Online Appendix to "Sequential Exchange with Stochastic Transaction Costs"

Yuichiro Kamada Neel Rao

August 15, 2020

Contents

I Violation of Inertia by Maximal Equilibrium **2**

II Alternative Specification of Payoffs **5**

III Intuitive Examples for Optimal Solution **6**

IV Efficiency Results **10**

IV.1 Properties of Efficient Strategies 10

IV.2 General Condition for Asymptotic Efficiency 14

V Variations of Basic Model **16**

V.1 Model with Stochastic Supplies of Goods 17

V.2 Model with Cost Proportional to Amount Transferred 19

I Violation of Inertia by Maximal Equilibrium

This section shows that the inertia condition excludes the optimal strategies. Below is a formal definition of the concept of inertia as it relates to the strategies of agents in our model. A strategy is said to satisfy inertia if there is an interval of time after each transaction during which the agents cannot make further transfers.

Definition 1. For $i \in \{1, 2\}$, the strategy $\pi_i \in \Pi_i$ is **inertial** if for every history $h_t = \big(\{\tilde{c}_\tau\}_{\tau \in [0,t]}, \{(\tilde{f}^1_\tau, \tilde{f}^2_\tau)\}_{\tau \in [0,t)}\big)$ up to an arbitrary time $t \in [0, \infty)$ such that $\pi_i(h_t) > 0$, there exists $\epsilon > 0$ such that $\pi_i(k_u) = 0$ for any history $k_u = \big(\{\tilde{g}_\tau\}_{\tau \in [0,u]}, \{(\tilde{b}^1_\tau, \tilde{b}^2_\tau)\}_{\tau \in [0,u)}\big)$ up to an arbitrary time $u \in (t, t+\epsilon)$ satisfying $\{\tilde{g}_\tau\}_{\tau \in [0,t]} = \{\tilde{c}_\tau\}_{\tau \in [0,t]}$ and $\{(\tilde{b}^1_\tau, \tilde{b}^2_\tau)\}_{\tau \in [0,t)}$ $= \{(\tilde{f}^1_\tau, \tilde{f}^2_\tau)\}_{\tau \in [0,t)}$.

The following result shows that the optimal strategies violate the inertia condition. If the cost decreases sufficiently after a transaction, then it is optimal for the agents to make another transfer, but inertia may prevent such an exchange.

Proposition A. *Assume that the cost process $\{c_t\}_{t \in [0,\infty)}$ follows a geometric Brownian motion with arbitrary drift μ and positive volatility σ. Then any maximal symmetric SPE in grim-trigger strategies is not inertial.*

Proof. Let $\pi^* = (\pi^*_1, \pi^*_2)$ be any maximal symmetric SPE in grim-trigger strategies. From the main text, the strategy profile π^* is characterized by a sequence $\{c^*_k, f^*_k\}^\infty_{k=1}$ such that, with probability one, the k^{th} transaction is made when the cost reaches c^*_k for the first time, and the amount f^*_k is transferred by each agent at this transaction. We prove by contradiction that π^* is not inertial.

Suppose to the contrary that π^* is inertial. Choose any positive integer n. Let $h = \{g_t, (b^1_t, b^2_t)\}_{t \in [0,\infty)}$ be any history for which there exists a time $u \in [0, \infty)$ such that $g_u = c^*_n$, $g_t > c^*_n$ for all $t \in [0, u)$, and $\pi^*_i(h_t) = b^i_t$ for all $t \in [0, u)$ and each $i \in \{1, 2\}$. It follows that $\pi^*_i(h_u) > 0$ for $i \in \{1, 2\}$. Hence, the inertia property implies that there exists $\epsilon > 0$ such that $\phi^i_t(h_u, \{\tilde{g}_\tau\}_{\tau \in (u,\infty)}, \pi^*) = 0$ for each $i \in \{1, 2\}$, all $t \in (u, u+\epsilon)$, and any $\{\tilde{g}_\tau\}_{\tau \in (u,\infty)}$. Given that the realization of the cost process up to time u is such that $\{c_t\}_{t \in [0,u]} = \{g_t\}_{t \in [0,u]}$, there is a positive conditional probability of the cost process $\{c_t\}_{t \in (u,\infty)}$ after time u being such that there exists a time $v \in (u, u+\epsilon)$ satisfying $c_v = c^*_{n+1}$, $c_t > c^*_{n+1}$ for all $t \in [0, v)$, and $\phi^i_v(h_0, \{c_t\}_{t \in (0,\infty)}, \pi^*) = 0$ for $i \in \{1, 2\}$.

Note that there is positive probability of the realization of the cost process $\{c_t\}_{t\in[0,\infty)}$ being such that there exists a time $\tilde{u} \in [0,\infty)$ satisfying $c_{\tilde{u}} = c_n^*$, $c_t > c_n^*$ for all $t \in [0,\tilde{u})$, and $\phi_{\tilde{u}}^i(h_0, \{c_t\}_{t\in(0,\infty)}, \pi^*) > 0$ for $i \in \{1,2\}$. It follows from the argument in the preceding paragraph that there is positive probability of the realization of the cost process $\{c_t\}_{t\in[0,\infty)}$ being such that there exists $\tilde{v}$ satisfying $c_{\tilde{v}} = c_{n+1}^*$, $c_t > c_{n+1}^*$ for all $t \in [0,\tilde{v})$, and $\phi_{\tilde{v}}^i(h_0, \{c_t\}_{t\in(0,\infty)}, \pi^*) = 0$ for $i \in \{1,2\}$. Hence, the path of play induced by π^* cannot be such that, with probability one, the $(n+1)^{\text{th}}$ transaction is made when the cost reaches c_{n+1}^* for the first time. This contradicts the definition of π^*. $\square$

In addition, Bergin and MacLeod (1993) define a less restrictive condition based on the completion of the set of inertial strategies. However, their methodology cannot be easily adapted to our setting. Below we show that the maximal equilibrium of our model cannot be expressed as the limit of a Cauchy sequence of inertial strategies. In doing so, we employ a class of metrics on the strategy space that includes the one used by Bergin and Macleod (1993). This class of metrics is reasonably large, suggesting that a simple modification of their technique does not apply here, which makes it necessary to introduce a new set of restrictions as we did in the main text.

Let $\tilde{\Pi}$ be the set consisting of any strategy profile $\pi = (\pi_1,\pi_2) \in \Pi$ such that the following holds. Choose any path of transfers $b = \{(b_t^1, b_t^2)\}_{t\in[0,u)}$ up to an arbitrary time u. For any cost realization $g = \{g_t\}_{t\in[0,u]}$ up to time u, there is conditional probability one given $\{c_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$ that there exists a unique path of transfers $\{\phi_t^i(k_u, \{c_\tau\}_{\tau\in(u,\infty)}, \pi)\}_{t\in[0,\infty)}$ for each agent $i \in \{1,2\}$ such that $\phi_t^i(k_u, \{c_\tau\}_{\tau\in(u,\infty)}, \pi) = g_t$ for all $t \in [0,u)$ and each $i \in \{1,2\}$ and such that $\pi_i\big(\{c_v\}_{v\in[0,t]}, \{[\phi_v^1(k_u, \{c_\tau\}_{\tau\in(u,\infty)}, \pi), \phi_v^2(k_u, \{c_\tau\}_{\tau\in(u,\infty)}, \pi)]\}_{v\in[0,t)}\big) = \phi_t^i(k_u, \{c_\tau\}_{\tau\in(u,\infty)}, \pi)$ for all $t \in [u,\infty)$ and each $i \in \{1,2\}$, where $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b_t^1, b_t^2)\}_{t\in[0,u)}\big)$ is the history up to time u. Furthermore, for any $g = \{g_t\}_{t\in[0,u]}$, the set $\{t \in [u,v] : \phi_t^i(k_u, \{c_\tau\}_{\tau\in(u,\infty)}, \pi) > 0\}$ is finite for all $v \geq u$ and each $i \in \{1,2\}$ with conditional probability one given $\{c_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$. Finally, the stochastic process $\xi_b^i(\pi_1,\pi_2)$ defined as follows is progressively measurable for $i \in \{1,2\}$.[1] At any time $t \in [0,u)$, the value of $\xi_b^i(\pi_1,\pi_2)$ is 0. Let $g = \{g_t\}_{t\in[0,u]}$ represent the cost realization until time u, and denote the resulting history up to time u by $k_u = (b,g)$. Given the realization of the cost $\{c_\tau\}_{\tau\in(u,\infty)}$ after time u, the value of $\xi_b^i(\pi_1,\pi_2)$ at each time $t \in [u,\infty)$ is

[1]Progressive measurability ensures that the conditional probability in Theorem A is well defined.

$\phi_t^i(k_u, \{c_\tau\}_{\tau\in(u,\infty)}, \pi)$.

The result below provides a general condition under which the maximal equilibrium cannot be approximated by a sequence of inertial strategies. The set $\tilde{\Pi}$ is endowed with the metric d. Let $\Sigma_i(k_u; \{g_t\}_{t\in[u,v]}; u, v; \pi)$ be the sum of the transfers that agent $i \in \{1,2\}$ would make during the time interval $[u, v]$ with $u \leq v$ if $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b_t^1, b_t^2)\}_{t\in[0,u)}\big)$ is the history up to time u, the cost path $\{g_t\}_{t\in[u,v]}$ is realized between times u and v, and strategy profile $\pi \in \tilde{\Pi}$ is played by the agents.

Theorem A. *If d has the following property, then there is no sequence $\{\psi_n\}_{n=1}^{\infty}$ with $\lim_{n\to\infty} d(\psi_n, \pi^*) = 0$ such that $\psi_n \in \tilde{\Pi}$ is a profile of inertial strategies for all n, where π^* is any maximal symmetric SPE in grim-trigger strategies. For any $s > 0$, $l > 0$, and $p \in (0, 1]$, there exists $z > 0$ satisfying $d(\pi^a, \pi^b) > z$ for any $\pi^a, \pi^b \in \tilde{\Pi}$ such that one can find $k_u = \big(\{g_t\}_{t\in[0,u]}, \{(b_t^1, b_t^2)\}_{t\in[0,u)}\big)$ and $i \in \{1,2\}$ for which $\big|\Sigma_i(k_u; \{c_t\}_{t\in[u,u+l]}; u, u+l; \pi^b) - \Sigma_i(k_u; \{c_t\}_{t\in[u,u+l]}; u, u+l; \pi^a)\big| > s$ with conditional probability no less than p given that $\{c_t\}_{t\in[0,u]} = \{g_t\}_{t\in[0,u]}$.*

Proof. Let d be a metric on $\tilde{\Pi}$ having the property in the statement of the theorem. Let $\pi^* = (\pi_1^*, \pi_2^*)$ be any maximal symmetric SPE in grim-trigger strategies. Let $\{\psi_n\}_{n=1}^{\infty}$ be any sequence for which $\psi_n = (\psi_{n1}, \psi_{n2}) \in \tilde{\Pi}$ is a profile of inertial strategies for all n. Choose any integer $x > 1$. For each index $n \geq 1$, there exists a history $k_u^n = \big(\{g_t^n\}_{t\in[0,u]}, \{(b_t^n, b_t^n)\}_{t\in[0,u)}\big)$ up to some time u such that $\pi_i^*(k_u^n) = f_x^* > 0$ and $\psi_{ni}(k_u^n) = 0$ for each $i \in \{1,2\}$. Choose any integer $j \in \{1,2\}$. There are two cases to consider.

Suppose first that there exists $s > 0$, $l > 0$, and $p \in (0,1]$ such that for any index m, one can find an index $n_m > m$ for which there exists a history $\tilde{k}_{\tilde{u}_m}^{n_m} = \big(\{\tilde{g}_t^{n_m}\}_{t\in[0,\tilde{u}_m]}, \{(\tilde{b}_t^{n_m}, \tilde{b}_t^{n_m})\}_{t\in[0,\tilde{u}_m)}\big)$ up to some time $\tilde{u}_m > u$ satisfying $\{\tilde{g}_t^{n_m}\}_{t\in[0,u]} = \{g_t^{n_m}\}_{t\in[0,u]}$, $\{\tilde{b}_t^{n_m}\}_{t\in[0,u)} = \{b_t^{n_m}\}_{t\in[0,u)}$, $\tilde{b}_u^{n_m} = 0$, and the following condition. Given that $\{c_t\}_{t\in[0,\tilde{u}_m]} = \{\tilde{g}_t^{n_m}\}_{t\in[0,\tilde{u}_m]}$, there is conditional probability no less than p that $\Sigma_j(\tilde{k}_{\tilde{u}_m}^{n_m}; \{c_t\}_{t\in[\tilde{u}_m,\tilde{u}_m+l]}; \tilde{u}_m, \tilde{u}_m + l; \psi_{n_m}) > s$. It follows from $\pi_j^*(k_u^{n_m}) \neq \tilde{b}_u^{n_m}$ that $\Sigma_j(\tilde{k}_{\tilde{u}_m}^{n_m}; \{c_t\}_{t\in[\tilde{u}_m,\tilde{u}_m+l]}; \tilde{u}_m, \tilde{u}_m + l; \pi^*) = 0$ with conditional probability one given that $\{c_t\}_{t\in[0,\tilde{u}_m]} = \{\tilde{g}_t^{n_m}\}_{t\in[0,\tilde{u}_m]}$. Therefore, $\big|\Sigma_j(\tilde{k}_{\tilde{u}_m}^{n_m}; \{c_t\}_{t\in[\tilde{u}_m,\tilde{u}_m+l]}; \tilde{u}_m, \tilde{u}_m+l; \psi_{n_m}) - \Sigma_j(\tilde{k}_{\tilde{u}_m}^{n_m}; \{c_t\}_{t\in[\tilde{u}_m,\tilde{u}_m+l]}; \tilde{u}_m, \tilde{u}_m+l; \pi^*)\big| > s$ with conditional probability no less than p given that $\{c_t\}_{t\in[0,\tilde{u}_m]} = \{\tilde{g}_t^{n_m}\}_{t\in[0,\tilde{u}_m]}$. This implies that there exists $z > 0$ such that for any index m, one can find an index $n_m > m$ satisfying $d(\psi_{n_m}, \pi^*) > z$. Hence, $\lim_{n\to\infty} d(\psi_n, \pi^*) = 0$ cannot hold.

Suppose next that no such $s > 0$, $l > 0$, and $p \in (0, 1]$ exist. Then there exists an index m such that for any index $n > m$ along with any history $\tilde{k}^n_{\tilde{u}} = \big(\{\tilde{g}^n_t\}_{t\in[0,\tilde{u}]}, \{(\tilde{b}^n_t, \tilde{b}^n_t)\}_{t\in[0,\tilde{u})}\big)$ up to any time $\tilde{u} > u$ satisfying $\{\tilde{g}^n_t\}_{t\in[0,u]} = \{g^n_t\}_{t\in[0,u]}$, $\{\tilde{b}^n_t\}_{t\in[0,u)} = \{b^n_t\}_{t\in[0,u)}$, and $\tilde{b}^n_u = 0$, the conditional probability that $\Sigma_j(\tilde{k}^n_{\tilde{u}}; \{c_t\}_{t\in[\tilde{u},\tilde{u}+1]}; \tilde{u}, \tilde{u}+1; \psi_n) \geq f^*_x/2$ is less than $1/2$ given that $\{c_t\}_{t\in[0,\tilde{u}]} = \{\tilde{g}^n_t\}_{t\in[0,\tilde{u}]}$. Thus, if $n > m$, then $\big|\Sigma_j(k^n_u; \{c_t\}_{t\in[u,u+1]}; u, u+1; \psi_n) - \Sigma_j(k^n_u; \{c_t\}_{t\in[u,u+1]}; u, u+1; \pi^*)\big| > f^*_x/2$ with conditional probability greater than $1/2$ given that $\{c_t\}_{t\in[0,u]} = \{g^n_t\}_{t\in[0,u]}$. It follows that there exists $z > 0$ such that $d(\psi_n, \pi^*) > z$ for any index $n > m$. Hence, $\lim_{n\to\infty} d(\psi_n, \pi^*) = 0$ cannot hold. □

II Alternative Specification of Payoffs

The analysis in the main text was based on a payoff function in which each transaction resulted in a discrete cost and benefit. As mentioned in the body of the paper, there is an equivalent formulation in which each transfer induces a stream of flow benefits while the transaction cost is discrete. This version of the model may better fit some applications.

The proposition below formalizes this notion. Let $\hat{s}^{i,j}_t = \lim_{\tau\to t^-} s^{i,j}_\tau$ represent the amount of good j that agent i possesses immediately before time t.[2] The transfer made by agent i at time t satisfies $f^i_t = \hat{s}^{i,i}_t - s^{i,i}_t$.

Proposition B. *For $h = \{c_\tau, (f^1_\tau, f^2_\tau)\}_{\tau\in[0,\infty)}$, let*

$$W^i_t(h) = \int_t^\infty e^{-\rho\cdot(\tau-t)} s^{i,-i}_\tau d\tau \;\; - \sum_{\{\tau\in[t,\infty):f^i_\tau>0\}} e^{-\rho\cdot(\tau-t)} C_\tau,$$

where $C_t = c_t/\rho$. Then

$$V^i_t(h) = \rho W^i_t(h) - \hat{s}^{i,-i}_t.$$

Proof. For any $h = \{c_\tau, (f^1_\tau, f^2_\tau)\}_{\tau\in[0,\infty)}$, the following holds:

$$\begin{aligned} W^i_t(h) &= \int_t^\infty e^{-\rho\cdot(\tau-t)} s^{i,-i}_\tau d\tau \;\; - \sum_{\{\tau\in[t,\infty):f^i_\tau>0\}} e^{-\rho\cdot(\tau-t)} C_\tau \\ &= \int_t^\infty e^{-\rho\cdot(\tau-t)} \hat{s}^{i,-i}_\tau d\tau + \sum_{\{\tau\in[t,\infty):f^{-i}_\tau>0\}} \int_\tau^\infty e^{-\rho\cdot(\upsilon-t)} (s^{i,-i}_\upsilon - \hat{s}^{i,-i}_\upsilon) d\upsilon \end{aligned}$$

[2] We define $\hat{s}^{i,j}_0 = s^{i,j}_0 = q$ if $i = j$ and $\hat{s}^{i,j}_0 = s^{i,j}_0 = 0$ if $i \neq j$. Note that $\lim_{\tau\to t^-} s^{i,j}_\tau$ is well defined because $s^{i,j}_\tau$ is monotonic over time.

$$
\begin{aligned}
&\quad - \sum_{\{\tau\in[t,\infty):f^i_\tau>0\}} e^{-\rho\cdot(\tau-t)}C_\tau \\
&= \hat{s}^{i,-i}_t \cdot -\frac{1}{\rho}e^{-\rho\cdot(\tau-t)}\Big|_t^\infty + \sum_{\{\tau\in[t,\infty):f^{-i}_\tau>0\}} (s^{i,-i}_\tau - \hat{s}^{i,-i}_\tau)\cdot -\frac{1}{\rho}e^{-\rho\cdot(\upsilon-t)}\Big|_\tau^\infty \\
&\quad - \sum_{\{\tau\in[t,\infty):f^i_\tau>0\}} e^{-\rho\cdot(\tau-t)}C_\tau \\
&= \frac{1}{\rho}\hat{s}^{i,-i}_t + \frac{1}{\rho}\sum_{\{\tau\in[t,\infty):f^{-i}_\tau>0\}} (s^{i,-i}_\tau - \hat{s}^{i,-i}_\tau)e^{-\rho\cdot(\tau-t)} \\
&\quad - \frac{1}{\rho}\sum_{\{\tau\in[t,\infty):f^i_\tau>0\}} e^{-\rho\cdot(\tau-t)}\rho C_\tau \\
&= \frac{1}{\rho}\left(\hat{s}^{i,-i}_t + \sum_{\{\tau\in[t,\infty):f^{-i}_\tau>0\}} e^{-\rho\cdot(\tau-t)}f^{-i}_\tau - \sum_{\{\tau\in[t,\infty):f^i_\tau>0\}} e^{-\rho\cdot(\tau-t)}c_\tau\right) \\
&= \frac{1}{\rho}\left(\hat{s}^{i,-i}_t + V^i_t(h)\right).
\end{aligned}
$$

Rearranging, we obtain the desired result. □

The preceding result suggests the following alternative formulation of the game. Agent i pays the fixed cost C_t if she makes a transfer at time t, where we define $C_t = c_t/\rho$. Moreover, if agent i transfers the amount f^i_t of her good to agent $-i$ at time t, then the transfer gives agent $-i$ a flow benefit of f^i_t at each instant from time t onwards. Note that the present discounted value of this flow benefit is f^i_t/ρ at the time of the transaction. In the modified model with flow benefits and discrete costs, the realized payoff to agent i at time t is $W^i_t(h)$ when the history is $h = \{c_\tau, (f^1_\tau, f^2_\tau)\}_{\tau\in[0,\infty)}$. The proposition above implies that any SPE π of the original model is an SPE of the modified model, and vice versa.

III Intuitive Examples for Optimal Solution

This section contains examples demonstrating the basic properties of the maximal symmetric equilibrium of the model. Assume that $\{c_t\}_{t\in[0,\infty)}$ is a continuous Markov cost process and that each random variable c_t for $t \geq 0$ takes values in the state space $S \subseteq \mathbb{R}_{++}$. The following example illustrates why a non-stationary symmetric SPE is weakly Pareto dominated by a stationary symmetric SPE.

Example A. Let $\{\tilde{c}_k\}_{k=1}^\infty$ be a positive decreasing sequence with $\tilde{c}_1 < c_0$. For $j \in$

$\{a, b\}$, let $\{\tilde{f}_k^j\}_{k=1}^{\infty}$ be a positive sequence with $\sum_{k=1}^{\infty} \tilde{f}_k^j \leq q$. Assume that $\tilde{f}_k^a \neq \tilde{f}_k^b$ for some index k. Choose any $\tilde{t} > 0$. Let π be a symmetric SPE in grim-trigger strategies. Suppose that the path of play induced by π is as follows. If the first time that the cost reaches $\tilde{c}_1$ is greater than $\tilde{t}$, then the k^{th} transaction is made when the cost reaches $\tilde{c}_k$ for the first time, and the amount $\tilde{f}_k^a$ is transferred by each agent at this transaction. Otherwise, the k^{th} transaction is made when the cost reaches $\tilde{c}_k$ for the first time, and the amount $\tilde{f}_k^b$ is transferred by each agent at this transaction. If an agent deviates from the specified path of play, then neither agent makes any transactions following the deviation.

Note that π is a non-stationary strategy profile. For $j \in \{a, b\}$, define the stationary symmetric SPE π^j in grim-trigger strategies as follows. The k^{th} transaction is made when the cost reaches $\tilde{c}_k$ for the first time, and the amount $\tilde{f}_k^j$ is transferred by each agent at this transaction. If an agent deviates from the specified path of play, then neither agent makes any transactions following the deviation. Since the cost follows a Markov process, the conditional distribution of future values of the cost given that the current value of the cost is $\tilde{c}_1$ does not vary based on whether the current time is greater than $\tilde{t}$. Given that π is played, let $\tilde{V}^a$ and $\tilde{V}^b$ respectively denote the expected payoffs to each agent upon reaching the cost $\tilde{c}_1$ for the first time when this time is strictly greater than and weakly less than $\tilde{t}$.[3] If $\tilde{V}^a$ is no less than $\tilde{V}^b$, then the expected payoff to each agent is at least as high under π^a as under π. If $\tilde{V}^b$ is no less than $\tilde{V}^a$, then the expected payoff to each agent is at least as high under π^b as under π. Thus, there exists at least one stationary symmetric SPE that weakly Pareto dominates the non-stationary symmetric SPE π. $\square$

Next is an example that helps to explain why a stationary symmetric SPE with some incentive constraint slack is not Pareto optimal.

Example B. The agents are playing a symmetric SPE π in grim-trigger strategies. Letting $0 < \tilde{c}_2 < \tilde{c}_1 < c_0$, assume that the path of play induced by π is such that the agents transfer the positive amount $\tilde{f}_1$ at the first time that the cost reaches $\tilde{c}_1$ and transfer the positive amount $\tilde{f}_2$ at the first time that the cost reaches $\tilde{c}_2$. Suppose that the incentive constraint at the first transaction is slack, meaning that the cost

[3]These values are calculated before any transaction that happens at the aforementioned time. For $j \in \{a, b\}$, the expected payoff $\tilde{V}^j$ is equal to the sum of the difference between the amount transferred $\tilde{f}_1^j$ and the cost incurred $\tilde{c}_1$ on the first transaction and the continuation value after reaching the cost $\tilde{c}_1$ for the first time.

incurred at the first transaction is less than the continuation value.

Then there exists a symmetric SPE π' in grim-trigger strategies such that each agent receives a higher expected payoff when playing π' than when playing π and such that the following property holds. For some $\epsilon > 0$, the strategy profile π' induces a path of play in which the agents transfer the positive amount $\tilde{f}_1 + \epsilon$ at the first time that the cost reaches $\tilde{c}_1$ and transfer the amount $\tilde{f}_2 - \epsilon$ at the first time that the cost reaches $\tilde{c}_2$. This perturbation of the original strategy profile enables some of each good to be transferred sooner rather than later. The expected payoff can be increased without violating the incentive constraints. It follows that π cannot be a maximal symmetric SPE. $\square$

The example below provides intuition for why a stationary symmetric SPE with a nondecreasing sequence of costs incurred is strongly Pareto dominated.

Example C. Suppose that the agents are playing a symmetric SPE π in grim-trigger strategies. Let $0 < \tilde{c}_1 < \tilde{c}_2 < c_0$. The path of play induced by π is such that the first transaction occurs at the first time that the cost reaches $\tilde{c}_1$, and the agents transfer the positive amount $\tilde{f}_1$ on this transaction. The second transaction occurs at the first time after the first transaction that the cost reaches $\tilde{c}_2$, and the agents transfer the positive amount $\tilde{f}_2$ on the second transaction. Assume that the incentive constraints at the first two transactions are binding, meaning that the costs incurred at these transaction are equal to the respective continuation values after these transactions.

Then there exists a symmetric SPE π' in grim-trigger strategies such that each agent receives a higher expected payoff when playing π' than when playing π and such that the following property holds. The strategy profile π' induces a path of play in which the first transaction occurs at the first time that the cost reaches $\tilde{c}_2$, and the agents transfer the amount $\tilde{f}_1 + \tilde{f}_2$ on this transaction. In other words, the first and second transactions in the original strategy profile are combined into a single transaction in the revised strategy profile. Moreover, this combined transaction occurs sooner than the first two transactions originally occur. Noting that the incentive constraints at these two transactions are binding in the original strategy profile, the revised strategy profile can increase the expected payoff without violating the incentive constraints. It follows that π cannot be a maximal symmetric SPE. $\square$

Assume further that the cost process $\{c_t\}_{t\in[0,\infty)}$ follows a geometric Brownian motion with arbitrary drift μ and positive volatility σ. The following is an example of a maximal symmetric SPE not in grim-trigger strategies.

Example D. Let π^* denote the maximal symmetric SPE in grim-trigger strategies as specified in the main text. Suppose that π^* is played with the following exception. If agent 1 transfers 0 but agent 2 transfers f_1^* at the first time the cost reaches c_1^*, then the agents do not make any transfers until the next time the cost reaches $q - f_1^*$.[4] The strategy profile specifies that neither agent makes a transfer at the next time the cost reaches $q - f_1^*$. However, if agent 1 happens to transfer the amount f_1^* at the next time the cost reaches $q - f_1^*$, then the agents play π^* starting from the next time that the cost reaches c_2^*, behaving as if no agent previously deviated from the specified path of play. Otherwise, no further transactions occur. □

The example below helps to demonstrate why the model does not have a non-stationary maximal symmetric SPE.

Example E. Suppose that π is a non-stationary maximal symmetric SPE in grim-trigger strategies. Recall the definitions of c_k^* and f_k^* in the main text. Let $c_2^* < \tilde{c}_1 < c_1^*$. Choose any $\tilde{t} > 0$. The path of play induced by π is as follows. If the first time that the cost reaches c_1^* is greater than $\tilde{t}$, then the first transaction is made when the cost reaches $\tilde{c}_1$, and the amount f_k^* is transferred by each agent at this transaction. Otherwise, the first transaction is made when the cost reaches c_1^* for the first time, and the amount f_k^* is transferred by each agent at this transaction. For any positive integer $k > 1$, the k^{th} transaction is made when the cost reaches c_k^* for the first time, and the amount f_k^* is transferred by each agent at this transaction. If an agent deviates from the specified path of play, then neither agent makes any transactions following the deviation.

Noting that the cost follows a Markov process, the conditional distribution of future values of the cost given that the current value of the cost is c_1^* does not vary based on whether the current time is greater than $\tilde{t}$. Because π is a maximal symmetric SPE, the expected payoff to each agent upon reaching the cost c_1^* for the first time is the same regardless of whether this time is greater than $\tilde{t}$. Hence, the following is a stationary maximal symmetric SPE in grim-trigger strategies. The first transaction is made when the cost reaches $\tilde{c}_1$ for the first time, and the amount f_k^* is transferred by each agent at this transaction. For any positive integer $k > 1$, the k^{th} transaction

[4] Note that $c_1^* < q - f_1^*$, where c_1^* and f_1^* are defined in the main text, which characterizes the maximal symmetric SPE.

is made when the cost reaches c_k^* for the first time, and the amount f_k^* is transferred by each agent at this transaction. If an agent deviates from the specified path of play, then neither agent makes any transactions following the deviation. However, this contradicts the result that any stationary maximal symmetric SPE must induce the uniquely optimal path of play described in the main text. Thus, the non-stationary strategy profile π cannot be a maximal symmetric SPE. $\square$

IV Efficiency Results

This section extends the welfare analysis. Appendix IV.1 derives some properties of efficient strategy profiles. Appendix IV.2 identifies a general condition on the cost process under which the efficient outcome can be approximated as the discount rate approaches zero.

IV.1 Properties of Efficient Strategies

We begin with some comparative statics for the efficient path of play. The following results are immediate given the expression for the cost cutoff $\bar{c}$ in the main text. Therefore, their proofs are omitted. We start by describing how the parameters of the model affect the efficient cost incurred.

Corollary A. *Assume that the cost process $\{c_t\}_{t\in[0,\infty)}$ follows a geometric Brownian motion with arbitrary drift μ and positive volatility σ. The efficient cost cutoff $\bar{c}$ is increasing in μ and ρ but decreasing in σ.*

If μ decreases or σ increases, then a low realization of cost process becomes more likely. Hence, it is profitable for the agents to wait for the cost to become low before making a transaction. As ρ decreases, agents become more patient and so prefer waiting for a low cost before transacting. Accordingly, the cost cutoff $\bar{c}$ is small. These comparative statics differ from those for a maximal symmetric equilibrium. In the presence of incentive constraints, the costs paid on later transactions are increasing in μ and ρ as well as decreasing in σ, but the opposite may hold for earlier transactions depending on the parameter values.

We next examine the efficient behavior as the discount rate respectively approaches zero and infinity.

Corollary B. *Assume that the cost process* $\{c_t\}_{t\in[0,\infty)}$ *follows a geometric Brownian motion with arbitrary drift* μ *and positive volatility* σ*. If* $\mu \leq \sigma^2/2$*, then* $\lim_{\rho\to 0} \bar{c} = 0$*. If* $\mu > \sigma^2/2$*, then* $\lim_{\rho\to 0} \bar{c} > 0$*. In addition,* $\lim_{\rho\to\infty} \bar{c} = q$

If $\mu \leq \sigma^2/2$, then a decrease in the transaction cost is likely, and so infinitely patient agents wait for the cost to become extremely low before transacting. If $\mu > \sigma^2/2$, then the transaction cost becomes prohibitively high if the agents wait indefinitely, and so it is not efficient for even infinitely patient agents to wait for the cost to become negligible. As agents become infinitely impatient, they transfer the good as soon as it is possible to obtain a positive payoff, thereby minimizing the effect of discounting. In the limit, the transaction cost incurred approaches the total stock of each good, causing the expected payoff of each agent to converge to zero.

We now study the relationship between the efficient solution and the maximal equilibrium in the case of a continuous Markov cost process. It is further assumed that the conditional distribution of future values of the cost divided by the current value does not depend on the cost realization up to the current time. That is, the transaction cost obeys the scaling rule below, whereby the incremental change in the cost is proportional to the current value of the cost.[5]

Definition 2. The positive cost process $\{c_t\}_{t\in[0,\infty)}$ is said to be **proportional** if the conditional distribution of $\{c_\tau/c_\kappa\}_{\tau\in(\kappa,\infty)}$ given $\{c_\upsilon\}_{\upsilon\in[0,\kappa]}$ does not vary with $\{c_\upsilon\}_{\upsilon\in[0,\kappa]}$.

Observe that any proportional cost process has the Markov property. Now consider a stationary efficient symmetric strategy profile and a stationary maximal symmetric SPE each of which induces a transaction with positive probability. The result below implies that the sole transaction when playing an efficient strategy profile happens sooner than all the transactions when playing maximal equilibrium strategies.

Theorem B. *Let* $\{c_t\}_{t\in[0,\infty)}$ *be a continuous and proportional cost process. Let* $\pi^{'}$ *be any efficient symmetric strategy profile for which there exists* $c^{'} > 0$ *such that, with positive probability, the realization of the cost process* $\{c_\tau\}_{\tau\in[0,\infty)}$ *satisfies* $c_t = c^{'}$ *and* $\phi_t^i(h_0, \{c_\tau\}_{\tau\in(0,\infty)}, \pi^{'}) > 0$ *for* $i \in \{1,2\}$ *and some* $t \in [0,\infty)$*. Let* $\pi^{''}$ *be any maximal symmetric SPE for which there exists* $c^{''} > 0$ *such that, with positive probability, the realization of the cost process* $\{c_\tau\}_{\tau\in[0,\infty)}$ *satisfies* $c_t = c^{''}$ *and* $\phi_t^i(h_0, \{c_\tau\}_{\tau\in(0,\infty)}, \pi^{''}) > 0$ *for* $i \in \{1,2\}$ *and some* $t \in [0,\infty)$*. Then* $c^{'} > c^{''}$*.*

[5]Note that a geometric Brownian motion has this property.

Proof. Suppose first that $c^{'} < c^{''}$. We show that this assumption leads to a contradiction. Let $\delta < 1$ denote the expected discounted value of an asset that pays 1 at the first time the cost process reaches $c^{'}$ given that the current value of the cost is $c^{''}$.

Because $\pi^{''}$ is a maximal symmetric SPE such that a transaction occurs at the cost $c^{''}$ with positive probability, there exists a cost realization $\{c^*_\tau\}_{\tau\in[0,\infty)}$ as well as a time t^* such that $c^*_{t^*} = c^{''}$, $\phi^i_{t^*}(h_0, \{c^*_\tau\}_{\tau\in(0,\infty)}, \pi^{''}) > 0$, and $V[h_{t^*}(\{c^*_\tau\}_{\tau\in[0,\infty)}, \pi^{''}), \pi^{''}] \geq V[h_{t^*}(\{c^*_\tau\}_{\tau\in[0,\infty)}, \pi^{''}), \tilde{\pi}]$ for every symmetric SPE $\tilde{\pi}$. Denote $u = V[h_{t^*}(\{c^*_\tau\}_{\tau\in[0,\infty)}, \pi^{''}), \pi^{''}]$.

It must be that $u - c^{''} \geq \delta(u - c^{'})$. Otherwise, the proportionality condition on the transaction cost would imply that the following symmetric SPE yields a higher expected payoff to each agent than $\pi^{''}$. The agents play $\pi^{''}$ until the first time t^{1*} that the current value of the cost process is $c^{''}$ and $\pi^{''}$ requires each agent to make a transaction at this time. Thereafter, the agents do not make any transactions until the first time t^{2*} greater than t^{1*} that the current value of the cost process is $c^{'}$. At time t^{2*}, each agent transfers the amount $\phi^i_{t^*}(h_0, \{c^*_\tau\}_{\tau\in(0,\infty)}, \pi^{''}) > 0$. Thereafter, the agents play according to strategy profile $\pi^{''}$, behaving as if the history at the time of this transaction were $h_{t^*}(\{c^*_\tau\}_{\tau\in[0,\infty)}, \pi^{''})$ and the value of the cost process at any successive time were $c^{''}/c^{'}$ multiplied by its actual value. That is, if the cost realization during the time interval of length l following this transaction is $\{g^*_\tau\}_{\tau\in(0,l]}$, then the players act as if the cost realization during this time interval were $\{(c^{''}/c^{'}) \cdot g^*_\tau\}_{\tau\in(0,l]}$.

Because $\pi^{'}$ is an efficient symmetric strategy profile such that a transaction occurs at the cost $c^{'}$ with positive probability, it must be that $q - c^{''} \leq \delta(q - c^{'})$. Otherwise, the strategy profile $\pi^{'}$ would yield a lower expected payoff to each agent than the symmetric strategy profile that requires each agent to transfer the amount q at the first time the cost reaches $c^{''}$.

The conditions $u - c^{''} \geq \delta(u - c^{'})$ and $q - c^{''} \leq \delta(q - c^{'})$ imply that $(q - c^{''})/(q - c^{'}) \leq \delta \leq (u - c^{''})/(u - c^{'})$. However, $u < q$ because the incentive compatibility of $\pi^{''}$ requires that $\phi^i_{t^*}(h_0, \{c^*_\tau\}_{\tau\in(0,\infty)}, \pi^{''}) < q$. It follows that $(q - c^{''})/(q - c^{'}) > (u - c^{''})/(u - c^{'})$. Thus, no value of δ satisfies the inequalities $(q - c^{''})/(q - c^{'}) \leq \delta \leq (u - c^{''})/(u - c^{'})$, resulting in a contradiction.

Suppose next that $c^{'} = c^{''}$. We show that this assumption leads to a contradiction. Denote $\tilde{c} = c^{'} = c^{''}$. Since $\pi^{''}$ is a maximal symmetric SPE such that a transaction occurs at the cost $\tilde{c}$ with positive probability, there exists a cost realization $\{c^\dagger_\tau\}_{\tau\in[0,\infty)}$ as well as a time $t^\dagger$ such that $c^\dagger_{t^\dagger} = \tilde{c}$, $\phi^i_{t^\dagger}(h_0, \{c^\dagger_\tau\}_{\tau\in(0,\infty)}, \pi^{''}) > 0$, and $V[h_{t^\dagger}(\{c^\dagger_\tau\}_{\tau\in[0,\infty)},\pi''),\pi''] \geq V[h_{t^\dagger}(\{c^\dagger_\tau\}_{\tau\in[0,\infty)},\pi''),\tilde{\pi}]$ for every symmetric SPE $\tilde{\pi}$. Denote $v = V[h_{t^\dagger}(\{c^\dagger_\tau\}_{\tau\in[0,\infty)},\pi''),\pi'']$. Because the incentive compatibility of π'' requires that $\phi^i_{t^\dagger}(h_0,\{c^\dagger_\tau\}_{\tau\in(0,\infty)},\pi'') < q$, it must be that $v < q$.

Choose any $\kappa > 1$ such that $\kappa v < q$. Let $\lambda_1 < 1$ denote the expected discounted value of an asset that pays 1 at the first time the cost reaches $\tilde{c}$ given that the current value of the cost is $\kappa\tilde{c}$. Let $\lambda_2 < 1$ denote the expected discounted value of an asset that pays 1 at the first time the cost reaches $\tilde{c}/\kappa$ given that the current value of the cost is $\tilde{c}$. Note that $\lambda_1 = \lambda_2$ because of the proportionality condition on the cost process. Let $\lambda = \lambda_1 = \lambda_2$.

Because π' is an efficient symmetric strategy profile such that a transaction occurs at the cost $\tilde{c}$ with positive probability, it must be that $q - \kappa\tilde{c} \leq \lambda(q - \tilde{c})$. Otherwise, the strategy profile π' would yield a lower expected payoff to each agent than the symmetric strategy profile that requires each agent to transfer the amount q at the first time the cost reaches $\kappa\tilde{c}$.

Because π'' is a maximal symmetric SPE such that a transaction occurs at the cost $\tilde{c}$ with positive probability, it must be that $v - \tilde{c} \geq \lambda(v - \tilde{c}/\kappa)$. Otherwise, the proportionality condition on the transaction cost would imply that the following symmetric SPE yields a higher expected payoff to each agent than π''. The agents play π'' until the first time $t^{1\dagger}$ that the current value of the cost process is $\tilde{c}$ and π'' requires each agent to make a transaction at this time. Thereafter, the agents do not make any transactions until the first time $t^{2\dagger}$ greater than $t^{1\dagger}$ that the current value of the cost process is $\tilde{c}/\kappa$. At time $t^{2\dagger}$, each agent transfers the amount $\phi^i_{t^\dagger}(h_0,\{c^\dagger_\tau\}_{\tau\in(0,\infty)},\pi'') > 0$. Thereafter, the agents play according to strategy profile π'', behaving as if the history at the time of this transaction were $h_{t^\dagger}(\{c^\dagger_\tau\}_{\tau\in[0,\infty)},\pi'')$ and the value of the cost process at any successive time were κ multiplied by its actual value. That is, if the cost realization during the time interval of length l following this transaction is $\{g^\dagger_\tau\}_{\tau\in(0,l]}$, then the players act as if the cost realization during this time interval were $\{\kappa\cdot g^\dagger_\tau\}_{\tau\in(0,l]}$.

The conditions $q-\kappa\tilde{c} \leq \lambda(q-\tilde{c})$ and $v-\tilde{c} \geq \lambda(v-\tilde{c}/\kappa)$ imply that $(v-\tilde{c})/(v-\tilde{c}/\kappa) \geq \lambda \geq (q-\kappa\tilde{c})/(q-\tilde{c})$. However, it follows from $\kappa v < q$ that $(v-\tilde{c})/(v-\tilde{c}/\kappa) < (q-\kappa\tilde{c})/(q-\tilde{c})$. Thus, no value of λ satisfies the inequalities $(v-\tilde{c})/(v-\tilde{c}/\kappa) \geq \lambda \geq (q-\kappa\tilde{c})/(q-\tilde{c})$, resulting in a contradiction. □

The proof is by contradiction. As in the statement of the theorem, let c' and c'' be the respective cost thresholds at which a transaction occurs with positive probability in the efficient solution $\pi^{'}$ and maximal equilibrium $\pi^{''}$. First, consider the possibility that $c^{'} < c^{''}$. In this case, the strategy profile $\pi^{'}$ would be Pareto dominated by a symmetric strategy profile in which the agents instead transact at the cost $c^{''}$. Next, consider the possibility that $c^{'} = c^{''}$. In this case, the SPE $\pi^{''}$ would be Pareto dominated by a symmetric SPE in which the agents instead transact at a cost slightly lower than $c^{'}$.

IV.2 General Condition for Asymptotic Efficiency

The result below establishes a general property of the cost process under which the efficient outcome can be approximated in equilibrium as discounting frictions disappear. In the limit as agents become infinitely patient, all the potential gains from trade are realized.

Theorem C. *Assume that $\{c_t\}_{t\in[0,\infty)}$ is an arbitrary right-continuous cost process and that each random variable c_t for $t \geq 0$ takes values in the state space $S \subseteq \mathbb{R}_{++}$. Suppose that there exists $r < 1$ such that for any $\epsilon > 0$, one can find $p > 1 - \epsilon$ and $v > 0$ for which given any realization of the cost process $\{c_t\}_{t\in[0,u]}$ up to an arbitrary time u, there is conditional probability no less than p that the cost process $\{c_t\}_{t\in(u,\infty)}$ after time u satisfies $c_\tau \leq rc_u$ for some $\tau \in (u, u+v)$. Then for any sequence $\{\rho_n\}_{n=1}^{\infty}$ of discount rates with $\lim_{n\to\infty} \rho_n = 0$, there exists a sequence $\{\psi_n\}_{n=1}^{\infty}$ such that ψ_n is a symmetric SPE when the discount rate is ρ_n and such that the expected payoff to each agent when ψ_n is played and the discount rate is ρ_n converges to q in the limit as n goes to infinity.*

Proof. We begin by constructing a sequence $\{\tilde{\psi}_n\}_{n=1}^{\infty}$ of symmetric strategy profiles and a sequence $\{\tilde{\rho}_n\}_{n=1}^{\infty}$ of discount rates such that $\tilde{\psi}_n$ is an SPE when the discount rate is no greater than $\tilde{\rho}_n$ and such that the expected payoff to each agent when $\tilde{\psi}_n$ is played and the discount rate is $\tilde{\rho}_n$ converges to q in the limit as n goes to infinity. Let $r < 1$ be as defined in the statement of the theorem. For each index n, one can find $p_n > 1 - 1/n^2$ and $v_n > 0$ such that given any realization of the cost process $\{c_t\}_{t\in[0,u]}$ up to an arbitrary time u, there is conditional probability no less than p_n that the realization of the cost process $\{c_t\}_{t\in(u,\infty)}$ after time u satisfies $c_\tau \leq rc_u$ for some $\tau \in (u, u + v_n)$. Hence, given any realization of the cost process $\{c_t\}_{t\in[0,u]}$ up to an arbitrary time u, there is conditional probability no less than p_n^n that the realization of the cost process $\{c_t\}_{t\in(u,\infty)}$ after time u satisfies $c_\tau \leq r^n c_u$ for some $\tau \in (u, u + n \cdot v_n)$.

For each index n, define the discount rate as $\tilde{\rho}_n = (n^2 v_n)^{-1}$. For each index n, construct the symmetric SPE $\tilde{\psi}_n$ as follows. Choose c_n^* equal to the smaller of c_0 and $\exp(-\tilde{\rho}_n \cdot n \cdot v_n) \cdot p_n^n \cdot q \cdot r^n \cdot (1 - r^n)$. The first transaction occurs at the first time that the current value of the cost is less than or equal to c_n^*. If the previous transaction occurred at cost $\hat{c}$, then the next transaction occurs at the first time that the cost is less than or equal to $r^n \hat{c}$. For every positive integer k, each agent transfers the amount $r^{n(k-1)} q(1 - r^n)$ on the k^{th} transaction. If an agent deviates from the path of play described above, then neither agent makes any transactions following the deviation. It can be shown as in the proofs of the results in the main text that the strategy profile $\tilde{\psi}_n$ is an SPE when the discount rate is no greater than $\tilde{\rho}_n$.

Consider the sequence $\{\tilde{\psi}_n\}_{n=1}^\infty$ of symmetric SPE and the sequence $\{\tilde{\rho}_n\}_{n=1}^\infty$ of discount rates. Note that the cost c_n^* paid on the first transaction converges to zero in the limit as n goes to infinity. Note that the amount $q(1-r^n)$ transferred on the first transaction converges to q in the limit as n goes to infinity. Moreover, when playing strategy profile $\tilde{\psi}_n$, the continuation value after transaction n cannot be negative.

Let m_n be the least integer greater than or equal to $\log_r(c_n^*/c_0)$. In the limit as n goes to infinity, the probability that the first transaction occurs by time $m_n \cdot v_n$ converges to a number no less than:

$$\begin{aligned}\lim_{n\to\infty} p_n^{m_n} &\geq \lim_{n\to\infty} (1 - 1/n^2)^{m_n} = \lim_{n\to\infty} (1 - 1/n^2)^{\log_r(c_n^*/c_0)+1} = \lim_{n\to\infty} (1 - 1/n^2)^{\log_r(c_n^*)} \\ &= \lim_{n\to\infty} (1 - 1/n^2)^{\log_r[\exp(-\tilde{\rho}_n \cdot n \cdot v_n)\cdot p_n^n \cdot q \cdot r^n \cdot (1-r^n)]} = \lim_{n\to\infty} (1 - 1/n^2)^{\log_r(r^n)} \\ &= \lim_{n\to\infty} (1 - 1/n^2)^n = 1.\end{aligned}$$

In the limit as n goes to infinity, the discount factor at time $m_n \cdot v_n$ converges to:

$$\begin{aligned}\lim_{n\to\infty} \exp(-\tilde{\rho}_n m_n v_n) &= \lim_{n\to\infty} \exp\{-(n^2 v_n)^{-1}[\log_r(c_n^*/c_0) + 1] v_n\} \\ &= \lim_{n\to\infty} \exp\{-n^{-2}[\log_r(c_n^*) - \log_r(c_0) + 1]\} \\ &= \lim_{n\to\infty} \exp\big(-n^{-2}\{\log_r[\exp(-\tilde{\rho}_n n v_n) p_n^n q r^n (1 - r^n)] - \log_r(c_0) + 1\}\big) \\ &= \lim_{n\to\infty} \exp\big(-n^{-2}\{n + \log_r[\exp(-\tilde{\rho}_n n v_n) p_n^n q(1 - r^n)] - \log_r(c_0) + 1\}\big) = 1.\end{aligned}$$

It follows from the observations in the preceding two paragraphs that the expected payoff to each agent when strategy profile $\tilde{\psi}_n$ is played and the discount rate is no greater than $\tilde{\rho}_n$ converges to q in the limit as n goes to infinity.

Let $\{\rho_n\}_{n=1}^{\infty}$ be an arbitrary sequence of discount rates with $\lim_{n\to\infty} \rho_n = 0$. For any index n, let m_n denote the greatest positive integer k such that $\tilde{\rho}_k$ is no less than ρ_n. If ρ_n is greater than $\tilde{\rho}_k$ for all k, then let ψ_n be the symmetric strategy profile that requires each agent to never make a transfer at any history. Otherwise, let ψ_n be the same as the symmetric strategy profile $\tilde{\psi}_{m_n}$. Note that ψ_n is an SPE when the discount rate is ρ_n.

We now argue that the expected payoff to each agent when ψ_n is played and the discount rate is ρ_n converges to q in the limit as n approaches infinity. Choose any $\epsilon > 0$. By construction, there exists an index k^* such that the expected payoff to each agent when $\tilde{\psi}_n$ is played and the discount rate is $\tilde{\rho}_n$ is greater than $q - \epsilon$ for all $n \geq k^*$. In addition, there exists an index l^* such that $\rho_n < \tilde{\rho}_{k^*}$ for all $n \geq l^*$. Note that the expected payoff to each agent when ψ_n is played and the discount rate is ρ_n is greater than $q - \epsilon$ for all $n \geq l^*$. $\square$

The following is a summary of the proof. First, a sequence of discount rates converging to zero is specified. Next, a nondegenerate symmetric equilibrium is constructed for each discount rate in the sequence. The path of play involves a potentially infinite sequence of transactions with decreasing amounts transferred and costs incurred. The resulting sequence of equilibria is such that the cost paid on the first transaction converges to zero while the size of the initial transfer approaches the total stock of each good. Moreover, the probability of a transaction happening converges to one, and the first transaction occurs sufficiently rapidly that discounting becomes negligible in the limit. Hence, the expected payoff to each player approximates the total stock of each good.

V Variations of Basic Model

This section analyzes variations of the baseline framework. Appendix V.1 extends the setup to allow the supply of each good to follow a stochastic process. Appendix V.2 presents an extension in which the transfer size affects the transaction cost.

V.1 Model with Stochastic Supplies of Goods

This appendix analyzes an environment in which the supplies of the goods vary randomly over time. The framework here is the same as that presented in the main text, except that the cost of making a transaction is fixed at some positive constant and the stock of each good is assumed to follow a stochastic process between transactions.[6] In particular, assume that the transaction cost is constant at $\chi > 0$. For any positive integer k, let t_k^i denote the time of the k^{th} transaction by agent i. In addition, let x_k^i be the amount of good i transferred by agent i on the k^{th} transaction as a fraction of the amount of good i remaining immediately before the k^{th} transaction. Let $\{g_t\}_{t\in[0,\infty)}$ be a strictly positive stochastic process. Assume that the amount of good i remaining at the end of time $t \geq 0$ is given by $s_t^{i,i} = g_t \cdot \prod_{\{k:\, 0\leq t_k^i \leq t\}}(1 - x_k^i)$.

The stochastic process $\{g_t\}_{t\in[0,\infty)}$ captures random growth or decay in the stock of each good. The following theorem establishes a condition on this process under which the model has a nondegenerate equilibrium. Assume that the process is right-continuous. Suppose that there exist constants $p > 0$, $r > 1$, and $v > 0$ for which there is conditional probability p of the process growing by a factor of r during a time interval of length v.[7] Then the result below shows that the extended model has an equilibrium in which a transaction occurs with positive probability.

Theorem D. *Assume that $\{g_t\}_{t\in[0,\infty)}$ is an arbitrary right-continuous stochastic process and that each random variable g_t for $t \geq 0$ takes values in the state space $S \subseteq \mathbb{R}_{++}$. Suppose that one can find $p > 0$, $r > 1$, and $v > 0$ for which given any realization of the process $\{g_t\}_{t\in[0,u]}$ up to an arbitrary time u, there is conditional probability no less than p that the process $\{g_t\}_{t\in(u,\infty)}$ after time u is such that $g_\tau \geq r g_u$ for some $\tau \in (u, u+v)$. Then the model in this section has a symmetric SPE π in which there is positive probability of each agent making a positive transfer at some time.*

Proof. Consider the symmetric grim-trigger strategy profile ψ defined as follows. Letting $\delta = e^{-\rho v}$, choose any κ no less than $[r/(r-1)][1 + (1-\delta p)/(\delta p)]$. The first transaction occurs at the first time that the amount of each good remaining is greater

[6] In addition, it is assumed as in the analysis of the model in the main text that each agent consumes the good from the other agent as soon as it is received.

[7] This property is satisfied by a geometric Brownian motion with arbitrary drift and positive volatility.

than or equal to $\kappa\chi$. If the stock of each good was $\tilde{s}$ immediately before the previous transaction, then the next transaction occurs at the first time that the stock of each good is greater than or equal to $\tilde{s}$. On each transaction, agent i transfers a fraction $(r-1)/r$ of the amount of good i remaining before the transaction. If an agent deviates from the path of play described above, then neither agent makes any transactions following the deviation.

We argue that strategy profile ψ is an SPE. It suffices to show that the incentive compatibility constraint is satisfied at each transaction when playing ψ. Suppose that the agents have followed strategy profile ψ up to the current time and that the next transaction will occur at the first time the stock of each good is at least $\tilde{s}$. If the agents follow strategy profile ψ, then the cost incurred by each agent on the next transaction is χ, and the expected payoff to each agent immediately after the next transaction is no less than $\sum_{m=1}^{\infty} \delta^m p^m \{[(r-1)/r]\kappa\chi - \chi\}$. Hence, the incentive compatibility constraint is satisfied for the next transaction if the following holds:

$$\chi \leq \sum_{m=1}^{\infty} \delta^m p^m \{[(r-1)/r]\kappa\chi - \chi\},$$

which reduces to:

$$\chi \leq [\delta p/(1-\delta p)]\{[(r-1)/r]\kappa - 1\}\chi \quad \Leftrightarrow \quad \kappa \geq [r/(r-1)][1 + (1-\delta p)/(\delta p)].$$

The last inequality is true by assumption, confirming that the incentive compatibility constraint is satisfied. $\square$

The proof is straightforward. A nondegenerate equilibrium can be constructed using grim-trigger strategies. It has the following form. Each agent transfers a specified fraction of the good at every time that the remaining stock of each good meets or exceeds a particular proportion of the transaction cost. If an agent deviates from this path of play, then no further transfers are made. The critical value of the stock and the size of each transfer can be chosen such that the continuation value from the relationship is at least as large as the transaction cost. Consequently, neither agent has an incentive to deviate.

The theorem is valid even if the stock of each good has a very high growth rate. For example, suppose that the stochastic process $\{g_t\}_{t\in[0,\infty)}$ evolves according to a geometric Brownian motion. If the drift parameter μ is greater than the discount rate

ρ, then there is no upper bound on the expected payoff attainable in equilibrium. As previously explained, an equilibrium can be supported using grim-trigger strategies. Following a deviation from the path of play, the maximum payoff an agent can secure is zero, given the strategy played by the other agent.

The theorem holds even when the discount rate of each agent is very high. An equilibrium can be implemented as described above. In the case where the future is heavily discounted, the agents are required to wait for the stock to become relatively large before making a transaction. When there is a sizeable remaining stock, the continuation value of the relationship is substantial. In particular, there is a positive conditional probability of the stochastic process $\{g_t\}_{t\in[0,\infty)}$ growing by a factor of $r > 1$ during a certain length of time. By allowing the stock to become sufficiently high, the absolute increase in the stock due to a given proportional increase can be made arbitrarily large. Thus, the continuation value can be made big enough to prevent the agents from deviating.

Also observe that the theorem above offers a counterpoint to the impossibility result in the main text. Suppose that the transaction cost is fixed at a positive constant. If the total supply of each good remains constant over time, then no transactions can occur in an equilibrium of the model. However, if the stock of each good evolves randomly as in the preceding theorem, then a transaction can occur in equilibrium with positive probability.

Finally, the model with uncertainty in the supply of each good can be applied to some of the examples in the introduction, particularly negotiations between countries over the release of prisoners. The stocks may vary over time because of the capture of additional combatants during military operations, the death of prisoners while in custody, or the escape of captives from detention facilities. An immediate payoff is received when prisoners are repatriated. This payoff might represent a rise in public support for elected officials or a psychic benefit from the homecoming of missing family members. The cost of releasing detainees might include criticism from political opponents or a worsening of national security. The next appendix studies the case where this cost depends on the size of the transfer.

V.2 Model with Cost Proportional to Amount Transferred

This appendix examines the effect of letting the transaction cost paid depend on the amount transferred. The setup here is the same as that described in the main text, except that the cost paid at each transaction includes a term that is proportional to the amount transferred.[8] Specifically, if agent $i \in \{1, 2\}$ transfers the amount x_t^i at time $t \in [0, \infty)$ and the fixed cost of making a transfer is c_t at time t, then agent i incurs the transaction cost $c_t + \phi \cdot x_t^i$ at time t, where $\phi \in (0, 1)$.[9]

The following result demonstrates that the game may have a nondegenerate equilibrium. Consider a right-continuous cost process. Suppose that there exist constants $p > 0$, $r < 1$, and $v > 0$ for which there is conditional probability p of the fixed cost becoming a fraction r of its current value during a time interval of length v.[10] According to the theorem below, there exists an equilibrium of the extended model in which a transaction occurs with positive probability, provided that the component of the cost proportional to the amount transferred is not excessively large.[11]

Theorem E. *Assume that $\{c_t\}_{t \in [0,\infty)}$ is an arbitrary right-continuous cost process and that each random variable c_t for $t \geq 0$ takes values in the state space $S \subseteq \mathbb{R}_{++}$. Suppose that one can find $p > 0$, $r < 1$, and $v > 0$ for which given any realization of the cost process $\{c_t\}_{t \in [0,u]}$ up to an arbitrary time u, there is conditional probability no less than p that the cost process $\{c_t\}_{t \in (u,\infty)}$ after time u is such that $c_\tau \leq rc_u$ for some $\tau \in (u, u+v)$. Then there exists $\bar{\phi}$ such that for $\phi < \bar{\phi}$, the model in this section has a symmetric SPE π in which there is positive probability of each agent making a positive transfer at some time.*

Proof. Letting $\delta = e^{-\rho v}$, define $\bar{\phi} = \delta pr$. Assume that $\phi < \bar{\phi}$. Consider the symmetric grim-trigger strategy profile ψ defined as follows. Recall that q denotes the initial stock of each good. Choose any value of c^* no greater than $q(1-r)(\delta pr - \phi)$. The first transaction occurs at the first time that the current value of the cost is less than or equal to c^*. If the previous transaction occurred at cost $\hat{c}$, then the next transaction occurs at the first time that the cost is less than or equal to $r\hat{c}$. For every positive integer k, each agent transfers the amount $r^{k-1}q(1-r)$ on the k^{th} transaction. If an

[8] In addition, it is assumed as in the analysis of the model in the main text that each agent consumes the good from the other agent as soon as it is received.

[9] Note that the current model would reduce to the basic model if $\phi = 0$.

[10] This property is satisfied by a geometric Brownian motion with arbitrary drift and positive volatility.

[11] If the fixed cost c_t follows a geometric Brownian motion with any drift and positive volatility, then Theorem E implies that the component of the cost proportional to the amount transferred can be arbitrary close to the amount transferred. This holds because we can respectively choose p and v to be sufficiently close to one and zero and then let r approach one.

agent deviates from the path of play described above, then neither agent makes any transactions following the deviation.

We argue that strategy profile ψ is an SPE. Note that ψ is feasible because $\sum_{k=1}^{\infty} r^{k-1} q(1-r) = q$. We next show that the incentive compatibility constraint is satisfied at each transaction when playing ψ. Choose any positive integer l. Suppose that the agents have followed strategy profile ψ up to the current time, $l-1$ transactions have happened in the past, and ψ specifies transaction l will occur at the first time the cost is at most $\hat{c}$. If the agents follow strategy profile ψ, then the cost incurred by each agent on transaction l is no greater than $r^{l-1}c^* + r^{l-1}q(1-r)\phi$, and the expected payoff to each agent immediately after transaction l is no less than $\sum_{m=1}^{\infty} \delta^m p^m [r^{l+m-1} q(1-r)(1-\phi) - r^{l+m-1}c^*]$. Hence, the incentive compatibility constraint is satisfied for transaction l if the following holds:

$$r^{l-1}c^* + r^{l-1}q(1-r)\phi \leq \sum_{m=1}^{\infty} \delta^m p^m [r^{l+m-1} q(1-r)(1-\phi) - r^{l+m-1}c^*],$$

which reduces to:

$$c^* + q(1-r)\phi \leq \delta pr[q(1-r)(1-\phi) - c^*]/(1-\delta pr) \quad \Leftrightarrow \quad c^* \leq q(1-r)(\delta pr - \phi).$$

The last inequality is true by assumption, confirming that the incentive compatibility constraint is satisfied. □

The assumption regarding the fixed cost is the same as for the corresponding result in the main text, which identifies a condition such that the model has a nondegenerate equilibrium. The proof is also similar. If the part of the cost dependent on the amount transferred is sufficiently small, then grim-trigger strategies can be used to support a nondegenerate equilibrium. The sequence of transactions is potentially infinite, and the amount transferred and the cost incurred are gradually decreasing. Hence, it is not crucial to assume that the transaction cost is insensitive to the amount transferred in order to support an equilibrium with positive gains from trade.