

STRATEGIES IN STOCHASTIC CONTINUOUS-TIME GAMES*

Yuichiro Kamada[†] Neel Rao[‡]

December 31, 2023

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 Ken-Ichi Akao, Eduardo Faingold, Drew Fudenberg, Oliver Hart, Michihiro Kandori, Akitada Kasahara, Jay Lu, Bentley MacLeod, Eilon Solan, Phillip Strack, Takuo Sugaya, Satoru Takahashi, Jean Tirole, Juuso Toikka, and Yasutora Watanabe 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 working paper “Sequential Exchange with Stochastic Transaction Costs” by the same authors.

[†]Haas School of Business, University of California Berkeley, 2220 Piedmont Avenue, Berkeley, CA 94720-1900, USA, and University of Tokyo, Faculty of Economics, 7-3-1 Hongo, Bunkyo, Tokyo, 113-0033, Japan, e-mail: y.cam.24@gmail.com

[‡]University at Buffalo, SUNY, Department of Economics, 415 Fronczak Hall, Buffalo, NY 14260-1520, USA, 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 agents 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 been successful in the analysis of stochastic environments as well, 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.

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 when the other agents do not move in the future, there exists a history consistent with an agent's strategy. *Frictionality* eliminates the latter possibility, requiring that in any history that is consistent with an agent's strategy, the agent 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.

A further issue regards the measurability of the action process. 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. Simply requiring the action process of an agent to be progressively measurable (as is standard in the literature) is problematic because the strategy of one agent affects whether the strategy of another agent induces a measurable action process under

perfect monitoring.¹ Thus, simply restricting the action process of an agent to be progressively measurable even when the behavior of another agent is nonmeasurable (which would correspond to the standard approach) would result in an extremely limited set of strategies as explained in section 4.1, ruling out a variety of situations in which the actions of one agent are influenced by the behavior of another agent.

Therefore, we invent a two-step 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 that includes the set of all quantitative strategies and implies measurable actions.

A few methods have previously been proposed for defining strategy spaces in continuous time when the environment is deterministic, but they are not directly applicable to the stochastic case. In particular, the “inertia” assumption from Bergin and MacLeod (1993) would ensure that each strategy profile induces a unique path of play. This condition essentially requires that at each history up to a given time t , there exists a time interval of positive length during which agents do not change their actions.² There are at least two ways to extend this definition to a stochastic setting. In one approach, the aforementioned time interval can depend on the history up to time t but not on the realization of uncertainty after time t . The other approach allows this time interval to also depend on the realization of uncertainty after time t .

The former version of inertia may be too strong in stochastic environments, where given a history up to time t and any $\epsilon > 0$, the state variable may change quickly within the time 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 $(t, t + \epsilon)$. We present various examples to illustrate that such a situation naturally arises in stochastic settings.³ The latter version is weaker but does not guarantee

¹This type of complexity does not arise in models of continuous-time games with imperfectly observable actions such as Sannikov (2007).

²Bergin and MacLeod (1993) also consider a weaker condition involving the completion of the set of inertia strategies, and we discuss it in footnote 43.

³When comparing our approach with existing techniques, we consider many examples that involve rapid transitions between states, but as the applications in section 5 illustrate, such a property is not necessary for our methodology to apply in situations that are not covered by existing techniques in the literature.

measurable behavior. In a deterministic environment, the inertia condition would essentially suffice to ensure that the resulting path of play is measurable, but more is involved in a stochastic setting because the actions of the agents can be contingent on moves of nature.⁴ Accordingly, we introduce the notion of calculability. Furthermore, neither version of inertia implies nor is implied by traceability and frictionality, as shown in section 3.3. We present an example to demonstrate that even the weaker formulation of inertia rules out some applications of interest covered by traceability and frictionality in which the timing of moves by one agent depends on the actions of another agent.⁵

Relatedly, the “admissibility” criterion in Simon and Stinchcombe (1989) when extended to a stochastic environment would imply existence and uniqueness of the path of play. However, this condition requires the number of moves by an agent to be bounded.⁶ As will be seen from our examples in section 5, imposing the bound uniformly so as not to depend on the realization of uncertainty can sometimes be too restrictive. Specifically, the equilibria that we analyze entail no upper bound on the number of moves during any finite time interval. Agents may make infinitely many moves over an infinite horizon with probability one.

Allowing the bound to depend on the realization of uncertainty does not suffice to cover some relevant applications, particularly those in which the number of moves by one agent depends on the behavior of the other agent. Furthermore, since admissibility applies off as well as on the path of play, it eliminates some simple strategies that require an agent to move multiple times following a deviation. By contrast, traceability and frictionality are weaker requirements and do not preclude strategies in which the number of moves in a given finite time interval may be arbitrarily large.⁷

⁴For instance, the strategies in example 8 satisfy inertia as well as traceability and frictionality but do not induce measurable behavior.

⁵Although the inertia condition may be insufficient to model optimal behavior in some of the applications we consider, it may be enough to accommodate ϵ -optimal strategies. In item 3 of remark 3, we explain that restricting attention to ϵ -optimal strategies without defining strategy spaces more generally creates a problem of circularity.

⁶Simon and Stinchcombe (1989) also restrict how agents 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.

⁷More precisely, fix an interval of the nonnegative real line with positive but finite length as well as any realization of uncertainty. 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 consistent with this strategy such that the

As shown in section 3.3, admissibility when applied to our model implies traceability and frictionality, but not vice versa. Moreover, even if agents use admissible strategies, the behavior that arises may not be measurable, necessitating an additional condition like calculability.⁸

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 move at every moment in continuous time⁹ (which are not allowed by Simon and Stinchcombe (1989) or Bergin and MacLeod (1993), either). As the reader will hopefully see, even though our framework is not totally comprehensive, defining strategy spaces in a stochastic environment is a nontrivial task, and our examples show that our model encompasses a number of economically relevant situations. We emphasize that our objective is not to undermine the usefulness of existing restrictions in the literature like those of Bergin and MacLeod (1993) and Simon and Stinchcombe (1989). We discuss those restrictions simply to underscore the additional complications that arise when the environment is stochastic.

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 major class of problems involving the timing of moves is search models, in which agents must choose between trading at the present time and waiting for a better opportunity to trade. Section 5.2 examines a supply chain in the petroleum industry, where the price of oil varies over time, and a well is deciding when to extract each unit oil and transfer it to a refinery, which then processes the input and delivers the output to a customer. The oil well

number of moves by the agent during the given time interval is equal to n .

⁸See example 8 for a particular instance of admissible strategies that do not generate a measurable path of play.

⁹Such 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.

faces a search problem, and its optimal strategy is characterized by a reservation price like in the classic model of McCall (1970). A topic related to investment under uncertainty 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 5.3, 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 agents 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 agents can make offers at random times and must reach an agreement by a specified deadline. Another class of models is revision games as formalized by Kamada and Kandori (2020), where agents have opportunities to alter their actions before interacting at a designated time. In section 5.4, we apply our restrictions to a finite-horizon game that has a deadline by which a buyer must place an order with a seller. The buyer's taste for a good is stochastically changing over time, and the seller receives opportunities at random times to fulfill the order.

Further applications of our framework are considered in the online appendix. Ortner (2019) specifies a bargaining model in continuous time, where the bargaining power of each agent 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 a distributor in a model of inventory adjustment with a randomly changing stock as well as a model of exchange where prices evolve stochastically.

Given the difficulty of restricting the strategy space to prevent nonmeasurable behavior, a pertinent question is whether the issue can simply be avoided by allowing nonmeasurable behavior and exogenously assigning it a (for example, low) expected

payoff. In section 6, we examine this question, and relate the equilibria under this alternative approach to those under our original methodology. A problem is that the assignment of expected payoffs to nonmeasurable behavior is inherently arbitrary as the conditional expectation may be undefined and the set of equilibria depends on the particular assignment of expected payoffs to nonmeasurable behavior. Under some conditions, a type of “folk theorem” can even be obtained, whereby arbitrary behavior can be sustained in equilibrium by punishing deviations with nonmeasurable behavior that is assigned an extremely low payoff. In addition, we show that any equilibrium under the calculability restriction can be supported as an equilibrium under the payoff assignment approach for some assignment of expected payoffs to nonmeasurable behavior. A partial converse of this result is also obtained.

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 non-measurable behavior. Proofs of the major results are in the appendix. The online appendix provides further applications, extensions, derivations, and discussions.

2 Model

There are a finite number 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 (Ω, \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 $\mathcal{B}(S)$.¹⁰ The state space S can be any set but in applications is typically the product of the set of variables representing the environment (e.g., the amount of resources, the price, the cost) and \mathbb{R}_+ representing calendar time.¹¹ The shock process $\{s_t\}_{t \in [0, T)}$ is assumed

¹⁰The notation $\mathcal{B}(X)$ represents the Borel sigma-algebra on a topological space X .

¹¹The former set could also be a binary set representing whether there is a Poisson hit at a given moment.

to be progressively measurable.¹²

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_t^i represent the action that agent i chooses at time t . The collection $\{a_t^i\}_{t \in [0, T)}$ of agent i 's actions indexed by time $t \in [0, T)$ is called an **action path**. Letting $u \in [0, T)$, the collection $\{a_t^i\}_{t \in [0, u]}$ is said to be an action path up to time u . Neither the probability space nor the shock process depends in any way on the actions of the agents.¹³

A history of the game is represented as $h = \{s_t, (a_t^i)_{i \in I}\}_{t \in [0, T)}$. 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 $(\{s_t\}_{t \in [0, u]}, \{(a_t^i)_{i \in I}\}_{t \in [0, u]})$. 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 $\bar{A}_i : H \rightarrow 2^{A_i}$. 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 $\pi_i : H \rightarrow A_i$.¹⁴ Let Π_i with generic element π_i represent the set of all strategies for agent i .¹⁵ In addition,

¹²A 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 $\mathcal{B}([0, v]) \times \mathcal{F}_v$.

¹³The assumption that the shock process does not depend on the actions of the agents is not particularly restrictive. For example, in order to model a situation where uncertainty evolves according to an Itô process that depends on the actions of the agents, the shock process could be specified as a Wiener process, and the value of the Itô process at each time could be encoded in the action space, which may be history-dependent. A related example is the application in section 5.1, where the volume of a forest, which is affected by the actions of agents to harvest wood and evolves according to an arithmetic Brownian motion between harvests, is specified in the set of feasible actions, which we define shortly.

¹⁴Simon and Stinchcombe (1989) allow for sequential moves at a single moment of time. The definition of the strategy space here rules out such behavior. We can thereby express histories in a natural way as specifying a profile of actions at each time and define strategies in the standard way as a mapping from histories to actions. In addition, it may not be realistic to allow players to move sequentially without any elapse of time. 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. If agents could immediately respond to a deviation at the current Poisson jump instead of waiting for the next Poisson jump to respond, such an inefficiency would not arise.

¹⁵The 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 $\pi_i(h_t) \in \bar{A}_i(h_t)$ for any $h_t \in H$. Let $\bar{\Pi}_i$ denote the set of all feasible strategies for agent i . For any profile of sets $(D_j)_{j \in I}$ and any $i \in I$, let $D_{-i} = \times_{j \in I \setminus \{i\}} D_j$.

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.¹⁶

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 no 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 $1/m \in (0, t)$. \square

In the preceding example, an action path consistent with the specified strategy profile fails to exist because a given set of times may not have a least element in continuous time. Similarly, restricting attention to Markov strategies is not enough to ensure that strategy spaces are well defined because in continuous time there may not exist a least time at which a state variable satisfies a given condition.¹⁷

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).

¹⁶These problems arise because the strategy of an agent is a contingent plan in which the action of an agent at a given time may be conditional on past actions. Similar examples can be constructed with just one agent, but in that case such difficulties could be avoided by defining a strategy as a mapping from each time to an action at that time without reference to previous actions.

¹⁷A strategy is said to be Markov if the action that it specifies at any history up to a given time

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 the current time t is equal to $1/m$ for some positive integer m and another agent has chosen action x at time $1/n$ for every positive integer n greater than m , then each agent chooses x at time t . 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 time t is equal to $1/m$ for some positive integer m , and they choose z otherwise. It is straightforward to check by inspection that all the agents are following the given strategy profile in each case. \square

To overcome the difficulties with specifying strategies in continuous-time games as in the above examples, we use a framework in which strategies are defined as mappings that assign to each history an action at that history. However, a question is whether there is an alternative formulation that bypasses these difficulties. A seemingly simple possibility may be to define strategies using a sequence of stopping times. Nonetheless, this approach involves a number of complications, as explained below.

Remark 1. (Formulating strategies in terms of stopping times) Suppose that each agent chooses a positive stopping time at the null history as well as a move to make once that time is reached. Thereafter, whenever some agent moves, each agent chooses another stopping time and move to make at that time.

An issue with such a formulation of strategies is that the resulting path of play may not be determined for all times during the game. In particular, the sequence of stopping times may be infinite and have an upper bound that is less than the time at which the game ends. Deriving the path of play after this upper bound would involve taking the limit of this sequence and specifying another sequence in a well-defined manner.

This formulation also leaves open the question of how to define behavior at histories where the current time is already greater than the applicable stopping time. One solution to this problem may be to specify the stopping time as a function that depends only on the current value of a state variable that summarizes the history of the game up to that time.

depends on the current time as well, but that would be no simpler than our approach where a strategy is just a mapping from histories to actions. Moreover, such a specification in which the stopping time depends on the current time would necessitate an additional condition to guarantee the consistency of stopping times across different histories.

In addition, note that under this formulation, a strategy is not simply a sequence of stopping times but also entails specifying a stochastic process encoding the move at each stopping time. As in section 4, the action process would need to be suitably restricted so as to be progressively measurable, thereby ensuring that expected payoffs are well defined.

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 $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ 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 $\{a_t^i\}_{t \in [0, u]} = \{b_t^i\}_{t \in [0, u]}$. Given any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u , let $W_i^{TR}(k_u; \pi_i)$ be the set of all $\omega \in \Omega$ such that for any $\{a_t^{-i}\}_{t \in [0, T]} \in \Gamma_{-i}(\{b_t^{-i}\}_{t \in [0, u]})$ with $a_t^j = z$ for all $t > u$ and $j \neq i$, there exists $\{a_t^i\}_{t \in [0, T]} \in \Gamma_i(\{b_t^i\}_{t \in [0, u]})$ for which the history $h = \{s_t(\omega), (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i at each $t \in [u, T]$. In other words, $W_i^{TR}(k_u; \pi_i)$ represents the set of all sample paths of the shock process such that given the action paths of the agents up to time u , there exists a history consistent with agent i 's strategy π_i from time u onwards, assuming that agent i 's opponents do not move after time u .

Definition 2. Given $i \in I$, the strategy $\pi_i \in \Pi_i$ is **traceable** if for any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u , the set $W_i^{TR}(k_u; \pi_i)$ is measurable with respect to \mathcal{F} and has conditional probability one given that $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$.

Intuitively, a strategy for a given agent is said to be traceable if there exists a history that is consistent with the strategy when the other agents do not move in the future. Traceability excludes the strategy in example 1.

Let $\Xi_i(t)$ denote the set consisting of every action path $\{a_\tau^i\}_{\tau \in [0, T]}$ of agent $i \in I$ for which there exists no $u > t$ such that the set $\{\tau \in [t, u] : a_\tau^i \neq z\}$ contains infinitely many elements. Given any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u , let $W_i^{FR}(k_u; \pi_i)$ be the set of all $\omega \in \Omega$ such that $\{a_t^i\}_{t \in [0, T]} \in \Xi_i(u)$ for all $\{a_t^i\}_{t \in [0, T]} \in \Gamma_i(\{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ such that there exists $\{a_t^{-i}\}_{t \in [0, T]} \in \Gamma_{-i}(\{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ for which the history $h = \{s_t(\omega), (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i at each $t \in [u, T]$. In other words, $W_i^{FR}(k_u; \pi_i)$ represents the set of all sample paths of the shock process such that given the action paths of the agents up to time u , agent i 's strategy π_i induces only finitely many moves in any finite interval of time in the future, regardless of the behavior of agent i 's opponents from time u onwards.

Definition 3. Given $i \in I$, the strategy $\pi_i \in \Pi_i$ is **frictional** if for any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u , the set $W_i^{FR}(k_u; \pi_i)$ is measurable with respect to \mathcal{F} and has conditional probability one given that $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$.

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, where the number of moves can depend on the sample path of the shock process as well as the action paths of the other agents. Frictionality excludes the strategy in example 2. It also rules out flow moves.

Traceability and frictionality restrict the strategy space for each agent. 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.¹⁸

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

¹⁸The statement of theorem 1 is cumbersome, but it is inevitably so. For example, we cannot simplify the statement of the theorem by making reference to the conditional probability of an event given the history k_u up to time u because the action path $\{(b_t^j)_{j \in I}\}_{t \in [0, u]}$ of the agents up to time u , which is part of k_u , cannot be treated as a well defined random variable at this point. Hence, we first fix the history k_u in the statement of the theorem and then consider the conditional expectation given the shock process $\{s_t\}_{t \in [0, u]}$ 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 $(\{a_t^j\}_{t \in [0, T]})_{j \in I}$ of action paths with $\{a_t^j\}_{t \in [0, T]} \in \Gamma_j(\{b_t^j\}_{t \in [0, u]})$ for each $j \in I$ such that the history $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T]$.
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 agents' strategies at all times.¹⁹ 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 . Below we comment on the consequences of relaxing the frictionality assumption.

Remark 2. (Weak frictionality and existence of a unique outcome) Frictionality requires that for any action path of the other agents, an agent moves only finitely many times in a finite interval of time. A less restrictive way to specify frictionality is to impose this requirement only when the other agents move just finitely often in a finite time interval. Formally, a strategy is **weakly frictional** if it meets the definition of frictionality when $W_i^{FR}(k_u)$ is restricted to be the set of all $\omega \in \Omega$ such that $\{a_t^i\}_{t \in [0, T]} \in \Xi_i(u)$ for all $\{a_t^j\}_{t \in [0, T]} \in \Gamma_j(\{b_t^j\}_{t \in [0, u]})$ such that there exists

¹⁹See the online appendix for an example of such a strategy profile.

$\{a_t^{-i}\}_{t \in [0, T]} \in \Gamma_{-i}(\{b_t^{-i}\}_{t \in [0, u]}) \cap \Xi_{-i}(u)$ for which the history $h = \{s_t(\omega), (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i at each $t \in [u, T]$.

Suppose first that all the agents use traceable and weakly frictional strategies. Then the resulting path of play may not be unique. For instance, the strategies in example 2 satisfy traceability and weak frictionality, but there are multiple action paths consistent with this strategy profile. Suppose next that all the agents use traceable and weakly frictional strategies and no more than one agent uses a strategy that is not frictional. Then both items in theorem 1 continue to hold.²⁰ In particular, the continuation path at any history up to a given time is unique, with each agent moving only finitely many times in any finite time interval. This is a useful extension because it enables modelling situations where one agent quickly responds to every move by another agent (see section 5.2 for an example).

3.3 Comparison with Previous Literature

A question is how traceability and frictionality relate to restrictions on the strategy space that were previously developed for a deterministic setting. Comparing our methodology to existing approaches helps to assess how robust it is and clarify what its limits may be. One condition is inertia from Bergin and MacLeod (1993), 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 pathologies related to the nonexistence or nonuniqueness of the path of play. The inertia condition might be extended to a stochastic environment in at least two ways.

One definition involves imposing inertia uniformly at each history. This version of the condition requires that given any history h_t up to an arbitrary time t , there exists $\epsilon > 0$ such that an agent does not move in the time interval $(t, t + \epsilon)$, where ϵ cannot depend on the realization of the shock process $\{s_\tau\}_{\tau \in (t, \infty)}$ after time t . Formally, a strategy $\pi_i \in \Pi_i$ is said to be **uniformly inertial** if for every history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u , there exists $\epsilon > 0$ such that $\pi_i(h_\tau) = z$ for all $\tau \in (u, u + \epsilon)$ and every history $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ with $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$ and $\{a_t^j\}_{t \in [0, u]} = \{b_t^j\}_{t \in [0, u]}$ for all $j \in I$. Nonetheless, this approach is too restrictive in the current setting. For instance, in the tree harvesting problem (section 5.1), the optimal behavior of the agents is such that for every $\epsilon > 0$, there is positive conditional probability of the agents cutting trees in the time interval $(u, u + \epsilon)$ given

²⁰The proof is in appendix A.1.

that the agents cut trees at time u . This violates the uniform 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 $(u, u + \epsilon)$.

Another approach would be to apply inertia pathwise from each history. Formally, a strategy $\pi_i \in \Pi_i$ is said to be **pathwise inertial** if for every history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u and any realization $\{g_t\}_{t \in (u, T)}$ of the shock process after time u , there exists $\epsilon > 0$ such that $\pi_i(h_\tau) = z$ for all $\tau \in (u, u + \epsilon)$ and every history $h = \{g_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ with $\{a_t^j\}_{t \in [0, u]} = \{b_t^j\}_{t \in [0, u]}$ for all $j \in I$. In other words, for every history h_t up to an arbitrary time t and any realization of the shock process $\{s_\tau\}_{\tau \in (t, \infty)}$ after time t , there is required to exist $\epsilon > 0$ such that an agent does not move in the time interval $(t, t + \epsilon)$.²¹ This formulation enables arbitrarily quick responses to changes in the shock level. However, it does not in itself ensure measurability of the path of play in a stochastic setting. Some additional restriction is needed, but the iterative procedure that we use to prove theorem 2 would not be applicable if the traceability and frictionality restrictions were replaced with either version of inertia.²² Example 8 in section 4.1 describes a strategy profile satisfying both versions of inertia (as well as traceability and frictionality) while the resulting behavior is not measurable.

Irrespective of whether the former or the latter version of inertia is used, the set of strategies satisfying inertia is neither a subset nor a superset of the set of traceable and frictional strategies. As the result below states, even the pathwise formulation of inertia would imply traceability if applied to the current setting because inertia guarantees the existence of a history consistent with a given strategy at each time.

Proposition 1. *Any pathwise inertial strategy is traceable.*

The next example shows that frictionality is not satisfied by even the uniform specification of inertia because inertia allows for some strategies that induce infinitely many moves in a finite time interval.

²¹This condition could be modified so that zero probability events would never cause it to be violated. In particular, we could require that for every history h_t up to an arbitrary time t , there is conditional probability one given $\{s_\tau\}_{\tau \in [0, t]}$ of $\{s_\tau\}_{\tau \in (t, T)}$ being such that there exists $\epsilon > 0$ for which an agent does not move in the time interval $(t, t + \epsilon)$. However, such a redefinition would not affect the analysis of the applications that we present, particularly given that diffusion processes are modelled as a canonical Brownian motion as explained in footnote 38.

²²For this reason, it is unclear to us whether calculability in conjunction with either uniform or pathwise inertia would ensure that the action processes of the agents are measurable.

Example 3. (Uniform inertia does not imply frictionality) Suppose that $\bar{A}_i(h_t) = \{x, z\}$ for every $h_t \in H$. Consider the strategy in which agent i chooses x if and only if the current time t is such that $t = 1 - 1/2^m$ for some positive integer m . This strategy satisfies uniform inertia because for any time t , agent i does not move in the time interval (t, ∞) if $t \geq 1$, and agent i does not move in the time interval $(t, 1 - 1/2^l)$ if $t < 1$, where l is the least integer k such that $1 - 1/2^k > t$. This strategy violates frictionality because any history consistent with it is such that agent i moves at each of the infinitely many times $t < 1$ at which $t = 1 - 1/2^m$ for some positive integer m . \square

There also exist traceable and frictional strategies that do not satisfy even the pathwise definition of inertia. In particular, inertia eliminates some traceable and frictional strategies in which the response time of one agent varies with the actions of another agent. An example of such a strategy is given below.

Example 4. (Traceability and frictionality do not imply pathwise inertia) Suppose $I = \{1, 2\}$ and that $\bar{A}_i(h_t) = \{x, z\}$ for each $i \in I$ and every $h_t \in H$. Consider the strategy $\tilde{\pi}_i$ in which for any $t > 0$, agent i chooses x at time t if and only if agent i has not chosen x before time t , agent $-i$ has chosen x at time $t/2$, and agent $-i$ has not chosen x before time $t/2$. The strategy $\tilde{\pi}_i$ does not satisfy pathwise inertia because for any $\epsilon > 0$, there exists a history consistent with this strategy in which agent i chooses x at time $\epsilon/2$ and agent $-i$ chooses x at time $\epsilon/4$.

We argue that $\tilde{\pi}_i$ satisfies traceability and frictionality. Choose any history up to an arbitrary time u as well as any action path for agent $-i$. If there exists a least time \tilde{t}_{-i} such that agent $-i$ chooses x , if \tilde{t}_{-i} satisfies $2\tilde{t}_{-i} \geq u$, and if x is not chosen by agent i before time u , then a history h is consistent with $\tilde{\pi}_i$ from time u onwards if and only if h is such that x is chosen by agent i at time $2\tilde{t}_{-i}$ and z is chosen by agent i at every other time $t \geq u$. Otherwise, a history h is consistent with $\tilde{\pi}_i$ from time u onwards if and only if h is such that z is chosen by agent i at every time $t \geq u$. Since there exists a history consistent with $\tilde{\pi}_i$ at and after time u and in any such history x is chosen by agent i at most once during this interval, the strategy $\tilde{\pi}_i$ is traceable and frictional. \square

Another condition specified for deterministic models is admissibility in Simon and Stinchcombe (1989), which would ensure the existence of a unique path of play if appropriately applied to our framework. This condition comprises three restrictions

that we explain in what follows. The first is F1, which bounds the number of moves a strategy can induce. One approach to extending F1 to a stochastic setting is to impose the bound uniformly. Formally, a strategy $\pi_i \in \Pi_i$ satisfies **uniform F1** if for any time $\tilde{u} \in [0, T)$, there exists an integer m such that the following holds. Choose any history of the form $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T)}$ as well as any time $\hat{u} \in [0, \tilde{u}]$ such that h is consistent with π_i for $t \in [\hat{u}, \tilde{u}]$. Suppose that the set consisting of each time $t < \hat{u}$ such that $a_t^i \neq z$ is nonempty and has a maximum denoted by u . If $\pi_i(h_u) = a_u^i$ and there are at least m distinct values of $t < \hat{u}$ such that $a_t^i \neq z$, then $a_t^i = z$ for $t \in [\hat{u}, \tilde{u}]$.

An alternative is to specify the requirement pathwise. Formally, a strategy $\pi_i \in \Pi_i$ satisfies **pathwise F1** if for any realization $\{g_t\}_{t \in [0, T)}$ of the shock process as well as any time $\tilde{u} \in [0, T)$, there exists an integer m such that the following holds. Choose any history of the form $h = \{g_t, (a_t^j)_{j \in I}\}_{t \in [0, T)}$ as well as any time $\hat{u} \in [0, \tilde{u}]$ such that h is consistent with π_i for $t \in [\hat{u}, \tilde{u}]$. Suppose that the set consisting of each time $t < \hat{u}$ such that $a_t^i \neq z$ is nonempty and has a maximum denoted by u . If $\pi_i(h_u) = a_u^i$ and there are at least m distinct values of $t < \hat{u}$ such that $a_t^i \neq z$, then $a_t^i = z$ for $t \in [\hat{u}, \tilde{u}]$.

Intuitively, a strategy for agent i satisfies uniform F1 if for any finite interval of time, there exists an upper bound on the number of times agent i can move during that time interval. If, however, the bound has been reached or exceeded due to a deviation by agent i , then agent i is permitted to move at most one more time. The pathwise version of F1 is defined similarly except that the bound on the number of moves in each interval of time can also depend on the sample path of the shock process. The uniform version of F1 rules out some applications of interest. For example, the maximal equilibrium of the tree harvesting problem (section 5.1) is such that given any proper interval of time along with a nonnegative integer m , there is positive probability that trees are cut more than m times during this interval. Pathwise F1 is less restrictive, but does not cover every situation. For example, the equilibrium of the supply chain example (section 5.2) is such that given any proper interval of time and any price path along with a nonnegative integer m , there exists a history consistent with the strategy of the oil refinery in which it processes more than m batches of oil during this interval in response to oil being extracted more than m times.

The second criterion for admissibility is F2, which requires a strategy to in-

duce a well defined action path when the other agents do not move in the future. Formally, a strategy satisfies **F2** if the following holds. Choose any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ 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 $\{a_t^{-i}\}_{t \in [0, T]} \in \Gamma_{-i}(\{(b_t^{-i})_{t \in [0, u]})$ such that $a_t^j = z$ for all $t > u$ and $j \neq i$, there exists a unique action path $\{a_t^i\}_{t \in [0, T]} \in \Gamma_i(\{(b_t^i)_{t \in [0, u]})$ for which the history $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i at each $t \in [u, T)$, and this action path satisfies $\{a_t^i\}_{t \in [0, T]} \in \Xi_i(u)$. Intuitively, a strategy for agent i satisfies F2 if there exists a unique history that is consistent with the strategy when the other agents do not move in the future and this history is such that agent i moves only finitely many times in any finite time interval in the future.

The third requirement is F3, which imposes a special form of continuity on strategies. In particular, if one history up to a given time specifies the same moves as another history but at times that are just slightly greater, then an agent is required to take the same action at these two histories. This condition is violated by the solutions to all the applications in section 5 and the online appendix except for the ones in sections 5.3 and 5.4.²³ However, Simon and Stinchcombe (1989) note that this condition is not needed to prove the existence and uniqueness of an outcome consistent with a given strategy profile. It is used only to relate continuous-time outcomes to the limits of outcomes in discrete-time games.²⁴ For our purposes, we define **uniformly admissible** strategies as those that satisfy uniform F1 and F2 and **pathwise admissible** strategies as those that satisfy pathwise F1 and F2.

The result below shows that pathwise and hence also uniform F1 imply frictionality and that F2 implies traceability. Intuitively, suppose that agent i plays an admissible strategy π_i and that the other agents do not move in the future. Then F2 guarantees the existence of an action path for agent i that is consistent with π_i . Hence, traceability is satisfied. Moreover, both uniform and pathwise F1 ensure that π_i induces only finitely many moves in any finite interval of time. Hence, frictionality

²³For instance, consider two histories up to a given time specifying the same moves but at slightly different times, where one history but not the other is on the equilibrium path of play. If an equilibrium in grim-trigger strategies is played as in section 5.1, then the behavior prescribed by these strategies may differ between these two histories.

²⁴Since traceability and frictionality do not imply F3, the approach in Simon and Stinchcombe (1989) cannot easily be extended to relate the continuous-time equilibria that we consider to the equilibria of a discretized version of the game that is played on an increasingly fine grid. We believe this should not be considered a deficiency as continuous-time and discrete-time methods involve different models that do not necessarily have the same implications, and it is unclear a priori which approach is more plausible.

is satisfied.

Proposition 2. *If π_i satisfies pathwise F1, then π_i is frictional. If π_i satisfies F2, then π_i is traceable.*

Hence, the set of admissible strategies, regardless of whether defined pathwise or uniformly, is a subset of the set of traceable and frictional strategies. The converse of this result does not hold for two reasons. First, not only do both versions of property F1 apply on the path of play, but they also place important restrictions on the behavior of the agents off the path of play. In particular, if an agent violates the upper bound on the number of moves in an interval of time because of a deviation, then the agent can move only once more in that interval. As the example below illustrates, this rules out some simple strategies that require an agent always to move at specified times.

Example 5. (Strategy violating pathwise F1 off the path of play) Suppose that $\bar{A}_i(h_t) = \{x, z\}$ for every $h_t \in H$. Consider the following strategy $\tilde{\pi}_i$. Let t' and t'' be any two times with $t'' > t' > 0$. Agent i chooses x if and only if the time is currently t' or t'' . Note that $\tilde{\pi}_i$ is traceable and frictional. For any positive integer m , let h^m be a history in which agent i chooses x if and only if the time is equal to t'' or to $t' \cdot k/m$ for some positive integer $k \leq m$. The history h^m is consistent with $\tilde{\pi}_i$ at time t' , but there is no value of m such that $\tilde{\pi}_i(h_{t'}^m) = z$. Hence, strategy $\tilde{\pi}_i$ does not satisfy pathwise F1. \square

Second, both versions of F1 preclude some traceable and frictional strategies in which the number of moves by one agent varies with the timing of moves by another agent. Below is an example of a traceable and frictional strategy that violates pathwise F1, even when it is applied only on the path of play.

Example 6. (Traceability and frictionality do not imply pathwise F1) Suppose $I = \{1, 2\}$ and that $\bar{A}_i(h_t) = \{x, z\}$ for each $i \in I$ and every $h_t \in H$. Consider the following strategy $\tilde{\pi}_i$. Agent i does not choose x at any time in the interval $[0, 1]$. If there exists a positive integer m such that agent $-i$ chooses x at time $1/m$ and agent $-i$ chooses z at every time $t < 1$ such that $t \neq 1/m$, then agent i chooses x at time $t > 1$ if and only if there exists a positive integer $k \leq m$ such that $t = 1 + k/m$. Otherwise, agent i chooses z at each time $t > 1$. This strategy does not satisfy pathwise F1 because for every integer $m > 0$, there exists a history consistent with this

strategy in which agent $-i$ chooses x at time $1/m$ and agent i chooses x at m distinct times during the interval $(1, 2]$, meaning that there is no upper bound on the number of moves by agent i in this interval.

We argue that $\tilde{\pi}_i$ satisfies traceability and frictionality. Choose any history up to an arbitrary time u as well as any action path for agent $-i$. If there exists an integer $m > 0$ such that agent $-i$ chooses x at time $1/m$ and agent $-i$ does not choose x at any time $t < 1$ such that $t \neq 1/m$, then a history h is consistent with $\tilde{\pi}_i$ from time u onwards if and only if h is such that x is chosen by agent i at every time $t \geq u$ for which there exists a positive integer $k \leq m$ such that $t = 1 + k/m$ and z is chosen by agent i at every other time $t \geq u$. Otherwise, a history h is consistent with $\tilde{\pi}_i$ from time u onwards if and only if h is such that z is chosen by agent i at every time $t \geq u$. Since there exists a history consistent with $\tilde{\pi}_i$ at and after time u and in any such history x is chosen by agent i finitely many times during this interval of time, the strategy $\tilde{\pi}_i$ is traceable and frictional. \square

Traceability alone does not imply F2. Although example 2 satisfies F2, this restriction is violated by the following single-agent variant of that example. There exist traceable strategies that induce more than one path of play or an action path with infinitely many moves in a finite interval of time.

Example 7. (Traceable strategy violating F2) Suppose that $\bar{A}_i(h_t) = \{x, z\}$ for every $h_t \in H$. Consider the traceable strategy $\tilde{\pi}_i$ defined as follows. If the current time t is equal to $1/k$ for some positive integer k and agent i has chosen action x at time $1/m$ for every positive integer m greater than k , then agent i chooses x at time t . Otherwise, agent i chooses z . A history h is consistent with $\tilde{\pi}_i$ if agent i chooses z at all times or if agent i chooses x if and only if the current time t is equal to $1/k$ for some positive integer k . Noting that there exists more than one history consistent with $\tilde{\pi}_i$ (including one in which agent i moves infinitely many times in a finite interval of time), $\tilde{\pi}_i$ does not satisfy F2. \square

When taken together, traceability and frictionality imply F2, as the result below shows. Intuitively, suppose that agent i plays a traceable and frictional strategy π_i and that each of the other agents does not move in the future, which is also a traceable and frictional strategy. By theorem 1, there exists a unique action path for agent i that is consistent with π_i and this action path specifies only finitely many moves in

any finite time interval. Hence, F2 is satisfied. This result and its proof also hold if the frictionality requirement is relaxed to weak frictionality as defined in remark 2.

Proposition 3. *Any traceable and frictional strategy satisfies F2.*

Finally, there are uniformly and pathwise admissible strategies, like $\tilde{\pi}_1$ in example 8, that may not induce measurable behavior. A further restriction is necessary, as discussed in the next section.

4 Measurability of the Action Process

4.1 An Example

Another matter 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) and Simon and Stinchcombe (1989). 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 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 agents 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.

In continuous-time games with imperfect monitoring like Sannikov (2007), it is again unproblematic to define the strategy space so as to ensure the measurability of the actions taken by agents. The strategy of each agent can be specified as a stochastic process that is progressively measurable with respect to an exogenously given filtration. 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 8. (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 (Ω, \mathcal{F}, P) .²⁵ 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 $\tilde{t} > 0$, 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 (Ω, \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.

The following two simple ways of restricting the strategy space do not succeed in removing $\tilde{\pi}_1$ while retaining $\tilde{\pi}_2$, which is why we introduce a more complex procedure in section 4.2. 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 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 $\tilde{\pi}_2$. \square

²⁵For 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} .

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 $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$, 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 $(\{\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)\}_{t \in [0, T]})_{j \in I}$ 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 $\{s_t, [\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)]_{j \in I}\}_{t \in [0, T]}$ 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 $\tilde{k}_u = (\tilde{g}, b)$. 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.²⁶

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

²⁶For 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 .

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 $\pi_{-i} \in \Pi_{-i}^{TF}$.

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 8 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 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 $\pi_{-i} \in \Pi_{-i}^Q$.

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 (Ω, \mathcal{F}_1) to A_{-i} , and so the action by agent i at time 2 is a measurable function from (Ω, \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 agent. We then do the same for the second non- z action, the third non- z action, and so on.²⁷ Intuitively, if no agent has chosen a non- z action yet, then it is as if each agent's opponents are following the strategy of always choosing z , which belongs to the set Π_{-i}^Q . Hence, if an agent 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.²⁸

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 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 4. *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

²⁷The 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 .

²⁸The 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.

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 $v_i : (\times_{j \in I} A_j) \times S \rightarrow \mathbb{R}$ 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.²⁹

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 .³⁰ 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, if one wishes, it can be modeled by including the current time as part of the shock.³¹ For example, there may exist w_i such that $v_i[(a_\tau^j)_{j \in I}, s_\tau] = e^{-\rho\tau} w_i[(a_\tau^j)_{j \in I}]$. We also define:

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

²⁹The model can be extended to allow the utility function v_i to depend on the time t at which the payoff is calculated. Such an extension, which would allow for time-inconsistent behavior, would leave the analysis in this section essentially unchanged.

³⁰The 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. The model could be extended to allow the action profile in which all agents choose z to generate a flow utility for each agent that depends arbitrarily on the shock process. Such an extension would not affect the equilibria of the model in calculable strategies as long as the value of this flow utility does not depend on the past actions of the agents.

³¹In the applications in section 5 and the online appendix, the shock is formally defined to include calendar time as one of its components.

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

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. The space of sample paths of $\{s_t\}_{t \in [0,u]}$ is equipped with an arbitrary sigma-algebra \mathcal{A} . We impose the following **one-sided boundedness** condition.³² For any history $k_u = (\{g_t\}_{t \in [0,u]}, \{(b_t^j)_{j \in I}\}_{t \in [0,u]})$ up to time u and any strategy profile $\pi \in \times_{j \in I} \bar{\Pi}_j^{TF}$ such that the process $\xi_b^i(\pi)$ with $b = \{(b_t^j)_{j \in I}\}_{t \in [0,u]}$ 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 $o \in \{n, p\}$:

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

The conditional expectation is taken with respect to $\{s_t\}_{t \in (u,T)}$ given that $\{s_t\}_{t \in [0,u]} = \{g_t\}_{t \in [0,u]}$.³³

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

$$V_i(k_u, \pi) = V_i^p(k_u, \pi) + V_i^n(k_u, \pi). \quad (1)$$

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 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.

A potential issue is the nonuniqueness of the conditional expectation, which is a function from the sample space Ω to the set of sample paths of $\{s_t\}_{t \in [0,u]}$. Math-

³²All of the applications considered in section 5 and the online appendix satisfy this property.

³³The 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 $h = \{s_t, (a_t^i)_{i \in I}\}_{t \in [0,T]}$ such that $\{a_t^i\}_{t \in [0,T]} \in \Xi_i(u)$ for each $i \in I$ and such that $a_t^i \in \bar{A}_i(\{s_\tau\}_{\tau \in [0,t]}, \{(a_\tau^j)_{j \in I}\}_{\tau \in [0,t]})$ for all $t \geq [u, T]$ and each $i \in I$, we have either $\sup_{h \in \bar{H}} \sum_{t \in M_u(h)} \max\{0, v_i[(a_t^j)_{j \in I}, s_t]\} < \infty$ or $\inf_{h \in \bar{H}} \sum_{t \in M_u(h)} \min\{0, v_i[(a_t^j)_{j \in I}, s_t]\} > -\infty$ for every $i \in I$.

ematically, any two conditional expectation functions that differ on a set of probability measure zero may be regarded as the same random variable. This ambiguity is resolved for each strategy profile $\pi \in \times_{j \in I} \bar{\Pi}_j^C$ and every time u by choosing any conditional expectation function and holding this choice fixed. The expected payoff $V_i(k_u, \pi)$ can thereby be specified uniquely for each π and every k_u . This enables us to use a sure version of optimality instead of an almost sure notion when defining an equilibrium.

Having specified 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 $k_u = (\{g_t\}_{t \in [0, u]}, \{b_t^j\}_{j \in I, t \in [0, u]})$ 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 $\pi'_i \in \bar{\Pi}_i^C$.³⁴ Although proving the existence of an SPE is difficult in general because of the complexity of the strategy space that we consider, we show that an SPE exists in all of the applications that we analyze in section 5 and the online appendix.³⁵

5 Applications

The four examples in this section and the five additional ones in the online appendix illustrate the applicability of our methodology to a broad range of settings. The existence of a unique solution to each model facilitates the derivation of comparative statics, which demonstrates that our methodology enables substantive economic analysis. Although other methods could be used to analyze each specific application, *our point is that our method simultaneously covers all of the examples that we present, hence providing a unified approach to a variety of stochastic continuous-time games.* Moreover, even though the equilibria we study often involve Markov strategies, it is not sufficient, as explained in section 3.1, to restrict attention to Markov strategies for strategy spaces to be well defined. These applications also illustrate how to adapt our framework to accommodate particular situations that are seemingly beyond the scope of its assumptions.³⁶

³⁴The 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. This definition is useful in applications because it provides a simpler alternative to determining whether a potential deviating strategy is traceable, frictional, and calculable.

³⁵This issue is not specific to our approach: for example, Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) do not demonstrate the existence of equilibria in their general model.

³⁶For example, the action A in section 5.3 may be interpreted as a way of patching the model to allow the agents to move sequentially at the same time. In the online appendix, we also consider an

More specifically, in each application, we compare our restrictions with existing conditions that have been developed in the literature for a deterministic environment. We explain this comparison in detail for the first two examples, in which existing conditions are most problematic to apply. The uniform version of the inertia condition from Bergin and MacLeod (1993) is not satisfied in any application, and pathwise inertia is violated in section 5.2. As for the admissibility restriction of Simon and Stinchcombe (1989), uniform F1 does not hold except in sections 5.3 and 5.4 and in one example in the online appendix. Moreover, pathwise F1 is violated by the example in section 5.2 and also by one of the applications in the online appendix. While any traceable and frictional strategies have property F2, property F3 fails in every application except those in sections 5.3 and 5.4.

It is straightforward to confirm that the equilibrium strategies we study satisfy the calculability restriction in each application (verifying that they are traceable and frictional is simple as well). Doing so basically involves checking that each agent's behavior up to and including the time of the next move is progressively measurable given that the behavior of all the agents up to the current time is progressively measurable and the other agents are using quantitative strategies.³⁷ In addition, it is possible to demonstrate that a profile of calculable strategies is an SPE without having to determine whether potential deviating strategies satisfy all of the requirements for calculability (see also footnote 34).

While the applications we present have not been previously studied due to the lack of suitable techniques, the examples in sections 5.1 and 5.3 as well as those in the online appendix where the shock follows a diffusion process can be regarded as multiplayer extensions of the single-agent model of investment under uncertainty in McDonald and Siegel (1986).

5.1 Tree Harvesting Problem

There are $n \geq 2$ woodcutters harvesting trees in a common forest. The volume of forest resources evolves stochastically over time. It can increase due to natural growth,

example that involves flow payoffs and explain how it can be formulated within our framework.

³⁷For example, consider an equilibrium in grim-trigger strategies like in section 5.1 or two of the applications in the online appendix. Suppose that agent i plays its equilibrium strategy and the other agents play arbitrary quantitative strategies. Then the behavior of agent i will be progressively measurable up to and including the first time that one of the other agents deviates from its equilibrium strategy, after which the behavior of agent i simply involves not moving. Since this behavior is progressively measurable, it follows that the equilibrium strategy of agent i is calculable.

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 : $db_t = \mu dt + \sigma dz_t$.³⁸ 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.

At every moment of time $t \in [0, \infty)$, each woodcutter i decides whether to cut trees and, if so, a total amount of wood from the forest to harvest $e_t^i \in (0, q_t]$ as well as an amount of wood from the harvest to claim for itself $f_t^i > 0$. The woodcutters need help from each other in order to cut trees, so that the forest is cut if and only if all the woodcutters choose to harvest the same positive amount of wood. Suppose that at time t , every woodcutter i chooses to harvest the total amount of wood $d_t > 0$. Then the utility of each woodcutter i at that time is $(f_t^i / \sum_j f_t^j)(d_t - \kappa)$, which represents the difference between the benefit and cost of cutting trees. The amount of wood from the harvest d_t and the total cost of harvesting trees $\kappa > 0$ are rationed proportionally among the woodcutters according to the amount each of them claims.³⁹ If the forest is not cut at time t (which happens when not all the woodcutters choose to cut the same positive amount of trees), then each woodcutter 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 .⁴⁰ 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 an SPE 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 agent.

Let

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

³⁸In order to simplify the exposition throughout the paper, we consider the canonical version of Brownian motion, in which every sample path is continuous, thereby reducing the need to refer to zero probability events. For example, if it were possible for a sample path of the shock process to be discontinuous in the application here, then a maximal equilibrium in which the path of play is always as specified in the statement of proposition 5 might violate pathwise inertia, unless the definition of this condition were extended as described in footnote 21.

³⁹The analysis would not change qualitatively if each woodcutter who harvests trees were instead required to bear the full cost κ of cutting trees.

⁴⁰Formal definitions of histories and strategy spaces are provided in the online appendix.

where $\alpha = (-\mu + \sqrt{\mu^2 + 2\sigma^2\rho})/\sigma^2$, and W is the principal branch of the Lambert W function.⁴¹ Assume that $\kappa < \bar{x}$, and define $\hat{x} = \min\{x^*, \bar{x}\}$.

Proposition 5. *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.*

In a maximal SPE, the continuation payoff from deviation is zero, which is the minmax payoff of each agent.⁴² To understand the intuition for why it is optimal to cut the trees to zero, consider a strategy profile in which the trees are cut to the volume $y > 0$ whenever 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 3. 1. (Relationship to inertia) In a maximal SPE, each agent's strategy violates uniform inertia but satisfies pathwise inertia. Consider any history up to an arbitrary time t in which no agent has deviated in the past. For any $\epsilon > 0$, there is positive conditional probability that $q_\tau = \hat{x}$ for some $\tau \in (t, t + \epsilon)$, in which case the agents are required to cut trees. Thus, there does not exist $\epsilon > 0$ such that the agents do not move during the time interval $(t, t + \epsilon)$, meaning that uniform inertia is violated.⁴³ However, pathwise inertia holds given the continuity of the Brownian path. For any realization of the shock process $\{b_\tau\}_{\tau \in (t, \infty)}$ after time t , there exists $\epsilon > 0$ such that $q_\tau \neq \hat{x}$ for all $\tau \in (t, t + \epsilon)$, in which case the agents do not move during the time interval $(t, t + \epsilon)$.

⁴¹That is, $W(-e^{-1-\alpha\kappa})$ is the larger of the two values of y satisfying $-e^{-1-\alpha\kappa} = ye^y$.

⁴²There are multiple possibilities for off-path strategies, as described in the proof of proposition 5 in the online appendix.

⁴³Bergin and MacLeod (1993) also expand the strategy space to the completion of the set of inertia strategies, but this extension turns out not to apply for the same reason as why the uniform version of inertia does not. As shown in the online appendix of Kamada and Rao (2018) in the context of their model, the outcomes of some strategy profiles of interest satisfying our restrictions cannot be expressed as the limit of the outcomes generated by a Cauchy sequence of strategy profiles satisfying the uniform version of inertia, even when the approach in Bergin and MacLeod (1993) is extended to the stochastic case. We are currently working on a paper about the problem of generally defining the completion of the set of inertia strategies.

2. (Relationship to admissibility) In any maximal SPE, uniform F1 is violated because there is no upper bound on the number of times that trees may be harvested in any proper time interval. However, pathwise F1 would hold because for any realization of the shock process $\{b_\tau\}_{\tau \in [0,t]}$ up to time t , the agents are required to harvest trees only finitely many times during the time interval $[0, t]$. According to proposition 3, any traceable and frictional strategy has property F2. Nonetheless, property F3 in Simon and Stinchcombe (1989), which in some sense requires behavior to depend continuously on the history, is violated by a maximal equilibrium, in which agents use grim-trigger strategies where even a small deviation incurs a large punishment.

3. (ϵ -optimal strategies) In the current example, the complexities of defining strategy spaces in continuous time could be mitigated by focusing on a class of ϵ -optimal strategies that is relatively well behaved. For example, the strategies of the agents might be discretized so that agents can move if and only if the current time is a multiple of $\eta > 0$, where η can be made arbitrarily small so as to approximate the expected payoff to each player in a maximal SPE. Such strategies would satisfy both uniform and pathwise inertia. Nevertheless, this approach is subject to a problem of circularity. In order to determine whether a strategy is ϵ -optimal, the supremum expected payoff must be computed among a more general class of strategies such as those satisfying traceability, frictionality, and calculability. However, doing so requires defining expected payoffs for such a class of strategies in the first place. This circularity problem is not specific to this example but may arise in general when trying to restrict attention to a set of ϵ -optimal strategies without first having been able to compute the expected payoffs given a larger space of strategies. We will not repeat the same point in the subsequent examples where we solve for an optimal strategy profile.

4. (Comparative statics) The cutoff \hat{x} is the lesser of x^* , which represents the volume at which it is socially efficient to harvest wood, and \bar{x} , which represents the highest volume at which it is incentive compatible to do so.⁴⁴

The threshold x^* is increasing in μ , σ , and κ and decreasing in ρ . Intuitively, a higher value of μ or σ reduces the time required for the forest to grow by a

⁴⁴Proofs of the comparative statics results that follow are provided in the online appendix.

given amount, and a lower value of ρ makes it less costly to postpone the gains from harvesting wood, which makes it optimal to wait for the forest to grow to a larger volume before cutting it. A higher value of κ means that harvesting wood is more costly, so the associated benefits must be larger for it to be optimal to do so. In addition, the payoff maximizing behavior as represented by x^* does not depend on n because the expected payoff to each agent from any given profile of symmetric strategies is inversely proportional to the number of agents.

The threshold \bar{x} is increasing in μ and σ and decreasing in ρ and n . Intuitively, a higher value of μ or σ or a lower value of ρ raises the present value of the future benefits from cooperation by the agents in harvesting wood, and when n is lower, each agent receives a bigger share of these benefits. This makes it possible to enforce cooperation among the agents while harvesting a larger amount of wood on each cutting. In addition, the incentive compatibility constraint as represented by \bar{x} does not depend on κ because both the expected payoff to each agent in a maximal equilibrium and the supremum of the expected payoffs from a deviation are proportional to $\hat{x} - \kappa$.

5.2 Petroleum Supply Chain

Another application of our restrictions is search models, which involve the timing of exchanges. Opportunities to transact occur randomly over time, and an agent may delay trade in pursuit of a better opportunity. In the current example, one agent reacts to the behavior of another agent that is engaged in a search problem. The equilibrium strategy of the former agent in responding to the latter agent violates both pathwise inertia and pathwise admissibility (and hence it also violates uniform inertia and uniform admissibility).

There are two agents: an oil well W and an oil refinery R . Let p_t denote the price of oil at time t . The price of oil changes at the jump times of a Poisson process with parameter $\lambda > 0$. The resulting price is independently and identically distributed according to the cumulative distribution function G on the positive real line and has a finite expectation. Both agents discount the future at the rate $\rho > 0$.

The well contains a known quantity q of oil. Let x_t denote the amount of oil withdrawn before time t . At each moment in time $t \in [0, \infty)$, the well decides whether to extract oil and, if so, an amount $e_t \in (0, q - x_t]$ to remove. Let $c(x_t)$ represent

the marginal cost of extracting oil when x_t units have already been withdrawn, where $c : [0, q] \rightarrow \mathbb{R}_+$ is an increasing function. Correspondingly, the cost to the oil well at time t of extracting the quantity e_t when x_t is the amount that has already been extracted is given by $\int_{x_t}^{x_t+e_t} c(\xi)d\xi$.

Any oil extracted at time t is immediately transferred to the refinery, and the refinery pays the well the amount $p_t e_t$ at time t . Upon receiving the input at time t , the refinery must process it, which takes an amount of time $d(e_t)$, where the measurable function $d : (0, q] \rightarrow \mathbb{R}_{++}$ satisfies $\lim_{e \downarrow 0} d(e) = 0$. Let $f(e_t) \geq 0$ represent the associated processing costs discounted to time t . Having finished processing the oil received at time t , the refinery decides at each time $\tau \geq t + d(e_t)$ whether or not to deliver the output to a customer, if it has not done so already.

The refinery can be compensated in a variety of ways. In general, if the material that it started processing at time t is delivered to the customer at time $t' \geq t + d(e_t)$, then the refinery is paid the amount $y(e_t, p_t) > \exp[\rho d(e_t)][p_t e_t + f(e_t)]$ by the customer at time t' , where $y : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}_{++}$ is a measurable function. In the case where d is bounded, for example, the payment to the refinery may have the cost-plus form $\tilde{y} \cdot [p_t e_t + f(e_t)]$ with $\tilde{y} > \exp(\rho \bar{d})$, where $\bar{d} = \sup_{e \in (0, q]} d(e)$. If in addition G has a bounded support and $f(e)/e$ is bounded for $e \in (0, q]$, then the refinery may be paid a constant unit price $\hat{y} > \exp(\rho \bar{d})(\bar{p} + \bar{f})$, so that it receives $\hat{y} \cdot e_t$, where \bar{p} is the supremum of the support of G , and $\bar{f} = \sup_{e \in (0, q]} f(e)/e$.

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_W^C$, for the oil well and traceable, weakly frictional, calculable, and feasible strategies, denoted $\hat{\Pi}_R^C$, for the oil refinery.⁴⁵ Call this game with such strategy spaces the *supply chain model*. It is characterized by $(G, c, d, y, q, \lambda, \rho)$. The analysis in sections 3 and 4 implies that an SPE is well defined.

The equilibrium of the model is presented in the result below. Let $\varsigma : \mathbb{R} \rightarrow \mathbb{R}$ be

⁴⁵Formal definitions of histories and strategies are provided in the online appendix. In the formal specification of the model, W 's action at a given time is defined as a pair representing the respective amounts of oil extracted at the current time and in the past, and R 's action at a given time is defined as a function representing the payments it obtains from delivering the output produced from the oil received at each time. This way of defining action spaces enables the instantaneous utility of each agent to be expressed solely as a function of the actions of the agents at the current time, so that the expected payoffs in this application can be calculated using the general approach in section 4.3.

the function that solves the following reservation price equation for all $\kappa \in \mathbb{R}$:

$$\zeta(\kappa) = \kappa + \frac{\lambda}{\rho} \int_{\zeta(\kappa)}^{\infty} p - \zeta(\kappa) dG(p). \quad (2)$$

It is straightforward to show that there exists a unique value of $\zeta(\kappa)$ satisfying the above equation for each value of κ and that ζ is continuous, increasing, and surjective. Denoting the inverse of ζ by ζ^{-1} , let ζ_t denote the supremum of the set $c^{-1}([0, \zeta^{-1}(p_t)])$, where $c^{-1}([0, \zeta^{-1}(p_t)])$ is the preimage of the interval $[0, \zeta^{-1}(p_t)]$ under c . Let $\{\theta_{t,k}\}_{k=1}^l$ be a finite sequence of times at which oil was extracted from the well before time t , and for each index k , let $\xi_{t,k}$ be an indicator that is equal to 1 if the refinery has by time t processed the oil received at time $\theta_{t,k}$ but has not before time t delivered the output from the oil received at time $\theta_{t,k}$ and that is equal to 0 otherwise. If the refinery has received infinitely many batches of oil by time t , then assume that it is infeasible for the refinery to continue operating from time t onwards.

Proposition 6. *In the supply chain model with $(G, c, d, y, q, \lambda, \rho)$, the following hold in any SPE:*

1. *The strategy of W is such that:*
 - (a) *If $x_t < \zeta_t$, then W extracts the amount $e_t = \zeta_t - x_t$ at time t .*
 - (b) *If $x_t \geq \zeta_t$, then W does not extract oil at time t .*
2. *The strategy of R is such that for any k satisfying $\xi_{t,k} = 1$, R delivers the product of the oil extracted at time $\theta_{t,k}$ to the customer at time t .*

The well faces a search problem in deciding when to sell each unit of oil, and its equilibrium strategy has the reservation price property. In particular, each unit of oil is sold once the price reaches or exceeds a certain threshold, which is increasing in the marginal cost of extracting that unit of oil. The best response of the refinery is to deliver each batch of output to the customer as soon as the relevant input has been processed.

Remark 4. 1. (Relationship to inertia) R 's equilibrium strategy violates both the uniform and pathwise versions of inertia. In particular, consider any history up to an arbitrary time t and any $\epsilon > 0$. Since $\lim_{e \downarrow 0} d(e) = 0$, there exists $\chi > 0$

such that $d(\chi) < \epsilon/2$. If W extracts the amount χ at any time $u \in (t, t + \epsilon/2)$, then R delivers the resulting output to the customer at time $u + d(\chi) < t + \epsilon$. Thus, there is no $\epsilon > 0$ such that R does not move in the interval $(t, t + \epsilon)$.

W 's equilibrium strategy does not in general satisfy uniform inertia. For example, suppose that G has full support on \mathbb{R}_+ and that $c(0) > 0$. Consider any history up to an arbitrary time t such that $p_t < \varsigma[c(0)]$ and no oil has been extracted yet. For any $\epsilon > 0$, there is positive conditional probability that $p_\tau > \varsigma[c(q)]$ for some $\tau \in (t, t + \epsilon)$, in which case W extracts oil. Thus, there does not exist $\epsilon > 0$ such that W does not move during the time interval $(t, t + \epsilon)$. However, W has an equilibrium strategy that satisfies pathwise inertia. For any price path $\{p_\tau\}_{\tau \in (t, \infty)}$ after time t , there exists $\epsilon > 0$ such that $p_\tau = p_t$ for all $\tau \in (t, t + \epsilon)$, in which case W is not required to extract oil during the time interval $(t, t + \epsilon)$.

2. (Relationship to admissibility) R 's equilibrium strategy violates both uniform and pathwise F1. Choose any integer $m \geq 0$ and any $t > 0$. Since $\lim_{e \downarrow 0} d(e) = 0$, there exists $\chi > 0$ such that $d(\chi) < t/(m + 1)$. If W extracts the amount χ at any time τ for which there exists an integer $k \geq 0$ such that $\tau = kd(\chi)$, then R makes more than m deliveries during the time interval $[0, t]$. Hence, there is no upper bound on the number of moves by R up to a given positive time. Moreover, W 's equilibrium strategy does not generally satisfy uniform F1. For example, assume that G has full support on \mathbb{R}_+ and that c is positive, continuous, and increasing. Suppose that $p_0 < \varsigma[c(0)]$, and choose any integer $m \geq 1$ and any $t > 0$. There is positive probability that the price path is nondecreasing, has m discontinuities during the interval $[0, t]$, and is such that for every positive integer $k \leq m$, there exists $\tau_k \in [0, t]$ satisfying $\varsigma\{c[(k - 1)q/m]\} < p_{\tau_k} < \varsigma[c[kq/m]]$, in which case W extracts oil m times during the time interval $[0, t]$. Hence, the number of moves by W in a finite time interval is not bounded. However, W has an equilibrium strategy that satisfies pathwise F1. The number of times that W needs to extract oil during any finite time interval $[0, t]$ is bounded above by one plus the number of discontinuities up to time t in the price path $\{p_\tau\}_{\tau \in [0, \infty)}$, and the number of discontinuities is finite. F2 is satisfied by any traceable and weakly frictional strategy.

R 's equilibrium strategy violates the strong continuity property F3 because a

small difference in when W extracts a given amount of oil results in some difference in when R delivers the resulting output. W has an equilibrium strategy that satisfies F3. In particular, W 's behavior depends only on the amount of oil remaining and the current price of oil, so the specific times when oil was previously extracted or processed do not affect W 's behavior.

3. (Role of weak frictionality) While W 's equilibrium strategy is frictional, R 's is not. For example, suppose that for every positive integer k , W extracts the amount $q/2^k$ at time $1/k$. Then since $\lim_{e \downarrow 0} d(e) = 0$, R moves infinitely many times in the finite time interval $[0, \epsilon)$ for any $\epsilon > 0$. However, R 's equilibrium strategy is weakly frictional. This is because the number of times that R moves up to a given time t is bounded above by the number of times that W moves up to that time. Hence, if W moves only finitely many times in any finite interval of time, then the same is true of R .
4. (Comparative statics) It is straightforward to show from equation (2) that the reservation price $\varsigma[c(x_t)]$ of a unit of oil with marginal extraction cost $c(x_t)$ is nondecreasing in the rate parameter λ and nonincreasing in the discount rate ρ . If $\varsigma[c(x_t)]$ is less than the supremum of the support of G , then $\varsigma[c(x_t)]$ is strictly increasing in λ and strictly decreasing in ρ . Intuitively, a higher value of λ makes the price more volatile, raising the option value of waiting for the price to rise, and a higher value of ρ means that an agent is more impatient and less willing to delay sale until the price rises.

5.3 Entry Game between Competing Firms

Our restrictions on strategy spaces are useful 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. In the following example, investment is modelled as entry into a market, which is a common setting in industrial organization. Uniform inertia is violated, although this condition holds when defined pathwise.

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: I (“get in”) and A (“accommodate”). 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 F (“follow the other firm by entering”) or not. 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 agent $i = 1, 2$.⁴⁶ 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 an SPE is well defined.

Below we characterize Markov perfect equilibria, which we define as SPE in Markov strategies.⁴⁷ 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 7. *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 $-i$ has not.*

⁴⁶Formal definitions of histories and strategy spaces are provided in the online appendix.

⁴⁷There exist non-Markov 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.

- (a) If $\kappa_2 < c_t$, then both firms choose z .
- (b) If $c_t \leq \kappa_2$, then i chooses z while $-i$ 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 .

- Remark 5.** 1. (Relationship to inertia and admissibility) Uniform inertia is violated in any Markov perfect equilibrium, but pathwise inertia is satisfied. The example satisfies all of the requirements for admissibility, regardless of whether F1 is defined pathwise or uniformly. The details are discussed in the online appendix.
2. (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 b_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 .⁴⁸

5.4 Model of Ordering with a Deadline

Our methodology applies to environments like, for example, standard discrete-time repeated games and revision games as formalized by Kamada and Kandori (2020), in which each agent is restricted to change actions at discrete times or Poisson opportunities and the payoff to each agent is determined by the actions and shock levels

⁴⁸Proofs of these comparative statics results are provided in the online appendix.

at these moving times. The following is an example in which one agent can take a non- z action (i.e., has an opportunity to move) only at the arrival times of a Poisson process. Requiring each agent 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 agent who can move at any time in an interval of the real line, and restricting that agent 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.⁴⁹

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each agent $i = B, S$.⁵⁰ 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

⁴⁹It 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.

⁵⁰Formal definitions of histories and strategy spaces are provided in the online appendix.

implies that an SPE is well defined.⁵¹

Proposition 8. *In the finite-horizon ordering game with $(\lambda, T, \sigma, p, v)$, there exists a unique SPE. In this SPE, 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 SPE, B places an order at time t^* and S sells the good at the first Poisson hit after time t^* .⁵²

Remark 6. 1. (Relationship to inertia and admissibility) B 's equilibrium strategy satisfies both uniform and pathwise inertia, and S 's equilibrium strategy is pathwise inertial but not uniformly inertial. The example satisfies all of the requirements for admissibility, regardless of whether F1 is defined pathwise or uniformly. The details are discussed in the online appendix.

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 changes in the shock. There is probability zero that the shock discretely changes at time t^* , where the set of times at which the shock discretely changes is defined as the set of Poisson arrival times.

3. (Comparative statics) The existence of a unique equilibrium facilitates the derivation of comparative statics.⁵³ 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 ,

⁵¹In 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 the online appendix, the game here can be reinterpreted so as to conform with our general framework.

⁵²Note 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 $t > t^*$.

⁵³Proofs of these comparative statics results are provided in the online appendix.

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 at least 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.⁵⁴ With some restrictions, we show that any behavior on the path of play can be sustained in an SPE under a 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 provide conditions under which the converse is also true (section 6.5).

⁵⁴For 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.

6.1 Formulation of Expected Payoffs

We consider the assignment of expected payoffs to strategy profiles that induce non-measurable behavior.⁵⁵ For each agent $i \in I$, define a function $\chi_i : H \times \bar{\Pi}_i^{TF} \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$.

Choose any strategy profile $\pi \in \times_{i \in I} \bar{\Pi}_i^{TF}$. Let $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^i)_{i \in I}\}_{t \in [0, u]})$ 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) in section 4.3. Otherwise, the expected payoff to agent i at k_u is given by $U_i(k_u, \pi) = \chi_i(k_u, \pi)$.⁵⁶

Given a strategy space $\times_{i \in I} \hat{\Pi}_i \subseteq \times_{i \in I} \bar{\Pi}_i^{TF}$, 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 (see footnote 56), the standard one-shot deviation principle does not hold in general. Thus, it is crucial for the definition of SPE to consider deviations to a strategy in the entire subgame. The following example illustrates.

Example 9. (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 $\tilde{t} > 0$ such that $\{\omega \in \Omega : s_{\tilde{t}}(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space (Ω, \mathcal{F}, P) .⁵⁷ 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

⁵⁵We 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. We could also consider the assignment of expected payoffs to nondefined behavior, but 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.

⁵⁶The 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.

⁵⁷For example, let $\{s_t\}_{t \in [0, T]}$ be a standard Brownian motion, and suppose that the set consisting of every continuous function $c : [0, T] \rightarrow \mathbb{R}$ with $c(0) = 0$ and $c(\tilde{t}) \in \tilde{S}$ is not Wiener measurable.

that induces measurable behavior. However, there is no history up to a given time at which a one-shot deviation would increase the expected payoff of the agent. \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 $\Gamma(\hat{\Pi}, (\chi_i)_{i \in I})$, 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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^C, (\chi_i)_{i \in I})$. 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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^C, (\chi_i)_{i \in I})$ by $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$.

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 10. (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 (Ω, \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, $U_1(h_0, \pi^\emptyset) = 0 > -1 = U_1(h_0, \pi^{\tilde{S}})$ 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 11. (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 (Ω, \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 $(\{s_t\}_{t \in [0, u]}, \{(a_t^j)_{j \in I}\}_{t \in [0, u]}) = k_u$, $\{a_t^i\}_{t \in [0, T]} \in \Xi_i(u)$ for each $i \in I$, and $a_\tau^i \in \bar{A}_i(\{s_t\}_{t \in [0, \tau]}, \{(a_t^j)_{j \in I}\}_{t \in [0, \tau]})$ 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 $\zeta_i(k_u) = \inf_{h \in \bar{H}(k_u)} \sum_{\tau \in M_u(h)} v_i[(a_\tau^j)_{j \in I}, s_\tau]$ denote the greatest lower bound on the feasible payoffs to agent i at k_u .⁵⁸ In this section, we consider the possibility of letting $\chi_i(k_u, \pi) \leq \zeta_i(k_u)$ for all $\pi \in \times_{j \in I} \bar{\Pi}_j^{TF}$ and each $i \in I$. That is, nonmeasurable

⁵⁸Recall that $M_u(h)$ denotes the set of times at and after u where some agent moves under history h .

behavior is assigned an expected payoff no greater than the infimum of the set of feasible payoffs.⁵⁹

A motivation for this specification of payoffs is that assigning an extremely low payoff to nonmeasurable behavior disincentivizes agents from pursuing such behavior, thereby ensuring that the path of play is measurable, in which case the computation of expected payoffs is straightforward. In what follows, 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 12. (Mutual Nonmeasurability) 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 (Ω, \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 $\pi', \pi'' \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$.

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 13. (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 $\tilde{S} \subseteq S$ such that $\{\omega \in \Omega : s_1(\omega) \in \tilde{S}\}$ is not a measurable subset of the probability space (Ω, \mathcal{F}, P) . The utility functions satisfy $v_1[(a_1, a_2), s] \geq v_1[(z, a_2), s] \geq 0$ and

⁵⁹Similar results hold under an alternative definition in which $\zeta_i(k_u) = \inf_{\pi \in \bar{\Pi}_b^{TF}} V_i(k_u, \pi)$ at any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u , where $\bar{\Pi}_b^{TF}$ denotes the set consisting of any strategy profile $\pi \in \times_{j \in I} \bar{\Pi}_j^{TF}$ such that the process $\xi_b^i(\pi)$ with $b = \{(b_t^j)_{j \in I}\}_{t \in [0, u]}$ is progressively measurable for all $i \in I$.

$v_2[(a_1, a_2), s] \geq v_2[(a_1, z), s] \geq 0$ 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 agent has an incentive to deviate again for the same reason as in example 12. \square

The preceding example illustrates how any profile of actions at the null history can be implemented in equilibrium.⁶⁰ 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 9. *Let $|I| \geq 2$ and $T = \infty$. Consider the game $\Gamma(\times_{j \in I} \bar{\Pi}_j^{TF}, (\chi_j)_{j \in I})$ 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 (Ω, \mathcal{F}, P) . Choose any $\pi \in \times_{i \in I} \bar{\Pi}_i^{TF}$ such that for any profile of action paths $(\{b_t^i\}_{t \in [0, u]})_{i \in I}$ up to an arbitrary time $u > 0$, there exists with probability one some $t < u$ such that $\pi_i(\{s_\tau\}_{\tau \in [0, t]}, \{(b_\tau^j)_{j \in I}\}_{\tau \in [0, t]}) \neq b_t^i$ for some $i \in I$. Then there exists an SPE $\pi' \in \times_{i \in I} \bar{\Pi}_i^{TF}$ such that $(\{\phi_t^i(h_0, \{s_\tau\}_{\tau \in (0, T)}, \pi')\}_{t \in [0, T]})_{i \in I} = (\{\phi_t^i(h_0, \{s_\tau\}_{\tau \in (0, T)}, \pi)\}_{t \in [0, T]})_{i \in I}$ with probability one.*

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

⁶⁰In 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.

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.⁶¹

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 10. *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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$, then π is an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$ with $\chi_i(h_t, \bar{\pi}) \leq \zeta_i(h_t)$ for all $\bar{\pi} \in \times_{j \in I} \bar{\Pi}_j^{TF}$, 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 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.

⁶¹We 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 26 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 26, the main results in section 4 are valid under both definitions of the action process.

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 $\theta : \Omega \rightarrow [0, T]$ be a stopping time.⁶² 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. $\psi_i(\pi', \pi'', \theta)(\{s_t(\omega)\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]}) = \pi'_i(\{s_t(\omega)\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ 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. $\psi_i(\pi', \pi'', \theta)(\{s_t(\omega)\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]}) = \pi''_i(\{s_t(\omega)\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ 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.⁶³

⁶²That is, it satisfies $\{\omega \in \Omega : \theta(\omega) \leq t\} \in \mathcal{F}_t$ for all $t \in [0, T]$.

⁶³The synchronicity assumption is satisfied by the maximal equilibrium of the tree harvesting problem (section 5.1) 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 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.

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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$, then π is an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$.*

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 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. This is illustrated by the following example, in which there are two agents, two times at which the agents can move, and two possible actions at each of these times. We specify a strategy profile π^* that is an SPE for a given assignment of payoffs to nonmeasurable behavior but that is not measurably attachable because the action of each agent at time 2 is not measurable if one or both agents choose a non- z action at time 1. We explain that π_i^* is not calculable because agent i 's behavior at time 2 under π_i^* is not measurable if agent $-i$ follows the quantitative strategy $\tilde{\pi}_{-i}$ of always choosing a non- z action at time 1 and always choosing z at time 2.

Example 14. (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 (Ω, \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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$ 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. In the online appendix, we provide a formal definition of this concept and demonstrate that any synchronous strategy profile satisfying this notion of equilibrium 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 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. Existing techniques developed for a deterministic environment do not guarantee all of these properties in a stochastic setting, and they are not sufficient to cover some applications of interest. We also compared our method to an alternative approach in which specific payoffs are assigned to strategies inducing nonmeasurable behavior, and found a certain relationship between our method and this alternative. A variety of economic examples were presented to illustrate the applicability of our framework.

As we mentioned in the introduction, our methodology does not cover every possible situation. For example, although it applies to timing games in which agents choose when to take actions, 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 focused on the case in which past shocks and actions are perfectly observable. Allowing for imperfect information about them would enable the analysis of a wider range of applications. One example would be strategic experimentation with multiple agents, in which the actions of the agents affect the acquisition of information.⁶⁴ Other examples may include models of repeated games as well as bargaining with incomplete information as in Daley and Green (2012) and Ortner (2017) in which at least some actions are perfectly observed instantaneously. The entry game in section 5.3 could also be extended to allow firms to have private information about their own costs and benefits of entering the market. Finally, we restricted attention to pure strategies. With mixed strategies, measurability problems may arise even

⁶⁴For example, Keller and Rady (2015) consider a multi-agent bandit problem in continuous time but restrict agents to Markov strategies.

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 $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ 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 $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ 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 $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ 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 $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ 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 $h^p = \{s_t, (p_t^j)_{j \in I}\}_{t \in [0, T]}$ and $h^q = \{s_t, (q_t^j)_{j \in I}\}_{t \in [0, T]}$ 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 $\{q_t^j\}_{t \in [u, T)}$.⁶⁵ Let h_v^p be

⁶⁵To 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

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 $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T)}$ 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 $a_t^{j,0} = z$ 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 $t \in [u, T)$, where $d_t^{j,k} = a_t^{j,k}$ for $t \in [u, v_k]$ if $k \geq 1$. 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 $\{d_t^{j,k}\}_{t \in [u, T)}$ 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 $a_{v_k}^{j,k+1} = d_{v_k}^{j,k+1}$ for $t = v_k$. The history $\{s_t, (a_t^{j,k+1})_{j \in I}\}_{t \in [0, T)}$ is consistent with π_i for each $i \in I$ at every $t \in [u, v_k]$. Redefine k as $k + 1$. Return to Stage 1.

many times in X greater than $\inf(X)$ but less than $\inf(X) + \epsilon$. This implies that at least one of the $2|I|$ 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.

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, $\lim_{k \rightarrow \infty} a_t^{j,k}$ is well defined for all t . Consider the history $h_T = \lim_{k \rightarrow \infty} \{s_t, (a_t^{j,k})_{j \in I}\}_{t \in [0, T]}$. The history h_T is consistent with π_i for each $i \in I$ for $t \in [u, \lim_{k \rightarrow \infty} v_k)$.

Suppose that $\lim_{k \rightarrow \infty} v_k < T$. 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 $\{(d_t^{j, \infty})_{j \in I}\}_{t \in [0, T]}$ 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 $\{s_t, (d_t^{j, \infty}, \lim_{k \rightarrow \infty} (a_t^{i, k})_{i \neq j})\}_{t \in [0, T]}$ is consistent with π_j at each $t \in [u, T)$, where $d_t^{j, \infty} = \lim_{k \rightarrow \infty} a_t^{j, k}$ for $j \in I$ and $t \in [u, \lim_{k \rightarrow \infty} v_k)$. Since the algorithm does not terminate after a finite number of iterations, there exists $j \in I$ such that $\{d_t^{j, \infty}\}_{[0, T]} \notin \Xi_j(u)$, which implies π_j violates the frictionality assumption for some $j \in I$. Thus, $\lim_{k \rightarrow \infty} v_k = T$.

Hence, the history h_T is such that (i) $s_t = g_t$ for all $t \in [0, u]$, (ii) $\lim_{k \rightarrow \infty} a_t^{j, k} = b_t^j$ 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 $t \in [u, T)$. \square

Proof of Remark 2. Suppose that the first sentence of the statement of theorem 1 is modified so that $(\pi_j)_{j \in I}$ is a profile of traceable and weakly frictional strategies for which there exists l such that π_j is frictional for all $j \neq l$. We show that the remainder of the statement of theorem 1 continues to hold. The proof of theorem 1 still applies with the changes below.

The fourth sentence of the third paragraph should be replaced with the following. 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 $j \neq l$. Therefore, from weak frictionality, there is zero probability that $\{p_t^l\}_{t \in [u, T)}$ or $\{q_t^l\}_{t \in [u, T)}$ has infinitely many non- z actions in any finite time interval.

The second pair of sentences in the fourth paragraph should be replaced with the following. 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 $\{a_t^j\}_{t \in [0, T]} \notin \Xi_j(u)$ for some $j \neq l$ and for which the history $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ 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 $j \neq l$. Since no such $\{(a_t^j)_{j \in I}\}_{t \in [0, T]}$ can exist with positive probability, it follows

from the weak frictionality assumption that there is zero 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 $\{a_t^l\}_{t \in [0, T]} \notin \Xi_l(u)$ and for which the history $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i for each $i \in I$ at every $t \in [u, T]$.

In the second sentence of item 2 of the iterative procedure, frictionality should be replaced with weak frictionality. In the second to last paragraph, the second to last sentence should be replaced with the following. Since the algorithm does not terminate after a finite number of iterations, there exists $j \in I$ such that $d_t^{j, \infty} \notin \Xi_j(u)$. If $d_t^{j, \infty} \notin \Xi_j(u)$ for some $j \neq l$, then π_j violates the frictionality assumption for some $j \neq l$. Thus, it must be that $\{d_t^{j, \infty}\}_{t \in [0, T]} \in \Xi_j(u)$ for all $j \neq l$, so that $\{\lim_{k \rightarrow \infty} a_t^{j, k}\}_{t \in [0, T]} \in \Xi_j(u)$ for all $j \neq l$. Hence, if $d_t^{l, \infty} \notin \Xi_l(u)$, then π_l violates the weak frictionality assumption. \square

Proof of Proposition 1. Assume that the strategy π_i is pathwise inertial. For each $j \neq i$, let $\tilde{\pi}_j \in \Pi_j$ be the pathwise inertial strategy that prescribes z at every history up to an arbitrary time. Choose any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u as well as any realization $\{g_t\}_{t \in (u, T)}$ of the shock process after time u . From theorem 1 in Bergin and MacLeod (1993), there exists a unique profile $(\{\tilde{a}_t^j\}_{t \in [0, T]})_{j \in I}$ of action paths with $\{\tilde{a}_t^j\}_{t \in [0, T]} \in \Gamma_j(\{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ for each $j \in I$ such that the history $h = \{g_t, (\tilde{a}_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i and $\tilde{\pi}_{-i}$ at every $t \in [u, T]$. For each $j \neq i$, from the definition of $\tilde{\pi}_j$, the action path $\{\tilde{a}_t^j\}_{t \in [0, T]}$ satisfies $\tilde{a}_t^j = z$ for all $t \in [u, \infty)$. Combining the two preceding statements, π_i is traceable. \square

Proof of Proposition 2. Fix an arbitrary history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u . Assume that $\pi_i \in \Pi_i$ satisfies F2. Then given that $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$, there is conditional probability one that for any $\{a_t^{-i}\}_{t \in [0, T]} \in \Gamma_{-i}(\{(b_t^{-i})_{t \in [0, u]}\})$ such that $a_t^j = z$ for all $t > u$ and $j \neq i$, there exists an action path $\{a_t^i\}_{t \in [0, T]} \in \Gamma_i(\{(b_t^i)_{t \in [0, u]}\})$ for which the history $h = \{s_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i at each $t \in [u, T]$. Hence, π_i is traceable.

Choose any realization $\{g_t\}_{t \in (u, T)}$ of the shock process after time u and any time $u' > u$. Assume that $\pi_i \in \Pi_i$ satisfies pathwise F1. Then given any $\{a_t^{-i}\}_{t \in [0, T]} \in \Gamma_{-i}(\{(b_t^{-i})_{t \in [0, u]}\})$ for which there exists $\{a_t^i\}_{t \in [0, T]} \in \Gamma_i(\{(b_t^i)_{t \in [0, u]}\})$ such that the history $h = \{g_t, (a_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i at each $t \in [u, T]$, there exists an integer m such that there are at most m distinct values of $t \in [u, u']$ at which $a_t^i \neq z$. Hence, π_i is frictional. \square

Proof of Proposition 3. Assume that $\pi_i \in \Pi_i^{TF}$. For each $j \neq i$, let $\tilde{\pi}_j \in \Pi_j^{TF}$ be

a strategy that prescribes z at every history up to an arbitrary time. Choose any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to time u . The following hold with conditional probability one given that $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$. From the first item of theorem 1, there exists a unique profile $(\{\tilde{a}_t^j\}_{t \in [0, T]})_{j \in I}$ of action paths with $\{\tilde{a}_t^j\}_{t \in [0, T]} \in \Gamma_j(\{b_t^j\}_{t \in [0, u]})$ for each $j \in I$ such that the history $h = \{s_t, (\tilde{a}_t^j)_{j \in I}\}_{t \in [0, T]}$ is consistent with π_i and $\tilde{\pi}_{-i}$ at every $t \in [u, T]$. For each $j \neq i$, from the definition of $\tilde{\pi}_j$, the action path $\{\tilde{a}_t^j\}_{t \in [0, T]}$ satisfies $\tilde{a}_t^j = z$ for all $t \in [u, \infty)$. From the second item of theorem 1, $\{\tilde{a}_t^i\}_{t \in [0, T]} \in \Xi_i(u)$. Combining the three preceding statements, π_i has property F2. \square

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 $\tilde{k}_u(\{s_\tau\}_{\tau \in [0, u]}) = (\{s_\tau\}_{\tau \in [0, u]}, \{(b_\tau^j)_{j \in I}\}_{\tau \in [0, u]})$.

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 $h_t^p = (\{g_\tau^p\}_{\tau \in [0, t]}, \{(p_\tau^j)_{j \in I}\}_{\tau \in [0, t]})$ and $h_t^q = (\{g_\tau^q\}_{\tau \in [0, t]}, \{(q_\tau^j)_{j \in I}\}_{\tau \in [0, t]})$ 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 $t \in \Lambda(\{s_\tau\}_{\tau \in [0, T]})$.

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 $\phi_{\tilde{t}}^i[\tilde{k}_u(\{s_\tau\}_{\tau \in [0, u]}), \{s_\tau\}_{\tau \in (u, T)}, \pi] \neq z$ 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 $\pi'_{-i} \in \Pi_{-i}^{TF}$, the strategy profiles (π_i^2, π'_{-i}) and (π_i, π_{-i}^1) 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 $\xi_b^i(\pi_i, \pi_{-i}^1)$ is progressively measurable. Hence, $\pi_i^2 \in \Pi_i^Q$ 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 $\phi_{\tilde{t}}^i[\tilde{k}_u(\{s_\tau\}_{\tau \in [0, u]}), \{s_\tau\}_{\tau \in (u, T)}, \pi] \neq z$ for some $i \in I$. That is, $\Lambda^k(\{s_\tau\}_{\tau \in [0, T]})$ 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_i^k 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_i^k 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 $\pi_i^k \in \Pi_i^Q$ 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_i^1 at any time $t \in \Lambda^1(\{s_\tau\}_{\tau \in [0, T]})$. That is, agent i follows π_i^1 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_i^k 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_i^k 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 $\xi_b^i(\psi_i, \psi_{-i})$ is progressively measurable for $i \in I$. It will then follow that $\xi_b^i(\pi_i, \pi_{-i})$ 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 we have $\phi_{\tilde{t}}^i[\tilde{k}_u(\{s_\tau(\omega)\}_{\tau \in [0, u]}), \{s_\tau(\omega)\}_{\tau \in (u, T)}, \psi] \neq z$ 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.⁶⁶ 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 $i \in I$. \square

Proof of Proposition 4. 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 . \square

A.2 Proofs for Section 6

Proof of Proposition 9. 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 $k_u = (\{s_t\}_{t \in [0, u]}, \{(a_t^i)_{i \in I}\}_{t \in [0, u]})$ be any history up to an arbitrary time u . If $a_t^i = \pi_i(\{s_\tau\}_{\tau \in [0, t]}, \{(a_\tau^j)_{j \in I}\}_{\tau \in [0, t]})$ for each $i \in I$ and 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 $a_t^i \neq \pi_i(\{s_\tau\}_{\tau \in [0, t]}, \{(a_\tau^j)_{j \in I}\}_{\tau \in [0, t]})$, 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.

⁶⁶Given 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 $\mathcal{B}([0, v]) \times \mathcal{F}_v$.

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 $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to an arbitrary time $u > 0$. By assumption, there exists with probability one some $t < u$ such that $\pi_i(\{s_\tau\}_{\tau \in [0, t]}, \{(b_\tau^j)_{j \in I}\}_{\tau \in [0, t]}) \neq b_t^i$ for some $i \in I$. It follows that for each $i \in I$, the function $\phi_{n\tilde{t}}^i[\{s_t(\omega)\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]}, \{s_t(\omega)\}_{t \in (u, \infty)}, \pi']$ 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 10. Let π be an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$, and consider the game $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$ 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 $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ 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 $\pi''_i[\{s_\tau\}_{\tau \in [0, t]}, \{(d_\tau^j)_{j \in I}\}_{\tau \in [0, t]}] = \phi_t^i[\{s_\tau\}_{\tau \in [0, u]}, \{(b_\tau^j)_{j \in I}\}_{\tau \in [0, u]}, \{s_\tau\}_{\tau \in (u, T)}, (\pi'_i, \pi_{-i})]$ 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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$, 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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$. \square

Proof of Theorem 3. We first show that π is an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$, assuming that π is an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$ 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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$, 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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$, 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 $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$.

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 $\pi'_{-i} \in \Pi_{-i}^Q$. 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

$$\phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u)}_{j \in I} \right), \{s_\tau(\omega)\}_{\tau \in (u, T)}, \pi \right] \neq z$$

for some $e \in I$, and let $\bar{\Upsilon}^0(\omega)$ denote the set consisting of all $t \in [u, T)$ such that

$$\phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u)}_{j \in I} \right), \{s_\tau(\omega)\}_{\tau \in (u, T)}, (\pi_i, \pi'_{-i}) \right] \neq z$$

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

$$\phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u)}_{j \in I} \right), \{s_\tau(\omega)\}_{\tau \in (u, T)}, \psi[(\pi_i, \pi'_{-i}), \pi, \theta^{k-1}] \right] \neq z$$

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

$$\phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u)}_{j \in I} \right), \{s_\tau(\omega)\}_{\tau \in (u, T)}, (\pi_i, \pi'_{-i}) \right] \neq z$$

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 $\xi_b^i(\psi\{\psi[(\pi_i, \pi'_{-i}), \pi, \theta^{k-1}], \pi^z, \theta^k\})$ 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 $P[\{\omega \in \Omega : \lim_{k \rightarrow \infty} \theta^k(\omega) < \infty\}] \neq 0$. For each $\omega \in \Omega$, let $\Xi(\omega)$ denote the set consisting of all $t \in [u, T)$ such that

$$\phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u)}_{j \in I} \right), \{s_\tau(\omega)\}_{\tau \in (u, T)}, (\pi_i, \pi'_{-i}) \right] \neq z$$

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 $\lim_{k \rightarrow \infty} \theta^k(\omega) < \infty$, there exists $\tilde{t}(\omega) \in [u, \lim_{k \rightarrow \infty} \theta^k(\omega))$ such that

$$\phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u]} \right)_{j \in I}, \{s_\tau(\omega)\}_{\tau \in (u, T)}, (\pi_i, \pi'_{-i}) \right] = z$$

for each $e \in I$ and all $t \in [\tilde{t}(\omega), \lim_{k \rightarrow \infty} \theta^k(\omega))$.

For any $\omega \in E$ such that $\lim_{k \rightarrow \infty} \theta^k(\omega) < \infty$, 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,

$$\begin{aligned} \phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u]} \right)_{j \in I}, \{s_\tau(\omega)\}_{\tau \in (u, T)}, \psi[(\pi_i, \pi'_{-i}), \pi, \theta^{\tilde{k}(\omega)-1}] \right] \\ = \phi_t^e \left[\left(\{s_\tau(\omega)\}_{\tau \in [0, u]}, \{b_\tau^j\}_{\tau \in [0, u]} \right)_{j \in I}, \{s_\tau(\omega)\}_{\tau \in (u, T)}, (\pi_i, \pi'_{-i}) \right] = z \end{aligned}$$

for each $e \in I$, any $t \in (\theta^{\tilde{k}(\omega)-1}(\omega), \lim_{k \rightarrow \infty} \theta^k(\omega))$, and all $\omega \in E$ such that $\lim_{k \rightarrow \infty} \theta^k(\omega) < \infty$. This implies that $\theta^{\tilde{k}(\omega)}(\omega) \geq \lim_{k \rightarrow \infty} \theta^k(\omega)$ for all $\omega \in E$ such that $\lim_{k \rightarrow \infty} \theta^k(\omega) < \infty$, from which it follows that there is a set of nonzero measure consisting of $\omega \in E$ with $\lim_{k \rightarrow \infty} \theta^k(\omega) < \infty$ for which there exists $l \geq 1$ such that $\theta^l(\omega) > \lim_{k \rightarrow \infty} \theta^k(\omega)$. 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

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.

- DALEY, B., AND B. GREEN (2012): “Waiting for News in the Market for Lemons,” *Econometrica*, 80(4), 1433–1504.
- 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 (2020): “Revision Games,” *Econometrica*, 88(4), 1599–1630.
- KAMADA, Y., AND N. RAO (2018): “Sequential Exchange with Stochastic Transaction Costs,” Mimeo, University of California, Berkeley.
- KELLER, G., AND S. RADY (2015): “Breakdowns,” *Theoretical Economics*, 10(1), 175–202.
- MCCALL, J. J. (1970): “Economics of Information and Job Search,” *Quarterly Journal of Economics*, 84(1), 113–126.
- MCDONALD, R., AND D. SIEGEL (1986): “The Value of Waiting to Invest,” *The Quarterly Journal of Economics*, 101(4), 707–727.
- ORTNER, J. (2017): “Durable Goods Monopoly with Stochastic Costs,” *Theoretical Economics*, 12(2), 817–861.
- (2019): “A Continuous-Time Model of Bilateral Bargaining,” *Games and Economic Behavior*, 113(C), 720–733.
- 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.