Online Appendix to: "Strategies in Stochastic Continuous-Time Games"

Yuichiro Kamada[†] Neel Rao[‡]

March 27, 2025

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

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

B Formal Details for Section 5

B.1 Application in Section 5.1

B.1.1 Formal Definitions of Histories and Strategy Spaces

Choose any time $t \in [0, \infty)$ and Brownian motion $\{b_\tau\}_{\tau\in[0,t]}$ up to that time. Let $I = \{1, \ldots, n\}$ denote the set of woodcutters. A history up to time t is represented by $\big(\{b_\tau\}_{\tau\in[0,t]}, \{(a^i_\tau)_{i\in I}\}_{\tau\in[0,t)}\big)$, where $\{a^i_\tau\}_{\tau\in[0,t)}$ denotes the action path of woodcutter i up to time t with the action space being $\mathbb{R}^2_{++} \cup \{z\}$. The action $a^i_\tau = (e^i_\tau, f^i_\tau) \in \mathbb{R}^2_{++}$ means that woodcutter i seeks to harvest the amount e^i_τ and claim the amount f^i_τ at time τ. The action $a^i_\tau = z$ stands for choosing not to cut trees at that time.

The set of all histories up to an arbitrary time is denoted by H. Choose an arbitrary $h_t \in H$. Let X represent the set consisting of each time $\tau \in [0, t)$ for which there is no $i \in I$ such that $a^i_\tau = z$ and there exists $d_\tau > 0$ such that $e^i_\tau = d_\tau$ for all $i \in I$. If the set X has only finitely many elements, then let $\{t_k\}_{k=1}^K$ be the increasing sequence consisting of all the elements of X. For each $k \in \{1, \ldots, K\}$, define the volume of the forest right before the k^{th} cutting by $q_{t_k} = b_{t_k} - \sum_{l=1}^{k-1} r_{t_l}$, and define the amount harvested on the k^{th} cutting by $r_{t_k} = \min(q_{t_k}, d_{t_k})$. The volume of the forest at time t is given by $q_t = b_t - \sum_{l=1}^{k} r_{t_l}$. If the set X has only finitely many elements, then the feasibility constraint is $\bar{A}_i(h_t) = (0, q_t] \times \mathbb{R}_{++} \cup \{z\}$ for each $i \in I$. Otherwise, the feasibility constraint is simply $\bar{A}_i(h_t) = \{z\}$ for any $i \in I$.

The set of feasible strategies is for each $i \in I$:

$$\bar{\Pi}_i = \{\pi_i : H \to \mathbb{R}^2_{++} \cup \{z\} \mid \pi_i(h_t) \in \bar{A}_i(h_t) \text{ for all } h_t \in H\}.$$

The set of traceable, frictional, calculable, and feasible strategies can then be defined and is denoted by $\bar{\Pi}^C_i$ for woodcutter $i \in I$.

The shock process s_t is formally defined as a pair comprising the Brownian motion b_t and calendar time t. The instantaneous utility function v_i is specified for each $i \in I$ as $v_i[(a^i_\tau, a^{-i}_\tau), s_\tau] = 0$ if $a_j = z$ for some $j \in I$ or else if $e_j \neq e_k$ for some $j, k \in I$ and as

$$v_i[(a^i_\tau, a^{-i}_\tau), s_\tau] = \exp(-\rho\tau)\left(f^i_\tau \Big/ \sum_{j\in I} f^j_\tau\right)(d_\tau - \kappa)$$

if there is no $j \in I$ such that $a_j = z$ and there exists $d_\tau > 0$ such that $e^i_\tau = d_\tau$ for all $i \in I$.

B.1.2 Proofs

Proof of Proposition 5. The proof consists of three parts. We first assume the Markov property on the path of play and solve for the unique optimum, where a symmetric SPE is said to be Markov on the path of play if the action prescribed by each strategy at any history up to an arbitrary time on the path of play depends only on the volume q_t at that time. Second, we show that any maximal equilibrium must be Markov on the path of play. Third, we show that the supremum of the set of expected payoffs attainable in a symmetric SPE can be approximated arbitrarily closely by a symmetric SPE that is Markov on the path of play. These three results imply the existence of a maximal equilibrium.

Lemma 9. *For any profile $(n, \mu, \sigma, \kappa, \rho)$, the tree harvesting game has a symmetric SPE that is Markov on the path of play and weakly Pareto dominates any symmetric SPE that is Markov on the path of play. Moreover, on the path of play of any such SPE, the m^{th} cutting of trees occurs with probability one 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.*

Proof. Note first that at any history up to an arbitrary time, the minmax continuation payoff to each agent is zero, which can be obtained under the symmetric Markov strategy profile in which no woodcutter ever chooses to harvest trees. Hence, we restrict attention without loss of generality to strategy profiles in which after any deviation from the path of play, a symmetric Markov strategy profile is played in which no woodcutter ever chooses to harvest trees.

Let $U(b_t)$ denote the value of an asset that pays r at the first time the Brownian motion reaches $c \geq b_t$ when the current value of the Brownian motion is b_t. The function $U(b_t)$ satisfies the Bellman equation $\rho U(b_t) = \mathbb{E}(dU)$ subject to the boundary condition $U(c) = r$. Using Ito's lemma, the Bellman equation can be expressed as $\rho U(b_t) = \mu U^{'}(b_t) + \frac{1}{2}\sigma^2 U^{''}(b_t)$. It has the unique solution $U(b_t) = re^{\alpha(b_t - c)}$, where $\alpha = (-\mu + \sqrt{\mu^2 + 2\sigma^2\rho})/\sigma^2$.

In any symmetric SPE that is Markov on the path of play, there exist $y \geq 0$ and $z > y$ such that with probability one on the equilibrium path, the trees are cut if and only if the volume of the forest is currently $z > y$, with the volume being $y \geq 0$ after each cutting. Consider any symmetric SPE in grim-trigger strategies that is Markov on the path of play in which the equilibrium path is such that with probability one, the trees are cut if and only if the volume of the forest is currently $z > y$, with the volume being $y > 0$ after each cutting. There exists a symmetric SPE in grim-trigger strategies with a higher expected payoff to each agent in which the equilibrium path is such that with probability one, the trees are cut if and only if the volume of the forest is currently $z - y$, with the volume being 0 after each cutting. Noting that such an SPE is Markov on the path of play, we restrict attention to symmetric SPE in grim-trigger strategies for which there exists $x > 0$ such that with probability one on the equilibrium path, the trees are cut if and only if the volume of the forest is currently $x > 0$, with the volume of the forest being 0 after each cutting.

The expected payoff to each agent from playing such a strategy profile is given by $V(x) = [(x-\kappa)/n+V(x)]e^{-\alpha x}$, which yields $V(x) = (x-\kappa)/[n(e^{\alpha x}-1)]$. The optimization problem is to choose $x \geq \kappa$ so as to maximize $V(x)$ subject to the constraint $(x-\kappa)/n + V(x) \geq x - \kappa$. The left-hand side of the incentive constraint represents the expected payoff from following the prescribed strategy profile when cutting trees, and the right-hand side represents the payoff to an agent that unilaterally deviates in the limit as the amount of wood that it claims becomes arbitrarily large.

The derivative of $V(x)$ with respect to x is given by $V'(x) = \{e^{\alpha x}[1 - \alpha(x-\kappa)] - 1\}/[n(e^{\alpha x} - 1)^2]$, which satisfies $V'(\kappa) > 0$, $V'(\infty) < 0$, and $V'(y) < 0$ if $V'(x) \leq 0$ and $y > x$. Hence, the unconstrained maximization problem has a unique solution given by $V'(x) = 0$. The closed form expression for the value of x that solves $V'(x) = 0$ is $x^* = [1 + \alpha\kappa + W(-e^{-1-\alpha\kappa})]/\alpha$. In addition, the constraint can be expressed as $x \leq \bar{x}$, where $\bar{x} = \ln[n/(n-1)]/\alpha$.

Hence, the solution for x is the minimum of x^* and $\bar{x}$. ■

Lemma 10. *Up to zero probability events, any maximal equilibrium must be Markov on the path of play, with the path of play in a maximal equilibrium being unique.*

Proof. Suppose that there exists a maximal equilibrium. Then one can find $z > 0$ and $y < z$ such that there exists a maximal equilibrium in which with probability one, the first cutting occurs at the first time the volume reaches z, and the trees are cut to volume y at the first cutting. Since such a strategy profile is optimal, there exists a maximal equilibrium in which with probability one, the first cutting occurs at the first time the volume reaches z with the trees being cut to volume y, and the second cutting occurs at the second time the volume reaches z with the trees being cut to volume z. Continuing in this way, there exists a maximal equilibrium in which there is probability one that for any positive integer k, the k^{th} cutting occurs at the k^{th} time the volume reaches z, with the trees being cut to volume y at each cutting.

If $y > 0$, then such a strategy profile would be Pareto dominated by an SPE in which the path of play is such that with probability one, the trees are cut if and only if the volume of the forest is currently $z - y$, with the volume being 0 after each cutting. It follows that the volume of the forest after the first cutting is zero with probability one in any maximal equilibrium.

Let V denote the expected payoff to each agent when a maximal equilibrium is played starting at the null history. Then the continuation payoff to each agent after the first cutting on the equilibrium path should be V with probability one when a maximal equilibrium is played. Now consider the following optimization problem. The value at volume 0 of an asset that pays $V + (x-\kappa)/n$ at the first time that the volume reaches x is maximized with respect to x subject to the constraint that $V + (x-\kappa)/n \geq x - \kappa$. It is straightforward to show that this problem has a unique maximizer x'. Hence, up to zero probability events, a maximal equilibrium must be Markov on the equilibrium path up to the first cutting, which happens at the first time the volume reaches x'. We can iteratively apply a similar argument to each successive cutting on the equilibrium path to show that with probability one in any maximal equilibrium, the trees are cut if and only if the volume of the forest is currently x', with the trees being cut to the volume 0 at each cutting. ■

Lemma 11. *Given any symmetric SPE π, there exists a symmetric SPE that is Markov on the path of play and that yields no lower an expected payoff to each agent than does π.*

Proof. Let V denote the supremum of the expected payoffs to each agent that can be supported in a symmetric SPE. We show that there exists a symmetric SPE that is Markov on the path of play and that yields an expected payoff arbitrarily close to V, which proves the desired claim given lemma 10.

Let $V(q)$ denote the supremum of the expected payoffs that can be supported in a symmetric SPE at any history up to an arbitrary time in which the volume is currently q. Consider an asset $\mathcal{A}$ that pays $(x - q - \kappa)/n + V(q)$ at the first time that the volume reaches x. The value V is equal to the supremum of the value of this asset at the null history over $x \geq 0$ and $q \in [0, x]$ subject to the constraint that $(x-q-\kappa)/n+V(q) \geq x-\kappa$. Call this optimization problem $\mathcal{P}$. Note that the function $V(q)$ is continuous in q because for any $\gamma > 0$, one can find $\delta > 0$ such that there is probability greater than $1 - \gamma$ of the volume reaching q in a time interval of length γ when the current volume is $q - \delta$. We begin by proving the following claim.

Claim 12. *The value of asset $\mathcal{A}$ at volume c is bounded above by the sum of c/n and a constant.*

Proof. Consider a revised model that is identical to the tree harvesting game, except that the cost of cutting trees is zero if the volume has increased by at least the amount κ since right after the previous cutting. At any history up to an arbitrary time, the supremum in the tree harvesting game of the expected payoffs to each agent over all symmetric strategy profiles is no greater than the supremum in the revised model of the expected payoffs to each agent over all symmetric strategy profiles. In addition, the following implies that the latter supremum is no greater than the sum of $(c+2\kappa)/n$ and the value of an asset at the null history that for every positive integer p, pays $2\kappa/n$ when the Brownian motion reaches $p\kappa$ for the first time. This sum can be expressed as c/n plus a constant.

First, we observe that given any symmetric strategy profile in which trees are not harvested until the volume is at least $c + 2\kappa$, there exists in the revised model when the volume is currently c a symmetric strategy profile yielding a higher expected payoff to each agent in which trees are harvested before the volume reaches $c + 2\kappa$. To see this, choose any volume $l \geq c + 2\kappa$, and let m denote the greatest integer no larger than

$(l - c)/\kappa - 1$. Given any symmetric strategy profile in which the trees are cut at the next time the volume reaches l, there exists a symmetric strategy profile in the revised model yielding a higher expected payoff to each agent in which the trees are cut at the next time that the volume reaches $l - m\kappa$ and at the m successive times that the volume increases by the amount κ since right after the previous cutting.

Second, given any symmetric strategy profile in which the volume right after the next cutting is greater than zero, there exists a symmetric strategy profile yielding a higher expected payoff to each agent in which the volume right after the next cutting is zero. In particular, consider any symmetric strategy profile π in which the trees are cut at time u to a volume $z > 0$. There exists a symmetric strategy profile yielding a higher expected payoff to each agent at time u in which the trees are cut to zero at time u, the agents do not cut the trees at any time u' at which the total amount cut after time u up to and including time u' when playing π would be no greater than z, the agents cut the amount $y - z$ at the first time u' at which the total amount cut y after time u up to and including time u' when playing π would be greater than z, and the agents thereafter play strategy profile π behaving as if strategy profile π had always been played from time u onwards.

Letting S denote the supremum over all symmetric strategy profiles of the expected payoff to each agent at volume c in the revised model, the two preceding observations imply that for any $\epsilon > 0$, there exists a symmetric strategy profile yielding an expected payoff to each agent greater than $S - \epsilon$ in which the trees are harvested before the volume first reaches $c + 2\kappa$, the trees are always harvested again before the volume reaches 2κ, and the volume right after each cutting is zero. To compute an upper bound on the expected payoff to each agent when such a strategy profile is played, note that the utility of each agent at the first cutting is at most $(c + 2\kappa)/n$. Second, note that each cutting thereafter occurs when the volume is at least κ and yields a utility to each agent no greater than $2\kappa/n$. Hence, an upper bound on the continuation value after the first cutting can be computed by assuming that for every positive integer p, the trees are harvested when the Brownian motion reaches $p\kappa$ for the first time with the amount $2\kappa/n$ being harvested by each agent at every cutting. ∎

Since the upper bound on the value at volume c is less than $c - \kappa$ for c sufficiently high, the values of x satisfying the constraint are bounded above. The values of q satisfying the constraint are consequently bounded above. It is also straightforward to confirm that the values of x and q satisfying the constraint form a closed set. Since the objective function is continuous and the admissible values of x and q form a compact set, there exist values of x and q that achieve the supremum in problem $\mathcal{P}$. Let x^* and q^* denote these maximizers. Note that q^* cannot be equal to x^* because the contradiction $V(q^*) = V(q^*) - \kappa/n$ would otherwise result. There are two cases to consider. In the first case, the constraint in problem $\mathcal{P}$ is not binding. In the second case, the constraint in problem $\mathcal{P}$ is binding.

Consider the first case. Choose any $\epsilon > 0$. There exists a symmetric SPE ϕ_1 in grim-trigger strategies with the following properties that yields an expected payoff greater than $V - \epsilon$. With probability one, the first cutting on the equilibrium path occurs at the first time the volume reaches the threshold x^*, the trees are cut to the volume q^* at the first cutting, and the agents after the first cutting on the equilibrium path play a strategy profile that yields a continuation payoff W that does not depend on the history up to the time of the first cutting. Let Y denote the expected payoff that each agent receives with probability one at the first time the volume reaches q^* when playing strategy profile ϕ_1. Note that $V(q^*) - Y \leq V(q^*) - W$ because the behavior up to the first cutting when playing strategy profile ϕ_1 is the same as the behavior in problem $\mathcal{P}$.

Since $Y \geq W$, there exists a symmetric SPE ϕ_2 in grim-trigger strategies with the following properties that yields an expected payoff greater than $V - \epsilon$. With probability one on the equilibrium path, the first cutting occurs at the first time the volume reaches the threshold x^*, the second cutting occurs at the first time after the first cutting that the volume reaches the threshold x^*, the trees are cut to the volume q^* at the first and second cutting, and the agents after the second cutting play a strategy profile that yields a continuation payoff W that does not depend on the history up to the time of the second cutting. In particular, with probability one, the agents start by playing ϕ_1, and then after any history up to an arbitrary time on the equilibrium path after the first cutting, the agents play ϕ_1 behaving after the first cutting on the equilibrium path as if the volume q^* were reached for the first time.

Applying this procedure iteratively, one can show that there exists a symmetric SPE ϕ in grim-trigger strategies with the following properties that yields an expected payoff greater than $V - \epsilon$. There is probability one of the equilibrium path being such that for any positive integer m, the m^{th} cutting occurs at the first time after the $(m-1)^{\text{th}}$ cutting that the volume reaches the threshold x^* and the volume after each positive cutting is q^*, where the 0^{th} cutting is said to occur at time 0. This shows for the first case that there exists a symmetric SPE that is Markov on the path of play and yields an expected payoff arbitrarily close to V.

Consider the second case. Choose any $\epsilon > 0$. There exists a symmetric SPE ψ_1 in grim-trigger strategies with the following properties such that the expected payoff Y_1 at the first time the volume reaches q^* is greater than $V(q^*) - \epsilon$. With probability one on the equilibrium path, the first cutting occurs at the first time the volume reaches the threshold x_1, the trees are cut to q^* at the first cutting, and the agents after the first cutting play a strategy profile that yields a continuation payoff W_1 that does not depend on the history up to the time of the first cutting. Moreover, because the constraint in Problem $\mathcal{P}$ is binding, the threshold x_1 can be chosen such that $(x_1 - q^* - \kappa)/n + W_1 = x_1 - \kappa$ by choosing x_1 to maximize the expected payoff under ψ_1 given the continuation payoff W_1 and the volume q^* after the first cutting.

Applying such an argument to any subgame after the first cutting on the equilibrium path, there exists a symmetric SPE ψ_2' in grim-trigger strategies with the following properties such that the expected payoff at the first time the volume reaches q^* is greater than $V(q^*) - \epsilon$. With probability one on the equilibrium path, the first cutting occurs at the first time the volume reaches the threshold x_1, the second cutting occurs at the first time after the first cutting that the volume reaches a threshold x_2, the trees are cut to the volume q^* at the first and second cutting, and the agents after the second cutting play a strategy profile that yields a continuation payoff W_2 that does not depend on the history up to the time of the second cutting. Moreover, because the constraint in problem $\mathcal{P}$ is binding, the threshold x_2 can be chosen such that $(x_2 - q^* - \kappa)/n + W_2 = x_2 - \kappa$ by choosing x_2 to maximize the expected payoff under ψ_2' given the first threshold x_1, the continuation payoff W_2, and the volume q^* after the first and second cutting. Let Y_2 be the continuation payoff that each agent receives with probability one immediately after the first cutting on the equilibrium path when playing ψ_2'.

Note that $W_1 > W_2$ if $x_1 > x_2$, $W_1 < W_2$ if $x_1 < x_2$, and $W_1 = W_2$ if $x_1 = x_2$. In addition, $Y_1 > Y_2$ if $x_1 > x_2$, $Y_1 < Y_2$ if $x_1 < x_2$, and $Y_1 = Y_2$ if $x_1 = x_2$. If $x_2 > x_1$, then let $\psi_2 = \psi_2'$. If $x_2 \leq x_1$, then let ψ_2 be a strategy profile in which with probability one, the agents start by playing ψ_1, and then after any history on the equilibrium path after the first cutting, the agents play ψ_1 behaving as if the game just started after the first cutting on the equilibrium path.

Continuing in this way, one can show that there exists a symmetric SPE ψ in grim-trigger strategies with the following properties such that the expected payoff Y_1 at the first time the volume reaches q^* is greater than $V(q^*) - \epsilon$. There is probability one of the equilibrium path being such that for any positive integer m, the m^{th} cutting occurs at the first time after the $(m-1)^{\text{th}}$ cutting that the volume reaches the threshold x_m and the volume after each positive cutting is q^*, where the 0^{th} cutting is said to occur at time 0. Moreover, x_m is nondecreasing in m, and the continuation payoff Q_m that each agent receives with probability one after the m^{th} cutting is greater than $V(q^*) - \epsilon$.

Let y denote the limit of the sequence $\{x_m\}$. Consider a symmetric SPE ξ in which there is probability one of the equilibrium path being such that for any positive integer m, the m^{th} cutting occurs at the first time after the $(m-1)^{\text{th}}$ cutting that the volume reaches the threshold y and the volume after each positive cutting is q^*, where the 0^{th} cutting is said to occur at time 0. With probability one, the expected payoff R under strategy profile ξ at the first time the volume reaches q^* is no less than $V(q^*) - \epsilon$ because Q_m is greater than $V(q^*) - \epsilon$ for all m, where R is the limit of the sequence $\{Q_m\}$. Hence, the expected payoff under strategy profile ξ at the null history is no less than $V - \epsilon$. Moreover, the incentive constraint $(y - q^* - \kappa)/n + R \geq y - \kappa$ is satisfied because the incentive constraint $(x_m - q^* - \kappa)/n + Q_m \geq x_m - \kappa$ is satisfied for all m. This shows for the second case that there exists a symmetric SPE that is Markov on the path of play and yields an expected payoff arbitrarily close to V. ■

In a maximal SPE, there are multiple possibilities for off-path strategies, but in any off-path strategies, the continuation payoff from deviation is zero, which is the minmax payoff of each agent. One possibility for off-path strategies is for each agent never to move. Another possibility is for each woodcutter to cut trees at time t if and only if $q_t = \kappa$ and $q_\tau = 0$ for some time $\tau \in (\hat{t}, t)$, where $\hat{t}$ is the supremum of the set of times before t at which some agent moved. Yet another possibility is as follows. Let M be a positive integer, and let $c \in (0, \kappa)$. The agents do not move until reaching a time t such that $q_t = 0$. Subsequently, the m^{th} cutting for any $m \leq M$ occurs when the current time t is such that the volume reaches c for the m^{th} time, and the trees are cut to zero on each cutting. After the M^{th} cutting, the woodcutters play a maximal equilibrium. If there is any deviation from this path of play, then the agents never move. The values of M and c are chosen so that the ex ante expected payoff of each agent is equal to zero. □

Proof of Item 4 in Remark 2. It is straightforward to show that α is decreasing in μ and σ and increasing in ρ. Since $\bar{x}$ is decreasing in $\alpha > 0$, it follows that $\bar{x}$ is increasing in μ and σ and decreasing in ρ. Clearly, $\bar{x}$ is decreasing in n, and x^* is increasing in κ. We argue below that x^* is decreasing in α, from which it follows that x^* is increasing in μ and σ and decreasing in ρ.

Defining $\tilde{W}(\alpha) = W(-e^{-1-\alpha\kappa})$, the cutoff x^* can be expressed as follows:

$$x^* = 1/\alpha + \kappa + \tilde{W}(\alpha)/\alpha.$$

The partial derivative of x^* with respect to α is given by:

$$\partial x^*/\partial\alpha = [-1 + \alpha\tilde{W}'(\alpha) - \tilde{W}(\alpha)]/\alpha^2.$$

Differentiating $-1 + \alpha\tilde{W}'(\alpha) - \tilde{W}(\alpha)$ with respect to α yields:

$$\alpha\tilde{W}''(\alpha) + \tilde{W}'(\alpha) - \tilde{W}'(\alpha) = \alpha\tilde{W}''(\alpha),$$

where $\tilde{W}''(\alpha)$ is given by:

$$\tilde{W}''(\alpha) = e^{-2-2\alpha\kappa}\kappa^2[-e^{1+\alpha\kappa}W'(-e^{-1-\alpha\kappa}) + W''(-e^{-1-\alpha\kappa})],$$

which is negative because W is increasing and concave. It follows that $-1 + \alpha\tilde{W}'(\alpha) - \tilde{W}(\alpha)$ is decreasing in α. In order to demonstrate that $\partial x^*/\partial\alpha < 0$, it suffices to show that $\lim_{\alpha\downarrow 0} -1 + \alpha\tilde{W}'(\alpha) - \tilde{W}(\alpha) = 0$.

Using the formula $W'(\ell) = W(\ell)/\{\ell[1+W(\ell)]\}$ with $\ell = -e^{-1-\alpha\kappa}$, we obtain $\alpha\tilde{W}'(\alpha) = \big(\alpha W(\ell)/\{\ell[1 + W(\ell)]\}\big)\partial\ell/\partial\alpha$, which simplifies to $-\alpha\kappa\tilde{W}(\alpha)/[1 + \tilde{W}(\alpha)]$. Applying

L'Hôpital's Rule, we have $\lim_{\alpha\downarrow 0} -\alpha\kappa\tilde{W}(\alpha)/[1+\tilde{W}(\alpha)] = \lim_{\alpha\downarrow 0} -\kappa\tilde{W}(\alpha)/\tilde{W}'(\alpha)$, which equals 0 since $\lim_{\alpha\downarrow 0} \tilde{W}(\alpha) = -1$ and $\lim_{\alpha\downarrow 0} \tilde{W}'(\alpha) = \infty$. It follows that $\lim_{\alpha\downarrow 0} -1 + \alpha\tilde{W}'(\alpha) - \tilde{W}(\alpha) = 0$. □

B.2 Application in Section 5.2

B.2.1 Formal Definitions of Histories and Strategy Spaces

Choose any time $t \in [0,\infty)$ and any price process $\{p_t\}_{\tau\in[0,t]}$ up to that time. A history up to time t is represented by $\big(\{p_\tau\}_{\tau\in[0,t]}, \{(a^i_\tau)_{i\in\{W,R\}}\}_{\tau\in[0,t)}\big)$, where $\{a^W_\tau\}_{\tau\in[0,t)}$ and $\{a^R_\tau\}_{\tau\in[0,t)}$ respectively denote the action paths of the oil well and oil refinery up to time t. Agent W's action space is $\mathbb{R}_{++}\times\mathbb{R}_+\cup\{z\}$. In the case where $a^W_\tau \in \mathbb{R}_{++}\times\mathbb{R}_+$, the first element of a^W_τ, denoted by e_τ, represents the amount of oil extracted by the oil well at time τ, and the second element, denoted by x_τ, records the total amount extracted before time τ. The action z stands for not extracting any oil at time τ. Agent R's action space is $\mathbb{R}_+^{\mathbb{R}_+}\cup\{z\}$. In the case where a^R_τ is a function from $\mathbb{R}_+$ to itself, $a^R_\tau(\tau') > 0$ represents the payment to the oil refinery at time τ from delivering the output produced from the oil received at time τ', where $a^R_\tau(\tau') = 0$ indicates that no such delivery was made by the oil refinery at time τ. The action $a^R_\tau = z$ means that the oil refinery does not deliver any output at time τ.

The set of all histories up to an arbitrary time is denoted by H. Choose any $h_t \in H$. Let X represent the set consisting of each time $\tau \in [0,t)$ such that $a^W_\tau \neq z$. If the set X has infinitely many elements, then the feasibility constraints are simply $\bar{A}_W(h_t) = \bar{A}_R(h_t) = \{z\}$. Consider the case where the set X has only finitely many elements, and let $\{t_k\}_{k=1}^K$ be the sequence consisting of all the elements of X. The set of W's feasible actions is $\bar{A}_W(h_t) = (0, q - x_t]\times\{x_t\}\cup\{z\}$, where $x_t = \sum_{k=1}^K e_{t_k}$. The set of R's feasible actions $\bar{A}_R(h_t)$ is such that $a^R_t \in \bar{A}_R(h_t)$ if and only if $a^R_t = z$ or a^R_t satisfies the following. Choose any time $\tau \in [0,\infty)$. If there exists k such that $t_k = \tau$ and $\tau + d(e_\tau) \le t$ and there is no $t' < t$ such that $a^R_{t'}(\tau) > 0$, then $a^R_t(\tau) \in \{0, y(p_\tau, e_\tau)\}$. Otherwise, $a^R_t(\tau) = 0$. In addition, $a^R_t(\tau) > 0$ for some $\tau \in [0,t)$.

The sets of feasible strategies are:

$$\begin{aligned}\bar{\Pi}_W &= \{\pi_W : H \to \mathbb{R}_{++}\times\mathbb{R}_+\cup\{z\} \mid \pi_W(h_t)\in\bar{A}_W(h_t) \text{ for all } h_t\in H\}\\ \bar{\Pi}_R &= \{\pi_R : H \to \mathbb{R}_+^{\mathbb{R}_+}\cup\{z\} \mid \pi_R(h_t)\in\bar{A}_R(h_t) \text{ for all } h_t\in H\}\end{aligned}.$$

For agent W, the set of traceable, frictional, calculable, and feasible strategies can be defined and is denoted by $\bar{\Pi}^C_W$. For agent R, the set of traceable, weakly frictional, calculable, and feasible strategies can be defined and is denoted by $\hat{\Pi}^C_R$.

The shock process s_t is formally defined as a pair comprising the price p_t and calendar time t. The instantaneous utility function v_i is specified as follows for $i = W$:

$$v_W[(a_\tau^W, a_\tau^R), s_\tau] = \begin{cases} [p_\tau e_\tau - \int_{x_\tau}^{x_\tau + e_\tau} c(\xi) d\xi] \exp(-\rho\tau) & \text{if } a_\tau^W \neq z \\ 0 & \text{if } a_\tau^W = z \end{cases},$$

and as follows for $i = R$:

$$v_R[(a_\tau^W, a_\tau^R), s_\tau] = \begin{cases} [\sum_{\{\tau' : a_\tau^R(\tau') > 0\}} a_\tau^R(\tau') - p_\tau e_\tau] \exp(-\rho\tau) & \text{if } a_\tau^W \neq z \text{ and } a_\tau^R \neq z \\ [\sum_{\{\tau' : a_\tau^R(\tau') > 0\}} a_\tau^R(\tau')] \exp(-\rho\tau) & \text{if } a_\tau^W = z \text{ and } a_\tau^R \neq z \\ -p_\tau e_\tau \exp(-\rho\tau) & \text{if } a_\tau^W \neq z \text{ and } a_\tau^R = z \\ 0 & \text{if } (a_\tau^W, a_\tau^R) = (z, z) \end{cases}.$$

B.2.2 Proofs

Proof of Proposition 6. Consider the problem faced by an oil well deciding when to sell a single unit of oil whose extraction cost is κ where the price evolves according to the stochastic process $\{p_t\}_{t \in [0,\infty)}$. This is a basic search problem in continuous time.[1] Letting $B(\kappa)$ denote the expected payoff to an oil well that chooses to retain the oil at the current time, the optimal policy of the oil well is to extract the oil at time t if $p_t > B(\kappa) + \kappa$, to retain the oil at time t if $p_t < B(\kappa) + \kappa$, and either if $p_t = B(\kappa) + \kappa$. The solution is characterized by the Bellman equation:

$$\rho B(\kappa) = \lambda \int_{-\infty}^{\infty} \max[p - B(\kappa) - \kappa, 0] dG(p). \tag{3}$$

Defining the reservation price $\varsigma(\kappa) = B(\kappa) + \kappa$, the preceding equation can be expressed as:

$$\varsigma(\kappa) = \kappa + \frac{\lambda}{\rho} \int_{\varsigma(\kappa)}^{\infty} p - \varsigma(\kappa) dG(p). \tag{4}$$

It is straightforward to show that there exists a unique value of $\varsigma(\kappa)$ satisfying the above equation and that $\varsigma(\kappa)$ is increasing and continuous in κ.

For any $\kappa \in \mathbb{R}_+$, let $S(\kappa)$ be the supremum of the set $\{e/q : c(e) \leq \kappa\}$ if $c(0) \leq \kappa$, and let $S(\kappa) = 0$ otherwise. It follows from the analysis so far that the optimal policy of an oil well that has a measure q of oil with extraction cost distributed according to the cdf S is to extract at time t any remaining unit of oil with extraction cost κ satisfying $\varsigma(\kappa) < p_t$, to retain at time t any remaining unit of oil with extraction cost κ satisfying $\varsigma(\kappa) > p_t$, and either if $\varsigma(\kappa) = p_t$. Hence, the equilibrium strategy of the oil well in the supply chain model is as specified in the statement of the proposition.

For any k such that $\xi_{t,k} = 1$, consider the oil extracted at time $\theta_{t,k}$. If the refinery

[1] Rogerson, Shimer, and Wright (2005) present a similar problem in their review of search models of the labor market.

delivers the resulting output at time $t' \geq t$, then its payoff at time t from the delivery is $\exp[-\rho(t'-t)] \cdot y(p_{\theta_{t,k}}, e_{\theta_{t,k}}) > 0$. Since this expression is decreasing in t', the refinery maximizes its payoff by delivering the output immediately. Hence, the equilibrium strategy of the oil refinery is as specified in the statement of the proposition. $\square$

B.3 Application in Section 5.3

B.3.1 Formal Definitions of Histories and Strategy Spaces

Choose any time $t \in [0, \infty)$ and cost process $\{c_\tau\}_{\tau\in[0,t]}$ up to that time. A history up to time t is represented by $\big(\{c_\tau\}_{\tau\in[0,t]}, \{(a_\tau^i)_{i\in\{1,2\}}\}_{\tau\in[0,t)}\big)$, where $\{a_\tau^i\}_{\tau\in[0,t)}$ denotes the action path of firm $i \in \{1, 2\}$ up to time t with the action space being $\{I, A, F, z\}$.

The set of all histories up to an arbitrary time is denoted by H. We partition it as follows.

1. Let $H^{\text{no,no}}$ be the set consisting of every history up to any time t that has the form $\big(\{c_\tau\}_{\tau\in[0,t]}, \{(a_\tau^i)_{i\in\{1,2\}}\}_{\tau\in[0,t)}\big)$ where $a_\tau^i = z$ for each $i = 1, 2$ and all $\tau \in [0, t)$.
2. Let $H^{\text{yes,no}}$ be the set consisting of every history up to any time t that has the form $\big(\{c_\tau\}_{\tau\in[0,t]}, \{(a_\tau^i)_{i\in\{1,2\}}\}_{\tau\in[0,t)}\big)$ where there exists $\tau' \in [0, t)$ such that $(a_\tau^1, a_\tau^2) = (z, z)$ for all $\tau \in [0, t) \setminus \{\tau'\}$ and $(a_{\tau'}^1, a_{\tau'}^2)$ is (I, z), (I, A), or (A, z).
3. Let $H^{\text{no,yes}}$ be the set consisting of every history up to any time t that has the form $\big(\{c_\tau\}_{\tau\in[0,t]}, \{(a_\tau^i)_{i\in\{1,2\}}\}_{\tau\in[0,t)}\big)$ where there exists $\tau' \in [0, t)$ such that $(a_\tau^1, a_\tau^2) = (z, z)$ for all $\tau \in [0, t) \setminus \{\tau'\}$ and $(a_{\tau'}^1, a_{\tau'}^2)$ is (z, I), (A, I), or (z, A).
4. Let $H^{\text{yes,yes}}$ be the set consisting of every history up to any time t that has the form $\big(\{c_\tau\}_{\tau\in[0,t]}, \{(a_\tau^i)_{i\in\{1,2\}}\}_{\tau\in[0,t)}\big)$ where either of the following holds:
   (a) There exists $\tau' \in [0, t)$ such that $(a_{\tau'}^1, a_{\tau'}^2) \in \{(I, I), (A, A)\}$, and $(a_\tau^1, a_\tau^2) = (z, z)$ for all $\tau \in [0, t)$ with $\tau \neq \tau'$.
   (b) There exist $\tau', \tau'' \in [0, t)$ with $\tau' < \tau''$ such that $(a_\tau^1, a_\tau^2) = (z, z)$ for all $\tau \in [0, t)$ with $\tau \notin \{\tau', \tau''\}$ and either of the following holds:
      i. $(a_{\tau'}^1, a_{\tau'}^2) \in \{(I, z), (I, A), (A, z)\}$ and $(a_{\tau''}^1, a_{\tau''}^2) = (z, F)$.
      ii. $(a_{\tau'}^1, a_{\tau'}^2) \in \{(z, I), (A, I), (z, A)\}$ and $(a_{\tau''}^1, a_{\tau''}^2) = (F, z)$.

The feasibility constraints are as follows. For firm 1,

$$\bar{A}_1(h_t) = \begin{cases} \{I, A, z\} & \text{if } h_t \in H^{\text{no,no}} \\ \{F, z\} & \text{if } h_t \in H^{\text{no,yes}} \\ \{z\} & \text{otherwise} \end{cases}.$$

For firm 2,

$$\bar{A}_2(h_t) = \begin{cases} \{I, A, z\} & \text{if } h_t \in H^{\text{no,no}} \\ \{F, z\} & \text{if } h_t \in H^{\text{yes,no}} \\ \{z\} & \text{otherwise} \end{cases}.$$

The set of feasible strategies is for each $i = 1, 2$:

$$\bar{\Pi}_i = \{\pi_i : H \to \{I, A, F, z\} \mid \pi_i(h_t) \in \bar{A}_i(h_t) \text{ for all } h_t \in H\}.$$

The set of traceable, frictional, calculable, and feasible strategies can then be defined and is denoted by $\bar{\Pi}_i^C$ for firm $i = 1, 2$.

The shock process s_t is formally defined as a pair comprising the entry cost c_t and calendar time t. The instantaneous utility function v_i is specified as follows for each firm $i = 1, 2$:

$$v_i[(a_\tau^1, a_\tau^2), s_\tau] = \begin{cases} 0 & \text{if } (a_\tau^i, a_\tau^{-i}) \in \{(z,z), (z,I), (z,A), (z,F), (A,I)\} \\ (b_1 - c_\tau)e^{-\rho\tau} & \text{if } (a_\tau^i, a_\tau^{-i}) \in \{(I,z), (I,A), (A,z)\} \\ (b_2 - c_\tau)e^{-\rho\tau} & \text{if } (a_\tau^i, a_\tau^{-i}) \in \{(F,z), (I,I), (A,A)\} \end{cases}.$$

B.3.2 Proofs

Proof of Proposition 7. Define a parameter $\beta = \frac{1}{2} - \mu/\sigma^2 - \sqrt{(\mu/\sigma^2 - \frac{1}{2})^2 + 2\rho/\sigma^2} < 0$. For $c > 0$, let κ_2 be the value of $\kappa > 0$ that maximizes the expression $(b_2 - \kappa)(c/\kappa)^\beta$, which for $\kappa \leq c$ is the value of an asset that pays $b_2 - \kappa$ at the first time the cost reaches κ when the current cost is c. The maximizer is $\kappa_2 = [\beta/(\beta - 1)]b_2$, and the maximized value is $b_2^{1-\beta}c^\beta(-\beta)^{-\beta}(1-\beta)^{\beta-1}$.

Next let κ_1 be the value of $\kappa > \kappa_2$ that solves the equation $b_1 - \kappa = (b_2 - \kappa_2)(\kappa/\kappa_2)^\beta = b_2^{1-\beta}\kappa^\beta(-\beta)^{-\beta}(1-\beta)^{\beta-1}$. The left-hand side is bigger than the right-hand side in the limit as κ goes to κ_2, and the right-hand side is bigger than the left-hand side in the limit as κ goes to ∞. The derivative of the left-hand side minus the right-hand side with respect to κ is given by $-1 + [-\beta/(1-\beta)]^{1-\beta}(b_2/\kappa)^{1-\beta}$, which is decreasing in κ. Hence, there exists a unique value of κ that satisfies the preceding equation.

Now we characterize the SPE. First consider any history up to an arbitrary time t at which firm $i \in \{1, 2\}$ is the only firm not in the market. In any SPE, action F will be chosen by firm i if and only if $c_t \leq \kappa_2$.

Next consider the case in which neither firm has yet entered the market. In an SPE, the firms will both choose I or both choose A if the history up to the current time t is such that $c_t \leq \kappa_2$. Moreover, there cannot be an SPE in which a firm chooses I or A at a history up to a given time t satisfying $c_t > \kappa_1$. A firm that enters the market at such a history could increase its expected payoff by deviating to a strategy in which it chooses z whenever the cost is currently greater than κ_2 and it enters whenever the cost is currently no greater than κ_2. Finally, there cannot be an SPE in which the firms both choose I or both choose A at a history up to a given time t satisfying $c_t > \kappa_2$. A firm that enters the market at such a history could increase its expected payoff by deviating to a strategy in which it chooses z whenever the cost is currently greater than κ_2 and it enters whenever the cost is currently no greater than κ_2.

Now suppose that the firms are playing an SPE in Markov strategies. Consider the set of histories up to any time in which no firm has entered yet. For $i \in \{1, 2\}$, let ξ_i denote the maximum cost c_t at which agent i chooses I or A. Such a value of the cost exists due to the traceability assumption.

It must be that $\xi_1 = \kappa_1$ or $\xi_2 = \kappa_1$. Suppose to the contrary that $\xi_i < \kappa_1$ for each $i \in \{1, 2\}$. If $\xi_i < \xi_{-i}$, then there exists $\chi \in (\xi_{-i}, \kappa_1)$ such that firm i could increase its expected payoff by deviating and choosing I whenever the current value of the cost is χ. If $\xi_1 = \xi_2$ and firm i chooses A whenever the cost is currently ξ_i and there is no prior entry, then there exists $\chi \in (\xi_i, \kappa_1)$ such that firm i could increase its expected payoff by deviating and choosing I whenever the cost is currently χ.

It must further be that $\xi_1 = \xi_2 = \kappa_1$. Suppose to the contrary that $\xi_i = \kappa_1$ but $\xi_{-i} < \kappa_1$. There exists $\chi \in (\xi_{-i}, \kappa_1)$ such that firm i could increase its expected payoff by deviating and choosing z whenever the cost is currently greater than χ and choosing I whenever the cost is currently no greater than χ. It follows that one firm will choose I and the other firm will choose A whenever the cost is currently equal to κ_1.

This completes the desired characterization of Markov perfect equilibrium. □

Proof of Item 2 in Remark 4. Note that $\beta < 0$ is decreasing in μ and ρ but increasing in σ. The cutoff κ_2 is given by $[\beta/(\beta - 1)]b_2$, which is increasing in b_2 and decreasing in β. Hence, κ_2 is increasing in μ and ρ but decreasing in σ.

The cutoff κ_1 is defined by the implicit function $f(b_1, b_2, \kappa_1, \beta) = b_1 - \kappa_1 - b_2^{1-\beta}\kappa_1^{\beta}(-\beta)^{-\beta}(1-\beta)^{\beta-1} = 0$. It follows from the proof of proposition 7 that $\partial f/\partial\kappa_1 < 0$. It is also clear that $\partial f/\partial b_1 > 0$ and $\partial f/\partial b_2 < 0$. In addition, we have:

$$\partial f/\partial\beta = b_2^{1-\beta}\kappa_1^{\beta}(-\beta)^{-\beta}(1-\beta)^{\beta-1}\big(\log\{[\beta/(\beta-1)]b_2\} - \log(\kappa_1)\big) < 0,$$

observing that $[\beta/(\beta - 1)]b_2 < \kappa_1$.

The partial derivative of the threshold κ_1 with respect to a parameter $\alpha \in \{b_1, b_2, \beta\}$ can be signed as follows:

$$\text{sgn}(\partial\kappa_1/\partial\alpha) = \text{sgn}[-(\partial f/\partial\alpha)/(\partial f/\partial\kappa_1)] = \text{sgn}(\partial f/\partial\alpha).$$

Hence, $\partial\kappa_1/\partial b_1 > 0$, $\partial\kappa_1/\partial b_2 < 0$, and $\partial\kappa_1/\partial\beta < 0$. It follows that $\partial\kappa_1/\partial\mu > 0$, $\partial\kappa_1/\partial\rho > 0$, and $\partial\kappa_1/\partial\sigma < 0$. □

B.3.3 Discussion of Item 1 in Remark 4

Noting that each agent can move at most twice, it is straightforward to confirm that any Markov perfect equilibrium satisfies both uniform and pathwise admissibility. In any Markov perfect equilibrium, uniform inertia is violated. To see this, fix a history up to an arbitrary time t in which the cost is currently $c_t \in (\kappa_2, \kappa_1]$ and there has been no previous entry. Consider a firm that takes action A at such a history. For any $\epsilon > 0$, there is positive conditional probability that $c_\tau \leq \kappa_2$ for some $\tau \in (t, t+\epsilon)$, which implies that this firm takes action F in the time interval $(t, t+\epsilon)$. Thus, there cannot exist $\epsilon > 0$ such that this firm does not move during the time interval $(t, t+\epsilon)$.[2] However, pathwise inertia is satisfied because the cost process has continuous sample paths. To see this, consider any history up to time t and any realization of the cost process $\{c_\tau\}_{\tau \in (t,\infty)}$ after time t. If $c_t > \kappa_1$, there exists $\epsilon > 0$ such that $c_\tau \neq \kappa_1$ for all $\tau \in (t, t+\epsilon)$. If $\kappa_1 \geq c_t > \kappa_2$, there exists $\epsilon > 0$ such that $c_\tau \neq \kappa_2$ for all $\tau \in (t, t+\epsilon)$. In each case, the agents do not move during the time interval $(t, t+\epsilon)$. If $\kappa_2 \geq c_t$, then there is no $\epsilon > 0$ such that the agents move during the time interval $(t, t+\epsilon)$.

B.4 Application in Section 5.4

B.4.1 Formal Definitions of Histories and Strategy Spaces

Choose any time $t \in [0, T)$, the taste process $\{x_\tau\}_{\tau \in [0,t]}$ up to time t, and the sequence $(t^k)_{k=1}^K$ of Poisson arrival times no greater than t, where $t^K = t$ if there is a Poisson hit at time t. A history up to time t is represented by $\left(\{x_\tau\}_{\tau \in [0,t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i \in \{B,S\}}\}_{\tau \in [0,t)}\right)$, where $\{(a_\tau^i)_{i \in \{B,S\}}\}_{\tau \in [0,t)}$ denotes the action path of agent $i \in \{B, S\}$ up to time t with the action space of each agent i being $\mathbb{R} \cup \{z\}$.

The set of all histories up to an arbitrary time is denoted by H. We partition it as follows.

1. Let $H^\emptyset$ be the set consisting of every history up to any time t that has the form $\left(\{x_\tau\}_{\tau \in [0,t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i \in \{B,S\}}\}_{\tau \in [0,t)}\right)$ where $a_\tau^i = z$ for each $i = B, S$ and all $\tau \in [0, t)$.

2. For any $c \in \mathbb{R}$, let H^c be the set consisting of every history up to any time t that has the form $\left(\{x_\tau\}_{\tau \in [0,t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i \in \{B,S\}}\}_{\tau \in [0,t)}\right)$ where $a_\tau^S = z$ for all $\tau \in [0, t)$ and there exists $\tau' \in [0, t)$ such that $a_{\tau'}^B = c$ and $a_\tau^B = z$ for all $\tau \in [0, t) \setminus \{\tau'\}$.

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

3. For any $c \in \mathbb{R}$, let $H^{c,c}$ be the set consisting of every history up to any time t that has the form $\big(\{x_\tau\}_{\tau\in[0,t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i\in\{B,S\}}\}_{\tau\in[0,t)}\big)$ where there exist $\tau', \tau'' \in [0,t)$ with $\tau' < \tau''$ such that $a_{\tau'}^B = c$, $a_\tau^B = z$ for all $\tau \in [0,t) \setminus \{\tau'\}$, $a_{\tau''}^S = c$, and $a_\tau^S = z$ for all $\tau \in [0,t) \setminus \{\tau''\}$.

The feasibility constraints are as follows. For B, $\bar{A}_B(h_t) = \{z\} \cup \mathbb{R}$ if $h_t \in H^\emptyset$, and $\bar{A}_B(h_t) = \{z\}$ otherwise. For S, $\bar{A}_S(h_t) = \{z, c\}$ if $t^K = t$ and there exists $c \in \mathbb{R}$ such that $h_t \in H^c$, and $\bar{A}_S(h_t) = \{z\}$ otherwise.

The sets of feasible strategies are:

$$\begin{array}{l} \bar{\Pi}_B = \{\pi_B : H \to \{z\} \cup \mathbb{R} \mid \pi_B(h_t) \in \bar{A}_B(h_t) \text{ for all } h_t \in H\} \\ \bar{\Pi}_S = \{\pi_S : H \to \{z\} \cup \mathbb{R} \mid \pi_S(h_t) \in \bar{A}_S(h_t) \text{ for all } h_t \in H\} \end{array}.$$

The set of traceable, frictional, calculable, and feasible strategies can then be defined and is denoted by $\bar{\Pi}_i^C$ for agent $i = B, S$.

The shock process s_t is formally defined as a triple comprising the taste x_t, the calendar time t, and an indicator for there being a Poisson hit at that time. The instantaneous utility function v_i is specified as follows for $i = B$:

$$v_B[(a_\tau^B, a_\tau^S), s_\tau] = \begin{cases} 0 & \text{if } a_\tau^S = z \\ v - p - \mathbb{E}[(c - x_T)^2 | x_\tau] & \\ \quad = v - p - (c - x_\tau)^2 - \sigma^2(T - \tau) & \text{if } a_\tau^S = c \in \mathbb{R} \end{cases},$$

and as follows for $i = S$:

$$v_S[(a_\tau^B, a_\tau^S), s_\tau] = \begin{cases} 0 & \text{if } a_\tau^S = z \\ p & \text{if } a_\tau^S = c \in \mathbb{R} \end{cases}.$$

B.4.2 Proofs

Proof of Proposition 8. First, since $p > 0$, it is a strictly dominant strategy for S to sell the good as soon as he obtains a chance to do so after an order is placed. Second, since the only choice B effectively makes is the time of placing an order, her maximization problem can be written as:

$$\max_{\tau\in(0,T]} u(\tau) = (1 - e^{-\lambda\tau})\mathbb{E}[v - (s - x_T)^2 - p] = (1 - e^{-\lambda\tau})(v - \sigma^2\tau - p),$$

where τ represents the amount of time remaining until the deadline at time T. The first-order condition is:

$$u'(\tau) = \lambda e^{-\lambda\tau}(v - \sigma^2\tau - p) - \sigma^2(1 - e^{-\lambda\tau}) = 0. \tag{5}$$

The second derivative is given by $u''(\tau) = -\lambda e^{-\lambda\tau}[\lambda(v - \sigma^2\tau - p) + 2\sigma^2]$. Note that $u'(0) > 0$ since $v > p$. In addition, $u''(\tau) < 0$ whenever $u'(\tau) \geq 0$. Hence, the objective function has a unique global maximizer in $[0, T)$. Let τ' be the unique value of τ satisfying equation (5), and define $\tau^* = \min\{\tau', T\}$ and $t^* = T - \tau^*$. This completes the desired characterization of the unique equilibrium strategy profile. □

Proof of Item 3 in Remark 5. We apply the implicit function theorem to (5) to conduct comparative statics, focusing on the case where $\tau^* = \tau'$. That is, (5) holds when $\tau = \tau^*$.

Recall that if (5) holds, then the derivative of its left-hand side is strictly negative. Denoting the left-hand side with $\tau = \tau^*$ by $f(\tau^*, \sigma, v, p, \lambda)$, it follows that:

$$\text{sign}(\partial t^*/\partial\phi) = -\text{sign}(\partial\tau^*/\partial\phi) = -\text{sign}[-(\partial f/\partial\phi)/(\partial f/\partial\tau^*)] = -\text{sign}(\partial f/\partial\phi),$$

where $\phi \in \{\sigma, v, p, \lambda\}$.

First, t^* is increasing in σ because $\partial f/\partial\sigma = 2\sigma[-\lambda e^{-\lambda\tau^*}\tau^* - (1 - e^{-\lambda\tau^*})] < 0$. Second, t^* is decreasing in v because $\partial f/\partial v = \lambda e^{-\lambda\tau^*} > 0$. Third, t^* is increasing in p since the effect of p is opposite to the effect of v by (5). Fourth, t^* is increasing in λ as:

$$\begin{aligned}\partial f/\partial\lambda &= e^{-\lambda\tau^*}(v - \sigma^2\tau^* - p) - \lambda\tau^* e^{-\lambda\tau^*}(v - \sigma^2\tau^* - p) - \sigma^2\tau^* e^{-\lambda\tau^*} \\ &= \sigma^2(1 - e^{-\lambda\tau^*})/\lambda - \sigma^2\tau^*(1 - e^{-\lambda\tau^*}) - \sigma^2\tau^* e^{-\lambda\tau^*} = \sigma^2(1 - e^{-\lambda\tau^*} - \lambda\tau^*)/\lambda < 0,\end{aligned}$$

where the second step applies the equality $\lambda e^{-\lambda\tau^*}(v - \sigma^2\tau^* - p) = \sigma^2(1 - e^{-\lambda\tau^*})$. □

B.4.3 Discussion of Item 1 in Remark 5

Noting that each agent can move at most once, it can easily be seen that both uniform and pathwise admissibility are satisfied in an SPE. B's equilibrium strategy, which simply involves moving at a predetermined time, satisfies both uniform and pathwise inertia. S's unique equilibrium strategy violates uniform inertia. For any $\epsilon > 0$, there is positive probability of a Poisson hit in the time interval $(t^*, t^* + \epsilon)$, in which case S sells the good. Thus, there is no $\epsilon > 0$ such that S does not move in the time interval $(t^*, t^* + \epsilon)$. However, S's strategy in equilibrium is pathwise inertial. Given any time t as well as any realization of the Poisson process, there exists $\epsilon > 0$ such that there is no Poisson hit in the time interval $(t, t + \epsilon)$. Since S can move only at the arrival times of the Poisson process, S does not move during this interval of time.

C Payoff Assignment with Nonmeasurable Behavior

Section C.1 describes how to assign expected payoffs to nonmeasurable behavior. Section C.2 points out a few problems with this methodology. In section C.3, we show that with some restrictions, any behavior on the path of play can be sustained in an SPE under a particular assignment of expected payoffs. Section C.4 demonstrates that any SPE under the calculability assumption is an SPE for a certain way of assigning payoffs, and section C.5 provides conditions under which the converse is also true.

C.1 Formulation of Expected Payoffs

We consider the assignment of expected payoffs to strategy profiles that induce nonmeasurable behavior.[3] For each agent $i \in I$, define a function $\chi_i : H \times \bar{\Pi}_i^{TF} \to \mathbb{R} \cup \{-\infty, \infty\}$.

Choose any strategy profile $\pi \in \times_{i \in I} \bar{\Pi}_i^{TF}$. Let $k_u = \big(\{g_t\}_{t \in [0,u]}, \{(b_t^i)_{i \in I}\}_{t \in [0,u)}\big)$ be any history up to time u, and denote $b = \{(b_t^i)_{i \in I}\}_{t \in [0,u)}$. If the process $\xi_b^i(\pi)$ is progressively measurable for each $i \in I$, then the expected payoff to agent i at k_u is given by $U_i(k_u, \pi) = V_i(k_u, \pi)$, where $V_i(k_u, \pi)$ is as specified in equation (1) in section 4.3. Otherwise, the expected payoff to agent i at k_u is given by $U_i(k_u, \pi) = \chi_i(k_u, \pi)$.[4]

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 4), 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 $(\Omega, \mathcal{F}, P)$.[5] Suppose $I = \{1\}$ and that $\bar{A}_1(h_t) = \{x, z\}$ for every $h_t \in H$. The utility function satisfies $v_1(x, s) = 0$ for all $s \in S$. Let $\chi_1(h_t, \pi_1) = -1$ for all $h_t \in H$ and any $\pi_1 \in \bar{\Pi}_1^{TF}$. A strategy in which agent 1 chooses action x at time $\tilde{t}$ if and only if $s_{\tilde{t}}$ is in $\tilde{S}$ is not optimal at the null history because agent 1 can deviate to a strategy that induces measurable behavior. However, there is no history up to a given time at which a one-shot deviation would increase the expected payoff of the agent. □

[3] We assume traceability and frictionality, so that the behavior of the agents is well defined in the sense that there exists a unique action path consistent with each strategy profile. 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.

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

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

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

C.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 $(\Omega, \mathcal{F}, P)$. The utility function satisfies $v_1(x, s) = 1$ for all $s \in S$. Let $\chi_1(h_t, \pi_1) = -1$ for all $h_t \in H$ and any $\pi_1 \in \bar{\Pi}_1^{TF}$.

Consider a class of strategies, each of which is indexed by a set $C \subseteq S$, where π^C prescribes action z at any time $t \neq 1$ and action x at time $t = 1$ if and only if s_1 is in C. It may be natural for the expected payoffs to satisfy the following monotonicity condition: $U_1(h_0, \pi^{S''}) \leq U_1(h_0, \pi^{S'})$ if $S'' \subseteq S'$. That is, the expected payoff is monotonic in the realized payoffs in the sense of statewise dominance. However, $U_1(h_0, \pi^{\emptyset}) = 0 > -1 = U_1(h_0, \pi^{\tilde{S}})$ even though $\emptyset \subseteq \tilde{S}$. Hence, the monotonicity condition fails. □

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 $(\Omega, \mathcal{F}, P)$. The utility functions satisfy $v_1[(x, z), s] = 1$, $v_1[(z, x), s] = 0$, and $v_2[(x, z), s] = v_2[(z, x), s] = 0$ for all $s \in S$.

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

Second, suppose that $\chi_1(h_t, \pi) = \chi_2(h_t, \pi) = -1$ for all $h_t \in H$ and any $\pi \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$. Then there does not exist an SPE in which agent 1 receives an expected payoff of 0.

The reason is that there is no history h_1 up to time 1 for which agent 2 has an incentive to choose a strategy π_2 such that $U_2[h_1,(\pi_1,\pi_2)] = \chi_2[h_1,(\pi_1,\pi_2)]$ for some strategy π_1 of agent 1. Hence, it is always optimal for agent 1 to choose x at time 0, so that agent 1 receives an expected payoff of 1. $\square$

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

C.3 Arbitrary Behavior in Equilibrium

Now we examine the implications for agents' incentives of assigning payoffs to nonmeasurable behavior. Given any history k_u up to time u, let $\bar{H}(k_u)$ denote the set consisting of every history $h = \{s_t, (a_t^j)_{j\in I}\}_{t\in[0,T)}$ such that $\big(\{s_t\}_{t\in[0,u]}, \{(a_t^j)_{j\in I}\}_{t\in[0,u)}\big) = k_u$, $\{a_t^i\}_{t\in[0,T)} \in \Xi_i(u)$ for each $i \in I$, and $a_\tau^i \in \bar{A}_i\left(\{s_t\}_{t\in[0,\tau]}, \{(a_t^j)_{j\in I}\}_{t\in[0,\tau)}\right)$ 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.[6] 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 behavior is assigned an expected payoff no greater than the infimum of the set of feasible payoffs.[7]

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 measureable, 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 $(\Omega, \mathcal{F}, P)$.

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

[7] Similar 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 $k_u = \left(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\right)$ 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 $i \in I$.

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 $\chi_i(h_t,\tilde{\pi}) \leq 0$. □

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 $(\Omega, \mathcal{F}, P)$. The utility functions satisfy $v_1[(a_1,a_2),s] \geq v_1[(z,a_2),s] \geq 0$ and $v_2[(a_1,a_2),s] \geq v_2[(a_1,z),s] \geq 0$ for all $s \in S$ and any $(a_1,a_2) \in A_1 \times A_2$ such that $(a_1,a_2) \neq (z,z)$. For each $i \in \{1,2\}$, let $\chi_i(h_t,\pi') = \chi_i(h_t,\pi'') \leq 0$ for all $h_t \in H$ and any $\pi',\pi'' \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$.

Choose any pair of actions $(\tilde{a}_1,\tilde{a}_2) \in A_1 \times A_2$. Let $\tilde{\pi}$ be the strategy profile defined as follows. Each agent $i \in \{1,2\}$ chooses $\tilde{a}_i$ at time 0. If every agent i takes action $\tilde{a}_i$ at time 0, then the agents choose x at time 1. If some agent i does not take action $\tilde{a}_i$ at time 0, then the agents choose x at time 1 if and only if s_1 is in $\tilde{S}$. This strategy profile is an SPE. A unilateral deviation at h_0 would result in an expected payoff of $\chi_i(h_0,\tilde{\pi}) \leq 0$ to agent i. If strategy profile $\tilde{\pi}$ is followed at time 0, then playing x at time 1 is a best response for each agent to the action of the other agent. If strategy profile $\tilde{\pi}$ is not followed at time 0, then neither agent has an incentive to deviate again for the same reason as in example 12. □

The preceding example illustrates how any profile of actions at the null history can be implemented in equilibrium.[8] 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 13. *Let $|I| \geq 2$ and $T = \infty$. Consider the game $\Gamma\left(\times_{j\in I}\bar{\Pi}_j^{TF}, (\chi_j)_{j\in I}\right)$ where $\chi_i(h_t,\hat{\pi}) = \chi_i(h_t,\bar{\pi}) \leq \zeta_i(h_t)$ for all $\hat{\pi},\bar{\pi} \in \times_{j\in I}\bar{\Pi}_j^{TF}$, any $h_t \in H$, and every $i \in I$. Assume that for each $i \in I$ and any $h_t \in H$, there exists $\tilde{a} \in \bar{A}_i(h_t)$ such that $\tilde{a} \neq z$.*

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

Suppose that there exists $\tilde{t} > 0$ along with a collection of sets $\{\tilde{S}_t\}_{t\in[0,\tilde{t}]}$ such that $\{\omega \in \Omega : s_t(\omega) \in \tilde{S}_t,\ \forall\, t \in [0,\tilde{t}]\}$ is not a measurable subset of the probability space $(\Omega, \mathcal{F}, P)$. Choose any $\pi \in \times_{i\in I}\bar{\Pi}_i^{TF}$ such that for any profile of action paths $\left(\{b_t^i\}_{t\in[0,u)}\right)_{i\in I}$ up to an arbitrary time $u > 0$, there exists with probability one some $t < u$ such that $\pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(b_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right) \neq b_t^i$ for some $i \in I$. Then there exists an SPE $\pi' \in \times_{i\in I}\bar{\Pi}_i^{TF}$ such that $\left(\{\phi_t^i(h_0, \{s_\tau\}_{\tau\in(0,T)}, \pi')\}_{t\in[0,T)}\right)_{i\in I} = \left(\{\phi_t^i(h_0, \{s_\tau\}_{\tau\in(0,T)}, \pi)\}_{t\in[0,T)}\right)_{i\in I}$ with probability one.

Proof. Choose any strategy profile $\pi \in \times_{i\in I}\bar{\Pi}_i^{TF}$. Define the strategy profile $\pi' \in \times_{i\in I}\bar{\Pi}_i^{TF}$ as follows. Let $k_u = \left(\{s_t\}_{t\in[0,u]}, \{(a_t^i)_{i\in I}\}_{t\in[0,u)}\right)$ be any history up to an arbitrary time u. If $a_t^i = \pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(a_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right)$ for each $i \in I$ and all $t \in [0,u)$, then let $\pi_i'(k_u) = \pi_i(k_u)$ for all $i \in I$. If there exists $t \in [0,u)$ and $i \in I$ such that $a_t^i \neq \pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(a_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right)$, then for each $i \in I$, let $\pi_i'(k_u) \neq z$ if s_t is in $\tilde{S}_t$ for all $t \in [0,\tilde{t}]$ and $u = n\tilde{t}$ for some positive integer n, and let $\pi_i'(k_u) = z$ otherwise.

The strategy profile π' is an SPE because if agent $i \in I$ deviates at h_0, then the expected payoff to agent i is $\chi_i(h_0, \pi') \leq \zeta_i(h_0)$. Moreover, consider any history $k_u = \left(\{g_t\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\right)$ up to an arbitrary time $u > 0$. By assumption, there exists with probability one some $t < u$ such that $\pi_i\left(\{s_\tau\}_{\tau\in[0,t]}, \{(b_\tau^j)_{j\in I}\}_{\tau\in[0,t)}\right) \neq b_t^i$ for some $i \in I$. It follows that for each $i \in I$, the function $\phi_{n\tilde{t}}^i\left[\left(\{s_t(\omega)\}_{t\in[0,u]}, \{(b_t^j)_{j\in I}\}_{t\in[0,u)}\right), \{s_t(\omega)\}_{t\in(u,\infty)}, \pi'\right]$ 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'. □

The result is proved by constructing a strategy profile π' with the following properties. At the null history h_0, agent i has no incentive to deviate because doing so would result in an expected payoff of $\chi_i(h_0, \pi)$, which is no more than the infimal feasible payoff $\zeta_i(h_0)$ based on the utility function. At any history k_u up to a positive time u, the action path b up to time u is such that the process $\xi_b^i(\pi')$ is not progressively measurable given the restriction on strategy profile π in the statement of the proposition as well as the assumption that $\xi_b^i(\pi')$ records the continuation path of play when the action path up to time u is fixed at b. Hence, the expected payoff of agent i at history k_u is constant at $\chi_i(k_u, \pi')$ when the other agents play π', so that agent i has no incentive to deviate.[9]

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

[9]We assume only for simplicity when defining $\xi_b^i(\pi')$ that the action path up to time u is fixed irrespective of the shock realization up to time u. The reasoning in the proof does not entirely apply under the alternative definition in footnote 28 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 28, the main results in section 4 are valid under both definitions of the action process.

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.

C.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 C.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 14. *Let $\bar{A}_i(h_t) = A_i$ for every $h_t \in H$ and each $i \in I$. If π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^C\right)$, then π is an SPE of $\Gamma\left(\times_{i\in I}\bar{\Pi}_i^{TF}, (\chi_i)_{i\in I}\right)$ with $\chi_i(h_t, \tilde{\pi}) \leq \zeta_i(h_t)$ for all $\tilde{\pi} \in \times_{j\in I}\bar{\Pi}_j^{TF}$, any $h_t \in H$, and every $i \in I$.*

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

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

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. 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.

C.5 From Payoff Assignment to Calculability Restriction

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

We begin by defining a set of strategy profiles with certain measurability properties. Let the random variable $\theta : \Omega \to [0, T]$ be a stopping time.[10] Given any $\pi', \pi'' \in \times_{i\in I}\Pi_i$, let $\psi(\pi', \pi'', \theta)$ be the strategy profile satisfying the following two properties for each $i \in I$:

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

In other words, the strategy $\psi_i(\pi', \pi'', \theta)$ plays π_i' until and including the stopping time, and plays π_i'' thereafter. A strategy profile $\pi \in \times_{i\in I}\Pi_i^{TF}$ is said to be **measurably attachable** if for each action path b, every stopping time θ, and any $\tilde{\pi} \in \times_{i\in I}\Pi_i^{TF}$ such that $\xi_b^i(\tilde{\pi})$ is progressively measurable for all $i \in I$, the strategy $\psi_i(\tilde{\pi}, \pi, \theta)$ is traceable and frictional for all $i \in I$ and the process $\xi_b^i[\psi(\tilde{\pi}, \pi, \theta)]$ is progressively measurable for all $i \in I$. That is, π is required to induce progressively measurable behavior after any progressively measurable behavior up to and including an arbitrary random time. Let $\Pi^A \subseteq \times_{i\in I}\Pi_i^{TF}$ be the set of measurably attachable strategy profiles. In addition, a strategy profile $\pi \in \times_{i\in I}\Pi_i$ is said to be **synchronous** if for any $h_t \in H$, $\pi_j(h_t) = z$ for all $j \in I$ whenever $\pi_i(h_t) = z$ for some $i \in I$. That is, π requires the agents to move at the same time as each other.[11]

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

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

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

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

[11] The 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 supplementary information as well as by a Markov perfect equilibrium of the inventory restocking application in the supplementary information. 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 supplementary information satisfy synchronicity after such a reformulation.

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

Given any $i \in I$, we now confirm that $\pi_i \in \Pi_i^C$. Define $\pi_j^z(h_t) = z$ for each $j \in I$ and every $h_t \in H$. Choose any $b = \{(b_t^j)_{j\in I}\}_{t\in[0,u)}$ as well as any $\pi_{-i}' \in \Pi_{-i}^Q$. 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

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

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

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

for some $e \in I$. Let $\theta^0(\omega)$ be the lesser of the infimum of $\hat{\Upsilon}^0(\omega)$ and the infimum of $\bar{\Upsilon}^0(\omega)$. Note that $\xi_b^i[\psi(\pi,\pi^z,\theta^0)]$ and $\xi_b^i\{\psi[(\pi_i,\pi_{-i}'),\pi^z,\theta^0]\}$ are the same progressively measurable stochastic process.

Apply the following procedure iteratively for every positive integer k. Define the stopping time θ^k as follows. For any $\omega \in \Omega$, let $\hat{\Upsilon}^k(\omega)$ denote the set consisting of all $t \in (\theta^{k-1}(\omega),T)$ such that

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

for some $e \in I$, and let $\bar{\Upsilon}^k(\omega)$ denote the set consisting of all $t \in (\theta^{k-1}(\omega),T)$ such that

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

for some $e \in I$. Let $\theta^k(\omega)$ be the lesser of the infimum of $\hat{\Upsilon}^k(\omega)$ and the infimum of $\bar{\Upsilon}^k(\omega)$, where $P(\{\omega\in\Omega : \theta^k(\omega) > \theta^{k-1}(\omega)\} \cup \{\omega\in\Omega : \theta^{k-1}(\omega) = \infty\}) = 1$ because $\psi[(\pi_i,\pi_{-i}'),\pi,\theta^{k-1}], (\pi_i,\pi_{-i}') \in \times_{j\in I}\Pi_j^{TF}$. Note that $\xi_b^i\big(\psi\{\psi[(\pi_i,\pi_{-i}'),\pi,\theta^{k-1}],\pi^z,\theta^k\}\big)$ and $\xi_b^i\{\psi[(\pi_i,\pi_{-i}'),\pi^z,\theta^k]\}$ are the same progressively measurable stochastic process.

Suppose that the sequence $\{\theta^k\}_{k=1}^\infty$ does not converge almost surely to ∞. Then $P[\{\omega\in\Omega : \lim_{k\to\infty}\theta^k(\omega) < \infty\}] \neq 0$. For each $\omega \in \Omega$, let $\Xi(\omega)$ denote the set consisting of all $t \in [u,T)$ such that

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

for some $e \in I$. Letting E denote the set consisting of all $\omega \in \Omega$ such that $\{t \in \Xi(\omega) : t \leq c\}$ contains only finitely many elements for any $c \in [u,\infty)$, we have $P(\{\omega \in \Omega : \omega \in$

$E\}) = 1$ because $(\pi_i, \pi'_{-i}) \in \times_{j\in I}\Pi_j^{TF}$. The definition of E also implies that for all $\omega \in E$ such that $\lim_{k\to\infty} \theta^k(\omega) < \infty$, there exists $\tilde{t}(\omega) \in [u, \lim_{k\to\infty} \theta^k(\omega))$ such that

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

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

For any $\omega \in E$ such that $\lim_{k\to\infty} \theta^k(\omega) < \infty$, choose any $\tilde{k}(\omega) \geq 1$ such that $\theta^{\tilde{k}(\omega)-1}(\omega) \geq \tilde{t}(\omega)$. Assuming now that π is synchronous,

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

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

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 $(\Omega, \mathcal{F}, P)$. For all $s \in S$, the utility function of each agent $i \in \{1, 2\}$ satisfies $v_i[(x, x), s] = 1$ and $v_i[(a_1, a_2), s] = 0$ for any $(a_1, a_2) \in A_1 \times A_2$ such that $(a_1, a_2) \neq (x, x)$. Let $\chi_i(h_t, \pi) = 0$ for each $i \in I$, all $h_t \in H$, and any $\pi \in \bar{\Pi}_1^{TF} \times \bar{\Pi}_2^{TF}$.

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

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 supplementary information, 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.

References for the Online Appendix

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

Rogerson, R., R. Shimer, and R. Wright (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Literature*, 43(4), 959–988.