\{\hat{t}_a, \hat{t}_b\}$. Assume WLOG that $\hat{t}_b > \hat{t}_a$ and denote the entry time of b 's best response to a 's strategy as $t^{\text{BR}}(\hat{t}_a)$. If b 's original strategy can be supported in ε -equilibrium, she sacrifices $U(t^{\text{BR}}(\hat{t}_a), \hat{t}_a) - U(\hat{t}_b, \hat{t}_a) \leq \varepsilon$. Replacing b 's target entry time with $\hat{t}_a$ changes the sacrificed payoff to $U(t^{\text{BR}}(\hat{t}_a), \hat{t}_a) - U(\hat{t}_a, \hat{t}_a)$ which is also less than ε because $U(\hat{t}_b, \hat{t}_a) < U(\hat{t}_a, \hat{t}_a)$ for $\hat{t}_b > \hat{t}_a$.

We next show that all possible target times can be characterized using a subset of strategies called *trigger strategies*: if there is an ε -equilibrium in which at least one player is playing a non-trigger strategy with target entry time $\hat{t}$, we can always replace that player's strategy with a trigger strategy with the same target and still have an ε -equilibrium. Define the trigger strategy, denoted trigger- $\hat{t}$, as: for $t \leq \hat{t}$, enter at $\min\{\hat{t}, t' + \delta\}$ where t' is the time of one's opponent's entry (if any), and, for $t > \hat{t}$, enter immediately (if not already in the market). Note that i 's best response to an opponent playing trigger- $\hat{t}$ is to enter at $t^{\text{BR}}(\hat{t}) = \max\{0, \hat{t} - \delta\}$ —it is easy to verify that i 's payoff is strictly increasing in t below this entry time, and strictly decreasing above. WLOG, fix b 's (trigger- $\hat{t}$ or non-trigger- $\hat{t}$) strategy and replace a 's non-trigger- $\hat{t}$ strategy with a trigger- $\hat{t}$ strategy.

Off the equilibrium path, under a trigger strategy, a enters immediately if out of the market beyond $\hat{t}$ and as soon as possible at $\min\{\hat{t}, t' + \delta\}$ if b has entered at some $t' < \hat{t}$. These are, in both cases, best responses (a 's payoffs are strictly decreasing the longer she delays entering after b) and therefore a 's strategy continues to be within ε of best response payoffs off the equilibrium path (a 's change in strategy off the equilibrium path has no impact on b 's best response and therefore does not invalidate the ε -equilibrium). On the equilibrium path, b 's best response is to enter weakly *earlier* than $\hat{t} - \delta$ (the best response to trigger- $\hat{t}$) when a plays a non-trigger strategy by possibly delaying longer than δ (recall, by definition of inertia, a cannot enter earlier than δ). The payoff from this earlier entry in response to non-trigger- $\hat{t}$ is weakly greater than the payoff from $\hat{t} - \delta$ in response to trigger- $\hat{t}$. Thus, the deviation from best response necessary to support $\hat{t}$ *decreases* when a switches to a trigger strategy so that the strategy profile continues to be an ε -equilibrium. (Clearly, a 's change in strategy has no impact on her own best response.)

We now characterize the set of ε -equilibria in symmetric trigger strategies as a function of δ and, as shown above, with it the full set of first entry times, t_f , supportable as an ε -equilibrium. Under trigger strategies, subgame perfection is automatic because each agent best responds off the equilibrium path as shown above.

If a deviation to $\hat{t} - \delta$ is possible (i.e., $\hat{t} - \delta \geq 0$), it is the firm's best response as noted above. In this case, a pair of trigger- $\hat{t}$ strategies form a subgame perfect ε -equilibrium if and only if

$$U(\hat{t} - \delta, \hat{t}) - U(\hat{t}, \hat{t}) = \delta \left(\frac{3 - \hat{t}}{2} \right) \Pi_F - c(1 - \hat{t} + \delta)^2 + c(1 - \hat{t})^2 \leq \varepsilon. \quad (\text{A.1})$$

Using the fact that $\Pi_F = 4c$, (A.1) becomes $(\delta - 2 - \sqrt{4 - \frac{\varepsilon}{c}})(\delta - 2 + \sqrt{4 - \frac{\varepsilon}{c}}) \geq 0$ which, because $\delta \in (0, 1)$, is satisfied if and only if

$$\delta \leq 2 - \sqrt{4 - \frac{\varepsilon}{c}} \equiv \bar{\delta}. \quad (\text{A.2})$$

$\bar{\delta}$ is guaranteed to exist because $\varepsilon < 3c$ by assumption and, furthermore, is independent of $\hat{t}$. Therefore, if $\delta \leq \bar{\delta}$, we have an ε -equilibrium when $\hat{t} - \delta \geq 0$. If, instead, $\hat{t} - \delta < 0$, the firm's best response is to enter at time zero, which also provides a deviation payoff less than ε because $U(0, \hat{t}) - U(\hat{t}, \hat{t}) < U(0, \bar{\delta}) - U(\bar{\delta}, \bar{\delta}) = \varepsilon$. Thus, we have an ε -equilibrium for all $\hat{t} \in [0, t^*]$ when $\delta \leq \bar{\delta}$, establishing part (1) of the proposition.

For $\delta > \bar{\delta}$, an ε -equilibrium with joint entry time $\hat{t}$ can only exist if $\hat{t} - \bar{\delta} \leq 0$ (otherwise, the payoff deviation to $\max\{\hat{t} - \delta, 0\}$ is greater than ε). If $\hat{t} - \bar{\delta} \leq 0$, the firm's best response is to enter at zero, which again provides a deviation payoff less than ε , guaranteeing a symmetric ε -equilibrium exists for all $\hat{t} \leq \bar{\delta}$ when $\delta > \bar{\delta}$. Finally, because with $\delta > \bar{\delta}$ no symmetric ε -equilibrium exists with joint entry at $\hat{t} > \bar{\delta}$, no ε -equilibrium with first entry at this time exists either (which follows from the contrapositive of our statement regarding symmetry in the first paragraph of the proof), establishing part (2) of the proposition. Q.E.D.

REFERENCES

- AMANO, K., N. GODA, S. NISHIDA, Y. EJIMA, T. TAKEDA, AND Y. OHTANI (2006): "Estimation of the Timing of Human Visual Perception From Magnetoencephalography," *The Journal of Neuroscience* 26 (15), 3981–3991. [925]
- BERGIN, J., AND W. B. MACLEOD (1993): "Continuous Time Repeated Games," *International Economic Review*, 34 (1), 21–37. [915-917,920,921]
- BERNINGHAUS, S., K.-M. EHRHART, AND M. OTT (2006): "A Network Experiment in Continuous Time: The Influence of Link Costs," *Experimental Economics*, 9, 237–251. [916]
- BIGONI, M., M. CASARI, A. SKRZYPACZ, AND G. SPAGNOLO (2015): "Time Horizon and Cooperation in Continuous Time," *Econometrica*, 83 (2), 587–616. [916]
- BLONSKI, M., P. OCKENFELS, AND G. SPAGNOLO (2011): "Equilibrium Selection in the Repeated Prisoner's Dilemma: Axiomatic Approach and Experimental Evidence," *American Economic Journal: Microeconomics*, 3 (8), 164–192. [927,928]
- DAL BÓ, P., AND G. FRECHETTE (2011): "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence," *American Economic Review*, 101, 411–429. [918,927]
- (2016): "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey," Report. [918, 927,928]
- CALFORD, E., AND R. OPREA (2017): "Supplement to 'Continuity, Inertia, and Strategic Uncertainty: A Test of the Theory of Continuous Time Games'," *Econometrica Supplemental Material*, 85, <http://dx.doi.org/10.3982/ECTA14346>. [918]
- DECK, C., AND N. NIKIFORAKIS (2012): "Perfect and Imperfect Real-Time Monitoring in a Minimum-Effort Game," *Experimental Economics*, 15, 71–88. [916]
- EMBREY, M., G. R. FRECHETTE, AND S. F. LEHRER (2015): "Bargaining and Reputation: An Experiment on Bargaining in the Presence of Behavioural Types," *Review of Economic Studies*, 82, 608–631. [916]
- EMBREY, M., G. R. FRECHETTE, AND S. YUKSEL (2016): "Cooperation in the Finitely Repeated Prisoner's Dilemma," Report. [918,927,928]
- EVDOKIMOV, P., AND D. RAHMAN (2014): "Cooperative Institutions," Report. [916]
- FRIEDMAN, D., AND R. OPREA (2012): "A Continuous Dilemma," *American Economic Review*, 102 (1), 337–363. [916,925,929-931]
- GREINER, B. (2015): "Subject Pool Recruitment Procedures: Organizing Experiments With ORSEE," *Journal of the Economic Science Association*, 1, 114–125. [924]
- HARSANYI, J. C., AND R. SELTEN (1988): *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press. [928]

- KAPLAN, E. L., AND P. MEIER (1958): "Nonparametric Estimation From Incomplete Observations," *Journal of the American Statistical Association*, 53 (282), 457–481. [925]
- MILNOR, J. W. (1954): "Games Against Nature," in *Decision Processes*, ed. by R. M. Thrall, C. H. Coombs, and R. L. Davis. New York: Wiley. [931]
- MURPHY, R. O., A. RAPOPORT, AND J. E. PARCO (2006): "The Breakdown of Cooperation in Iterative Real-Time Trust Dilemmas," *Experimental Economics*, 9 (2), 147–166. [916]
- OPREA, R., G. CHARNES, AND D. FRIEDMAN (2014): "Continuous Time and Communication in a Public-Goods Experiment," *Journal of Economic Behaviour and Organization*, 108, 212–223. [916]
- OPREA, R., K. HENWOOD, AND D. FRIEDMAN (2011): "Separating the Hawks From the Doves: Evidence From Continuous Time Laboratory Games," *Journal of Economic Theory*, 146, 2206–2225. [916]
- SAVAGE, L. J. (1951): "The Theory of Statistical Decision," *Journal of the American Statistical Association*, 46 (253), 55–67. [931]
- SIMON, L. K., AND M. B. STINCHCOMBE (1989): "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 57 (5), 1171–1214. [915-921, 926, 932]
- STOYE, J. (2011): "Axioms for Minimax Regret Choice Correspondences," *Journal of Economic Theory*, 146, 2226–2251. [931]
- VESPA, E., AND A. J. WILSON (2016): "Experimenting With Equilibrium Selection in Dynamic Games," Report. [918, 927]

Dept. of Economics, Krannert School of Management, Purdue University, West Lafayette, IN 47907, U.S.A.; evancalford@gmail.com

and

Economics Department, University of California, Santa Barbara, Santa Barbara, CA 95064, U.S.A.; roprea@gmail.com.

Co-editor Dirk Bergemann handled this manuscript.

Manuscript received 26 April, 2016; final version accepted 30 January, 2017; available online 6 February, 2017.