Strategies in Stochastic Continuous-Time Games*

Yuichiro Kamada† Neel Rao‡

July 4, 2018

Abstract

We develop a new set of restrictions on strategy spaces for continuous-time games in stochastic environments. These conditions imply that there exists a unique path of play with probability one given a strategy profile, and also ensure that agents' behavior is measurable. Various economic examples are provided to illustrate the applicability of the new restrictions. We discuss how our methodology relates to an alternative approach in which certain payoffs are assigned to nonmeasurable behavior.

JEL Classification: C73, C78, D25

Keywords: dynamic games, continuous time, stochastic processes, timing

*We thank Eduardo Faingold, Drew Fudenberg, Oliver Hart, Michihiro Kandori, Akitada Kasahara, Phillip Strack, Takuo Sugaya, Satoru Takahashi, Jean Tirole, Juuso Toikka, Yasutora Watanabe, three anonymous referees, and the co-editor Dirk Bergemann for helpful suggestions. We also appreciate the comments of seminar participants at Chinese University of Hong Kong, Harvard University/Massachusetts Institute of Technology, Hitotsubashi University, University of California at Davis, University of Tokyo, and the Summer Workshop on Economic Theory 2018 at Otaru University of Commerce. Some of the results reported in this paper were included in an earlier version of another paper "Sequential Exchange with Stochastic Transaction Costs" by the same authors.

†University of California, Berkeley, Haas School of Business, Berkeley, CA 94720. E-mail: `y.cam.24@gmail.com`

‡University at Buffalo, SUNY, Department of Economics, Buffalo, NY 14260. E-mail: `neelrao@buffalo.edu`

1 Introduction

Continuous-time models have proven useful for the analysis of dynamic strategic interactions because closed-form solutions can be obtained while such results may be difficult to derive in discrete-time models. However, the specification of a game in continuous time entails some technical complexities. As discussed by Simon and Stinchcombe (1989) and Bergin and MacLeod (1993), many of these difficulties stem from the ability of players to instantaneously react to the actions of other agents. In order to resolve such problems, the strategy spaces need to be suitably defined, and those authors propose techniques for doing so in a deterministic environment. Although those methods are useful in that context, continuous-time modeling has also been successful in the analysis of stochastic environments, and hence it is desirable to develop suitable restrictions for strategy spaces in such models. This paper is devoted to the discussion of restrictions on strategy spaces for continuous-time games in stochastic settings.

One conceivable approach might be to directly employ the existing conditions defined for deterministic environments. For example, the "inertia" assumption proposed by Bergin and MacLeod (1993) essentially requires that at each history, there exists a time interval of (history-dependent) positive length during which agents do not change their actions. However, this restriction may be too strong in stochastic environments, where given a history up to time t and any small length of time $\epsilon > 0$, the state variable may change quickly within the interval $(t, t+\epsilon)$, so that the analyst may want to consider the possibility of an equilibrium in which agents change their actions during the interval $(t, t+\epsilon)$.[1] We present various examples to illustrate that such a situation naturally exists in stochastic settings.

Consequently, we develop a new set of conditions to restrict strategy spaces in continuous time. We begin by defining the concept of *consistency* between a strategy and a history, which means that a given strategy can generate a given history. One problem in continuous time is that there are cases where zero or multiple histories are consistent with a specific profile of strategies. In order to rule out these two possibilities, we develop two novel concepts. *Traceability* helps prevent the former situation, requiring that for any behavior of the other player, there exists a history

[1] Bergin and MacLeod (1993) also consider a weaker condition involving the completion of the set of inertia strategies, and we discuss it in footnote 22.

consistent with a player's strategy. *Frictionality* eliminates the latter possibility, requiring that in any history that is consistent with a player's strategy, the player can move only finitely many times during any finite time interval. By applying these criteria, we show that each strategy profile induces a unique path of play.

Unlike the approach in Simon and Stinchcombe (1989), which requires the number of moves to be uniformly bounded, our restrictions do not preclude strategies in which the number of moves in a given finite time interval may be arbitrarily large.[2] In fact, the equilibria that we analyze in some of our examples entail no upper bound on the number of moves during any finite time interval. Moreover, agents may make infinitely many moves over an infinite horizon with probability one.[3]

A further issue regards the measurability of the action process. In a deterministic environment like in Bergin and MacLeod (1993), restrictions like those described above would essentially suffice to ensure that the resulting path of play is measurable. However, more is involved in our stochastic setting because the actions of the agents can be contingent on moves of nature. For expected payoffs to be well defined, the stochastic process describing the moves of agents should be progressively measurable. In addition, specifying the strategy space so as to ensure the measurability of actions is complicated in our model due to perfect monitoring. The difficulty is that the strategy of one player affects whether the strategy of the other player induces a measurable action process.[4]

Therefore, we invent a two-stage technique for delineating the strategy space. A strategy is called quantitative if the behavior of an agent is measurable regardless of the strategy of its opponents. A strategy is called *calculable* if the behavior of an agent is measurable when its opponents play quantitative strategies. We show that each agent's behavior is measurable if every agent uses a calculable strategy. Moreover, the set of calculable strategies is the largest strategy space for each agent

[2] More precisely, fix an interval of the nonnegative real line with positive but finite length. For each agent, our model allows for the possibility that there exists a strategy satisfying the restrictions of the model such that the following holds. For any nonnegative integer n, one can find a history that is always consistent with this strategy such that the number of moves by the agent during the given time interval is equal to n.

[3] Simon and Stinchcombe (1989) also restrict how players may condition their behavior on the past by requiring strategies not to distinguish between two histories in which the same actions occur at different times that are sufficiently close to each other. Some of the equilibria we consider in our examples violate this assumption.

[4] This type of complexity does not arise in models of continuous-time games with imperfectly observable actions such as Sannikov (2007).

that includes the set of all quantitative strategies and implies measurable actions.

Although we repeatedly explain that the existing restrictions in the literature like that of Bergin and MacLeod (1993) may not be appropriate for our stochastic environment, we emphasize that our objective is not to undermine their usefulness. Our view is that those techniques and ours apply to different contexts. We discuss those restrictions simply to underscore the additional complications that arise when the environment is stochastic. The conditions in the literature are developed primarily for a deterministic environment, and it is natural that they may not extend to the stochastic case. Similarly, our model is not intended to cover all possible stochastic environments, and our focus is on games in which agents strategically choose the timing of discrete moves. This, in particular, precludes games in which agents' actions change continuously over time.[5] As the reader will hopefully see, even though our framework is not totally comprehensive, defining strategy spaces is a nontrivial task, and our examples show that our model encompasses a number economically relevant situations.

Continuous-time modeling has been widely employed in stochastic settings, and its use is growing. Section 5 of our paper considers various applications. An important question is the timing of investment under uncertainty, which has been studied by McDonald and Siegel (1986) in continuous time with a single agent. In section 5.1, we consider a model of forest management in which multiple agents decide when to harvest trees whose volumes evolve stochastically over time. A related topic is entry and exit by a firm, which is studied by Dixit (1989) in continuous time where the price follows a diffusion process. In section 5.2, we analyze a situation in which two firms contemplate the timing of entry into a market, where the cost of entry varies with time and the benefit depends on entry by the competitor.

Some existing models limit players to moving at random points in time. For instance, there is a literature on bargaining games in continuous time, including the model in Ambrus and Lu (2015), where players can make offers at random times and must reach an agreement by a specified deadline. Another class of models is

[5]Such models have been applied to study, for example, repeated games (Sannikov, 2007), contracting (Sannikov, 2008), reputation (Faingold and Sannikov, 2011), signaling (Dilme, 2017), and oligopoly (Bonatti, Cisternas, and Toikka, 2017). These papers assume the imperfect observability of opponents' actions and an inability to condition on one's own past actions, which simplify the description of strategies, thereby preventing issues related to the existence, uniqueness, and measurability of behavior that are discussed in our paper.

revision games as formalized by Kamada and Kandori (2018a,b), where players have opportunities to alter their actions before interacting at a designated time. In section 5.3, we apply our restrictions to a finite-horizon game that has a deadline by which a buyer, whose taste for a good is stochastically changing over time, must place an order with a seller, who has randomly arriving opportunities to fulfill the order.

Further applications of our framework are considered in the online appendix. Ortner (2016) specifies a bargaining model in continuous time, where the bargaining power of each player is governed by a Brownian motion. As mentioned in the online appendix, Kamada and Rao (2018) develop a model of bargaining and trade in which there is a transaction cost that may evolve according to a geometric Brownian motion. We also present an example in which a pair of criminals play a prisoner's dilemma at times that arrive according to a Poisson process, as well as continuously deciding whether to remain partners in crime. Another issue related to the timing of investment is technology adoption, which has been studied in continuous time by Fudenberg and Tirole (1985). The online appendix contains an example where two agents repeatedly decide on when to adopt a technology that has positive externalities and a stochastically evolving cost. A continuous-time formulation has also been used for models of adjustment costs, including the analysis of pricing by Caplin and Spulber (1987). In the online appendix, we examine interactions between a retailer and distributor in a model of inventory adjustment with a randomly changing stock as well as a model of exchange where prices evolve stochastically.

The rest of the paper proceeds as follows. In section 2, we specify the model. Section 3 shows that the concepts of traceability and frictionality together imply the existence of a unique path of play given a strategy profile. Section 4 deals with measurability issues and shows that the calculability assumption implies that agents' behavior is measurable. Section 5 contains the aforementioned applications of our methodology. In section 6, we compare our approach to an alternative approach in which the strategy space is not restricted but a certain payoff is assigned to nonmeasurable behavior. We investigate the relationship between the sets of equilibria under the two approaches. Section 7 concludes. All proofs are in the appendix. The online appendix provides further applications, extensions, and discussions.

2 Model

There are a finite number $n \geq 1$ of agents, and the set of agents is denoted by I. Time runs continuously in $[0, T)$ where $T \in (0, \infty]$. Each agent $i \in I$ has a measurable action space A_i. Let $(\Omega, \mathcal{F}, P)$ be a probability space, and let $\{\mathcal{F}_t\}_{t\in[0,T)}$ be a filtration of the sigma-algebra $\mathcal{F}$. The shock level at time t, which is denoted by s_t, evolves according to a stochastic process whose state space is S endowed with the sigma-algebra $\mathcal{B}(S)$.[6] The shock process $\{s_t\}_{t\in[0,T)}$ is assumed to be progressively measurable.[7] Moreover, $\{s_t\}_{t\in[0,T)}$ is action-independent. The probability space and state space do not depend in any way on the actions of the agents.

In every instant of time, each agent i observes the current realization of the shock level and chooses an action from a subset of A_i that can depend on the history. Formally, for each $i \in I$ and $t \in [0, T)$, let a^i_t represent the action that agent i chooses at time t. The collection $\{a^i_t\}_{t\in[0,T)}$ of agent i's actions indexed by time $t \in [0, T)$ is called an **action path**. Letting $u \in [0, T)$, the collection $\{a^i_t\}_{t\in[0,u)}$ is said to be an action path up to time u.

A history of the game is represented as $h = \{s_t, (a^i_t)_{i\in I}\}_{t\in[0,T)}$. That is, a history consists of the realization of the shock process along with the actions chosen by the agents at each time. Given a history h, a history h_u up to time u is defined as $\big(\{s_t\}_{t\in[0,u]}, \{(a^i_t)_{i\in I}\}_{t\in[0,u)}\big)$. Note that h_u includes information about the shock at time u but does not contain information about the actions at time u. By convention, h_0 is used to denote the null history $(s_0, \{\})$ at the start of the game. Letting H_t be the set of all histories up to time t, define $H = \bigcup_{t\in[0,T)} H_t$.

For each $i \in I$, define the feasible set of actions at every history by the function $\bar{A}_i : H \to 2^{A_i}$. We assume that there exists an action $z \in A_i$ such that $z \in \bar{A}_i(h_t)$ for any $h_t \in H$. The action z can be interpreted as "not moving," whereas an action other than z is regarded as a "move." A strategy for agent $i \in I$ is a map $\pi_i : H \to A_i$.[8] Let Π_i with generic element π_i represent the set of all strategies for agent i.[9] In addition,

[6] The notation $\mathcal{B}(C)$ represents the Borel sigma-algebra on a subset C of the real line.

[7] A stochastic process $\{x_t\}_{t\in[0,T)}$ can be treated for any $v \in [0, T)$ as a function $x(t, \omega)$ on the product space $[0, v] \times \Omega$. It is said that $\{x_t\}_{t\in[0,T)}$ is progressively measurable if for any $v \in [0, T)$ the function $x(t, \omega)$ is measurable with respect to the product sigma-algebra $\mathcal{B}([0, v]) \times \mathcal{F}_v$.

[8] Simon and Stinchcombe (1989) allow for sequential moves at a single moment of time. The definition of the strategy space here rules out such behavior. This restriction is innocuous except in the partnership and cooperation game between criminals in the online appendix, where we argue that restricting moves to happen at Poisson opportunities may cause an inefficient delay.

[9] The definition of strategies thus far does not eliminate the possibility of flow moves, whereby

a strategy π_i for agent $i \in I$ is said to be **feasible** if it satisfies the restriction that $\pi_i(h_t) \in \bar{A}(h_t)$ for any $h_t \in H$. Let $\bar{\Pi}_i$ denote the set of all feasible strategies for agent i.

3 Existence and Uniqueness of the Action Path

3.1 Examples

Since the model is formulated in continuous time, the definition of the strategy space is not trivial as noted by Simon and Stinchcombe (1989) and Bergin and MacLeod (1993). As explained below, we need to develop a novel approach for restricting the strategy space in our context. The following two examples illustrate the problems that our restrictions help eliminate.

Example 1. (No action path consistent with a given strategy profile) Suppose that $\bar{A}_i(h_t) = \{x, z\}$ for each $i \in I$ and every $h_t \in H$. We argue that there is no action path consistent with the following strategy profile. If there is a positive integer m such that the current time t is equal to $1/m$ and neither agent has chosen action x before time t, then each agent chooses x at time t. Otherwise, all agents choose action z. To see that there is no action path consistent with these strategies, notice that on any path of play of the given strategy profile, there must exist exactly one time $t > 0$ at which action x is taken. However, if the agents were to choose x at this time t, then they would be deviating from the given strategy profile because there exists $m < \infty$ large enough that $1/m \in (0, t)$. □

Example 2. (Multiple action paths consistent with a given strategy profile) Suppose again that $\bar{A}_i(h_t) = \{x, z\}$ for each $i \in I$ and every $h_t \in H$. We argue that there is more than one action path consistent with the following strategy profile. If each agent has chosen action x at time $1/n$ for every positive integer n greater than m, then each agent chooses x at time $1/m$. Otherwise, each agent chooses z.

The following two action paths can be outcomes of this strategy profile. First, all agents choose action z at all $t \in [0, T)$. Second, each agent chooses x if the current

---

actions other than z are taken for a positive measure of times. For example, when harvesting trees, agents may cut a tree continuously over time so that the amount cut at each instant of time is zero, but the total amount cut over an interval of time may be positive. The subsequent sections restrict the strategy space so as to avoid flow moves, thereby preventing such behavior in the tree harvesting problem (section 5.1).

time t is equal to $1/m$ for some positive integer m, and all agents choose z otherwise. It is straightforward to check by inspection that the agents are following the given strategy profile in each case. $\square$

In order to resolve these sorts of issues, Bergin and MacLeod (1993) develop an inertia condition, which requires agents to wait for a certain interval of time after changing actions. A suitable application of their restriction to our model would eliminate the pathologies above. Nonetheless, their approach is too restrictive in the current setting. While those authors study a deterministic environment, we consider a stochastic game in continuous time. For instance, in the tree harvesting problem (section 5.1), the proposed equilibrium of our model is such that for any $\epsilon > 0$, there is positive conditional probability of the agents being required to cut trees in the time interval $(u, u+\epsilon)$ given that the agents cut trees at time u. This violates the inertia condition, which requires that if the agents cut trees at time u, then there exists $\epsilon > 0$ such that they do not cut trees during the time interval $(u, u+\epsilon)$.

3.2 Traceability and Frictionality

We introduce a series of concepts so as to eliminate strategies like those in the examples above. We define three conditions, which we call consistency, traceability, and frictionality. Consistency is a property of histories, whereas traceability and frictionality are properties of strategies.

Definition 1. Given $i \in I$, the history $h = \left\{s_t, (a_t^j)_{j\in I}\right\}_{t\in[0,T)}$ is said to be **consistent** with strategy π_i at time t if $\pi_i(h_t) = a_t^i$.

Roughly speaking, a history is said to be consistent with a strategy for a given agent if the history is a possible outcome when that agent plays the strategy.

For any action path $\{b_t^i\}_{t\in[0,u)}$ of agent $i \in I$ up to an arbitrary time u, let $\Gamma_i(\{b_t^i\}_{t\in[0,u)})$ be the set consisting of any action path $\{a_t^i\}_{t\in[0,T)}$ such that $\{a_t^i\}_{t\in[0,u)} = \{b_t^i\}_{t\in[0,u)}$.

Definition 2. Given $i \in I$, the strategy $\pi_i \in \Pi_i$ is **traceable** if the following holds. Choose any history $k_u = \left(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\right)$ up to time u. Given that $\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$, there is conditional probability one that for any $\{a_t^{-i}\}_{t\in[0,T)} \in \Gamma_{-i}(\{b_t^{-i}\}_{t\in[0,u)})$, there exists $\{a_t^i\}_{t\in[0,T)} \in \Gamma_i(\{b_t^i\}_{t\in[0,u)})$ for which the history $h = \left\{s_t, (a_t^j)_{j\in I}\right\}_{t\in[0,T)}$ is consistent with π_i at each $t \in [u, T)$.

Intuitively, a strategy for a given agent is said to be traceable if for any behavior by the other agent, there exists a history that is consistent with the strategy. Traceability excludes the strategy in example 1.

Let $\Xi_i(t)$ denote the set consisting of every action path $\{a^i_\tau\}_{\tau\in[0,T)}$ of agent $i \in I$ for which there exists no $u > t$ such that the set $\{\tau \in [t,u] : a^i_\tau \neq z\}$ contains infinitely many elements.

Definition 3. Given $i \in I$, the strategy $\pi_i \in \Pi_i$ is **frictional** if the following holds. Choose any history $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b^j_t)_{j\in I}\}_{t\in[0,u)}\big)$ up to time u. Given that $\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$, there is conditional probability one that $\{a^i_t\}_{t\in[0,T)} \in \Xi_i(u)$ for all $\{a^i_t\}_{t\in[0,T)} \in \Gamma_i(\{b^i_t\}_{t\in[0,u)})$ such that there exists $\{a^{-i}_t\}_{t\in[0,T)} \in \Gamma_{-i}(\{b^{-i}_t\}_{t\in[0,u)})$ for which the history $h = \{s_t, (a^j_t)_{j\in I}\}_{t\in[0,T)}$ is consistent with π_i at each $t \in [u,T)$.

Intuitively, a strategy for a given agent is said to be frictional if any history that is consistent with that strategy has the property that the agent moves (i.e., takes an action other than z) only a finite number of times in any finite time interval. Frictionality excludes the strategy in example 2. It also rules out flow moves.

Traceability and frictionality restrict the strategy space for each agent.[10] Note that these conditions are defined in terms of the strategy of an individual agent as opposed to the strategy profile of all agents. In addition, observe that these conditions impose requirements on a strategy not only after the null history but after every history up to any time. We can now state a main result.

Theorem 1. *Choose any profile $(\pi_j)_{j\in I}$ of strategies that satisfy traceability and frictionality. Choose any history $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b^j_t)_{j\in I}\}_{t\in[0,u)}\big)$ up to time u.*

1. *Given that $\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$, there is conditional probability one that there exists a unique profile $(\{a^j_t\}_{t\in[0,T)})_{j\in I}$ of action paths with $\{a^j_t\}_{t\in[0,T)} \in \Gamma_j(\{a^j_t\}_{t\in[0,u)})$ for each $j \in I$ such that the history $h = \big\{s_t, (a^j_t)_{j\in I}\big\}_{t\in[0,T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u,T)$.*

[10]The inertia condition in Bergin and MacLeod (1993) would imply traceability but not frictionality if applied to the current setting. There would exist a history consistent with a given strategy at each time, but such a history might have infinitely many non-z actions in a finite time interval. For example, consider a strategy that requires a non-z action to be chosen if and only if the time t is such that $t = 1 - 1/2^n$ for some positive integer n. This strategy, which would satisfy inertia, is traceable but not frictional.

2. *Given that* $\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$*, there is conditional probability one that the action paths in the first part of the theorem satisfy* $\{a_t^j\}_{t\in[0,T)} \in \Xi_j(u)$ *for each* $j \in I$*.*

To paraphrase, suppose that each agent uses a strategy that satisfies the restrictions. For any realization of the shock process following a history up to a given time, the game has a unique action path. Moreover, the action path is such that there is only a finite number of non-z actions in any finite time interval. The proof entails the following complication. The strategies of an agent's opponents can specify actions that depend on the agent's behavior. However, the definitions of traceability and frictionality refer not to the opponents' strategies, but only to the opponents' action paths. Thus, the existence and uniqueness of action paths consistent with a given strategy profile are not immediate consequences of these assumptions. Indeed, there may exist a profile of traceable strategies such that no history is consistent with all players' strategies at all times.[11] The proof of theorem 1 employs frictionality as well to identify one by one each time a non-z action is taken, utilizing the two restrictions at every step.

Hereafter, we restrict attention to strategies that satisfy traceability and frictionality. Henceforth, let Π_i^{TF} denote the set consisting of every traceable and frictional strategy for agent i.

4 Measurability of the Action Process

4.1 An Example

Another question concerning the specification of the model involves the measurability of the path of play. This issue causes little trouble in nonstochastic situations including Bergin and MacLeod (1993). Indeed, if there were no uncertainty about the shock level, then constraints such as traceability and frictionality would be enough to define the payoff to each agent. However, the stochastic nature of the shock complicates matters since agents can condition their behavior on moves by Nature. In order to compute expected payoffs, the action process should be progressively measurable. In continuous-time games with imperfect monitoring like Sannikov (2007), it is unproblematic to define the strategy space so as to ensure the measurability of the actions

[11] See the online appendix for an example of such a strategy profile.

taken by agents. Nonetheless, this task is less straightforward in the current context where the actions induced by an agent's strategy may or may not be measurable depending on the other agents' behavior. The example below demonstrates this sort of interdependence in terms of the measurability of actions.

Example 3. (Actions induced by a strategy may or may not be measurable) Let $\{s_t\}_{t\in[0,\infty)}$ be an arbitrary stochastic process with state space $S \subseteq \mathbb{R}_{++}$. Assume that there exists $\tilde{S} \subseteq S$ along with $\tilde{t} > 0$ such that $\{\omega \in \Omega : s_{\tilde{t}}(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$.[12] Suppose $I = \{1, 2\}$ and that $\bar{A}_i(h_t) = \{x, z\}$ for each $i \in I$ and every $h_t \in H$.

Suppose that agent 1 plays the strategy $\tilde{\pi}_1$ defined as follows. If agent 2 chooses z at time $\tilde{t}$, then agent 1 is required to choose x at time $2\tilde{t}$. Otherwise, agent 1 is required to choose x at time $2\tilde{t}$ if the shock $s_{\tilde{t}}$ at time $\tilde{t}$ is in the set $\tilde{S}$ and to choose x at time $\tilde{t} + s_{\tilde{t}}$ if the shock $s_{\tilde{t}}$ at time $\tilde{t}$ is not in the set $\tilde{S}$. Suppose that agent 2 plays the strategy $\tilde{\pi}_2$ defined as follows. Agent 2 is required to choose x at time $\tilde{t}$. If agent 1 chooses x at time $2\tilde{t}$, then agent 2 is required to choose x at time $3\tilde{t}$. The agents do not take action x except as specified above.

For $i \in \{1, 2\}$ and $t \in [0, T)$, the action a_t^i of agent i at time t can be treated as a function from the probability space $(\Omega, \mathcal{F}, P)$ to $\{x, z\}$. Even if the action a_t^2 of agent 2 is measurable at each time $t \in [0, T)$, the strategy $\tilde{\pi}_1$ can induce a nonmeasurable action $a_{2\tilde{t}}^1$ by agent 1 at time $2\tilde{t}$. By contrast, the strategy $\tilde{\pi}_2$ induces a measurable action a_t^2 by agent 2 at every time $t \in [0, T)$, unless the action $a_{2\tilde{t}}^1$ of agent 1 at time $2\tilde{t}$ is nonmeasurable.

In the ensuing analysis, we seek to eliminate strategies like $\tilde{\pi}_1$ because such strategies can generate nonmeasurable behavior by one agent even if the actions of the other agents are measurable. Nonetheless, we wish to retain strategies like $\tilde{\pi}_2$ that ensure the measurability of one agent's actions given the measurability of the other agents' actions.

There is no simple way to restrict the strategy space so as to remove $\tilde{\pi}_1$ but not $\tilde{\pi}_2$. To see this, consider the following two possibilities. First, suppose that we delete strategies such that for some strategy of one's opponent, one's behavior is not measurable. However, this would rule out both $\tilde{\pi}_1$ and $\tilde{\pi}_2$. Second, suppose that we

[12] For instance, let the probability space be the interval $(0, 1)$ with the Borel sigma-algebra and the Lebesgue measure, and let the state space be the interval $(0, 1)$. Suppose that $s_{\tilde{t}}$ is the uniform random variable defined by $s_{\tilde{t}}(\omega) = \omega$. Then a Vitali set can be used to provide an example of $\tilde{S}$.

delete strategies such that for every strategy of one's opponent, one's behavior is not measurable. However, this would rule out neither $\tilde{\pi}_1$ nor $\tilde{\pi}_2$. □

In discrete time, it is relatively uncomplicated to specify a game so that the path of play is measurable. The analyst can simply require the strategies of the players to be measurable functions from the history up to each time to the actions at that time. Under such an assumption, if the path of play is measurable up to and including any given time, then the behavior of the agents will be measurable in the following period as well. Hence, the measurability of the path of play up to every time can be shown by induction. In continuous time, however, such an iterative procedure is not applicable because the next period after any given time is not well defined.

4.2 Two-Step Procedure

In order to suitably phrase the definition, a two-step procedure is adopted. We first specify a very restrictive set of strategies that always induce measurable behavior, and the first set is then used to construct a more inclusive second set. The elements of the resulting strategy space are said to be calculable.

We begin by providing a formal definition of the action processes. Let $\pi = (\pi_j)_{j\in I}$ with $\pi_j \in \Pi_j^{TF}$ for $j \in I$ be a profile of traceable and frictional strategies. Choose an arbitrary time u. Fix any path $\{(b_t^j)_{j\in I}\}_{t\in[0,u)}$ of actions for the agents up to this time. Define the history up to time u by $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\big)$, where $\{g_t\}_{t\in[0,u]}$ is any realization of shock levels from time 0 to u. Given that $\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$, let $\big(\{\phi_t^j(k_u, \{s_\tau\}_{\tau\in(u,T)}, \pi)\}_{t\in[0,T)}\big)_{j\in I}$ with $\{\phi_t^j(k_u, \{s_\tau\}_{\tau\in(u,T)}, \pi)\}_{t\in[0,T)} \in \Gamma_j(\{b_t^j\}_{t\in[0,u)})$ for each $j \in I$ be a profile of action paths for which the history $\big\{s_t, [\phi_t^j(k_u, \{s_\tau\}_{\tau\in(u,T)}, \pi)]_{j\in I}\big\}_{t\in[0,T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T)$. According to theorem 1, such a profile of action paths exists and is unique with conditional probability one. Furthermore, these action paths satisfy $\{\phi_t^j(k_u, \{s_\tau\}_{\tau\in(u,T)}, \pi)\}_{t\in[0,T)} \in \Xi_j(u)$ for each $j \in I$ with conditional probability one.

Denoting $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$ and $\pi = (\pi_j)_{j\in I}$, the **action process** $\xi_b^i(\pi)$ for $i \in I$ represents the stochastic process defined as follows. At any time $t \in [0, u)$, $\xi_b^i(\pi) = z$ holds. Let $\tilde{g} = \{\tilde{g}_t\}_{t\in[0,u]}$ represent the realization of shock levels until time u, and denote the resulting history up to time u by $\tilde{k}_u = (b, \tilde{g})$. Given the realization of the shock $\{s_\tau\}_{\tau\in(u,T)}$ after time u, $\xi_b^i(\pi_1, \pi_2) = \phi_t^i(\tilde{k}_u, \{s_\tau\}_{\tau\in(u,T)}, \pi)$ holds at each time $t \in [u, T)$. The state space of $\xi_b^i(\pi)$ is A_i with sigma-algebra $\mathcal{B}(A_i)$. To paraphrase, let

b signify any path of actions for the agents up to an arbitrary time u. The stochastic process $\xi_b^i(\pi)$ is simply equal to z up to this time. Thereafter, $\xi_b^i(\pi)$ records the actions chosen by agent i when the strategy profile π is played.[13]

We now restrict the strategy space so as to ensure the progressive measurability of the action processes.

First step: quantitative strategies

The first step is to define a set of traceable and frictional strategies that induce measurable behavior by one agent whenever the other agents play traceable and frictional strategies. The elements of this set are said to be quantitative.

Definition 4. For $i \in I$, the strategy $\pi_i \in \Pi_i^{TF}$ is **quantitative** if it belongs to the set Π_i^Q consisting of any $\pi_i \in \Pi_i^{TF}$ such that the stochastic process $\xi_b^i(\pi_i, \pi_{-i})$ is progressively measurable for all b and every $\pi_{-i} \in \times_{j \neq i} \Pi_j^{TF}$.

The composition of the set Π_i^Q can vary with the specification of the shock process. It contains traceable and frictional strategies that depend neither on the realized values of the shock process nor on the actions of one's opponents. An example of a quantitative strategy would be the strategy that requires agent i to choose action $x \neq z$ at time 1 regardless of the actions of agents $-i$ and that specifies action z by agent i at other times. As shown in the online appendix, the set Π_i^Q may also contain some strategies that are contingent on the realization of the shock and the behavior of one's opponents. Nonetheless, the set Π_i^Q is extremely restrictive as a strategy space. In particular, the strategies $\tilde{\pi}_1$ and $\tilde{\pi}_2$ in example 3 are both excluded.

Second step: calculable strategies

The second step is to define a set of traceable and frictional strategies that induce measurable behavior by one agent whenever the other agents play quantitative strategies. That is, the set Π_i^Q is used to construct a larger strategy space Π_i^C, the elements of which are said to be calculable.

Definition 5. For $i \in I$, the strategy $\pi_i \in \Pi_i^{TF}$ is **calculable** if it belongs to the

[13]For simplicity in defining the process $\xi_b^i(\pi)$, we consider the continuation path of play for each agent i under strategy profile π when the agents follow the fixed action path b up to time u. The results in this section would not change under an alternative definition in which at a history k_u up to time u on the path of play of π, we instead consider the continuation path of play under π when the agents follow strategy profile π up to time u.

set Π_i^C consisting of any $\pi_i \in \Pi_i^{TF}$ such that the stochastic process $\xi_b^i(\pi_i, \pi_{-i})$ is progressively measurable for all b and every $\pi_{-i} \in \times_{j \neq i} \Pi_j^Q$.

The strategy space Π_i^C admits a relatively broad range of behavior. It allows for strategies that depend on the actions of one's opponents, including many grim-trigger strategies. For example, consider the strategy that requires agent i to choose action x at time 2 if and only if all the other agents choose action x at time 1 and that specifies no other non-z actions by agent i. This is a calculable strategy. If all the opponents of i play quantitative strategies, then the action profile by those opponents at time 1 is a measurable function from $(\Omega, \mathcal{F}_1)$ to $\times_{j \neq i} A_j$, and so the action by agent i at time 2 is a measurable function from $(\Omega, \mathcal{F}_2)$ to A_i.

The result below shows that restricting the strategy space ensures measurable behavior. If each agent plays a calculable strategy, then the stochastic process encoding the actions of each agent is progressively measurable.

Theorem 2. *If $\pi_i \in \Pi_i^C$ for $i \in I$, then the stochastic process $\xi_b^i(\pi_i, \pi_{-i})$ is progressively measurable for all b and each $i \in I$.*

An iterative argument is used to establish the preceding result. Let (π_i, π_{-i}) be a profile of calculable strategies, and let b denote the past actions of the agents. We start by constructing a progressively measurable stochastic process that is the same as $\xi_b^i(\pi_i, \pi_{-i})$ up to and including the time of the first non-z action by some player. We then do the same for the second non-z action, third non-z action, and so on.[14] Intuitively, if no agent has chosen a non-z action yet, then it is as if each player's opponents are following the strategy of always choosing z, which belongs to the set $\times_{j \neq i} \Pi_j^Q$. Hence, if a player is using a strategy in Π_i^C, then its action process will satisfy the requirements for progressive measurability up to and including the time when a non-z action is chosen.[15]

The final result justifies our restriction on strategies. In order to guarantee the measurability of the action process, we seek to eliminate strategies that generate

[14] The second part of theorem 1 implies that, with probability one, the time of the k^{th} non-z action is well defined for every positive integer k.

[15] The set of calculable strategies cannot be defined using wording similar to the traceability and frictionality assumptions. In particular, suppose that agent i is simply required to play a strategy that for any action path of the opponents, induces a progressively measurable action process for agent i. If the opponents play quantitative strategies, then agent i's action process may not be progressively measurable because, for example, the time of a non-z action by agent i may depend on the opponents choosing non-z actions at some time in a nonmeasurable set.

nonmeasurable behavior. There is no obvious reason to delete a quantitative strategy because it induces measurable actions virtually regardless of the strategies played by the other agents. Given that any strategy in Π_i^Q is permitted, the set Π_i^C of calculable strategies is the largest strategy space for each agent that ensures the measurability of actions.

Proposition 1. *For each $i \in I$, let Ψ_i be any strategy space with $\Pi_i^Q \subseteq \Psi_i \subseteq \Pi_i^{TF}$. Suppose that the stochastic process $\xi_b^i(\pi)$ is progressively measurable for all b, any $\pi \in \times_{j \in I} \Psi_j$, and each $i \in I$. Then $\Psi_i \subseteq \Pi_i^C$ for $i \in I$.*

The logic is simple. Suppose that agent i uses some non-calculable strategy π_i'. By definition, there exists a quantitative strategy π_{-i}' such that the actions of agent i may not be progressively measurable when i's opponents play π_{-i}'. Hence, measurable behavior is not ensured if agent i can use some non-calculable strategy and the opponents can play any quantitative strategies.

The remainder of the paper considers for each $i \in I$ primarily strategies in Π_i^C.

4.3 Expected Payoffs

This section defines expected payoffs and equilibrium concepts. For each $i \in I$, let $v_i : (\times_{j \in I} A_j) \times S \to \mathbb{R}$ be a measurable utility function. Choose any history $h = \{s_\tau, (a_\tau^j)_{j \in I}\}_{\tau \in [0,T)}$ such that $\{a_\tau^j\}_{\tau \in [0,T)} \in \Xi_j(t)$ for each $j \in I$. The realized payoff to agent i at time t is given by:

$$V_t^i(h) = \sum_{\tau \in M_t(h)} v_i[(a_\tau^j)_{j \in I}, s_\tau],$$

where the set $M_t(h) = \{\tau \in [t, T) : \exists\, j \text{ s.t. } a_\tau^j \neq z\}$ represents the set of times from t onwards at which some agent moves under the given history.

That is, the payoff is the sum of discrete utilities from the times at which at least one agent chooses an action other than z.[16] It follows from the second part of theorem 1 that, with probability one, the number of non-z actions chosen by each agent is countable. Note also that discounting is not explicitly modeled. However, one can include it as part of the shock. For example, there may exist w_i such that

[16]The specification of the realized payoff reflects the normalization that for each element of S, the discrete utility from the action profile in which all agents choose z is 0.

$v_i[(a^j_\tau)_{j\in I}, s_\tau] = e^{-\rho(\tau-t)} w_i[(a^j_\tau)_{j\in I}]$. We also define:

$$V_t^{i,p}(h) = \sum_{\tau\in M_t(h)} \max\{0, v_i[(a^j_\tau)_{j\in I}, s_\tau]\} \text{ and } V_t^{i,n}(h) = \sum_{\tau\in M_t(h)} \min\{0, v_i[(a^j_\tau)_{j\in I}, s_\tau]\},$$

whence the realized payoff can be expressed as $V^i_t(h) = V_t^{i,p}(h) + V_t^{i,n}(h)$.

Under the preceding specification, no agent experiences a flow payoff. However, some of our examples in the online appendix illustrate how a game can be straightforwardly reformulated so that agents experience a stream of flow payoffs. We also interpret the different formulations, explaining how they emphasize different features of the economic setting.

Let $\bar\Pi_i^{TF} = \Pi_i^{TF} \cap \bar\Pi_i$ denote the set of strategies for each agent $i \in I$ that are feasible as well as traceable and frictional. We impose the following **one-sided boundedness** condition.[17] For any history $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b^j_t)_{j\in I}\}_{t\in[0,u)}\big)$ up to time u and any strategy profile $\pi \in \times_{j\in I}\bar\Pi_j^{TF}$ such that the process $\xi^i_b(\pi)$ with $b = \{(b^j_t)_{j\in I}\}_{t\in[0,u)}$ is progressively measurable for all $i \in I$, we have either $V_i^p(k_u,\pi) < \infty$ or $V_i^n(k_u,\pi) > -\infty$ for every $i \in I$, where we define for $o \in \{n,p\}$:

$$\begin{aligned} V_i^o(k_u,\pi) = \mathbb{E}_{\{s_t\}_{t\in(u,T)}}\big[V_u^{i,o}\big(\big\{s_t, [\phi^j_t(k_u, \{s_\tau\}_{\tau\in(u,T)}, \pi)]_{j\in I}\big\}_{t\in[0,T)}\big) \\ |\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}\big]. \end{aligned}$$

The conditional expectation is taken with respect to $\{s_t\}_{t\in(u,T)}$ given that $\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$.[18]

The expected payoff to agent i at k_u is specified as:

$$\begin{aligned} V_i(k_u,\pi) = \mathbb{E}_{\{s_t\}_{t\in(u,T)}}\big[V_u^{i}\big(\big\{s_t, [\phi^j_t(k_u, \{s_\tau\}_{\tau\in(u,T)}, \pi)]_{j\in I}\big\}_{t\in[0,T)}\big) \\ |\{s_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}\big], \end{aligned} \tag{1}$$

where the conditional expectation can be expressed as $V_i(k_u,\pi) = V_i^p(k_u,\pi) + V_i^n(k_u,\pi)$. Let $\bar\Pi_i^C = \Pi_i^C \cap \bar\Pi_i$ denote the set of strategies for each agent $i \in I$ that are feasible as

[17] All of the applications considered in section 5 and the online appendix satisfy this property.

[18] The following property is sufficient for the one-sided boundedness condition to hold. For any u, letting $\bar H$ denote the set consisting of every history $h = \{s_t, (a^i_t)_{i\in I}\}_{t\in[0,T)}$ such that $\{a^i_t\}_{t\in[0,T)} \in \Xi_i(u)$ for each $i \in I$ and such that $a^i_t \in \bar A_i\big(\{s_\tau\}_{\tau\in[0,t]}, \{(a^j_\tau)_{j\in I}\}_{\tau\in[0,t)}\big)$ for all $t \geq [u,T)$ and each $i \in I$, we have either $\sup_{h\in\bar H}\sum_{t\in M_u(h)} \max\{0, v_i[(a^j_t)_{j\in I}, s_t]\} < \infty$ or $\inf_{h\in\bar H}\sum_{t\in M_u(h)} \min\{0, v_i[(a^j_t)_{j\in I}, s_t]\} > -\infty$ for every $i \in I$.

well as calculable. For any $\pi \in \times_{j\in I}\bar{\Pi}_j^C$, the expected payoff is well defined because theorem 1 implies that the realized payoff can be uniquely computed with conditional probability one, and theorem 2 along with the one-sided boundedness property ensures the existence of the conditional expectation.

Having defined the expected payoffs, we can now define subgame-perfect equilibrium (SPE). Formally, a strategy profile π with $\pi_j \in \bar{\Pi}_j^C$ for $j \in I$ is a **subgame-perfect equilibrium** if for any history $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\big)$ up to time u, the expected payoff to agent $i \in I$ at k_u satisfies $V_i(k_u, \pi) \geq V_i[k_u, (\pi_i', \pi_{-i})]$ for any $\pi_i' \in \bar{\Pi}_i^C$.[19]

5 Applications

The examples below and in the online appendix illustrate the applicability of our methodology to a broad range of settings. Although other methods could be used to analyze each specific application, there is no alternative approach that covers all of the examples that we present.

5.1 Tree Harvesting Problem

There are n woodcutters harvesting trees in a common forest. The volume of forest resources evolves stochastically over time. It can increase due to natural growth, which depends on the weather. It can also decrease because of damage to trees by wind or fire. The forest is large, making it reasonable to assume that the volume follows a continuous process. Specifically, we assume that it follows an arithmetic Brownian motion with a lower bound of zero. Formally, let b_t with $b_0 = 0$ be a Brownian motion having arbitrary drift μ and positive volatility σ: $db_t = \mu dt + \sigma dz_t$. The volume q_t of the forest at time t is given by the greater of $b_t - p$ and 0, where p is the total amount cut in the past. The decision to harvest trees is described in what follows.

Each woodcutter i decides on an amount f_t^i to cut at every moment of time $t \in [0, \infty)$. Suppose that at time t, the set of woodcutters who cut trees is M where

[19] The online appendix provides an alternative but equivalent definition that simplifies checking whether a strategy profile is an SPE given the restrictions on the strategy space.

$|M| \geq 1$. Then the utility of each woodcutter $i \in M$ at that time is:

$$\frac{f_t^i}{\sum_{j\in M} f_t^j}\left[\min\left\{\sum_{j\in M} f_t^j, q_t\right\} - \kappa\right],$$

which represents the difference between the benefit and cost of cutting trees. If the amount of trees that all the woodcutters are planning to cut does not exceed the current volume of the forest, then each woodcutter harvests its planned amount. Otherwise, the amount cut by each woodcutter is rationed proportionally. The total cost of cutting trees is κ, which is split among the woodcutters harvesting trees at time t proportionally to the amount each of them cuts. Any woodcutter who does not participate in cutting trees at time t receives the payoff 0 at that time. The woodcutters discount the future at rate $\rho > 0$.

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each woodcutter $i = 1, 2$.[20] Call this game with such strategy spaces the *tree harvesting game*. It is characterized by $(n, \mu, \sigma, \kappa, \rho)$. The analysis in sections 3 and 4 implies that a subgame-perfect equilibrium is well defined. A symmetric SPE is said to be maximal if there is no symmetric SPE that yields a higher expected payoff to each player.

Let

$$x^* = [1 + \alpha\kappa + W(-e^{-1-\alpha\kappa})]/\alpha \quad \text{and} \quad \bar{x} = \ln[n/(n-1)]/\alpha,$$

where W is the principal branch of the Lambert W function.[21] Define:

$$\hat{x} = \begin{cases} \kappa & \text{if } \kappa \geq \bar{x} \\ \min\{x^*, \bar{x}\} & \text{if } \kappa < \bar{x} \end{cases}.$$

Proposition 2. *The tree harvesting game has a maximal equilibrium for any $(n, \mu, \sigma, \kappa, \rho)$. Moreover, the path of play in any maximal equilibrium is such that, with probability one, the m^{th} cutting of trees occurs at the m^{th} time the volume reaches $\hat{x}$ for every positive integer m, where the trees are cut to volume 0 on each cutting.*

There is a simple intuition for why cutting the trees to zero is optimal. To see this, consider a strategy profile in which the trees are cut to the volume $y > 0$ whenever

[20] Formal definitions of histories and strategy spaces are provided in appendix B.1.1.

[21] The Lambert W function is defined by $z = W(z)e^{W(z)}$ for any complex number z.

they reach a volume $x > y$ on the equilibrium path. This policy would be dominated by one in which the trees are cut to zero when they first reach a volume x and are thereafter cut to zero whenever they reach a volume of $x - y$.

Remark 1.

1. (Violation of inertia) In any symmetric SPE, each agent's strategy violates inertia. As shown in the proof of the proposition, the supremum over all symmetric strategy profiles of the expected payoff to each player at any history up to an arbitrary time with a current volume of c is bounded above by the sum of c/n and a constant. Hence, there exists a threshold $\bar{c} > \kappa$ such that there is no symmetric SPE that yields an expected payoff to each player of at least $c - \kappa$ when played at a history up to an arbitrary time in which the volume is currently $c > \bar{c}$. Fix any history up to an arbitrary time t with $q_t > 0$. For any $\epsilon > 0$, there is positive conditional probability of there being a time $\tau \in (t, t + \epsilon)$ such that $b_\tau - b_t > \bar{c}$. If such a time τ exists, then it can be argued that trees must be harvested during the time interval $(t, t + \epsilon)$ in a symmetric SPE, for otherwise a woodcutter would have an incentive to deviate by unilaterally cutting trees in this interval. Thus, there cannot exist $\epsilon > 0$ such that a player does not move during the time interval $(t, t + \epsilon)$, meaning that inertia is not satisfied.[22]

2. (Off-path strategies) In a maximal equilibrium, there are multiple possibilities for off-path strategies, but in any off-path strategies, the continuation payoff is zero, which is the minmax payoff of each player. One possibility is that each woodcutter cuts trees at time t if and only if $q_t \geq \kappa$. Another possibility is as follows. Let M be a positive integer, and let $c \in (0, \kappa)$. For $m \leq M$, the m^{th} cutting occurs when the volume reaches c for the m^{th} time, and the trees are cut to zero on each cutting. After the M^{th} cutting, the woodcutters play a maximal equilibrium. If there is any deviation from this path of play, then the agents play the aforementioned SPE in which trees are cut at time t if and only if $q_t \geq \kappa$. The values of M and c are chosen so that the ex ante expected payoff of each player is equal to zero.

[22] Bergin and MacLeod (1993) also expand the strategy space to the completion of the set of inertia strategies, but this extension does not necessarily apply for the same reason as why the inertia assumption does not. Some strategy profiles satisfying our restrictions cannot be expressed as the limit of a Cauchy sequence of strategy profiles satisfying inertia, as shown in the online appendix of Kamada and Rao (2018) in the context of their model. We are currently working on a paper about the problem of generally defining the completion of the set of inertia strategies.

3. (Relationship to admissibility) Simon and Stinchcombe (1989) classify a strategy as admissible if it satisfies three properties F1-F3. Assumption F1, which restricts the number of moves to be uniformly bounded, is violated by any symmetric SPE, in which there is no upper bound on the number of times that trees may be harvested in any proper time interval. Assumption F2, which involves the piecewise continuity of the path of play, is essentially satisfied because trees are cut only finitely many times in a finite time interval with probability one. Assumption F3, which in some sense requires behavior to depend continuously on the history, is violated by a maximal equilibrium, in which agents use trigger strategies where even a small deviation incurs a large punishment.

5.2 Entry Game between Competing Firms

The inertia condition may not be suitable for analyzing investment games, in which agents optimally time their decisions based on a shock process that may be modeled by a diffusion process such as geometric Brownian motion. Our restrictions on strategy spaces are useful in such situations. The following example, in which investment is modelled as an entry into a market, illustrates the violation of inertia and the usefulness of our approach.

Two firms, 1 and 2, are deciding whether to enter a market. Assume that both firms are initially out of the market. At each moment of time $t \in [0, \infty)$ such that no firm has entered, each firm can choose between two moves: get in I and accommodate A. Choosing I means that the firm enters the market. When a firm chooses A at time t, it enters the market if and only if the other firm does not choose I at the same time t. Once a firm has entered, it cannot move any longer. At any time such that one firm has entered but not the other, then the latter firm chooses whether or not to follow F the other firm by entering. In any event, not moving (i.e., choosing action z) means that the firm does not enter for now or has already entered.

Let c_t be the discrete cost incurred by a firm when it enters the market at time t. The entry cost evolves according to a geometric Brownian motion: $dc_t = \mu c_t dt + \sigma c_t dz_t$, with initial condition $c_0 = \tilde{c}$ for some $\tilde{c} \in \mathbb{R}_{++}$. Let b_1 be the discrete benefit to a firm from entering the market at time t if it is the only firm to enter the market up to and including time t. Let b_2 be the discrete benefit from entering at time t if both firms enter the market at or before time t. Assume that $b_1 > b_2 > 0$ and $c_0 > b_1$. Agents discount the future at rate $\rho > 0$.

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each player $i = 1, 2$.[23] Call the game with such strategy spaces the *entry game*. It is characterized by $(b_1, b_2, \mu, \sigma, \rho)$. The analysis in sections 3 and 4 implies that a subgame-perfect equilibrium is well defined.

Below we characterize Markov perfect equilibria, which we define as SPE in Markov strategies.[24] A strategy is said to be Markov if the action prescribed at any history up to a given time depends only on the current value of the cost and whether or not each firm has already entered.

Proposition 3. *In the entry game with $(b_1, b_2, \mu, \sigma, \rho)$, there exist κ_1 and κ_2 with $0 < \kappa_2 < \kappa_1 < \infty$ such that in any Markov perfect equilibrium, the following hold at any time t.*

1. *Suppose that no firm has entered yet.*
   (a) *If $\kappa_1 < c_t$, then both firms choose z.*
   (b) *If $\kappa_2 < c_t \leq \kappa_1$, then one firm chooses I, and the other firm chooses A.*
   (c) *If $c_t \leq \kappa_2$, then both firms choose I, or both firms choose A.*
2. *Suppose that firm i has already entered but j has not.*
   (a) *If $\kappa_2 < c_t$, then both firms choose z.*
   (b) *If $c_t \leq \kappa_2$, then i chooses z while j chooses F.*

That is, on the path of play of any Markov perfect equilibrium, one firm enters when the cost reaches κ_1 for the first time, and the other firm enters when the cost reaches κ_2 for the first time. The cutoffs κ_1 and κ_2 are such that each firm is indifferent between being the first and second entrant (i.e., taking action I and A when the cost c_t reaches κ_1 for the first time). This is because if being the first entrant is better, then the second entrant would deviate by choosing I or A just before c_t reaches κ_1 for the first time, and if being the second entrant is better, then the first entrant would deviate by choosing z whenever $c_t > \kappa_2$ and choosing F at the first time c_t reaches κ_2.

[23] Formal definitions of histories and strategy spaces are provided in appendix B.2.1.

[24] Without the restriction to Markov perfect equilibrium, there exist SPE in which the identity of the first entrant depends on the path of the cost process in an arbitrary manner. Here we rule out such complications by assuming the Markov property to focus on key issues.

Remark 2. 1. (Violation of inertia) In any Markov perfect equilibrium, there is a firm whose strategy does not satisfy inertia. To see this, fix a history up to an arbitrary time t in which the cost is currently $c_t \in (\kappa_2, \kappa_1]$ and there has been no previous entry. Consider a firm that takes action A at such a history. For any $\epsilon > 0$, there is positive conditional probability that $c_\tau \leq \kappa_2$ for some $\tau \in (t, t+\epsilon)$, which implies that this firm takes action F in the time interval $(t, t+\epsilon)$. Thus, there cannot exist $\epsilon > 0$ such that this firm does not move during the time interval $(t, t+\epsilon)$, meaning that inertia is violated.[25]

2. (Relationship to admissibility) Assumptions F1 and F2 of Simon and Stinchcombe (1989) are satisfied by any feasible strategies. Noting that each firm can move at most twice, the number of moves is uniformly bounded, and the path of play is piecewise continuous. Moreover, the equilibrium strategies have property F3, which requires a particular form of continuity for actions in histories. The behavior of each firm depends on which moves were taken in the past but not on the times when they were taken.

3. (Comparative statics) For each $j = 1, 2$, κ_j is increasing in the benefit b_j at entry because the j^{th} entry is more attractive if κ_j is higher. The cutoff κ_1 is decreasing in b_2 because the expected payoff of the second entrant is greater when the value of b_2 is higher, so that the pressure to be the first entrant is lower and the first entrant can wait for the cost to fall. The cutoff κ_2 clearly does not depend on b_1. Both cutoffs are increasing in μ and ρ because the future gains from waiting for the entry cost to fall are lower when these parameters are higher. The cutoffs are decreasing in σ since the effect of σ on the probability of the cost falling is opposite to that of μ.[26]

[25] If action A were not available, then a Markov perfect equilibrium would not exist. For example, there cannot be an equilibrium in which when neither firm has entered yet, one firm chooses I if the current cost is no greater than κ_1 and chooses z otherwise, and the other firm chooses I if the current cost is no greater than κ_2 and chooses z otherwise. In such a strategy profile, if the cost were currently κ_1 for the first time and neither firm has entered yet, then the former firm could profitably deviate by choosing I at the first time the cost reaches κ and choosing z otherwise, where $\kappa \in (\kappa_2, \kappa_1)$.

[26] Proofs of these comparative statics results are provided in the appendix.

5.3 Model of Ordering with a Deadline

The following is an example in which one player can take a non-z action (i.e., has an opportunity to move) only at the arrival times of a Poisson process. Requiring each player to move only at discrete state changes (i.e., Poisson hits) may perhaps suffice for the purpose of restricting strategy spaces so as to ensure a well defined outcome. Nonetheless, in the example below, there is another player who can move at any time in an interval of the real line, and restricting that player to move only at Poisson arrival times is problematic due to the nonstationarity of the model.

Consider a buyer B who is contemplating the timing of an order to a seller S, who faces a predetermined deadline for providing a good. The game is played in continuous time, where time is denoted by $t \in [0, T)$ with $T > 0$. The buyer B's taste x_t evolves according to a Brownian motion with zero drift and positive volatility σ: $dx_t = \sigma dz_t$, with initial condition $x_0 = \tilde{x}$ for some $\tilde{x} \in \mathbb{R}$. The deadline may, for example, correspond to Christmas day, and the good may be a Christmas present. B's taste may vary over time because, for example, the gift is intended for child who has fickle preferences.

At each moment in time, B observes her taste, and if she has not yet ordered, chooses whether or not to place an order indicating her preferred specification of the good. Once an order is placed, B is unable to revise it. After learning the specification, S has stochastic chances arriving according to a Poisson process with parameter $\lambda > 0$ to produce and supply the ordered good to B. Specifically, at each Poisson arrival time, S chooses whether to supply the good or not.

If B places an order and the order is fulfilled, then her utility is $v - (s - x_T)^2 - p$, where $v > 0$ is her valuation for a good whose specification perfectly matches her taste, s is the specification of the good actually purchased, and $p \in (0, v)$ is the fixed price that B pays. Otherwise, her utility is 0. The seller S's payoff is p if he fulfills an order and 0 otherwise.[27]

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each player $i = B, S$.[28] Call the game with such strategy spaces the *finite-horizon ordering game*. It is characterized by $(\lambda, T, \sigma, p, v)$. The analysis in sections 3 and 4

[27] It is assumed without loss of generality that S's cost of producing and supplying the good is zero. The equilibrium strategies would not change if S were to face a constant cost strictly less than the price.

[28] Formal definitions of histories and strategy spaces are provided in appendix B.3.1.

implies that a subgame-perfect equilibrium is well defined.[29]

Proposition 4. *In the finite-horizon ordering game with $(\lambda, T, \sigma, p, v)$, there exists a unique subgame-perfect equilibrium. In this subgame-perfect equilibrium, there exists $t^* \in [0, T)$ such that B places an order at time t if and only if she has not done so yet and $t \geq t^*$, and S sells the good at time t if and only if he has not done so yet, an order has already been placed, and there is Poisson hit at time t.*

The result implies that on the equilibrium path of the unique subgame-perfect equilibrium, B places an order at time t^* and S sells the good at the first Poisson hit after time t^*.[30]

Remark 3.

1. (Violation of inertia) S's unique equilibrium strategy does not satisfy inertia. For any $\epsilon > 0$, there is positive probability of a Poisson hit in the time interval $(t^*, t^* + \epsilon)$, in which case S sells the good. Thus, there is no $\epsilon > 0$ such that S does not move in the time interval $(t^*, t^* + \epsilon)$.

2. (Non-z action at a time without a Poisson hit) B's unique equilibrium strategy would not satisfy a condition requiring that a non-z action be taken only at the times of discrete state changes. There is probability zero that the state discretely changes at time t^*, where the set of times at which the state discretely changes is defined as the set of Poisson arrival times.

3. (Relationship to admissibility) Assumptions F1 and F2 in Simon and Stinchcombe (1989) are satisfied by any feasible strategies. Since each agent can move only once, the number of moves is uniformly bounded, and the path of play is piecewise continuous. Moreover, the equilibrium strategies have property F3. Because the action of the buyer is the same at any time until an order has been placed and the action of the seller after an order has been placed does not depend on the specific timing of the move by the buyer, equilibrium behavior is strongly continuous in histories as required.

[29] In our general framework, the shock level, action profile, and instantaneous utility are not defined at the terminal time T. However, our description of the finite-horizon ordering game suggests that B receives a payoff at the deadline that depends on action profiles before the deadline as well as the shock level at the deadline. As detailed in appendix B.3.1, the game here can be reinterpreted so as to conform with our general framework.

[30] Note that the traceability restriction would be violated if B were to follow a strategy of placing an order if and only if she has not done so yet and $t > t^*$.

4. (Comparative statics) The existence of a unique equilibrium facilitates the derivation of comparative statics.[31] We focus on the case where $t^* > 0$. First, t^* is increasing in σ. Intuitively, as her taste becomes more volatile, B wants to wait longer to better match the specification to her taste. Second, t^* is decreasing in v. The reason is that as her valuation for the product becomes higher, B increasingly wants S to have an opportunity to fulfill her order. Third, t^* is increasing in p since the effect of p is opposite to that of v. Finally, t^* is increasing in λ. As opportunities to provide the product arrive more frequently to S, B wants to wait longer so as to set the specification closer to the realization of her taste at time T.

6 Payoff Assignment with Nonmeasurable Behavior

In section 4, we observed that the expected payoffs are not well defined when a strategy profile generates nonmeasurable behavior. There are two approaches to resolving this problem. The first method is to restrict the strategy space. The calculability assumption considered in section 4 is an example of the first approach and can be justified as being the most inclusive restriction in a certain sense. The second method is to assign an expected payoff to nonmeasurable behavior. This section is devoted to an exploration of the second approach.

We begin by specifying how to assign expected payoffs to nonmeasurable behavior (section 6.1). Then we point out a few properties of this methodology that may be problematic (section 6.2). Ultimately, the disadvantage of the second approach is the lack of clarity about what expected payoff to assign to nonmeasurable behavior. The choice of an expected payoff is necessarily arbitrary because the conditional expectation may be undefined when the behavior of the agents is nonmeasurable, and such arbitrariness may be problematic because the set of equilibria depends on the choice of an expected payoff. Hence, there is not an obvious way to argue that one choice of expected payoff is better than another.[32] With some restrictions, we show that any behavior on the path of play can be sustained in an SPE under a

[31] Proofs of these comparative statics results are provided in the appendix.

[32] For a particular class of games, some payoff assignments might be more reasonable than others. For example, an agent may be able to secure a certain payoff by choosing a specific action. In such a case, this payoff could be assigned to nonmeasurable behavior. More generally, nonmeasurable behavior might be assigned the minmax payoff. However, this approach suffers from circularity because the minmax payoff may not be defined until the expected payoffs have already been defined.

particular assignment of expected payoffs (section 6.3). Given the absence of an ideal assignment, the last parts of this section examine when the second approach would provide the same solution as the first approach. We first show that any SPE under the calculability assumption is an SPE for a certain way of assigning payoffs (section 6.4). Then we prove under some restrictions that the converse is also true (section 6.5).

6.1 Formulation of Expected Payoffs

We consider the assignment of expected payoffs to strategy profiles that induce nonmeasurable behavior.[33] For each agent $i \in I$, define a function $\chi_i : H \times \bar{\Pi}_i^{TF} \to \mathbb{R} \cup \{-\infty, \infty\}$.

Choose any strategy profile $\pi \in \times_{i \in I} \bar{\Pi}_i^{TF}$. Let $k_u = \big(\{g_t\}_{t \in [0,u]}, \{(b_t^i)_{i \in I}\}_{t \in [0,u)}\big)$ be any history up to time u, and denote $b = \{(b_t^i)_{i \in I}\}_{t \in [0,u)}$. If the process $\xi_b^i(\pi)$ is progressively measurable for each $i \in I$, then the expected payoff to agent i at k_u is given by $U_i(k_u, \pi) = V_i(k_u, \pi)$, where $V_i(k_u, \pi)$ is as specified in equation (1). Otherwise, the expected payoff to agent i at k_u is given by $U_i(k_u, \pi) = \chi_i(k_u, \pi)$.[34]

Given a strategy space $\times_{i \in I} \hat{\Pi}_i \in \times_{i \in I} \bar{\Pi}_i^{TF}$, we say that $\pi \in \times_{i \in I} \hat{\Pi}_i$ is a **subgame-perfect equilibrium** if for any history k_u up to time u, the expected payoff to agent $i \in I$ at k_u satisfies $U_i(k_u, \pi) \geq U_i[k_u, (\pi_i', \pi_{-i})]$ for any $\pi_i' \in \hat{\Pi}_i$. Since the assignment of expected payoffs to nonmeasurable behavior is not based on an extensive form, the standard one-shot deviation principle may not hold. Thus, it is crucial for the definition of SPE to consider deviations to a strategy in the entire subgame. The following example illustrates.

Example 4 (Deviation to Measurable Behavior)**.** Let $\{s_t\}_{t \in [0,T)}$ be an arbitrary stochastic process with state space S. Assume that there exists $\tilde{S} \subseteq S$ along with

[33] We assume traceability and frictionality, so that the behavior of the agents is well defined in the sense that there exists a unique action path consistent with each strategy profile. However, there is little difference between assigning payoffs to nondefined and nonmeasurable behavior. In order to define equilibria, every element in a given set of strategy profiles should be mapped to an expected payoff, and nondefined behavior like nonmeasurability would ordinarily preclude such a mapping. The results here can be extended to allow for strategy profiles that induce zero or multiple action paths, but these additional results are not stated so as to simplify the exposition.

[34] The extensive form of the game, which assigns a payoff profile to each deterministic history, is well defined. However, if a profile of traceable and frictional strategies induces nonmeasurable behavior, then there is no standard method to compute the expected payoff, even though the strategy profile is associated with a unique payoff under each realization of the shock process. In this sense, the assignment of expected payoffs is not based on an extensive form.

$\tilde{t} > 0$ such that $\{\omega \in \Omega : s_{\tilde{t}}(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$.[35] Suppose $I = \{1\}$ and that $\bar{A}_1(h_t) = \{x, z\}$ for every $h_t \in H$. The utility function satisfies $v_1(x, s) = 0$ for all $s \in S$. Let $\chi_1(h_t, \pi_1) = -1$ for all $h_t \in H$ and any $\pi_1 \in \bar{\Pi}_1^{TF}$. A strategy in which agent 1 chooses action x at time $\tilde{t}$ if and only if $s_{\tilde{t}}$ is in $\tilde{S}$ is not optimal at the null history because agent 1 can deviate to a strategy that induces measurable behavior. $\square$

In what follows, we will compare SPE under the calculability restriction with SPE under a payoff assignment. Whenever there is ambiguity about the strategy space or payoff assignment being considered, we identify the problem as $\Gamma\big(\hat{\Pi}, (\chi_i)_{i \in I}\big)$, where $\hat{\Pi} \subseteq \times_{i \in I} \bar{\Pi}_i^{TF}$ is the space of strategy profiles in consideration. For example, our analysis in the main sections corresponds to considering the game with $\Gamma\left(\times_{i \in I} \bar{\Pi}_i^{C}, (\chi_i)_{i \in I}\right)$. Since $U_i(h_t, \pi) = V_i(h_t, \pi)$ for all $i \in I$ whenever $\pi \in \times_{i \in I} \bar{\Pi}_i^{C}$, the specification of $(\chi_i)_{i \in I}$ is irrelevant in this case, so that we denote $\Gamma\left(\times_{i \in I} \bar{\Pi}_i^{C}, (\chi_i)_{i \in I}\right)$ by $\Gamma\left(\times_{i \in I} \bar{\Pi}_i^{C}\right)$.

6.2 Problems with Payoff Assignment

Assigning a payoff to nonmeasurable behavior may result in a model with objectionable properties. We first observe that assigning expected payoffs to nonmeasurable behavior may lead to a non-monotonic relationship between expected and realized payoffs.

Example 5 (Non-Monotonicity of Expected Payoffs in Realized Payoffs)**.** Suppose $I = \{1\}$. Let $\bar{A}_1(h_t) = \{x, z\}$ if $t = 1$, and let $\bar{A}_1(h_t) = \{z\}$ otherwise. Let $\{s_t\}_{t \in [0,T)}$ be an arbitrary stochastic process with state space S. Assume that there exists $\tilde{S} \subseteq S$ such that $\{\omega \in \Omega : s_1(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$. The utility function satisfies $v_1(x, s) = 1$ for all $s \in S$. Let $\chi_1(h_t, \pi_1) = -1$ for all $h_t \in H$ and any $\pi_1 \in \bar{\Pi}_1^{TF}$.

Consider a class of strategies, each of which is indexed by a set $C \subseteq S$, where π^C prescribes action z at any time $t \neq 1$ and action x at time $t = 1$ if and only if s_1 is in C. It may be natural for the expected payoffs to satisfy the following monotonicity condition: $U_1(h_0, \pi^{S''}) \leq U_1(h_0, \pi^{S'})$ if $S'' \subseteq S'$. That is, the expected payoff is monotonic in the realized payoffs in the sense of statewise dominance. However,

[35] For example, let $\{s_t\}_{t \in [0,T)}$ be a standard Brownian motion, and suppose that the set consisting of every continuous function $c : [0, T) \to \mathbb{R}$ with $c(0) = 0$ and $c(\tilde{t}) \in \tilde{S}$ is not Wiener measurable.

$U_1(h_0, \pi^{\emptyset}) = 0 > -1 = U(h_0, \pi^{\tilde{S}})$ even though $\emptyset \subseteq \tilde{S}$. Hence, the monotonicity condition fails. $\square$

As shown by the example below, the specific assignment of expected payoffs to nonmeasurable behavior affects the set of payoffs that can be supported in an SPE.

Example 6 (Dependence of Equilibrium Set on Payoff Assignment)**.** Suppose $I = \{1, 2\}$. For each $i \in \{1, 2\}$, let $\bar{A}_i(h_t) = \{x, z\}$ if $t = i - 1$, and let $\bar{A}_i(h_t) = \{z\}$ otherwise. Let $\{s_t\}_{t\in[0,T)}$ be an arbitrary stochastic process with state space S. Assume that there exists $\tilde{S} \subseteq S$ such that $\{\omega \in \Omega : s_1(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$. The utility functions satisfy $v_1[(x, z), s] = 1$, $v_1[(z, x), s] = 0$, and $v_2[(x, z), s] = v_2[(z, x), s] = 0$ for all $s \in S$.

First, suppose that $\chi_1(h_t, \pi) = -1$ and $\chi_2(h_t, \pi) = 0$ for all $h_t \in H$ and any $\pi \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$. Then there exists an SPE in which agent 1 receives an expected payoff of 0. For example, agent 2 may use a strategy of choosing action x at time 1 if and only if agent 1 chooses action x at time 0 and s_1 is in $\tilde{S}$.

Second, suppose that $\chi_1(h_t, \pi) = \chi_2(h_t, \pi) = -1$ for all $h_t \in H$ and any $\pi \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$. Then there does not exist an SPE in which agent 1 receives an expected payoff of 0. The reason is that there is no history h_1 up to time 1 for which agent 2 has an incentive to choose a strategy π_2 such that $U_2[h_1, (\pi_1, \pi_2)] = \chi_2[h_1, (\pi_1, \pi_2)]$ for some strategy π_1 of agent 1. Hence, it is always optimal for agent 1 to choose x at time 0, so that agent 1 receives an expected payoff of 1. $\square$

The general problem is that when the agents' behavior is nonmeasurable, there is not a well defined probability distribution over future paths of play. Hence, the expected payoff assigned by the function χ_i to a strategy profile involving nonmeasurable behavior does not have any natural relationship with the realized payoffs at future times as determined by the function v_i. Despite such a problem, the usual concept of SPE is well defined.

6.3 Arbitrary Behavior in Equilibrium

Now we examine the implications for agents' incentives of assigning payoffs to nonmeasurable behavior. Given any history k_u up to time u, let $\bar{H}(k_u)$ denote the set consisting of every history $h = \{s_t, (a_t^j)_{j\in I}\}_{t\in[0,T)}$ such that $\big(\{s_t\}_{t\in[0,u]}, \{(a_t^j)_{j\in I}\}_{t\in[0,u)}\big) = k_u$, $\{a_t^i\}_{t\in[0,T)} \in \Xi_i(u)$ for each $i \in I$, and $a_\tau^i \in \bar{A}_i\big(\{s_t\}_{t\in[0,\tau]}, \{(a_t^j)_{j\in I}\}_{t\in[0,\tau)}\big)$ for all $\tau \geq u$ and each $i \in I$. That is, $\bar{H}(k_u)$ includes every feasible history with finitely many moves in a finite time interval for which k_u is a subhistory. Let $\zeta_i(k_u) = \inf_{h \in \bar{H}(k_u)} \sum_{\tau \in M_u(h)} v_i[(a^j_\tau)_{j \in I}, s_\tau]$ denote the greatest lower bound on the feasible payoffs to agent i at k_u.[36] In this section, we consider the possibility of letting $\chi_i(k_u, \pi) \leq \zeta_i(k_u)$ for all $\pi \in \times_{j \in I} \bar{\Pi}^{TF}_j$ and each $i \in I$. That is, nonmeasurable behavior is assigned an expected payoff no greater than the infimum of the set of feasible payoffs.[37]

Below we demonstrate that such a payoff assignment may have a significant impact on the set of SPE. Intuitively, if an extremely low payoff is supportable in an SPE, then it may be used to severely punish deviations. In fact, we can prove a type of "folk theorem" under certain conditions. To illustrate this point, we first describe how both agents using nonmeasurable behavior can be supported as an SPE.

Example 7 (Mutual Monmeasurability)**.** Suppose $I = \{1, 2\}$. For each $i \in \{1, 2\}$, let $\bar{A}_i(h_t) = \{x, z\}$ if $t = 1$, and let $\bar{A}_i(h_t) = \{z\}$ otherwise. Let $\{s_t\}_{t \in [0,T)}$ be an arbitrary stochastic process with state space S. Assume that there exists $\tilde{S} \subseteq S$ such that $\{\omega \in \Omega : s_1(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$. The utility function of each agent $i \in \{1, 2\}$ satisfies $v_i[(a_1, a_2), s] \geq 0$ for all $s \in S$ and any $(a_1, a_2) \in A_1 \times A_2$ such that $(a_1, a_2) \neq (z, z)$. For each $i \in \{1, 2\}$, let $\chi_i(h_t, \pi') = \chi_i(h_t, \pi'') \leq 0$ for all $h_t \in H$ and any $\pi', \pi'' \in \bar{\Pi}^{TF}_1 \times \bar{\Pi}^{TF}_2$.

Let $\tilde{\pi}$ be a strategy profile in which each agent chooses x at time 1 if and only if s_1 is in $\tilde{S}$. This strategy profile is an SPE because at any history h_t up to a time $t \leq 1$, there is no unilateral deviation that would enable an agent to obtain an expected payoff greater than $\chi_i(h_t, \tilde{\pi}) \leq 0$. $\square$

We next show how this behavior may be used as a punishment to support other paths of play.

Example 8 (Folk Theorem)**.** Suppose $I = \{1, 2\}$. For each $i \in \{1, 2\}$, let $\bar{A}_i(h_0) = A_i$, $\bar{A}_i(h_t) = \{x, z\}$ if $t = 1$, and $\bar{A}_i(h_t) = \{z\}$ if $t \notin \{0, 1\}$. Let $\{s_t\}_{t \in [0,T)}$ be an arbitrary stochastic process with state space S. Assume that there exists $\tilde{S} \subseteq$

[36] Recall that $M_u(h)$ denotes the set of times at and after u where some agent moves under history h.

[37] Similar results hold under an alternative definition in which $\zeta_i(k_u) = \inf_{\pi \in \bar{\Pi}^{TF}_b} V_i(k_u, \pi)$ at any history $k_u = (\{g_t\}_{t \in [0,u]}, \{(b^j_t)_{j \in I}\}_{t \in [0,u)})$ up to time u, where $\bar{\Pi}^{TF}_b$ denotes the set consisting of any strategy profile $\pi \in \times_{j \in I} \bar{\Pi}^{TF}_j$ such that the process $\xi^i_b(\pi)$ with $b = \{(b^j_t)_{j \in I}\}_{t \in [0,u)}$ is progressively measurable for all $i \in I$.

S such that $\{\omega \in \Omega : s_1(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$. The utility functions satisfy $v_1[(a_1, a_2), s] \geq v_1[(z, a_2), s] \geq 0$ and $v_2[(a_1, a_2), s] \geq v_2[(a_1, z), s] \geq 0$ for all $s \in S$ and any $(a_1, a_2) \in A_1 \times A_2$ such that $(a_1, a_2) \neq (z, z)$. For each $i \in \{1, 2\}$, let $\chi_i(h_t, \pi') = \chi_i(h_t, \pi'') \leq 0$ for all $h_t \in H$ and any $\pi', \pi'' \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$.

Choose any pair of actions $(\tilde{a}_1, \tilde{a}_2) \in A_1 \times A_2$. Let $\tilde{\pi}$ be the strategy profile defined as follows. Each agent $i \in \{1, 2\}$ chooses $\tilde{a}_i$ at time 0. If every agent i takes action $\tilde{a}_i$ at time 0, then the agents choose x at time 1. If some agent i does not take action $\tilde{a}_i$ at time 0, then the agents choose x at time 1 if and only if s_1 is in $\tilde{S}$. This strategy profile is an SPE. A unilateral deviation at h_0 would result in an expected payoff of $\chi_i(h_0, \tilde{\pi}) \leq 0$ to agent i. If strategy profile $\tilde{\pi}$ is followed at time 0, then playing x at time 1 is a best response for each agent to the action of the other agent. If strategy profile $\tilde{\pi}$ is not followed at time 0, then neither player has an incentive to deviate again for the same reason as in example 7. $\square$

The preceding example illustrates how any profile of actions at the null history can be implemented in equilibrium.[38] Now we identify general conditions under which arbitrary behavior after the null history can also be supported in equilibrium by suitably assigning payoffs to nonmeasurable behavior.

Proposition 5. *Let $n \geq 2$ and $T = \infty$. Consider the game $\Gamma\left(\times_{j\in I}\bar{\Pi}_j^{TF}, (\chi_j)_{j\in I}\right)$ where $\chi_i(h_t, \hat{\pi}) = \chi_i(h_t, \bar{\pi}) \leq \zeta_i(h_t)$ for all $\hat{\pi}, \bar{\pi} \in \times_{j\in I}\bar{\Pi}_j^{TF}$, any $h_t \in H$, and every $i \in I$. Assume that for each $i \in I$ and any $h_t \in H$, there exists $\tilde{a} \in \bar{A}_i(h_t)$ such that $\tilde{a} \neq z$. Suppose that there exists $\tilde{t} > 0$ along with a collection of sets $\{\tilde{S}_t\}_{t\in[0,\tilde{t}]}$ such that $\{\omega \in \Omega : s_t(\omega) \in \tilde{S}_t,\ \forall\ t \in [0, \tilde{t}]\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$. Choose any $\pi \in \times_{i\in I}\bar{\Pi}_i^{TF}$ such that for any profile of action paths $\left(\{b_t^i\}_{t\in[0,u)}\right)_{i\in I}$ up to an arbitrary time $u > 0$, there exists with probability one some $t < u$ such that $\pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(b_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right) \neq b_t^i$ for some $i \in I$. Then there exists an SPE $\pi' \in \times_{i\in I}\bar{\Pi}_i^{TF}$ such that $\left(\{\phi_t^i(h_0, \{s_\tau\}_{\tau\in(0,T)}, \pi')\}_{t\in[0,T)}\right)_{i\in I} = \left(\{\phi_t^i(h_0, \{s_\tau\}_{\tau\in(0,T)}, \pi)\}_{t\in[0,T)}\right)_{i\in I}$ with probability one.*

[38]In our model, actions are assumed to be perfectly observable. Hence, given a strategy profile in which behavior on the path of play is measurable, assigning an extremely low payoff to nonmeasurable behavior at information sets that can be reached only after a deviation does not affect expected payoffs at the null history. Bonatti, Cisternas, and Toikka (2017) consider a related approach in which a payoff of negative infinity is assigned to strategy profiles with undesirable properties. However, their model assumes imperfect monitoring, so that punishment using an infinitely negative payoff results in an infinitely negative payoff at the null history.

The result is proved by constructing a strategy profile π' with the following properties. At the null history h_0, agent i has no incentive to deviate because doing so would result in an expected payoff of $\chi_i(h_0, \pi)$, which is no more than the infimal feasible payoff $\zeta_i(h_0)$ based on the utility function. At any history k_u up to a positive time u, the action path b up to time u is such that the process $\xi_b^i(\pi')$ is not progressively measurable given the restriction on strategy profile π in the statement of the proposition as well as the assumption that $\xi_b^i(\pi')$ records the continuation path of play when the action path up to time u is fixed at b. Hence, the expected payoff of agent i at history k_u is constant at $\chi_i(k_u, \pi')$ when the other agents play π', so that agent i has no incentive to deviate.[39]

Note that simply requiring progressive measurability of the shock process does not enable an arbitrary path of play to be implemented as an equilibrium. To see this, suppose that the shock can take values only in a finite set and can change values only at discrete times. Then the shock process would be progressively measurable, but the finiteness of the state space and the discrete timing of state changes make it impossible for a strategy profile to induce nonmeasurable behavior.

6.4 From Calculability Restriction to Payoff Assignment

We examine the relationship between the SPE under the calculability restriction and the SPE when payoffs are assigned to nonmeasurable behavior. As in section 6.3, we associate nonmeasurable behavior with an expected payoff no more than the greatest lower bound on the feasible payoffs. The following result shows that the set of SPE in this case is at least as large as the set of SPE under the calculability restriction.

Proposition 6. *Let $\bar{A}_i(h_t) = A_i$ for every $h_t \in H$ and each $i \in I$. If π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^C\right)$, then π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$ with $\chi_i(h_t, \tilde{\pi}) \leq \zeta_i(h_t)$ for all $\tilde{\pi} \in \times_{j\in I}\bar{\Pi}_j^{TF}$, any $h_t \in H$, and every $i \in I$.*

Therefore, if the model has an SPE under the calculability restriction, then there exists an SPE under the approach where nonmeasurable behavior is assigned an

[39] We assume only for simplicity when defining $\xi_b^i(\pi')$ that the action path up to time u is fixed irrespective of the shock realization up to time u. The reasoning in the proof does not entirely apply under the alternative definition in footnote 13 where behavior up to time u may be determined by π'. However, it can still be shown that at any history up to a given time that is reached with probability one when playing π, π' specifies the same action profile as π, and π'_i designates a best response to π'_{-i}. As mentioned in footnote 13, the main results in section 4 are valid under both definitions of the action process.

expected payoff no greater than the infimal feasible payoff, even when we do not associate each instance of nonmeasurable behavior with the same expected payoff. This result implies that insofar as the concept of SPE is concerned, the calculability restriction is not picking up strategy profiles that would be ruled out by every assignment of expected payoffs to nonmeasurable behavior.

6.5 From Payoff Assignment to Calculability Restriction

Here we identify when an SPE under the method of assigning payoffs to nonmeasurable behavior is also an SPE under the calculability restriction.

We begin by defining a set of strategy profiles with certain measurability properties. Let the random variable $\theta : \Omega \to [0, T]$ be a stopping time.[40] Given any $\pi', \pi'' \in \times_{i \in I} \Pi_i$, let $\psi(\pi', \pi'', \theta)$ be the strategy profile satisfying the following two properties for each $i \in I$:

1. $\psi_i(\pi', \pi'', \theta)\big(\{s_t(\omega)\}_{t \in [0,u]}, \{(b_t^j)_{j \in I}\}_{t \in [0,u)}\big) = \pi_i'\big(\{s_t(\omega)\}_{t \in [0,u]}, \{(b_t^j)_{j \in I}\}_{t \in [0,u)}\big)$ for each $i \in I$, all $u \geq 0$, every $\{(b_t^j)_{j \in I}\}_{t \in [0,u)}$, and any $\omega \in \Omega$ with $u \leq \theta(\omega)$;
2. $\psi_i(\pi', \pi'', \theta)\big(\{s_t(\omega)\}_{t \in [0,u]}, \{(b_t^j)_{j \in I}\}_{t \in [0,u)}\big) = \pi_i''\big(\{s_t(\omega)\}_{t \in [0,u]}, \{(b_t^j)_{j \in I}\}_{t \in [0,u)}\big)$ for each $i \in I$, all $u \geq 0$, every $\{(b_t^j)_{j \in I}\}_{t \in [0,u)}$, and any $\omega \in \Omega$ with $u > \theta(\omega)$.

In other words, the strategy $\psi_i(\pi', \pi'', \theta)$ plays π_i' until and including the stopping time, and plays π_i'' thereafter. A strategy profile $\pi \in \times_{i \in I} \Pi_i^{TF}$ is said to be **measurably attachable** if for each action path b, every stopping time θ, and any $\tilde{\pi} \in \times_{i \in I} \Pi_i^{TF}$ such that $\xi_b^i(\tilde{\pi})$ is progressively measurable for all $i \in I$, the strategy $\psi_i(\tilde{\pi}, \pi, \theta)$ is traceable and frictional for all $i \in I$ and the process $\xi_b^i[\psi(\tilde{\pi}, \pi, \theta)]$ is progressively measurable for all $i \in I$. That is, π is required to induce progressively measurable behavior after any progressively measurable behavior up to and including an arbitrary random time. Let $\Pi^A \subseteq \times_{i \in I} \Pi_i^{TF}$ be the set of measurably attachable strategy profiles. In addition, a strategy profile $\pi \in \times_{i \in I} \Pi_i$ is said to be **synchronous** if for any $h_t \in H$, $\pi_j(h_t) = z$ for all $j \in I$ whenever $\pi_i(h_t) = z$ for some $i \in I$. That is, π requires the agents to move at the same time as each other.[41]

[40] That is, it satisfies $\{\omega \in \Omega : \theta(\omega) \leq t\} \in \mathcal{F}_t$ for all $t \in [0, T)$.

[41] The synchronicity assumption is satisfied by the maximal equilibrium of the tree harvesting problem and of the sequential exchange model and technology adoption game in the online appendix as well as by a Markov perfect equilibrium of the inventory restocking application in the online appendix. In addition, any asynchronous strategy profile can be expressed as a synchronous strategy profile by adding a payoff irrelevant action to the action space of each agent and requiring this action

According to the following result, any synchronous and measurably attachable strategy profile that is an SPE when payoffs are assigned to nonmeasurable behavior is also an SPE under the calculability restriction.

Theorem 3. *If the synchronous strategy profile* $\pi \in \Pi^A$ *is an SPE of* $\Gamma\big(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\big)$*, then* π *is an SPE of* $\Gamma\big(\times_{i\in I}\bar{\Pi}_i^{C}\big)$*.*

The theorem implies that the restriction to calculable strategies does not exclude from the set of SPE any synchronous and measurably attachable strategy profile that is supported as an SPE under some assignment of expected payoffs to nonmeasurable behavior.

To prove this result, we first show that any profile of calculable strategies that is an SPE under an assignment of payoffs to nonmeasurable behavior is also an SPE under the calculability restriction. Intuitively, when the other agents are playing calculable strategies, a deviation by an agent from one calculable strategy to another calculable strategy produces the same change in expected payoffs under the calculability restriction as under the payoff assignment method.

We then confirm that any synchronous and measurably attachable strategy profile π is also a profile of calculable strategies. This part of the proof involves an iterative procedure as in the proof of the main theorem stating that calculable strategies generate a measurable path of play. Specifically, we let π'_{-i} be a profile of quantitative strategies for the agents other than i. If the agents are playing (π_i, π'_{-i}), then the behavior induced by (π_i, π'_{-i}) is progressively measurable up to and including the first time that π or (π_i, π'_{-i}) prescribes a move. Because π induces progressively measurable behavior after any progressively measurable behavior up to and including an arbitrary random time, the behavior induced by (π_i, π'_{-i}) is progressively measurable up to and including the next time that π or (π_i, π'_{-i}) prescribes a move. We can apply this argument iteratively in order to show that the behavior induced by (π_i, π'_{-i}) is progressively measurable. This sort of reasoning establishes that π_i is calculable.

The restriction to synchronous SPE ensures that the aforesaid iterative procedure characterizes the path of play of (π_i, π'_{-i}) over the entire course of the game. To see this, choose any time $t \in [0, T)$. By the traceability and frictionality assumptions, strategy profile (π_i, π'_{-i}) with probability one induces only a finite number of moves

to be chosen by an agent that does not move when another agent moves. All the equilibria studied in section 5 and the online appendix satisfy synchronicity after such a reformulation.

before time t. Moreover, the synchronicity assumption implies that the procedure is such that π prescribes a move only if (π_i, π'_{-i}) does so. Hence, the iterations with probability one reach time t after only finitely many steps.

Without the assumption that π is measurably attachable, the proposition fails. The following example illustrates.

Example 9 (Role of Measurable Attachability)**.** Suppose $I = \{1, 2\}$. For each $i \in \{1, 2\}$, let $\bar{A}_i(h_t) = \{w, z\}$ if $t = 1$, $\bar{A}_i(h_t) = \{x, z\}$ if $t = 2$, and let $\bar{A}_i(h_t) = \{z\}$ otherwise. Let $\{s_t\}_{t\in[0,T)}$ be an arbitrary stochastic process with state space S. Assume that there exists $\tilde{S} \subseteq S$ such that $\{\omega \in \Omega : s_2(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$. For all $s \in S$, the utility function of each agent $i \in \{1, 2\}$ satisfies $v_i[(x, x), s] = 1$ and $v_i[(a_1, a_2), s] = 0$ for any $(a_1, a_2) \in A_1 \times A_2$ such that $(a_1, a_2) \neq (x, x)$. Let $\chi_i(h_t, \pi) = 0$ for each $i \in I$, all $h_t \in H$, and any $\pi \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$.

Let π^* be a strategy profile in which both agents choose z at time 1 and choose z at time 2 if and only if some agent $i \in \{1, 2\}$ chooses w at time 1 and s_2 is in $\tilde{S}$. First, π^* is not measurably attachable since the process $\xi^i_{\{\}}[\psi(\tilde{\pi}, \pi^*, \tilde{\theta})]$ is not progressively measurable, where $\tilde{\pi}$ is a strategy profile in which both agents always choose w at time 1 and choose z at time 2, and the stopping time $\tilde{\theta}$ is equal to the constant 1. Second, π^* is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$ because for all $i \in I$ and $\pi'_i \in \bar{\Pi}_i^{TF}$, we have $U_i(h_t, \pi^*) = V_i(h_t, \pi^*) = 1 \geq U_i[h_t, (\pi'_i, \pi^*_{-i})]$ for every history h_t up to time 1 and every history h_t up to time 2 where z is chosen by both agents at time 1 and because for all $i \in I$ and $\pi'_i \in \bar{\Pi}_i^{TF}$, we have $U_i(h_t, \pi^*) = \chi_i(h_t, \pi^*) = 0 = \chi_i[h_t, (\pi'_i, \pi^*_{-i})] = U_i[h_t, (\pi'_i, \pi^*_{-i})]$ for every history h_t up to time 2 where w is chosen by some agent at time 1. Third, π^*_i is not calculable since the process $\xi^i_{\{\}}(\pi^*_i, \tilde{\pi}_{-i})$ is not progressively measurable, where $\tilde{\pi}_{-i}$ is the quantitative strategy of always choosing w at time 1 and choosing z at time 2. $\square$

Note that measurable attachability is not a restriction on the strategy space of an individual agent but on the space of strategy profiles. Hence, Π^A does not necessarily have a product structure. Although we find it unsatisfactory, one could define a notion of SPE under the restriction of measurable attachability, where the set of strategies to which an agent can deviate depends on the strategy profile of its opponents: We say that $\pi \in \Pi^A \cap \bar{\Pi}$ is a **pseudo-SPE** of $\Gamma(\Pi^A \cap \bar{\Pi})$ if $V_i(k_u, \pi) \geq V_i[k_u, (\pi'_i, \pi_{-i})]$ for each $i \in I$, any history $k_u = \left(\{g_t\}_{t\in[0,u]}, \{(b^j_t)_{j\in I}\}_{t\in[0,u)}\right)$ up to an arbitrary time u, and every

$\pi_i' \in \bar{\Pi}_i^{TF}$ for which there exists $\pi'' \in \Pi^A$ such that $\xi_b^e(\pi'')$ with $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$ is the same stochastic process as $\xi_b^e[(\pi_i', \pi_{-i})]$ for all $e \in I$.

The proposition below states that if the synchronous strategy profile π is a pseudo-SPE under the restriction to measurably attachable strategy profiles, then π is an SPE under the calculability restriction.

Proposition 7. *If the synchronous strategy profile $\pi \in \Pi^A$ is a pseudo-SPE of $\Gamma(\Pi^A \cap \bar{\Pi})$, then π is an SPE of $\Gamma\big(\times_{i\in I}\bar{\Pi}_i^C\big)$.*

To prove this result, we first show that if the synchronous strategy profile π is an SPE under the restriction to measurably attachable strategy profiles, then π is an SPE when nonmeasurable behavior is assigned an expected payoff no greater than the infimal feasible payoff. It then follows from the theorem in this section that π is an SPE under the calculability restriction.

7 Conclusion

This paper considers the problem of defining strategy spaces for continuous-time games in a stochastic environment. We showed through various examples that existing techniques developed for a deterministic environment do not necessarily apply because they prevent agents from quickly responding to changes in the state of nature. We introduced a new set of restrictions on individual strategy spaces that guarantee the existence of a unique action path as well as the measurability of the induced behavior. Specifically, traceability and frictionality ensure the former, and calculability ensures the latter. We compared our method to an alternative approach in which specific payoffs are assigned to strategies inducing nonmeasurable behavior, and found a certain equivalence between our method and this alternative.

Our methodology does not cover every possible situation. For example, it applies to timing games in which agents choose when to take actions, while it is not applicable to settings where agents continuously change their actions. We hope that future research will address such settings as well. In addition, we restricted attention to pure strategies. With mixed strategies, measurability problems may arise even in the absence of an exogenous shock because one agent's behavior may condition on the realization of another agent's action in a nonmeasurable way. Despite those limitations, we showed the relevance of our methodology to a wide range of problems, and we hope that it will be useful in future work on these applications and others for appropriately defining strategy spaces.

A Appendix

A.1 Proofs for Sections 3 and 4

Proof of Theorem 1. We divide the proof of the theorem into three steps. The first step shows uniqueness in part 1, the second shows part 2, and the third shows existence in part 1. Fix an arbitrary history $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\big)$ up to time u. Assume that the realization of the shock process $\{s_t\}_{t\in[0,T)}$ is such that $s_t = g_t$ for $t \in [0, u]$.

For $t \in [0, u)$, let $a_t^j = b_t^j$ for each $j \in I$. First, we show that there is probability one that $\{s_t\}_{t\in(u,T)}$ is such that there exists at most one profile $\{(a_t^j)_{j\in I}\}_{t\in[0,T)}$ of action paths for which the history $h = \big\{s_t, (a_t^j)_{j\in I}\big\}_{t\in[0,T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T)$. Second, we show that there is probability one that $\{s_t\}_{t\in(u,T)}$ is such that $\{a_t^j\}_{t\in[0,T)} \in \Xi_j(u)$ for each $j \in I$ if the history $h = \big\{s_t, (a_t^j)_{j\in I}\big\}_{t\in[0,T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T)$. Third, we show that there is probability one that $\{s_t\}_{t\in(u,T)}$ is such that there exists a profile $\{(a_t^j)_{j\in I}\}_{t\in[0,T)}$ of action paths for which the history $h = \big\{s_t, (a_t^j)_{j\in I}\big\}_{t\in[0,T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T)$.

For the first step, we use the definition of frictionality. For $t \in [0, u)$ and $j \in I$, let $p_t^j = b_t^j$ and $q_t^j = b_t^j$. Suppose to the contrary that there is positive probability that $\{s_t\}_{t\in(u,T)}$ is such that there exist two distinct profiles $\{(p_t^j)_{j\in I}\}_{t\in[0,T)}$ and $\{(q_t^j)_{j\in I}\}_{t\in[0,T)}$ of action paths for which the histories $h^p = \big\{s_t, (p_t^j)_{j\in I}\big\}_{t\in[0,T)}$ and $h^q = \big\{s_t, (q_t^j)_{j\in I}\big\}_{t\in[0,T)}$ are both consistent with π_i for each $i \in I$ at every $t \in [u, T)$. The frictionality assumption implies that there is zero probability that $\{s_t\}_{t\in[u,T)}$ is such that $\{p_t^j\}_{t\in[u,T)}$ or $\{q_t^j\}_{t\in[u,T)}$ has infinitely many non-z actions in any finite interval of time for some $j \in I$. It follows that one can find a first time $v \geq u$ such that there exists $j \in I$ such that $\{p_t^j\}_{t\in[u,T)}$ is different from $\{q_t^j\}_{t\in[u,T)}$.[42] Let h_v^p be

[42] To see this, let X be the set of times at which $\{p_t^j\}_{t\in[u,T)}$ is different from $\{q_t^j\}_{t\in[u,T)}$ for some $j \in I$. If $\inf(X) = \infty$, it means there is no such time, which is a contradiction. Hence, $\inf(X) < \infty$. If the time $\inf(X)$ does not belong to the set X, then there exists $\epsilon > 0$ such that there are infinitely many times in X greater than $\inf(X)$ but less than $\inf(X) + \epsilon$. This implies that at least one of the $2n$ paths $\{(p_t^j)_{j\in I}\}_{t\in[u,T)}$, $\{(q_t^j)_{j\in I}\}_{t\in[u,T)}$ differs from z at infinitely many points in the time interval between $\inf(X)$ and $\inf(X) + \epsilon$. In this case, there exists $j \in I$ such that π_j violates the frictionality assumption.

the history up to time v when h^p is the history between times 0 and T, and let h_v^q be the history up to time v when h^q is the history between times 0 and T. By the definition of v, it must be that $h_v^p = h_v^q$. It follows that $\pi_i(h_v^p) = \pi_i(h_v^q)$ for each $i \in I$. Moreover, since both h^p and h^q are consistent with π_i for each $i \in I$ at every $t \in [u, T)$, it must be that $\pi_i(h_v^p) = p_v^i$ and $\pi_i(h_v^q) = q_v^i$ for each $i \in I$. Hence, we have $p_v^i = q_v^i$ for each $i \in I$, which contradicts the fact that v is the first time no less than u such that $\{p_t^j\}_{t\in[u,T)}$ is different from $\{q_t^j\}_{t\in[u,T)}$ for some $j \in I$.

The second step is straightforward. We again use the definition of frictionality. Suppose to the contrary that there is positive probability that $\{s_t\}_{t\in(u,T)}$ is such that there exists $\{(a_t^j)_{j\in I}\}_{t\in[0,T)}$ for which $\{a_t^j\}_{t\in[0,T)}$ is not in $\Xi_j(u)$ for some $j \in I$ and for which the history $h = \left\{s_t, (a_t^j)_{j\in I}\right\}_{t\in[0,T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T)$. In this case, π_j would clearly violate the frictionality assumption for some $j \in I$.

For the third step, we use the following construction that relies on an iterative argument. We apply the definitions of both traceability and frictionality. With probability one, $\{s_t\}_{t\in(u,T)}$ is such that the following algorithm can be applied. For $j \in I$, define the action path $\{a_t^{j,0}\}_{t\in[0,T)}$ so that $a_t^{j,0} = b_t^j$ for all $t \in [0, u)$ and $a_t^{j,0} = 0$ for all $t \in [u, T)$. Let $k = 0$.

1. If the history $\{s_t, (a_t^{j,k})_{j\in I}\}_{t\in[0,T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T)$, then we are finished. Otherwise, continue to Stage 2.

2. From traceability, for each $j \in I$, one can find an action path $\{d_t^{j,k}\}_{t\in[0,T)}$ with $d_t^{j,k} = b_t^j$ for $t \in [0, u)$ such that the history $\{s_t, (d_t^{j,k}, (a_t^{i,k})_{i\neq j})\}_{t\in[0,T)}$ is consistent with π_j at each $t \in [u, T)$. From frictionality, $\{d_t^{j,k}\}_{t\in[0,T)}$ can be treated as having only finitely many non-z actions in any finite interval of time that is a subset of $[u, T)$. It follows that one can find a first time $v_k \geq u$ such that $\{a_t^{j,k}\}_{t\in[u,T)}$ is different from $\{d_t^{j,k}\}_{t\in[u,\infty)}$ for some $j \in I$. For $j \in I$, define the action path $\{a_t^{j,k+1}\}_{t\in[0,T)}$ so that $a_t^{j,k+1} = a_t^{j,k}$ for $t \neq v_k$ and so that $a_t^{j,k+1} = d_t^{j,k+1}$ for $t = v_k$. Redefine k as $k + 1$. Return to Stage 1.

Consider the case where the preceding algorithm does not terminate after a finite number of iterations. By construction, v_k is increasing in k. Note also that for $j \in I$ along with any fixed value of t, there exists l such that $a_t^{j,k}$ is constant in k for $k > l$. Thus, $\lim_{k\to\infty} a_t^{j,k}$ is well defined for all t. Consider the history $h_T =$

$\lim_{k\to\infty}\{s_t, (a_t^{j,k})_{j\in I}\}_{t\in[0,T)}$. The history h_T is consistent with π_i for each $i \in I$ for $t \in [u, \lim_{k\to\infty} v_k)$.

Suppose that $\lim_{k\to\infty} v_k < T$. With probability one, $\{s_t\}_{t\in(u,T)}$ is such that the following argument can be applied. From traceability, one can find action paths $\{(d_t^{j,\infty})_{j\in I}\}_{t\in[0,\infty)}$ with $d_t^{j,\infty} = b_t^j$ for every $j \in I$ and all $t \in [0, u)$ such that for each $j \in I$, the history $\{s_t, (d_t^{j,\infty}, \lim_{k\to\infty}(f_t^{i,k})_{i\neq j})\}_{t\in[0,T)}$ is consistent with π_j at each $t \in [u, \infty)$. Moreover, it can be shown that $d_t^{j,\infty} = \lim_{k\to\infty} a_t^{j,k}$ for $j \in I$ and $t \in [u, \lim_{k\to\infty} v_k)$. Since the algorithm does not terminate after a finite number of iterations, there exists $j \in I$ such that $d_t^{j,\infty} \notin \Xi_j(u)$, which implies π_j violates the frictionality assumption for some $j \in I$. Thus, $\lim_{k\to\infty} v_k = T$.

Hence, the history h_T is such that (i) $s_t = g_t$ for all $t \in [0, u]$, (ii) $\lim_{k\to\infty} a_t^{j,k} = b_t^j$ for $j \in I$ and $t \in [0, u)$, and (iii) h_T is consistent with π_i for each $i \in I$ at every $t \in [u, T)$. □

Proof of Theorem 2. For $i \in I$, choose any $\pi_i \in \Pi_i^C$. Let $\{(b_\tau^j)_{j\in I}\}_{\tau\in[0,u)}$ be any path of actions by the agents up to an arbitrary time u. Given the realization of the shock $\{s_\tau\}_{\tau\in[0,u]}$ until time u, denote the history up to time u by $\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]}) = \big(\{s_\tau\}_{\tau\in[0,u]}, \{(b_\tau^j)_{j\in I}\}_{\tau\in[0,u)}\big)$.

It is helpful to define a set of strategies that depend only on the realization of the shock and not on the behavior of the agents. For $i \in I$, a strategy $\tilde{\pi}_i \in \Pi_i^{TF}$ is said to be individualistic if $\tilde{\pi}_i(h_t^p) = \tilde{\pi}_i(h_t^q)$ for any two histories $h_t^p = \big(\{g_\tau^p\}_{\tau\in[0,t]}, \{(p_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\big)$ and $h_t^q = \big(\{g_\tau^q\}_{\tau\in[0,t]}, \{(q_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\big)$ up to an arbitrary time t such that $\{g_\tau^p\}_{\tau\in[0,t]} = \{g_\tau^q\}_{\tau\in[0,t]}$.

Next some additional terminology regarding the action process is introduced. For any realization of the shock process $\{s_\tau\}_{\tau\in[0,T)}$, let $\Lambda(\{s_\tau\}_{\tau\in[0,T)})$ be an arbitrary subset of the interval $[u, T)$. Choose any $\pi' = (\pi_j')_{j\in I}$ and $\pi'' = (\pi_j'')_{j\in I}$ with $\pi_j', \pi_j'' \in \Pi_j^{TF}$ for $j \in I$. The strategy profiles π' and π'' are said to almost surely induce the same path of play by agent $i \in I$ for all $t \in \Lambda(\{s_\tau\}_{\tau\in[0,T)})$ if the following holds. There is probability one of the realization of the shock process $\{s_\tau\}_{\tau\in[0,T)}$ being such that $\phi_t^i[\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]}), \{s_\tau\}_{\tau\in(u,T)}, \pi'] = \phi_t^i[\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]}), \{s_\tau\}_{\tau\in(u,T)}, \pi'']$ for all $t \in \Lambda(\{c_\tau\}_{\tau\in[0,T)})$.

Now observe that the actions of each agent depend only on the realization of the shock if the agents play a fixed profile of traceable and frictional strategies. In particular, theorem 1 implies that there exists a profile $(\pi_j^*)_{j\in I}$ of individualistic strategies such that $\pi = (\pi_j)_{j\in I}$ and $\pi^* = (\pi_j^*)_{j\in I}$ almost surely induce the same path of play by every agent for all $t \in [u, T)$. We prove below that $\xi_b^i(\pi_i, \pi_{-i})$ is progressively measurable for $i \in I$.

For $i \in I$, let π_i^0 be the strategy that requires agent i to choose action z at every history, and note that $\pi_i^0 \in \Pi_i^Q$. Moreover, the stochastic process $\xi_b^i(\pi_i, \pi_{-i}^0)$ is progressively measurable for $i \in I$. Given any realization of the shock process $\{s_\tau\}_{\tau \in [0,T)}$, let $\Lambda^1(\{s_\tau\}_{\tau \in [0,T)})$ denote the set consisting of every time $t \in [u, T)$ for which there does not exist a time $\tilde{t} \in [u, t)$ such that $\phi_{\tilde{t}}^i[\tilde{k}_u(\{s_\tau\}_{\tau \in [0,u]}), \{s_\tau\}_{\tau \in (u,T)}, \pi] \neq z$ for some $i \in I$. That is, $\Lambda^1(\{s_\tau\}_{\tau \in [0,T)})$ represents the interval consisting of all times no less than u and no greater than the time that the first non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau \in [0,u]})$ would be chosen when playing π. Note that the strategy profiles (π_i, π_{-i}) and (π_i, π_{-i}^0) almost surely induce the same path of play by agent $i \in I$ for all $t \in \Lambda^1(\{s_\tau\}_{\tau \in [0,T)})$.

For $i \in I$, let π_i^1 be the strategy that requires agent i to follow π_i^* at any time $t \in \Lambda^1(\{s_\tau\}_{\tau \in [0,T)})$ and to choose action z at any time $t \notin \Lambda^1(\{s_\tau\}_{\tau \in [0,T)})$. That is, the strategy π_i^1 requires agent i to choose action z at each time less than u, to follow the strategy π_i^* at any time no less than u and no greater than the time that the first non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau \in [0,u]})$ would be chosen when playing π, and to choose action z thereafter. Note that for any $\pi'_{-i} \in \Pi_{-i}^{TF}$, the strategy profiles (π_i^1, π'_{-i}) and (π_i, π_{-i}^0) almost surely induce the same path of play by agent i for all $t \in \Lambda^1(\{s_\tau\}_{\tau \in [0,T)})$. Moreover, it was noted above that $\xi_b^i(\pi_i, \pi_{-i}^0)$ is progressively measurable. Hence, $\pi_i^1 \in \Pi_i^Q$ for $i \in I$.

For $i \in I$, the strategy π_i is such that $\xi_b^i(\pi_i, \pi_{-i}^1)$ is progressively measurable. Given any realization of the shock process $\{s_\tau\}_{\tau \in [0,T)}$, let $\Lambda^2(\{s_\tau\}_{\tau \in [0,T)})$ denote the set consisting of every time $t \in [u, T)$ for which there exists at most one time $\tilde{t} \in [u, t)$ such that $\phi_{\tilde{t}}^i[\tilde{k}_u(\{s_\tau\}_{\tau \in [0,u]}), \{s_\tau\}_{\tau \in (u,\infty)}, \pi] \neq z$ for some $i \in I$. That is, $\Lambda^2(\{s_\tau\}_{\tau \in [0,T)})$ represents the interval consisting of all times no less than u and no greater than the time that the second non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau \in [0,u]})$ would be chosen when playing π. Note that the strategy profiles (π_i, π_{-i}) and (π_i, π_{-i}^1) almost surely induce the same path of play by agent $i \in I$ for all $t \in \Lambda^2(\{s_\tau\}_{\tau \in [0,T)})$.

For $i \in I$, let π_i^2 be the strategy that requires agent i to follow π_i^* at any time $t \in \Lambda^2(\{s_\tau\}_{\tau \in [0,T)})$ and to choose action z at any time $t \notin \Lambda^2(\{s_\tau\}_{\tau \in [0,T)})$. That is, the strategy π_i^2 requires agent i to choose action z at each time less than u, to follow the strategy π_i^* at any time no less than u and no greater than the time that the second non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau \in [0,u]})$ would be chosen when playing π, and to choose action z thereafter. Note that for any $\pi'_{-i} \in \Pi^{TF}_{-i}$, the strategy profiles (π^2_i, π'_{-i}) and (π_i, π^1_{-i}) almost surely induce the same path of play by agent i for all $t \in \Lambda^2(\{s_\tau\}_{\tau\in[0,T)})$. Moreover, it was noted above that $\xi^i_b(\pi_i, \pi^1_{-i})$ is progressively measurable. Hence, $\pi^2_i \in \Pi^Q_i$ for $i \in I$.

Let k be an arbitrary positive integer. Given any realization of the shock process $\{s_\tau\}_{\tau\in[0,T)}$, let $\Lambda^k(\{s_\tau\}_{\tau\in[0,T)})$ denote the set consisting of every time $t \in [u, T)$ for which there exist at most $k-1$ values of $\tilde{t} \in [u, t)$ such that $\phi^i_{\tilde{t}}[\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]}), \{s_\tau\}_{\tau\in(u,T)}, \pi] \neq z$ for some $i \in I$. That is, $\Lambda^k(\{s_\tau\}_{\tau\in[0,T)})$ represents the interval consisting of all times no less than u and no greater than the time that the k^{th} non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]})$ would be chosen when playing π. For $i \in I$, let π^k_i be the strategy that requires agent i to follow π^*_i at any time $t \in \Lambda^k(\{s_\tau\}_{\tau\in[0,T)})$ and to choose action z at any time $t \notin \Lambda^k(\{s_\tau\}_{\tau\in[0,T)})$. That is, the strategy π^k_i requires agent i to choose action z at each time less than u, to follow the strategy π^*_i at any time no less than u and no greater than the time that the k^{th} non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]})$ would be chosen when playing π, and to choose action z thereafter. Proceeding as above, it follows that $\pi^k_i \in \Pi^Q_i$ for $i \in I$.

For $i \in I$, let ψ_i be the strategy that requires agent i to behave as follows. Agent i chooses action z at each time before u. Agent i plays strategy π^1_i at any time $t \in \Lambda^1(\{s_\tau\}_{\tau\in[0,T)})$. That is, agent i follows π^1_i at any time no less than u and no greater than the time that the first non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]})$ would be chosen when playing π. For every integer $k \geq 2$, agent i plays strategy π^k_i at any time t satisfying $t \notin \Lambda^{k-1}(\{s_\tau\}_{\tau\in[0,T)})$ and $t \in \Lambda^k(\{s_\tau\}_{\tau\in[0,T)})$. That is, agent i follows π^k_i between the times that the $(k-1)^{\text{th}}$ and k^{th} non-z actions after reaching $\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]})$ would be chosen when playing π. Note that $\psi = (\psi_j)_{j\in I}$ and $\pi = (\pi_j)_{j\in I}$ almost surely induce the same path of play by every agent for all $t \in [u, T)$. We prove below that $\xi^i_b(\psi_i, \psi_{-i})$ is progressively measurable for $i \in I$. It will then follow that $\xi^i_b(\pi_i, \pi_{-i})$ is progressively measurable for $i \in I$.

For any positive integer k, let Θ_k denote the set that consists of each pair $(t, \omega) \in [u, T) \times \Omega$ for which there exist exactly $k-1$ values of $\tilde{t} \in [u, t)$ such that $\phi^i_{\tilde{t}}[\tilde{k}_u(\{s_\tau(\omega)\}_{\tau\in[0,u]}), \{s_\tau(\omega)\}_{\tau\in(u,T)}, \psi] \neq z$ for some $i \in I$. For $k = 1$, this condition means that time t is no less than u and no greater than the time that the first non-z action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]})$ would be chosen when the sample point is ω and strategy profile ψ is played by the agents. For $k > 1$, this condition means that time t is greater than the time of the $(k-1)^{\text{th}}$ but no greater than the time of the k^{th} non-z

action after reaching $\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]})$ when the sample point is ω and strategy profile ψ is played by the agents.

Recall that $\pi_i^k \in \Pi_i^Q$ for $i \in I$ and every positive integer k, which implies that $\xi_b^i(\pi_i^k, \pi_{-i}^k)$ is progressively measurable. It follows that for each positive integer k, the set Θ_k is progressively measurable.[43] For $i \in I$ and any positive integer k, let $\tilde{\xi}_b^i[(\pi_j^k)_{j\in I}]$ denote the stochastic process that is equal to $\xi_b^i[(\pi_j^k)_{j\in I}]$ on the set Θ_k and is equal to zero elsewhere. It also follows that for each positive integer k, the stochastic process $\tilde{\xi}_b^i[(\pi_j^k)_{j\in I}]$ is progressively measurable for $i \in I$.

For any $t \geq u$ and positive integer k, let Θ_k^t denote the set consisting of every pair $(\tau, \omega) \in \Theta_k$ such that $\tau \leq t$. Define $\Theta^t = \bigcup_{k=1}^{\infty} \Theta_k^t$. For any $t \geq u$, let Υ^t denote the set consisting of every pair $(\tau, \omega) \in [u, t] \times \Omega$. Recall that $\pi_i \in \Pi_i^C$ is a traceable and frictional strategy for $i \in I$. Hence, theorem 1 implies that given any realization of the shock process $\{s_\tau\}_{\tau\in[0,u]}$ up to time u, there is conditional probability one that there exists only finitely many values of $\tilde{t} \in [u, T)$ such that $\phi_{\tilde{t}}^i[\tilde{k}_u(\{s_\tau\}_{\tau\in[0,u]}), \{s_\tau\}_{\tau\in(u,T)}, \pi] \neq z$ for some $i \in I$. It follows that for any $t \geq u$, the set consisting of each pair (τ, ω) such that $(\tau, \omega) \in \Upsilon^t$ and $(\tau, \omega) \notin \Theta^t$ has measure zero with respect to the product measure on $\mathcal{B}([0, t]) \times \mathcal{F}_t$.

Note that for any positive integer k and $i \in I$, the stochastic process $\xi_b^i[(\psi_j)_{j\in I}]$ is equal to the stochastic process $\tilde{\xi}_b^i[(\pi_j^k)_{j\in I}]$ on the set Θ_k. Recall that each set Θ_k along with every stochastic process $\tilde{\xi}_b^i[(\pi_j^k)_{j\in I}]$ is progressively measurable. Hence, $\xi_b^i(\psi_i, \psi_{-i})$ is progressively measurable for $i \in I$. □

Proof of Proposition 1. Assume that for some $j \in I$, there exists $\psi_j \in \Psi_j$ such that $\psi_j \notin \Pi_j^C$. By definition, there exists $\psi_{-j} \in \Pi_{-j}^Q$ along with v such that the stochastic process $\xi_v^j(\psi_j, \psi_{-j})$ is not progressively measurable. It follows from $\Pi_{-j}^Q \subseteq \Psi_{-j}$ that $\psi_{-j} \in \Psi_{-j}$. Hence, there exists $i \in I$ such that the stochastic process $\xi_b^i(\pi_i, \pi_{-i})$ is not progressively measurable for some $\pi_i \in \Psi_i$, $\pi_{-i} \in \Psi_{-i}$, and b. □

A.2 Proofs for Section 6

Proof of Proposition 5. Choose any strategy profile $\pi \in \times_{i\in I}\bar{\Pi}_i^{TF}$. Define the strategy profile $\pi' \in \times_{i\in I}\bar{\Pi}_i^{TF}$ as follows. Let $k_u = \left(\{s_t\}_{t\in[0,u]}, \{(a_t^i)_{i\in I}\}_{t\in[0,u)}\right)$ be any history up to an arbitrary time u. If $a_t^i = \pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(a_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right)$ for each $i \in I$ and

[43] Given any $\Theta \subseteq [0, T) \times \Omega$, let χ_Θ denote the indicator function of Θ. The set Θ is said to be progressively measurable if for any $v \geq 0$ the function χ_Θ is measurable with respect to the product sigma-algebra $\mathcal{B}([0, v]) \times \mathcal{F}_v$.

all $t \in [0, u)$, then let $\pi_i'(k_u) = \pi_i(k_u)$ for all $i \in I$. If there exists $t \in [0, u)$ and $i \in I$ such that $a_t^i \neq \pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(a_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right)$, then for each $i \in I$, let $\pi_i'(k_u) \neq z$ if s_t is in $\tilde{S}_t$ for all $t \in [0, \tilde{t}]$ and $u = n\tilde{t}$ for some positive integer n, and let $\pi_i'(k_u) = z$ otherwise.

The strategy profile π' is an SPE because if agent $i \in I$ deviates at h_0, then the expected payoff to agent i is $\chi_i(h_0, \pi') \leq \zeta_i(h_0)$. Moreover, consider any history $k_u = \left(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\right)$ up to an arbitrary time $u > 0$. By assumption, there exists with probability one some $t < u$ such that $\pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(b_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right) \neq b_t^i$ for some $i \in I$. It follows that for each $i \in I$, the function $\phi_{n\tilde{t}}^i\big[\left(\{s_t(\omega)\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\right)$, $\{s_t(\omega)\}_{t\in(u,\infty)}, \pi'\big]$ from Ω to A_i is not measurable, where n is any integer such that $n\tilde{t} \geq u$. Hence, the process $\xi_b^i(\pi')$ with $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$ is not progressively measurable, so that the expected payoff to agent i is $\chi_i(k_u, \pi') \leq \zeta_i(k_u)$ when the other agents follow strategy profile π'. $\square$

Proof of Proposition 6. Let π be an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^C\right)$, and consider the game $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$ with $\chi_i(h_t, \tilde{\pi}) \leq \zeta_i(h_t)$ for all $\tilde{\pi} \in \times_{j\in I}\bar{\Pi}_j^{TF}$, $h_t \in H$, and $i \in I$. Let $k_u = \left(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\right)$ be any history up to an arbitrary time u, and denote $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$.

For any $i \in I$, choose any $\pi_i' \in \bar{\Pi}_i^{TF}$. If there is some $j \in I$ such that $\xi_b^j(\pi_i', \pi_{-i})$ is not progressively measurable, then $U_i[k_u, (\pi_i', \pi_{-i})] = \chi_i[k_u, (\pi_i', \pi_{-i})] \leq \zeta_i(k_u)$, whereas $U_i(k_u, \pi) = V_i(k_u, \pi) \geq \zeta_i(k_u)$. Suppose now that $\xi_b^j(\pi_i', \pi_{-i})$ is progressively measurable for all $j \in I$. Let π_i'' with $\pi_i''(h_t) = z$ for $t < u$ be defined such that $\pi_i''\big[\left(\{s_\tau\}_{\tau\in[0,t]}, \{(d_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right)\big] = \phi_t^i\big[\left(\{s_\tau\}_{\tau\in[0,u]}, \{(b_\tau^j)_{j\in I}\}_{\tau\in[0,u)}\right), \{s_\tau\}_{\tau\in(u,T)}, (\pi_i', \pi_{-i})\big]$ for each realization of the shock process $\{s_\tau\}_{\tau\in[0,T)}$ and any action path $\{(d_\tau^j)_{j\in I}\}_{\tau\in[0,t)}$ up to an arbitrary time $t \geq u$. Note that $\pi_i'' \in \bar{\Pi}_i$ given the assumption that $\bar{A}_i(h_t) = A_i$ for all $h_t \in H$ and $i \in I$. By the definition of π_i'', $\pi_i'' \in \Pi_i^{TF}$, and the stochastic process $\xi_b^j(\pi_i'', \pi_{-i})$ is the same as $\xi_b^j(\pi_i', \pi_{-i})$ for all $j \in I$, which implies that $U_i[k_u, (\pi_i'', \pi_{-i})] = V_i[k_u, (\pi_i'', \pi_{-i})] = V_i[k_u, (\pi_i', \pi_{-i})] = U_i[k_u, (\pi_i', \pi_{-i})]$. Moreover, π'' is quantitative and hence calculable. Since π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^C\right)$, it must be that $U_i(k_u, \pi) \geq U_i[k_u, (\pi_i'', \pi_{-i})]$. It follows that $U_i(k_u, \pi) \geq U_i[k_u, (\pi_i', \pi_{-i})]$. Hence, no agent i has an incentive to deviate from π_i to π_i' at k_u, which proves that π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$. $\square$

Proof of Theorem 3. We first show that π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^C\right)$, assuming that π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$ and π_i is calculable for each $i \in I$. Then we confirm that π_i is calculable for each $i \in I$ given that π is synchronous and measurably attachable.

To show that π is an SPE of $\Gamma\big(\times_{i\in I}\bar{\Pi}_i^C\big)$, consider any $h_t \in H$. Given any $i \in I$, choose any $\pi_i' \in \bar{\Pi}_i^C$. Since π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$, it must be that $U_i(h_t, \pi) \geq U_i[h_t, (\pi_i', \pi_{-i})]$, where $U_i(h_t, \pi) = V_i(h_t, \pi)$ and $U_i[h_t, (\pi_i', \pi_{-i})] = V_i[h_t, (\pi_i', \pi_{-i})]$ because $\pi \in \times_{j\in I}\Pi_j^C$. It follows that π is an SPE of $\Gamma\big(\times_{i\in I}\bar{\Pi}_i^C\big)$.

Given any $i \in I$, we now confirm that $\pi_i \in \Pi_i^C$. Define $\pi_j^z(h_t) = z$ for each $j \in I$ and every $h_t \in H$. Choose any $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$ as well as any $\pi_{-i}' \in \times_{j\neq i}\Pi_j^Q$. Define the stopping time θ^0 as follows. For any $\omega \in \Omega$, let $\hat{\Upsilon}^0(\omega)$ denote the set consisting of all $t \in [u, T)$ such that $\phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, \pi\big] \neq z$ for some $e \in I$, and let $\bar{\Upsilon}^0(\omega)$ denote the set consisting of all $t \in [u, T)$ such that $\phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, (\pi_i, \pi_{-i}')\big] \neq z$ for some $e \in I$. Let $\theta^0(\omega)$ be the lesser of the infimum of $\hat{\Upsilon}^0(\omega)$ and the infimum of $\bar{\Upsilon}^0(\omega)$. Note that $\xi_b^i[\psi(\pi, \pi^z, \theta^0)]$ and $\xi_b^i\{\psi[(\pi_i, \pi_{-i}'), \pi^z, \theta^0]\}$ are the same progressively measurable stochastic process.

Apply the following procedure iteratively for every positive integer k. Define the stopping time θ^k as follows. For any $\omega \in \Omega$, let $\hat{\Upsilon}^k(\omega)$ denote the set consisting of all $t \in (\theta^{k-1}(\omega), T)$ such that $\phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, \psi[(\pi_i, \pi_{-i}'), \pi, \theta^{k-1}]\big] \neq z$ for some $e \in I$, and let $\bar{\Upsilon}^k(\omega)$ denote the set consisting of all $t \in (\theta^{k-1}(\omega), T)$ such that $\phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, (\pi_i, \pi_{-i}')\big] \neq z$ for some $e \in I$. Let $\theta^k(\omega)$ be the lesser of the infimum of $\hat{\Upsilon}^k(\omega)$ and the infimum of $\bar{\Upsilon}^k(\omega)$, where $P(\{\omega \in \Omega : \theta^k(\omega) > \theta^{k-1}(\omega)\} \cup \{\omega \in \Omega : \theta^{k-1}(\omega) = \infty\}) = 1$ because $\psi[(\pi_i, \pi_{-i}'), \pi, \theta^{k-1}], (\pi_i, \pi_{-i}') \in \times_{j\in I}\Pi_j^{TF}$. Note that $\xi_b^i\big(\psi\{\psi[(\pi_i, \pi_{-i}'), \pi, \theta^{k-1}], \pi^z, \theta^k\}\big)$ and $\xi_b^i\{\psi[(\pi_i, \pi_{-i}'), \pi^z, \theta^k]\}$ are the same progressively measurable stochastic process.

Suppose that the sequence $\{\theta^k\}_{k=1}^\infty$ does not converge almost surely to ∞. Then $P[\{\omega \in \Omega : \lim_{k\to\infty}\theta^k(\omega) < \infty\}] \neq 0$. For each $\omega \in \Omega$, let $\Xi(\omega)$ denote the set consisting of all $t \in [u, T)$ such that $\phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, (\pi_i, \pi_{-i}')\big] \neq z$ for some $e \in I$. Letting E denote the set consisting of all $\omega \in \Omega$ such that $\{t \in \Xi(\omega) : t \leq c\}$ contains only finitely many elements for any $c \in [u, \infty)$, we have $P(\{\omega \in \Omega : \omega \in E\}) = 1$ because $(\pi_i, \pi_{-i}') \in \times_{j\in I}\Pi_j^{TF}$. The definition of E also implies that for all $\omega \in E$ such that $\lim_{k\to\infty}\theta^k(\omega) < \infty$, there exists $\tilde{t}(\omega) \in [u, \lim_{k\to\infty}\theta^k(\omega))$ such that $\phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, (\pi_i, \pi_{-i}')\big] = z$ for each $e \in I$ and all $t \in [\tilde{t}(\omega), \lim_{k\to\infty}\theta^k(\omega))$.

For any $\omega \in E$ such that $\lim_{k\to\infty}\theta^k(\omega) < \infty$, choose any $\tilde{k}(\omega) \geq 1$ such that

$\theta^{\tilde{k}(\omega)-1}(\omega) \geq \tilde{t}(\omega)$. Assuming now that π is synchronous, $\phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, \psi[(\pi_i,\pi'_{-i}),\pi,\theta^{\tilde{k}(\omega)-1}]\big] = \phi_t^e\big[\big(\{s_\tau(\omega)\}_{\tau\in[0,u]}, (\{b_\tau^j\}_{\tau\in[0,u)})_{j\in I}\big), \{s_\tau(\omega)\}_{\tau\in(u,T)}, (\pi_i,\pi'_{-i})\big] = z$ for each $e \in I$, any $t \in (\theta^{\tilde{k}(\omega)-1}(\omega), \lim_{k\to\infty}\theta^k(\omega))$, and all $\omega \in E$ such that $\lim_{k\to\infty}\theta^k(\omega) < \infty$. This implies that $\theta^{\tilde{k}(\omega)}(\omega) \geq \lim_{k\to\infty}\theta^k(\omega)$ for all $\omega \in E$ such that $\lim_{k\to\infty}\theta^k(\omega) < \infty$, from which it follows that there is a set of nonzero measure consisting of $\omega \in E$ with $\lim_{k\to\infty}\theta^k(\omega) < \infty$ for which there exists $l \geq 1$ such that $\theta^l(\omega) > \lim_{k\to\infty}\theta^k(\omega)$. However, $\theta^k(\omega)$ is nondecreasing in k by construction, so this is a contradiction. Thus, the sequence $\{\theta^k\}_{k=1}^\infty$ must converge almost surely to ∞. Since the stochastic process $\xi_b^i\{\psi[(\pi_i,\pi'_{-i}),\pi^z,\theta^k]\}$ is progressively measurable for all $k \geq 1$ and the sequence $\{\theta^k\}_{k=1}^\infty$ converges almost surely to ∞, the stochastic process $\xi_b^i(\pi_i,\pi'_{-i})$ is progressively measurable. It follows that $\pi_i \in \Pi_i^C$. $\square$

Proof of Proposition 7. Let the synchronous strategy profile $\pi \in \Pi^A$ be a pseudo-SPE of $\Gamma(\Pi^A \cap \bar{\Pi})$. Assume that $\chi_i(h_t,\tilde{\pi}) \leq \zeta_i(h_t)$ for all $\tilde{\pi} \in \times_{j\in I}\bar{\Pi}_j^{TF}$, $h_t \in H$, and $i \in I$. It suffices to show that π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$ because it will then follow from theorem 3 that π is an SPE of $\Gamma\big(\times_{i\in I}\bar{\Pi}_i^C\big)$.

To show this, let $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\big)$ be any history up to an arbitrary time u, and denote $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$. For any $i \in I$, choose any $\pi'_i \in \bar{\Pi}_i^{TF}$. If there is some $j \in I$ such that $\xi_b^j(\pi'_i,\pi_{-i})$ is not progressively measurable, then $U_i[k_u,(\pi'_i,\pi_{-i})] = \chi_i[k_u,(\pi'_i,\pi_{-i})] \leq \zeta_i(k_u)$, whereas $U_i(k_u,\pi) = V_i(k_u,\pi) \geq \zeta_i(k_u)$. Hence, $U_i(k_u,\pi) \geq U_i[k_u,(\pi'_i,\pi_{-i})]$ holds in this case.

Suppose now that $\xi_b^j(\pi'_i,\pi_{-i})$ is progressively measurable for all $j \in I$. First, we construct $\pi'' \in \Pi^A$ such that $\xi_b^e(\pi'')$ with $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$ is the same stochastic process as $\xi_b^e[(\pi'_i,\pi_{-i})]$ for all $e \in I$. For each $e \in I$, let π''_e with $\pi''_e(h_t) = z$ for $t < u$ be defined such that $\pi''_e\big[\big(\{s_\tau\}_{\tau\in[0,t]}, \{(d_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\big)\big] = \phi_t^e\big[\big(\{s_\tau\}_{\tau\in[0,u]}, \{(b_\tau^j)_{j\in I}\}_{\tau\in[0,u)}\big), \{s_\tau\}_{\tau\in(u,T)}, (\pi'_i,\pi_{-i})\big]$ for each realization of the shock process $\{s_\tau\}_{\tau\in[0,T)}$ and any action path $\{(d_\tau^j)_{j\in I}\}_{\tau\in[0,t)}$ up to an arbitrary time $t \geq u$. By the definition of π'', $\pi''_e \in \Pi_e^{TF}$ for each $e \in I$, the stochastic process $\xi_b^j(\pi'')$ is the same as $\xi_b^j(\pi'_i,\pi_{-i})$ for all $j \in I$, and π'' is a measurably attachable strategy profile.

Second, note that since π is a pseudo-SPE of $\Gamma(\Pi^A \cap \bar{\Pi})$, $\pi'' \in \Pi^A$ and the property that $\xi_b^e(\pi'')$ and $\xi_b^e[(\pi'_i,\pi_{-i})]$ are the same stochastic process for all $e \in I$ imply that $V_i(k_u,\pi) \geq V_i[k_u,(\pi'_i,\pi_{-i})]$. Since $U_i(k_u,\pi) = V_i(k_u,\pi)$ and $U_i[k_u,(\pi'_i,\pi_{-i})] = V_i[k_u,(\pi'_i,\pi_{-i})]$, we conclude that $U_i(k_u,\pi) \geq U_i[k_u,(\pi'_i,\pi_{-i})]$ in this case, too.

Overall, no agent i has an incentive to deviate from π_i to any π'_i at k_u, which proves that π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$. $\square$

References

Ambrus, A., and S. E. Lu (2015): "A Continuous-Time Model of Multilateral Bargaining," *American Economic Journal: Microeconomics*, 7(1), 208–249.

Bergin, J., and W. B. MacLeod (1993): "Continuous Time Repeated Games," *International Economic Review*, 34(1), 1993.

Bonatti, A., G. Cisternas, and J. Toikka (2017): "Dynamic Oligopoly with Incomplete Information," *Review of Economic Studies*, 84(2), 503–546.

Caplin, A. S., and D. F. Spulber (1987): "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, 102(4), 703–726.

Dilme, F. (2017): "Dynamic Quality Signaling with Hidden Actions," Mimeo, University of Bonn.

Dixit, A. (1989): "Entry and Exit Decisions under Uncertainty," *Journal of Political Economy*, 97(3), 620–638.

Faingold, E., and Y. Sannikov (2011): "Reputation in Continuous-Time Games," *Econometrica*, 79(3), 773–876.

Fudenberg, D., and J. Tirole (1985): "Preemption and Rent Equalization in the Adoption of New Technology," *Review of Economic Studies*, 52(3), 383–401.

Kamada, Y., and M. Kandori (2018a): "Revision Games, Part I: Theory," Mimeo, University of California, Berkeley.

——— (2018b): "Revision Games, Part II: Applications and Robustness," Mimeo, University of California, Berkeley.

Kamada, Y., and N. Rao (2018): "Sequential Exchange with Stochastic Transaction Costs," Mimeo, University of California, Berkeley.

McDonald, R., and D. Siegel (1986): "The Value of Waiting to Invest," *The Quarterly Journal of Economics*, 101(4), 707–727.

Ortner, J. (2016): "A Continuous-Time Model of Bilateral Bargaining," Mimeo, Boston University.

Sannikov, Y. (2007): "Games with Imperfectly Observable Actions in Continuous Time," *Econometrica*, 75(5), 1285–1329.

——— (2008): "A Continuous-Time Version of the Principal-Agent Problem," *Review of Economic Studies*, 75(3), 957–984.

Simon, L. K., and M. B. Stinchcombe (1989): "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 57(5), 1171–1214.