

SEQUENTIAL EXCHANGE WITH STOCHASTIC TRANSACTION COSTS*

Yuichiro Kamada[†] Neel Rao[‡]

August 22, 2018

Abstract

We analyze the bilateral trade of divisible goods in the presence of stochastic transaction costs. Mutual best-response conditions are applied to a model of optimal investment under uncertainty. The first-best solution involves a single transaction, but such behavior is not incentive compatible without court-enforceable contracts. We solve for a second-best policy in which some gains from trade can be realized through a gradual transfer. When the transaction cost follows a geometric Brownian motion, the optimal path of play in a subgame-perfect equilibrium is unique, and a closed-form solution can be obtained. A number of comparative statics and welfare implications are presented as well as real-world examples.

JEL Classification: C73, C78, D23, D86, L14

Keywords: dynamic games, continuous time, transaction costs, stochastic processes

*We thank Eduardo Faingold, Drew Fudenberg, Oliver Hart, Michihiro Kandori, Akitada Kasahara, Phillip Strack, Takuo Sugaya, Satoru Takahashi, Juuso Toikka, Yasutora Watanabe, and seminar participants at Chinese University of Hong Kong, Harvard University/Massachusetts Institute of Technology, Hitotsubashi University, University of California at Davis, University of Tokyo, and the Summer Workshop on Economic Theory 2018 at Otaru University of Commerce for helpful comments.

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

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

1 Introduction

This paper is motivated by situations like the following. Two firms operating in different markets find it potentially profitable to exchange trade secrets, but there is a cost for transferring knowledge from one firm to the other. This cost might represent the expense of encrypting data to protect secrets from outsiders or the resources spent on administrative paperwork. Due to the intangible nature of information, it is infeasible for the two parties to write a court-enforceable contract specifying the goods to be traded. Moreover, if one party immediately reveals all of its information to the other party, then the latter would have no incentive to reveal its information to the former, because such transactions are costly. In this setting, how can the two parties arrange trade with each other?

We study the exchange of divisible goods in a continuous-time environment. Each of two agents is endowed with an equal amount of a good that only the other party values. In order to transfer some of its good to the other party, an agent must pay a transaction cost that evolves according to a stochastic process. The first-best solution requires each agent to make a one-shot transfer of the entire stock of its good to the other agent. Nonetheless, such a policy cannot be supported in a subgame-perfect equilibrium. It is not incentive compatible because the agents are unwilling to incur the transaction cost if positive future transfers are impossible.

Our ultimate objective is to solve for a second-best transaction scheme that enables the agents to realize positive gains from trade. We establish conditions under which the game has a nondegenerate equilibrium, and the general form of a solution is determined. Assuming that the cost process follows a geometric Brownian motion, a closed-form expression can be obtained for the unique path of play in a maximal equilibrium.¹ It involves a potentially infinite and gradually decreasing sequence of transfers. We derive a number of comparative statics and welfare implications.

The basic idea behind the implementation of a self-enforcing agreement is simple. In equilibrium, when the two agents make transfers, they withhold some amount of the goods. The players use grim-trigger strategies, whereby they continue to transfer the withheld goods when the transaction cost reaches a certain threshold if and only if both parties have made the prescribed transfers in the past. Because the contin-

¹A symmetric equilibrium is said to be maximal if no symmetric equilibrium yields a strictly higher expected payoff to each agent.

uation value is positive if and only if each party transfers its goods as stipulated, the agents are willing to incur a positive transaction cost, enabling trade to occur in equilibrium. Thus, the fundamental mechanism for sustaining cooperation is akin to that in repeated games, where future rewards and punishments are used to prevent deviation.²

Nonetheless, our analysis differs from much existing work on stochastic games in that the state transition in our model is non-irreducible.³ In particular, as agents make transfers, the remaining stock of each good decreases irreversibly. A related complication is that a single agent can obstruct a state transition. For example, a player can prevent the supply of its good from falling simply by refusing to make any further transfers.

This sort of non-irreducibility generates a key tradeoff between *the size of the next transfer* and *the waiting time until the transfer*. In order for agents to anticipate a high payoff, they must exchange a large quantity of each good at the next transaction. Hence, the amount withheld must be low at that transaction. Because the resulting continuation value is low, the agents are willing to incur only a small cost when making the transaction. However, the expected waiting time for a small cost to realize is high. Since agents discount the future, a lengthy waiting time reduces the expected payoff. We solve for a subgame-perfect equilibrium that optimizes the tradeoff between the transfer size and waiting time.

Methodologically, we combine the theory of investment under uncertainty with the theory of repeated games by imposing *mutual* best-response constraints on investment decisions.⁴ That is, the set of admissible investment policies is restricted by the incentive-compatibility requirements of a subgame-perfect equilibrium. Comparing the first- and second-best outcomes of our model, we show that the problem of incentive compatibility delays the realization of the gains from trade.

The game is specified in continuous time, which has a number of advantages.⁵ First, a closed-form solution for the maximal equilibrium is obtained with further assumptions on the cost process. The resulting expression uniquely determines the

²For example, see Fudenberg and Maskin (1986).

³See Dutta (1995), Fudenberg and Yamamoto (2011), and Hörner, Sugaya, Takahashi, and Vielle (2011) for folk theorems in stochastic games at varying levels of generality.

⁴See footnote 9 for references to the literature on investment under uncertainty.

⁵Continuous-time modeling also involves some complications as noted by Simon and Stinchcombe (1989). The methodology newly developed by Kamada and Rao (2018) is applied to suitably define strategy spaces in the game here.

cost incurred and amount transferred on each transaction. This facilitates the analysis of comparative statics. In addition, if the true model is one in which the transaction cost evolves continuously over time, then welfare might be lower when players are limited to moving at discrete times instead of responding instantly. Rational agents may be unwilling to constrain their behavior in this way.

Besides the opening example, our framework fits a number of social situations in the real world where exchanges are involved. For instance, consider the transfer of prisoners or hostages between groups that are enemies or countries that are fighting. The two agents in the model would be representatives of the parties who are negotiating for the release of detainees. Each party wishes to trade captives from the other side for the return of prisoners from its own side. The stochastic transaction cost might be a measure of the political climate or degree of tensions between the two groups. When tensions are high and the relationship is hostile, each group has a high cost of releasing prisoners of war. On the other hand, the cost is low when tensions are low and the relationship is not hostile.

In many cases, exchanges of prisoners occur in a gradual manner, as our model predicts. One example comes from the negotiations over the repatriation of prisoners towards the end of the Korean War. A release from the Associated Press (1953) describes a plan under which “the United Nations will return 500 disabled North Koreans and Chinese daily in exchange for 100 United Nations sick and wounded,” resulting in a total of “over 600 United Nations men in exchange for 5,800 Communists.” Likewise, the transfer of captives during the Civil War in El Salvador proceeded in several stages. According to a report by LeMoyne (1984), “The wounded guerillas left in four groups, numbering 15 each. As each group was permitted to leave El Salvador the rebels freed two army officers.” Finally, a similar principle was instrumental in the resolution of the Lebanon hostage crisis. In an article by Haberman (1991), a senior official is quoted as stating that “this will be the first of a series of stages at the end of which the question of the hostages, of the prisoners of war, of the missing in action, will come to a conclusion.” The chief hostage negotiator also remarks, “We believe that this cannot be done in one shot. This is a step-by-step process. We must take it one step at a time.”

A further prediction of our model is that the size of transfers will decrease over time. An example of this pattern in the real world is provided by Article Two of the Second Treaty of Indian Springs (1825). The following extract describes how the

United States intended to compensate the Creek Indians for lands ceded to the state of Georgia:

The United States agree to pay to the nation emigrating from the lands herein ceded the sum of four hundred thousand dollars; of which amount there shall be paid to said party of the second part, as soon as practicable after the ratification of this treaty, the sum of two hundred thousand dollars. And as soon as the said party of the second part shall notify the Government of the United States of their readiness to commence their removal, there shall be paid the further sum of one hundred thousand dollars. And the first year after said emigrating party shall have settled in their new country, they shall receive, of the amount first above named, the further sum of twenty-five thousand dollars; and the second year, the sum of twenty-five thousand dollars; and annually thereafter, the sum of five thousand dollars, until the whole is paid.

According to the agreement, the United States would make a gradually decreasing series of monetary payments to the Creek Indians as they progressively surrendered their existing lands and relocated to a different territory.

The remainder of this paper is organized as follows. Immediately after the introduction is a review of the related literature. Section 2 outlines our basic model. Section 3 analyzes the model in several steps. We start by identifying general conditions on the cost process under which a nondegenerate equilibrium does and does not exist. Thereafter, we characterize the basic properties of a maximal equilibrium. In the case where the cost process follows a geometric Brownian motion, a closed-form expression can be obtained for the optimal path of play. Section 4 presents comparative statics for the unique solution. Section 5 investigates the welfare implications of the model, and section 6 discusses some modeling extensions and robustness checks in the online appendix. Section 7 concludes. The proofs of the results in the body of the paper are provided in the main appendix.

1.1 Related Literature

Our paper connects with a line of research on gradualism in contribution games and concession bargaining.⁶ In this literature, parties arrive at an agreement in a step-

⁶See Admati and Perry (1991); Compte and Jehiel (1995, 2003, 2004); Marx and Matthews (2000). See also Gale (1995, 2001) and Gueron (2015) for monotone games.

by-step fashion, and there is an efficiency loss due to delay in reaching an agreement. Likewise, cooperation between the two players in our model is sustained through a gradual sequence of transactions over time. There is a key modeling difference between our paper and, for example, Admati and Perry (1991). In their setup, a benefit from cooperation is realized only when a joint project is completed, whereas in our framework, a benefit is received every time a transaction takes place. Additionally, our model differs from much of this literature in that transactions continue indefinitely in equilibrium; so that, the game we study cannot be solved using an iterated-dominance procedure.

In our model, agents incur costs whenever they take actions. Thus, another point of comparison is the literature on repeated interaction with switching costs. In particular, Lipman and Wang (2000, 2009) show that equilibrium dynamics are nontrivial in repeated games with finite and infinite horizons where switching costs are large compared to stage-game payoffs. Caruana and Einav (2008) derive similar results for the case where agents can revise their actions over a finite horizon with increasing switching costs. Nonetheless, the role of switching costs in those papers is different from the effect of transaction costs in our model. In those papers, the fact that switching costs are large in the future is important for the results, as such costs can serve as a commitment device. By contrast, no cooperation occurs in the equilibrium of our model if the transaction cost is bounded away from zero with probability one.

Exchange mechanisms are extensively studied in many different contexts that involve trading favors. The seminal paper on this topic is Mobius (2001). Later work includes: Abdulkadiroglu and Bagwell (2012) and Hauser and Hopenhyan (2005) for the “chips mechanism”; Johnson, Simchi-Levi, and Sun (2014) for the “scrip system”; and Wolitzky (2015) for the exchange of tokens. The distinguishing feature of our mechanism is that devices such as chips or scrips or tokens are not needed to sustain cooperation. It is the remaining supply of each good that is used to reward cooperators.

Pitchford and Snyder (2004) consider a holdup problem between a buyer and a seller in which no investment occurs in the equilibrium of the static game. In a dynamic version, positive investment can be supported in an equilibrium where the seller’s investment and the buyer’s repayment take place alternately. Those authors observe that such an equilibrium would not exist in their deterministic setting if agents must incur a fixed cost for transacting that does not depend on the amount invested.

We resolve this issue by allowing for uncertainty in the transaction cost, in which case agents may be able to realize positive gains from trade in a subgame-perfect equilibrium.⁷ The stochastic nature of the cost may be more realistic in practice as economic conditions or the political climate may vary over time.⁸ Furthermore, the welfare implications of our model are somewhat more nuanced than in Pitchford and Snyder (2004), where the equilibrium converges to the efficient outcome as discounting frictions disappear. In our setup, the fixed cost of making a transfer creates an additional friction, and so such convergence is not guaranteed. We identify conditions on the cost process under which the efficient outcome can be approximated in equilibrium as agents become infinitely patient.

Several theoretical results are presented in the context of general stochastic processes in continuous time. To obtain an explicit solution for the optimal strategies, we also consider the special case where the transaction cost follows a geometric Brownian motion. The assumption of geometric Brownian motion is standard in the finance literature on investment under uncertainty.⁹ Similarly, a number of papers on the principal-agent problem in continuous time, including Holmstrom and Milgrom (1987) and Sannikov (2008), assume that the output process follows an arithmetic Brownian motion.

Although a continuous-time setting is useful in obtaining closed-form solutions and comparative statics, there are subtle technicalities in defining strategy spaces due to the possibility of instantaneous responses. Simon and Stinchcombe (1989) and Bergin and MacLeod (1993) introduced methods to overcome such a problem in deterministic models, but they do not fit our stochastic environment. For this reason, we use the framework recently introduced by Kamada and Rao (2018), which also resolves technical issues related to the measurability of stochastic processes in continuous time.

⁷As noted in section 6, the online appendix also examines a setting in which the quantity of goods available for trade can vary randomly over time. With uncertainty in the supplies of goods, positive transfers may be sustainable in equilibrium even if the transaction cost is a positive constant.

⁸We give further justification for stochastic transaction costs in section 6.

⁹For example, geometric Brownian motion is used by McDonald and Siegel (1986) to model the market value of an investment, by Dixit (1991) to model a demand parameter, and by Bertola and Caballero (1994) to model an index of business conditions. He (2009) studies optimal managerial compensation under the assumption of a geometric Brownian motion in firm size. Empirical evidence is presented in support of the model. Merton (1971) derives optimal portfolio rules under the benchmark hypothesis of a geometric Brownian motion in asset prices. He also discusses the effect of alternative modeling assumptions.

2 Model

There are two agents, 1 and 2, who take actions in continuous time $t \in [0, \infty)$. The discount rate is $\rho > 0$. There are two divisible goods, 1 and 2. The allocation of the goods at time t is represented by $[(s_t^{1,1}, s_t^{1,2}), (s_t^{2,1}, s_t^{2,2})]$, where $s_t^{i,j}$ denotes the amount of good j that agent i possesses at time t .¹⁰ The total supply $q > 0$ of each good is constant over time; so that, $s_t^{1,j} + s_t^{2,j} = q > 0$ for $j \in \{1, 2\}$ and $t \in [0, \infty)$.¹¹ The initial endowment vector is $s_0 = [(q, 0), (0, q)]$.¹² That is, agent 1 is endowed with all of good 1, and agent 2 is endowed with all of good 2.¹³

In addition, there is a positive transaction cost c_t for transferring goods between the two parties, which changes over time according to some stochastic process.¹⁴ For ease of exposition, let $c_0 > q$.¹⁵ We assume that the cost process evolves independently from the actions of the agents.¹⁶

In every instant of time, each agent observes the current realization of the cost and chooses an amount of her good to transfer to the other agent. For $i \in \{1, 2\}$ and $t \in [0, \infty)$, let $f_t^i \in [0, q]$ represent the amount of good i that agent i transfers at time t . A history of the game is represented as $h = \{c_t, (f_t^1, f_t^2)\}_{t \in [0, \infty)}$. That is, a history consists of the realization of the cost process along with the transfers made by the agents at each time. Given a history h , a history h_u up to time u is defined as

¹⁰If a transaction occurs at time t , then $s_t^{i,j}$ is the amount of good j that agent i possesses immediately after this transaction.

¹¹As described in section 6, the online appendix presents a model in which the supply of each good can vary over time.

¹²The equilibria of the model do not necessarily depend on the assumption that the agents have symmetric endowments of goods. If π is an SPE when each agent i is endowed with the quantity q of good i , then a model in which each agent i is endowed with any quantity $q_i \geq q$ of good i also has the SPE π as well as an SPE that is the same as π except that each agent i transfers the additional amount $q_i - q$ at the first transaction on the equilibrium path.

¹³This assumption is without loss of generality provided that $s_0^{1,1} = s_0^{2,2}$.

¹⁴The equilibria of the model are robust to some relaxations of the assumption that both agents face the same transaction cost. Let $\{c_t^i\}_{t \in [0, \infty)}$ for each $i \in \{1, 2\}$ be a cost process such that $c_t^i(\omega) \leq c_t(\omega)$ for all $t \in [0, \infty)$ and every $\omega \in \Omega$. Any SPE in grim-trigger strategies when the transaction cost is $\{c_t\}_{t \in [0, \infty)}$ for each party is also an SPE of a model in which the cost is $\{c_t^1\}_{t \in [0, \infty)}$ for agent 1 and $\{c_t^2\}_{t \in [0, \infty)}$ for agent 2.

¹⁵This condition will help ensure that the initial value of the transaction cost is sufficiently high that it is a strictly dominated strategy for an agent to make a transaction right at the beginning of the game.

¹⁶To be clear, $\{c_t\}_{t \in [0, \infty)}$ is a stochastic process on the probability space (Ω, \mathcal{F}, P) , where each random variable c_t for $t \in [0, \infty)$ takes values in the state space $S \subseteq \mathbb{R}_{++}$. The probability space and state space do not depend in any way on the transactions of the agents. As described in section 6, the online appendix presents a model in which the transaction cost depends on the amount transferred.

$(\{c_t\}_{t \in [0,u]}, \{(f_t^1, f_t^2)\}_{t \in [0,u]})$. Note that h_u includes information about the cost at time u but does not contain information about the transfers at time u . By convention, h_0 is used to denote the null history $(c_0, \{\})$ at the start of the game. Letting H_t be the set of all histories up to time t , define $H = \bigcup_{t \in [0, \infty)} H_t$. A strategy for agent $i \in \{1, 2\}$ is a map $\pi_i : H \rightarrow \mathbb{R}_+$ that assigns an amount to transfer to each history up to a given time.¹⁷ Let Π_i with generic element π_i represent the set of all possible strategies for agent i .¹⁸

Throughout the analysis, we restrict attention to feasible strategies that satisfy traceability, frictionality, and calculability as defined by Kamada and Rao (2018). The set of agent i 's strategies satisfying these properties is denoted by $\bar{\Pi}_i^C$. Restricting attention to $\bar{\Pi}_i^C$ avoids technical difficulties arising from instantaneous responses and nonmeasurable behavior in continuous time.¹⁹ With this restriction, each strategy profile in $\bar{\Pi}_1^C \times \bar{\Pi}_2^C$ induces a unique path of play that can be represented by a progressively measurable stochastic process, and the expected payoffs to each agent are well defined (Kamada and Rao, 2018).²⁰ A formal definition of the histories and strategy spaces for which we define calculable strategies is provided in Appendix A.1.

It is helpful to introduce some notation to describe the path of play. Let $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ be a profile of calculable strategies. Let $k_u = (\{g_t\}_{t \in [0,u]}, \{(b_t^1, b_t^2)\}_{t \in [0,u]})$ be any history up to an arbitrary time u . Given that $\{c_t\}_{t \in [0,u]} = \{g_t\}_{t \in [0,u]}$, let $\{\phi_t^i(k_u, \{c_\tau\}_{\tau \in (u, \infty)}, \pi)\}_{t \in [0, \infty)}$ be a path of transfers 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 $\pi_i(\{c_v\}_{v \in [0,t]}, \{[\phi_v^1(k_u, \{c_\tau\}_{\tau \in (u, \infty)}, \pi), \phi_v^2(k_u, \{c_\tau\}_{\tau \in (u, \infty)}, \pi)]\}_{v \in [0,t]}) = \phi_t^i(k_u, \{c_\tau\}_{\tau \in (u, \infty)}, \pi)$ for all $t \in [u, \infty)$ and each $i \in \{1, 2\}$. With conditional prob-

¹⁷Simon and Stinchcombe (1989) allow for sequential moves at a single moment of time. The definition of the strategy space here rules out such behavior. This restriction is innocuous in the present setting because of the maximality assumption that we will impose: given a symmetric equilibrium in which sequential transactions by either or both agents occur at some history up to a given instant in time, one can find a symmetric equilibrium generating a weakly higher expected payoff to each agent and requiring only a one-shot simultaneous transaction at any such history.

¹⁸The definition of strategies thus far does not eliminate the possibility of a flow transfer, whereby the amount transferred at each instant of time is zero, but the total amount transferred over an interval of time may be positive. However, current behavior cannot be conditioned on such flows given the set of admissible histories. The restriction to frictional strategies precludes flow exchange.

¹⁹Since the cost process can be specified such that it changes only at a certain countable set of times and the agents can be restricted to moving only at these times, our continuous-time setup encompasses discrete-time models as a special case.

²⁰Another possible definition of strategy spaces might involve the notion of inertia from Bergin and MacLeod (1993), but the maximal equilibrium for which we solve violates inertia, as explained in section 3.3.

ability one, the traceability and frictionality assumptions guarantee the existence and uniqueness of such a path of transfers, and 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 $i \in \{1, 2\}$. In addition, denote by $h(k_u, \{c_t\}_{t \in (u, \infty)}, \pi) = \{c_t, [\phi_t^1(k_u, \{c_\tau\}_{\tau \in (u, \infty)}, \pi), \phi_t^2(k_u, \{c_\tau\}_{\tau \in (u, \infty)}, \pi)]\}_{t \in [0, \infty)}$ the history that results when k_u is the history up to time u , $\{c_t\}_{t \in (u, \infty)}$ is the cost after time u , and strategy profile π is played from time u onwards.

To define expected payoffs and equilibrium concepts, we begin by describing the evolution of the stocks of the goods. Choose any history $h = \{c_t, (f_t^1, f_t^2)\}_{t \in [0, \infty)}$ such that the set $\{t \in [0, \infty) : f_t^i > 0\}$ is finite for each $i \in \{1, 2\}$. Recall that $s_t^{i,j}$ denotes the amount of good j that agent i possesses immediately after time t . The relationship of the stocks of the goods to the transfer paths is specified as $s_t^{i,j} = q - \sum_{\{\tau \in [0, t] : f_\tau^i > 0\}} f_\tau^i$ if $i = j$ and as $s_t^{i,j} = \sum_{\{\tau \in [0, t] : f_\tau^j > 0\}} f_\tau^j$ if $i \neq j$.

We define payoffs as follows. Letting $h = \{c_t, (f_t^1, f_t^2)\}_{t \in [0, \infty)}$ be a history for which the set $\{t \in [0, \infty) : f_t^i > 0\}$ is finite for each $i \in \{1, 2\}$, the realized payoff to agent i at time t is given by:

$$V_t^i(h) = \sum_{\{\tau \in [t, \infty) : f_\tau^{-i} > 0\}} e^{-\rho(\tau-t)} f_\tau^{-i} - \sum_{\{\tau \in [t, \infty) : f_\tau^i > 0\}} e^{-\rho(\tau-t)} c_\tau. \quad (1)$$

Note that each agent i values not its own good i , but the good $-i$ possessed by the other agent $-i$, so that there exist gains to trade, provided that the transaction cost is sufficiently small.

The first term in the preceding expression represents the discrete benefits generated by the good received from the other agent. The second term captures the fixed costs incurred by an agent when making transactions. The restriction of each agent i 's strategy to $\bar{\Pi}_i^C$ implies that, with probability one, the number of transactions made by each agent is countable (Kamada and Rao, 2018).

Under this specification, each agent seemingly consumes the good from the other agent immediately upon receipt. In section 6, however, we argue that the game can be straightforwardly reformulated so that the goods being exchanged are regarded as durable and trade yields a stream of flow benefits.²¹ We also interpret these different formulations in some detail, explaining how they capture different features of the transaction process.

²¹The online appendix formally establishes the equivalence between the original model with discrete benefits and the modified game with flow benefits.

Next choose any strategy profile $\pi = (\pi_1, \pi_2)$ such that $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. Let $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ be any history up to time u for which the set $\{t \in [0, u) : b_t^i > 0\}$ is finite for each $i \in \{1, 2\}$. The expected payoff to agent i at k_u can be expressed as:

$$V_i(k_u, \pi) = \mathbb{E}_{\{c_t\}_{t \in (u, \infty)}} \{V_u^i[h(k_u, \{c_t\}_{t \in (u, \infty)}, \pi)] | \{c_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}\}. \quad (2)$$

The left-hand side of the equation above is the expected payoff at k_u . On the right-hand side, the conditional expectation is taken with respect to $\{c_t\}_{t \in (u, \infty)}$ given that $\{c_t\}_{t \in [0, u]} = \{g_t\}_{t \in [0, u]}$. The expected payoff at k_u is well defined because the restriction of each agent i 's strategy to $\bar{\Pi}_i^C$ ensures that the realized payoff can be uniquely computed with conditional probability one and that the conditional expectation exists.

Since there is no uncertainty regarding past play and past events, we use subgame-perfect equilibrium (SPE) as our equilibrium concept. Formally, a strategy profile (π_1, π_2) with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ is a **subgame-perfect equilibrium** if for any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ up to time u such that the set $\{t \in [0, u) : b_t^i > 0\}$ is finite for each $i \in \{1, 2\}$, the expected payoff to agent $i \in \{1, 2\}$ at k_u satisfies $V_i[k_u, (\pi_i, \pi_{-i})] \geq V_i[k_u, (\pi'_i, \pi_{-i})]$ for any $\pi'_i \in \bar{\Pi}_i^C$.

In some parts of the analysis, we restrict attention to symmetric strategies. A strategy profile (π_1, π_2) with $\pi_i \in \Pi_i$ for $i \in \{1, 2\}$ is **symmetric** if $\pi_1(k_u) = \pi_2(k_u)$ for any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ up to time u such that $b_t^1 = b_t^2$ for $t \in [0, u)$. The symmetry assumption requires that the amount of good 1 transferred by agent 1 at time u is equal to the amount of good 2 transferred by agent 2 at time u given that players 1 and 2 have made symmetric transfers up to time u .²² Given any symmetric strategy profile (π_1, π_2) with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ as well as any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ up to time u such that the set $\{t \in [0, u) : b_t^i > 0\}$ is finite for each $i \in \{1, 2\}$, we define $V(k_u, \pi) = V_1(k_u, \pi) = V_2(k_u, \pi)$.

Let Π^* denote the nonempty set of all symmetric SPE.²³ We next define optimal behavior. A symmetric SPE $\pi \in \Pi^*$ is **maximal** in the class of symmetric SPE if there is no $\pi' \in \Pi^*$ such that $V(h_0, \pi') > V(h_0, \pi)$.

²²Note that this definition does not impose any restriction on behavior following a unilateral deviation from a symmetric path of play.

²³The strategy profile in which each agent never makes a transfer conditional on any history up to an arbitrary time is an element of Π^* ; therefore, the set Π^* is nonempty.

Most of our analysis focuses on maximal equilibrium because we consider a situation in which the two parties have made an informal agreement with each other. Hence, it is reasonable to assume that agents can coordinate their play so as to induce their preferred outcome as long as the incentive constraints of each party are not violated. Restricting attention to such an equilibrium makes it possible to obtain a unique solution, which enables us to derive meaningful comparative statics.

3 Analysis of Model

The solution of the model proceeds as follows. Section 3.1 presents conditions under which the model does and does not have a nondegenerate equilibrium. Section 3.2 characterizes the general form of a maximal equilibrium. Section 3.3 explicitly derives the optimal policy subject to the incentive constraints under the assumption that the cost process follows a geometric Brownian motion.

3.1 Nondegenerate Equilibrium

The following theorem is an impossibility result. If the transaction cost is bounded below by a positive number and the quantity of each good available for trade is fixed, then there is no equilibrium in which an agent receives a positive expected payoff.

Theorem 1. *Assume that $\{c_t\}_{t \in [0, \infty)}$ is an arbitrary cost process and that each random variable c_t for $t \geq 0$ takes values in the state space $S \subseteq \mathbb{R}$ with $\inf(S) > 0$. If the strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ is an SPE, then there is probability one that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) = 0$ for $i \in \{1, 2\}$ and all $t \in [0, \infty)$.*

Note that this result does not depend on assuming symmetric or maximal equilibrium strategies. Furthermore, the proof only relies on an induction argument. That is, for any given S in the statement of the theorem, only a finite hierarchy of knowledge about rationality is needed to establish the theorem. In what follows, we relax the assumption that the transaction cost is bounded below by a positive number.²⁴

The next result identifies a general condition under which the game has a nondegenerate equilibrium.²⁵

²⁴As explained in section 6, the online appendix outlines an alternative approach. We examine a model with uncertainty in the supply of each good, demonstrating that the impossibility result may not hold in such a setting, even with a positive lower bound on the cost.

²⁵In the following statement, a stochastic process is said to be right continuous if its sample path is almost surely continuous from the right everywhere.

Theorem 2. *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 a symmetric SPE π such that, with positive probability, $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0$ for $i \in \{1, 2\}$ and some $t \in [0, \infty)$.*

An example of a cost process with the specified property is geometric Brownian motion. Note that it may be permissible for there to be zero probability of the cost approximating zero. That is, the condition admits some cases in which the probability of a cost less than c ever being reached approaches 0 in the limit as c goes to 0.²⁶

The proof constructs a grim-trigger equilibrium strategy profile in which no further transfers occur following a deviation. On the path of play, each agent makes a transaction at the first time that the cost process reaches a sufficiently low value. Thereafter, each agent makes a transaction whenever the cost process reaches a value no greater than a fraction r of the cost incurred on the preceding transaction. Moreover, the amount transferred by each agent on a given transaction is a fraction r of the previous transfer.

3.2 Optimal Solution

This section describes the basic structure of a maximal symmetric SPE of the game. First, we formally define the concept of a grim-trigger strategy profile, and we justify restricting the analysis to this class of strategies. Next, we define stationary strategy profiles and identify some important properties of stationary maximal symmetric SPE when the cost process is continuous and has the Markov property.

Given a strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \Pi_i$ for $i \in \{1, 2\}$, the strategy π_i is a **grim trigger** if $\pi_i(h_t) = 0$ for every history h and any time t such that h is not consistent with π_1 or π_2 at some time $u < t$. The next result provides justification for restricting attention to SPE in grim-trigger strategies. Given any SPE, there

²⁶In addition, it is straightforward to generalize the preceding theorem to allow the state space for the cost process to include zero. However, this result does not necessarily extend to the case where the state space for the cost process includes negative numbers. If the transaction cost is negative, then it may be possible for an agent to secure an arbitrarily high expected payoff by making a sufficiently large number of transfers. An SPE may not exist in this situation.

exists an SPE in grim-trigger strategies that achieves the same path of play.²⁷ Thus, when characterizing the equilibrium strategies on the path of play, there is no loss of generality from limiting the analysis to grim-trigger strategies.

Proposition 1. *Given an arbitrary SPE π , there exists an SPE π' in grim-trigger strategies such that, with probability one, $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) = \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi')$ for all $t \in [0, \infty)$ and $i \in \{1, 2\}$.*

Note that the proposition above does not require symmetric or maximal equilibrium strategies to be played. The basic idea behind the proof is that each agent can always obtain a continuation value of zero by transferring nothing. If an opponent uses a grim-trigger strategy, then zero is the maximum payoff that can be achieved after any deviation. Thus, the grim-trigger strategy is the severest punishment available.

Hereafter, we restrict attention to symmetric SPE in grim-trigger strategies. In order to characterize the basic properties of a maximal symmetric SPE, it is helpful to focus on stationary strategy profiles. However, we ultimately show that no non-stationary maximal symmetric SPE exists given that the cost process follows a geometric Brownian motion. Hence, there is no loss of generality from requiring stationarity given the assumption about the form of the cost process.

Stationary strategy profiles are formally defined as follows. For any $c \in (0, \infty)$ and $s_1, s_2 \in [0, q]$, let $\tilde{H}_\tau(c, s_1, s_2)$ denote the set consisting of every history $k_\tau = (\{g_v\}_{v \in [0, \tau]}, \{(b_v^1, b_v^2)\}_{v \in [0, \tau]})$ up to time τ such that the set $\{v \in [0, \tau) : b_v^i > 0\}$ is finite for each $i \in \{1, 2\}$, $q - \sum_{\{v \in [0, \tau) : b_v^1 > 0\}} b_v^1 = s_1$, $\sum_{\{v \in [0, \tau) : b_v^2 > 0\}} b_v^2 = s_2$, and $g_\tau = c$. Define $\tilde{H}(c, s_1, s_2) = \bigcup_{\tau \in [0, \infty)} \tilde{H}_\tau(c, s_1, s_2)$.

The strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ is **stationary** if there exists a strategy profile $\pi^\dagger = (\pi_1^\dagger, \pi_2^\dagger)$ with $\pi_i^\dagger \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ satisfying the following two conditions. For any $c \in (0, \infty)$ and $s_1, s_2 \in [0, q]$, $\pi_1^\dagger(k'_u) = \pi_1^\dagger(k''_u)$ and $\pi_2^\dagger(k'_u) = \pi_2^\dagger(k''_u)$ for all $k'_u, k''_u \in \tilde{H}(c, s_1, s_2)$. There is probability one that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) = \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^\dagger)$ for all $t \in [0, \infty)$ and $i \in \{1, 2\}$.

In words, a state is a triple (c, s_1, s_2) , where c is the current value of the cost process, and s_1 and s_2 are the amounts of the goods remaining untransferred. Ac-

²⁷By path of play, we mean the following. Consider any strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. The path of play induced by π is the function that maps each realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ to transfer paths $\{\phi_t^1(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi)\}_{t \in [0, \infty)}$ and $\{\phi_t^2(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi)\}_{t \in [0, \infty)}$ for agents 1 and 2.

cordingly, two histories up to given times are regarded as being in the same state if the current value of the transaction cost process and the total amount of each good previously transferred are the same at these two histories. The set $\tilde{H}(c, s_1, s_2)$ consists of every history belonging to the state (c, s_1, s_2) . The set of stationary strategy profiles includes any strategy profile that requires the agents to transfer the same amount at any two histories belonging to the same state. This set also includes any strategy profile that almost surely induces the same path of play as some strategy profile that satisfies the condition in the previous sentence.²⁸

For example, the maximal symmetric SPE in theorem 4 is a stationary strategy profile. Noting that the agents play grim-trigger strategies, this SPE requires the agents to transfer different amounts at some two histories belonging to the same state. However, this SPE almost surely induces the same path of play as some strategy profile that requires the agents to transfer the same amount at any two histories belonging to the same state. Moreover, the preceding definition is such that zero probability events do not affect the stationarity of a strategy profile.²⁹

The next theorem establishes some basic properties of an optimal solution. The result requires the cost to follow a continuous Markov process.³⁰ It does not rely on the assumption of geometric Brownian motion, which is introduced later. The theorem shows that in any stationary maximal symmetric equilibrium, there is probability one that the incentive constraints are binding and the costs incurred are decreasing.

In order to explain what is meant by binding incentive constraints, we must first define the continuation value after each history. Let $\pi = (\pi_1, \pi_2)$ be any strategy profile such that $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. Let $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ be any history up to time u for which the set $\{t \in [0, u] : b_t^i > 0\}$ is finite for each $i \in \{1, 2\}$. Define $Y_i(k_u, \pi) = V_i(k_u, \pi) - [\pi_{-i}(k_u) - g_u]$. That is, $Y_i(k_u, \pi)$ represents the expected payoff to agent i immediately after any transaction at k_u has been made. If the strategy profile π is also symmetric, then we denote $Y(k_u, \pi) = Y_1(k_u, \pi) = Y_2(k_u, \pi)$,

²⁸Two strategy profiles π^a and π^b are said to almost surely induce the same path of play if there is zero probability that the realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that there exist $t \in [0, \infty)$ and $i \in \{1, 2\}$ for which $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^a) \neq \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^b)$.

²⁹The insensitivity of stationarity to zero probability events is relevant when showing that a non-stationary maximal symmetric SPE does not exist when the cost follows a geometric Brownian motion. Without this insensitivity property, it may be trivial to construct a non-stationary maximal symmetric SPE given a stationary maximal symmetric SPE.

³⁰A stochastic process is said to be Markov if the conditional probability distribution of future values of the process depends only on its current value. A stochastic process is said to be continuous if its sample path is almost surely continuous everywhere.

provided that k_u is such that $b_t^1 = b_t^2$ for $t \in [0, u)$.

In the case where π is symmetric, let p be the probability that the realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that there exists a time t for which $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0$ and $Y[h_t(\{c_\tau\}_{\tau \in [0, \infty)}, \pi), \pi] > c_t$. If $p > 0$, then π is said to have a positive probability of the incentive constraint being **slack** at some transaction. If $p = 0$, then π is said to have probability one of the incentive constraint being **binding** at every transaction. In essence, an incentive constraint is binding if the continuation value after a transaction is equal to the cost paid on the transaction.

The meaning of decreasing costs incurred is straightforward. Let $\pi = (\pi_1, \pi_2)$ be any symmetric strategy profile such that $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. Suppose that there is zero probability of the realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that there exist times t_1 and t_2 with $t_1 < t_2$ for which $\phi_{t_1}^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0$, $\phi_{t_2}^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0$, and $c_{t_2} > c_{t_1}$. Then we say that there is probability one of the cost incurred being decreasing between any two consecutive transactions.

Theorem 3. *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 $S \subseteq \mathbb{R}_{++}$. Then any stationary maximal symmetric SPE in grim-trigger strategies is such that there is probability one of the incentive constraint being binding at every transaction and of the cost incurred being decreasing between any two consecutive transactions.*

The theorem is a consequence of three lemmata, which are stated and proved in the main appendix.³¹ The first of these lemmata exploits the Markov property of the cost process to show that every non-stationary symmetric SPE is weakly Pareto dominated by a stationary symmetric SPE. To understand the intuition for the proof, suppose that the agents are playing a symmetric SPE in which the next transaction may occur at multiple cost levels, each of which is less than the current value of the cost. Since the cost is Markov, another symmetric SPE with a weakly higher continuation payoff can be constructed so that the next transaction can occur at only one such cost level. This construction can be performed iteratively so as to define a stationary symmetric SPE with an expected payoff no lower than the original symmetric SPE. The actual proof is more complicated because the next transaction in a symmetric SPE may occur at a cost level greater than or equal to the current value of the cost.

³¹For each lemma, the online appendix provides an example that illustrates the intuition behind the proof.

The second lemma in the proof states that any stationary symmetric SPE with positive probability of the incentive constraint being slack at some transaction is strictly Pareto dominated by a stationary symmetric SPE in which every incentive constraint is binding. Intuitively, if the incentive constraint were slack at some transaction in a stationary symmetric SPE, then a symmetric SPE with a higher expected payoff could be constructed by raising the amount transferred on that transaction and lowering the amount transferred on successive transactions.

Given binding incentive constraints, the third lemma shows that any stationary symmetric SPE with positive probability of the cost incurred being nondecreasing between successive transactions is strictly Pareto dominated by a stationary symmetric SPE in which the cost incurred is decreasing. The intuition is that if, for example, the costs incurred were increasing between some two consecutive transactions at which the incentive constraints were binding, then a symmetric SPE with a higher expected payoff could be generated by combining the two transfers into a single transaction at the greater of the two costs. Thereby, some of the goods would be transferred earlier, and the probability of such a transfer occurring may rise as well.

Let Π' denote the set consisting of every stationary symmetric SPE $\pi \in \Pi^*$ in grim-trigger strategies for which there is probability one of the incentive constraint being binding at every transaction and there is probability one of the cost incurred being decreasing between any two consecutive transactions. According to theorem 3, any stationary maximal symmetric SPE in grim-trigger strategies belongs to the set Π' . Given any strategy profile $\pi \in \Pi'$ for which there is positive probability of a transaction occurring, let $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^\infty$ be the unique sequence of cost cutoffs and amounts transferred such that $\tilde{c}_k(\pi) \in (0, c_0)$ and $\tilde{f}_k(\pi) \in (0, q)$ for all k , and the following holds. The path of play induced by π is such that, with probability one, the k^{th} transaction is made when the cost reaches $\tilde{c}_k(\pi)$ for the first time, and the amount $\tilde{f}_k(\pi)$ is transferred by each agent at this transaction.

Some further properties of a maximal symmetric equilibrium can be established as corollaries of theorem 3. The result below shows that the sequences of potential amounts transferred and costs incurred converge to zero in a stationary maximal symmetric equilibrium in which a transaction occurs with positive probability.

Corollary 1. *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 $S \subseteq \mathbb{R}_{++}$. Let $\pi \in \Pi^*$ be a stationary maximal symmetric SPE in grim-trigger strategies for which there is*

positive probability of the cost realization $\{c_t\}_{t \in [0, \infty)}$ being such that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0$ for $i \in \{1, 2\}$ and some $t \in [0, \infty)$. Then $\lim_{k \rightarrow \infty} \tilde{c}_k(\pi) = 0$ and $\lim_{k \rightarrow \infty} \tilde{f}_k(\pi) = 0$.

Intuitively, if the amount transferred did not converge to zero, then the stock of each good would be completely exhausted with positive probability. Consequently, if the cost incurred did not converge to zero while the amount transferred did, then there would exist some transaction at which the cost incurred is greater than the amount remaining to be potentially transferred. However, the incentive constraint for that transaction would be violated because the continuation value from the relationship would necessarily be smaller than the transaction cost.

The next result shows that in a stationary maximal symmetric equilibrium in which a transaction occurs with positive probability, the quantities of each good potentially transferred sum up to the total quantity q of each good available.

Corollary 2. *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 $S \subseteq \mathbb{R}_{++}$. Let $\pi \in \Pi^*$ be a stationary maximal symmetric SPE in grim-trigger strategies for which there is positive probability of the cost realization $\{c_t\}_{t \in [0, \infty)}$ being such that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0$ for $i \in \{1, 2\}$ and some $t \in [0, \infty)$. Then $\sum_{k=1}^{\infty} \tilde{f}_k(\pi) = q$.*

This property is a consequence of the maximality condition. For example, consider an equilibrium in which a transaction occurs with positive probability but the sum of the amounts of each good potentially transferred is at most $\tilde{q} < q$. A Pareto superior equilibrium can be constructed by requiring the agents to transfer the additional amount $q - \tilde{q}$ at the first transaction.

The preceding corollary has a further implication. Let π be a stationary maximal symmetric SPE. The set consisting of each transaction time when the agents play π and the realized cost is $\{c_\tau\}_{\tau \in [0, \infty)}$ is denoted by $J(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = \{v \in [0, \infty) : \phi_v^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0\}$. Some continuous Markov cost processes have the property that the probability of the cost ever reaching the value c is one for all $c \in (0, c_0)$.³² This condition is necessary and sufficient for $\sum_{v \in J(\{c_\tau\}_{\tau \in [0, \infty)}, \pi)} \phi_v^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi)$ to be equal to q almost surely. That is, the entire stock of each good is almost surely

³²A geometric Brownian motion with drift μ and volatility σ satisfies this condition for $\mu \leq \sigma^2/2$ but not for $\mu > \sigma^2/2$.

transferred in the limit as time goes to infinity if and only if there is probability one of the transaction cost ever approximating zero.

3.3 Explicit Formula

This section derives the optimal strategies given the form of the cost process. Specifically, we assume that the cost process $\{c_t\}_{t \in [0, \infty)}$ follows a geometric Brownian motion $dc_t = \mu c_t dt + \sigma c_t dz_t$ with arbitrary drift μ and positive volatility σ .³³ It is useful to define the parameter

$$\beta = \frac{1}{2} - \frac{\mu}{\sigma^2} - \sqrt{\left(\frac{\mu}{\sigma^2} - \frac{1}{2}\right)^2 + 2\frac{\rho}{\sigma^2}}.$$

We now provide a closed-form solution. The subsequent sections analyze the solution in greater detail by calculating comparative statics and characterizing welfare properties.

Theorem 4. *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 there exists a stationary maximal symmetric SPE in grim-trigger strategies. Moreover, any maximal symmetric SPE in grim-trigger strategies is characterized by a sequence $\{c_k^*, f_k^*\}_{k=1}^\infty$ satisfying*

$$c_k^* = \left(\frac{q}{1-\beta}\right) \left(\frac{\beta}{\beta-1}\right)^{k-\beta} \quad \text{and} \quad f_k^* = \left(\frac{q}{1-\beta}\right) \left(\frac{\beta}{\beta-1}\right)^{k-1}$$

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.*

The preceding result is derived by solving a constrained optimization problem. The three lemmata in the proof of theorem 3 are used in order to express the objective function and constraints in a tractable form. Specifically, the first of these lemmata allows attention to be restricted to stationary symmetric SPE without lowering the expected payoffs that can be attained. In addition, the second and third lemmata imply that only stationary symmetric SPE with binding incentive constraints and decreasing costs incurred need to be considered when solving the maximization

³³By specifying that the value of the cost process evolves according to a geometric Brownian motion, we are assuming that the growth in the transaction cost follows a Brownian motion with drift. This is an important baseline case because a Brownian motion with drift is the only continuous Lévy process. As noted by Sannikov and Skrzypacz (2010), a Lévy process has independent and identically distributed increments. It is the continuous-time analog of a random walk.

problem. The incentive compatibility conditions provide one set of constraints, which require that the cost incurred on each transaction equals the continuation value after that transaction. Another constraint is feasibility, whereby the sum of the amounts transferred cannot exceed the total amount of each good available. Because the incentive constraints are binding, the objective function is simply the expected payoff to each agent from the transfer at the first transaction.

The solution to this maximization problem delivers an expression for the uniquely optimal path of play in any stationary maximal symmetric SPE. Furthermore, we show that there is no non-stationary maximal symmetric SPE given that the cost process follows a geometric Brownian motion. Therefore, the path of play specified in the statement of the theorem applies to any maximal symmetric SPE.

The intuition behind the non-existence of a non-stationary maximal symmetric SPE is illustrated by an example in the online appendix. If there were to exist a non-stationary maximal symmetric SPE, then such a strategy profile could be used to find two stationary maximal symmetric SPE that do not almost surely induce the same path of play. However, this would contradict the result that the path of play in a stationary maximal symmetric SPE is unique up to probability zero events.

As justified above, attention is restricted to grim-trigger strategies. Although there exist maximal symmetric SPE not in grim-trigger strategies, any such SPE induces the path of play described in the preceding theorem.³⁴

The optimal solution in theorem 4 has a proportional structure, in which the cost incurred decreases by a factor of $(1 - \beta)/\beta$ from the current transaction to the next. The amount transferred is likewise becoming smaller.

Corollary 3. $f_n^* < f_m^*$ if $n > m$.

Intuitively, the incentive to deviate is falling over successive transactions because the prescribed cost payment is decreasing, making cooperation sustainable with smaller future transfers. The decline in the transaction size is consistent with the example of land transferral in the introduction.

Note that the strategy profile in theorem 4 violates the inertia condition of Bergin and MacLeod (1993), which would require in our context that for any history up to an arbitrary time t , there exists $\epsilon > 0$ such that no transactions occur in the time

³⁴To understand how to construct a maximal symmetric SPE not in grim-trigger strategies, see the online appendix for an example.

interval $(t, t + \epsilon)$. To see this, fix a history up to time t such that neither player has deviated in the past, each player has transacted k times, and the cost is currently greater than c_{k+1}^* . Then for any $\epsilon > 0$, there is positive probability that the cost reaches c_{k+1}^* during the time interval $(t, t + \epsilon)$, in which case a transaction takes place under the maximal equilibrium.³⁵

4 Comparative Statics

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 path of play in a maximal symmetric SPE is unique. Using the closed-form expression in theorem 4, a number of comparative statics can be obtained. In particular, we describe how the drift and volatility of the cost process affect the sequence of costs incurred and amounts transferred. The role of the discount rate is also investigated.

The corollary below describes how the size of each transfer changes with the discount rate and the parameters of the cost process.

Corollary 4. *If $k < 1 + |\beta|$, then f_k^* is decreasing in μ and ρ but increasing in σ . If $k > 1 + |\beta|$, then f_k^* is increasing in μ and ρ but decreasing in σ .*

This is an intuitive result.³⁶ If the drift μ of the cost process decreases, then the cost is more likely to fall enough in the near future for the agents to make another transaction. The greater proximity of a future transaction raises the continuation value of the relationship and relaxes the incentive constraints for the problem. Likewise, a lower discount rate ρ increases the continuation value, thereby weakening the incentive constraints. Thus, agents can make larger transfers at early stages as stated in the first part of the corollary. However, if more of each good is transferred at earlier stages, then less of each good is remaining at later stages. Therefore, the transfers at later stages must be smaller as stated in the second part of the corollary. If the volatility σ of the cost process increases, then both extremely high and low

³⁵Bergin and MacLeod (1993) also propose a weaker condition involving the completion of the set of inertia strategies. However, this approach does not apply here either. As shown in the online appendix, the maximal equilibrium of our model cannot be represented as the limit of a Cauchy sequence of strategy profiles satisfying inertia.

³⁶Note that the ratio of the amount transferred by an agent on each transaction to the stock of each good remaining before the transaction is $(1 - \beta)^{-1}$. Because $\beta < 0$ is decreasing in μ and ρ as well as increasing in σ , the fraction of the remaining stock transferred on each transaction decreases with μ and ρ but increases with σ .

realizations of the cost process become more likely. Because the solution has a cutoff form, the favorable impact of low cost realizations dominates the adverse impact of high cost realizations. This option-value argument suggests that a high volatility σ has a similar effect on the solution as a low drift μ .

The next corollary characterizes the impact of the cost parameters and discount rate on the sequence of transaction costs paid.

Corollary 5. *There exists an increasing function $n(|\beta|)$ with $\lim_{|\beta| \rightarrow 0} n(|\beta|) = 0$ and $\lim_{|\beta| \rightarrow \infty} n(|\beta|) = \infty$ such that the following holds. If $k < n(|\beta|)$, then c_k^* is decreasing in μ and ρ but increasing in σ . If $k > n(|\beta|)$, then c_k^* is increasing in μ and ρ but decreasing in σ .*

This result exemplifies a tradeoff between two effects. Suppose that the drift μ decreases. On the one hand, low values of the transaction cost become more likely, which encourages the agents to wait longer before transacting. Thereby, the cost payment can be reduced. This suggests that the cost cutoff c_k^* should decrease. On the other hand, the continuation value of the relationship rises, which enables the agents to transact at a higher cost without violating the incentive constraints. Thereby, the waiting time can be reduced. This suggests that the cost cutoff c_k^* should increase. As has already been noted, a high volatility σ has an impact similar to a low drift μ . A smaller discount rate ρ also exerts two effects. More patient agents are willing to wait longer for a low realization of the transaction cost, but greater patience increases the continuation value of the relationship, enabling higher cost payments to be sustained in equilibrium.

The dominant effect depends on the parameters of the model. Intuitively, larger absolute values of β reflect greater frictions. If $|\beta|$ is low, then the first effect is important, in which case the cost thresholds are increasing in μ and ρ but decreasing in σ . If $|\beta|$ is high, then the second effect is relatively strong. Hence, the cost thresholds early in the relationship are decreasing in μ and ρ but increasing in σ . However, a decrease in μ , an increase in σ , or a decrease in ρ implies that the supplies of untransferred goods are depleted earlier in the relationship. Given a small remaining stock and a low continuation value, the cost cutoffs tend to decline at later stages.

An implication of the foregoing result is that each cost threshold is nonmonotonic in the drift parameter and possibly also the volatility and discount rate. As the drift μ increases from $-\infty$ to ∞ , the absolute value of β increases from 0 to ∞ . Hence, each

cost cutoff c_k^* first rises with μ and then falls with μ . As the volatility σ decreases from ∞ to 0, the absolute value of β increases from 0 to ∞ when $\mu \geq 0$ and from 0 to $-\rho/\mu$ when $\mu < 0$. Therefore, the threshold c_k^* exhibits an inverted U-shaped relationship with σ if μ is nonnegative. Otherwise, c_k^* can be decreasing or nonmonotonic in σ , depending on specific parameter values. An increase in the discount rate ρ from 0 to ∞ causes β to rise in absolute value from 0 to ∞ for $\mu \leq \sigma^2/2$ and from $2\mu/\sigma^2 - 1$ to ∞ for $\mu > \sigma^2/2$. It follows that the relationship between c_k^* and ρ displays an inverted U-shape if $\mu \leq \sigma^2/2$. Otherwise, c_k^* decreases with ρ for some parameter values, but the effect is nonmonotonic in other cases.

This nonmonotonicity is on display in figure 1. The first cost threshold is plotted against the drift parameter μ when the volatility is fixed at a given level σ and the initial stock of each good is 1. Different values of the discount rate ρ are considered in comparison to the limiting case as ρ approaches 0. As explained above, the cutoff c_1^* is first rising and then falling in μ . In addition, the respective curves depicting c_1^* for $\rho = 0.01$ and $\rho = 0.1$ intersect. Accordingly, c_1^* is increasing in ρ when μ is low and so $|\beta|$ is small, whereas the opposite holds for high values of μ , in which case $|\beta|$ is large. This property generates the potentially nonmonotonic relationship between the cost threshold and the discount rate mentioned above. Furthermore, c_1^* essentially becomes insensitive to ρ when μ is high relative to σ .³⁷ The intuition is that as μ increases, the upward drift of the cost becomes a bigger friction in the exchange process than the discounting of future transactions. Hence, the discount rate plays less of a role in determining the optimal trading policy.

The following corollary examines how the transfers made and costs incurred behave in the limits as the discount rate goes to zero and to infinity.

Corollary 6. *If $\mu \leq \sigma^2/2$, then $\lim_{\rho \rightarrow 0} f_1^* = q$, $\lim_{\rho \rightarrow 0} f_k^* = 0$ for $k \geq 2$, and $\lim_{\rho \rightarrow 0} c_k^* = 0$ for all k . If $\mu > \sigma^2/2$, then $\lim_{\rho \rightarrow 0} f_1^* < q$, $\lim_{\rho \rightarrow 0} f_k^* > 0$ for $k \geq 2$, and $\lim_{\rho \rightarrow 0} c_k^* > 0$ for all k . In addition, $\lim_{\rho \rightarrow \infty} f_k^* = \lim_{\rho \rightarrow \infty} c_k^* = 0$ for all k .*

These results follow immediately from theorem 4. First, consider the case where $\mu \leq \sigma^2/2$, meaning that the drift of the cost process is low relative to the volatility.

³⁷Given any positive integer k , let $c_k^*(\rho, \mu, \sigma)$ denote the cost incurred on the k^{th} transaction in the maximal symmetric SPE when the discount rate is ρ and the cost process follows a geometric Brownian motion with drift μ and volatility σ . It is straightforward to show that for any $\rho', \rho'' > 0$ along with any $\sigma > 0$, the ratio $c_k^*(\rho', \mu, \sigma)/c_k^*(\rho'', \mu, \sigma)$ converges to one in the limit as μ approaches ∞ .




Figure 1: c_1^* vs. μ for fixed $\sigma > 0$ and $q = 1$

As the agents become infinitely patient, it is optimal to wait for an extremely low cost realization before making a transfer. Furthermore, a large initial transfer can be supported in equilibrium, because the incentive constraints are weak for low values of the transaction cost and so a high continuation value is not needed to prevent deviation. Next, suppose that $\mu > \sigma^2/2$, in which case the drift is high compared with the volatility. The transaction cost is trending upward, and a fortuitously low realization of the cost is improbable. Because a favorable trading opportunity is unlikely, it is not optimal for even infinitely patient agents to wait an extremely long time before transacting. Moreover, since transaction costs are still significant in the limit, a substantial continuation value is needed to ensure that cooperation is incentive compatible, which restricts the size of the initial transfer.

Finally, the corollary states that as the agents become infinitely impatient, it becomes impossible to induce them to incur any strictly positive transaction cost. In this situation, the continuation value from the relationship is low and so the incentive to deviate is high. Therefore, the size of each transfer must also be small, in order to ensure that the continuation value is sufficiently high to sustain cooperation.

Several of the comparative statistics in this section can be interpreted in terms of the applications described in the introduction. For instance, consider the example of prisoner exchange between two hostile parties, where the transaction cost represents

the level of tensions between the opponents. The solution to the model suggests a prisoner exchange protocol in which the prevailing level of tensions and the quantity of prisoners traded are smaller at the next transfer than at the current transfer. If the level of tensions decreases more quickly or becomes more volatile, then larger transfers of prisoners would take place at earlier stages of the relationship, and smaller transfers would happen at later stages. These changes may cause either a rise or a fall in the level of tensions at which earlier transfers occur, but tensions at later transfers would decrease.³⁸ An interesting insight to arise from the analysis is that greater volatility in the level of tensions often has the same effect on the transaction scheme as a more rapid decrease in the level of tensions.

5 Welfare Properties

Section 5.1 describes the efficient trading policy in the absence of incentive constraints. Section 5.2 examines whether the efficient outcome can be approximated as agents become infinitely patient. A number of comparative statics regarding welfare levels are presented.

5.1 Efficient Outcome

This section characterizes the efficient outcome of the model. It thereby establishes a benchmark for the welfare analysis. When the cost process follows a geometric Brownian motion, a closed-form expression can be obtained.

We begin with a formal definition of efficiency. A symmetric strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ is **efficient** if there is no symmetric strategy profile $\pi' = (\pi'_1, \pi'_2)$ with $\pi'_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ such that $V(h_0, \pi') > V(h_0, \pi)$. Efficiency is defined similarly to maximality, but non-equilibrium strategy profiles must also be considered when identifying an optimum.

The result below shows that efficiency requires there to be probability one that at most one transaction occurs. This is in contrast to the path of play in any nondegenerate equilibrium, in which there is no bound on the number of transactions that can occur with positive probability.

³⁸Note that the cutoff $n(|\beta|)$ in corollary 5 satisfies $\lim_{|\beta| \rightarrow 0} n(|\beta|) = 0$ and $\lim_{|\beta| \rightarrow \infty} n(|\beta|) = \infty$, where the parameter β is such that $\lim_{\mu \rightarrow -\infty} |\beta| = 0$ and $\lim_{\mu \rightarrow \infty} |\beta| = \infty$.

Proposition 2. *Assume that $\{c_t\}_{t \in [0, \infty)}$ is an arbitrary cost process and that each random variable c_t for $t \geq 0$ takes values in the state space $S \subseteq \mathbb{R}_{++}$. Let π be any efficient symmetric strategy profile. Then there is probability one of the realization of the cost process $\{c_t\}_{t \in [0, \infty)}$ being such that either $\phi_i^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) = 0$ for $i \in \{1, 2\}$ and all $t \in [0, \infty)$ or $\phi_i^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) = q$ for $i \in \{1, 2\}$ and some $t \in [0, \infty)$.*

The explanation is straightforward. Suppose that the agents reach a history at which a transaction occurs but some of each good is remaining after the transaction. Then the agents can increase the expected payoffs at that history by instead transferring all of the goods. Hence, efficiency requires there to be zero probability that goods remain untransferred after a transaction.

Assuming that the transaction cost evolves according to a geometric Brownian motion, the result below solves for the path of play induced by any efficient symmetric strategy profile.

Proposition 3. *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 there exists an efficient symmetric strategy profile. Moreover, any efficient symmetric strategy profile is characterized by a cost cutoff \bar{c} given by*

$$\bar{c} = \left(\frac{\beta}{\beta - 1} \right) q$$

such that, with probability one, a transaction is made when the cost reaches \bar{c} for the first time, and the amount q is transferred by each agent at this transaction.

Note that there is a critical value of the cost such that the agents almost surely transfer all of their goods at the first time the cost reaches that value. The online appendix documents the following comparative statics. First, the efficient cost cutoff \bar{c} is increasing in the drift parameter and discount rate but decreasing in the volatility of the cost process. Second, this threshold approaches the initial endowment of each good in the limit as the discount rate goes to infinity. Finally, it may or may not converge to zero as agents become infinitely patient, depending on the parameters of the cost process.

Comparing proposition 3 with theorem 4, the efficient cost threshold is greater than the first cost cutoff in the maximal symmetric SPE.

Corollary 7. $\bar{c} > c_1^*$.

Intuitively, an efficient path of play requires the agents to transfer all of the goods at the first time cost reaches the threshold \bar{c} . However, such a policy is not incentive compatible because the transaction cost incurred, which is positive, is less than the quantity remaining untransferred, which is zero. Since it is profitable for each agent to deviate by transferring nothing, this is not an equilibrium. In order to satisfy the incentive constraint in a maximal equilibrium, the amount remaining after the first transaction is increased, and the cost at which the first transaction occurs is decreased. The online appendix presents a theorem that generalizes this remark to the case of a continuous Markov cost process. We relax the specific assumption of geometric Brownian motion, requiring instead that the cost obeys a scaling rule according to which the incremental change in the cost is proportional to the current value of the cost.

5.2 Approximating Efficiency

This section evaluates the welfare implications of the model. We begin by presenting formulae for the first- and second-best expected payoffs, which are then used to derive comparative statics. Although the analysis in this section is confined to the case in which the transaction cost follows a geometric Brownian motion, the online appendix relaxes the restriction on the functional form.

Consider the closed-form solutions for the maximal symmetric equilibrium in theorem 4 and for the efficient symmetric strategy profile in proposition 3. We use the term second-best to refer to the expected payoff of each agent when the path of play specified in theorem 4 is followed. In this case, the behavior of the players is limited by incentive constraints. We use the term first-best to refer to the expected payoff of each agent when the path of play specified in proposition 3 is followed. The proposition below provides closed-form expressions for the expected payoffs to the agents in the first- and second-best cases.

Proposition 4. *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 the first-best expected payoff M^{fb} and the second-best expected payoff M^{sb} are given by*

$$M^{fb} = \theta q^{1-\beta} c_0^\beta \quad \text{and} \quad M^{sb} = \theta^{1-\beta} q^{1-\beta} c_0^\beta,$$

where

$$\theta = (-\beta)^{-\beta}(1 - \beta)^{-(1-\beta)}.$$

The aforementioned formulae are used to derive a number of corollaries. The result below provides comparative statics for the effects of the parameters on the expected payoffs in the first- and second-best cases.

Corollary 8. *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 M^{fb} and M^{sb} are both decreasing in μ and ρ but increasing in σ .*

The intuition is straightforward. If the drift μ decreases or the volatility σ increases, then the probability of a low realization of the transaction cost rises. Hence, an opportunity for profitable trade becomes more likely, and so the expected payoff to each agent increases in both the first- and second-best cases. In addition, more patient agents receive higher expected payoffs because they assign greater value to future benefits.

The next corollary confirms that the second-best expected payoff converges to the first-best expected payoff in the limit as the discount rate approaches zero, provided that the drift of the cost process is not excessively high relative to the volatility.

Corollary 9. *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 M^{fb} and M^{sb} both converge to q in the limit as ρ goes to zero. If $\mu > \sigma^2/2$, then M^{fb} and M^{sb} converge neither to q nor to each other in the limit as ρ goes to zero.*

On the one hand, if trading frictions are low in the sense that the drift of the cost process is sufficiently small relative to the volatility, then the first-best outcome can be approximated as the agents become infinitely patient. On the other hand, if the drift is high and the volatility is low, then the trading environment tends to deteriorate over time, and there is uncertainty about whether a future transaction will take place. In this case, the second-best solution fails to converge to the first-best outcome.

The condition $\mu \leq \sigma^2/2$ is critical to the welfare analysis because it determines whether there is probability one of the transaction cost ever approximating zero. The probability of the cost c ever being reached is one for all $c \in (0, c_0)$ if and only if $\mu \leq \sigma^2/2$. For $\mu \leq \sigma^2/2$, the cost incurred on the first transaction converges to zero in both

the first- and second-best solutions as agents become infinitely patient. In the limiting case, the amount of each good transferred on the first transaction in the second-best solution approximates the entire amount, just as in the first-best solution.³⁹ For $\mu > \sigma^2/2$, the respective costs incurred on the first transaction converge to different positive real numbers in the first- and second-best cases. In addition, the incentive constraint requires the continuation value after the first transaction to be no less than the cost incurred on the first transaction. Hence, the amount transferred by each agent on the first transaction does not converge to the total stock of each good in the second-best solution.

The following corollary investigates the situation in which the first- and second-best outcomes differ when agents are infinitely patient. It examines how the drift and volatility of the cost process affect the relative values of the first- and second-best expected payoffs.

Corollary 10. *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 $\mu > \sigma^2/2$, then in the limit as ρ approaches zero, M^{sb}/M^{fb} is decreasing in μ and increasing in σ .*

Intuitively, a lower drift or a higher volatility results in a more favorable trading environment in which future transactions are more likely to occur. Consequently, the continuation value from the relationship is higher, and larger transfers can be sustained in equilibrium. The result above shows that the second-best expected payoff increases as a fraction of the first-best expected payoff.

To illustrate the welfare properties of the model, figure 2 displays how the expected payoffs vary with μ for a given value of σ . The graph fixes the initial value of the cost at $3/2$ and the original amount of each good at 1. In the limit as the discount rate approaches zero, M^{fb} and M^{sb} coincide with each other if $\mu \leq \sigma^2/2$. If $\mu > \sigma^2/2$, then M^{sb} becomes an increasingly small fraction of M^{fb} as μ increases relative to σ . With strictly positive discounting, M^{fb} is greater than M^{sb} . Finally, the diagram suggests that the discount rate has little effect on welfare when μ is high compared to σ .⁴⁰ Intuitively, when the transaction cost is rapidly trending upward with little

³⁹An implication of the preceding corollary for $\mu \leq \sigma^2/2$ is that given any profile of expected payoffs $(V_1, V_2) \in (0, M^{fb})^2$, there exists $\bar{\rho} > 0$ such that for all $\rho < \bar{\rho}$, (V_1, V_2) can be attained in an SPE. This follows because the incentive constraints on the first transaction are binding with probability one in the maximal equilibrium, and so any equilibrium payoff profile $(V_1, V_2) \in (0, M^{sb})^2$ can be achieved by appropriately lowering the amounts transferred on the first transaction.

⁴⁰More precisely, let $M^{fb}(\rho, \mu, \sigma)$ and $M^{sb}(\rho, \mu, \sigma)$ respectively denote the first- and second-best




Figure 2: M^{fb} and M^{sb} vs. μ for fixed $\sigma > 0$ and $(c_0, q) = (3/2, 1)$

chance of dropping, then a transaction, if any, is likely to occur early in the game. This tends to lessen the impact of discounting on the expected payoffs.

The online appendix proves that the following general condition on the cost process suffices for all the potential gains from trade to be realized in the limit as the discount rate approaches zero. Consider a right-continuous cost process. Suppose that there exists a constant $r < 1$ for which the conditional probability of the cost becoming a fraction r of its current value during a time interval of length v approaches one in the limit as v goes to infinity. Then one can find a discount rate and a symmetric equilibrium such that the expected payoff to each agent is arbitrarily close to the total stock of each good. Note that a geometric Brownian motion satisfies the aforementioned condition if and only if its drift and volatility parameters are such that $\mu \leq \sigma^2/2$. Hence, the first part of corollary 9 is simply an application of the relevant theorem in the online appendix.

6 Discussions and Extensions

Alternative specification of payoffs

expected payoffs when the discount rate is ρ and the cost process follows a geometric Brownian motion with drift μ and volatility σ . It is straightforward to show that for any $\rho', \rho'' > 0$ along with any $\sigma > 0$, the ratios $M^{fb}(\rho', \mu, \sigma)/M^{fb}(\rho'', \mu, \sigma)$ and $M^{sb}(\rho', \mu, \sigma)/M^{sb}(\rho'', \mu, \sigma)$ both converge to one in the limit as μ approaches ∞ .

As mentioned in section 2, the payoff function in equation (1) has an alternative formulation. The online appendix considers a modified model in which the transferred good provides the recipient with a stream of flow payoffs. In particular, if an agent transfers the amount f_t^i of her good at time t , then she incurs the fixed cost $C_t = c_t/\rho$ at time t , and her opponent enjoys a flow benefit of f_t^i starting from time t . We show that any SPE of the original game is also an SPE of the model with flow payoffs, and vice versa. That is, the equilibrium set remains unchanged, and a maximal equilibrium in the original game is an optimal solution to the modified model. This idea is formalized in the online appendix.

These different specifications of the game can be related to some of the examples from the introduction, such as prisoner exchange between countries. In the modified framework, goods are treated as durable, and a country enjoys a flow payoff from having possession of returned prisoners. This flow payoff might represent a stream of emotional benefits to individuals in each country from being reunited with family members who were being held captive by the other country. Alternatively, each country might receive a stream of economic benefits as former captives reenter the labor force. The original setup effectively requires the consumption of goods instantly upon delivery. In this case, a country obtains a single discrete payoff when prisoners from its side are released. For example, the return of prisoners might increase the popularity of political leaders, or individuals might experience a sudden feeling of happiness upon meeting their relatives again.

Variations of basic model

The online appendix describes two variations of the basic model. First, we allow for growth or volatility in the supply of each good available for trade. We show that even if the transaction cost is bounded away from zero, positive transfers may be supported in equilibrium, provided that there is some variability in the supply of each good. In particular, grim-trigger strategies are used to generate an equilibrium path in which the agents make a transaction whenever the remaining stock of each good reaches a given threshold. Second, we study a setting in which the cost incurred at a given transaction depends on the amount of each good transferred. Even in such a situation, positive transfers may be supported in equilibrium. On the path of play that we specify, the number of transactions is potentially infinite, and the transfer size and transaction cost are gradually decreasing.

Additional explanation of transaction cost

Transactions may involve various costs related to transportation, communication, and monitoring. In the example of prisoner exchange between countries, the transaction cost may be interpreted in political terms as a loss of popular support for the government. This political cost may be increasing in the number of prisoners released and may fluctuate depending on the level of public approval for the administration. Under a democratic system, if public approval is closer to the threshold for electoral victory, then the political cost of releasing prisoners may be higher as the election outcome is more likely to be adversely affected. In the case where the cost follows a geometric Brownian motion, the drift and volatility parameters may reflect the pattern of public approval ratings for the government over time. These parameters may also be related to changes in economic factors that influence public opinion, such as the unemployment rate, wage growth, stock market returns, and inflation. Such factors could be correlated across countries to some extent, and the common transaction cost, which might be linked to indices of global economic conditions, serves as a simple way to capture this correlation.

7 Conclusion

This paper studied the exchange of divisible goods between two agents facing a stochastic transaction cost. We developed a model of trade that applies techniques from the literature on investment under uncertainty to the theory of repeated games. In our setup, the first-best policy requires a single transaction, whereas the second-best solution involves a decreasing sequence of transfers of each good. Several comparative statics were presented that illustrate how the properties of the cost process affect optimal behavior. We also examined the convergence of equilibria to the efficient outcome as discounting disappears. Our framework can be applied to real-world situations involving the gradual exchange of trade secrets, captured prisoners, or land claims.

Continuous-time modeling made the analysis of a stochastic game tractable. In particular, a closed-form solution can be obtained given that the cost process evolves according to a geometric Brownian motion, and we used the resulting formula to study comparative statics. Interestingly, we found that the waiting time until the first transaction is nonmonotonic in the drift and maybe also the volatility of the

stochastic process. Moreover, the waiting time is longer than in the first-best case, where the entire stock of each good is transferred at once.

The analysis suggests a number of possibilities for further research. For example, the framework could be extended to incorporate asymmetries between players in the key elements of the model, such as the payoff functions, initial endowments, cost processes, and equilibrium strategies. Allowing for mixed strategies might increase the gains from trade in equilibrium. It would also be interesting to study incentives to cooperate when players have imperfect information about past actions. Although these questions may be beyond the scope of the current paper, we hope our analysis serves as a basis for tackling them.

References

- ABDULKADIROGLU, A., AND K. BAGWELL (2012): “The Optimal Chips Mechanism in a Model of Favors,” Mimeo, Stanford University.
- ADMATI, A. R., AND M. PERRY (1991): “Joint Projects without Commitment,” *Review of Economic Studies*, 58(2), 259–276.
- ASSOCIATED PRESS (1953): “Prisoner Timetable Proposed,” *New York Times*, p. 2, April 18.
- BERGIN, J., AND W. B. MACLEOD (1993): “Continuous Time Repeated Games,” *International Economic Review*, 34(1), 1993.
- BERTOLA, G., AND R. CABALLERO (1994): “Irreversibility and Aggregate Investment,” *Review of Economic Studies*, 61(2), 223–246.
- CARUANA, G., AND L. EINAV (2008): “A Theory of Endogenous Commitment,” *Review of Economic Studies*, 75(1), 99–116.
- COMPTE, O., AND P. JEHIEL (1995): “On the Role of Arbitration in Negotiations,” Mimeo, C.E.R.A.S.-E.N.P.C.
- (2003): “Voluntary Contributions to a Joint Project with Asymmetric Agents,” *Journal of Economic Theory*, 112(2), 334–342.
- (2004): “Gradualism in Bargaining and Contribution Games,” *Review of Economic Studies*, 71(4), 975–1000.

- DIXIT, A. (1991): “Irreversible Investment with Price Ceilings,” *Journal of Political Economy*, 99(3), 541–557.
- DUTTA, P. K. (1995): “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, 66(1), 1–32.
- FUDENBERG, D., AND E. MASKIN (1986): “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54(3), 533–554.
- FUDENBERG, D., AND Y. YAMAMOTO (2011): “The Folk Theorem for Irreducible Stochastic Games with Imperfect Public Monitoring,” *Journal of Economic Theory*, 146(4), 1664–1683.
- GALE, D. (1995): “Dynamic Coordination Games,” *Economic Theory*, 5(1), 1–18.
- (2001): “Monotone Games with Positive Spillovers,” *Games and Economic Behavior*, 37(2), 295–320.
- GUERON, Y. (2015): “Failure of Gradualism under Imperfect Monitoring,” *Journal of Economic Theory*, 157, 128–145.
- HABERMAN, C. (1991): “Israel Releases 51 Arab Prisoners; Hopes Rise for Lebanon Hostages,” *New York Times*, p. A1, September 12.
- HAUSER, C., AND H. A. HOPENHYAN (2005): “Trading Favors: Optimal Exchange and Forgiveness,” Mimeo, University of California, Los Angeles.
- HE, Z. (2009): “Optimal Executive Compensation when Firm Size Follows Geometric Brownian Motion,” *Review of Financial Studies*, 22(2), 859–892.
- HOLMSTROM, B., AND P. MILGROM (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica*, 55(2), 303–328.
- HÖRNER, J., T. SUGAYA, S. TAKAHASHI, AND N. VIELLE (2011): “Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem,” *Econometrica*, 79(4), 1277–1318.
- JOHNSON, K., D. SIMCHI-LEVI, AND P. SUN (2014): “Analyzing Scrip Systems,” *Operations Research*, 62(3), 524–534.

- KAMADA, Y., AND N. RAO (2018): “Strategies in Stochastic Continuous-Time Games,” Mimeo, University of California, Berkeley.
- LEMOYNE, J. (1984): “Captives Released in Salvador War,” *New York Times*, p. 3, October 6.
- LIPMAN, B. L., AND R. WANG (2000): “Switching Costs in Frequently Repeated Games,” *Journal of Economic Theory*, 93(2), 149–190.
- (2009): “Switching Costs in Infinitely Repeated Games,” *Games and Economic Behavior*, 66(1), 292–314.
- MARX, L. M., AND S. A. MATTHEWS (2000): “Dynamic Voluntary Contribution to a Public Project,” *Review of Economic Studies*, 67(2), 327–358.
- MCDONALD, R., AND D. SIEGEL (1986): “The Value of Waiting to Invest,” *Quarterly Journal of Economics*, 101(4), 707–727.
- MERTON, R. C. (1971): “Optimum Consumption and Portfolio Rules in a Continuous-Time Model,” *Journal of Economic Theory*, 3(4), 373–413.
- MOBIUS, M. (2001): “Trading Favors,” Mimeo, Harvard University.
- PITCHFORD, R., AND C. M. SNYDER (2004): “A Solution to the Hold-Up Problem Involving Gradual Investment,” *Journal of Economic Theory*, 114(1), 88–103.
- SANNIKOV, Y. (2008): “A Continuous-Time Version of the Principal-Agent Problem,” *Review of Economic Studies*, 75(3), 957–984.
- SANNIKOV, Y., AND A. SKRZYPACZ (2010): “The Role of Information in Repeated Games with Frequent Actions,” *Econometrica*, 78(3), 847–882.
- SECOND TREATY OF INDIAN SPRINGS (1825): “Treaty with the Creeks at the Indian Springs,” *American State Papers: Indian Affairs*, 2(222), 563–584.
- SIMON, L. K., AND M. B. STINCHCOMBE (1989): “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, 57(5), 1171–1214.
- WOLITZKY, A. (2015): “Communication with Tokens in Repeated Games on Networks,” *Theoretical Economics*, 10(1), 67–101.

A Appendix

A.1 Formal Definitions of Histories and Strategy Spaces

This section provides formal definitions of histories and strategy spaces. In doing so, we follow the notation and terminology in Kamada and Rao (2018).

The action z stands for “no move.” In the present paper’s model, we replace the action of transferring the amount 0 with action z , so that the transaction amount f_t^i is defined over $\{z\} \cup (0, q]$. 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 $(\{c_\tau\}_{\tau \in [0, t]}, \{f_\tau^1, f_\tau^2\}_{\tau \in [0, t]})$, where $\{f_\tau^i\}_{\tau \in [0, t]}$ denotes the action path of agent $i \in \{1, 2\}$ up to time t .

The set of all histories up to an arbitrary time is denoted by H . For each $i \in \{1, 2\}$ and any $q' \in \mathbb{R}_{++}$, let $H^{i, q'} \subseteq H$ be the set consisting of every history up to any time t of the form $(\{c_\tau\}_{\tau \in [0, t]}, \{f_\tau^1, f_\tau^2\}_{\tau \in [0, t]})$ where $q' = q - \sum_{\{\tau \in [0, t]: f_\tau^i \neq z\}} f_\tau^i$.

Let $\bar{A}_i(h_t)$ denote the set of feasible actions for agent $i \in \{1, 2\}$ at history $h_t \in H$. If $h_t \in H^{i, q'}$ for some $q' \in \mathbb{R}_{++}$, then let $\bar{A}_i(h_t) = \{z\} \cup (0, q']$. Otherwise, let $\bar{A}_i(h_t) = \{z\}$.

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

$$\bar{\Pi}_i = \{\pi_i : H \rightarrow \{z\} \cup (0, q] \mid \pi_i(h_t) \in \bar{A}_i(h_t) \text{ for all } h_t \in H\}.$$

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

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 used in the definition of the expected payoff is specified as follows for each $i = 1, 2$:

$$v_i[(f_\tau^1, f_\tau^2), s_\tau] = \begin{cases} 0 & \text{if } f_\tau^i = f_\tau^{-i} = z \\ -c_\tau & \text{if } f_\tau^i \neq z \text{ and } f_\tau^{-i} = z \\ f_\tau^{-i} & \text{if } f_\tau^i = z \text{ and } f_\tau^{-i} \neq z \\ f_\tau^{-i} - c_\tau & \text{if } f_\tau^i \neq z \text{ and } f_\tau^{-i} \neq z \end{cases}.$$

A.2 Proof of Theorem 1

Proof. Denote $\inf(S) = \underline{c} > 0$. Let \tilde{H}_τ denote the set consisting of every history $k_\tau = (\{g_v\}_{v \in [0, \tau]}, \{(b_v^1, b_v^2)\}_{v \in [0, \tau]})$ up to time τ such that the set $\{v \in [0, \tau) : b_v^i > 0\}$

is finite for each $i \in \{1, 2\}$. Define $\tilde{H} = \bigcup_{\tau \in [0, \infty)} \tilde{H}_\tau$. We prove by induction that no strategies in which some agent i makes a positive transfer f_t^i at some $h_t \in \tilde{H}$ can be played in any SPE.

Consider any $h_t \in \tilde{H}$. Define $\hat{s}_t^{i,i} = s_t^{i,i} + f_t^i$. That is, $\hat{s}_t^{i,i} \in [0, q]$ represents the amount of good i that agent i possesses immediately before time t . If $\hat{s}_t^{1,1} < \underline{c}$, then no strategy where agent 2 makes a positive transfer at h_t can be played in any SPE, because such a strategy would give agent 2 an expected payoff no greater than $\hat{s}_t^{1,1} - \underline{c} < 0$, whereas agent 2 could obtain an expected payoff of at least zero by making no transfers after reaching h_t . Because no strategy where agent 2 makes a positive transfer at some h_t with $\hat{s}_t^{1,1} < \underline{c}$ can be played in any SPE, agent 1 would obtain an expected payoff no greater than $-\underline{c} < 0$ by making a positive transfer at an h_t with $\hat{s}_t^{1,1} < \underline{c}$ but would obtain an expected payoff of zero by making no transfers after such an h_t . Thus, no strategy where agent 1 makes a positive transfer at some $h_t \in \tilde{H}$ in which $\hat{s}_t^{1,1} < \underline{c}$ can be played in any SPE. A symmetric argument shows that no strategy where agent 1 or 2 makes a positive transfer at some $h_t \in \tilde{H}$ in which $\hat{s}_t^{2,2} < \underline{c}$ can be played in any SPE.

Suppose now that for some integer $n \geq 1$, no strategies where agent 1 or 2 makes a positive transfer at some $h_t \in \tilde{H}$ in which $\hat{s}_t^{1,1} < n\underline{c}$ or $\hat{s}_t^{2,2} < n\underline{c}$ can be played in any SPE. Given this assumption, we show that no strategies where agent 1 or 2 makes a positive transfer at some $h_t \in \tilde{H}$ in which $n\underline{c} \leq \hat{s}_t^{1,1} < (n+1)\underline{c}$ or $n\underline{c} \leq \hat{s}_t^{2,2} < (n+1)\underline{c}$ can be played in any SPE. Consider in particular any $h_t \in \tilde{H}$ in which $n\underline{c} \leq \hat{s}_t^{1,1} < (n+1)\underline{c}$. If agent 1 is using a strategy that can be played in an SPE, then the greatest amount of the good that agent 1 can transfer at h_t is $\hat{s}_t^{1,1} - n\underline{c}$, because if agent 1 made a transfer greater than $\hat{s}_t^{1,1} - n\underline{c}$ at h_t , then with probability one, the remaining amount $\hat{s}_u^{1,1}$ of agent 1's good for $u > t$ would be such that agent 2 makes no further transfers, implying that agent 1 could obtain a higher expected payoff by instead making no transfers after reaching h_t . Thus, if agent 1 is using a strategy that can be played in an SPE and h_t is the history up to time t , then with probability one, the history h_u up to any time $u > t$ must be such that $\hat{s}_u^{1,1} \geq n\underline{c}$, in which case the total amount transferred by agent 1 after reaching h_t is at most $\hat{s}_t^{1,1} - n\underline{c} < \underline{c}$.

It follows that no strategy in which agent 2 makes a positive transfer at h_t can be played in any SPE, because making a positive transfer would give agent 2 an expected payoff no greater than $\hat{s}_t^{1,1} - n\underline{c} - \underline{c} < 0$, whereas agent 2 could obtain an expected payoff of at least zero by making no transfers after reaching h_t . Thus, no strategy

where agent 1 makes a transfer at h_t can be played in any SPE, because agent 1 would obtain an expected payoff no greater than $-\underline{c} < 0$ by making a transfer at h_t but would obtain an expected payoff of zero by making no transfers after reaching h_t . Thus, no strategies where agent 1 or 2 makes a transfer at some $h_t \in \tilde{H}$ in which $n\underline{c} \leq \hat{s}_t^{1,1} < (n+1)\underline{c}$ can be played in any SPE. A symmetric argument holds for any $h_t \in \tilde{H}$ in which $n\underline{c} \leq \hat{s}_t^{2,2} < (n+1)\underline{c}$. This completes the induction. \square

A.3 Proof of Theorem 2

Proof. Consider the symmetric grim-trigger strategy profile ψ defined as follows. Recall that q denotes the amount of good i that agent i initially possesses. Letting $\delta = e^{-\rho v}$, choose any $c^* > 0$ no greater than $\delta pqr(1-r)$. 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 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 the strategies in ψ are 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 reaches \hat{c} . If the agents follow strategy profile ψ , then the cost incurred by each agent on transaction l is no greater than $r^{l-1}c^*$, and the expected payoff to each agent immediately after transaction l is no less than $\sum_{m=1}^{\infty} \delta^m p^m [r^{l+m-1}q(1-r) - r^{l+m-1}c^*]$. Hence, the incentive compatibility constraint is satisfied for transaction l if the following holds:

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

which reduces to:

$$c^* \leq \frac{\delta pr[q(1-r) - c^*]}{1 - \delta pr} \Leftrightarrow c^* \leq \delta pqr(1-r).$$

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

A.4 Proof of Proposition 1

Proof. Let \tilde{H}_τ denote the set comprising every history $k_\tau = (\{g_\nu\}_{\nu \in [0, \tau]}, \{(b_\nu^1, b_\nu^2)\}_{\nu \in [0, \tau]})$ up to time τ such that the set $\{\nu \in [0, \tau] : b_\nu^i > 0\}$ is finite for each $i \in \{1, 2\}$. Define $\tilde{H} = \bigcup_{\tau \in [0, \infty)} \tilde{H}_\tau$. Given any history $h_t = (\{c_\tau\}_{\tau \in [0, t]}, \{(f_\tau^1, f_\tau^2)\}_{\tau \in [0, t]})$ up to an arbitrary time t , let $h = \{g_\tau, (b_\tau^1, b_\tau^2)\}_{\tau \in [0, \infty)}$ be any history such that $\{g_\tau\}_{\tau \in [0, t]} = \{c_\tau\}_{\tau \in [0, t]}$ and $\{(b_\tau^1, b_\tau^2)\}_{\tau \in [0, t]} = \{(f_\tau^1, f_\tau^2)\}_{\tau \in [0, t]}$. Then h_t is said to be on the equilibrium path induced by π if h is consistent with π_1 and π_2 at each time $\tau \in [0, t)$.

Fix any SPE $\pi = (\pi_1, \pi_2)$. We construct an SPE that induces the same equilibrium path of play using grim-trigger strategies. Given any $h_t \in \tilde{H}$, recall that $V_i(h_t, \pi)$ is the expected payoff to agent i when both agents follow the strategy profile π from time t onwards, and let $X_i(h_t, \pi)$ denote the supremum of the expected payoffs to agent i from any set of deviations when π_{-i} is fixed. Because π is an SPE, it must be that $V_i(h_t, \pi) \geq X_i(h_t, \pi)$.

Now consider the grim-trigger strategy profile $\pi' = (\pi'_1, \pi'_2)$ such that $[\pi'_1(h_t), \pi'_2(h_t)] = [\pi_1(h_t), \pi_2(h_t)]$ for any $h_t \in H$ on the equilibrium path induced by π and such that $[\pi'_1(h_t), \pi'_2(h_t)] = (0, 0)$ for any history $h_t \in H$ not on the equilibrium path induced by π . We show that π' is an SPE strategy profile.⁴¹

Suppose that $h_t \in \tilde{H}$ is on the equilibrium path induced by π . If $\pi'_i(h_t) = 0$, then the expected payoff to agent i from a one-shot deviation at h_t must be negative. However, the expected payoff to agent i when both agents follow the strategy profile π' from h_t onwards must be nonnegative, because $\pi'(k_u) = \pi(k_u)$ for any $k_u \in H$ on the equilibrium path induced by π , where π is an SPE. Hence, if $\pi'_i(h_t) = 0$, then agent i does not have an incentive to make a one-shot deviation at h_t .

Assume instead that $\pi'_i(h_t) > 0$. By definition, if agent $-i$ is using strategy π'_{-i} , then the expected payoff to agent i from following π'_i from h_t onwards is $V_i(h_t, \pi)$; so that, $V_i(h_t, \pi') = V_i(h_t, \pi)$. In addition, if agent $-i$ is using strategy π'_{-i} , then a one-shot deviation at h_t gives agent i an expected payoff of at most $\pi'_{-i}(h_t)$, because agent $-i$ makes no transfers at any history not on the equilibrium path induced by π . Furthermore, agent i can ensure that she receives an expected payoff of $\pi'_{-i}(h_t)$ by

⁴¹When confirming that the grim-trigger strategy profile π' is an SPE, it is sufficient to consider only one-shot deviations from π' .

transferring nothing from time t onward. Hence, $\tilde{X}_i(h_t, \pi') = \pi'_{-i}(h_t)$, where $\tilde{X}_i(h_t, \pi')$ is the supremum of the expected payoffs to agent i at h_t from any one-shot deviation when π'_{-i} is fixed.

Now note that $X_i(h_t, \pi) \geq \pi_{-i}(h_t) = \pi'_{-i}(h_t)$ because if agent $-i$ is using strategy π_{-i} , then agent i can always obtain at history h_t an expected payoff of at least $\pi_{-i}(h_t)$ by transferring nothing from time t onward. Thus, we have $V_i(h_t, \pi') = V_i(h_t, \pi) \geq X_i(h_t, \pi) \geq \pi'_{-i}(h_t) = \tilde{X}_i(h_t, \pi')$ for any $h_t \in \tilde{H}$ on the equilibrium path induced by π such that $\pi'_i(h_t) > 0$. Moreover, for any $h_t \in \tilde{H}$ not on the equilibrium path induced by π , we have $V_i(h_t, \pi') \geq X_i(h_t, \pi')$ because $V_i(h_t, \pi') = 0 = X_i(h_t, \pi')$. Hence, the strategy profile π' constitutes an SPE. Since π' is defined so as to agree with π on the equilibrium path, this completes the proof. \square

A.5 Lemmata in Proof of Theorem 3

It is helpful to introduce some additional terminology and notation. Consider any strategy profiles $\pi^a = (\pi_1^a, \pi_2^a)$ and $\pi^b = (\pi_1^b, \pi_2^b)$ with $\pi_i^a \in \bar{\Pi}_i^C$ and $\pi_i^b \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. The strategy profiles π^a and π^b are said to induce different paths of play with probability p if p is the probability that the realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that there exist $t \in [0, \infty)$ and $i \in \{1, 2\}$ for which $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^a) \neq \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^b)$. If there is probability zero that the strategy profiles π^a and π^b induce different paths of play, then π^a and π^b are said to induce the same path of play with probability one.

Next consider any strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. We denote:

$$h_t(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = (\{c_\tau\}_{\tau \in [0, t]}, \{[\phi_\tau^1(h_0, \{c_\nu\}_{\nu \in (0, \infty)}, \pi), \phi_\tau^2(h_0, \{c_\nu\}_{\nu \in (0, \infty)}, \pi)]\}_{\tau \in [0, t]}).$$

That is, $h_t(\{c_\tau\}_{\tau \in [0, \infty)}, \pi)$ represents the history up to time t when the strategy profile π is played and the realization of the cost is $\{c_\tau\}_{\tau \in [0, \infty)}$. We also define:

$$\sigma_t^{i,j}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = \begin{cases} q - \sum_{\{v \in [0, t]: \phi_v^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0\}} \phi_v^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi), & \text{if } i = j \\ \sum_{\{v \in [0, t]: \phi_v^j(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) > 0\}} \phi_v^j(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi), & \text{if } i \neq j \end{cases}.$$

We also denote $\hat{\sigma}_t^{i,j}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = \lim_{v \rightarrow t^-} \sigma_v^{i,j}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi)$, letting $\hat{\sigma}_0^{i,j}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = \sigma_0^{i,j}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = q$ if $i = j$ and $\hat{\sigma}_0^{i,j}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = \sigma_0^{i,j}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = 0$

if $i \neq j$. That is, $\hat{\sigma}_t^{i,j}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)$ and $\sigma_t^{i,j}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)$ represent the amount of good j that agent i possesses immediately before and after time t , respectively.

We now define the concept of a dual cutoff form. Some further notation is useful. Given a symmetric strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$, let $t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)$ signify the k^{th} smallest time t such that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0,\infty)}, \pi) > 0$, where k is a positive integer. If no such time exists, then let $t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) = \infty$. That is, $t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)$ is the time of the k^{th} transaction when the strategy profile π is played and the realization of the cost is $\{c_\tau\}_{\tau \in [0,\infty)}$. For notational purposes, let $t_0(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) = 0$.

In addition, we abbreviate:

$$\begin{aligned}\hat{\sigma}^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k) &= \hat{\sigma}_{t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)}^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) \\ \sigma^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k) &= \sigma_{t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)}^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) \\ \phi^i(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k) &= \phi_{t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)}^i(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) \\ c(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k) &= c_{t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi)},\end{aligned}$$

assuming that $t_k(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) < \infty$. That is, $\hat{\sigma}^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k)$, $\sigma^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k)$, $\phi^i(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k)$, and $c(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k)$ respectively denote the amount of good i possessed by agent i immediately before the k^{th} transaction, the amount of good i possessed by agent i immediately after the k^{th} transaction, the amount of good i transferred by agent i on the k^{th} transaction, and the value of the cost process at the k^{th} transaction when the strategy profile π is played and the realization of the cost is $\{c_\tau\}_{\tau \in [0,\infty)}$.

Definition 1. Choose any symmetric strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. Let k be a positive integer. Transaction k has a **dual cutoff form** when playing π if there is probability one that $t_{k-1}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) = \infty$ or both $t_{k-1}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi) < \infty$ and the following holds. Given the realizations of $c(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, n)$ and $\hat{\sigma}^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, n)$ for $n \in \{0, 1, \dots, k-1\}$, there exist:

1. c^a with $0 \leq c^a \leq c(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k-1)$,
2. ϕ^a with $0 < \phi^a \leq \sigma^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k-1)$,
3. c^b with $c(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k-1) \leq c^b \leq \infty$,
4. ϕ^b with $0 < \phi^b \leq \sigma^{i,i}(\{c_\tau\}_{\tau \in [0,\infty)}, \pi, k-1)$.

for which there is conditional probability one that the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that one of the three sets of requirements below is satisfied:

1. $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) < \infty$, $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi, k) = c^a$, $\phi^i(\{c_\tau\}_{\tau \in [0, \infty)}, \pi, k) = \phi^a$, there does not exist u with $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) < u < t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi)$ such that $c_u = c^a$ or $c_u = c^b$.
2. $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) < \infty$, $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi, k) = c^b$, $\phi^i(\{c_\tau\}_{\tau \in [0, \infty)}, \pi, k) = \phi^b$, there does not exist u with $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) < u < t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi)$ such that $c_u = c^a$ or $c_u = c^b$.
3. $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi) = \infty$, there does not exist $u > t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi)$ such that $c_u = c^a$ or $c_u = c^b$.

The following is a verbal description of this concept. The first transaction has a dual cutoff form when playing a symmetric strategy profile π if there exists $c^a \in [0, c_0]$, $c^b \in [c_0, \infty]$ and $\phi^a \in (0, s_0^{i,i}]$, $\phi^b \in (0, s_0^{i,i}]$ such that, with probability one, the agents transfer the amount ϕ^a at the first positive time the cost reaches c^a if and only if c^b has not been reached earlier, the agents transfer the amount ϕ^b at the first positive time the cost reaches c^b if and only if c^a has not been reached earlier, and the agents do not make any transfers until a cost of c^a or c^b has been reached.⁴² For any positive integer k , the $(k+1)^{\text{th}}$ transaction has a dual cutoff form when playing a symmetric strategy profile π if there is probability one that a k^{th} transaction does not occur or both a k^{th} transaction occurs and the first transaction in the subgame immediately following the k^{th} transaction has a dual cutoff form when playing π .

The concept of a dual cutoff form is closely related to the concept of stationarity for strategy profiles. The stationarity property implies a dual cutoff form, but not vice versa. The result below shows that each transaction has a dual cutoff form when playing a stationary symmetric strategy profile.

Claim 1. *If $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ is a stationary symmetric strategy profile, then transaction k has a dual cutoff form when playing π , where k is an arbitrary positive integer.*

Proof. Let $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ be a stationary symmetric strategy profile. By definition, there exists a symmetric strategy profile $\pi^\dagger = (\pi_1^\dagger, \pi_2^\dagger)$ with

⁴²If $c^a = 0$, then there is zero probability of the first transaction occurring at a cost less than c_0 . If $c^b = \infty$, then there is zero probability of the first transaction occurring at a cost higher than c_0 .

$\pi_i^\dagger \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ satisfying the following two conditions. For any $c \in (0, \infty)$ and $s_1, s_2 \in [0, q]$, $\pi_1^\dagger(k'_{u'}) = \pi_1^\dagger(k''_{u''})$ and $\pi_2^\dagger(k'_{u'}) = \pi_2^\dagger(k''_{u''})$ for all $k'_{u'}, k''_{u''} \in \tilde{H}(c, s_1, s_2)$. There is probability one that $\phi_i^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) = \phi_i^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^\dagger)$ for all $t \in [0, \infty)$ and $i \in \{1, 2\}$. Choose any positive integer k . If $k > 1$, then assume that for every positive integer $n < k$, transaction n has a dual cutoff form when playing π^\dagger . We argue that transaction k has a dual cutoff form when playing π^\dagger . Because π^\dagger induces the same path of play as π with probability one, it then follows that transaction k has a dual cutoff form when playing π .

If $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger) = \infty$ with probability one, then transaction k has a dual cutoff form when playing π^\dagger . Therefore, assume that $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger) < \infty$ with positive probability. Choose any c^\dagger and s^\dagger for which there is positive probability of the realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that both $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger, k-1) = c^\dagger$ and $\hat{\sigma}^{i,i}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger, k-1) = s^\dagger$. Let $\{c_\tau^*\}_{\tau \in [0, \infty)}$ be any cost realization such that $c(\{c_\tau^*\}_{\tau \in [0, \infty)}, \pi^\dagger, k-1) = c^\dagger$ and $\hat{\sigma}^{i,i}(\{c_\tau^*\}_{\tau \in [0, \infty)}, \pi^\dagger, k-1) = s^\dagger$. Let T be the set consisting of every c for which there exists $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ such that $t_k(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger) < \infty$, $c(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger, k) = c$, and $\tilde{c}_t = c_t^*$ for $t \leq t_{k-1}(\{c_\tau^*\}_{\tau \in [0, \infty)}, \pi^\dagger)$.

Let c^a denote the supremum of the set consisting of every $c \leq c^\dagger$ such that $c \in T$. Let c^b denote the infimum of the set consisting of every $c \geq c^\dagger$ such that $c \in T$. If there is no $c \in T$ with $c \leq c^\dagger$, then let $c^a = 0$. If there is no $c \in T$ with $c \geq c^\dagger$, then let $c^b = \infty$. Note that there is no $c \in T$ with $c < c^a$ and there is no $c \in T$ with $c > c^b$. Otherwise, there would exist some $c \in (0, \infty)$ and $s \in [0, q]$ such that $\pi_1^\dagger(k'_{u'}) \neq \pi_1^\dagger(k''_{u''})$ and $\pi_2^\dagger(k'_{u'}) \neq \pi_2^\dagger(k''_{u''})$ for some $k'_{u'}, k''_{u''} \in \tilde{H}(c, s, s)$. This would contradict the properties of the strategy profile π^\dagger .

For $x \in \{a, b\}$, let $\{c_\tau^x\}_{\tau \in [0, \infty)}$ be any cost realization such that $t_k(\{c_\tau^x\}_{\tau \in [0, \infty)}, \pi^\dagger) < \infty$, $c(\{c_\tau^x\}_{\tau \in [0, \infty)}, \pi^\dagger, k) = c^x$, and $c_t^x = c_t^*$ for $t \leq t_{k-1}(\{c_\tau^x\}_{\tau \in [0, \infty)}, \pi^\dagger)$. For $x \in \{a, b\}$, define $\phi^x = \phi^i(\{c_\tau^x\}_{\tau \in [0, \infty)}, \pi^\dagger, k)$, provided that the cost realization $\{c_\tau^x\}_{\tau \in [0, \infty)}$ in the previous sentence exists. Suppose that the realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ satisfies $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger, k-1) = c^\dagger$ and $\hat{\sigma}^{i,i}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger, k-1) = s^\dagger$. Then $\{c_\tau\}_{\tau \in [0, \infty)}$ is either such that $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger) = \infty$ or such that $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger) < \infty$, $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger, k) = c^x$, and $\phi^i(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger, k) = \phi^x$ for some $x \in \{a, b\}$. Moreover, if $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger) < \infty$, then there exists no u with $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger) < u < t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^\dagger)$ such that $c_u = c^x$ for some $x \in \{a, b\}$. Hence, transaction k has a dual cutoff form when playing π^\dagger . \square

The next result provides an example to show that not every symmetric strategy

profile in which each transaction has a dual cutoff form is stationary.

Claim 2. *There exists a non-stationary symmetric strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ such that for every positive integer k , transaction k has a dual cutoff form when playing π .*

Proof. Consider the non-stationary symmetric grim-trigger strategy profile π defined as follows. First, choose any $c^1 \in (0, c_0)$ and $\phi^1 \in (0, q)$. The agents each transfer the amount ϕ^1 at the first time that the cost reaches c^1 , and the agents do not make any transfers until a cost of c^1 has been reached.

Second, choose any $c^{2,a} \in (0, c^1)$, $c^{2,b} \in (c^1, \infty)$, and $\phi^2 \in (0, q - \phi^1)$. Each agent transfers the amount ϕ^2 at the first time the cost reaches $c^{2,a}$ after the first transaction by each agent if and only if $c^{2,b}$ has not been reached earlier following the first transaction, each agent transfers the amount ϕ^2 at the first time the cost reaches $c^{2,b}$ after the first transaction by each agent if and only if $c^{2,a}$ has not been reached earlier following the first transaction, and the agents do not make any transfers after the first transaction until a cost of $c^{2,a}$ or $c^{2,b}$ has been reached following the first transaction.

Third, choose any $c^3 \in \mathbb{R}_{++}$ with $c^3 \neq c^{2,a}$ and $c^3 \neq c^{2,b}$, and choose any $\phi^{3,a} \in (0, q - \phi^1 - \phi^2)$ and $\phi^{3,b} \in (0, q - \phi^1 - \phi^2)$ with $\phi^{3,a} \neq \phi^{3,b}$. If the cost incurred on the second transaction is $c^{2,a}$, then the agents each transfer the amount $\phi^{3,a}$ at the first time after the second transaction by each agent that the cost reaches c^3 , and the agents do not make any transfers after the second transaction until a cost of c^3 has been reached following the second transaction. If the cost incurred on the second transaction is $c^{2,b}$, then the agents each transfer the amount $\phi^{3,b}$ at the first time after the second transaction by each agent that the cost reaches c^3 , and the agents do not make any transfers after the second transaction until a cost of c^3 has been reached following the second transaction.

Each agent makes at most three transactions. If an agent deviates from the path of play described above, then neither agent makes any transactions following the deviation. The strategy profile π is such that for every positive integer k , transaction k has a dual cutoff form when playing π . \square

The following are the statements and proofs of the three lemmata that imply theorem 3. The result below shows that given any non-stationary symmetric SPE

π , one can find a stationary symmetric SPE π' that generates at least as high an expected payoff to each agent.

Here is an outline of the proof. Let π be an arbitrary symmetric SPE. We begin by using the Markov property of the cost process to construct a symmetric SPE in which the payoff to each agent is at least as high as under π and the first transaction has a dual cutoff form. The next step is to find a symmetric SPE that satisfies the conditions in the preceding sentence as well as the additional requirement that the incentive constraint is binding on the first transaction. It is then straightforward to specify a symmetric SPE in which the first transaction has a dual cutoff form, the incentive constraint is binding on the first transaction, and the payoff to each agent is the same as under π . An iterative procedure is next used along with the Markov property of the cost to define a symmetric SPE in which every transaction has a dual cutoff form, the incentive constraint is binding on every transaction, and the payoff to each agent is the same as under π . Finally, a stationary symmetric SPE π' is derived in which the payoff to each agent is at least as high as under π . In addition, the cost incurred between any two consecutive transactions is decreasing when playing π' .

Lemma 1. *Given any symmetric SPE $\pi \in \Pi^*$ in grim-trigger strategies, there exists a stationary symmetric SPE $\pi' \in \Pi^*$ in grim-trigger strategies such that $V(h_0, \pi') \geq V(h_0, \pi)$. Moreover, π' can be chosen such that the cost incurred between any two consecutive transactions is decreasing.*

Proof. Let the strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ be an arbitrary symmetric SPE in grim-trigger strategies. Define the symmetric SPE $\pi^0 = (\pi_1^0, \pi_2^0)$ in grim-trigger strategies as follows. For any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ up to an arbitrary time u , let $\pi_i^0(k_u) = \pi_i(k_u)$ if $\pi_i(k_u) \geq g_u$, and let $\pi_i^0(k_u) = 0$ if $\pi_i(k_u) < g_u$, where $i \in \{1, 2\}$. Note that $V(h_0, \pi^0) \geq V(h_0, \pi)$.

We begin by finding a symmetric SPE $\pi^* \in \Pi^*$ in grim-trigger strategies with $V(h_0, \pi^*) \geq V(h_0, \pi^0)$ such that the first transaction has a dual cutoff form when playing π^* . There are three cases to consider. First, π^0 is a stationary symmetric SPE in grim-trigger strategies. Second, π^0 is a non-stationary maximal symmetric SPE in grim-trigger strategies. Third, π^0 is a non-stationary strategy profile that is not a maximal symmetric SPE in grim-trigger strategies. The strategy profile π^* is constructed as follows in each case. In the first case, where π^0 is a stationary symmetric SPE in grim-trigger strategies, simply let $\pi^* = \pi^0$.

Next, consider the second case, where π^0 is a non-stationary maximal symmetric SPE in grim-trigger strategies. Either there is zero probability that $t_1(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^0) < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^0, 1) < c_0$, or there exists a realization $\{c_\tau^a\}_{\tau \in [0, \infty)}$ of the cost process as well as a time t^a such that $t_1(\{c_\tau^a\}_{\tau \in [0, \infty)}, \pi^0) = t^a$, $c_{t^a}^a < c_0$, and $V[h_t(\{c_\tau^a\}_{\tau \in [0, \infty)}, \pi^0), \pi^0] \geq V[h_t(\{c_\tau^a\}_{\tau \in [0, \infty)}, \pi^0), \hat{\pi}]$ for every symmetric SPE $\hat{\pi}$ in grim-trigger strategies along with any time $t \leq t^a$. If such a cost realization $\{c_\tau^a\}_{\tau \in [0, \infty)}$ exists, then let $c^{a'} = c_{t^a}^a$. Otherwise, let $c^{a'} = 0$. Either there is zero probability that $t_1(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^0) < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^0, 1) \geq c_0$, or there exists a realization $\{c_\tau^b\}_{\tau \in [0, \infty)}$ of the cost process as well as a time t^b such that $t_1(\{c_\tau^b\}_{\tau \in [0, \infty)}, \pi^0) = t^b$, $c_{t^b}^b \geq c_0$, and $V[h_t(\{c_\tau^b\}_{\tau \in [0, \infty)}, \pi^0), \pi^0] \geq V[h_t(\{c_\tau^b\}_{\tau \in [0, \infty)}, \pi^0), \hat{\pi}]$ for every symmetric SPE $\hat{\pi}$ in grim-trigger strategies along with any time $t \leq t^b$. If such a cost realization $\{c_\tau^b\}_{\tau \in [0, \infty)}$ exists, then let $c^{b'} = c_{t^b}^b$. Otherwise, let $c^{b'} = \infty$. Define the strategy profile $\pi^* \in \Pi^*$ in grim-trigger strategies as follows. The agents do not make any transactions until the first time t^* that the current value of the cost process is $c^{a'}$ or $c^{b'}$. If no such time t^* exists, then the agents never make a transaction. Consider the case where there exists such a time t^* . If $c^{a'} \in (0, c_0)$ and $c_{t^*} = c^{a'}$, then each agent transfers the amount $\phi_{t^*}^i(h_0, \{c_\tau^a\}_{\tau \in (0, \infty)}, \pi^0)$ at time t^* . After this transaction, the agents play according to strategy profile π^0 , behaving as if the history up to the time of this transaction were $h_{t^*}(\{c_\tau^a\}_{\tau \in [0, \infty)}, \pi^0)$ and strategy profile π^0 had always been followed from the beginning of the game. If $c^{b'} \in [c_0, \infty)$ and $c_{t^*} = c^{b'}$, then each agent transfers the amount $\phi_{t^*}^i(h_0, \{c_\tau^b\}_{\tau \in (0, \infty)}, \pi^0)$ at time t^* . After this transaction, the agents play according to strategy profile π^0 , behaving as if the history up to the time of this transaction were $h_{t^*}(\{c_\tau^b\}_{\tau \in [0, \infty)}, \pi^0)$ and strategy profile π^0 had always been followed from the beginning of the game.

Now, consider the third case, where π^0 is a non-stationary strategy profile that is not a maximal symmetric SPE in grim-trigger strategies. In this case, one can find a symmetric SPE $\tilde{\pi} = (\tilde{\pi}_1, \tilde{\pi}_2)$ in grim-trigger strategies such that $V(h_0, \tilde{\pi}) > V(h_0, \pi^0)$. Define the symmetric SPE $\tilde{\pi}^0 = (\tilde{\pi}_1^0, \tilde{\pi}_2^0)$ in grim-trigger strategies as follows. For any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ up to an arbitrary time u , let $\tilde{\pi}_i^0(k_u) = \tilde{\pi}_i(k_u)$ if $\tilde{\pi}_i(k_u) \geq g_u$, and let $\tilde{\pi}_i^0(k_u) = 0$ if $\tilde{\pi}_i(k_u) < g_u$, where $i \in \{1, 2\}$. Note that $V(h_0, \tilde{\pi}^0) \geq V(h_0, \tilde{\pi})$. Let $\delta = V(h_0, \tilde{\pi}^0) - V(h_0, \pi^0)$.

Consider the strategy profile $\tilde{\pi}^0$. Let T be the set consisting of each cost level c for which there exists a realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ of the cost process as well as a time \tilde{u} such that $\tilde{c}_{\tilde{u}} = c$, $\hat{\sigma}_{\tilde{u}}^{i,i}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}^0) = s_0^{i,i}$, and $\phi_{\tilde{u}}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}^0) > 0$. If the set

T is empty, then the expected payoff to each agent from playing strategy profile $\tilde{\pi}^0$ is zero, because no transfers are made on the equilibrium path. However, this would contradict the fact that π^0 is an SPE, because $\tilde{\pi}^0$ generates a strictly higher expected payoff to each agent than π^0 , and no SPE can generate an expected payoff to each agent lower than zero. Therefore, it can be assumed that the set T is non-empty.

For any cost level c in T , let $K(c)$ be the set consisting of every history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ up to an arbitrary time u for which $g_u = c$, $\tilde{\pi}_i^0(k_u) > 0$, $b_v^i = \tilde{\pi}_i^0[\{g_t\}_{t \in [0, v]}, \{(b_t^1, b_t^2)\}_{t \in [0, v]}] = 0$ for all $v < u$. Define $Q(c) = \sup_{k_u \in K(c)} V(k_u, \tilde{\pi}^0)$. A state is a pair (c, s) , where c is the current value of the cost process, and s is the amount of each good remaining untransferred. For any cost levels c^a and c^b in T , let $R(c^a, c^b)$ denote the value of an asset at state $(c_0, s_0^{i,i})$ that pays $Q(c^a)$ at the first time the cost reaches c^a if and only if c^b has not been reached earlier and that pays $Q(c^b)$ at the first time the cost reaches c^b if and only if c^a has not been reached earlier. Letting $X = \sup_{c^a, c^b \in T} R(c^a, c^b)$, it follows from the properties of the cost process that $X \geq V(h_0, \tilde{\pi}^0)$.

To understand the last claim, consider the following argument. If the agents play strategy profile $\tilde{\pi}^0$, then $Q(c)$ is an upper bound on the expected payoff to each agent when a state $(c, s_0^{i,i})$ with c in T is reached and the agents choose to make a transfer at this time. Hence, $V(h_0, \tilde{\pi}^0) \leq B$, where B denotes the supremum of the set of possible expected payoffs to each agent in the following problem. Assume that the initial state is $(c_0, s_0^{i,i})$. At each time a state $(c, s_0^{i,i})$ with $c \in T$ is reached, the agents can choose either to accept or to decline, provided that the agents have not chosen to accept in the past. If the agents accept, then each agent receives an immediate payoff of $Q(c)$, and the problem ends. If the agents decline, then each agent receives an immediate payoff of zero, but the problem continues. Given any state $(c, s_0^{i,i})$ with $c \in T$, let $H(c)$ be the supremum of the set of possible expected payoffs to each agent upon reaching state $(c, s_0^{i,i})$ if the agents follow policies in which they choose to decline at this time.

Let Z be the set consisting of each cost level $c \in T$ such that $Q(c) \geq H(c)$. Let z^l denote the supremum of the set consisting of every cost level $c \in Z$ such that $c \leq c_0$. Let z^h denote the infimum of the set consisting of every cost level $c \in Z$ such that $c \geq c_0$. For any $\epsilon > 0$, let $Z^l(\epsilon)$ denote the set consisting of every cost level $c \in Z$ for which $c \in [z^l - \epsilon, z^l]$, and let $Z^h(\epsilon)$ denote the set consisting of every cost level $c \in Z$ for which $c \in [z^h, z^h + \epsilon]$. For any $\epsilon > 0$, define $Q^l(\epsilon) = \sup_{c \in Z^l(\epsilon)} Q(c)$ and

$Q^h(\epsilon) = \sup_{c \in Z^h(\epsilon)} Q(c)$. Let $D(\epsilon)$ denote the value of an asset at state $(c_0, s_0^{i,i})$ that pays $Q^h(\epsilon)$ at the first time the cost reaches $z^h + \epsilon$ if and only if $z^l - \epsilon$ has not been reached earlier and that pays $Q^l(\epsilon)$ at the first time the cost reaches $z^l - \epsilon$ if and only if $z^h + \epsilon$ has not been reached earlier. Note that $\lim_{\epsilon \rightarrow 0} D(\epsilon) = B$. For any $\epsilon > 0$ and $\eta > 0$, there exist cost levels $c^{a''}$ and $c^{b''}$ for which the expected payoff to each agent is greater than $D(\epsilon) - \eta$ when the agents follow a policy of accepting at the first time the cost reaches $c^{a''}$ if and only if $c^{b''}$ has not been reached earlier and accepting at the first time the cost reaches $c^{b''}$ if and only if $c^{a''}$ has not been reached earlier.

Define the strategy profile π^* as follows. Choose any realizations $\{c_\tau^p\}_{\tau \in [0, \infty)}$ and $\{c_\tau^q\}_{\tau \in [0, \infty)}$ of the cost process and times t^p and t^q such that $\hat{\sigma}_{t^p}^{i,i}(\{c_\tau^p\}_{\tau \in [0, \infty)}, \tilde{\pi}^0) = s_0^{i,i}$, $\hat{\sigma}_{t^q}^{i,i}(\{c_\tau^q\}_{\tau \in [0, \infty)}, \tilde{\pi}^0) = s_0^{i,i}$, $\phi_{t^p}^i(h_0, \{c_\tau^p\}_{\tau \in (0, \infty)}, \tilde{\pi}^0) > 0$, $\phi_{t^q}^i(h_0, \{c_\tau^q\}_{\tau \in (0, \infty)}, \tilde{\pi}^0) > 0$, and $X - P(\{c_\tau^p\}_{\tau \in [0, \infty)}, \{c_\tau^q\}_{\tau \in [0, \infty)}) < \delta/2$, where $P(\{c_\tau^p\}_{\tau \in [0, \infty)}, \{c_\tau^q\}_{\tau \in [0, \infty)})$ is the value of an asset at state $(c_0, s_0^{i,i})$ that pays $V[h_{t^p}(\{c_\tau^p\}_{\tau \in [0, \infty)}, \tilde{\pi}^0), \tilde{\pi}^0]$ at the first time the cost reaches $c_{t^p}^p$ if and only if $c_{t^q}^q$ has not been reached earlier and that pays $V[h_{t^q}(\{c_\tau^q\}_{\tau \in [0, \infty)}, \tilde{\pi}^0), \tilde{\pi}^0]$ at the first time the cost reaches $c_{t^q}^q$ if and only if $c_{t^p}^p$ has not been reached earlier. The strategy profile π^* requires the agents to make no transfers until the cost level $c_{t^p}^p$ or $c_{t^q}^q$ is reached. If the cost level $c_{t^p}^p$ is reached and the cost level $c_{t^q}^q$ has not been reached earlier, then π^* requires the agents to play according to the strategy profile $\tilde{\pi}^0$ from the first time that the cost level $c_{t^p}^p$ is reached, behaving as if the history up to this time were $h_{t^p}(\{c_\tau^p\}_{\tau \in [0, \infty)}, \tilde{\pi}^0)$. If the cost level $c_{t^q}^q$ is reached and the cost level $c_{t^p}^p$ has not been reached earlier, then π^* requires the agents to play according to the strategy profile $\tilde{\pi}^0$ from the first time that the cost level $c_{t^q}^q$ is reached, behaving as if the history up to this time were $h_{t^q}(\{c_\tau^q\}_{\tau \in [0, \infty)}, \tilde{\pi}^0)$. The strategy profile π^* requires the agents to make no transfers following a deviation from the path of play described above.

Consider the three cases above. In the first case, the stationarity of π^0 implies that the first transaction has a dual cutoff from when playing π^* . Recall that $c_0 > q$. In the second case, the feasibility of π_i implies that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) < c_0$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and time t such that $\tilde{c}_t = c_0$. It follows that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^0) = 0$ for any such $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and t . In the third case, the feasibility of $\tilde{\pi}_i$ implies that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \tilde{\pi}) < c_0$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and time t such that $\tilde{c}_t = c_0$. It follows that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \tilde{\pi}^0) = 0$ for any such $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and t . Hence, it must be that $\pi_i^*(h_0) = 0$ in the second and third cases. It follows that the first transaction has a dual cutoff form when playing π^* . Moreover,

$\pi^* \in \Pi^*$ is a symmetric SPE in grim-trigger strategies with $V(h_0, \pi^*) \geq V(h_0, \pi^0)$.

It is helpful to introduce some terminology. Given a symmetric SPE $\hat{\pi} \in \Pi^*$ as well as any positive integer k , let \hat{p} be the probability that $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}) < \infty$ and there exists a time t with $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}) < t < t_{k+1}(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi})$ for which $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \hat{\pi}) > 0$ and $Y[h_t(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}), \hat{\pi}] > c_t$. If $\hat{p} > 0$, then $\hat{\pi}$ is said to have a positive probability of the incentive constraint being slack at transaction k . If $\hat{p} = 0$, then $\hat{\pi}$ is said to have probability one of the incentive constraint being binding at transaction k .

We next construct a symmetric SPE $\pi^{**} \in \Pi^*$ in grim-trigger strategies with $V(h_0, \pi^{**}) \geq V(h_0, \pi^*)$ such that the first transaction has a dual cutoff form when playing π^{**} and such that π^{**} has probability one of the incentive constraint being binding at the first transaction. If π^* has probability one of the incentive constraint being binding at the first transaction, then simply let $\pi^{**} = \pi^*$. If π^{**} has a positive probability of the incentive constraint being slack at the first transaction, then the strategy profile π^{**} is constructed as described below.

It is helpful to introduce some notation. Let \hat{t}_1 and \hat{t}_2 be any two nonnegative real numbers such that $\hat{t}_1 \leq \hat{t}_2$. Let $\hat{\pi} \in \Pi^*$ be a symmetric SPE in grim-trigger strategies. Given a realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ of the cost process, let $\Sigma[h_{\hat{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}); \{\tilde{c}_\tau\}_{\tau \in [\hat{t}_1, \hat{t}_2]}; \hat{t}_1, \hat{t}_2; \hat{\pi}]$ be the sum of the transfers that an agent would make between times \hat{t}_1 and \hat{t}_2 inclusive if $h_{\hat{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \hat{\pi})$ is the history up to time \hat{t}_1 , the cost path $\{\tilde{c}_\tau\}_{\tau \in [\hat{t}_1, \hat{t}_2]}$ is realized between times \hat{t}_1 and \hat{t}_2 , and the strategy profile $\hat{\pi}$ is played by the agents.

The following is how π^{**} is defined in this case. The agents play according to strategy profile π^* until the first time t for which the realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ of the cost process is such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) > 0$ and $Y[h_t(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^*] > \tilde{c}_t$. If no such time t exists, then the agents simply follow strategy profile π^* . Otherwise, let \tilde{t}_1 denote the first time t that this condition holds. At time \tilde{t}_1 , each agent transfers the amount $\phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) + \zeta(\tilde{c}_{\tilde{t}_1})$, where $\zeta(\tilde{c}_{\tilde{t}_1})$ is a positive real number no greater than $s_0^{i,i} - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*)$ that can depend on $\tilde{c}_{\tilde{t}_1}$. After time \tilde{t}_1 , the agents do not make any transfers at any time $t > \tilde{t}_1$ such that the realization of the cost process $\{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, t]}$ between times \tilde{t}_1 and t satisfies $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) = 0$ or $\Sigma[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*); \{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, t]}; \tilde{t}_1, t; \pi^*] - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) - \zeta(\tilde{c}_{\tilde{t}_1}) < \tilde{c}_t$.

Let \tilde{t}_2 denote the first time t such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) > 0$ and $\Sigma[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*); \{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, t]}; \tilde{t}_1, t; \pi^*] - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) - \zeta(\tilde{c}_{\tilde{t}_1}) \geq \tilde{c}_t$. If such a time t does not exist, then the agents do not make any further transactions. Other-

wise, proceed as follows. At time \tilde{t}_2 , each agent transfers an amount equal to the lesser of $\phi_{\tilde{t}_2}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*)$ and $\Sigma[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*); \{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, \tilde{t}_2]}; \tilde{t}_1, \tilde{t}_2; \pi^*] - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) - \zeta(\tilde{c}_{\tilde{t}_1})$. Thereafter, the agents play according to strategy profile π^* , behaving as if the history up to time \tilde{t}_1 were $h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*)$, the realization of the cost process between times \tilde{t}_1 and \tilde{t}_2 were $\{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, \tilde{t}_2]}$, and strategy profile π^* were followed up to and including time \tilde{t}_2 . The agents do not make any transfers following a deviation from the path of play specified above.

If each $\zeta(\tilde{c}_{\tilde{t}_1})$ is chosen appropriately, then π^{**} is a symmetric SPE in grim-trigger strategies with $V(h_0, \pi^{**}) > V(h_0, \pi^*)$, and π^{**} has probability one of the incentive constraint being binding at the first transaction. Such a choice of $\zeta(\tilde{c}_{\tilde{t}_1})$ is possible because $V[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ is increasing in $\zeta(\tilde{c}_{\tilde{t}_1})$, $Y[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ approaches $Y[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^*]$ as $\zeta(\tilde{c}_{\tilde{t}_1})$ goes to 0, $Y[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ equals 0 for $\zeta(\tilde{c}_{\tilde{t}_1}) = s_0^{i,i} - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*)$, and $Y[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ is a continuous function of $\zeta(\tilde{c}_{\tilde{t}_1})$.

We now observe that it is possible to construct a symmetric SPE π^{***} in grim-trigger strategies with $V(h_0, \pi^{***}) = V(h_0, \pi)$ such that the first transaction has a dual cutoff form when playing π^{***} and such that π^{***} has probability one of the incentive constraint being binding at the first transaction. If there is zero probability of a transaction occurring when playing π^{**} , then simply let $\pi^{***} = \pi^{**}$. If there is a positive probability of a transaction occurring when playing π^{**} , then a symmetric SPE π^{***} satisfying the prescribed conditions can be constructed from π^{**} by appropriately lowering or leaving unchanged the amount of each good transferred at the first transaction.

Let $\pi^{1,***} = \pi^{***}$. Choose any positive integer k . Assume that there exists a symmetric SPE $\pi^{k,***}$ in grim-trigger strategies with $V(h_0, \pi^{k,***}) = V(h_0, \pi)$ such that the n^{th} transaction has a dual cutoff form when playing $\pi^{k,***}$ and such that $\pi^{k,***}$ has probability one of the incentive constraint being binding at the n^{th} transaction, where n can be any positive integer no greater than k . Furthermore, assume that $\pi^{k,***}$ is such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^{k,***}) \geq \tilde{c}_t$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and time $t > t_k(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***})$ at which $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^{k,***}) > 0$. Note that this condition holds for $k = 1$. It will be shown below that there exists a symmetric SPE $\pi^{k+1,***}$ in grim-trigger strategies with $V(h_0, \pi^{k+1,***}) = V(h_0, \pi)$ such that the n^{th} transaction has a dual cutoff form when playing $\pi^{k+1,***}$ and such that $\pi^{k+1,***}$ has probability one of the incentive constraint being bind-

ing at the n^{th} transaction, where n can be any positive integer no greater than $k + 1$. Moreover, $\pi^{k+1,***}$ can be constructed such that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^{k+1,***}) = \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^{k,***})$ for all $t \leq t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***})$. Additionally, $\pi^{k+1,***}$ will be such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^{k+1,***}) \geq \tilde{c}_t$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and time $t > t_{k+1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^{k+1,***})$ with $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^{k+1,***}) > 0$. It will then follow that there exists a symmetric SPE π'' in grim-trigger strategies with $V(h_0, \pi'') = V(h_0, \pi)$ such that the n^{th} transaction has a dual cutoff form when playing π'' and such that π'' has probability one of the incentive constraint being binding at the n^{th} transaction, where n can be any positive integer.

Given the strategy profile $\pi^{k,***}$ described in the preceding paragraph, a strategy profile $\pi^{k+1,***}$ with the aforementioned properties can be constructed by applying the procedure above to each subgame on the equilibrium path immediately following the first transaction when playing $\pi^{k,***}$. Let $\bar{\pi}^{k,***} = \pi^{k,***}$. In particular, the strategy profile $\pi^{k+1,***}$ is derived as follows. Either $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \bar{\pi}^{k,***}) = \infty$ with probability one, or $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \bar{\pi}^{k,***}) < \infty$ with positive probability. If $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \bar{\pi}^{k,***}) = \infty$ with probability one, then simply let $\pi^{k+1,***} = \bar{\pi}^{k,***}$. Otherwise, proceed as below. Let $\{c_n^\dagger\}_{n=1}^k$ be any sequence of costs incurred on the first k transactions for which $\{c(\{c_\tau\}_{\tau \in [0, \infty)}, \bar{\pi}^{k,***}, n)\}_{n=1}^k = \{c_n^\dagger\}_{n=1}^k$ with positive probability. Let $\{f_n^\dagger\}_{n=1}^k$ be any sequence of amounts transferred on the first k transactions for which there is positive probability of $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that both $\{c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***}, n)\}_{n=1}^k = \{c_n^\dagger\}_{n=1}^k$ and $\{\phi^i(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***}, n)\}_{n=1}^k = \{f_n^\dagger\}_{n=1}^k$. Let $\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}$ be any realization of the cost process such that $\{c(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***}, n)\}_{n=1}^k = \{c_n^\dagger\}_{n=1}^k$, $\{\phi^i(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***}, n)\}_{n=1}^k = \{f_n^\dagger\}_{n=1}^k$, and $Y_{u^\dagger}[h_{u^\dagger}(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***}), \pi^{k,***}] = c_k^\dagger$, where $u^\dagger = t_k(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***})$.

Consider a modified game in which the initial value of the cost process is c_k^\dagger instead of c_0 and the original amount of each good is $q - \sum_{n=1}^k f_n^\dagger$ instead of q . Let h'_0 denote the null history in this modified game. Define the symmetric SPE ψ in grim-trigger strategies as follows. At time 0, the agents do not make any transactions. After time 0, the agents play according to strategy profile $\pi^{k,***} \in \Pi^*$, behaving as if the history at time 0 were $h_{u^\dagger}(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***})$. The agents do not make any transfers following a deviation from the path of play described above. Note that ψ is such that $\psi_i(h_t) \geq g_t$ for any history $h_t = (\{g_\tau\}_{\tau \in [0, t]}, \{(b_\tau^1, b_\tau^2)\}_{\tau \in [0, t]})$ up to an arbitrary time t for which $\psi_i(h_t) > 0$. Suppose that a procedure similar to that described in the first nine paragraphs in the proof of the lemma is applied to the strategy profile ψ . By doing

so, it is possible to construct a symmetric strategy profile ψ^* in grim-trigger strategies such that the following statement holds. Either ψ^* is a symmetric SPE in grim-trigger strategies with $V(h'_0, \psi^*) \geq V(h'_0, \psi)$ such that the first transaction has a dual cutoff form when playing ψ^* , or each agent makes a transfer at time 0 with probability one. Note that ψ^* is such that $\psi_i^*(h_t) \geq g_t$ for any history $h_t = (\{g_\tau\}_{\tau \in [0,t]}, \{(b_\tau^1, b_\tau^2)\}_{\tau \in [0,t]})$ up to an arbitrary time t for which $\psi_i^*(h_t) > 0$.

In the former case, a procedure similar to that described in the eleventh through sixteenth paragraphs can be used to construct a symmetric SPE ψ^{**} in grim-trigger strategies with $V(h'_0, \psi^{**}) \geq V(h'_0, \psi)$ such that the first transaction has a dual cutoff form when playing ψ^{**} and such that ψ^{**} has probability one of the incentive constraint being binding at the first transaction. As in the seventeenth paragraph, the strategy profile ψ^{**} can then be adjusted so as to generate a symmetric SPE ψ^{***} in grim-trigger strategies with $V(h'_0, \psi^{***}) = V(h'_0, \psi)$ such that the first transaction has a dual cutoff form when playing ψ^{***} and such that ψ^{***} has probability one of the incentive constraint being binding at the first transaction. If this case holds, then the next three paragraphs should be skipped after constructing the strategy profile ψ^{***} . Note that ψ^{***} is such that $\phi_t^i(h'_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \psi^{***}) \geq \tilde{c}_t$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}$ and time $t > t_1(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \psi^{***})$ at which $\phi_t^i(h'_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \psi^{***}) > 0$.

In the latter case, proceed as follows. Consider another modified game in which the initial value of the cost process is equal to c_0 and the original amount of each good is $q - \sum_{n=1}^k f_n^\dagger$. Define the symmetric SPE $\psi^{*'}$ in grim-trigger strategies as follows. The agents do not make any transactions until the first time $t^{*'}$ that the current value of the cost process is c_k^\dagger . At time $t^{*'}$, each agent transfers the amount $\psi_i^*(h'_0)$. After this transaction, the agents play according to strategy profile ψ^* , behaving as if the history at the time of this transaction were h'_0 and strategy profile ψ^* had always been followed. The agents do not make any transfers following a deviation from the path of play described above. A procedure similar to that described in the eleventh through sixteenth paragraphs can be used to construct a symmetric SPE $\psi^{**'}$ in grim-trigger strategies such that the first transaction has a dual cutoff form when playing $\psi^{**'}$, such that $\psi^{**'}$ has probability one of the incentive constraint being binding at the first transaction, and for which there is probability one of $\{c_\tau\}_{\tau \in [0,\infty)}$ being such that either $t_1(\{c_\tau\}_{\tau \in [0,\infty)}, \psi^{**'}) = \infty$ or both $t_1(\{c_\tau\}_{\tau \in [0,\infty)}, \psi^{**'}) < \infty$ and $V[h_t(\{c_\tau\}_{\tau \in [0,\infty)}, \psi^{**'})] \geq V(h'_0, \psi)$ for $t = t_1(\{c_\tau\}_{\tau \in [0,\infty)}, \psi^{**'})$. Note that $\psi^{*'}$ is such that $\psi_i^{*'}(h_t) \geq g_t$ for any history $h_t = (\{g_\tau\}_{\tau \in [0,t]}, \{(b_\tau^1, b_\tau^2)\}_{\tau \in [0,t]})$ up to

an arbitrary time t for which $\psi_i^{**'}(h_t) > 0$ and that $\psi^{**'}$ is such that $\psi_i^{**'}(h_t) \geq g_t$ for any history $h_t = (\{g_\tau\}_{\tau \in [0,t]}, \{(b_\tau^1, b_\tau^2)\}_{\tau \in [0,t]})$ up to an arbitrary time t for which $\psi_i^{**'}(h_t) > 0$.

Let $\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}$ be any realization of the cost process such that $t_1(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{**'}) < \infty$ and both $Y[h_t(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{**'}], \psi^{**'}] = \tilde{p}_t$ and $V[h_t(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{**'}], \psi^{**'}] \geq V(h'_0, \psi)$ for $t = t_1(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{**'})$. Now consider the original game in which the initial value of the cost process is c_0 and the original amount of each good is q . Define the strategy profile $\pi^{k,***'}$ as follows. The agents play according to strategy profile $\pi^{k,***}$ until the history up to the current time is such that each agent has made $k-1$ transactions in the past, the respective transaction costs incurred on these transactions were $\{c_n^\dagger\}_{n=1}^{k-1}$, the respective amounts transferred on these transactions were $\{f_n^\dagger\}_{n=1}^{k-1}$, and the current value of the cost process is c_k^\dagger . Let $t^{k,***'}$ denote the first such time if such a time exists. Let $f_k^{\dagger'} = f_k^\dagger + \phi^i(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{k,***'}, 1)$. Note that $\phi^i(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{k,***'}, 1) \geq c_k^\dagger$. At time $t^{k,***'}$, each agent transfers the amount $f_k^{\dagger'}$. After time $t^{k,***'}$, the agents play according to strategy profile $\psi^{**'}$, behaving as if the history at this transaction were $h_t(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{**'})$ for $t = t_1(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{**'})$ and as if strategy profile $\psi^{**'}$ had always been followed. The agents do not make any transfers following a deviation from the path of play described above.

Return to the nineteenth paragraph in the proof of the lemma. Redefine the strategy profile $\pi^{k,***}$ as $\pi^{k,***'}$. However, do not redefine $\bar{\pi}^{k,***}$. Choose $\{c_n^\dagger\}_{n=1}^k$ and $\{f_n^\dagger\}_{n=1}^{k-1}$ to be the same as before, but redefine f_k^\dagger as $f_k^{\dagger'}$. Given these redefinitions, let $\{\tilde{g}_\tau\}_{\tau \in [0,\infty)}$ be any realization of the cost process such that $\{c(\{\tilde{g}_\tau\}_{\tau \in [0,\infty)}, \bar{\pi}^{k,***}, n)\}_{n=1}^k = \{c_n^\dagger\}_{n=1}^k$, $\{\phi^i(\{\tilde{g}_\tau\}_{\tau \in [0,\infty)}, \bar{\pi}^{k,***}, n)\}_{n=1}^k = \{f_n^\dagger\}_{n=1}^k$, and $Y_{u^\dagger}[h_{u^\dagger}(\{\tilde{g}_\tau\}_{\tau \in [0,\infty)}, \bar{\pi}^{k,***}), \bar{\pi}^{k,***}] = c_k^\dagger$, where $u^\dagger = t_k(\{\tilde{g}_\tau\}_{\tau \in [0,\infty)}, \bar{\pi}^{k,***})$. Follow the instructions from the twentieth paragraph of the proof.

Recall that $\phi^i(\{\tilde{p}_\tau\}_{\tau \in [0,\infty)}, \psi^{k,***'}, 1) \geq c_k^\dagger$ in the twenty-third paragraph from the proof of the lemma. Hence, it would be the case that $q - \sum_{n=1}^k f_n^\dagger < c_k^\dagger$ after only finitely many repetitions of the procedure in the nineteenth to twenty-fourth paragraphs. In this case, the strategy profile ψ^* described in the twentieth paragraph could not be such that each agent makes a transaction at time zero with probability one. It follows that the procedure in the nineteenth to twenty-fourth paragraphs needs to be repeated only finitely many times in order to generate the strategy profile ψ^{***} described in the twentieth paragraph. In addition, this procedure can be applied to any sequence $\{c_n^\dagger\}_{n=1}^k$ of costs incurred on the first k transactions for which

$\{c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k, **}, n)\}_{n=1}^k = \{c_n^\dagger\}_{n=1}^k$ with positive probability. Given such a sequence $\{c_n^\dagger\}_{n=1}^k$, let $\psi^{***, \{c_n^\dagger\}_{n=1}^k}$ denote the strategy profile ψ^{***} derived by applying this procedure to $\{c_n^\dagger\}_{n=1}^k$.

Now define the strategy profile $\pi^{***, k+1}$ as follows. The agents play according to strategy profile $\bar{\pi}^{***, k}$ until k transactions have occurred. Let $t^{***, k}$ denote the time of the k^{th} transaction if such a time exists, and proceed as follows in this case. Let $\{\tilde{c}_n^\dagger\}_{n=1}^k$ denote the sequence of costs incurred on the first k transactions. After time $t^{***, k}$, the agents play according to strategy profile $\psi^{***, \{\tilde{c}_n^\dagger\}_{n=1}^k}$, behaving as if the history at time $t^{***, k}$ were h_0 and strategy profile $\psi^{***, \{c_n^\dagger\}_{n=1}^k}$ had always been followed. The agents do not make any transfers following a deviation from the path of play described above. As explained in the eighteenth paragraph in the proof of the lemma, it follows that there exists a symmetric SPE π'' in grim-trigger strategies with $V(h_0, \pi'') = V(h_0, \pi)$ such that the n^{th} transaction has a dual cutoff form when playing π'' and such that π'' has probability one of the incentive constraint being binding at the n^{th} transaction, where n can be any positive integer.

In the sixteenth paragraph in the proof, note that the strategy profile π^{**} is such that $\pi_i^{**}(h_t) \geq g_t$ for any history $h_t = (\{g_\tau\}_{\tau \in [0, t]}, \{(b_\tau^1, b_\tau^2)\}_{\tau \in [0, t]})$ up to an arbitrary time t for which $\pi_i^{**}(h_t) > 0$. Similarly, the strategy profile ψ^{**} in the twenty-first paragraph satisfies $\psi_i^{**}(h_t) \geq g_t$ for any history $h_t = (\{g_\tau\}_{\tau \in [0, t]}, \{(b_\tau^1, b_\tau^2)\}_{\tau \in [0, t]})$ up to an arbitrary time t for which $\psi_i^{**}(h_t) > 0$. Given the procedure used to construct π'' , it follows that $\sum_{n=k+1}^\infty \phi^i(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi'', n) \leq q - \sum_{n=1}^k c(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi'', n)$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and any positive integer k , where we define $\phi^i(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi'', n) = 0$ and $c(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi'', n) = 0$ if $t_n(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi'') = \infty$.

We finally explain how to construct a stationary symmetric SPE $\pi' \in \Pi^*$ in grim-trigger strategies with $V(h_0, \pi') \geq V(h_0, \pi)$ for which the cost incurred between any two consecutive transactions is decreasing. Recall the strategy profile π'' constructed above. To simplify the exposition, it is helpful to define the symmetric SPE π''' in grim-trigger strategies as follows. With probability one, $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi''') = \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi'')$ for all $t \geq 0$. Moreover, the strategy profile π''' satisfies $\phi_{t_1}^i(h_0, \{\tilde{c}_\tau^1\}_{\tau \in (0, \infty)}, \pi''') = \phi_{t_2}^i(h_0, \{\tilde{c}_\tau^2\}_{\tau \in (0, \infty)}, \pi''')$ for any pair of cost realizations $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ and $\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}$ along with times $t_1 \geq 0$ and $t_2 \geq 0$ such that $\tilde{c}_{t_1}^1 = \tilde{c}_{t_2}^2$ and $\hat{\sigma}_{t_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi''') = \hat{\sigma}_{t_2}^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \pi''')$ and such that $c(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi''', k) = c(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \pi''', k)$ and $\sigma^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi''', k) = \sigma^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \pi''', k)$ for any integer $k \geq 0$ satisfying both $t_k(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi''') < t_1$ and $t_k(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \pi''') < t_2$.

Consider any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0,t]}$ up to an arbitrary time t . Let $R(\{\tilde{c}_\tau\}_{\tau \in [0,t]})$ denote the set consisting of every cost level c such that the following holds. Given that $\{\tilde{c}_\tau\}_{\tau \in [0,t]}$ is the realization of the cost process up to time t , there is positive conditional probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in (t,\infty)}$ after time t is such that there exists $u > t$ with $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \pi''') > 0$ and $\tilde{c}_u = c$. The strategy profile π' is defined as follows. If π''' is such that there is zero probability of a transaction occurring, then simply let $\pi' = \pi'''$. In the case where π''' is such that a transaction occurs with positive probability, the strategy profile π' is constructed using the procedure below.

The first transaction when playing strategy profile π' is specified as follows. Let r_1 denote the supremum of the set $R(c_0)$. If $r_1 \in R(c_0)$, then let $\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}$ be any realization of the cost process for which there exists a time \tilde{t}_1 such that $\phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau^1\}_{\tau \in (0,\infty)}, \pi''') > 0$ and $\tilde{c}_{\tilde{t}_1}^1 = r_1$ for $t = \tilde{t}_1$ and such that every element of the set $R(\{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]})$ is less than r_1 . For the case where $r_1 \in R(c_0)$, the strategy profile π' requires each agent to make a transaction at the first time that the cost reaches r_1 , and each agent should transfer the amount $q - \sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}, \pi''')$ on this transaction.

If $r_1 \notin R(c_0)$, then let \tilde{g}_1 denote the unique cost level g for which there is positive probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}$ is such that there exists a time t satisfying $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \pi''') > 0$, $\tilde{c}_t = g$, and $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \pi''') = 0$ for all $u \leq t$. Choose any cost level $\tilde{r}_1 \in [\tilde{g}_1, r_1)$ for which there exists a cost realization $\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}$ along with a time \tilde{t}_1 such that $\phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau^1\}_{\tau \in (0,\infty)}, \pi''') > 0$ and $\tilde{c}_{\tilde{t}_1}^1 = \tilde{r}_1$ for $t = \tilde{t}_1$ and such that every element of the set $R(\{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]})$ is less than \tilde{r}_1 . For the case where $r_1 \notin R(c_0)$, the strategy profile π' requires each agent to make a transaction at the first time that the cost reaches \tilde{r}_1 , and each agent should transfer the amount $q - \sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}, \pi''')$ on this transaction.

In each case, such a choice of $\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}$ and \tilde{t}_1 is possible because $\sum_{n=k+1}^{\infty} \phi^i(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi''', n) \leq q - \sum_{n=1}^k c(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi''', n)$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}$ and any positive integer k , where $\phi^i(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi''', n) = 0$ and $c(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi''', n) = 0$ if $t_n(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi''') = \infty$. Otherwise, if no such $\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}$ and \tilde{t}_1 existed, then there would exist k such that $\sum_{n=k+1}^{\infty} \phi^i(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi''', n) < 0$ for some $\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}$, which is impossible.

The second transaction when playing strategy profile π' is specified as follows. Let r_2 denote the supremum of the set $R(\{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]})$. If $r_2 \in R(\{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]})$, then let $\{\tilde{c}_\tau^2\}_{\tau \in [0,\tilde{t}_2]}$ be any realization of the cost process with $\{\tilde{c}_\tau^2\}_{\tau \in [0,\tilde{t}_1]} = \{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]}$ for which there exists a time $\tilde{t}_2 > \tilde{t}_1$ such that $\phi_{\tilde{t}_2}^i(h_0, \{\tilde{c}_\tau^2\}_{\tau \in (0,\infty)}, \pi''') > 0$ and $\tilde{c}_{\tilde{t}_2}^2 = r_2$ for

$t = \tilde{t}_2$ and such that every element of the set $R(\{\tilde{c}_\tau^2\}_{\tau \in [0, \tilde{t}_2]})$ is less than r_2 . For the case where $r_2 \in R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$, the strategy profile π' requires each agent to make a transaction at the first time that the cost reaches r_2 , and each agent should transfer the amount $\sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi''') - \sigma_{\tilde{t}_2}^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \pi''')$ on this transaction.

If $r_2 \notin R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$, then let \tilde{g}_2 denote the unique cost level g such that the following holds. Given that the realization of the cost process $\{\tilde{c}_\tau\}_{\tau \in [0, \tilde{t}_1]}$ up to time \tilde{t}_1 is such that $\{\tilde{c}_\tau\}_{\tau \in [0, \tilde{t}_1]} = \{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]}$, there is positive conditional probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in (\tilde{t}_1, \infty)}$ after time \tilde{t}_1 is such that there exists a time t satisfying $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi''') > 0$, $\tilde{c}_t = g$, and $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi''') = 0$ for all $u \in (\tilde{t}_1, t)$. Choose any cost level $\tilde{r}_2 \in [\tilde{g}_2, r_2)$ for which there exists a cost realization $\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}$ along with a time $\tilde{t}_2 > \tilde{t}_1$ such that $\{\tilde{c}_\tau^2\}_{\tau \in [0, \tilde{t}_1]} = \{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]}$, such that $\phi_t^i(h_0, \{\tilde{c}_\tau^2\}_{\tau \in (0, \infty)}, \pi''') > 0$ and $\tilde{c}_t^2 = \tilde{r}_2$ for $t = \tilde{t}_2$, and such that every element of the set $R(\{\tilde{c}_\tau^2\}_{\tau \in [0, \tilde{t}_2]})$ is less than \tilde{r}_2 . For the case where $r_2 \notin R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$, the strategy profile π' requires each agent to make a transaction at the first time that the cost reaches \tilde{r}_2 , and each agent should transfer the amount $\sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi''') - \sigma_{\tilde{t}_2}^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \pi''')$ on this transaction. In each case, such a choice of $\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}$ and \tilde{t}_2 is possible because of the property described above.

Proceeding in this way, the k^{th} transaction when playing strategy profile π' is specified as follows. Let r_k denote the supremum of the set $R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]})$. If $r_k \in R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]})$, then let $\{\tilde{c}_\tau^k\}_{\tau \in [0, \tilde{t}_k]}$ be any realization of the cost process with $\{\tilde{c}_\tau^k\}_{\tau \in [0, \tilde{t}_{k-1}]} = \{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]}$ for which there exists a time $\tilde{t}_k > \tilde{t}_{k-1}$ such that $\phi_t^i(h_0, \{\tilde{c}_\tau^k\}_{\tau \in (0, \infty)}, \pi''') > 0$ and $\tilde{c}_t^k = r_k$ for $t = \tilde{t}_k$ and such that every element of the set $R(\{\tilde{c}_\tau^k\}_{\tau \in [0, \tilde{t}_k]})$ is less than r_k . For the case where $r_k \in R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]})$, the strategy profile π' requires each agent to make a transaction at the first time that the cost reaches r_k , and each agent should transfer the amount $\sigma_{\tilde{t}_{k-1}}^{i,i}(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \infty)}, \pi''') - \sigma_{\tilde{t}_k}^{i,i}(\{\tilde{c}_\tau^k\}_{\tau \in [0, \infty)}, \pi''')$ on this transaction.

If $r_k \notin R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]})$, then let \tilde{g}_k denote the unique cost level g such that the following holds. Given that the realization of the cost process $\{\tilde{c}_\tau\}_{\tau \in [0, \tilde{t}_{k-1}]}$ up to time \tilde{t}_{k-1} is such that $\{\tilde{c}_\tau\}_{\tau \in [0, \tilde{t}_{k-1}]} = \{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]}$, there is positive conditional probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in (\tilde{t}_{k-1}, \infty)}$ after time \tilde{t}_{k-1} is such that there exists a time t satisfying $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi''') > 0$, $\tilde{c}_t = g$, and $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi''') = 0$ for all $u \in (\tilde{t}_{k-1}, t)$. Choose any cost level $\tilde{r}_k \in [\tilde{g}_k, r_k)$ for which there exists a cost realization $\{\tilde{c}_\tau^k\}_{\tau \in [0, \infty)}$ along with a time $\tilde{t}_k > \tilde{t}_{k-1}$ such that $\{\tilde{c}_\tau^k\}_{\tau \in [0, \tilde{t}_{k-1}]} = \{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]}$, such that $\phi_t^i(h_0, \{\tilde{c}_\tau^k\}_{\tau \in (0, \infty)}, \pi''') > 0$ and $\tilde{c}_t^k = \tilde{r}_k$ for $t = \tilde{t}_k$, and

such that every element of the set $R(\{\tilde{c}_\tau^k\}_{\tau \in [0, \tilde{t}_k]})$ is less than \tilde{r}_k . For the case where $r_k \notin R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \tilde{t}_{k-1}]})$, the strategy profile π' requires each agent to make a transaction at the first time that the cost reaches \tilde{r}_k , and each agent should transfer the amount $\sigma_{\tilde{t}_{k-1}}^{i,i}(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0, \infty)}, \pi''') - \sigma_{\tilde{t}_k}^{i,i}(\{\tilde{c}_\tau^k\}_{\tau \in [0, \infty)}, \pi''')$ on this transaction. In each case, such a choice of $\{\tilde{c}_\tau^k\}_{\tau \in [0, \infty)}$ and \tilde{t}_k is possible because of the property described above.

The agents do not make any transfers following a deviation from the path of play described above. It is straightforward to confirm that π' is a stationary symmetric SPE in grim-trigger strategies with $V(h_0, \pi') \geq V(h_0, \pi)$ for which the cost incurred between any two consecutive transactions is decreasing. \square

The next result shows that given a stationary symmetric SPE π for which there is positive probability of the incentive constraint being slack at some transaction, one can find a stationary symmetric SPE π' in which each agent receives a higher expected payoff and there is probability one of the incentive constraint being binding at every transaction.

Lemma 2. *Given a stationary symmetric SPE $\pi \in \Pi^*$ in grim-trigger strategies for which there is a positive probability of the incentive constraint being slack at some transaction, there exists a stationary symmetric SPE $\pi' \in \Pi^*$ in grim-trigger strategies with $V(h_0, \pi') > V(h_0, \pi)$ for which there is probability one of the incentive constraint being binding at every transaction.*

Proof. Consider any stationary symmetric SPE $\pi \in \Pi^*$ in grim-trigger strategies for which there is a positive probability of the incentive constraint being slack at some transaction. We first explain how to construct a symmetric SPE $\tilde{\pi} \in \Pi^*$ in grim-trigger strategies such that $V(h_0, \tilde{\pi}) > V(h_0, \pi)$.

It is helpful to introduce some notation. Let \hat{t}_1 and \hat{t}_2 be any two nonnegative real numbers such that $\hat{t}_1 \leq \hat{t}_2$. Let $\hat{\pi} \in \Pi^*$ be a symmetric SPE in grim-trigger strategies. Given a realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ of the cost process, let $\Sigma[h_{\hat{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}); \{\tilde{c}_\tau\}_{\tau \in [\hat{t}_1, \hat{t}_2]}; \hat{t}_1, \hat{t}_2; \hat{\pi}]$ be the sum of the transfers that an agent would make between times \hat{t}_1 and \hat{t}_2 inclusive if $h_{\hat{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \hat{\pi})$ is the history up to time \hat{t}_1 , the cost path $\{\tilde{c}_\tau\}_{\tau \in [\hat{t}_1, \hat{t}_2]}$ is realized between times \hat{t}_1 and \hat{t}_2 , and the strategy profile $\hat{\pi}$ is played by the agents.

The following is how $\tilde{\pi}$ is defined in this case. The agents play according to strategy profile π until the first time t for which the realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ of the cost process is such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi) > 0$ and $Y[h_t(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi), \pi] > \tilde{c}_t$. If no

such time t exists, then the agents simply follow strategy profile π . Otherwise, let \tilde{t}_1 denote the first time t that this condition holds. At time \tilde{t}_1 , each agent transfers the amount $\phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) + \zeta[h_t(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)]$, where $\zeta[h_t(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)]$ is a positive real number no greater than $\sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)$ that can depend on $h_t(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)$. After time \tilde{t}_1 , the agents do not make any transfers at any time $t > \tilde{t}_1$ such that the realization of the cost process $\{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, t]}$ between times \tilde{t}_1 and t satisfies $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) = 0$ or $\Sigma[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi); \{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, t]}; \tilde{t}_1, t; \pi] - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) - \zeta[h_t(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)] < \tilde{c}_t$.

Let \tilde{t}_2 denote the first time t such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) > 0$ and $\Sigma[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi); \{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, t]}; \tilde{t}_1, t; \pi] - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) - \zeta[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)] \geq \tilde{c}_t$. If such a time t does not exist, then the agents do not make any further transactions. Otherwise, proceed as follows. At time \tilde{t}_2 , each agent transfers an amount equal to the lesser of $\phi_{\tilde{t}_2}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi)$ and $\Sigma[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi); \{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, \tilde{t}_2]}; \tilde{t}_1, \tilde{t}_2; \pi] - \phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) - \zeta[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)]$. Thereafter, the agents play according to strategy profile π , behaving as if the history up to time \tilde{t}_1 were $h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)$, the realization of the cost process between times \tilde{t}_1 and \tilde{t}_2 were $\{\tilde{c}_\tau\}_{\tau \in [\tilde{t}_1, \tilde{t}_2]}$, and strategy profile π were followed up to and including time \tilde{t}_2 . The agents do not make any transfers following a deviation from the path of play specified above.

If each $\zeta[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)] > 0$ is chosen to be sufficiently small, then $\tilde{\pi}$ is a symmetric SPE in grim-trigger strategies such that $V(h_0, \tilde{\pi}) > V(h_0, \pi)$. Such a choice of $\zeta[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)]$ is possible because $V[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi), \tilde{\pi}]$ is increasing in $\zeta[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)]$, and $Y[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi), \tilde{\pi}]$ approaches $Y[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi), \pi]$ as $\zeta[h_{\tilde{t}_1}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi)]$ goes to 0.

Having constructed a symmetric SPE $\tilde{\pi}$ in grim-trigger strategies such that $V(h_0, \tilde{\pi}) > V(h_0, \pi)$, lemma 1 shows that there exists a stationary symmetric SPE $\tilde{\tilde{\pi}}$ in grim-trigger strategies with $V(h_0, \tilde{\tilde{\pi}}) \geq V(h_0, \tilde{\pi})$ for which the cost incurred between any two consecutive transactions is decreasing. To simplify the exposition, it is helpful to define the stationary symmetric SPE π^* in grim-trigger strategies as follows. With probability one, $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^*) = \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \tilde{\tilde{\pi}})$ for all $t \geq 0$. Moreover, the strategy profile π^* satisfies $\phi_{t_1}^i(h_0, \{\tilde{c}_\tau^1\}_{\tau \in (0, \infty)}, \pi^*) = \phi_{t_2}^i(h_0, \{\tilde{c}_\tau^2\}_{\tau \in (0, \infty)}, \pi^*)$ for any pair of cost realizations $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ and $\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}$ along with times $t_1 \geq 0$ and $t_2 \geq 0$ such that $\tilde{c}_{t_1}^1 = \tilde{c}_{t_2}^2$ and $\hat{\sigma}_{t_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi^*) = \hat{\sigma}_{t_2}^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \pi^*)$. We next explain how to construct a stationary symmetric SPE $\pi' \in \Pi^*$ in grim-trigger strategies with $V(h_0, \pi') = V(h_0, \pi^*)$ for which there is probability one of the incentive

constraint being binding on every transaction and the cost incurred between any two consecutive transactions is decreasing.

It is helpful to introduce some terminology. Given a symmetric SPE $\hat{\pi} \in \Pi^*$ as well as any positive integer k , let \hat{p} be the probability that $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}) < \infty$ and there exists a time t with $t_{k-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}) < t < t_{k+1}(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi})$ for which $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \hat{\pi}) > 0$ and $Y[h_t(\{c_\tau\}_{\tau \in [0, \infty)}, \hat{\pi}), \hat{\pi}] > c_t$. If $\hat{p} > 0$, then $\hat{\pi}$ is said to have a positive probability of the incentive constraint being slack at transaction k . If $\hat{p} = 0$, then $\hat{\pi}$ is said to have probability one of the incentive constraint being binding at transaction k .

Let the symmetric strategy profile π^{**} be defined as follows. If π^* is such that there is probability one of the incentive constraint being binding at the first transaction, then simply define $\pi^{**} = \pi^*$. If π^* has a positive probability of the incentive constraint being slack at the first transaction, then the strategy profile π^{**} is constructed as follows. The agents play according to strategy profile π^* until the first time t for which the realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ of the cost process is such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) > 0$ and $Y[h_t(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^*] > \tilde{c}_t$. If no such time t exists, then the agents simply follow strategy profile π^* . Otherwise, let t_1^* denote the first time t that this condition holds. At time t_1^* , each agent transfers the amount $\phi_{t_1^*}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) + \zeta$, where ζ is a positive real number no greater than $s_0^{i,i} - \phi_{t_1^*}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*)$. After time t_1^* , the agents do not make any transfers at any time $t > t_1^*$ such that the realization of the cost process $\{\tilde{c}_\tau\}_{\tau \in [t_1^*, t]}$ between times t_1^* and t satisfies $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) = 0$ or $\Sigma[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*); \{\tilde{c}_\tau\}_{\tau \in [t_1^*, t]}; t_1^*, t; \pi^*] - \phi_{t_1^*}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) - \zeta < \tilde{c}_t$.

Let t_2^* denote the first time t such that $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) > 0$ and $\Sigma[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*); \{\tilde{c}_\tau\}_{\tau \in [t_1^*, t]}; t_1^*, t; \pi^*] - \phi_{t_1^*}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) - \zeta \geq \tilde{c}_t$. If such a time t does not exist, then the agents do not make any further transactions. Otherwise, proceed as follows. At time t_2^* , each agent transfers an amount equal to the lesser of $\phi_{t_2^*}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*)$ and $\Sigma[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*); \{\tilde{c}_\tau\}_{\tau \in [t_1^*, t_2^*]}; t_1^*, t_2^*; \pi^*] - \phi_{t_1^*}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*) - \zeta$. Thereafter, the agents play according to strategy profile π^* , behaving as if the history up to time t_1^* were $h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*)$, the realization of the cost process between times t_1^* and t_2^* were $\{\tilde{c}_\tau\}_{\tau \in [t_1^*, t_2^*]}$, and strategy profile π^* were followed up to and including time t_2^* . The agents do not make any transfers following a deviation from the path of play specified above.

If each ζ is chosen appropriately, then π^{**} is a symmetric SPE in grim-trigger strategies with $V(h_0, \pi^{**}) > V(h_0, \pi^*)$, and π^{**} has probability one of the incentive

constraint being binding at the first transaction. Such a choice of ζ is possible because $V[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ is increasing in ζ , $Y[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ approaches $Y[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^*]$ as ζ goes to 0, $Y[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ equals 0 for $\zeta = s_0^{i,i} - \phi_{t_1^*}^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi^*)$, and $Y[h_{t_1^*}(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \pi^*), \pi^{**}]$ is a continuous function of ζ .

We now observe that it is possible to construct a stationary symmetric SPE π^{***} in grim-trigger strategies with $V(h_0, \pi^{***}) = V(h_0, \pi^*)$ such that π^{***} has probability one of the incentive constraint being binding at the first transaction and the cost incurred between any two consecutive transactions is decreasing. If there is zero probability of a transaction occurring when playing π^{**} , then simply let $\pi^{***} = \pi^{**}$. If there is a positive probability of a transaction occurring when playing π^{**} , then a symmetric SPE π^{***} satisfying the prescribed conditions can be constructed from π^{**} by appropriately lowering or leaving unchanged the amount of each good transferred at the first transaction.

Let $\pi^{1,***} = \pi^{***}$. Choose any positive integer k . Assume that there exists a stationary symmetric SPE $\pi^{k,***}$ in grim-trigger strategies with $V(h_0, \pi^{k,***}) = V(h_0, \pi^*)$ such that the cost incurred between any two consecutive transactions is decreasing and such that $\pi^{k,***}$ has probability one of the incentive constraint being binding at the n^{th} transaction, where n can be any positive integer no greater than k . Note that this property holds for $k = 1$. It will be shown below that there exists a stationary symmetric SPE $\pi^{k+1,***}$ in grim-trigger strategies with $V(h_0, \pi^{k+1,***}) = V(h_0, \pi^*)$ such that the cost incurred between any two consecutive transactions is decreasing and such that $\pi^{k+1,***}$ has probability one of the incentive constraint being binding at the n^{th} transaction, where n can be any positive integer no greater than $k + 1$. Moreover, $\pi^{k+1,***}$ can be constructed such that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^{k+1,***}) = \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi^{k,***})$ for all $t \leq t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k,***})$. It will then follow that there exists a stationary symmetric SPE π' in grim-trigger strategies with $V(h_0, \pi') > V(h_0, \pi)$ such that the cost incurred between any two consecutive transactions is decreasing and such that π' has probability one of the incentive constraint being binding at the n^{th} transaction, where n can be any positive integer.

Given the strategy profile $\pi^{k,***}$ described in the preceding paragraph, a strategy profile $\pi^{k+1,***}$ with the aforementioned properties can be constructed by applying the procedure above to each subgame on the equilibrium path immediately following the first transaction when playing $\pi^{k,***}$. In particular, the strategy profile $\pi^{k+1,***}$

is derived as follows. Either $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k, ***}) = \infty$ with probability one, or $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k, ***}) < \infty$ with positive probability. If $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{k, ***}) = \infty$ with probability one, then simply let $\pi^{k+1, ***} = \pi^{k, ***}$. Otherwise, proceed as below. Let $\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}$ be any realization of the cost process for which there exists a time $\tilde{t} < \infty$ such that $t_k(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k, ***}) = \tilde{t}$.

Consider a modified game in which the initial value of the cost process is $\tilde{g}_{\tilde{t}}$ instead of c_0 and the original amount of each good is $\sigma_{\tilde{t}}^{i, i}(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k, ***})$ instead of q . Let h'_0 denote the null history in this modified game. Define the symmetric SPE ψ^* in grim-trigger strategies as follows. At time 0, the agents do not make any transactions. After time 0, the agents play according to strategy profile $\pi^{k, ***} \in \Pi^*$, behaving as if the history at time 0 were $h_{\tilde{t}}(\{\tilde{g}_\tau\}_{\tau \in [0, \infty)}, \pi^{k, ***})$. The agents do not make any transfers following a deviation from the path of play described above. A procedure similar to that described above can be used to construct a stationary symmetric strategy profile ψ^{***} in grim-trigger strategies with $V(h'_0, \psi^{***}) = V(h'_0, \psi^*)$ such that ψ^{***} has probability one of the incentive constraint being binding at the first transaction and the cost incurred between any two consecutive transactions is decreasing.

Now define the strategy profile $\pi^{***, k+1}$ as follows. The agents play according to strategy profile $\pi^{***, k}$ until k transactions have occurred. Let $t^{***, k}$ denote the time of the k^{th} transaction if such a time exists, and proceed as follows in this case. After time $t^{***, k}$, the agents play according to strategy profile ψ^{***} , behaving as if the history at time $t^{***, k}$ were h'_0 and strategy profile ψ^{***} had always been followed. The agents do not make any transfers following a deviation from the path of play described above. As previously explained, it follows that there exists a stationary symmetric SPE π' in grim-trigger strategies with $V(h_0, \pi') = V(h_0, \pi^*)$ such that the cost incurred between any two consecutive transactions is decreasing and such that π' has probability one of the incentive constraint being binding at the n^{th} transaction, where n can be any positive integer. \square

The next result shows that given a stationary symmetric SPE π for which there is probability one of the incentive constraint being binding at every transaction and there is a positive probability of the cost incurred being nondecreasing between some two consecutive transactions, one can find a stationary symmetric SPE π' in which each agent receives a higher expected payoff, there is probability one of the incentive constraint being binding at every transaction, and there is probability one of the cost incurred being decreasing between any two consecutive transactions.

Lemma 3. *Given a stationary symmetric SPE $\pi \in \Pi^*$ in grim-trigger strategies for which there is probability one of the incentive constraint being binding at every transaction and there is a positive probability of the cost incurred being nondecreasing between some two consecutive transactions, there exists a stationary symmetric SPE $\pi' \in \Pi^*$ in grim-trigger strategies with $V(h_0, \pi') > V(h_0, \pi)$ for which there is probability one of the incentive constraint being binding at every transaction and there is probability one of the cost incurred being decreasing between any two consecutive transactions.*

Proof. Let $\pi \in \Pi^*$ be any stationary symmetric SPE in grim-trigger strategies for which there is probability one of the incentive constraint being binding at every transaction and there is a positive probability of the cost incurred being nondecreasing between some two consecutive transactions. We start by explaining how to construct a stationary symmetric SPE $\pi^* \in \Pi^*$ in grim-trigger strategies with $V(h_0, \pi^*) > V(h_0, \pi)$ for which the cost incurred between any two consecutive transactions is decreasing.

Suppose first that there is positive probability that the cost realization $\{\tilde{c}_t\}_{t \in [0, \infty)}$ is such that there exists a time t for which $0 < \phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) < \tilde{c}_t$. In this case, define the symmetric SPE π^0 in grim-trigger strategies as follows. For any history $k_u = (\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u]})$ up to an arbitrary time u , let $\pi_i^0(k_u) = \pi_i(k_u)$ if $\pi_i(k_u) \geq g_u$, and let $\pi_i^0(k_u) = 0$ if $\pi_i(k_u) < g_u$, where $i \in \{1, 2\}$. Note that $V(h_0, \pi^0) > V(h_0, \pi)$. In addition, lemma 1 implies that there exists a stationary symmetric SPE π^* in grim-trigger strategies with $V(h_0, \pi^*) \geq V(h_0, \pi^0)$ for which the cost incurred between any two consecutive transactions is decreasing.

Suppose next that there is zero probability that the cost realization $\{\tilde{c}_t\}_{t \in [0, \infty)}$ is such that there exists a time t for which $0 < \phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \pi) < \tilde{c}_t$. To simplify the exposition, it is helpful to define the stationary symmetric SPE $\tilde{\pi}$ in grim-trigger strategies as follows. With probability one, $\phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \tilde{\pi}) = \phi_t^i(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi)$ for all $t \geq 0$. Moreover, the strategy profile $\tilde{\pi}$ satisfies $\phi_{t_1}^i(h_0, \{\tilde{c}_\tau^1\}_{\tau \in (0, \infty)}, \tilde{\pi}) = \phi_{t_2}^i(h_0, \{\tilde{c}_\tau^2\}_{\tau \in (0, \infty)}, \tilde{\pi})$ for any pair of cost realizations $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ and $\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}$ along with times $t_1 \geq 0$ and $t_2 \geq 0$ such that $\tilde{c}_{t_1}^1 = \tilde{c}_{t_2}^2$ and $\hat{\sigma}_{t_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \tilde{\pi}) = \hat{\sigma}_{t_2}^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0, \infty)}, \tilde{\pi})$. Note that $\tilde{\pi}$ satisfies the condition $\sum_{n=k+1}^{\infty} \phi^i(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n) \leq q - \sum_{n=1}^k c(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n)$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and any positive integer k , where we define $\phi^i(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n) = 0$ and $c(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n) = 0$ if $t_n(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}) = \infty$. In this case, the strategy profile π^* is constructed using the procedure below.

Consider any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, t]}$ up to an arbitrary time t . Let $R(\{\tilde{c}_\tau\}_{\tau \in [0, t]})$ denote the set consisting of every cost level c such that the following holds. Given that $\{\tilde{c}_\tau\}_{\tau \in [0, t]}$ is the realization of the cost process up to time t , there is positive conditional probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in (t, \infty)}$ after time t is such that there exists $u > t$ with $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \tilde{\pi}) > 0$ and $\tilde{c}_u = c$.

The first transaction when playing strategy profile π^* is specified as follows. Let r_1 denote the supremum of the set $R(c_0)$. If $r_1 \in R(c_0)$, then let $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ be any realization of the cost process for which there exists a time \tilde{t}_1 such that $\phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau^1\}_{\tau \in (0, \infty)}, \tilde{\pi}) > 0$ and $\tilde{c}_{\tilde{t}_1}^1 = r_1$ for $t = \tilde{t}_1$ and such that every element of the set $R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$ is less than r_1 . For the case where $r_1 \in R(c_0)$, the strategy profile π^* requires each agent to make a transaction at the first time that the cost reaches r_1 , and each agent should transfer the amount $q - \sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \tilde{\pi})$ on this transaction.

If $r_1 \notin R(c_0)$, then let \tilde{g}_1 denote the unique cost level g for which there is positive probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ is such that there exists a time t satisfying $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \tilde{\pi}) > 0$, $\tilde{c}_t = g$, and $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0, \infty)}, \tilde{\pi}) = 0$ for all $u \leq t$. Choose any cost level $\tilde{r}_1 \in [\tilde{g}_1, r_1)$ for which there exists a cost realization $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ along with a time \tilde{t}_1 such that $\phi_{\tilde{t}_1}^i(h_0, \{\tilde{c}_\tau^1\}_{\tau \in (0, \infty)}, \tilde{\pi}) > 0$ and $\tilde{c}_{\tilde{t}_1}^1 = \tilde{r}_1$ for $t = \tilde{t}_1$ and such that every element of the set $R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$ is less than \tilde{r}_1 . For the case where $r_1 \notin R(c_0)$, the strategy profile π^* requires each agent to make a transaction at the first time that the cost reaches \tilde{r}_1 , and each agent should transfer the amount $q - \sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \tilde{\pi})$ on this transaction.

In each case, such a choice of $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ and \tilde{t}_1 is possible because $\sum_{n=k+1}^{\infty} \phi^i(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n) \leq q - \sum_{n=1}^k c(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n)$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$ and any positive integer k , where $\phi^i(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n) = 0$ and $c(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n) = 0$ if $t_n(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}) = \infty$. Otherwise, if no such $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ and \tilde{t}_1 existed, then there would exist k such that $\sum_{n=k+1}^{\infty} \phi^i(\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}, \tilde{\pi}, n) < 0$ for some $\{\tilde{c}_\tau\}_{\tau \in [0, \infty)}$, which is impossible.

The second transaction when playing strategy profile π^* is specified as follows. Let r_2 denote the supremum of the set $R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$. If $r_2 \in R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$, then let $\{\tilde{c}_\tau^2\}_{\tau \in [0, \tilde{t}_2]}$ be any realization of the cost process with $\{\tilde{c}_\tau^2\}_{\tau \in [0, \tilde{t}_1]} = \{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]}$ for which there exists a time $\tilde{t}_2 > \tilde{t}_1$ such that $\phi_{\tilde{t}_2}^i(h_0, \{\tilde{c}_\tau^2\}_{\tau \in (0, \infty)}, \tilde{\pi}) > 0$ and $\tilde{c}_{\tilde{t}_2}^2 = r_2$ for $t = \tilde{t}_2$ and such that every element of the set $R(\{\tilde{c}_\tau^2\}_{\tau \in [0, \tilde{t}_2]})$ is less than r_2 . For the case where $r_2 \in R(\{\tilde{c}_\tau^1\}_{\tau \in [0, \tilde{t}_1]})$, the strategy profile π^* requires each agent to make a transaction at the first time that the cost reaches r_2 , and each agent should transfer

the amount $\sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}, \tilde{\pi}) - \sigma_{\tilde{t}_2}^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0,\infty)}, \tilde{\pi})$ on this transaction.

If $r_2 \notin R(\{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]})$, then let \tilde{g}_2 denote the unique cost level g such that the following holds. Given that the realization of the cost process $\{\tilde{c}_\tau\}_{\tau \in [0,\tilde{t}_1]}$ up to time \tilde{t}_1 is such that $\{\tilde{c}_\tau\}_{\tau \in [0,\tilde{t}_1]} = \{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]}$, there is positive conditional probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in (\tilde{t}_1,\infty)}$ after time \tilde{t}_1 is such that there exists a time t satisfying $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \tilde{\pi}) > 0$, $\tilde{c}_t = g$, and $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \tilde{\pi}) = 0$ for all $u \in (\tilde{t}_1, t)$. Choose any cost level $\tilde{r}_2 \in [\tilde{g}_2, r_2)$ for which there exists a cost realization $\{\tilde{c}_\tau^2\}_{\tau \in [0,\infty)}$ along with a time $\tilde{t}_2 > \tilde{t}_1$ such that $\{\tilde{c}_\tau^2\}_{\tau \in [0,\tilde{t}_1]} = \{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]}$, such that $\phi_t^i(h_0, \{\tilde{c}_\tau^2\}_{\tau \in (0,\infty)}, \tilde{\pi}) > 0$ and $\tilde{c}_t^2 = \tilde{r}_2$ for $t = \tilde{t}_2$, and such that every element of the set $R(\{\tilde{c}_\tau^2\}_{\tau \in [0,\tilde{t}_2]})$ is less than \tilde{r}_2 . For the case where $r_2 \notin R(\{\tilde{c}_\tau^1\}_{\tau \in [0,\tilde{t}_1]})$, the strategy profile π^* requires each agent to make a transaction at the first time that the cost reaches \tilde{r}_2 , and each agent should transfer the amount $\sigma_{\tilde{t}_1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0,\infty)}, \tilde{\pi}) - \sigma_{\tilde{t}_2}^{i,i}(\{\tilde{c}_\tau^2\}_{\tau \in [0,\infty)}, \tilde{\pi})$ on this transaction. In each case, such a choice of $\{\tilde{c}_\tau^2\}_{\tau \in [0,\infty)}$ and \tilde{t}_2 is possible because of the property described above.

Proceeding in this way, the k^{th} transaction when playing strategy profile π^* is specified as follows. Let r_k denote the supremum of the set $R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]})$. If $r_k \in R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]})$, then let $\{\tilde{c}_\tau^k\}_{\tau \in [0,\tilde{t}_k]}$ be any realization of the cost process with $\{\tilde{c}_\tau^k\}_{\tau \in [0,\tilde{t}_{k-1}]} = \{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]}$ for which there exists a time $\tilde{t}_k > \tilde{t}_{k-1}$ such that $\phi_t^i(h_0, \{\tilde{c}_\tau^k\}_{\tau \in (0,\infty)}, \tilde{\pi}) > 0$ and $\tilde{c}_t^k = r_k$ for $t = \tilde{t}_k$ and such that every element of the set $R(\{\tilde{c}_\tau^k\}_{\tau \in [0,\tilde{t}_k]})$ is less than r_k . For the case where $r_k \notin R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]})$, the strategy profile π^* requires each agent to make a transaction at the first time that the cost reaches r_k , and each agent should transfer the amount $\sigma_{\tilde{t}_{k-1}}^{i,i}(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\infty)}, \tilde{\pi}) - \sigma_{\tilde{t}_k}^{i,i}(\{\tilde{c}_\tau^k\}_{\tau \in [0,\infty)}, \tilde{\pi})$ on this transaction.

If $r_k \notin R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]})$, then let \tilde{g}_k denote the unique cost level g such that the following holds. Given that the realization of the cost process $\{\tilde{c}_\tau\}_{\tau \in [0,\tilde{t}_{k-1}]}$ up to time \tilde{t}_{k-1} is such that $\{\tilde{c}_\tau\}_{\tau \in [0,\tilde{t}_{k-1}]} = \{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]}$, there is positive conditional probability that the cost realization $\{\tilde{c}_\tau\}_{\tau \in (\tilde{t}_{k-1},\infty)}$ after time \tilde{t}_{k-1} is such that there exists a time t satisfying $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \tilde{\pi}) > 0$, $\tilde{c}_t = g$, and $\phi_u^i(h_0, \{\tilde{c}_\tau\}_{\tau \in (0,\infty)}, \tilde{\pi}) = 0$ for all $u \in (\tilde{t}_{k-1}, t)$. Choose any cost level $\tilde{r}_k \in [\tilde{g}_k, r_k)$ for which there exists a cost realization $\{\tilde{c}_\tau^k\}_{\tau \in [0,\infty)}$ along with a time $\tilde{t}_k > \tilde{t}_{k-1}$ such that $\{\tilde{c}_\tau^k\}_{\tau \in [0,\tilde{t}_{k-1}]} = \{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]}$, such that $\phi_t^i(h_0, \{\tilde{c}_\tau^k\}_{\tau \in (0,\infty)}, \tilde{\pi}) > 0$ and $\tilde{c}_t^k = \tilde{r}_k$ for $t = \tilde{t}_k$, and such that every element of the set $R(\{\tilde{c}_\tau^k\}_{\tau \in [0,\tilde{t}_k]})$ is less than \tilde{r}_k . For the case where $r_k \notin R(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\tilde{t}_{k-1}]})$, the strategy profile π^* requires each agent to make a transaction at the first time that the cost reaches \tilde{r}_k , and each agent should transfer the amount

$\sigma_{\tilde{t}_{k-1}}^{i,i}(\{\tilde{c}_\tau^{k-1}\}_{\tau \in [0,\infty)}, \tilde{\pi}) - \sigma_{\tilde{t}_k}^{i,i}(\{\tilde{c}_\tau^k\}_{\tau \in [0,\infty)}, \tilde{\pi})$ on this transaction. In each case, such a choice of $\{\tilde{c}_\tau^k\}_{\tau \in [0,\infty)}$ and \tilde{t}_k is possible because of the property described above.

The agents do not make any transfers following a deviation from the path of play described above. Note that π^* is a stationary symmetric SPE in grim-trigger strategies with $V(h_0, \pi^*) > V(h_0, \tilde{\pi})$ for which the cost incurred between any two consecutive transactions is decreasing.

To understand why $V(h_0, \pi^*) > V(h_0, \tilde{\pi})$, consider the following argument. Because there is a positive probability when playing $\tilde{\pi}$ that the cost incurred is nondecreasing between some two consecutive transactions, there exists a positive integer m for which one can find a positive integer n greater than m such that there is a positive probability of the cost incurred at the n^{th} transaction being at least as large as the cost incurred at the m^{th} transaction. Let m^* denote the least such integer m . If $m^* = 1$, then the first transaction occurs at least as early when playing π^* as when playing $\tilde{\pi}$, and the amount transferred at the first transaction is larger when playing π^* than when playing $\tilde{\pi}$. Noting that the incentive constraint at the first transaction is binding when playing $\tilde{\pi}$, it follows that $V(h_0, \pi^*) > V(h_0, \tilde{\pi})$ if $m^* = 1$.

If $m^* > 1$, then both $\tilde{\pi}$ and π^* induce the same path of play up to the $(m^* - 1)^{\text{th}}$ transaction. That is, $t_{m-1}(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \tilde{\pi}) = t_{m-1}(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi^*)$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}$, and $\phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \tilde{\pi}) = \phi_t^i(h_0, \{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi^*)$ for any cost realization $\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}$ along with any time t such that $t \leq t_{m-1}(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \tilde{\pi}) = t_{m-1}(\{\tilde{c}_\tau\}_{\tau \in [0,\infty)}, \pi^*)$. Moreover, the m^{th} transaction occurs at least as early when playing π^* as when playing $\tilde{\pi}$, and the amount transferred at the m^{th} transaction is larger when playing π^* than when playing $\tilde{\pi}$. Noting that the incentive constraint at the m^{th} transaction is binding when playing $\tilde{\pi}$, it follows that $V(h_0, \pi^*) > V(h_0, \tilde{\pi})$ if $m^* > 1$.

Finally, there may be a positive probability of the incentive constraint being slack at some transaction when playing π^* . However, the procedure described in the last nine paragraphs from the proof of lemma 2 can be applied to π^* to generate a stationary symmetric SPE π' in grim-trigger strategies with $V(h_0, \pi') = V(h_0, \pi^*)$ for which there is probability one of the incentive constraint being binding at every transaction and there is probability one of the cost incurred being decreasing between any two transactions. \square

A.6 Proof of Corollaries to Theorem 3

A.6.1 Proof of Corollary 1

Proof. Suppose that the sequence $\{\tilde{f}_k(\pi)\}_{k=1}^{\infty}$ does not converge to zero. Then there exists n such that $\sum_{k=1}^n \tilde{f}_k(\pi) > q$. However, this violates the requirement for the strategies in π to be feasible. It follows that $\lim_{k \rightarrow \infty} \tilde{f}_k(\pi) = 0$.

Now suppose that $\lim_{k \rightarrow \infty} \tilde{f}_k(\pi) = 0$ but that the sequence $\{\tilde{c}_k(\pi)\}_{k=1}^{\infty}$ does not converge to zero. For any $\epsilon > 0$, there exists n such that $\sum_{k=m}^{\infty} \tilde{f}_k(\pi) < \epsilon$ for all $m \geq n$. Moreover, there exists $\eta > 0$ such that for any n , one can find $m \geq n$ for which $\tilde{c}_m(\pi) > \eta$. Hence, there exists l such that $\sum_{k=l}^{\infty} \tilde{f}_k(\pi) < \tilde{c}_l(\pi)$.

It follows that, with positive probability, the cost $\tilde{c}_l(\pi)$ is reached for the first time, and the expected payoff to each agent is negative at that time. However, an agent can secure an expected payoff of at least zero by transferring nothing at such a time. This contradicts the fact that π is an SPE. It follows that $\lim_{k \rightarrow \infty} \tilde{c}_k(\pi) = 0$. \square

A.6.2 Proof of Corollary 2

Proof. Suppose that $\sum_{k=1}^{\infty} \tilde{f}_k(\pi) < q$. Then there exists a symmetric SPE $\pi' \in \Pi'$ such that $\tilde{f}_k(\pi') = \tilde{f}_k(\pi)$ for $k \geq 2$, $\tilde{c}_k(\pi') = \tilde{c}_k(\pi)$ for $k \geq 1$, and $\tilde{f}_1(\pi') = \tilde{f}_1(\pi) + [q - \sum_{k=1}^{\infty} \tilde{f}_k(\pi)]$. Note that $V(h_0, \pi') > V(h_0, \pi)$. This contradicts the fact that π is a maximal symmetric SPE. It follows that $\sum_{k=1}^{\infty} \tilde{f}_k(\pi) = q$. \square

A.7 Proof of Theorem 4

We first restrict attention to stationary maximal symmetric SPE in grim-trigger strategies. In particular, we show that there exists a strategy profile $\pi^* \in \Pi'$ such that $V(h_0, \pi^*) > V(h_0, \pi)$ for every $\pi \in \Pi'$ that does not almost surely induce the same path of play as π^* . Moreover, we characterize the path of play induced by π^* . We then show that there does not exist a non-stationary maximal symmetric SPE in grim-trigger strategies.

Choose any strategy profile $\pi \in \Pi'$ for which there is positive probability of a transaction occurring. Let $\pi' \in \Pi'$ denote the strategy profile in which the k^{th} transaction is made when the cost reaches $\tilde{c}_k(\pi)$ for the first time and the amount $\tilde{f}_k(\pi)$ is transferred by each agent at this transaction. Note that π almost surely induces the same path of play as π' .

Consider the expected payoff to each agent when both agents play strategy profile π' . Assume that each agent has previously made $n \geq 0$ transactions and that the current value of the cost is $c \geq \tilde{c}_{n+1}(\pi)$. Let $M_n(c, \pi)$ be the expected payoff to each agent if both agents play strategy profile π' from the current time onwards. The value $M_n(c, \pi)$ is defined as follows. Let $\{g_t, (b_t^1, b_t^2)\}_{t \in [0, \infty)}$ be any history that is consistent with π'_1 and π'_2 at each $t \in [0, \infty)$ and such that there exists $u \geq 0$ for which $g_u = c$ and the set $\{t \in [0, u) : b_t^1 > 0 \text{ and } b_t^2 > 0\}$ contains exactly n elements. Then $M_n(c, \pi) = V[(\{g_t\}_{t \in [0, u]}, \{(b_t^1, b_t^2)\}_{t \in [0, u)}), \pi']$.

The sequence $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^\infty$ characterizing π must have the following properties. Because there is probability one of the cost incurred being decreasing between any two consecutive transactions, it must be that $\tilde{c}_{k+1}(\pi) < \tilde{c}_k(\pi)$ for all $k \geq 1$. In order for π to be feasible for the agents to play, it must be that $\sum_{k=1}^\infty \tilde{f}_k(\pi) \leq q$. Because there is probability one of the incentive constraint being binding at every transaction, it must be that $\tilde{c}_k(\pi) = M_k[\tilde{c}_k(\pi), \pi]$ for all $k \geq 1$.

It follows that $M_n(c, \pi)$ is simply the value of an asset that pays $\tilde{f}_{n+1}(\pi)$ at the first time that the cost reaches $\tilde{c}_{n+1}(\pi)$ when the current value of the cost is $c \geq \tilde{c}_{n+1}(\pi)$. Calculating $M_n(c, \pi)$ is a basic asset pricing problem.⁴³ The Bellman equation for this problem is:

$$\rho M_n(c, \pi) dt = \mathbb{E}(dM_n), \quad (3)$$

with the boundary condition:

$$M_n[\tilde{c}_{n+1}(\pi), \pi] = \tilde{f}_{n+1}(\pi). \quad (4)$$

The result below describes the solution to this asset pricing problem under the assumption that the cost process follows a geometric Brownian motion.

Lemma 4. *The unique solution to the Bellman equation (3) with the boundary condition (4) is:*

$$M_n(c, \pi) = \tilde{f}_{n+1}(\pi) \left(\frac{c}{\tilde{c}_{n+1}(\pi)} \right)^\beta,$$

where

$$\beta = \frac{1}{2} - \mu/\sigma^2 - \sqrt{(\mu/\sigma^2 - \frac{1}{2})^2 + 2\rho/\sigma^2}.$$

⁴³McDonald and Siegel (1986) perform a similar calculation under the assumption of geometric Brownian motion when solving for the optimal timing of a single irreversible investment.

Proof. A straightforward application of Ito's lemma to the right-hand side of equation (3) yields:

$$\rho M_n(c, \pi) = \mu c \frac{\partial M_n(c, \pi)}{\partial c} + \frac{1}{2} \sigma^2 c^2 \frac{\partial^2 M_n(c, \pi)}{\partial c^2},$$

which provides a second-order linear differential equation for $M_n(c, \pi)$. The boundary condition is $M_n[\tilde{c}_{n+1}(\pi), \pi] = \tilde{f}_{n+1}(\pi)$. We seek a solution of the form $g[c, \tilde{c}_{n+1}(\pi), \tilde{f}_{n+1}(\pi)] = B[\tilde{c}_{n+1}(\pi), \tilde{f}_{n+1}(\pi)]c^{\tilde{\beta}}$. The following quadratic equation is obtained by substituting the functional form into the differential equation:

$$\frac{1}{2} \sigma^2 \tilde{\beta}(\tilde{\beta} - 1) + \mu \tilde{\beta} - \rho = 0,$$

whose solution is given by:

$$\tilde{\beta} = \frac{1}{2} - \mu/\sigma^2 \pm \sqrt{(\mu/\sigma^2 - \frac{1}{2})^2 + 2\rho/\sigma^2}.$$

Letting $\tilde{\beta}^+$ and $\tilde{\beta}^-$ respectively denote the positive and negative roots of the quadratic, the general solution to the differential equation is $M_n(c, \pi) = B^+[\tilde{c}_{n+1}(\pi), \tilde{f}_{n+1}(\pi)]c^{\tilde{\beta}^+} + B^-[\tilde{c}_{n+1}(\pi), \tilde{f}_{n+1}(\pi)]c^{\tilde{\beta}^-}$. It must be the case that $B^+[\tilde{c}_{n+1}(\pi), \tilde{f}_{n+1}(\pi)] = 0$, because $M_n(c, \pi)$ would otherwise become unboundedly large in absolute value as c goes to ∞ . Moreover, the boundary condition $M_n[\tilde{c}_{n+1}(\pi), \pi] = \tilde{f}_{n+1}(\pi)$ yields $B^-[\tilde{c}_{n+1}(\pi), \tilde{f}_{n+1}(\pi)] = \tilde{f}_{n+1}(\pi)/\tilde{c}_{n+1}(\pi)^{\tilde{\beta}^-}$. Hence, the solution to the Bellman equation (3) is:

$$M_n(c, \pi) = \tilde{f}_{n+1}(\pi) \left(\frac{c}{\tilde{c}_{n+1}(\pi)} \right)^{\tilde{\beta}^-}.$$

Denoting $\beta = \tilde{\beta}^-$ gives the expression in the statement of the lemma. Uniqueness obtains because the solution to a second-order linear differential equation is unique in general. \square

We search for a strategy profile $\hat{\pi} \in \Pi'$ such that $V(h_0, \hat{\pi}) \geq V(h_0, \tilde{\pi})$ for every $\tilde{\pi} \in \Pi'$. The strategy profile π satisfies this condition if and only if the trading policy $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}$ solves the optimization problem:

$$\begin{aligned} & \max_{\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}} \tilde{f}_1(\pi) \left(\frac{c_0}{\tilde{c}_1(\pi)} \right)^{\beta}, \text{ subject to } \sum_{k=1}^{\infty} \tilde{f}_k(\pi) \leq q \\ & \text{and } \tilde{c}_n(\pi) = \tilde{f}_{n+1}(\pi) \left(\frac{\tilde{c}_n(\pi)}{\tilde{c}_{n+1}(\pi)} \right)^{\beta} \text{ for all } n \geq 1, \end{aligned} \tag{5}$$

where $\tilde{c}_k(\pi) \in (0, c_0)$ and $\tilde{f}_k(\pi) \in (0, q)$ for all k , and $\tilde{c}_k(\pi)$ is decreasing in k .

Our approach for solving the optimization problem above is as follows. We impose the additional constraint that $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$. We then identify the unique sequence that solves the maximization problem in expression (5) with this additional constraint. We finally argue that this sequence is the unique solution to the original maximization problem without the extra constraint.

The result below describes how the maximization problem in (5) can be simplified by applying an extra constraint.

Lemma 5. *Given the additional constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$, the trading policy $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}$ solves the maximization problem in expression (5) if and only if it solves:*

$$\begin{aligned} & \max_{\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}} \sum_{k=1}^{\infty} \left(\frac{\beta}{\beta - 1} \right)^{k-1} \ln[\tilde{f}_k(\pi)], \text{ subject to } \sum_{k=1}^{\infty} \tilde{f}_k(\pi) \leq q \\ & \text{and } \ln[\tilde{c}_n(\pi)] = -(\beta - 1)^{-1} \sum_{j=0}^{\infty} \left(\frac{\beta}{\beta - 1} \right)^j \ln[\tilde{f}_{n+j+1}(\pi)] \text{ for all } n \geq 1. \end{aligned} \quad (6)$$

Proof. Note that each incentive compatibility constraint in expression (5) can be rewritten as:

$$\ln[\tilde{c}_n(\pi)] = \frac{\beta \ln[\tilde{c}_{n+1}(\pi)] - \ln[\tilde{f}_{n+1}(\pi)]}{\beta - 1}.$$

Iterating the incentive compatibility constraint x times starting with transfer n , it can be shown by induction that $\tilde{c}_n(\pi)$ must satisfy:

$$\ln[\tilde{c}_n(\pi)] = \left(\frac{\beta}{\beta - 1} \right)^x \ln[\tilde{c}_{x+n}(\pi)] - (\beta - 1)^{-1} \sum_{j=0}^{x-1} \left(\frac{\beta}{\beta - 1} \right)^j \ln[\tilde{f}_{n+j+1}(\pi)].$$

Assuming that $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$, the term $\ln[\tilde{c}_n(\pi)]$ can be expressed as follows:

$$\ln[\tilde{c}_n(\pi)] = -(\beta - 1)^{-1} \sum_{j=0}^{\infty} \left(\frac{\beta}{\beta - 1} \right)^j \ln[\tilde{f}_{n+j+1}(\pi)].$$

Substituting for $\ln[\tilde{c}_1(\pi)]$ after taking the logarithm of the maximand in expression (5), we obtain the optimization problem in expression (6), where the constant term $\beta \ln(c_0)$ is eliminated from the logarithm of the maximand. \square

The maximization problem in (6) is solved in three parts. The result below provides the solution for the sequence of amounts transferred when the constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$ is not imposed.

Lemma 6. *Dropping the additional constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$, the unique sequence of amounts transferred that solves the maximization problem in expression (6) is given by:*

$$f_k^* = \left(\frac{q}{1 - \beta} \right) \left(\frac{\beta}{\beta - 1} \right)^{k-1} \text{ for all } k \geq 1.$$

Proof. The sequence $\{\tilde{f}_k(\pi)\}_{k=1}^{\infty}$ of amounts transferred solves the maximization problem in expression (6) if and only if it solves the simplified version of the problem below:

$$\max_{\{\tilde{f}_k(\pi)\}_{k=1}^{\infty}} \sum_{k=1}^{\infty} \left(\frac{\beta}{\beta - 1} \right)^{k-1} \ln[\tilde{f}_k(\pi)], \text{ subject to } \sum_{k=1}^{\infty} \tilde{f}_k(\pi) \leq q.$$

The Lagrangian for the maximization problem in the preceding expression is given by:

$$\mathcal{L} = \sum_{k=1}^{\infty} \left(\frac{\beta}{\beta - 1} \right)^{k-1} \ln[\tilde{f}_k(\pi)] + \lambda \left(q - \sum_{k=1}^{\infty} \tilde{f}_k(\pi) \right),$$

which provides a first-order condition for each index $k \geq 1$:

$$\tilde{f}_k(\pi) = \frac{1}{\lambda} \left(\frac{\beta}{\beta - 1} \right)^{k-1}.$$

Substitution of this result into the budget constraint yields:

$$q = \frac{1}{\lambda} \sum_{k=1}^{\infty} \left(\frac{\beta}{\beta - 1} \right)^{k-1} = \frac{1 - \beta}{\lambda}.$$

Thus, we obtain $\lambda = (1 - \beta)/q$, leading to the following sequence of amounts transferred:

$$\tilde{f}_k(\pi) = \frac{q}{1 - \beta} \left(\frac{\beta}{\beta - 1} \right)^{k-1} \text{ for all } k \geq 1.$$

□

Given the amount transferred at each transaction, the next result states the solu-

tion for the sequence of cost cutoffs when the constraint $\lim_{x \rightarrow \infty} [\beta/(\beta-1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$ is not imposed.

Lemma 7. *Dropping the additional constraint $\lim_{x \rightarrow \infty} [\beta/(\beta-1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$, the unique sequence of cost cutoffs that solves the maximization problem in expression (6) is given by:*

$$c_k^* = \left(\frac{q}{1-\beta} \right) \left(\frac{\beta}{\beta-1} \right)^{k-\beta} \text{ for all } k \geq 1.$$

Proof. Substituting the expression for $\tilde{f}_k(\pi)$ obtained in lemma 6 into the expression for $\ln[\tilde{c}_n(\pi)]$ stated in lemma 5, we have:

$$\begin{aligned} \ln[\tilde{c}_n(\pi)] &= (1-\beta)^{-1} \sum_{j=0}^{\infty} \left(\frac{\beta}{\beta-1} \right)^j \left[\ln \left(\frac{q}{1-\beta} \right) + \ln \left(\frac{\beta}{\beta-1} \right)^{j+n} \right] \\ &= \ln \left(\frac{q}{1-\beta} \right) + (n-\beta) \ln \left(\frac{\beta}{\beta-1} \right), \end{aligned}$$

where the standard formulas for geometric series and their derivatives provide:

$$(1-\beta)^{-1} \sum_{j=0}^{\infty} \left(\frac{\beta}{\beta-1} \right)^j \ln \left(\frac{q}{1-\beta} \right) = \ln \left(\frac{q}{1-\beta} \right)$$

and

$$\begin{aligned} (1-\beta)^{-1} \sum_{j=0}^{\infty} \left(\frac{\beta}{\beta-1} \right)^j \ln \left(\frac{\beta}{\beta-1} \right)^{j+n} &= (1-\beta)^{-1} \ln \left(\frac{\beta}{\beta-1} \right) \sum_{j=0}^{\infty} (j+n) \left(\frac{\beta}{\beta-1} \right)^j \\ &= (1-\beta)^{-1} \ln \left(\frac{\beta}{\beta-1} \right) [n(1-\beta) - \beta(1-\beta)] \\ &= (n-\beta) \ln \left(\frac{\beta}{\beta-1} \right). \end{aligned}$$

Thus, the sequence of cost cutoffs is given by:

$$\tilde{c}_k(\pi) = \left(\frac{q}{1-\beta} \right) \left(\frac{\beta}{\beta-1} \right)^{k-\beta} \text{ for all } k \geq 1.$$

□

Using the solution for the sequence of cost cutoffs, we can confirm that the constraint $\lim_{x \rightarrow \infty} [\beta/(\beta-1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$ is satisfied, even though the constraint was

not imposed when solving the maximization problem.

Lemma 8. *The solution $\{c_k^*, f_k^*\}_{k=1}^\infty$ in lemmata 6 and 7 satisfies the technical condition $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$.*

Proof. Substituting the expression for $\{c_k^*\}$ from lemma 7 into the technical condition, we have:

$$\begin{aligned}
& \lim_{x \rightarrow \infty} \left(\frac{\beta}{\beta - 1} \right)^x \ln(c_{x+n}^*) \\
&= \lim_{x \rightarrow \infty} \left(\frac{\beta}{\beta - 1} \right)^x \left[\ln \left(\frac{q}{1 - \beta} \right) + (x + n - \beta) \ln \left(\frac{\beta}{\beta - 1} \right) \right] \\
&= \lim_{x \rightarrow \infty} \left(\frac{\beta}{\beta - 1} \right)^x \ln \left(\frac{q}{1 - \beta} \right) + \lim_{x \rightarrow \infty} \left(\frac{\beta}{\beta - 1} \right)^x (x + n - \beta) \ln \left(\frac{\beta}{\beta - 1} \right) \\
&= \lim_{x \rightarrow \infty} \left(\frac{\beta}{\beta - 1} \right)^x (x + n - \beta) \ln \left(\frac{\beta}{\beta - 1} \right) \\
&= 0.
\end{aligned}$$

□

Hence, the sequence $\{c_k^*, f_k^*\}_{k=1}^\infty$ solves the maximization problem in expression (6) when the constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$ is imposed. The lemma below implies that $\{c_k^*, f_k^*\}_{k=1}^\infty$ is the unique solution to the optimization problem in expression (5).

Lemma 9. *If $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^\infty$ is the unique sequence that solves the maximization problem in expression (6) with the constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$, then $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^\infty$ is the unique solution to the maximization problem in expression (5).*

Proof. In the proof of lemma 5, it was shown that any trading policy $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^\infty$ satisfying the constraints for the maximization problem in expression (5) must solve the equation:

$$\ln[\tilde{c}_n(\pi)] = \left(\frac{\beta}{\beta - 1} \right)^x \ln[\tilde{c}_{x+n}(\pi)] - (\beta - 1)^{-1} \sum_{j=0}^{x-1} \left(\frac{\beta}{\beta - 1} \right)^j \ln[\tilde{f}_{n+j+1}(\pi)]. \quad (7)$$

First, note that there must exist an index $B > 1$ such that $\tilde{f}_b(\pi) < 1$ for all $b \geq B$. Otherwise, if there were no such B , then the feasibility constraint $\sum_{k=1}^\infty \tilde{f}_k(\pi) \leq q$

would be violated. Second, it must be that $\tilde{c}_b(\pi) < 1$ for all $b \geq B - 1$. Noting that $[\tilde{c}_n(\pi)/\tilde{c}_{n+1}(\pi)]^\beta < 1$ for all $n \geq 1$, this claim follows from the incentive compatibility constraint $\tilde{c}_n(\pi) = \tilde{f}_{n+1}(\pi)[\tilde{c}_n(\pi)/\tilde{c}_{n+1}(\pi)]^\beta$ for all $n \geq 1$ along with the fact that $\tilde{f}_b(\pi) < 1$ for all $b \geq B$.

Because $\ln[\tilde{f}_{n+j+1}(\pi)]$ is negative for j sufficiently large, there exists X such that the sum

$$-(\beta - 1)^{-1} \sum_{j=0}^{x-1} \left(\frac{\beta}{\beta - 1} \right)^j \ln[\tilde{f}_{n+j+1}(\pi)]$$

is decreasing in x for $x \geq X$. This fact along with equation (7) implies that $[\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)]$ is increasing in x for $x \geq X$. Because $\ln[\tilde{c}_{x+n}(\pi)]$ is negative for x sufficiently large, it follows from the monotone convergence theorem that the term $[\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)]$ converges to some number $L \leq 0$ as x goes to infinity.

Now suppose that the additional constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$ is not imposed in the proof of lemma 5. We take the logarithm of the maximand of the optimization problem in expression (5), after which we substitute for $\ln[\tilde{c}_1(\pi)]$ using equation (7) in the limit as x goes to infinity. Doing so and eliminating the constant term $\beta \ln(c_0)$ from the logarithm of the maximand, the optimization problem becomes:

$$\begin{aligned} & \max_{\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}} -\beta \lim_{x \rightarrow \infty} \left(\frac{\beta}{\beta - 1} \right)^x \ln[\tilde{c}_{x+n}(\pi)] + \sum_{k=1}^{\infty} \left(\frac{\beta}{\beta - 1} \right)^{k-1} \ln[\tilde{f}_k(\pi)], \\ & \text{subject to } \sum_{k=1}^{\infty} \tilde{f}_k(\pi) \leq q \text{ and } \tilde{c}_n = \tilde{f}_{n+1}(\pi) \left(\frac{\tilde{c}_n(\pi)}{\tilde{c}_{n+1}(\pi)} \right)^\beta \text{ for all } n \geq 1. \end{aligned} \quad (8)$$

Choose any trading policy $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}$ that satisfies the constraints in the preceding expression. Consider the objective function in optimization problem (8). The first term in this objective function cannot be greater than 0 when evaluated at $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}$. Moreover, the second term in this objective function when evaluated at $\{\tilde{c}_k(\pi), \tilde{f}_k(\pi)\}_{k=1}^{\infty}$ cannot be greater than the maximum of the problem in expression (6) when the additional constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$ is not imposed.

Recall that the solution $\{c_k^*, f_k^*\}_{k=1}^{\infty}$ in lemmata 6 and 7 is the unique trading policy that achieves the maximum of problem (6) without the additional constraint $\lim_{x \rightarrow \infty} [\beta/(\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$. Moreover, the sequence $\{c_k^*, f_k^*\}_{k=1}^{\infty}$ satisfies the

technical condition $\lim_{x \rightarrow \infty} [\beta / (\beta - 1)]^x \ln[\tilde{c}_{x+n}(\pi)] = 0$. Hence, $\{c_k^*, f_k^*\}_{k=1}^\infty$ achieves the upper bounds for the first and second terms of the objective function in expression (8). In addition, $\{c_k^*, f_k^*\}_{k=1}^\infty$ satisfies all of the constraints in expression (8). It follows that $\{c_k^*, f_k^*\}_{k=1}^\infty$ is the unique solution to maximization problem (8), which is equivalent to optimization problem (5). \square

It follows that a stationary maximal symmetric SPE in grim-trigger strategies exists and that the sequence $\{c_k^*, f_k^*\}_{k=1}^\infty$ characterizes any such strategy profile.

We now show that there does not exist a non-stationary maximal symmetric SPE. Hence, any maximal symmetric SPE induces the aforementioned path of play.

Lemma 10. *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 there does not exist a non-stationary maximal symmetric SPE in grim-trigger strategies.*

Proof. Let π^* be any stationary maximal symmetric SPE in grim-trigger strategies. Recall that any stationary maximal symmetric SPE induces the path of play specified in the statement of theorem 4. Suppose that π' is a non-stationary maximal symmetric SPE in grim-trigger strategies. Let g be the greatest positive integer for which there is probability one that the realization of the cost process $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that $\phi_t^i(h_0, \{c_\tau\}_{\tau \in [0, \infty)}, \pi') = \phi_t^i(h_0, \{c_\tau\}_{\tau \in [0, \infty)}, \pi^*)$ for all $t \leq t_{g-1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^*)$. Such an integer g must exist because π^* is a stationary SPE and π' is a non-stationary SPE; so that, π^* and π' must induce different paths of play with positive probability. Noting that the SPE π' is non-stationary, there must be positive probability that $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$. Recall the definitions of c_k^* and f_k^* in the statement of theorem 4. Let $c_0^* = c_0$ for definitional purposes. At least one of the four claims in the following paragraph must be true.

First, there is positive probability that $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) > c_g^*$. Second, there is positive probability that $\{c_\tau\}_{\tau \in [0, \infty)}$ is such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) < c_g^*$. Third, there is probability one of $\{c_\tau\}_{\tau \in [0, \infty)}$ being either such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') = \infty$ or such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) = c_g^*$, and there is positive probability of $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that there exists $v < t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi')$ for which $c_v = c_g^*$. Fourth, there is probability one of $\{c_\tau\}_{\tau \in [0, \infty)}$ being either such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') = \infty$ or such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) = c_g^*$, there is zero probability of $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that there exists $v < t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi')$,

π') for which $c_v = c_g^*$, and there is positive probability of $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $\phi^i(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) \neq f_g^*$. We show below that there is a contradiction if one or more of the preceding claims is true. It will then follow that there cannot exist a non-stationary maximal symmetric SPE π' .

Consider the case where the first of the four claims holds. We argue in the paragraph below that there is zero probability of $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) \geq c_{g-1}^*$; so that, there is probability one that either $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') = \infty$ or both $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) < c_{g-1}^*$. Because π' is a maximal symmetric SPE for which $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) \in (c_g^*, c_{g-1}^*)$ with positive probability, there exists a realization $\{c_\tau^{1*}\}_{\tau \in [0, \infty)}$ of the cost process as well as a time u^{1*} such that $t_g(\{c_\tau^{1*}\}_{\tau \in [0, \infty)}, \pi') = u^{1*}$, $c_{u^{1*}}^{1*} > c_g^*$, and $V[h_t(\{c_\tau^{1*}\}_{\tau \in [0, \infty)}, \pi'] \geq V[h_t(\{c_\tau^{1*}\}_{\tau \in [0, \infty)}, \pi'']$ for every symmetric SPE $\pi'' \in \Pi^*$ in grim-trigger strategies along with any time $t \leq u^{1*}$. A stationary maximal symmetric SPE $\pi^{1*} \in \Pi^*$ in grim-trigger strategies can be constructed as follows. The agents play strategy profile π' up to and including the time t^{1*} when transaction $g - 1$ occurs. Thereafter, the agents do not make any transactions until the first time t^{1**} that the current value of the cost process is $c_{u^{1*}}^{1*}$. At time t^{1**} , each agent transfers the amount $\phi_{u^{1*}}^{i,i}(h_0, \{c_\tau^{1*}\}_{\tau \in [0, \infty)}, \pi')$. After this transaction, the agents follow the path of play specified in the statement of theorem 4, behaving as if the game just started at the time of this transaction with the initial value of the cost process being $c_{u^{1*}}^{1*}$ and the amount of each good being $\sigma_{u^{1*}}^{i,i}(\{c_\tau^{1*}\}_{\tau \in [0, \infty)}, \pi')$. The agents do not make any transfers following a deviation from the path of play described above. The fact that π^{1*} and π^* induce different paths of play with positive probability contradicts the assumption that every stationary maximal symmetric SPE induces the path of play specified in the statement of theorem 4.

We now argue as noted above that there is probability one that either $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') = \infty$ or both $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) < c_{g-1}^*$. Suppose to the contrary that $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) \geq c_{g-1}^*$ with positive probability. For any $k \geq 1$, it follows from $c_0 > q$ that there is probability one that either $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') = \infty$ or both $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', k) < c_0$. Otherwise, if there were to exist $k \geq 1$ such that $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', k) > c_0$ with positive probability, then there would be positive probability of $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that $t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $V[h_u(\{c_\tau\}_{\tau \in [0, \infty)}, \pi')] < 0$ for $u = t_k(\{c_\tau\}_{\tau \in [0, \infty)}, \pi')$, which would con-

tradict the assumption that π' is a symmetric SPE. It is immediate that $g \neq 1$. Assume therefore that $g > 1$. Given that π' is a maximal symmetric SPE for which $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) \geq c_{g-1}^*$ with positive probability, there exists a realization $\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}$ of the cost process as well as a time \tilde{u}^1 such that $t_g(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi') = \tilde{u}^1$, $\tilde{c}_{\tilde{u}^1}^1 \geq c_{g-1}^*$, and $V[h_{\tilde{u}^1}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi'), \pi'] \geq V[h_{\tilde{u}^1}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi''), \pi'']$ for every symmetric SPE $\pi'' \in \Pi^*$ in grim-trigger strategies. Define the symmetric SPE $\pi'' \in \Pi^*$ in grim-trigger strategies as follows. The agents play strategy profile π' up to but not including the first time \tilde{t}^1 that the current value of the cost process is $\tilde{c}_{\tilde{u}^1}^1$. At time \tilde{t}^1 , the agents each transfer the amount $\hat{s}_{\tilde{t}^1}^{i,i} - \sigma_{\tilde{u}^1}^{i,i}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi')$. After this transaction, the agents play according to the strategy profile π' , behaving as if the history up to the time of this transaction were $h_{\tilde{u}^1}(\{\tilde{c}_\tau^1\}_{\tau \in [0, \infty)}, \pi')$ and strategy profile π' had always been followed from the beginning of the game. The agents do not make any transfers following a deviation from the path of play described above. It can easily be seen that $V(h_0, \pi'') > V(h_0, \pi^*)$, which contradicts the assumption that π^* is a maximal symmetric SPE.

Consider the case where the second of the four claims holds. Because π' is a maximal symmetric SPE for which $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) < c_g^*$ with positive probability, there exists a realization $\{c_\tau^{2*}\}_{\tau \in [0, \infty)}$ of the cost process as well as a time u^{2*} such that $t_g(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi') = u^{2*}$, $c_{u^{2*}}^{2*} < c_g^*$, and $V[h_t(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi'] \geq V[h_t(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi''), \pi'']$ for every symmetric SPE $\pi'' \in \Pi^*$ in grim-trigger strategies along with any time $t \leq u^{2*}$. A stationary maximal symmetric SPE $\pi^{2*} \in \Pi^*$ in grim-trigger strategies can be constructed as follows. The agents play strategy profile π' up to and including the time t^{2*} when transaction $g - 1$ occurs. Thereafter, the agents do not make any transactions until the first time t^{2**} that the current value of the cost process is $c_{u^{2*}}^{2*}$. At time t^{2**} , each agent transfers the amount $\phi_{u^{2*}}^i(h_0, \{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi')$. After this transaction, the agents follow the path of play specified in the statement of theorem 4, behaving as if the game just started at the time of this transaction with the initial value of the cost process being $c_{u^{2*}}^{2*}$ and the amount of each good being $\sigma_{u^{2*}}^{i,i}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi')$. The agents do not make any transfers following a deviation from the path of play described above. We argue in the paragraph below that there is probability one that either $t_{g+1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{2*}) = \infty$ or both $t_{g+1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{2*}) < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{2*}, g + 1) < c_{u^{2*}}^{2*}$. The fact that π^{2*} and π^* induce different paths of play with positive probability contradicts the assumption that every stationary maximal symmetric SPE induces the path of play

specified in the statement of the theorem 4.

We argue as mentioned in the previous paragraph that there is probability one of $\{c_\tau\}_{\tau \in [0, \infty)}$ being such that either $t_{g+1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{2*}) = \infty$ or both $t_{g+1}(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{2*}) < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi^{2*}, g+1) < c_{u^{2*}}^{2*}$. To do so, it suffices to show that $\sigma^{i,i}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) \leq \tilde{s}$, where $\tilde{s} = -c_{u^{2*}}^{2*} \beta [\beta / (\beta - 1)]^{\beta - 1}$ would almost surely be the amount of each good remaining untransferred after the first transaction if the agents follow the path of play specified in the statement of theorem 4 in a game where the initial value of the cost process is c_0 and the initial stock of each good is such that the first transaction almost surely occurs at a cost of $c_{u^{2*}}^{2*}$. Suppose to the contrary that $\sigma^{i,i}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) > \tilde{s}$. Let $z^{2'} = \sigma^{i,i}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) - \tilde{s}$. Define the symmetric SPE $\pi^{2'} \in \Pi^*$ in grim-trigger strategies as follows. The agents play strategy profile π' until the history up to the current time is $h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi')$. Upon reaching $h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi')$, the agents each transfer the amount $\phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) + z^{2'}$. After this transaction, the agents follow the path of play specified in the statement of theorem 4, behaving as if the game just started at the time of this transaction with the initial value of the cost process being $c_{u^{2*}}^{2*}$ and the amount of each good being \tilde{s} . The agents do not make any transfers following a deviation from the path of play described above. There exists a realization $\{c_\tau^{2'}\}_{\tau \in [0, \infty)}$ of the cost process along with a time v^{2*} such that $c_{v^{2*}}^{2'} = c_{u^{2*}}^{2*}$, $\hat{\sigma}_{v^{2*}}^{i,i}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}) = \tilde{s}$, and $V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}] \geq V[h_{u^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}]$ for every symmetric SPE $\pi^{2''} \in \Pi^*$. Observe that $V[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi^{2'}] = V[h_{u^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}] + [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) - c_{u^{2*}}^{2*}]$. We argue in the paragraph below that $V[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi^{2'}] - V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}] < [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) - c_{u^{2*}}^{2*}]$, which implies that $V[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi^{2'}] > V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}]$. This contradicts the fact that $V[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi^{2'}] \geq V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}]$ for every symmetric SPE $\pi^{2''} \in \Pi^*$ in grim-trigger strategies. It follows that $\sigma^{i,i}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) \leq \tilde{s}$.

We now argue as mentioned in the previous paragraph that $V[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi^{2'}] - V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}] < [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) - c_{u^{2*}}^{2*}]$. Let $\Sigma[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'); \{c_\tau\}_{\tau \in [v^{2*}, t]}; u^{2*}, u^{2*} + (t - v^{2*}); \pi^{2'}]$ be the sum of the transfers that an agent would make between times u^{2*} and $u^{2*} + (t - v^{2*})$ inclusive if $h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi')$ is the history up to time u^{2*} , the cost process $\{c_\tau\}_{\tau \in [v^{2*}, t]}$ is realized between times u^{2*} and $u^{2*} + (t - v^{2*})$, and strategy profile π' is played by the agents. Define the symmetric SPE $\tilde{\pi}' \in \Pi^*$ as follows. The agents do not make any transactions until

the history up to the current time is $h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'})$. After $h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'})$ has been reached, the agents do not make any transfers at any time t for which the realization of the cost process $\{c_\tau\}_{\tau \in [v^{2*}, t]}$ between times v^{2*} and t is such that $\Sigma[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi']; \{c_\tau\}_{\tau \in [v^{2*}, t]}; u^{2*}, u^{2*} + (t - v^{2*}); \pi'] \leq [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g)]$. Each agent transfers the amount $\Sigma[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi']; \{c_\tau\}_{\tau \in [v^{2*}, t]}; u^{2*}, u^{2*} + (t - v^{2*}); \pi'] - [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g)]$ at the first time t for which the realization of the cost process $\{c_\tau\}_{\tau \in [v^{2*}, t]}$ between times v^{2*} and t is such that $\Sigma[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi']; \{c_\tau\}_{\tau \in [v^{2*}, t]}; u^{2*}, u^{2*} + (t - v^{2*}); \pi'] > [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g)]$. Let w^{2*} denote the first time t that this condition holds. Thereafter, the agents play according to strategy profile π' , behaving as if this transaction happened at time $u^{2*} + (w^{2*} - v^{2*})$, the history up to time u^{2*} were $h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi')$, the cost process $\{c_\tau\}_{\tau \in [v^{2*}, w^{2*}]}$ were realized between times u^{2*} and $u^{2*} + (w^{2*} - v^{2*})$, and strategy profile π' were followed up to and including time $u^{2*} + (w^{2*} - v^{2*})$. The agents do not make any further transfers if either agent has deviated from the path of play described above after reaching $h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'})$. It can easily be seen that $V[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi'] - V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \tilde{\pi}'] < [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) - c_{u^{2*}}^{2*}]$. Noting that $V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}] \geq V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \tilde{\pi}']$ for every symmetric SPE $\pi'' \in \Pi^*$, it follows that $V[h_{u^{2*}}(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi'), \pi'] - V[h_{v^{2*}}(\{c_\tau^{2'}\}_{\tau \in [0, \infty)}, \pi^{2'}), \pi^{2'}] < [z^{2'} + \phi^i(\{c_\tau^{2*}\}_{\tau \in [0, \infty)}, \pi', g) - c_{u^{2*}}^{2*}]$.

Consider the case where the third of the four claims holds. Because π' is a maximal symmetric SPE for which $t_g(\{c_\tau\}_{\tau \in [0, \infty)}, \pi') < \infty$ and $c(\{c_\tau\}_{\tau \in [0, \infty)}, \pi', g) = c_g^*$ with positive probability, there exists a realization $\{c_\tau^{3*}\}_{\tau \in [0, \infty)}$ of the cost process as well as a time u^{3*} such that $t_g(\{c_\tau^{3*}\}_{\tau \in [0, \infty)}, \pi') = u^{3*}$, $c_{u^{3*}}^{3*} = c_g^*$, and $V[h_{u^{3*}}(\{c_\tau^{3*}\}_{\tau \in [0, \infty)}, \pi'), \pi'] \geq V[h_{u^{3*}}(\{c_\tau^{3*}\}_{\tau \in [0, \infty)}, \pi'), \pi'']$ for every symmetric SPE $\pi'' \in \Pi^*$ in grim-trigger strategies. Define the symmetric SPE $\pi^{3*} \in \Pi^*$ in grim-trigger strategies as follows. The agents play strategy profile π' up to and including the time t^{3*} when transaction $g - 1$ occurs. Thereafter, the agents do not make any transactions until the first time t^{3**} that the current value of the cost process is $c_{u^{3*}}^{3*}$. At time t^{3**} , each agent transfers the amount $\phi_{u^{3*}}^i(h_0, \{c_\tau^{3*}\}_{\tau \in [0, \infty)}, \pi')$. After this transaction, the agents play according to the strategy profile π' , behaving as if the history up to the time of this transaction were $h_{u^{3*}}(\{c_\tau^{3*}\}_{\tau \in [0, \infty)}, \pi')$ and strategy profile π' had always been followed from the beginning of the game. The agents do not make any transfers following a deviation from the path of play described above. It can easily be seen that $V(h_0, \pi^{3*}) > V(h_0, \pi')$, which contradicts the assumption that π' is a maximal symmetric SPE.

Consider the case where the last of the four claims holds. A stationary maximal symmetric SPE $\pi^{4*} \in \Pi^*$ in grim-trigger strategies can be constructed as follows. The agents play strategy profile π' up to and including the time t^{4*} when transaction g occurs. Thereafter, the agents follow the path of play specified in the statement of theorem 4, behaving as if the game just started with the initial value of the cost process being $c_{t^{4*}}^{4*}$ and the amount of each good being $s_{t^{4*}}^{i,i}$. The agents do not make any transfers following a deviation from the path of play described above. The fact that π^{4*} and π^* induce different paths of play with positive probability contradicts the assumption that every stationary maximal symmetric SPE induces the path of play specified in the statement of theorem 4. \square

A.8 Proof of Corollaries to Theorem 4

A.8.1 Proof of Corollary 4

Proof. The sign of $\partial f_k^*/\partial\beta$ can be determined as follows:

$$\begin{aligned} \operatorname{sgn} \left(\frac{\partial f_k^*}{\partial\beta} \right) &= \operatorname{sgn} \left\{ \partial \left[\frac{q}{1-\beta} \left(\frac{\beta}{\beta-1} \right)^{k-1} \right] / \partial\beta \right\} \\ &= \operatorname{sgn} \left[-(k-1)(\beta)^{k-2}(\beta-1)^k + k(\beta)^{k-1}(\beta-1)^{k-1} \right] \\ &= \operatorname{sgn} \left\{ [-(k-1)(\beta-1) + k\beta](\beta)^{k-2}(\beta-1)^{k-1} \right\} \\ &= \operatorname{sgn} [(k-1)(\beta-1) - k\beta] = \operatorname{sgn} (1 - k - \beta). \end{aligned}$$

Thus, $\partial f_k^*/\partial\beta$ is positive if $k < 1 + |\beta|$ and negative if $k > 1 + |\beta|$. In addition, we have $\partial\beta/\partial\mu < 0$, $\partial\beta/\partial\sigma > 0$, and $\partial\beta/\partial\rho < 0$. Combining these facts, we obtain the desired result. \square

A.8.2 Proof of Corollary 5

Proof. The sign of $\partial c_k^*/\partial\beta$ can be determined as follows:

$$\begin{aligned} \operatorname{sgn}\left(\frac{\partial c_k^*}{\partial\beta}\right) &= \operatorname{sgn}\left(\frac{\partial \ln(c_k^*)}{\partial\beta}\right) \\ &= \operatorname{sgn}\left(\frac{\partial\{\ln(q) - \ln(1-\beta) + (k-\beta)[\ln(-\beta) - \ln(1-\beta)]\}}{\partial\beta}\right) \\ &= \operatorname{sgn}\left[\frac{1}{1-\beta} + (k-\beta)\left(\frac{1}{\beta} + \frac{1}{1-\beta}\right) - \ln\left(\frac{-\beta}{1-\beta}\right)\right]. \end{aligned}$$

Let k^* be the value of k that solves the following equation:

$$\frac{1}{1-\beta} + (k-\beta)\left(\frac{1}{\beta} + \frac{1}{1-\beta}\right) - \ln\left(\frac{-\beta}{1-\beta}\right) = 0,$$

which gives:

$$k^* = -\beta(1-\beta) \ln\left(\frac{1-\beta}{-\beta}\right) > 0.$$

Note that $\partial c_k^*/\partial\beta$ is positive if $k < k^*$ and negative if $k > k^*$. In addition, we have $\partial\beta/\partial\mu < 0$, $\partial\beta/\partial\sigma > 0$, and $\partial\beta/\partial\rho < 0$. Hence, c_k^* is decreasing in μ and ρ as well as increasing in σ for $k < k^*$, and c_k^* is increasing in μ and ρ as well as decreasing in σ for $k > k^*$.

We next examine how k^* varies with β . The partial derivative with respect to β is given by:

$$\frac{\partial c_k^*}{\partial\beta} = 1 - (1-2\beta) \ln\left(\frac{1-\beta}{-\beta}\right) < 1 - (1-2\beta)\left(1 - \frac{-\beta}{1-\beta}\right) = \frac{\beta}{1-\beta} < 0,$$

where the general result $\ln(x) > 1 - 1/x$ for all $x > 1$ is used to establish the first inequality. Hence, k^* is increasing in $|\beta|$.

The limit as β goes to $-\infty$ can be calculated as:

$$\begin{aligned} \lim_{\beta \rightarrow -\infty} k^* &= \lim_{\beta \rightarrow -\infty} -\beta(1-\beta) \ln\left(\frac{1-\beta}{-\beta}\right) \geq \lim_{\beta \rightarrow -\infty} -\beta(1-\beta)\left(1 - \frac{-\beta}{1-\beta}\right), \\ &= \lim_{\beta \rightarrow -\infty} -\beta = \infty \end{aligned}$$

where the general result $\ln(x) > 1 - 1/x$ for all $x > 1$ is used to establish the inequality. Hence, $\lim_{|\beta| \rightarrow \infty} k^* = \infty$.

The limit as β goes to 0 can be calculated as:

$$\begin{aligned} \lim_{\beta \rightarrow 0^-} k^* &= \lim_{\beta \rightarrow 0^-} -\beta(1-\beta) \ln \left(\frac{1-\beta}{-\beta} \right) = \lim_{\beta \rightarrow 0^-} \ln \left(\frac{1-\beta}{-\beta} \right) / \left(\frac{1}{-\beta(1-\beta)} \right) \\ &= \lim_{\beta \rightarrow 0^-} \frac{\partial}{\partial \beta} \left[\ln \left(\frac{1-\beta}{-\beta} \right) \right] / \frac{\partial}{\partial \beta} \left(\frac{1}{-\beta(1-\beta)} \right), \\ &= \lim_{\beta \rightarrow 0^-} \left(\frac{1}{-\beta(1-\beta)} \right) / \left(\frac{1}{\beta^2} - \frac{1}{(1-\beta)^2} \right) = \lim_{\beta \rightarrow 0^-} \frac{(-\beta)(1-\beta)}{1-2\beta} = 0 \end{aligned}$$

where the third equality applies L'Hôpital's rule. Hence, $\lim_{|\beta| \rightarrow 0} k^* = 0$. \square

A.9 Proof of Proposition 2

Proof. Choose any symmetric strategy profile $\pi = (\pi_1, \pi_2)$ with $\pi_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$. Suppose that there is positive probability of the cost realization $\{c_t\}_{t \in [0, \infty)}$ being such that $\phi_i^1(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) = \phi_i^2(h_0, \{c_\tau\}_{\tau \in (0, \infty)}, \pi) \in (0, q)$ for some $t \in [0, \infty)$. Let $\pi' = (\pi'_1, \pi'_2)$ with $\pi'_i \in \bar{\Pi}_i^C$ for $i \in \{1, 2\}$ be a symmetric strategy profile that requires the agents to behave as follows. At any history where a transaction occurs for the first time when playing π , each agent transfers the amount q . At any other history, each agent transfers the amount 0. Note that $V(h_0, \pi') > V(h_0, \pi)$. Hence, π is not efficient. \square

A.10 Proof of Proposition 3

Proof. We identify the path of play induced by an efficient strategy profile. In the absence of incentive constraints, the efficient policy requires the agents to almost surely make at most one transfer. Attention can be restricted to a stationary threshold policy. The efficient path of play can be obtained by choosing $\hat{c} \in (0, q)$ so as to maximize the value $L(\hat{c})$ of an asset that pays $q - \hat{c}$ when the cost first reaches \hat{c} starting from c_0 . From the earlier analysis, the value of such an asset is:

$$L(\hat{c}) = (q - \hat{c})(\hat{c}/c_0)^{-\beta}.$$

Differentiating with respect to \hat{c} yields the following first-order condition for \bar{c} :

$$-\beta(q - \bar{c})c_0^{-1}(\bar{c}/c_0)^{-\beta-1} - (\bar{c}/c_0)^{-\beta} = 0 \Rightarrow \bar{c} = -\beta(1 - \beta)^{-1}q,$$

where the derivative is positive for $\hat{c} < \bar{c}$ and negative for $\hat{c} > \bar{c}$. □

A.11 Proofs of Proposition 4 and Its Corollaries

A.11.1 Proof of Proposition 4

Proof. We begin by computing the solution to the second-best problem. From the proof of theorem 4, the value function M^{sb} of each agent for the second-best problem is:

$$M^{sb} = \frac{f_1^*}{c_1^{*\beta}} c_0^\beta,$$

where f_1^* and c_1^* are defined as follows:

$$f_1^* = \frac{1}{1-\beta} \quad \text{and} \quad c_1^* = \frac{q}{1-\beta} \left(\frac{-\beta}{1-\beta} \right)^{1-\beta},$$

and $\beta < 0$ is given by:

$$\beta = \frac{1}{2} - \frac{\mu}{\sigma^2} - \sqrt{\left(\frac{\mu}{\sigma^2} - \frac{1}{2} \right)^2 + 2\frac{\rho}{\sigma^2}}.$$

Substituting for f_1^* and c_1^* in the expression for M^{sb} results in:

$$M^{sb} = \theta^{1-\beta} q^{1-\beta} c_0^\beta,$$

where the constant $\theta \in (0, 1)$ is given by:

$$\theta = (-\beta)^{-\beta} (1-\beta)^{-(1-\beta)}.$$

We next compute the solution to the first-best problem. From the proof of proposition 3, the value function M^{fb} of each agent for the first-best problem is:

$$M^{fb} = \frac{q - \bar{c}}{\bar{c}^\beta} c_0^\beta,$$

where \bar{c} is defined as follows:

$$\bar{c} = -\beta(1-\beta)^{-1}q.$$

Substituting for \bar{c} in the expression for M^{fb} results in:

$$M^{fb} = \theta q^{1-\beta} c_0^\beta.$$

□

A.11.2 Proof of Corollary 8

Proof. The logarithm of M^{fb} is:

$$\ln(M^{fb}) = -\beta \ln(-\beta) - (1 - \beta) \ln(1 - \beta) + (1 - \beta) \ln(q) + \beta \ln(c_0),$$

and the logarithm of M^{sb} is:

$$\ln(M^{sb}) = -\beta(1 - \beta) \ln(-\beta) - (1 - \beta)^2 \ln(1 - \beta) + (1 - \beta) \ln(q) + \beta \ln(c_0).$$

Differentiating $\ln(M^{fb})$ with respect to β yields:

$$\frac{\partial \ln(M^{fb})}{\partial \beta} = \ln(1 - \beta) - \ln(-\beta) + \ln(c_0) - \ln(q),$$

and differentiating $\ln(M^{sb})$ with respect to β yields:

$$\frac{\partial \ln(M^{sb})}{\partial \beta} = (2 - 2\beta) \ln(1 - \beta) - (1 - 2\beta) \ln(-\beta) + \ln(c_0) - \ln(q).$$

Noting that $0 < q < c_0$ and $\beta < 0$, we have $\partial \ln(M^{fb})/\partial \beta > 0$ and $\partial \ln(M^{sb})/\partial \beta > 0$. Hence, M^{fb} and M^{sb} are increasing in β . Since β is decreasing in μ and ρ as well as increasing in σ , this completes the proof. □

A.11.3 Proof of Corollary 9

Proof. We now show that M^{fb} and M^{sb} converge to q as ρ approaches zero from the right if and only if the condition $\mu \leq \sigma^2/2$ is satisfied. The limiting value of β is as follows:

$$\lim_{\rho \rightarrow 0^+} \beta = \begin{cases} 0, & \text{if } \mu \leq \sigma^2/2 \\ 1 - 2\mu/\sigma^2, & \text{if } \mu > \sigma^2/2 \end{cases}.$$

Note that $\lim_{\beta \rightarrow 0^-} \theta = \lim_{\beta \rightarrow 0^-} (-\beta)^{-\beta} \cdot \lim_{\beta \rightarrow 0^-} (1 - \beta)^{-(1-\beta)}$, where $\lim_{\beta \rightarrow 0^-} (1 - \beta)^{-(1-\beta)}$ is clearly equal to one, and $\lim_{\beta \rightarrow 0^-} (-\beta)^{-\beta}$ is easily shown to be one by taking the logarithm and applying L'Hôpital's rule. Thus, the limiting value of θ is given by:

$$\lim_{\rho \rightarrow 0^+} \theta = \begin{cases} 1, & \text{if } \mu \leq \sigma^2/2 \\ (2\mu/\sigma^2 - 1)^{2\mu/\sigma^2 - 1} (2\mu/\sigma^2)^{-2\mu/\sigma^2}, & \text{if } \mu > \sigma^2/2 \end{cases}.$$

It follows that:

$$\lim_{\rho \rightarrow 0^+} M^{fb} = \begin{cases} q, & \text{if } \mu \leq \sigma^2/2 \\ (\lim_{\rho \rightarrow 0^+} \theta) q^{(1 - \lim_{\rho \rightarrow 0^+} \beta)} c_0^{(\lim_{\rho \rightarrow 0^+} \beta)}, & \text{if } \mu > \sigma^2/2 \end{cases},$$

$$\lim_{\rho \rightarrow 0^+} M^{sb} = \begin{cases} q, & \text{if } \mu \leq \sigma^2/2 \\ (\lim_{\rho \rightarrow 0^+} \theta)^{(1 - \lim_{\rho \rightarrow 0^+} \beta)} q^{(1 - \lim_{\rho \rightarrow 0^+} \beta)} c_0^{(\lim_{\rho \rightarrow 0^+} \beta)}, & \text{if } \mu > \sigma^2/2 \end{cases}.$$

Note that $c_0 > q > 0$ by assumption. In addition, $-\infty < \lim_{\rho \rightarrow 0^+} \beta < 0$ and $0 < \lim_{\rho \rightarrow 0^+} \theta < 1$ for $\mu > \sigma^2/2$. Hence, $0 < \lim_{\rho \rightarrow 0^+} M^{sb} < \lim_{\rho \rightarrow 0^+} M^{fb} < q$ if $\mu > \sigma^2/2$. \square

A.11.4 Proof of Corollary 10

Proof. Assume that $\mu > \sigma^2/2$, and let $\gamma = \mu/\sigma^2$. Then the ratio of $\lim_{\rho \rightarrow 0^+} M^{sb}$ to $\lim_{\rho \rightarrow 0^+} M^{fb}$ is given by:

$$\left(\lim_{\rho \rightarrow 0^+} M^{sb} \right) / \left(\lim_{\rho \rightarrow 0^+} M^{fb} \right) = \left(\lim_{\rho \rightarrow 0^+} \theta \right)^{-(\lim_{\rho \rightarrow 0^+} \beta)} = [(2\gamma - 1)^{2\gamma - 1} / (2\gamma)^{2\gamma}]^{-(1 - 2\gamma)}.$$

Taking the logarithm of $\lim_{\rho \rightarrow 0^+} \theta$, we obtain:

$$\ln(\lim_{\rho \rightarrow 0^+} \theta) = (2\gamma - 1) \ln(2\gamma - 1) - (2\gamma) \ln(2\gamma).$$

Differentiating with respect to γ , we have:

$$\frac{\partial \ln(\lim_{\rho \rightarrow 0^+} \theta)}{\partial \gamma} = 2[\ln(2\gamma - 1) - \ln(2\gamma)] < 0.$$

Hence, $\lim_{\rho \rightarrow 0^+} \theta$ is decreasing in γ . Moreover, $\lim_{\rho \rightarrow 0^+} \beta = 1 - 2\gamma$ is decreasing in γ . Noting that $0 < \lim_{\rho \rightarrow 0^+} \theta < 1$ and $-\infty < \lim_{\rho \rightarrow 0^+} \beta < 0$, the ratio is decreasing in γ . Since γ is increasing in μ and decreasing in σ , this completes the proof. \square