



UTMD Working Paper

The University of Tokyo
Market Design Center

UTMD-040

[Appendix]

Strategies in Stochastic Continuous-Time Games

Yuichiro Kamada

University of California Berkeley

Neel Rao

University at Buffalo - SUNY

January 28, 2023

UTMD Working Papers can be downloaded without charge from:

<https://www.mdc.e.u-tokyo.ac.jp/category/wp/>

Working Papers are a series of manuscripts in their draft form. They are not intended for circulation or distribution except as indicated by the author. For that reason, Working Papers may not be reproduced or distributed without the written consent of the author.

ONLINE SUPPLEMENTARY APPENDIX TO:
“STRATEGIES IN STOCHASTIC
CONTINUOUS-TIME GAMES”

Yuichiro Kamada[†] Neel Rao[‡]

January 28, 2023

[†]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 $I = \{1, \dots, n\}$ denote the set of woodcutters. A history up to time t is represented by $(\{b_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in I}\}_{\tau \in [0, t]})$, where $\{a_\tau^i\}_{\tau \in [0, t]}$ denotes the action path of woodcutter i up to time t with the action space being $\mathbb{R}_{++}^2 \cup \{z\}$. The action $a_\tau^i = (e_\tau^i, f_\tau^i) \in \mathbb{R}_{++}^2$ means that woodcutter i seeks to harvest the amount e_τ^i and claim the amount f_τ^i at time τ . The action $a_\tau^i = z$ 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 $a_\tau^i = z$ and there exists $d_\tau > 0$ such that $e_\tau^i = d_\tau$ 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 $k \in \{1, \dots, K\}$, 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 $i \in I$:

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

The set of traceable, frictional, calculable, and feasible strategies can then be defined and is denoted by $\bar{\Pi}_i^C$ 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 $v_i[(a_\tau^i, a_\tau^{-i}), s_\tau] = 0$ 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

$$v_i[(a_\tau^i, a_\tau^{-i}), s_\tau] = \exp(-\rho\tau) \left(f_\tau^i / \sum_{j \in I} f_\tau^j \right) (d_\tau - \kappa)$$

if there is no $j \in I$ such that $a_j = z$ and there exists $d_\tau > 0$ such that $e_\tau^i = d_\tau$ 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 11. *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 $U(b_t) = r e^{\alpha(b_t - c)}$, 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 12. *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 13. *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 12.

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 14. *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 $V(q^*) = V(x^*) - \kappa/n$ 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 $\alpha\tilde{W}'(\alpha) = (\alpha W(\ell)/\{\ell[1 + W(\ell)]\})\partial\ell/\partial\alpha$, 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 $\lim_{\alpha\downarrow 0} -1 + \alpha\tilde{W}'(\alpha) - \tilde{W}(\alpha) = 0$. \square

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 $(\{p_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in \{W, R\}}\}_{\tau \in [0, t]})$, where $\{a_\tau^W\}_{\tau \in [0, t]}$ and $\{a_\tau^R\}_{\tau \in [0, t]}$ 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 $a_\tau^W \in \mathbb{R}_{++} \times \mathbb{R}_+$, the first element of a_τ^W , 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_τ^R is a function from \mathbb{R}_+ to itself, $a_\tau^R(\tau') > 0$ represents the payment to the oil refinery at time τ from delivering the output produced from the oil received at time τ' , where $a_\tau^R(\tau') = 0$ indicates that no such delivery was made by the oil refinery at time τ . The action $a_\tau^R = z$ 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 $a_\tau^W \neq z$. 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 $a_t^R \in \bar{A}_R(h_t)$ if and only if $a_t^R = z$ or a_t^R satisfies the following. Choose any time $\tau \in [0, \infty)$. If there exists k such that $t_k = \tau$ and $\tau + d(e_\tau) \leq t$ and there is no $t' < t$ such that $a_{t'}^R(\tau) > 0$, then $a_t^R(\tau) \in \{0, y(p_\tau, e_\tau)\}$. Otherwise, $a_t^R(\tau) = 0$. In addition, $a_t^R(\tau) > 0$ for some $\tau \in [0, t)$.

The sets of feasible strategies are:

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

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

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:

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

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

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

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 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 $(\{c_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in \{1, 2\}}\}_{\tau \in [0, t]})$, 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 $\{I, A, F, z\}$.

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

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 $(\{c_\tau\}_{\tau \in [0,t]}, \{(a_\tau^i)_{i \in \{1,2\}}\}_{\tau \in [0,t]})$ 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 $(\{c_\tau\}_{\tau \in [0,t]}, \{(a_\tau^i)_{i \in \{1,2\}}\}_{\tau \in [0,t]})$ 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 $(\{c_\tau\}_{\tau \in [0,t]}, \{(a_\tau^i)_{i \in \{1,2\}}\}_{\tau \in [0,t]})$ 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 $(\{c_\tau\}_{\tau \in [0,t]}, \{(a_\tau^i)_{i \in \{1,2\}}\}_{\tau \in [0,t]})$ 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 $i = 1, 2$:

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

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. \square

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:

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

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 $\partial\kappa_1/\partial\sigma < 0$. \square

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 $(t, t + \epsilon)$.

²If 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 $\kappa \in (\kappa_2, \kappa_1)$.

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 $(\{x_\tau\}_{\tau \in [0, t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i \in \{B, S\}}\}_{\tau \in [0, t]})$, 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 $(\{x_\tau\}_{\tau \in [0, t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i \in \{B, S\}}\}_{\tau \in [0, t]})$ 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 $(\{x_\tau\}_{\tau \in [0, t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i \in \{B, S\}}\}_{\tau \in [0, t]})$ 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 $\tau \in [0, t) \setminus \{\tau'\}$.
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 $(\{x_\tau\}_{\tau \in [0, t]}, (t^k)_{k=1}^K, \{(a_\tau^i)_{i \in \{B, S\}}\}_{\tau \in [0, t]})$ 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:

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

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 $i = B$:

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

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:

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

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. \square

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 $\lambda e^{-\lambda \tau^*} (v - \sigma^2 \tau^* - p) = \sigma^2 (1 - e^{-\lambda \tau^*})$. \square

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 Additional Applications

C.1 Sequential Exchange with Transaction Cost

Kamada and Rao (2018) consider a problem involving the bilateral trade of divisible goods. Each of two parties is endowed with the same amount of a different good. Each agent derives utility not from its own good but from the good that the other agent initially possesses. The decision facing each agent is when to transfer its good to the other agent and how much of the good to transfer on each transaction. There is

a strictly positive transaction cost evolving over time according to a diffusion process that does not depend on the size of a transaction. The framework can be applied to study the exchange of information, prisoners, and land.³

The first-best solution (or efficient outcome), which maximizes the expected payoff of each agent in the absence of incentive constraints, requires there to be at most one transaction and for the entire stock of each good to be transferred on that transaction. However, such a strategy profile is not an SPE because an agent does not have an incentive to incur the transaction cost. In particular, once each agent transfers all of its good to the other agent, the future value of the relationship is zero since there is no possibility of further exchange, so that there is no reward for cooperating by making a transfer.

In Kamada and Rao (2018), any SPE in which trade occurs must involve a potentially infinite sequence of transactions on the path of play. For the case where the cost process $\{c_t\}_{t \in [0, \infty)}$ follows a geometric Brownian motion, those authors further solve for the maximal equilibrium (or second-best solution), which is the strategy profile that maximizes the expected payoff of each agent over the class of SPE. The maximal equilibrium is characterized by a sequence $\{c_k^*, f_k^*\}_{k=1}^\infty$ such that on the path of play, the amount f_k^* of each good is exchanged between the two agents when the cost reaches c_k^* for the first time. Such an equilibrium can be supported using grim-trigger strategies, in which failure to follow the specified path of play results in a permanent suspension of trade.

The maximal SPE violates uniform inertia for the following reason. Consider any history up to an arbitrary time t in which no agent has deviated in the past and each agent has previously made a total of m transactions. For any $\epsilon > 0$, there is positive conditional probability that $c_\tau = c_{m+1}^*$ for some $\tau \in (t, t + \epsilon)$, in which case the maximal SPE requires each agent to transfer the amount f_{m+1}^* during the time interval $(t, t + \epsilon)$. Thus, there cannot exist $\epsilon > 0$ such that the agents do not move during this time interval. However, the maximal SPE is pathwise inertial. To see this, consider any time t and any realization of the cost process $\{c_\tau\}_{\tau \in (t, \infty)}$ after time t . Let l be the least index k such that $c_t > c_k^*$. Due to the continuity of the sample paths of geometric Brownian motion, there exists $\epsilon > 0$ such that $c_\tau \neq c_l^*$ for all $\tau \in (t, t + \epsilon)$, which implies that the agents do not move during the time interval

³Kamada and Rao (2018) provide a detailed discussion of how the model and results fit such real-life settings.

$(t, t + \epsilon)$.

In regard to admissibility, uniform F1 is violated because there is no upper bound on the number of transactions in any proper time interval. Nonetheless, the maximal SPE satisfies pathwise F1 because for any realization of the cost process $\{c_\tau\}_{\tau \in [0, t]}$ up to time t , the agents transact only finitely many times during the time interval $[0, t]$. From proposition 3, assumption F2 holds for any traceable and frictional strategy. However, the grim-trigger strategies used to support the maximal SPE are incompatible with assumption F3 in Simon and Stinchcombe (1989) because any deviation from the prescribed timing of a transaction changes the subsequent behavior of the agents.

For a formal definition of strategy spaces in this context as well as comparative statics for the maximal equilibrium, see Kamada and Rao (2018).

C.2 Partnership and Cooperation between Criminals

In the analysis of the finite-horizon ordering game in section 5.3, the evolution of a diffusion process induces agents to move at a time with no Poisson hit. In the example below, there is no such diffusion process, but restricting agents to move only at Poisson arrival times is still problematic. In particular, such a restriction may cause a delay in punishment after a deviation, thereby weakening the scope for punishment, which makes it difficult to enforce cooperation.

Consider a partnership game between two criminals, 1 and 2, where time t runs continuously in $[0, \infty)$. At each moment of time t , a criminal chooses between remaining in the partnership R and permanently leaving the partnership L . The criminals receive flow payoffs at each time t depending on the profile of choices at time t . In particular, the flow payoff to each agent is 1 if both choose R and is 0 if either chooses L . Moreover, at random points in time that arrive according to a Poisson process with arrival rate $\lambda > 0$, the criminals are apprehended by the police, and they play a prisoner's dilemma with discrete payoffs, where they need to choose between admitting guilt D and remaining silent C . The payoffs in the prisoner's dilemma are specified in Table 1. The discount rate is $\rho > 0$.

Even though our general model seemingly does not allow for flow payoffs, this game can be reformulated as follows so as to fit into our framework. The reformulated version will not have a flow payoff, while the discrete payoffs from the prisoner's dilemma are unchanged. At time 0, each criminal receives a discrete payoff $1/\rho$, which

	C	D
C	c, c	p, b
D	b, p	d, d

Table 1: Prisoner’s Dilemma ($p < d < c < b$)

is the integral over time of the flow payoff to a criminal from the partnership game if both criminals remain in the partnership perpetually. If either criminal changes from R to L at time t and neither criminal has changed to L before t , then each criminal receives the discrete payoff $-1/\rho$ at time t . If one criminal changes from R to L at time t and the other criminal has changed to L at some time before t , then each criminal receives the discrete payoff 0 at time t . The action z in our framework would correspond to “not changing from R to L and choosing C .”

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each agent $i = 1, 2$.⁴ Call the game with such strategy spaces the *criminal game*. It is characterized by $(c, b, p, d, \lambda, \rho)$. The analysis in sections 3 and 4 implies that a subgame-perfect equilibrium is well defined.

We introduce two strategy profiles and show that one of them is supportable as an SPE for a larger set of parameter values than the other. First, we define the following strategy profile, which we call the *optimal Poisson-revision strategy profile*. Suppose that the current time is t .

1. When D and L have never been played in the past, play R , and C if there is a Poisson hit at t .
2. If D or L has ever been played in the past, let t^* be the infimum of the times at which D or L is played.
 - (a) If there is a Poisson hit at t , play D .
 - (b) If there is no Poisson hit at any time in $(t^*, t]$, play R if feasible.
 - (c) If there is a Poisson hit at some time in $(t^*, t]$, play L .

Second, we define the following strategy profile for $\Delta > 0$, which we call the *optimal Δ -delay strategy profile*. Suppose that the current time is t .

⁴Formal definitions of histories and strategy spaces are provided in section C.2.1.

1. When D and L have never been played in the past, play R , and C if there is a Poisson hit at t .
2. If D or L has ever been played in the past, let t^* be the infimum of the times at which D or L is played.
 - (a) If there is a Poisson hit at t , play D .
 - (b) If $t \in (t^*, t^* + \Delta)$, play R if feasible.
 - (c) If $t \in [t^* + \Delta, \infty)$, play L .

Proposition 15. *In the criminal game with $(c, b, p, d, \lambda, \rho)$, the optimal Δ -delay strategy profile is an SPE whenever the optimal Poisson-revision strategy profile is an SPE. Moreover, for any given profile (c, p, d, λ, ρ) , there exists b such that in the criminal game with $(c, b, p, d, \lambda, \rho)$, the optimal Δ -delay strategy profile is an SPE for some $\Delta > 0$ while the optimal Poisson-revision strategy profile is not.*

There is a simple intuition behind this result. If the criminals are restricted to move only at Poisson arrival times, then punishment for a deviation at time t must be postponed until the first Poisson hit strictly after time t . The distribution of the first arrival time of a Poisson process is governed by the parameter λ , which bounds the scope for punishment in the optimal Poisson-revision strategy profile. In the optimal Δ -delay strategy profile, by setting the time lag $\Delta > 0$ for punishment small enough, the punishment for a deviation can be made arbitrarily close to immediate. Hence, for any λ , this severer punishment can potentially make cooperation in the prisoner's dilemma incentive compatible under the optimal Δ -delay strategy profile even when it is not under the optimal Poisson-revision strategy profile.

Remark 6. 1. (Relationship to inertia) Uniform inertia is not satisfied in either of the aforementioned strategy profiles. To see this, fix any time t and pair of action paths up to time t such that criminal i has not chosen D or L before time t , criminal j first chooses D or L at some time $t^* < t$, and there is no Poisson hit in the time interval $(t^*, t]$. For any $\epsilon > 0$, there is positive probability of there being a Poisson hit in the time interval $(t, t + \epsilon)$, in which case criminal i changes to D in the prisoner's dilemma at the time of the Poisson hit. Hence, there is no $\epsilon > 0$ such that criminal i does not move in the time interval $(t, t + \epsilon)$.

However, these two strategy profiles are pathwise inertial. To see this, fix any time t and action paths up to time t . For any realization of the Poisson process, there exists $\epsilon > 0$ such that no Poisson hit occurs in the time interval $(t, t + \epsilon)$. Under the optimal Poisson-revision strategy profile, the agents move only when there is a Poisson hit, so no move occurs during this time interval. Under the optimal Δ -delay strategy profile, agents can also move at time $t^* + \Delta$. If $t^* + \Delta > t$, then the optimal Δ -delay strategy profile prescribes no move at each time τ such that $t < \tau < \min(t + \epsilon, t^* + \Delta)$. Otherwise, it prescribes no move in the time interval $(t, t + \epsilon)$.

2. (Relationship to admissibility) Uniform F1 is violated by both strategy profiles since each agent defects in every prisoner's dilemma game following a deviation, where unboundedly many prisoner's dilemma games may be played in any proper time interval. However, each of these strategy profiles satisfies pathwise F1. Given any time t and any realization of the Poisson process, there are only finitely many Poisson hits in the time interval $[0, t]$. Under the optimal Poisson-revision strategy profile, agents can move only at the times of a Poisson hit, and under the optimal Δ -delay strategy profile, each agent can move only one additional time. Hence, the number of moves during the time interval $[0, t]$ is bounded given the realization of the Poisson process.

Each strategy profile has the piecewise continuity property F2, which according to proposition 3, is an implication of traceability and frictionality.

The strong continuity assumption F3 is satisfied by the optimal Poisson-revision strategy profile but not by the optimal Δ -delay strategy profile. Under the optimal Δ -delay strategy profile but not under the optimal Poisson-revision strategy profile, a slight difference in the time when D is first chosen may cause a difference in the time when each agent responds by choosing L .

3. (Limiting behavior) On the one hand, for any $\Delta > 0$, if $\rho < \infty$ is sufficiently large, then neither the optimal Poisson-revision strategy profile nor the optimal Δ -delay strategy profile is an SPE. This is because the present value of future punishment becomes very small. On the other hand, for any $\Delta > 0$, if $\rho > 0$ is sufficiently small, then both of these strategy profiles are SPE. This is because the criminals value the future very highly.

On the one hand, for any $\Delta > 0$, if $\lambda < \infty$ is sufficiently large, then both the optimal Poisson-revision strategy profile and the optimal Δ -delay strategy profile are SPE because the future punishment for a deviation is very frequent. On the other hand, if $\lambda > 0$ is sufficiently small, then the optimal Poisson-revision strategy profile is not an SPE because the future punishment is very infrequent. The evaluation of the optimal Δ -delay strategy profile in the case of small $\lambda > 0$ is ambiguous and depends on the size of $\rho(b - c)$. The reason is that the payoffs in the partnership game can also be used to punish a deviator, and the timing of the punishment under the optimal Δ -delay strategy profile does not depend on the frequency of the Poisson hits.⁵

C.2.1 Formal Definitions of Histories and Strategy Spaces

At every moment of time, each criminal $i \in \{1, 2\}$ chooses an action from the set $A_i = \{(L_1, D), (L_2, D), (L_1, \bar{z}), (L_2, \bar{z}), (\bar{z}, D), (\bar{z}, \bar{z})\}$, where we let $z = (\bar{z}, \bar{z})$. The interpretation is as follows. The first element of each action represents the choice in the partnership game, where \bar{z} means “not changing the relationship,” and L_1 or L_2 means permanently leaving the partnership. The subscript k on L_k indicates that the agent is the k^{th} criminal to leave the partnership. The second element denotes the choice in the prisoner’s dilemma, where \bar{z} signifies cooperation at the time of a Poisson hit and corresponds to no activity at a time without a Poisson hit, and D indicates defection at the time of a Poisson hit.

Choose any time $t \in [0, T)$, the sequence $(t^k)_{k=1}^K$ of past arrival times of the Poisson process, and the action path $\{(a_\tau^i)_{i \in \{1, 2\}}\}_{\tau \in [0, t)}$ up to that time. A history up to time t is represented by $((t^k)_{k=1}^K, w, \{(a_\tau^i)_{i \in \{1, 2\}}\}_{\tau \in [0, t)})$, where $w \in \{\text{yes, no}\}$. An interpretation is that $w = \text{yes}$ if and only if there is a Poisson hit at time t .

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

1. Let $H^{L, L, w}$ be the set consisting of every history up to any time t that has the form $((t^k)_{k=1}^K, w, \{(a_\tau^i)_{i \in \{1, 2\}}\}_{\tau \in [0, t)})$ with $a_\tau^i \in \{(\bar{z}, \bar{z}), (L_1, \bar{z}), (L_2, \bar{z})\}$ for each $i = 1, 2$ at any $\tau \in [0, t)$ such that $\tau \neq t^k$ for all $k = 1, 2, \dots, K$ and where either of the following holds:

⁵These results follow directly from the proof of proposition 15, and so their proofs are omitted.

- (a) There exists $\tau' \in [0, t)$ such that for each $i = 1, 2$, $a_{\tau'}^i \in \{(L_1, D), (L_1, \bar{z})\}$ and $a_{\tau'}^i \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for $\tau \in [0, t) \setminus \{\tau'\}$.
- (b) There exist $\tau', \tau'' \in [0, t)$ with $\tau' < \tau''$ such that for some $i \in \{1, 2\}$, $a_{\tau'}^i \in \{(L_1, D), (L_1, \bar{z})\}$, $a_{\tau''}^{-i} \in \{(L_2, D), (L_2, \bar{z})\}$ and $a_{\tau'}^i \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for $\tau \in [0, t) \setminus \{\tau'\}$, $a_{\tau''}^{-i} \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for $\tau \in [0, t) \setminus \{\tau''\}$.
2. Let $H^{L,R,w}$ be the set consisting of every history up to any time t that has the form $((t^k)_{k=1}^K, w, \{(a_{\tau}^i)_{i \in \{1,2\}}\}_{\tau \in [0,t)})$ with $a_{\tau}^i \in \{(\bar{z}, \bar{z}), (L_1, \bar{z})\}$ for each $i = 1, 2$ at any $\tau \in [0, t)$ such that $\tau \neq t^k$ for all $k = 1, 2, \dots, K$ and where there exists $\tau' < t$ such that $a_{\tau'}^1 \in \{(L_1, D), (L_1, \bar{z})\}$ and $a_{\tau'}^1 \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for $\tau \in [0, t) \setminus \{\tau'\}$ while $a_{\tau}^2 \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for all $\tau \in [0, t)$.
3. Let $H^{R,L,w}$ be the set consisting of every history up to any time t that has the form $((t^k)_{k=1}^K, w, \{(a_{\tau}^i)_{i \in \{1,2\}}\}_{\tau \in [0,t)})$ with $a_{\tau}^i \in \{(\bar{z}, \bar{z}), (L_1, \bar{z})\}$ for each $i = 1, 2$ at any $\tau \in [0, t)$ such that $\tau \neq t^k$ for all $k = 1, 2, \dots, K$ and where there exists $\tau' < t$ such that $a_{\tau'}^2 \in \{(L_1, D), (L_1, \bar{z})\}$ and $a_{\tau'}^2 \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for $\tau \in [0, t) \setminus \{\tau'\}$ while $a_{\tau}^1 \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for all $\tau \in [0, t)$.
4. Let $H^{R,R,w}$ be the set consisting of every history up to any time t that has the form $((t^k)_{k=1}^K, w, \{(a_{\tau}^i)_{i \in \{1,2\}}\}_{\tau \in [0,t)})$ with $a_{\tau}^i = (\bar{z}, \bar{z})$ for each $i = 1, 2$ at any $\tau \in [0, t)$ such that $\tau \neq t^k$ for all $k = 1, 2, \dots, K$ and where $a_{\tau}^i \in \{(\bar{z}, D), (\bar{z}, \bar{z})\}$ for all $\tau \in [0, t)$ and each $i = 1, 2$.

The feasibility constraints are as follows. For criminal 1,

$$\bar{A}_1(h_t) = \begin{cases} \{(\bar{z}, D), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{L,L,\text{yes}} \cup H^{L,R,\text{yes}} \\ \{(\bar{z}, \bar{z})\} & \text{if } h_t \in H^{L,L,\text{no}} \cup H^{L,R,\text{no}} \\ \{(L_2, D), (L_2, \bar{z}), (\bar{z}, D), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{R,L,\text{yes}} \\ \{(L_2, \bar{z}), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{R,L,\text{no}} \\ \{(L_1, D), (L_1, \bar{z}), (\bar{z}, D), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{R,R,\text{yes}} \\ \{(L_1, \bar{z}), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{R,R,\text{no}} \end{cases} .$$

Similarly, for criminal 2,

$$\bar{A}_2(h_t) = \begin{cases} \{(\bar{z}, D), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{L,L,\text{yes}} \cup H^{R,L,\text{yes}} \\ \{(\bar{z}, \bar{z})\} & \text{if } h_t \in H^{L,L,\text{no}} \cup H^{R,L,\text{no}} \\ \{(L_2, D), (L_2, \bar{z}), (\bar{z}, D), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{L,R,\text{yes}} \\ \{(L_2, \bar{z}), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{L,R,\text{no}} \\ \{(L_1, D), (L_1, \bar{z}), (\bar{z}, D), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{R,R,\text{yes}} \\ \{(L_1, \bar{z}), (\bar{z}, \bar{z})\} & \text{if } h_t \in H^{R,R,\text{no}} \end{cases}.$$

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

$$\bar{\Pi}_i = \{\pi_i : H \rightarrow A_i \mid \pi_i(h_t) \in \bar{A}_i(h_t) \text{ for all } h_t \in H\}.$$

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

The shock process s_t is formally defined as a pair comprising the calendar time t and an indicator $w \in \{\text{yes}, \text{no}\}$ for the existence of a Poisson hit at that time. The instantaneous utility function v_i is specified as follows for each agent $i = 1, 2$:

$$v_i[(a_\tau^1, a_\tau^2), s_\tau] = \begin{cases} \frac{1}{\rho} + g_i(a_\tau^1, a_\tau^2) + h_i(a_\tau^1, a_\tau^2) & \text{if } \tau = 0 \text{ and } w = \text{yes} \\ \frac{1}{\rho} + g_i(a_\tau^1, a_\tau^2) & \text{if } \tau = 0 \text{ and } w = \text{no} \\ [g_i(a_\tau^1, a_\tau^2) + h_i(a_\tau^1, a_\tau^2)]e^{-\rho\tau} & \text{if } \tau > 0 \text{ and } w = \text{yes} \\ g_i(a_\tau^1, a_\tau^2)e^{-\rho\tau} & \text{if } \tau > 0 \text{ and } w = \text{no} \end{cases},$$

where

$$g_i(a_\tau^1, a_\tau^2) = \begin{cases} -\frac{1}{\rho} & \text{if } a_\tau^i \in \{(L_1, D), (L_1, \bar{z})\} \text{ or } a_\tau^{-i} \in \{(L_1, D), (L_1, \bar{z})\} \\ 0 & \text{otherwise} \end{cases}$$

and

$$h_i(a_\tau^1, a_\tau^2) = \begin{cases} c & \text{if } a_\tau^j \in \{(L_1, \bar{z}), (L_2, \bar{z}), (\bar{z}, \bar{z})\} \text{ for each } j = 1, 2 \\ b & \text{if } a_\tau^i \in \{(L_1, D), (L_2, D), (\bar{z}, D)\} \text{ and } a_\tau^{-i} \in \{(L_1, \bar{z}), (L_2, \bar{z}), (\bar{z}, \bar{z})\} \\ p & \text{if } a_\tau^i \in \{(L_1, \bar{z}), (L_2, \bar{z}), (\bar{z}, \bar{z})\} \text{ and } a_\tau^{-i} \in \{(L_1, D), (L_2, D), (\bar{z}, D)\} \\ d & \text{if } a_\tau^j \in \{(L_1, D), (L_2, D), (\bar{z}, D)\} \text{ for each } j = 1, 2 \end{cases}.$$

C.2.2 Proofs

Proof of Proposition 15. The incentive constraint for the optimal Poisson-revision strategy profile is as follows:

$$c + \int_0^\infty 1 \cdot e^{-\rho t} dt + \int_0^\infty \lambda c e^{-\rho t} dt \geq b + \int_0^\infty e^{-\lambda t} 1 \cdot e^{-\rho t} dt + \int_0^\infty \lambda d e^{-\rho t} dt,$$

which by a simple manipulation, can be shown to be equivalent to $\rho(b - c) \leq \lambda/(\rho + \lambda) + \lambda(c - d)$.

Under the optimal Δ -delay strategy profile, there is clearly no incentive for a criminal to start a deviation when there is currently no Poisson hit. Suppose instead that there is a Poisson hit at the current time. If D or L has been played in the past, then it is again easy to see that each criminal has an incentive to follow the prescribed strategy. Assume now that D and L have not been played in the past. The incentive constraint is as follows:

$$c + \int_0^\infty 1 \cdot e^{-\rho t} dt + \int_0^\infty \lambda c e^{-\rho t} dt \geq b + \int_0^\Delta 1 \cdot e^{-\rho t} dt + \int_\Delta^\infty 0 \cdot e^{-\rho t} dt + \int_0^\infty \lambda d e^{-\rho t} dt,$$

which by a simple manipulation, can be shown to be equivalent to $\rho(b - c) \leq e^{-\rho\Delta} + \lambda(c - d)$.

Thus, the optimal Δ -delay strategy profile is an SPE but not the optimal Poisson-revision strategy profile if and only if $\lambda/(\rho + \lambda) + \lambda(c - d) < \rho(b - c) \leq e^{-\rho\Delta} + \lambda(c - d)$, which is satisfied for some $\Delta > 0$ if and only if $\lambda/(\rho + \lambda) < \rho(b - c) - \lambda(c - d) < 1$. \square

C.3 Repeated Technology Adoption⁶

There are two countries, 1 and 2, and a sequence of technologies $\{T^k\}_{k=1}^\infty$. At every time $t \in [0, \infty)$, each country decides whether or not to adopt technology T^k if and only if it has adopted each of the technologies T^1, \dots, T^{k-1} . Let t^k be the first time at which some country adopts T^k , and define t^0 to be 0. If country i adopts T^k , then it receives a private benefit $p > 0$ at the time of adoption, and the other country $-i$ receives an externality $q > 0$ at that time. In addition, country i incurs a cost when it adopts T^k . For each $k \in \mathbb{N}$, the cost is the sum of a base cost $P \in (p, p + q)$ that is time-invariant and a variable cost c_t^k , which evolves according to a geometric Brownian motion: $dc_t^k = \mu c_t^k dt + \sigma c_t^k dz_t$, with the initial condition $c_{t^{k-1}}^k = R$ for some $R \in \mathbb{R}_{++}$ such that $P + R > p + q$.⁷ The payoffs are discounted at rate $\rho > 0$.

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each agent $i = 1, 2$.⁸ Call the game with such strategy spaces the *technology adoption game*. It is characterized by $(p, q, P, R, \mu, \sigma, \rho)$. The analysis in sections 3 and 4 implies that an SPE is well defined.

Note that there is an SPE in which no country adopts any technology because $p < P$. Because there are multiple equilibria with different properties, we focus on the maximal equilibria. A symmetric SPE is said to be maximal if there is no symmetric SPE that yields a higher expected payoff to each agent.

Proposition 16. *The technology adoption game has a maximal equilibrium for any profile $(p, q, P, R, \mu, \sigma, \rho)$. Moreover, there exists $\bar{c} \in \mathbb{R}_+$ such that on the path of play of any maximal equilibrium, technology T^k is adopted with probability one by each country i at the first time that the cost c_t^k reaches \bar{c} . Additionally, the set consisting of each profile $(p, q, P, R, \mu, \sigma, \rho)$ such that $\bar{c} > 0$ is nonempty.*

Remark 7. 1. (Relationship to inertia) In a maximal equilibrium with $\bar{c} > 0$, each agent's strategy is not uniformly inertial but is pathwise inertial. Fix a history on the path of play up to an arbitrary time t such that $t \geq t^{k-1}$ but $c_\tau^k > \bar{c}$ for all $\tau \geq [t^{k-1}, t]$. For any $\epsilon > 0$, there is positive conditional probability that

⁶This example is structurally and analytically similar to the model in Kamada and Rao (2018), which is discussed in section C.1.

⁷Formally, let c_t be a cost process that evolves according to a geometric Brownian motion: $dc_t = \mu c_t dt + \sigma c_t dz_t$, with the initial condition $c_0 = R$ for some $R \in \mathbb{R}_{++}$ such that $P + R > p + q$. For each $k \in \mathbb{N}$, the cost process c_t^k is specified as $c_t^k = \chi_k c_t$ for $t \geq t^{k-1}$, where $\chi_k = R/c_{t^{k-1}}$.

⁸Formal definitions of histories and strategy spaces are provided in section C.3.1.

$c_\tau^k = \bar{c}$ for some $\tau \in (t, t + \epsilon)$, in which case the two countries adopt technology T^k in the time interval $(t, t + \epsilon)$. Thus, there cannot exist $\epsilon > 0$ such that either agent does not move during the time interval $(t, t + \epsilon)$, meaning that uniform inertia fails to hold. However, pathwise inertia holds due to the continuity of the sample path generated by geometric Brownian motion. For any realization of the cost process $\{c_\tau^k\}_{\tau \in (t, \infty)}$ after time t , there exists $\epsilon > 0$ such that $c_\tau^k \neq \bar{c}$ for all $\tau \in (t, t + \epsilon)$, in which case the agents do not move during the time interval $(t, t + \epsilon)$.

2. (Relationship to admissibility) Uniform F1 is not satisfied by a maximal equilibrium with $\bar{c} > 0$ because unboundedly many technologies may be adopted by each country in any proper time interval. However, any maximal equilibrium has pathwise F1. For any time t and any realization of the shock process, the adoption cost decreases by a factor of R/\bar{c} only finitely many times during the time interval $[0, t]$, meaning that each agent adopts only finitely many technologies during this interval. By proposition 3, any traceable and frictional strategy satisfies property F2. F3 is violated by the trigger strategies supporting a maximal equilibrium as any deviation from the prescribed path of play changes the subsequent pattern of technology adoption.

C.3.1 Formal Definitions of Histories and Strategy Spaces

Choose any time $t \in [0, \infty)$ and underlying cost process $\{c_\tau\}_{\tau \in [0, t]}$ up to that time. A history up to time t is represented by $(\{c_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in \{1, 2\}}\}_{\tau \in [0, t]})$, where $\{a_\tau^i\}_{\tau \in [0, t]}$ denotes the action path of country $i \in \{1, 2\}$ up to time t with the action space being $\mathbb{R}_{++} \cup \{z\}$. The number of technologies adopted by country $i \in \{1, 2\}$ in the time interval $[0, t)$ is denoted by l_i , which is the number of elements in the set $\{\tau \in [0, t) : a_\tau^i \in \mathbb{R}_{++}\}$.

The set of all histories up to an arbitrary time is denoted by H . Choose an arbitrary $i \in \{1, 2\}$ as well as any $h_t \in H$. If $l_j < \infty$ for each $j \in \{1, 2\}$, then define $\tau^* \in [0, t)$ as follows:

1. If $l_i > l_{-i}$, then τ^* is the unique value of $\tau \in [0, t)$ such that $a_\tau^i \in \mathbb{R}_{++}$ and $a_{\tau'}^i = z$ for all $\tau' \in (\tau, t)$.
2. If $l_i = l_{-i}$, then $\tau^* = \min\{\tau^{1*}, \tau^{2*}\}$, where for each $j \in \{1, 2\}$, τ^{j*} is the unique value of $\tau \in [0, t)$ such that $a_\tau^j \in \mathbb{R}_{++}$ and $a_{\tau'}^j = z$ for all $\tau' \in (\tau, t)$.

3. If $l_i < l_{-i}$, then for each $j \in \{1, 2\}$, there exists a unique set of times $\{\tau^{j,1}, \dots, \tau^{j,l_i}\}$ such that $a_{\tau^{j,k}}^j \in \mathbb{R}_{++}$ for each $k = 1, \dots, l_i$ and $a_{\tau'}^j = z$ for all $\tau' \in [0, \tau^{j,l_i}) \setminus \{\tau^{j,1}, \dots, \tau^{j,l_i-1}\}$. Let $\tau^* = \min\{\tau^{1,l_1}, \tau^{2,l_2}\}$.

Whenever $l_j < \infty$ for each $j \in \{1, 2\}$, the feasibility constraint is $\bar{A}_i(h_t) = \{r, z\}$, where $r = c_{\tau^*}$. Otherwise, if $l_j = \infty$ for some $j \in \{1, 2\}$, then let $\bar{A}_i(h_t) = \{z\}$.

The set of feasible strategies is for each $i \in \{1, 2\}$:

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

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

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

$$v_i[(a_\tau^i, a_\tau^{-i}), s_\tau] = \begin{cases} [p - (P + \frac{R}{r}c_\tau)]e^{-\rho\tau} & \text{if } a_\tau^i = r \in \mathbb{R}_{++} \text{ and } a_\tau^{-i} = z \\ qe^{-\rho\tau} & \text{if } (a_\tau^i, a_\tau^{-i}) \in \{z\} \times \mathbb{R}_{++} \\ [p + q - (P + \frac{R}{r}c_\tau)]e^{-\rho\tau} & \text{if } a_\tau^i = r \in \mathbb{R}_{++} \text{ and } a_\tau^{-i} \in \mathbb{R}_{++} \\ 0 & \text{if } (a_\tau^i, a_\tau^{-i}) = (z, z) \end{cases}.$$

C.3.2 Proofs

Proof of Proposition 16. As with Proposition 2, the proof consists of three parts. We first assume the Markov property on the path of play and solve for the unique optimum.⁹ 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 17. *For any profile $(p, q, P, R, \mu, \sigma, \rho)$, the technology adoption game has a symmetric SPE that is Markov on the path of play and weakly Pareto dominates any*

⁹A 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 cost c_t^k at that time where $k - 1$ is the number of technologies that have been adopted by each country up to then.

symmetric SPE that is Markov on the path of play. Moreover, there exists $\bar{c} \in \mathbb{R}_+$ such that on the path of play of any such SPE, technology T^k is adopted with probability one by each country i at the first time that the cost c_t^k reaches \bar{c} . Additionally, the set consisting of each profile $(p, q, P, R, \mu, \sigma, \rho)$ such that $\bar{c} > 0$ is nonempty.

Proof. At any history up to a given time, the least continuation payoff that each country can receive in an SPE is zero, which is achieved when each country follows a strategy of not making any further technology adoptions. Since $P > p$, it is an SPE for each country to never adopt any technology. Hence, assuming the Markov property on the path of play, there exists $\bar{c} \in (0, p + q)$ such that any symmetric SPE that maximizes the expected payoff of each agent has the following properties. With probability one, technology T_k is adopted at time t if and only if the history up to time t meets all of the following conditions:

1. There exists $k \in \mathbb{N}$ such that T^{k-1} has been adopted by both countries but T^k has not been adopted by either country.
2. For each $l = \{1, \dots, k-1\}$, T^l was adopted by both countries at a time t^l such that $c_\tau^l > \bar{c}$ for all $\tau \in (t^{l-1}, t^l)$ and $c_{t^l}^l = \bar{c}$.
3. $c_\tau^k > \bar{c}$ for all $\tau \in (t^{k-1}, t)$ and $c_t^k = \bar{c}$.

The expected payoff V of each agent in such an equilibrium is the value of an asset that pays $(p + q) - (P + \bar{c}) + V$ at the first time that the cost c_t reaches \bar{c} when the current cost is R . Letting $\beta = \frac{1}{2} - \mu/\sigma^2 - \sqrt{(\mu/\sigma^2 - \frac{1}{2})^2 + 2\rho/\sigma^2} < 0$, the value V satisfies the equation:

$$V = [(p + q) - (P + \bar{c}) + V](R/\bar{c})^\beta, \quad (6)$$

which yields the following expression for V :

$$V = (R/\bar{c})^\beta [(p + q) - (P + \bar{c})] / [1 - (R/\bar{c})^\beta]. \quad (7)$$

Note that \bar{c} must satisfy the incentive constraint $(p + q) - (P + \bar{c}) + V \geq q$, which is equivalent to:

$$V \geq (P - p) + \bar{c}. \quad (8)$$

We consider the problem of choosing the threshold $\bar{c} \in [0, p + q - P]$ to maximize the objective function in equation (7) given the constraint in equation (8). We show that the maximization problem has a unique solution whenever the constraint is satisfied for some \bar{c} . To do so, we first show that the log of the objective function is concave in \bar{c} whenever the objective function is nonincreasing, and we next show that the set containing each value of \bar{c} that satisfies the constraint is an interval. It is also noted that there exist parameter values such that this interval is nonempty.

The derivative of the log of the objective function with respect to \bar{c} is:

$$-1/[(p + q) - (\bar{c} + P)] + \beta/\{\bar{c}[(c_0/\bar{c})^\beta - 1]\},$$

which is nonpositive if and only if $-1/[(p + q) - (\bar{c} + P)] \leq -\beta/\{\bar{c}[(c_0/\bar{c})^\beta - 1]\}$. The second derivative of the log of the objective function with respect to \bar{c} is:

$$-1/[(p + q) - (\bar{c} + P)]^2 + [\beta + (c_0/\bar{c})^\beta(\beta - 1)\beta]/\{\bar{c}^2[(c_0/\bar{c})^\beta - 1]^2\},$$

which is no greater than the following whenever the derivative of the log of the objective function is nonpositive:

$$-\beta^2/\{\bar{c}^2[(c_0/\bar{c})^\beta - 1]^2\} + [\beta + (c_0/\bar{c})^\beta(\beta - 1)\beta]/\{\bar{c}^2[(c_0/\bar{c})^\beta - 1]^2\}. \quad (9)$$

Expression (9) has the same sign as $-\beta^2 + [\beta + (c_0/\bar{c})^\beta(\beta - 1)\beta] = [1 - (c_0/\bar{c})^\beta](\beta - \beta^2)$, which is negative. It follows that the log of the objective function is concave whenever the objective function is nonincreasing.

The constraint is equivalent to:

$$(c_0/\bar{c})^\beta q + p - (\bar{c} + P) \geq 0. \quad (10)$$

The first derivative of the expression on the left-hand side of (10) with respect to \bar{c} is $-\bar{c} + (c_0/\bar{c})^\beta q\beta/\bar{c}$, and the second derivative of the expression on the left-hand side of (10) with respect to \bar{c} is $(c_0/\bar{c})^\beta q\beta(1 + \beta)/\bar{c}^2$, which is negative for $\beta \in (-1, 0)$. The constraint is never satisfied for $\beta \leq -1$. It follows that the set containing each value of \bar{c} that satisfies the constraint is an interval.

The constraint is not satisfied for $\bar{c} = 0$ or $\bar{c} = p + q - P$. However, note that for any $\bar{c} \in (0, p + q - P)$, there exists $\bar{\beta} < 0$ such that the constraint is satisfied for

$\beta \in [\bar{\beta}, 0)$. This shows that there exist parameter values such that the aforementioned interval is nonempty. ■

Lemma 18. *Up to zero probability events, any maximal equilibrium must be Markov on the path of play.*

Proof. Suppose that there exists a maximal equilibrium. Let V denote the expected payoff to each agent at the null history when a maximal equilibrium is played. If $V = 0$, then there is a unique equilibrium in which no technology adoption occurs, and so the claim holds. Therefore, assume that $V > 0$, in which case technology adoption must occur with positive probability on the path of play in a maximal equilibrium. Note also that the continuation payoff to each country after the adoption of T^1 cannot differ from V with positive probability in a maximal equilibrium.

Now consider the following constrained optimization problem. The expression on the right-hand side of equation (6) is maximized with respect to \bar{c} subject to the constraint in (8), where V is treated as a constant. This problem has a unique maximizer c^* . In particular, there must exist $k > 0$ such that the constraint is satisfied for $\bar{c} \in [0, k]$. Otherwise, it would be impossible for an SPE to exist in which technology adoption occurs, contradicting the assumption that $V > 0$. Moreover, it can be shown that the derivative of the maximand with respect to \bar{c} is positive in the limit as \bar{c} approaches zero and changes sign only once as \bar{c} increases from zero.

This shows that any maximal equilibrium must be Markov on the equilibrium path up to the adoption of T^1 , which happens with probability one at the first time the cost reaches c^* . A similar argument can be applied to the adoption of T^2 , and so on. ■

Lemma 19. *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 18.

The value V cannot be greater than the value of the following constrained optimization problem \mathcal{M} . The value of an asset at cost R that pays $(p + q) - (P + c) + V$

at the first time that the cost reaches c is maximized with respect to c subject to the constraint $V \geq (P - p) + c$. If $V \leq P - p$, then there is a unique symmetric SPE in which no technology adoption occurs, and so the claim holds, with $V = 0$. Therefore, assume that $V > P - p$. Let c^* denote the maximizer in problem \mathcal{M} . There are two cases to consider. In the first case, the constraint in problem \mathcal{M} is not binding. In the second case, the constraint in problem \mathcal{M} 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 $V - \epsilon$. On the equilibrium path, the adoption of T^1 occurs at the first time the cost reaches the threshold c^* , and the agents after the adoption of T^1 play a strategy profile that yields a continuation payoff W that does not depend on the history up to the time when T^1 is adopted. Let Y denote the expected payoff to each agent at the null history when playing strategy profile ϕ_1 . Note that $V - Y \leq V - W$ because the behavior up to the adoption of T^1 when playing strategy profile ϕ_1 is the same as the behavior in problem \mathcal{M} .

Since $Y > W$, there exists a symmetric SPE ϕ_2 in grim-trigger strategies with the following properties that yields an expected payoff greater than $V - \epsilon$. On the equilibrium path, the adoption of T^1 occurs at the first time the cost reaches the threshold c^* , the adoption of T^2 occurs at the first time after the adoption of T^1 that the cost reaches the threshold c^* , and the agents after the adoption of T^2 play a strategy profile that yields a continuation payoff W that does not depend on the history up to the time when T^2 is adopted. In particular, the agents start by playing ϕ_1 , and then after any history up to an arbitrary time on the equilibrium path after the adoption of T^1 , the agents play ϕ_1 behaving as if the game just started after the adoption of T^1 .

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 $V - \epsilon$. For any positive integer m , the adoption of T^m occurs on the equilibrium path at the first time after the adoption of T^{m-1} that the cost reaches the threshold c^* , where the adoption of T^0 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 that 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 that yields an expected payoff

Y_1 greater than $V - \epsilon$. On the equilibrium path, the adoption of T^1 occurs at the first time the cost reaches the threshold c_1 , and the agents after the adoption of T^1 play a strategy profile that yields a continuation payoff W_1 that does not depend on the history up to the time when T^1 is adopted. Moreover, because the constraint in problem \mathcal{M} is binding, the threshold c_1 can be chosen such that $W_1 = (P - p) + c_1$ by choosing c_1 to maximize the expected payoff under ψ_1 given the continuation payoff W_1 .

Applying such an argument to any subgame after the adoption of T^1 on the equilibrium path, there exists a symmetric SPE ψ'_2 in grim-trigger strategies with the following properties that yields an expected payoff greater than $V - \epsilon$. On the equilibrium path, the adoption of T^1 occurs at the first time the cost reaches the threshold c_1 , the adoption of T^2 occurs at the first time after the adoption of T^1 that the cost reaches a threshold c_2 , and the agents after the adoption of T^2 play a strategy profile that yields a continuation payoff W_2 that does not depend on the history up to the time when T^2 is adopted. Moreover, because the constraint in problem \mathcal{M} is binding, the threshold c_2 can be chosen such that $W_2 = (P - p) + c_2$ by choosing c_2 to maximize the expected payoff under ψ'_2 given the first threshold c_1 and the continuation payoff W_2 . Let Y_2 be the continuation payoff after the adoption of T^1 on the equilibrium path when playing ψ'_2 .

Note that $W_1 > W_2$ if $c_1 > c_2$, $W_1 < W_2$ if $c_1 < c_2$, and $W_1 = W_2$ if $c_1 = c_2$. It must also be that $Y_1 > Y_2$ if $c_1 > c_2$, $Y_1 < Y_2$ if $c_1 < c_2$, and $Y_1 = Y_2$ if $c_1 = c_2$. If $c_2 > c_1$, then let $\psi_2 = \psi'_2$. If $c_2 \leq c_1$, then let ψ_2 be the strategy profile in which the agents start by playing ψ_1 , and then after any history up to an arbitrary time on the equilibrium path after the adoption of T^1 , the agents play ψ_1 behaving as if the game just started after the adoption of T^1 .

Continuing in this way, one can show that there exists a symmetric SPE ψ in grim-trigger strategies with the following properties that yields an expected payoff greater than $V - \epsilon$. For any positive integer m , the adoption of T^m occurs on the equilibrium path at the first time after the adoption of T^{m-1} that the cost reaches the threshold c_m , where the adoption of T^0 is said to occur at time 0. Moreover, c_m is nondecreasing in m , and the continuation payoff Q_m after the adoption of T^m is greater than $V - \epsilon$.

Let x denote the limit of the sequence $\{c_m\}$. Consider the grim-trigger strategy profile ξ in which for any positive integer m , the adoption of T^m occurs on the path

of play at the first time after the adoption of T^{m-1} that the cost reaches the threshold x , where the adoption of T^0 is said to occur at time 0. The expected payoff B under strategy profile ξ is no less than $V - \epsilon$ because $Q_m \geq V - \epsilon$ for all m , where B is the limit of the sequence $\{Q_m\}$. Moreover, the incentive constraint $B \geq (P - p) + x$ is satisfied because the incentive constraint $Q_m \geq (P - p) + c_m$ 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 that yields an expected payoff arbitrarily close to V . ■

□

C.4 Inventory Restocking Model

There is a retailer R and a distributor D of a good. At each moment of time $t \in [0, \infty)$, the retailer chooses between actions B and z , and the distributor chooses between actions S and z . The action B means that the retailer visits the distributor to buy the good, and z stands for the retailer not doing so. The action S means that the distributor is open to sell the good to the retailer, and z stands for the distributor being closed.

The retailer has a capacity constraint $f > 0$ on the amount of the good it can keep in stock, and the initial value of the stock is f . The good depreciates at the rate $\delta > 0$ in the inventory of the retailer. If B and S are simultaneously chosen at time t , then the retailer replenishes its stock so that its inventory reaches f , and the retailer pays the distributor a flat fee $p > 0$, which is exogenous and independent of the quantity purchased.¹⁰ Otherwise, the retailer and distributor do not transact at time t , in which case they do not incur any cost at that time.

A customer comes to the retailer according to a Poisson process with arrival rate $\lambda > 0$. Any customer that comes buys the entire stock that the retailer keeps, so that the stock reaches zero upon the arrival of the customer.¹¹ Let q be an exogenous unit price that the retailer charges a customer for the good. When a customer arrives, the retailer receives a revenue equal to q times the supply available at that time.¹² The

¹⁰It is assumed without loss of generality that the distributor can produce and supply the good at zero cost. The equilibria that we characterize would not change if the distributor were to face a constant cost strictly less than the price.

¹¹The actions of the retailer and distributor at time t are taken after learning whether a customer has arrived at time t .

¹²The supply available at time $t > 0$ is defined as the limit of the stock as time approaches t from the left. With probability one, such a left-hand limit exists at every time $t > 0$ given the traceability

discount rate is $\rho > 0$.

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each agent $i = R, D$.¹³ Call the game with such strategy spaces the *restocking game*. It is characterized by $(f, p, q, \delta, \lambda, \rho)$. The analysis in sections 3 and 4 implies that an SPE is well defined.

A Markov perfect equilibrium is defined as an SPE in Markov strategies. A strategy is said to be Markov if the action prescribed at any history up to a given time depends only on the stock at the current time. As in section C.5, the model has multiple equilibria. We consider the Markov perfect equilibrium that maximizes the expected payoff of the retailer.

Proposition 20. *For any profile $(f, \delta, \lambda, \rho)$, there exists $\alpha < \infty$ such that for any $q > \alpha p$, any Markov perfect equilibrium that maximizes the retailer's expected payoff in the restocking game with $(f, p, q, \delta, \lambda, \rho)$ satisfies the following. There exists $k \in (0, f)$ such that:*

1. *If the current stock level is 0 or k , then the retailer chooses B , and the distributor chooses S .*
2. *If the current stock level is greater than k , then the retailer chooses z .*

When the unit price q paid by consumers is high relative to the cost p of restocking the good, the retailer obtains a high payoff when the consumer buys a large quantity of the good, and the cost of maintaining a large supply of the good is relatively low. Therefore, the retailer has an incentive to replenish its inventory of the good even if the current stock is not zero. The distributor is willing to be open since it can do so at zero cost.

Remark 8. 1. (Relationship to inertia) In a Markov perfect equilibrium that maximizes the retailer's expected payoff, uniform inertia is violated. For any $\epsilon > 0$, there is positive probability that a customer arrives at some $\tau \in (t, t + \epsilon)$, in which case the retailer and distributor respectively take actions B and S in the time interval $(t, t + \epsilon)$. Thus, there cannot exist $\epsilon > 0$ such that either firm does not move during the time interval $(t, t + \epsilon)$. However, each agent's strategy is pathwise inertial in such an equilibrium. For any time t and any realization of

and frictionality assumptions on strategies.

¹³Formal definitions of histories and strategy spaces are provided in section C.4.1.

the Poisson process, there exists $\epsilon > 0$ such that the stock level is not equal to 0 or k at any time $\tau \in (t, t + \epsilon)$, and so the agents are not required to move during the time interval $(t, t + \epsilon)$.

2. (Relationship to admissibility) Uniform F1 is not satisfied by a Markov perfect equilibrium that maximizes the retailer's expected payoff. The number of moves is not uniformly bounded because the inventory may be restocked unboundedly many times in any proper time interval. Pathwise F1 is also violated. The requirements of pathwise F1 hold on but not off the path of play. Given any times t and \hat{t} with $\hat{t} < t$, any positive integer r , and any realization of the Poisson process with at least two arrival times in the interval $[\hat{t}, t]$, there exists a history up to time t such that the retailer and distributor have each chosen B or S no less than r times in the time interval $[0, \hat{t})$. Since the retailer and distributor respectively choose B and S at the time of a Poisson hit, there is no upper bound on the number of times each of them moves during the time interval $[0, t]$.

It follows from proposition 3 that property F2 applies.

The strong continuity assumption F3 is violated. Suppose that no customer arrives after time \hat{t} . When the retailer and distributor choose z at time \hat{t} and respectively choose B and S at time $\hat{t} + \epsilon$ with $\epsilon > 0$, the next restocking of the good occurs later than when the retailer and distributor respectively choose B and S at time \hat{t} . Hence, a small difference in the timing of moves affects future behavior.

3. (Non- z action at a time without a Poisson hit) The agents' strategies in the Markov perfect equilibrium described in proposition 20 would not satisfy a condition requiring that a non- z action be taken only at the times of discrete changes in the shock. Although the retailer and distributor transact at any time that the available supply reaches k , there is zero probability of a discrete change in the shock at such a time, where the set of times at which the shock discretely changes is defined as the set of Poisson arrival times.

C.4.1 Formal Definitions of Histories and Strategy Spaces

Choose any time $t \in [0, \infty)$ and sequence $(t^k)_{k=1}^K$ of past arrival times of a customer. A history up to time t is represented by $((t^k)_{k=1}^K, w, \{(a_\tau^i)_{i \in \{R, D, G\}}\}_{\tau \in [0, t)})$, where $\{(a_\tau^i)_{\tau \in [0, t)}$ denotes the action path of agent $i \in \{R, D\}$ up to time t with the action spaces being $\{B, z\}$ for R and $\{S, z\}$ for D . There is also a customer G , whose action path up to time t is $\{a_\tau^G\}_{\tau \in [0, t)}$ with action space $\mathbb{R}_+ \cup \{z\}$. A move by agent G represents the amount bought, and action z means no arrival by agent G . The term $w \in \{\text{yes}, \text{no}\}$ indicates whether or not a customer arrives at time t .

For any $u \leq t$, let $t^{B,S}(u)$ be the maximum of zero and the supremum of the set consisting of every time $\tau < u$ such that $a_\tau^R = B$ and $a_\tau^D = S$. For any $u \leq t$, let $t^G(u)$ be the maximum of zero and the supremum of the set consisting of every time $\tau < u$ such that $a_\tau^G \neq z$. Define the available supply of the good at time u as $x_u = fe^{-\delta[u-t^{B,S}(u)]}$ if $t^{B,S}(u) \geq t^G(u)$ and as $x_u = 0$ if $t^{B,S}(u) < t^G(u)$.

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

1. For any $c \in [0, f]$, let $H^{\text{yes}, c}$ be the set consisting of every history up to any time t that has the form $((t^k)_{k=1}^K, \text{yes}, \{(a_\tau^i)_{i \in \{R, D, G\}}\}_{\tau \in [0, t)})$ with $x_t = c$ where $a_\tau^G = z$ at any $\tau \in [0, t)$ such that $\tau \neq t^k$ for all $k \in \{1, 2, \dots, K\}$ and where $a_\tau^G = x_\tau$ at any $\tau \in [0, t)$ such that $\tau = t^k$ for some $k \in \{1, 2, \dots, K\}$.
2. Let H^{no} be the set consisting of every history up to any time t that has the form $((t^k)_{k=1}^K, \text{no}, \{(a_\tau^i)_{i \in \{R, D, G\}}\}_{\tau \in [0, t)})$ where $a_\tau^G = z$ at any $\tau \in [0, t)$ such that $\tau \neq t^k$ for all $k \in \{1, 2, \dots, K\}$ and where $a_\tau^G = x_\tau$ at any $\tau \in [0, t)$ such that $\tau = t^k$ for some $k \in \{1, 2, \dots, K\}$.

The feasibility constraints are as follows. For $i \in \{R, D\}$, $\bar{A}_R(h_t) = \{B, z\}$ and $\bar{A}_D(h_t) = \{S, z\}$, where $h_t \in H$. For $i = G$, $\bar{A}_G(h_t) = \{c\}$ if there exists $c \in \mathbb{R}_+$ such that $h_t \in H^{\text{yes}, c}$, and $\bar{A}_G(h_t) = \{z\}$ if $h_t \in H^{\text{no}}$.

The sets of feasible strategies are:

$$\begin{aligned} \bar{\Pi}_R &= \{\pi_R : H \rightarrow \{B, z\} \mid \pi_R(h_t) \in \bar{A}_R(h_t) \text{ for all } h_t \in H\} \\ \bar{\Pi}_D &= \{\pi_D : H \rightarrow \{S, z\} \mid \pi_D(h_t) \in \bar{A}_D(h_t) \text{ for all } h_t \in H\} \\ \bar{\Pi}_G &= \{\pi_G : H \rightarrow \mathbb{R}_+ \cup \{z\} \mid \pi_G(h_t) \in \bar{A}_G(h_t) \text{ for all } h_t \in H\} \end{aligned}$$

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

The shock process s_t is formally defined as a pair comprising the calendar time t and an indicator $w \in \{\text{yes, no}\}$ for the existence of a Poisson hit at that time. The instantaneous utility function v_i is specified as follows for $i = R$:

$$v_R[(a_\tau^R, a_\tau^D, a_\tau^G), s_\tau] = \begin{cases} (qc - p)e^{-\rho\tau} & \text{if } (a_\tau^R, a_\tau^D) = (B, S) \text{ and } a_\tau^G = c \in \mathbb{R}_+ \\ -pe^{-\rho\tau} & \text{if } (a_\tau^R, a_\tau^D, a_\tau^G) = (B, S, z) \\ qce^{-\rho\tau} & \text{if } (a_\tau^R, a_\tau^D) \neq (B, S) \text{ and } a_\tau^G = c \in \mathbb{R}_+ \\ 0 & \text{otherwise} \end{cases},$$

and as follows for $i = D$:

$$v_D[(a_\tau^R, a_\tau^D, a_\tau^G), s_\tau] = \begin{cases} pe^{-\rho\tau} & \text{if } (a_\tau^R, a_\tau^D) = (B, S) \\ 0 & \text{otherwise} \end{cases},$$

with v_G being arbitrarily defined.

C.4.2 Proofs

Proof of Proposition 20. Consider any strategy profile in which (S, B) is never chosen when the current stock is 0. The retailer's continuation payoff is 0 from following such a strategy profile when the current stock is 0. Let V denote the retailer's continuation payoff when the current stock is 0 from following a Markov strategy profile in which (S, B) is chosen if and only if the current stock is zero. The value V satisfies:

$$V = \int_0^\infty (qfe^{-\delta x} + V)(e^{-\rho x})(\lambda e^{-\lambda x})dx - p,$$

which gives $V = (\lambda + \rho)[fq\lambda - p(\delta + \lambda + \rho)]/[\rho(\delta + \lambda + \rho)]$. Note that for any profile $(f, \delta, \lambda, \rho)$, there exists $\phi < \infty$ such that $V > 0$ whenever $q > \phi p$.

Consider a Markov strategy profile in which (S, B) is chosen if and only if the current stock is 0. For any $\epsilon \in (0, 1)$, the following is the retailer's continuation payoff from following such a strategy profile when the current stock is ϵf :

$$\int_0^\infty (q\epsilon fe^{-\delta x} + W - p)(e^{-\rho x})(\lambda e^{-\lambda x})dx = q\epsilon f\lambda/(\delta + \lambda + \rho) + \lambda(W - p)/(\lambda + \rho),$$

where W is the retailer's expected payoff when the current stock is f from following a strategy profile in which (S, B) is chosen if and only if the current stock is 0. If the agents instead follow a strategy profile in which (S, B) is chosen at the current time and thereafter (S, B) is chosen if and only if the current stock is 0, then the following is the retailer's continuation payoff:

$$\int_0^\infty (qf e^{-\delta x} + W)(e^{-\rho x})(\lambda e^{-\lambda x}) dx - p = qf\lambda/(\delta + \lambda + \rho) + \lambda W/(\lambda + \rho) - p.$$

The latter expression is greater than the former if and only if $q/p > \rho(\delta + \lambda + \rho)/[f\lambda(\lambda + \rho)(1 - \epsilon)]$.

Hence, for any profile $(f, \delta, \lambda, \rho)$, one can find $\psi < \infty$ such that if $q > \psi p$, then there exists $\epsilon \in (0, 1)$ such that for all $\eta \in (0, \epsilon]$, the retailer obtains a higher continuation payoff when the current stock is ηf from a strategy profile in which (S, B) is chosen at the current time and thereafter (S, B) is chosen if and only if the current stock is 0 than from a Markov strategy profile in which (S, B) is chosen if and only if the current stock is 0. Iteratively applying this argument, it can be shown that for any profile $(f, \delta, \lambda, \rho)$, one can find $\psi < \infty$ such that if $q > \psi p$, then there exists $\epsilon \in (0, 1)$ such that for all $\eta \in (0, \epsilon]$, the retailer obtains a higher expected payoff when the current stock is ηf from a Markov strategy profile in which (S, B) is chosen if and only if the current stock is ηf or 0 than from a Markov strategy profile in which (S, B) is chosen if and only if the current stock is 0.

Fixing any profile $(f, \delta, \lambda, \rho)$, the preceding argument implies that one can find $\alpha < \infty$ such that if $q > \alpha p$, then there exists $\epsilon \in (0, 1)$ such that the following three statements hold for any $\eta \in (0, \epsilon]$: (i) There is a Markov perfect equilibrium in which (S, B) is chosen if and only if the current stock is 0. Denote by π^0 this Markov perfect equilibrium. (ii) There is a Markov perfect equilibrium in which (S, B) is chosen if and only if the current stock is 0 or ηf . Denote by π^η this Markov perfect equilibrium. (iii) The Markov perfect equilibrium π^0 yields a lower expected payoff to the retailer than π^η for $\eta < \epsilon$ and the same expected payoff to the retailer as π^η for $\eta = \epsilon$.

Consider a Markov perfect equilibrium in which (S, B) is chosen if and only if the current stock is 0 or ηf , where $\eta \in [0, \epsilon]$. Since the expected payoff to the retailer in such an equilibrium is a continuous function of η , it follows from the extreme value theorem that there exists $\omega \in [0, \epsilon]$ such that a strategy profile in which (S, B) is chosen if and only if the current stock is 0 or ωf is a Markov perfect equilibrium that

maximizes the expected payoff of the retailer. Noting that $\omega > 0$, the proposition follows from setting $k = \omega f$. \square

C.5 Retailer and Distributor with a Single Unit

The example below involves a retailer R and a distributor D who make decisions in continuous time. The distributor chooses when to open its salesroom for the retailer to buy a good, while the retailer chooses when to visit the distributor and when to sell the good to a consumer.

At each time $t \in [0, \infty)$, the distributor chooses between the actions S and z , and the retailer chooses among the actions B , C , V , and z . The action S means that the distributor is open to sell the good, whereas z stands for being closed. The action B means that the retailer visits the distributor to buy the good for inventory. When the retailer has already obtained but not yet sold the good, the action C means that the retailer sells the good to the consumer. When the retailer has not yet obtained the good from the distributor, the action V means that the retailer visits the distributor to buy the good for immediate sale to the consumer. The action z by the retailer stands for doing nothing. The retailer acquires the good at time t if and only if the distributor chooses S and the retailer chooses B or V at that time.

When the retailer buys the good, it pays an exogenously fixed price $p > 0$ to the distributor. The price q_t at which the retailer can sell the good to the consumer evolves according to a geometric Brownian motion: $dq_t = \mu q_t dt + \sigma q_t dz_t$, with initial condition $q_0 = \tilde{q}$ for some $\tilde{q} \in \mathbb{R}_{++}$. At the time of sale, the retailer incurs a fixed cost $c > 0$ for packaging the product for sale to the consumer. Both the distributor and retailer discount the future at rate $\rho > 0$.

Defining $\beta = \frac{1}{2} - \mu/\sigma^2 + \sqrt{(\mu/\sigma^2 - \frac{1}{2})^2 + 2\rho/\sigma}$, attention is restricted to the case where $\beta \in (1, 1 + c/p)$. The condition $\beta > 1$ ensures that the model has an SPE in which the retailer sells the good to the consumer with positive probability. The further restriction $\beta < 1 + c/p$ enables us to solve for an SPE of the form described in proposition 21, where there is a delay between when the retailer buys the good from the distributor and sells the good to the consumer.

We consider traceable, frictional, calculable, and feasible strategies, denoted $\bar{\Pi}_i^C$ for each agent $i = R, D$.¹⁴ Call the game with such strategy spaces the *retailer-distributor game*. It is characterized by $(p, c, \mu, \sigma, \rho)$. The analysis in sections 3 and

¹⁴Formal definitions of histories and strategy spaces are provided in section C.5.1.

4 implies that an SPE is well defined.

This model has multiple SPE because the retailer takes possession of the good from the distributor if and only if the two parties coordinate their behavior so that the retailer makes a visit at a time when the distributor is open. For this reason, we restrict attention to the SPE that maximizes the distributor's expected payoff.¹⁵

Proposition 21. *In the retailer-distributor game with $(p, c, \mu, \sigma, \rho)$, an SPE maximizing the distributor's expected payoff exists, and there exist ν_1 and ν_2 with $0 < \nu_1 < \nu_2 < \infty$ such that in any SPE that maximizes the distributor's expected payoff, the following hold on the path of play at any history up to an arbitrary time t .*

1. *If $q_t < \nu_1$, then the retailer chooses z .*
2. *If $q_t \in [\nu_1, \nu_2)$ and $q_\tau < \nu_1$ for all $\tau < t$, then the retailer chooses B , and the distributor chooses S .*
3. *If $q_t \in [\nu_1, \nu_2)$ and $q_\tau \geq \nu_1$ for some $\tau < t$, then the retailer and distributor each choose z .*
4. *If $q_t \geq \nu_2$ and $q_\tau < \nu_2$ for all $\tau < t$, then the retailer chooses C for $t > 0$ and V for $t = 0$, and the distributor chooses z for $t > 0$ and S for $t = 0$.*
5. *If $q_t \geq \nu_2$ and $q_\tau \geq \nu_2$ for some $\tau < t$, then the retailer and distributor each choose z .*

Remark 9. 1. (Relationship to inertia) In any SPE that maximizes the distributor's expected payoff, each agent's strategy violates uniform inertia if $q_0 < \nu_1$. To see this, fix a history on the equilibrium path up to an arbitrary time t with $q_t < \nu_1$ such that the firms have not yet transacted. For any $\epsilon > 0$, there is positive conditional probability that $q_\tau = \nu_1$ for some $\tau \in (t, t + \epsilon)$, in which case the retailer takes action B and the distributor takes action S in the time interval $(t, t + \epsilon)$. Thus, there cannot exist $\epsilon > 0$ such that either firm does not move during the time interval $(t, t + \epsilon)$. Similarly, the requirement that the retailer choose C at the first time on the equilibrium path that q_t reaches ν_2 also results in a violation of uniform inertia.

¹⁵In item 4 of remark 9, we consider the SPE that maximizes the retailer's expected payoff.

However, due to the continuity of the sample path of geometric Brownian motion, pathwise inertia is satisfied in an SPE that maximizes the distributor's expected payoff. To see this, consider any history up to time t and any realization of the retail price process $\{q_\tau\}_{\tau \in (t, \infty)}$ after time t . If $q_t < \nu_1$, then there exists $\epsilon > 0$ such that $q_\tau \neq \nu_1$ for all $\tau \in (t, t + \epsilon)$. If $\nu_1 \leq q_t < \nu_2$, then there exists $\epsilon > 0$ such that $q_\tau \neq \nu_2$ for all $\tau \in (t, t + \epsilon)$. In each case, the retailer and the distributor are not required to move during the time interval $(t, t + \epsilon)$. If $\nu_2 \leq q_t$, then there is no $\epsilon > 0$ such that the agents have to move during the time interval $(t, t + \epsilon)$.

2. (Relationship to admissibility) Noting that the retailer moves at most twice and the distributor moves at most once in equilibrium, there exists an SPE maximizing the distributor's expected payoff that satisfies both the uniform and pathwise versions of criterion F1. Condition F2 is satisfied given proposition 3. However, criterion F3 is violated. Suppose that $\hat{t} > 0$ is the least time t such that $q_t \geq \nu_2$. If the only non- z actions before time \hat{t} were that the retailer chose B and the distributor chose S at time $\bar{t} < \hat{t}$, then the retailer would choose C at time \hat{t} in equilibrium. If the only non- z actions before time \hat{t} were that the retailer chose B at time $\bar{t} < \hat{t}$ and the distributor chose S at time $\bar{t} + \epsilon < \hat{t}$ with $\epsilon > 0$, then the retailer could not choose C at time \hat{t} . Hence, the strong continuity requirement F3 does not hold because a small difference in the timing of past moves affects current behavior.
3. (Comparative statics) The threshold ν_1 is increasing in p , c , and ρ while being decreasing in μ and σ . The reason is that ν_1 is chosen so as to make the retailer indifferent between buying the good when the current retail price is q_t and never buying the good. If p or c increases, then the cost to the retailer of procuring or packaging the good becomes higher, making the retailer reluctant to buy the good unless it can be sold to the consumer at a higher price. When μ or σ increases, the prospect of a high retail price q_t in the future becomes better, so the retailer is more willing to buy the good. When ρ is high, the payoff from the future sale of the good to the consumer is heavily discounted, which discourages the retailer from buying the good. The threshold ν_2 is increasing in c , μ , and σ but decreasing in ρ . Intuitively, when the cost c for preparing the good for sale is higher, it is optimal for the retailer to wait for a higher price to sell the good

to the consumer. An increase in μ or σ raises the prospect of a high retail price in the future, which encourages the retailer to wait for a higher price. Finally, a higher value of ρ means that the retailer is more impatient and thus less willing to postpone the payoff from a sale. The threshold ν_2 clearly does not depend on p .¹⁶

4. (Maximization of retailer's expected payoff) An SPE maximizing the retailer's expected payoff exists as well, and on the path of play of any SPE that maximizes the retailer's expected payoff, there exists a threshold $\nu \in (0, \infty)$ such that the retailer would not choose a non- z action until the first time that q_t is no less than ν . At the first time on the equilibrium path that q_t is greater than or equal to ν , the retailer and distributor would respectively choose V and S . Such a strategy profile involves strategies for the retailer and distributor that belong to $\bar{\Pi}_R^C$ and $\bar{\Pi}_D^C$. It can be argued as in items 1 and 2 of this remark that uniform inertia is violated but pathwise inertia is satisfied and that both uniform and pathwise F1 as well as F2 are satisfied but F3 is violated.

C.5.1 Formal Definitions of Histories and Strategy Spaces

Choose any time $t \in [0, \infty)$ and retail price process $\{q_\tau\}_{\tau \in [0, t]}$ up to that time. A history up to time t is represented by $(\{q_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in \{R, D\}}\}_{\tau \in [0, t]})$, where $\{a_\tau^i\}_{\tau \in [0, t]}$ denotes the action path of agent $i \in \{R, D\}$ up to time t with the action spaces being $\{B, C, V, z\}$ for R and $\{S, z\}$ for D .

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 $(\{q_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in \{R, D\}}\}_{\tau \in [0, t]})$ where there is no $\tau < t$ such that $(a_\tau^R, a_\tau^D) = (B, S)$ or $(a_\tau^R, a_\tau^D) = (V, S)$.
2. Let H^B be the set consisting of every history up to any time t that has the form $(\{q_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in \{R, D\}}\}_{\tau \in [0, t]})$ where there exists $\tau' < t$ such that $a_{\tau'}^R = z$ or $a_{\tau'}^D = z$ for all $\tau < \tau'$, $a_\tau^R = z$ and $a_\tau^D = z$ for all $\tau > \tau'$, and $(a_{\tau'}^R, a_{\tau'}^D) = (B, S)$.
3. Let $H^{C, V}$ be the set consisting of every history up to any time t that has the form $(\{q_\tau\}_{\tau \in [0, t]}, \{(a_\tau^i)_{i \in \{R, D\}}\}_{\tau \in [0, t]})$ where either of the following holds:

¹⁶Proofs of these comparative statics results are provided in section C.5.2.

- (a) There exists $\tau' < t$ such that $a_\tau^R = z$ or $a_\tau^D = z$ for all $\tau < \tau'$, $a_\tau^R = z$ and $a_\tau^D = z$ for all $\tau > \tau'$, and $(a_{\tau'}^R, a_{\tau'}^D) = (V, S)$.
- (b) There exists (τ', τ'') with $\tau' < \tau'' < t$ such that $a_\tau^R = z$ or $a_\tau^D = z$ for all $\tau < \tau'$, $a_\tau^R = z$ and $a_\tau^D = z$ for all $\tau > \tau'$ with $\tau \neq \tau''$, and $(a_{\tau'}^R, a_{\tau'}^D) = (B, S)$ and $(a_{\tau''}^R, a_{\tau''}^D) = (C, z)$.

The feasibility constraints are as follows. For R ,

$$\bar{A}_R(h_t) = \begin{cases} \{B, V, z\} & \text{if } h_t \in H^\emptyset \\ \{C, z\} & \text{if } h_t \in H^B \\ \{z\} & \text{if } h_t \in H^{C,V} \end{cases} .$$

For D ,

$$\bar{A}_D(h_t) = \begin{cases} \{S, z\} & \text{if } h_t \in H^\emptyset \\ \{z\} & \text{if } h_t \in H^B \cup H^{C,V} \end{cases} .$$

The sets of feasible strategies are:

$$\begin{aligned} \bar{\Pi}_R &= \{\pi_R : H \rightarrow \{B, C, V, z\} \mid \pi_R(h_t) \in \bar{A}_R(h_t) \text{ for all } h_t \in H\} \\ \bar{\Pi}_D &= \{\pi_D : H \rightarrow \{S, z\} \mid \pi_D(h_t) \in \bar{A}_D(h_t) \text{ for all } h_t \in H\}. \end{aligned}$$

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

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

$$v_R[(a_\tau^R, a_\tau^D), s_\tau] = \begin{cases} -pe^{-\rho\tau} & \text{if } (a_\tau^R, a_\tau^D) = (B, S) \\ q_\tau e^{-\rho\tau} & \text{if } a_\tau^R = C \\ (-p + q_\tau)e^{-\rho\tau} & \text{if } (a_\tau^R, a_\tau^D) = (V, S) \\ 0 & \text{otherwise} \end{cases} ,$$

and as follows for $i = D$:

$$v_D[(a_\tau^R, a_\tau^D), s_\tau] = \begin{cases} pe^{-\rho\tau} & \text{if } (a_\tau^R, a_\tau^D) \in \{(B, S), (V, S)\} \\ 0 & \text{otherwise} \end{cases}.$$

C.5.2 Proofs

Proof of Proposition 21. For any $q_t > 0$, consider the problem of choosing $r \geq q_t$ so as to maximize the expression $(r - c)(q_t/r)^\beta$, which is the value of an asset that pays $r - c$ at the first time that the retail price reaches r when the retail price is currently q_t .¹⁷ The solution to the maximization problem is given by the greater of q_t and $\nu_2 = c[\beta/(\beta - 1)]$. Consider any history up to an arbitrary time such that the retailer has already bought but not yet sold the good. In any SPE, the retailer at such a history chooses C when $q_t \geq \nu_2$ and chooses z when $q_t < \nu_2$, and the expected payoff to the retailer at such a history is $q_t - c$ if $q_t \geq \nu_2$ and $(\beta - 1)^{\beta-1}\beta^{-\beta}c^{1-\beta}q_t^\beta$ if $q_t < \nu_2$. Next consider any SPE along with a history up to an arbitrary time at which the retailer and distributor transact. The retailer chooses V if $q_t \geq \nu_2$ and chooses B if $q_t < \nu_2$, and the expected payoff to the retailer at such a history is $q_t - c - p$ if $q_t \geq \nu_2$ and $(\beta - 1)^{\beta-1}\beta^{-\beta}c^{1-\beta}q_t^\beta - p$ if $q_t < \nu_2$.

Letting ν_1 be the value of q_t that solves the equation $p = (\beta - 1)^{\beta-1}\beta^{-\beta}c^{1-\beta}q_t^\beta$, we have $\nu_1 = [(p/c)(\beta - 1)]^{1/\beta}c[\beta/(\beta - 1)]$. Note that $0 < \nu_1 < \nu_2 < \infty$ holds by the parameter restriction $\beta \in (1, 1 + c/p)$. Consider any history up to an arbitrary time such that the retailer has not yet obtained the good from the distributor. If $q_t < \nu_1$, then the retailer and distributor cannot transact at this history because the retailer would get a negative expected payoff whereas the retailer could secure a payoff of zero by always choosing z . Hence, the SPE that maximize the expected payoff of the distributor at the null history have the following property on the path of play. The retailer and distributor transact at the first time the retail price satisfies $q_t \geq \nu_1$, and the retailer sells the good to the consumer at the first time that the retail price satisfies $q_t \geq \nu_2$. \square

Proof of Item 3 in Remark 9. Note that β is decreasing in μ and σ but increasing in ρ . The threshold ν_2 is increasing in c and decreasing in β . Hence, ν_2 is increasing in

¹⁷McDonald and Siegel (1986) solve a similar problem, and the analysis of the model in Kamada and Rao (2018) involves an infinite sequence of such problems.

μ and σ but decreasing in ρ . The threshold ν_1 is clearly increasing in p and c . The partial derivative of ν_1 with respect to β is given by $\partial\nu_1/\partial\beta = \nu_1[\log(\nu_2) - \log(\nu_1)]/\beta$. Hence, ν_1 is increasing in β as $\nu_2 > \nu_1$. Thus, the threshold ν_1 is decreasing in μ and σ but increasing in ρ . \square

D Traceable Strategies without Consistent History

As mentioned in footnote 19 in section 3.2, the following is an example of a profile of traceable strategies such that there does not exist a history that is consistent with them at every time.

Example 15. Suppose $I = \{1, 2\}$. Let $\bar{A}_i(h_t) = \{x, z\}$ for each $i \in \{1, 2\}$ and all $h_t \in H$. Consider the following strategies for agents 1 and 2. If there is no positive integer n such that $t = 1/n$, then neither strategy specifies a transaction at time t .

The strategy ψ_1 of agent 1 is as follows. Consider any time t for which there exists a positive integer c such that $t = 1/c$. Suppose first that c is odd. If there is no $u < t$ such that agent 2 chose x at time u , then agent 1 chooses x at time t . If agent 2 chose x at some time $v < t$ such that $v = 1/b$ for some odd positive integer b and agent 2 did not choose x at any time $u < t$ such that there exists an even positive integer d satisfying $u = 1/d$, then agent 1 chooses x at time t . If neither of the two previous cases holds, then agent 1 chooses z at time t . Suppose next that c is even. If agent 2 chose x at some time $v < t$ such that $v = 1/b$ for some even positive integer b and agent 2 did not choose x at any time $u < t$ such that there exists an odd positive integer d satisfying $u = 1/d$, then agent 1 chooses x at time t . Otherwise, agent 1 chooses z at time t .

The strategy ψ_2 of agent 2 is as follows. Consider any time t for which there exists a positive integer c such that $t = 1/c$. Suppose first that c is odd. If there is no $u < t$ such that agent 1 chose x at time u , then agent 2 chooses x at time t . If agent 1 chose x at some time $v < t$ such that $v = 1/b$ for some even positive integer b and agent 1 did not choose x at any time $u < t$ such that there exists an odd positive integer d satisfying $u = 1/d$, then agent 2 chooses x at time t . If neither of the two previous cases holds, then agent 2 chooses z at time t . Suppose next that c is even. If agent 1 chose x at some time $v < t$ such that $v = 1/b$ for some odd positive integer b and agent 1 did not choose x at any time $u < t$ such that there exists an even positive integer d satisfying $u = 1/d$, then agent 2 chooses x at time t . Otherwise, agent 2

chooses z at time t . Note that by definition, strategies ψ_1 and ψ_2 are traceable.

We now prove by contradiction that no history is consistent with the strategy profile $\psi = (\psi_1, \psi_2) \in \Pi$. Suppose to the contrary that the history $h = \{s_t, (a_t^1, a_t^2)\}_{t \in [0, T]}$ is consistent with ψ_1 and ψ_2 at every time. It must be that $a_t^1 = a_t^2 = z$ for any t such that there does not exist a positive integer n satisfying $t = 1/n$. Suppose that there exists b such that for any positive integer $d > b$, $a_t^1 = a_t^2 = z$ at time $t = 1/d$. Then for any odd positive integer $c > b$, h would not be consistent with ψ_1 and ψ_2 at time $1/c$.

Therefore, the history h must have at least one of the following four properties. First, there exists an increasing sequence $\{r_k^1\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^1 = x$ at time $t = 1/r_k^1$. Second, there exists an increasing sequence $\{r_k^2\}_{k=1}^\infty$ of positive odd integers such that for all k , $a_t^1 = x$ at time $t = 1/r_k^2$. Third, there exists an increasing sequence $\{r_k^3\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^2 = x$ at time $t = 1/r_k^3$. Fourth, there exists an increasing sequence $\{r_k^4\}_{k=1}^\infty$ of positive odd integers such that for all k , $a_t^2 = x$ at time $t = 1/r_k^4$.

Consider the first case, where there exists an increasing sequence $\{r_k^1\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^1 = x$ at time $t = 1/r_k^1$. In order for the history h to be consistent with ψ_1 at each time in this situation, there must exist an increasing sequence $\{r_k^3\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^2 = x$ at time $t = 1/r_k^3$. In order for the history h to be consistent with ψ_2 at each time given the existence of such a sequence $\{r_k^3\}_{k=1}^\infty$, there must exist p such that for any even positive integer $d > p$, $a_t^1 = z$ at time $t = 1/d$. This contradicts the first sentence of this paragraph.

Consider the second case, where there exists an increasing sequence $\{r_k^2\}_{k=1}^\infty$ of positive odd integers such that for all k , $a_t^1 = x$ at time $t = 1/r_k^2$. Suppose that there exists g such that for any positive integer $d > g$, $a_t^2 = z$ at time $t = 1/d$. In order for h to be consistent with ψ_2 at each time in this situation, there must exist an increasing sequence $\{r_k^1\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^1 = x$ at time $t = 1/r_k^1$. This contradicts the result that h cannot have the first property. Therefore, assume that no such g exists. In order for the history h to be consistent with ψ_2 at each time in this situation, there must exist an increasing sequence $\{r_k^3\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^2 = x$ at time $t = 1/r_k^3$. In order for the history h to be consistent with ψ_1 at each time given the existence of such a sequence $\{r_k^3\}_{k=1}^\infty$, there must exist p such that for any odd positive integer $d > p$, $a_t^1 = z$ at

time $t = 1/d$. This contradicts the first sentence of this paragraph.

Consider the third case, where there exists an increasing sequence $\{r_k^3\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^2 = x$ at time $t = 1/r_k^3$. Since the history h cannot have the first two properties, there exists g such that for any positive integer $d > g$, $a_t^1 = z$ at time $t = 1/d$. In order for the history h to be consistent with ψ_2 at each time in this situation, it must be that for any even positive integer $c > g$, $a_t^2 = z$ at time $t = 1/c$. This contradicts the first sentence of this paragraph.

Consider the fourth case, where there exists an increasing sequence $\{r_k^4\}_{k=1}^\infty$ of positive odd integers such that for all k , $a_t^2 = x$ at time $t = 1/r_k^4$. Since the history h cannot have the third property, there does not exist an increasing sequence $\{r_k^3\}_{k=1}^\infty$ of positive even integers such that for all k , $a_t^2 = x$ at time $t = 1/r_k^3$. In order for the history h to be consistent with ψ_1 at each time given the existence of such a sequence $\{r_k^4\}_{k=1}^\infty$ and the nonexistence of such a sequence $\{r_k^3\}_{k=1}^\infty$, there must exist an increasing sequence $\{r_k^2\}_{k=1}^\infty$ of positive odd integers such that for all k , $a_t^1 = x$ at time $t = 1/r_k^2$. This contradicts the result that h cannot have the second property.

Since h cannot have any of the four aforementioned properties, h cannot be consistent with both ψ_1 and ψ_2 at every time. This contradicts our starting assumption that h is consistent with those strategies at every time, completing the proof. \square

E Additional Example of Quantitative Strategy

As discussed in section 4.2, the following is an example of a strategy in Π_i^Q that is contingent on the realization of the shock and the behavior of one's opponent.

Example 16. Suppose $I = \{1, 2\}$. Let $\bar{A}_i(h_t) = \{x, z\}$ for each $i \in \{1, 2\}$ and all $h_t \in H$. Let $\bar{t}_2 > \bar{t}_1 > 0$ and $\bar{s}^b \neq \bar{s}^a$. Suppose that with probability $\frac{1}{2}$ the value of the shock s_t is \bar{s}^a for all $t \in [0, \bar{t}_1]$ and that with probability $\frac{1}{2}$ the value of the shock s_t is \bar{s}^b for all $t \in [0, \bar{t}_1]$. The strategy that requires agent i to behave as follows is quantitative. If agent $-i$ chooses x at time \bar{t}_1 and \bar{s}^a is the realized value of the shock at time \bar{t}_1 , then agent i chooses x at time \bar{t}_2 . Otherwise, agent i chooses z at time \bar{t}_2 . Agent i chooses z at any time $t \neq \bar{t}_2$. Given that agent i plays this strategy and that agent $-i$ plays any traceable and frictional strategy, the actions of agent $-i$ at time \bar{t}_1 and of agent i at time \bar{t}_2 are random variables with a one- or two-point distribution. \square

F Alternative Formulation of Equilibrium Conditions

As discussed in footnote 34 in section 4.3, we illustrate a simple method to check whether a given strategy profile is an SPE. To this end, we extend the notations that were introduced for traceable, frictional, and calculable strategies in the body of the paper.

Choose any strategy profile $\pi = (\pi_j)_{j \in I}$ with $\pi_j \in \Pi_j$ for $j \in I$ and 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 that satisfy the following. With conditional probability one given $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$, there exists a unique profile $(\{\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)\}_{t \in [0, T]})_{j \in I}$ of action paths with $\{\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)\}_{t \in [0, T]} \in \Gamma_j(\{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ for each $j \in I$ 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 $t \in [u, T]$, and 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. Whenever the conditional expectation is well defined, let $V_i(k_u, \pi) = \mathbb{E}_{\{s_t\}_{t \in (u, T)}} [V_u^i(\{s_t, [\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)]_{j \in I}\}_{t \in [0, T]} | \{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]})]$ denote the expected payoff to agent $i \in I$ at k_u , where $V_u^i(\{s_t, [\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)]_{j \in I}\}_{t \in [0, T]})$ is as specified in the main text.

Next pick any strategy profile $\pi = (\pi_j)_{j \in I} \in \times_{j \in I} \Pi_j$ and any action paths $b = \{(b_t^j)_{j \in I}\}_{t \in [0, u]}$ up to an arbitrary time u such that for any shock realization $g = \{g_t\}_{t \in [0, u]}$ until time u , the following holds with conditional probability one given $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$. There exists a unique profile $(\{\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)\}_{t \in [0, T]})_{j \in I}$ of action paths with $\{\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)\}_{t \in [0, T]} \in \Gamma_j(\{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ for each $j \in I$ 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 $t \in [u, T]$, and these action paths also satisfy $\{\phi_t^j(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)\}_{t \in [0, T]} \in \Xi_j(u)$ for each $j \in I$. Let $\xi_b^i(\pi)$ be the stochastic process defined as follows for $i \in I$. At any time $t \in [0, u]$, the value of $\xi_b^i(\pi)$ is z . Let $g = \{g_t\}_{t \in [0, u]}$ represent the shock realization until time u , and denote the resulting history up to time u by $k_u = (g, b)$. Given the realization of the shock $\{s_\tau\}_{\tau \in (u, T)}$ after time u , the value of $\xi_b^i(\pi)$ at each time $t \in [u, T]$ is $\phi_t^i(k_u, \{s_\tau\}_{\tau \in (u, T)}, \pi)$.

Proposition 22. *A strategy profile $(\pi_i)_{i \in I}$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in I$ is an SPE if for any $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$, $V_i[k_u, (\pi_i, \pi_{-i})] \geq V_i[k_u, (\psi_i, \pi_{-i})]$ holds for every $\psi_i \in \bar{\Pi}_i$ satisfying the conditions below.*

1. For any $\tilde{g} = \{\tilde{g}_t\}_{t \in [0, u]}$, there is conditional probability one given $\{s_t\}_{t \in [0, u]} = \{\tilde{g}_t\}_{t \in [0, u]}$ of the following:

- (a) There exists a unique profile $(\{a_t^j\}_{t \in [0, T)})_{j \in I}$ of action paths with $\{a_t^j\}_{t \in [0, T)} \in \Gamma_j(\{a_t^j\}_{t \in [0, u)})$ 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 ψ_i and π_{-i} at each $t \in [u, T)$.
- (b) These action paths satisfy $\{a_t^j\}_{t \in [0, T)} \in \Xi_j(u)$ for each $j \in I$.
2. Denoting $b = \{(b_t^j)_{j \in I}\}_{t \in [0, u)}$, $\xi_b^i(\psi_i, \pi_{-i})$ and $\xi_b^{-i}(\psi_i, \pi_{-i})$ are progressively measurable.

The converse holds in the case where $\bar{A}_i(h_t) = A_i$ for every $h_t \in H$ and each $i \in I$.

Proof. Fix $i \in I$, and choose any $\psi_i \in \bar{\Pi}_i$ as well as any $\pi_{-i} \in \bar{\Pi}_{-i}^C$. Let $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u)})$ be any history up to time u , and denote $b = \{(b_t^j)_{j \in I}\}_{t \in [0, u)}$.

We begin by noting that the strategy ψ_i is not calculable if ψ_i does not satisfy the conditions in the statement of the proposition. Suppose first that there exists a realization $\bar{g} = \{\bar{g}_t\}_{t \in [0, u]}$ of shock levels up to time u for which there is conditional probability not equal to one given $\{s_t\}_{t \in [0, u]} = \{\bar{g}_t\}_{t \in [0, u]}$ of there existing a unique profile $(\{a_t^j\}_{t \in [0, T)})_{j \in I}$ of action paths with $\{a_t^j\}_{t \in [0, T)} \in \Gamma_j(\{a_t^j\}_{t \in [0, u)})$ 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 ψ_i and π_{-i} at each $t \in [u, T)$ and of these action paths satisfying $\{a_t^j\}_{t \in [0, T)} \in \Xi_j(u)$ for each $j \in I$. Then the strategy ψ_i cannot be calculable because π_{-i} is calculable and it follows from the main text that for every profile of calculable strategies and any history up to a given time, there is conditional probability one of there existing a unique continuation path, which has finitely many moves in any finite interval of time. Suppose next that no such \bar{g} exists but that $\xi_b^i(\psi_i, \pi_{-i})$ or $\xi_b^{-i}(\psi_i, \pi_{-i})$ is not progressively measurable. Then the strategy ψ_i cannot be calculable because π_{-i} is calculable and the analysis in the main text implies that $\xi_b^i(\psi_i, \pi_{-i})$ and $\xi_b^{-i}(\psi_i, \pi_{-i})$ must be progressively measurable if ψ_i is calculable.

In the case where $\bar{A}_i(h_t) = A_i$ for every $h_t \in H$ and each $i \in I$, we now observe that if the strategy ψ_i satisfies the conditions in the statement of the proposition, then there exists a calculable and feasible strategy ψ'_i such that (ψ'_i, π_{-i}) induces the same continuation path as (ψ_i, π_{-i}) at k_u . Assume that the aforesaid \bar{g} does not exist and that $\xi_b^i(\psi_i, \pi_{-i})$ and $\xi_b^{-i}(\psi_i, \pi_{-i})$ are progressively measurable. Let ψ'_i with $\psi'_i(h_t) = z$ for $t < u$ be defined such that $\psi'_i[(\{s_\tau\}_{\tau \in [0, t]}, \{(d_\tau^j)_{j \in I}\}_{\tau \in [0, t)})] = \phi_t^i[(\{s_\tau\}_{\tau \in [0, u]}, \{(b_\tau^j)_{j \in I}\}_{\tau \in [0, u)}), \{s_\tau\}_{\tau \in [u, T)}, (\psi_i, \pi_{-i})]$ 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 $\psi'_i \in \bar{\Pi}_i$ given the assumption that $\bar{A}_i(h_t) = A_i$ for all $h_t \in H$ and $i \in I$. It follows from the definition of ψ'_i that $\psi'_i \in \Pi_i^{TF}$, that $\psi'_i \in \Pi^Q$ and hence $\psi'_i \in \Pi^C$, that the stochastic process $\xi_b^i(\psi'_i, \pi'_{-i})$ is the same as $\xi_b^i(\psi_i, \pi_{-i})$ for any $\pi'_{-i} \in \Pi_{-i}^C$, and that $\xi_b^{-i}(\psi'_i, \pi_{-i})$ is the same as $\xi_b^{-i}(\psi_i, \pi_{-i})$. \square

The next result identifies a sufficient condition for a strategy profile to be an SPE. It follows immediately from the foregoing analysis because any strategy $\psi_i \in \bar{\Pi}_i$ that has the properties stated in the above proposition also has the properties stated in the corollary below.

Corollary 23. *A strategy profile $(\pi_i)_{i \in I}$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in I$ is an SPE if for any $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$, $V_i[k_u, (\pi_i, \pi_{-i})] \geq V_i[k_u, (\psi_i, \pi_{-i})]$ holds for every $\psi_i \in \bar{\Pi}_i$ satisfying the conditions below.*

1. *There is conditional probability one given $\{s_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$ of the following:*

- (a) *There exists a unique profile $(\{a_t^j\}_{t \in [0, T]})_{j \in I}$ of action paths with $\{a_t^j\}_{t \in [0, T]} \in \Gamma_j(\{a_t^j\}_{t \in [0, u]})$ 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 ψ_i and π_{-i} at each $t \in [u, T]$.*
- (b) *These action paths satisfy $\{a_t^j\}_{t \in [0, T]} \in \Xi_j(u)$ for each $j \in I$.*

2. *$V_i[k_u, (\psi_i, \pi_{-i})]$ and $V_{-i}[k_u, (\psi_i, \pi_{-i})]$ are well defined.*

G From Measurable Attachability to Calculability Restriction

As stated in section 6.5, we define a weakened concept of equilibrium under the restriction to measurably attachable strategy profiles and prove that any synchronous strategy profile satisfying this notion of equilibrium is an SPE under the calculability restriction. We say that $\pi \in \Pi^A \cap \bar{\Pi}$ is a **pseudo-SPE** of $\Gamma(\Pi^A \cap \bar{\Pi})$ if $V_i(k_u, \pi) \geq V_i[k_u, (\pi'_i, \pi_{-i})]$ for each $i \in I$, any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^j)_{j \in I}\}_{t \in [0, u]})$ up to an arbitrary time u , and every $\pi'_i \in \bar{\Pi}_i^{TF}$ for which there exists $\pi'' \in \Pi^A$ such that $\xi_b^e(\pi'')$ with $b = \{(b_t^j)_{j \in I}\}_{t \in [0, u]}$ is the same stochastic process as $\xi_b^e[(\pi'_i, \pi_{-i})]$ for all $e \in I$.

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

Proposition 24. *If the synchronous strategy profile $\pi \in \Pi^A$ is a pseudo-SPE of $\Gamma(\Pi^A \cap \bar{\Pi})$, then π is an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$.*

Proof of Proposition 24. Let the synchronous strategy profile $\pi \in \Pi^A$ be a pseudo-SPE of $\Gamma(\Pi^A \cap \bar{\Pi})$. Assume that $\chi_i(h_t, \tilde{\pi}) \leq \zeta_i(h_t)$ for all $\tilde{\pi} \in \times_{j \in I} \bar{\Pi}_j^{TF}$, $h_t \in H$, and $i \in I$. It suffices to show that π is an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^{TF}, (\chi_i)_{i \in I})$ because it will then follow from theorem 3 that π is an SPE of $\Gamma(\times_{i \in I} \bar{\Pi}_i^C)$.

To show this, 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 $U_i(k_u, \pi) = V_i(k_u, \pi) \geq \zeta_i(k_u)$. Hence, $U_i(k_u, \pi) \geq U_i[k_u, (\pi'_i, \pi_{-i})]$ holds in this case.

Suppose now that $\xi_b^j(\pi'_i, \pi_{-i})$ is progressively measurable for all $j \in I$. First, we construct $\pi'' \in \Pi^A$ such that $\xi_b^e(\pi'')$ with $b = \{(b_t^j)_{j \in I}\}_{t \in [0, u]}$ is the same stochastic process as $\xi_b^e[(\pi'_i, \pi_{-i})]$ for all $e \in I$. For each $e \in I$, let π''_e with $\pi''_e(h_t) = z$ for $t < u$ be defined such that

$$\pi''_e[(\{s_\tau\}_{\tau \in [0, t]}, \{(d_\tau^j)_{j \in I}\}_{\tau \in [0, t]})] = \phi_t^e[(\{s_\tau\}_{\tau \in [0, u]}, \{(b_\tau^j)_{j \in I}\}_{\tau \in [0, u]}), \{s_\tau\}_{\tau \in (u, T)}, (\pi'_i, \pi_{-i})]$$

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 $t \geq u$. By the definition of π'' , $\pi''_e \in \Pi_e^{TF}$ for each $e \in I$, the stochastic process $\xi_b^j(\pi'')$ is the same as $\xi_b^j(\pi'_i, \pi_{-i})$ for all $j \in I$, and π'' is a measurably attachable strategy profile.

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

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

To prove this result, we first show that if the synchronous strategy profile π is an SPE under the restriction to measurably attachable strategy profiles, then π is an SPE when nonmeasurable behavior is assigned an expected payoff no greater than the infimal feasible payoff. It then follows from theorem 3 that π is an SPE under the calculability restriction.

References for the Online Appendix

- KAMADA, Y., AND N. RAO (2018): “Sequential Exchange with Stochastic Transaction Costs,” Mimeo, University of California, Berkeley.
- MCDONALD, R., AND D. SIEGEL (1986): “The Value of Waiting to Invest,” *The Quarterly Journal of Economics*, 101(4), 707–727.
- 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.
- SIMON, L. K., AND M. B. STINCHCOMBE (1989): “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, 57(5), 1171–1214.