Sequential Implementation of Unenforceable Contracts with Stochastic Transaction Costs*

Yuichiro Kamada
Neel Rao †

July 15, 2009

Preliminary and Incomplete

Abstract

This paper analyzes the problem of bilateral trade of a divisible good in a continuous-time environment with stochastic transaction costs. In our basic setup, each of two parties is endowed with a durable good that is valued only by the other party, and the cost of transferring the good to the other party follows a geometric Brownian motion. The first-best solution to the problem requires each agent to transfer all of her good to the other agent when the transaction cost reaches a certain threshold value. However, in the absence of court-enforceable contracts, such a policy is not incentive-compatible, because an agent is unwilling to transfer her own good once she has received all of the other agent's good. This paper provides a closed-form solution for a second-best transaction scheme in which agents can realize some gains from trade as part of a subgame-perfect equilibrium. In the second-best solution, agents transfer the good in a sequential manner with the transfers becoming smaller over time. Moreover, as the discount rate approaches zero, the expected discounted payoffs from the first- and second-best policies converge, provided that the drift of the cost process is not excessively high relative to its volatility. In addition to obtaining a number of comparative statics results for the basic model, we prove an equivalence result stating that a model with decay or volatility in the amount of the good can be analyzed in terms of our basic framework. An implication of this result is that as long as there is some volatility in the amount of the good, positive gains from trade can be realized, even if the transaction cost is bounded away from zero.

*We thank Drew Fudenberg, Oliver Hart, and Michihiro Kandori for helpful discussions.
†Department of Economics, Harvard University, Cambridge, MA 02138 (Kamada and Rao). e-mail: ykamada@fas.harvard.edu (Kamada), rao@fas.harvard.edu (Rao).

1 Introduction

This paper is motivated by the following situation. Two firms that operate 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 resources spent training employees of the other firm or the expense of encrypting data to protect secrets from outsiders. Because of 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 the first party immediately reveals all of its information to the second party, then the second party would have no incentive to reveal its information to the first party. In this situation, how can the two parties exchange their knowledge with each other?

This paper studies the bilateral exchange of a divisible good in a continuous-time setting. Each of two agents possesses an equal amount of a good that is valued only by the other agent. In order to transfer some amount of her good to the other agent, an agent must pay a transaction cost that evolves according to a geometric Brownian motion. Although the first-best solution requires each agent to make a single transfer of all her goods to the other agent once the cost reaches a certain cutoff value, such a policy cannot be supported as part of a subgame-perfect equilibrium, because each agent is unwilling to transfer her own good once she has received all of the other agent's good.

The main objective of this paper is to solve for an incentive-compatible transaction scheme in which agents can realize positive gains from trade by making a sequence of gradually decreasing transfers. We obtain a closed-form solution that enables us to derive a number of comparative-statics results. Moreover, as agents become infinitely patient, the expected discounted payoff from the second-best policy converges to the efficient outcome, provided that the drift of the cost process is not excessively high relative to its volatility. We also prove an equivalence result stating that a model with a decaying or volatile good can be analyzed in terms of our basic setup. An important implication of this result is that positive volatility in the amount of the good enables agents to realize positive gains from trade, even if the cost process is bounded away from zero.

Our model is related to a line of literature on gradualism in contribution games and concession bargaining.[1] In these models, parties arrive at an agreement in a step by step fashion, and there is an efficiency loss due to delay in reaching an agreement. Likewise, cooperation between the two parties in our model is sustained through a gradual sequence of transactions over time. Nonetheless, our model differs from much of the literature in that transactions continue indefinitely in equilibrium; so that, our model cannot be solved using an iterated dominance procedure. The paper most closely related to ours is Pitchford and Snyder (2004). Those authors 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 of their model, however, positive investment can be supported as part of an equilibrium in which the seller's investment and the buyer's repayment take place alternately.

There are a number of differences between the model in Pitchford and Snyder (2004) and the framework developed in this paper. Pitchford and Snyder (2004) consider a discrete-time model with no uncertainty in the cost of transacting with the other party or the gains from trade that can be realized. By contrast, this paper analyzes a continuous-

[1]See Admati and Perry (1991); Compte and Jehiel (1995, 2003, 2004); Marx and Matthews (2000).

time setting that accommodates volatility in both the cost of transferring goods between parties and the quantities of goods available for trade. Furthermore, as Pitchford and Snyder (2004) observe, the equilibrium of their model is not robust to the inclusion of a fixed cost of making a transaction. If the seller must incur a fixed cost for investing instead of a variable cost that decreases to zero with the size of the investment, then an equilibrium with positive investment by the seller cannot be supported in an equilibrium of their model. Our model provides two ways of resolving this issue. If there is uncertainty in either the fixed cost of making a transaction or the quantity of goods available for trade, then agents might be able to realize positive gains from trade as part of a subgame-perfect equilibrium.

In addition, the efficiency properties of our model are somewhat more nuanced than those of Pitchford and Snyder (2004). Those authors demonstrate that as discounting frictions disappear, the solution to their model converges to the efficient outcome. In our model, the fixed cost of making a transaction acts as an additional friction in the bargaining process. Consequently, in the limit as the discount rate approaches zero, the solution to our model converges to the efficient outcome if and only if the drift of the transaction cost process is sufficiently low relative to its volatility. Even if this condition is satisfied, the efficiency of the equilibrium is not obvious given the results in Pitchford and Snyder (2004), because the timing of transactions in our model is endogenous. As agents become infinitely patient, not only does the cost of the first transaction decrease but the expected waiting time until this transaction increases. What our efficiency result demonstrates is that the positive effect of a lower discount rate dominates the negative effect of a higher waiting time.

A few models of finite-horizon games in continuous time have also been developed recently. Ambrus and Lu (2009) study a coalitional bargaining model, and Kamada and Kandori (2009) consider a revision process for strategies. In these models, opportunities to move arrive according to a Poisson process, and payoffs are realized only after an agreement or deadline is reached. In Kamada and Kandori (2009), the actions taken on the equilibrium path converge to the static best responses, because there is an increasingly small probability that other agents will receive another opportunity to move as the deadline approaches. Unlike much of the literature on gradualism, these models are similar to ours in that the gains from deviating progressively decrease over time, making it possible to support a nontrivial equilibrium.

Finally, our model is related to the large literature on repeated games. In standard models of repeated games, agents play a stage game for infinitely many periods, and the payoff structure for the stage game is typically constant over time.[2] For such models, the folk theorem states that any individually rational and feasible payoff can be achieved in a subgame-perfect equilibrium of the supergame if the discount rate is sufficiently low. Such equilibria are sustained by using the threat of future punishment to enforce cooperation in the current period. Likewise, the loss of future utility from deviating provides an incentive for agents in our model to exchange goods with each other. Nonetheless, there are two important differences between our model and the standard framework. First, in the case where the total supply of each good is fixed, the present discounted value of the relationship must eventually decrease with time if each agent is transferring a positive quantity

[2]See, for example, Fudenberg and Maskin (1986).

of her good. Hence, it is not immediately obvious that a non-degenerate equilibrium can be sustained in our model. Indeed, we obtain an impossibility result demonstrating that a positive expected discounted payoff cannot be supported in a subgame-perfect equilibrium of our model if the transaction cost is bounded below by a positive number and the quantity of each good available for trade is fixed. Second, the timing of moves is exogenously specified in standard models of repeated games, but in our model it is an endogenous strategic response to the realization of the cost process.

The remainder of this paper is organized as follows. Section 2 outlines our basic model of bilateral trade with stochastic transaction costs and defines the concept of subgame-perfect equilibrium in the context of our model. In section 3, we analyze the model in several steps. First, we prove an impossibility result stating that if (a) the transaction cost is bounded below by a positive number and (b) the quantities of goods available for trade are fixed, then there is no subgame-perfect equilibrium in which an agent obtains a positive expected discounted payoff. We then relax assumption (a) by considering the case where the transaction cost follows a geometric Brownian motion. In this case, we demonstrate the existence of a subgame-perfect equilibrium in which the agents obtain a positive expected discounted payoff by exchanging the goods through a sequence of transfers of decreasing size. We provide a closed-form solution for the equilibrium strategies and derive a number of comparative statics for the solution. Finally, we prove an efficiency result stating that the first-best expected discounted payoff can be approximated as the agents become infinitely patient, provided that the drift of the cost process is not excessively high relative to its volatility. In section 4, we relax assumption (b) by extending the basic model to allow the quantities of the goods to follow a geometric Brownian motion. We prove an equivalence result stating that the extended model can be analyzed in terms of our basic framework. If there is some volatility in the quantities of goods, then positive gains from trade can be realized in a subgame-perfect equilibrium, even if the transaction cost is a positive constant. All proofs, except the one for the main result (theorem 2), are given in the appendix.

2 Model of Stochastic Transaction Costs

There are two agents, A and B, 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 $s_t = [(s_t^{A1}, s_t^{A2}), (s_t^{B1}, s_t^{B2})]$, where s_t^{ij} denotes the amount of good j that agent i possesses at time t. The total supply $s > 0$ of each good is taken to be constant over time; so that, $s_t^{Aj} + s_t^{Bj} = s > 0$ for $j \in \{1, 2\}$ and $t \in [0, \infty)$.[3] The initial endowment vector is assumed to be $s_0 = [(s, 0), (0, s)]$; that is, agent A is endowed with all of good 1, and agent B is endowed with all of good 2. This assumption is without loss of generality provided that $s_0^{A1} = s_0^{B2}$. In addition, there is a transaction cost C_t for transferring goods between the two parties.[4] For ease of exposition but without loss of generality, the initial state C_0 of the transaction cost is assumed to be sufficiently high (e.g., $C_0 > s/\rho$).

[3]Section 4 extends the model to the case where the supply of each good can vary over time.

[4]In much of the analysis, the transaction cost C_t is assumed to evolve according to a geometric Brownian motion. The impossibility result of theorem 1 indicates that some restriction on the cost process is necessary to support a non-degenerate equilibrium.

In every instant of time, each agent observes the current realization of the cost and chooses an amount to transfer to the other agent. Formally, let H_t denote the set of all histories up to time t, and let $h_t = (\{C_\tau, s_\tau\}_{\tau\in[0,t)}, C_t)$ be a generic element of this set. By convention, let h_0 be the null history consisting of a singleton set. Denote the set of all histories by $H = \cup_{t\in\mathbb{R}_+} H_t$. Then a strategy for agent i is a function $\pi_i : H \to \mathbb{R}_+$ satisfying $\pi_i(h_t) \in [0, \hat{s}_t^{ij}]$ for $h_t \in H$ and $\hat{s}_t^{ij} = \lim_{\tau\to t^-} s_\tau^{ij}$ with $j = 1$ if $i = A$ and $j = 2$ if $i = B$.[5,6] Hence, the transfer made at time t is given by $x_t^{ij} = \hat{s}_t^{ij} - s_t^{ij}$ with $j = 1$ for $i = A$ and $j = 2$ for $i = B$. Finally, let Π_i with generic element π_i denote the set of all possible strategies for player i.

We restrict attention to symmetric strategies. In particular, we assume that if agent A transfers a certain amount of good 1 after history h_t, then agent B transfers an equal amount of good 2 after this history. If $\{s_t\}$ is the time path of allocations induced by a given strategy profile (π_A, π_B) and a realization of the cost $\{C_t\}$, then let $t(k)$ denote the time when the k^{th} transfer is made;[7] so that, the present discounted payoff received by agent i is given by:

$$U_i(\pi_A, \pi_B, \{C_t\}) = \sum_k \int_{t(k)}^{\infty} e^{-\rho t} s_{t(k)}^{ij} dt - \sum_k e^{-\rho t(k)} C_{t(k)},$$

where $j = 2$ if $i = A$ and $j = 1$ if $i = B$.

The equilibrium concept used in this paper is subgame-perfect equilibrium. Because there is no uncertainty regarding past play and past events, this is a sensible choice of equilibrium concept. Let Π^* denote the nonempty set of all subgame-perfect equilibria.[8]

Definition 1. A subgame-perfect equilibrium π is **maximal** if there is no $\pi' \in \Pi^*$ such that $\mathbb{E}[U_i(\pi'_A, \pi'_B, \{C_t\})] \geq \mathbb{E}[U_i(\pi_A, \pi_B, \{C_t\})]$ for each $i \in \{A, B\}$ with at least one inequality strict.

There are several reasons for focusing on a maximal equilibrium. First, the equilibrium is salient. Second, 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 on their preferred outcome as long as it does not violate the incentive constraints of either party. Third, restricting attention to such an equilibrium makes it possible to obtain strong comparative statics results and derive meaningful economic implications.

As stated in proposition 2, every SPE payoff can be achieved by an SPE involving grim-trigger strategies in which no further transfers are made if an agent deviates from the prescribed strategy profile. Hence, there is no loss of generality from assuming that agents use grim-trigger strategies.

[5] Note that the limit exists because s_t^{ij} is a monotonic function of time. Defining strategies in continuous time involves a number of difficulties in general, as discussed in Simon and Stinchcombe (1989). Nonetheless, the monotonicity of s_t^{ij} implies that agents take only a countable number of actions along any path of play; so that, this definition of strategies is innocuous.

[6] We restrict attention here to pure strategies. We do not expect that any of our conclusions depend on this restriction.

[7] Recall that only a countable number of transfers can occur along a given path of play.

[8] The strategy profile in which each agent never makes a transfer is an element of Π^*; therefore, the set Π^* is nonempty.

3 Analysis of Model

We begin with the following lemma that greatly simplifies the analysis as well as exposition.

Proposition 1. *Maximizing the payoffs of the agents subject to the cost process C_t is equivalent to maximizing $\sum_k x_k^{A2} e^{-\rho t(k)} - \sum_k e^{-\rho t(k)} c_{t(k)}$ and $\sum_k x_k^{B1} e^{-\rho t(k)} - \sum_k e^{-\rho t(k)} c_{t(k)}$ subject to the cost process c_t with $c_t = \rho C_t$.*

In other words, with an appropriate rescaling of the cost process, each agent can be regarded as consuming the good as soon as it is received. Because of its tractability, we work directly with the process c_t in our analysis, but we always keep in mind the original maximization problem involving C_t.

The following theorem is an impossibility result. If the transaction cost is bounded below by a positive number and the quantities of goods available for trade are fixed, then there is no equilibrium in which an agent receives a positive expected discounted payoff. Note that this result does not depend on assuming symmetric or maximal equilibrium strategies. Furthermore, the theorem only relies on the notion of iterated dominance as opposed to the more restrictive concept of subgame perfection. The restriction to pure strategies made above does not play any role in proving the theorem.

Theorem 1. *Assume that $\{\tilde{c}_t\}$ is an arbitrary cost process defined on the probability space $(\Omega, \mathcal{F}, P)$ and that each random variable $\tilde{c}_t$ for $t \geq 0$ takes values in the state space $(S, \mathcal{J})$ with $S \subset \mathbb{R}$ and $\inf(S) = k > 0$. Then any strategy profile surviving iterated dominance satisfies $s_t^{ij} = s$ for all t with $j = 2$ (resp., $j = 1$) if $i = A$ (resp., $i = B$).*

There are two approaches to resolve the issue above and to proceed with the analysis. Section 4 analyzes a model with uncertainty in the quantities of goods available for trade and demonstrates that the impossibility result does not hold in such a setting. In the current section, we relax the assumption that the transaction cost is bounded below by a positive number. Specifically, we assume that the cost process c_t follows a geometric Brownian motion $dc_t = \mu c_t dt + \sigma c_t dz_t$. The cost process is assumed to be non-degenerate (i.e., $\sigma^2 > 0$).

The following proposition provides justification for restricting attention to SPE in grim-trigger strategies. The basic idea behind the proof is that each agent can always obtain a continuation payoff of zero by transferring nothing. If an opponent uses a grim-trigger strategy, then zero is the maximum payoff that can be achieved after a deviation. Thus, the grim-trigger strategy is the most severe punishment available. The full proof of the result is given in appendix.

Proposition 2. *Given an arbitrary SPE, there exists an SPE in grim-trigger strategies that achieves the same continuation payoff after any history on the equilibrium path.*

In particular, this proposition implies that given an arbitrary maximal SPE, there exists a maximal SPE in grim-trigger strategies that achieves the same continuation payoff after any history on the equilibrium path.

Hereafter, we restrict attention to maximal SPE in grim-trigger strategies. We assume that the equilibrium has the form $\{c_{t(k)}, x_k\}_k$. This notation indicates that the k^{th} transaction occurs when the cost is $c_{t(k)}$ and that each agent transfers the amount x_k in that state. From now on, we sometimes abuse notation by writing $c_k = c_{t(k)}$ when what is meant by the subscript is clear.

The maximality assumption imposes the following restriction on the form of the SPE. The intuition for the result is that when a transaction takes place, the continuation value must equal the transaction cost incurred. Otherwise, the agents could achieve a welfare gain by decreasing the size of future transfers and increasing the amount currently transferred.

Proposition 3. *There is an infinite, positive, decreasing sequence* $\{c_k\}_{k=1}^{\infty}$ *of transaction costs paid with positive probability in any maximal SPE.*

We can now state the main theorem of this paper, which provides a closed-form solution for the model.

Theorem 2. *There exists a maximal SPE of the form* $z = \{c_k, x_k\}_{k=1}^{\infty}$ *where*

$$x_k = \frac{s}{1-\beta^-}\left(\frac{\beta^-}{\beta^- - 1}\right)^{k-1}$$

and

$$c_k = \left(\frac{s}{1-\beta^-}\right)\left(\frac{\beta^-}{\beta^- - 1}\right)^{k-\beta^-}$$

with

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

Proof. Given that the incentive constraint must bind for every index k, the optimization problem can be solved by generating an infinite sequence of asset-pricing equations. In particular, suppose that each agent has made a total of k transfers and that the current transaction cost c is greater than c_{k+1}. Note that $s_t^{ij} = s_{t(k)}^{ij}$. Then the value $V_i(c, s_t; \pi)$ of the relationship to each agent is the same as the value of an asset that pays x_{k+1} when the transaction cost is c_{k+1}. Because the cost c_{k+1} must be equal to the continuation value $V_i(c_{k+1}, s_{t(k+1)}; \pi)$, it is only necessary to consider the instantaneous payoff x_{k+1} when calculating $V(c, s_t; \pi)$. Thus, the solution involves maximizing over a sequence of individual asset-pricing problems. The Bellman equation for asset-pricing problem $k+1$ is given by the following for $c > c_{k+1}$:

$$\rho V_i\left(c, s_t; \pi\right) = \mathbb{E}(dV_i). \tag{1}$$

Our proof relies on a sequence of five lemmata. All proofs are given in appendix 5.5.

Lemma 1. *The unique solution to the Bellman equation (1) is*

$$V_i\left(c, s_{t(k)}; \pi\right) = x_{k+1}\left(\frac{c}{c_{k+1}}\right)^{\beta^-}$$

where

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

Thus, the optimal trading policy $z = \{c_k, x_k\}_{k=1}^{\infty}$ is a solution to the optimization problem:

$$\max_{c_1,\{x_i\}_{i=1}^{\infty}} x_1 \left(\frac{c}{c_1}\right)^{\beta^-}, \text{ subject to } \sum_{k=1}^{\infty} x_k \leq s \text{ and } c_k = x_{k+1}\left(\frac{c_k}{c_{k+1}}\right)^{\beta^-} \text{ for all } k \geq 1. \quad (2)$$

Given a technical condition that is satisfied by the solution, we can further simplify the maximization problem in (2) by eliminating each c_k.

Lemma 2. *Under the assumption that* $\lim_{n\to\infty}[\beta^-/(\beta^- - 1)]^n \ln(c_{n+k}) = 0$*, maximization problem (2) is equivalent to*

$$\max_{\{x_k\}_{k=1}^{\infty}} \sum_{k=1}^{\infty} \left(\frac{\beta^-}{\beta^- - 1}\right)^{k-1} \ln(x_k), \text{ subject to } \sum_{k=1}^{\infty} x_k \leq s. \quad (3)$$

Next, we solve the maximization problem in (3).

Lemma 3. *The solution to maximization problem* (3) *is*

$$x_k = \frac{s}{1-\beta^-}\left(\frac{\beta^-}{\beta^- - 1}\right)^{k-1}.$$

In turn, the solution for the x_k yields the solution for the c_k.

Lemma 4. *Under the assumption that* $\lim_{n\to\infty}[\beta^-/(\beta^- - 1)]^n \ln(c_{n+k}) = 0$*, the optimal cutoff rule* c_k *is given by*

$$c_k = \left(\frac{s}{1-\beta^-}\right)\left(\frac{\beta^-}{\beta^- - 1}\right)^{k-\beta^-}.$$

Using the solution for the c_k, we can confirm that the technical assumption made in lemma 2 is indeed satisfied by the solution.

Lemma 5. *The assumption that* $\lim_{n\to\infty}[\beta^-/(\beta^- - 1)]^n \ln(c_{n+k}) = 0$ *is satisfied by the solution* $z = \{c_k, z_k\}_{k=1}^{\infty}$.

□

The preceding theorem demonstrates the existence of an SPE in which agents realize some gains from trade. Note, however, that we have not yet established the uniqueness of the solution. In particular, there might be a solution to the Bellman equation that does not satisfy the technical condition assumed in lemma 2. We have not yet been able to show uniqueness, although we expect this conclusion to be likely. The proof is left to future work.

The closed-form solution for the model enables us to obtain a number of comparative statics results, which are helpful in understanding how changes in the parameters of the model affect the properties of the equilibrium. These results are enumerated in the corollaries below. Proofs are given in appendix 5.6 unless otherwise noted.

Corollary 1. $\lim_{k\to\infty} x_k = \lim_{k\to\infty} c_k = 0.$

That is, both the amount of each good transferred and the cost paid at each transaction tend to zero over time.

Corollary 2. *If $\mu < \frac{\sigma^2}{2}$, then $\lim_{t\to\infty} s_t^{A2} = \lim_{t\to\infty} s_t^{B1} = s$ with probability one.*

The proof simply involves calculating the sums of infinite series and is therefore omitted. This corollary states that with probability one the agents eventually transfer all of the goods available for trade, provided that the drift of the cost process is not excessively high relative to its volatility. As can easily be seen, this conclusion is the consequence of the maximality assumption. Given an SPE in which the amount of goods transferred is $r < s$, a Pareto superior SPE in which the agents transfer all of the goods, can be constructed by requiring the agents to transfer the additional amount $s - r$ at the first transaction.

Corollary 3. *If $k - 1 > |\beta^-|$, then x_k is increasing in μ and decreasing in σ^2. If $k - 1 < |\beta^-|$, then x_k is decreasing in μ and increasing in σ^2.*

This is an intuitive result. 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; so that, agents can make larger transfers at early stages and smaller transfers at later stages. If the volatility σ^2 of the cost process increases, then both extremely high and low 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^2 has a similar effect on the solution as a low drift μ.

Corollary 4. *For all k, c_k is decreasing in μ and increasing in σ^2.*

Intuitively, if the drift μ decreases, then the continuation value of the relationship rises; so that, higher cost payments can be elicited from each agent without violating her incentive constraints. As has already been noted, a high volatility σ^2 has an effect similar to a low drift μ.

Corollary 5. *$\lim_{\rho\to 0} x_1 = s$ and $\lim_{\rho\to 0} c_k = 0$ for all k. In addition, $\lim_{\rho\to\infty} x_k^{ij} = \lim_{\rho\to\infty} c_k = 0$ for all k.*

The proof is omitted, because it follows immediately from theorem 2. On the one hand, as the agents become infinitely patient, it is optimal for them 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 when the transaction cost is low and a high continuation value is not needed to prevent deviations. On the other hand, as the agents become infinitely impatient, it is impossible to induce them to incur a large transaction cost, because the continuation value from the relationship is low and the incentive to deviate is high. The size of each transfer must also be small, in order to ensure that the continuation value is sufficiently high to sustain cooperation.

Note that both extremely high and low values of the discount rate result in a very small cost cutoff c_k for each k. Thus, the expected waiting time to reach the allocation

$s^*(K) = [(s-K, K), (K, s-K)]$ for any fixed K with $0 < K < s$ is very high in these extreme cases. This suggests that the waiting time is non-monotonic in the discount rate. Letting $c^*(K)$ denote the value of c_t at which the sum of the transferred good (s_t^{A2} or s_t^{B1}) exceeds K, figures 1-4 plot $c^*(K)$ against β^- for $K \in \{10, 20, 30, 40\}$ and $s = 50$. Note that β^- is decreasing in μ and ρ as well as increasing in σ^2. In these figures, the expected waiting time rises as $c^*(K)$ decreases and becomes infinite as $c^*(K)$ approaches zero.

Since the waiting time becomes infinite as the discount rate approaches zero, it is not immediate that the first-best outcome can be approximated as agents become infinitely patient. The next theorem confirms that the second-best payoff converges to the first-best payoff in the limit as the discount rate approaches zero, provided that the drift of the cost process is not excessively high relative to its volatility.

Theorem 3. *In the limit as ρ approaches zero, the expected discounted payoff of each agent in the SPE of theorem 2 converges to the first-best outcome if and only if $\mu \leq \frac{\sigma^2}{2}$.*

On the one hand, if trading frictions are low in the sense that the drift of the cost process is sufficiently small relative to its volatility, then the first-best outcome can be approximated as the 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 proof of preceding theorem results in the following corollary.

Corollary 6. *If $\mu > \frac{\sigma^2}{2}$, then in the limit as ρ approaches zero, the ratio of the second-best payoff to the first-best payoff is decreasing in μ and increasing in σ^2.*

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.

4 Model with Stochastic Supplies of Goods

This section extends the basic model to allow for decay or volatility in the supplies of the goods available for trade. The setup here is the same as that described in section 2, except that the remaining stock of each agent's good is assumed to follow a geometric Brownian motion. Letting $j = 1$ if $i = A$ and $j = 2$ if $i = B$, we assume that s_t^{ij} evolves according to the process $ds_t^{ij} = -\delta s_t^{ij} dt + \xi s_t^{ij} dz_t$, where z_t is a Wiener process, $\delta > 0$ is the decay rate of each good, and ξ^2 is the volatility in the stock of goods. The following result is an equivalence theorem that characterizes the solution of the extended model in terms of the basic model with a fixed supply of each good.

Theorem 4. *The maximal SPE of a model in which each good decays at rate δ and has volatility ξ^2, is equivalent to the maximal SPE of a model with a constant supply of each good in which the discount rate is $\rho + \delta$ instead of ρ, the drift of the cost is $\mu + \delta$ instead of μ, and the volatility of the cost is $\sigma^2 + \xi^2$ instead of σ^2.*

The result is intuitive. A greater depreciation rate δ lowers the value of the relationship by increasing the effective discount rate from ρ to $\rho + \delta$ and the effective drift from μ to $\mu + \delta$. Greater volatility ξ^2 in the stock of goods raises the value of the relationship by increasing the effective volatility of the cost from σ^2 to $\sigma^2 + \xi^2$.

Three remarks are in order. First, as the proof in appendix 5.8 demonstrates, the theorem is still valid whenever $\rho + \delta > 0$; that is, the result holds even when there is growth in goods, provided that the growth rate is sufficiently low. An analysis of the case where $\rho + \delta < 0$ is left to future work, although we conjecture that the problem is not well defined in this case. In particular, there might exist an SPE that generates an infinitely high expected discounted payoff. Second, note that even in the case where the goods decay over time, transactions do not end in finite time. As the goods decay, the size of the transfers decreases accordingly to ensure that agents receive a sufficient continuation value after the current transaction is completed.

Finally, the result offers a counterpoint to theorem 1 in the following sense. Consider a cost process with nonnegative drift and no volatility. If the supply of each good is fixed over time, this cost process satisfies the conditions of theorem 1; so that, no transactions can occur in any SPE of the model. However, if there is some volatility in the stock of goods as in the model developed in this section, then our equivalence result states that such a model can be transformed into a model with volatility in the cost but no volatility in the goods; thus, positive transfers can be supported in equilibrium.

References

Admati, A. R., and M. Perry (1991): "Joint Projects without Commitment," *The Review of Economic Studies*, 58(2), 259–276.

Ambrus, A., and S. E. Lu (2009): "A Continuous-Time Model of Multilateral Bargaining," Mimeo, Harvard University.

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," *The Journal of Economic Theory*, 112(2), 334–342.

——— (2004): "Gradualism in Bargaining and Contribution Games," *The Review of Economic Studies*, 71(4), 975–1000.

Fudenberg, D., and E. Maskin (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54(3), 533–554.

Kamada, Y., and M. Kandori (2009): "Revision Games," Mimeo, Harvard.

Marx, L. M., and S. A. Matthews (2000): "Dynamic Voluntary Contribution to a Public Project," *The Review of Economic Studies*, 67(2), 327–358.

Pitchford, R., and C. M. Snyder (2004): "A Solution to the Hold-Up Problem Involving Gradual Investment," *The Journal of Economic Theory*, 114(1), 88–103.

Simon, L. K., and M. B. Stinchcombe (1989): "Extensive Form Games in Continuous Time: Pure Strategies," *Econometrica*, 57(5), 1171–1214.

5 Appendix

5.1 Proof of Proposition 1

Proof. Because transactions costs are paid only a countable number of times, s_t^{A2} and s_t^{B1} are step functions in time, which we define to be continuous from the right. Thus, we have:

$$\begin{aligned}
\int_0^\infty e^{-\rho t} s_t^{A2} dt - \sum_k e^{-\rho t(k)} C_{t(k)} &= \sum_k \int_{t(k)}^\infty e^{-\rho\tau}[s^{A2}(k) - s^{A2}(k-1)]d\tau - \sum_k e^{-\rho t(k)} C_{t(k)} \\
&= \sum_k [s^{A2}(k) - s^{A2}(k-1)][-\frac{1}{\rho} e^{-\rho\tau}]_{t(k)}^\infty - \sum_k e^{-\rho t(k)} C_{t(k)} \\
&= \frac{1}{\rho} \sum_k [s^{A2}(k) - s^{A2}(k-1)] e^{-\rho t(k)} - \frac{1}{\rho} \sum_k e^{-\rho t(k)} \rho C_{t(k)} \\
&= \frac{1}{\rho} \left[\sum_k x_k^{A2} e^{-\rho t(k)} - \sum_k e^{-\rho t(k)} c_{t(k)} \right],
\end{aligned}$$

where $s^{A2}(k)$ denotes the allocation at time $t(k)$. That is, $s_{t(k)}^{A2} = s^{A2}(k)$. In addition, we set $x_k^{A2} = s^{A2}(k) - s^{A2}(k-1)$ and $c_k = \rho C_k$.

Similarly, we have:

$$\int_0^\infty e^{-\rho t} s_t^{B1} dt - \sum_k e^{-\rho t(k)} C_{t(k)} = \frac{1}{\rho} \left[\sum_k x_k^{B1} e^{-\rho t(k)} - \sum_k e^{-\rho t(k)} c_{t(k)} \right],$$

with the obvious notational change.

The equations above imply that maximizing the agents' payoffs is equivalent to maximizing the discounted sum of amounts transferred (the x_k^{ij}) minus the discounted sum of transaction costs paid (the $c_{t(k)}$). $\square$

5.2 Proof of Theorem 1

Proof. We prove by induction that there are no SPE strategies in which some agent i makes a positive transfer x_t^{ij} at some history h_t. Consider any history h_t with the current allocation being $s = [(s_t^{A1}, s_t^{A2}), (s_t^{B1}, s_t^{B2})]$. If $s_t^{B2} < k$, then there is no SPE strategy where agent A makes a positive transfer at history h_t, because such a strategy would give agent A an expected discounted payoff no greater than $s_t^{B2} - k < 0$, whereas agent A could obtain an expected discounted payoff of at least zero by making no transfers after history h_t. Because there is no SPE strategy where agent A makes a positive transfer at some history h_t with $s_t^{B2} < k$, agent B would obtain an expected discounted payoff no greater than $-k < 0$ by making a positive transfer at a history h_t with $s_t^{B2} < k$ but would obtain an expected discounted payoff of zero by making no transfers after such a history. Thus, there is no SPE strategy where agent B makes a positive transfer at some history

h_t in which $s_t^{B2} < k$. A symmetric argument shows that there are no SPE strategies where agent A or B makes a positive transfer at some history h_t in which $s_t^{A1} < k$.

Suppose now that for some integer $n \geq 1$, there are no SPE strategies where agent A or B makes a positive transfer at some history h_t in which $s_t^{A1} < nk$ or $s_t^{B2} < nk$. Given this assumption, we show that there are no SPE strategies where agent A or B makes a positive transfer at some history h_t in which $nk \leq s_t^{A1} < (n+1)k$ or $nk \leq s_t^{B2} < (n+1)k$. Consider in particular a history h_t in which $nk \leq s_t^{A1} < (n+1)k$. If agent A is using an SPE strategy, then the greatest amount of the good that agent A can transfer at history h_t is $s_t^{A1} - nk$, because if agent A made a transfer greater than $s_t^{A1} - nk$ at history h_t, then the remaining amount s_τ^{A1} of agent A's good for $\tau > t$ would be such that agent B makes no further transfers, implying that agent A could obtain a higher expected discounted payoff by instead making no transfers after history h_t. Thus, if agent A is using an SPE strategy and history h_t is reached, then it must be that $s_t^{A1} \geq nk$ for every history h_τ with $\tau > t$; so that, the total amount transferred by agent A after history h_t is at most $s_t^{A1} - nk < k$.

It follows that there is no SPE strategy in which agent B makes a positive transfer after history h_t, because making a positive transfer would give agent B an expected discounted payoff no greater than $s_t^{B2} - (n+1)k < 0$, whereas agent B could obtain an expected discounted payoff of at least zero by making no transfers after history h_t. Thus, there is no SPE strategy where agent A makes a transfer at history h_t, because agent A would obtain an expected discounted payoff no greater than $-k < 0$ by making a transfer at history h_t but would obtain an expected discounted payoff of zero by making no transfers after history h_t. Thus, there are no SPE strategies where agent A or B makes a transfer at some history h_t in which $nk \leq s_t^{B2} < (n+1)k$. A symmetric argument holds for any history h_t in which $nk \leq s_t^{A1} < (n+1)k$. This completes the induction. □

5.3 Proof of Proposition 2

Proof.

Consider any SPE $\pi = (\pi_A, \pi_B)$. We construct an SPE using grim-trigger strategies $\pi' = (\pi'_A, \pi'_B)$ that achieves the same continuation payoff after any history h_t on the equilibrium path. Given the history h_t, let $u_i(h_t, \pi)$ denote the expected discounted payoff to player i from following the equilibrium strategy from time t onwards, and let $v_i(h_t, \pi)$ denote the supremum of the expected discounted payoffs to player i from any sequence of deviations. Then it must be that $u_i(h_t, \pi) \geq v_i(h_t, \pi)$, because π is an SPE.

Now consider the grim-trigger strategy π'_i which is the same as π_i on the equilibrium path induced by π and requires both players to stop making transfers if either player deviates. We show that π' is an equilibrium strategy profile. Given history h_t from above, the payoff to player i from following the equilibrium strategy from time t onwards is again $u_i(h_t, \pi)$; so that, $u_i(h_t, \pi') = u_i(h_t, \pi)$.

If agent $-i$ uses the grim-trigger strategy π'_{-i}, then a deviation at time t gives agent i an expected discounted payoff of at most zero, because agent $-i$ makes no further transfers. Furthermore, agent i can ensure that she receives an expected discounted payoff of zero by transferring nothing herself. Hence, $v_i(h_t, \pi') = 0$.

Now note that $v_i(h_t, \pi) \geq 0$ for all h_t because by transferring nothing agent i can always ensure that she receives an expected discounted payoff of zero. Thus, we have

$u_i(h_t, \pi') = u_i(h_t, \pi) \geq v_i(h_t, \pi) \geq 0 = v_i(h_t, \pi')$ for all $h_t \in H$. Hence, the strategy profile π' constitutes a SPE. Because $u_i(h_t, \pi') = u_i(h_t, \pi)$ for all $h_t \in H$, the equilibrium payoff $u_i(h_0, \pi)$ can be achieved using grim-trigger strategies. □

5.4 Proof of Proposition 3

Proof.

Given the current endowment vector s_t and the current transaction cost c_t, let agent i's value function $V_i(c_t, s_t; \pi)$ represent the expected present discounted value of her payoffs if both players following strategy profile π. First, we prove the following lemma.

Lemma 6. *In any maximal SPE, $c_k = V_i(c_k, s_t; \pi)$ and $x_{k+1} = V_i(c_k, \hat{s}_t; \pi)$ hold for every $k \geq 1$.*

Proof.

In order to be technologically feasible, a strategy profile π must satisfy the budget constraint $\sum_{k=1}^{\infty} x_k \leq s$. Moreover, incentive compatibility requires π to satisfy the constraint $c_k \leq V_i(c_k, s_{t(k)}; \pi)$ for every $k \geq 1$, which states that the cost of making each transfer must be less than the expected present discounted value of the relationship after the transfer has been made. We show that the inequality must be satisfied with equality if π is a maximal SPE.

Suppose to the contrary that $c_k < V_i(c_k, s_{t(k)}; \pi)$ for some index $k \geq 1$. Note first that we must have $c_l > 0$ for each index $l \geq 1$ in any non-degenerate solution, because there is zero probability that the cost process reaches the state $c_t = 0$ in finite time. Otherwise, choose l to be the smallest index such that $c_l = 0$. If $l \geq 2$, then it must be that $V_i(c_{l-1}, s_{t(l-1)}; \pi) = 0$ with $c_{l-1} > 0$, which violates the incentive constraint for transaction $l - 1$. If $l = 1$, then $V_i(c_t, s_t; \pi) = 0$ for all t. Now note that transfer k cannot be the final transaction. Otherwise, it must be that $c_k > 0 \geq V(c_k, s_{t(k)}; \pi)$, which violates the incentive constraint for transaction k. Furthermore, it must be that $x_{k+1} > 0$. Otherwise, if $x_{k+1} = 0$ and $c_{k+1} > 0$, then the agents could get a higher payoff by deleting (x_{k+1}, c_{k+1}) from π.

Thus, $c_k < V_i(c_k, s_{t(k)}; \pi)$ implies that the present value of the relationship could be increased by lowering x_{k+1} and raising x_k, because this perturbation would move a future payoff forward in time, as well as possibly increase the probability that the payoff is realized. In a maximal SPE, it must therefore be that $c_k = V_i(c_k, s_{t(k)}; \pi)$. It is immediate that $c_k = V_i(c_k, s_{t(k)}; \pi)$ implies $x_{k+1} = V_i(c_{k+1}, \hat{s}_{t(k+1)}; \pi)$. □

If follows from the proof of the preceding lemma that there is an infinite, positive sequence $\{c_k\}_{k=1}^{\infty}$ of transaction costs paid with positive probability in any maximal SPE π. We show that the sequence $\{c_k\}_{k=1}^{\infty}$ must be decreasing. Suppose to the contrary that there exists an index l such that $c_{l+1} \geq c_l$. Let m denote the least index greater than $l + 1$ satisfying $c_m < c_l$. Note that such an integer m must exist. Otherwise, the result that $c_k = V_i(c_k, s_{t(k)}; \pi)$ for all $k \geq 1$ implies that all of each good is exhausted before transaction $l + n$ takes place, where n is the least integer greater than s/c_l.

Now consider the grim-trigger strategy profile π' obtained from π by requiring agents to transfer $x_l + \sum_{k=l+1}^{m-1} x_k$ instead of x_l at transaction l and to transfer zero instead of x_k at each transaction $k \in \{l+1, \ldots, m-1\}$. In order to confirm that π' is an SPE, it suffices to show that the incentive constraint $c_l \leq V_i(c_l, s_{t(l)}; \pi')$ for transaction l is satisfied. Because π is an SPE, the incentive constraint for transaction $m-1$ must hold; so that, $c_{m-1} \leq V_i(c_{m-1}, s_{t(m-1)}; \pi)$. By construction, we have $V_i(c_l, s_{t(m-1)}; \pi) = V_i(c_l, s_{t(l)}; \pi')$. Noting that c_t follows a continuous Markov process, it must be that $V_i(c_{m-1}, s_{t(m-1)}; \pi) \leq V_i(c_l, s_{t(m-1)}; \pi)$ because $c_m < c_l \leq c_{m-1}$. All together, we obtain:

$$c_l \leq c_{m-1} \leq V_i(c_{m-1}, s_{t(m-1)}; \pi) \leq V_i(c_l, s_{t(m-1)}; \pi) = V_i(c_l, s_{t(l)}; \pi'),$$

confirming that the incentive constraint for transaction l under π' is satisfied. Moreover, each agent receives a higher expected discounted payoff under π' than under π, which contradicts the assumption that π is maximal. Thus, the sequence of transaction costs $\{c_k\}_{k=1}^{\infty}$ must be decreasing in any maximal SPE.

□

5.5 Proof of Lemmata in Theorem 2

5.5.1 Proof of Lemma 1

Proof.

A straightforward application of Ito's lemma yields:

$$\rho V_i\left(c, s_t; \pi\right) = \mu c \frac{\partial V_i\left(c, s_t; \pi\right)}{\partial c} + \frac{1}{2}\sigma^2 c^2 \frac{\partial^2 V_i\left(c, s_t; \pi\right)}{\partial c^2},$$

which provides a second-order linear differential equation for $V_i(c, s_t; \pi)$. Seeking a solution of the form $g(c; c_{k+1}, x_{k+1}) = B(c_{k+1}, x_{k+1})c^{\beta}$, the following quadratic equation is obtained by substituting the functional form into the differential equation:

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

whose solution is given by:

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

Letting β^+ and β^- respectively denote the positive and negative roots of the quadratic, the general solution to the differential equation is given by $V_i\left(c, s_t; \pi\right) = B^+(c_{k+1}, x_{k+1})c^{\beta^+} + B^-(c_{k+1}, x_{k+1})c^{\beta^-}$. It must be the case that $B^+(c_{k+1}, x_{k+1}) = 0$, because $V_i(c, s_t; \pi)$ would otherwise become unboundedly large in absolute value as c goes to ∞. Moreover, the boundary condition $V_i\left(c_{k+1}, s_{t(k)}; \pi\right) = x_{k+1}$ yields $B^-(c_{k+1}, x_{k+1}) = x_{k+1}/c_{k+1}^{\beta^-}$. Hence, the solution to the Bellman equation (1) is:

$$V_i\left(c, s_{t(k)}; \pi\right) = x_{k+1}\left(\frac{c}{c_{k+1}}\right)^{\beta^-}.$$

Uniqueness obtains because the solution to a second-order linear differential equation is generally unique. □

5.5.2 Proof of Lemma 2

Proof.
Note that each incentive-compatibility constraint in (2) can be rewritten as:

$$\ln(c_k) = \frac{\beta^- \ln(c_{k+1}) - \ln(x_{k+1})}{\beta^- - 1}.$$

Iterating the incentive-compatibility constraint n times starting with transfer k, it can be shown by induction that c_k satisfies:

$$\ln(c_k) = \left(\frac{\beta^-}{\beta^- - 1}\right)^n \ln(c_{n+k}) - (\beta^- - 1)^{-1} \sum_{j=0}^{n-1} \left(\frac{\beta^-}{\beta^- - 1}\right)^j \ln(x_{k+j+1}).$$

Assuming that $\lim_{n\to\infty}[\beta^-/(\beta^- - 1)]^n \ln(c_{n+k}) = 0$, $\ln(c_k)$ can be expressed as follows:

$$\ln(c_k) = -(\beta^- - 1)^{-1} \sum_{j=0}^{\infty} \left(\frac{\beta^-}{\beta^- - 1}\right)^j \ln(x_{k+j+1}).$$

Substituting for $\ln(c_1)$ in the maximization problem after taking the logarithm of the maximand, we obtain maximization problem (3). □

5.5.3 Proof of Lemma 3

Proof.
The Lagrangian is given by:

$$\mathcal{L} = \sum_{k=1}^{\infty} \left(\frac{\beta^-}{\beta^- - 1}\right)^{k-1} \ln(x_k) + \lambda \left(s - \sum_{k=1}^{\infty} x_k\right),$$

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

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

The budget constraint provides:

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

Thus, we obtain $\lambda = (1 - \beta^-)/s$, leading to the following sequence of transfers:

$$x_k = \frac{s}{1 - \beta^-} \left(\frac{\beta^-}{\beta^- - 1}\right)^{k-1}.$$

□

5.5.4 Proof of Lemma 4

Proof.

Substituting x_k obtained in lemma 3 into the last expression in the proof of lemma 2, we have:

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

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

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

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

Thus, the sequence of cost thresholds is given by:

$$c_k=\left(\frac{s}{1-\beta^-}\right)\left(\frac{\beta^-}{\beta^- - 1}\right)^{k-\beta^-}.$$

□

5.5.5 Proof of Lemma 5

Proof.

$$\begin{aligned}
&\lim_{n\to\infty}\left(\frac{\beta^-}{\beta^- - 1}\right)^n\ln(c_{n+k}) \\
&=\lim_{n\to\infty}\left(\frac{\beta^-}{\beta^- - 1}\right)^n\left(\ln\left(\frac{s}{1-\beta^-}\right)+(n+k-\beta^-)\ln\left(\frac{\beta^-}{\beta^- - 1}\right)\right) \\
&=\lim_{n\to\infty}\left(\frac{\beta^-}{\beta^- - 1}\right)^n\ln\left(\frac{s}{1-\beta^-}\right)+\lim_{n\to\infty}\left(\frac{\beta^-}{\beta^- - 1}\right)^n(n+k-\beta^-)\ln\left(\frac{\beta^-}{\beta^- - 1}\right) \\
&=\lim_{n\to\infty}\left(\frac{\beta^-}{\beta^- - 1}\right)^n(n+k-\beta^-)\ln\left(\frac{\beta^-}{\beta^- - 1}\right) \\
&=0.
\end{aligned}$$

□

5.6 Proof of Corollaries

5.6.1 Proof of Corollary 3

Proof.

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

Thus, $\partial x_k/\partial\beta^-$ is negative if $k-1 > |\beta^-|$ and positive if $k-1 < |\beta^-|$. In addition, we have $\partial\beta^-/\partial\mu < 0$ and $\partial\beta^-/\partial(\sigma^2) > 0$. Combining, we obtain the desired result. □

5.6.2 Proof of Corollary 4

Proof.

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

In addition, we have $\partial\beta^-/\partial\mu < 0$ and $\partial\beta^-/\partial(\sigma^2) > 0$. Combining, we obtain the desired result. □

5.7 Proofs of Theorem 3 and Corollary 6

5.7.1 Proof of Theorem 3

Proof. We begin by computing the solution to the second-best problem. From the proof of theorem 2, the value function $V_i^{sb}[c_0, s_0; \beta(\rho)]$ of agent $i \in \{A, B\}$ for the second-best problem is as follows:

$$
V_i^{sb}[c_0, s_0; \beta(\rho)] = \frac{f[\beta(\rho)]s}{[k(s)]^{\beta(\rho)}}c^{\beta(\rho)},\ f[\beta(\rho)] = \frac{1}{1-\beta(\rho)},\ k[s;\beta(\rho)] = \frac{s}{1-\beta(\rho)}\left(\frac{-\beta(\rho)}{1-\beta(\rho)}\right)^{1-\beta(\rho)},
$$

where $\beta(\rho) < 0$ is given by:

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

Note that the notation here differs from that used in the proof of theorem 2 in that the value function here depends on β instead of strategy profile π. In addition, let $\beta(\rho)$ denote the value of β given the discount rate ρ.

Substituting for $f[\beta(\rho)]$ and $k[s;\beta(\rho)]$ in the expression for $V_i^{sb}[s, c; \beta(\rho)]$ results in:

$$V_i^{sb}[c_0, s_0; \beta(\rho)] = \{a[\beta(\rho)]\}^{1-\beta(\rho)} s^{1-\beta(\rho)} c^{\beta(\rho)},$$

where the constant $a[\beta(\rho)] \in (0,1)$ is given by:

$$a[\beta(\rho)] = [-\beta(\rho)]^{-\beta(\rho)}[1-\beta(\rho)]^{-[1-\beta(\rho)]}.$$

We next compute the solution to the first-best problem. In the absence of incentive-compatibility constraints, the optimal policy requires the agents to make at most one transfer. Given that the agents use a stationary threshold policy, the value function $V_i^{fb}[c_0, s_0; \beta(\rho)]$ for the first-best problem for $i \in \{A, B\}$ can be obtained by finding the number k that maximizes the value $V_i^{fb}[c, s, k; \beta(\rho)]$ of an asset that pays $s - k$ when the cost reaches k starting from c. From the previous results, the value of such an asset is:

$$V_i^{fb}[c, s, k; \beta(\rho)] = (s-k)(k/c)^{-\beta(\rho)}.$$

Note that the optimal threshold k^* must lie in the interval $(0, s)$. Differentiating with respect to k yields the following first-order condition for k^*:

$$-\beta(\rho)(s-k^*)c^{-1}(k^*/c)^{-\beta(\rho)-1} - (k^*/c)^{-\beta(\rho)} = 0 \;\Rightarrow\; k^* = -\beta(\rho)[1-\beta(\rho)]^{-1}s,$$

where the derivative is positive for $k < k^*$ and negative for $k > k^*$. Thus, the value function for the first-best problem is:

$$V_i^{fb}[c_0, s_0; \beta(\rho)] = a[\beta(\rho)] s^{1-\beta(\rho)} c^{\beta(\rho)}.$$

We now show that $V_i^{sb}[c_0, s_0; \beta(\rho)]$ converges to $V_i^{fb}[c_0, s_0; \beta(\rho)]$ as ρ approaches zero from the right if and only if the condition $\mu \le \sigma^2/2$ is satisfied. The limiting value of $\beta(\rho)$ is as follows:

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

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

$$\lim_{\rho \to 0^+} a[\beta(\rho)] = \begin{cases} 1, & \text{if } \mu \le \sigma^2/2 \\ a(1 - 2\mu/\sigma^2), & \text{if } \mu > \sigma^2/2 \end{cases}.$$

It follows that:

$$\lim_{\rho\to 0^+} V_i^{sb}[c_0, s_0; \beta(\rho)] = \begin{cases} s, & \text{if } \mu \le \sigma^2/2 \\ [a(1-2\mu/\sigma^2)]^{2\mu/\sigma^2} s^{2\mu/\sigma^2} c^{1-2\mu/\sigma^2}, & \text{if } \mu > \sigma^2/2 \end{cases},$$
$$\lim_{\rho\to 0^+} V_i^{fb}[c_0, s_0; \beta(\rho)] = \begin{cases} s, & \text{if } \mu \le \sigma^2/2 \\ a(1-2\mu/\sigma^2) s^{2\mu/\sigma^2} c^{1-2\mu/\sigma^2}, & \text{if } \mu > \sigma^2/2 \end{cases}.$$

Note that if $\mu > \sigma^2/2$, then:

$$[a(1-2\mu/\sigma^2)]^{2\mu/\sigma^2} < a(1-2\mu/\sigma^2);$$

so that, $V_i^{sb}[c_0, s_0; \beta(\rho)]$ does not converge to $V_i^{fb}[c_0, s_0; \beta(\rho)]$ in this case. □

5.7.2 Proof of Corollary 6

Proof. Let $r = \mu/\sigma^2$. Then the ratio of the second-best value to the first-best value is given by:

$$[a(1-2r)]^{-(1-2r)}.$$

Since $a(\beta) = (-\beta)^{-\beta}(1-\beta)^{-(1-\beta)}$, taking the logarithm of both sides results in:

$$\ln[a(\beta)] = -\beta \ln(-\beta) - (1-\beta)\ln(-\beta).$$

Differentiating with respect to β, we have:

$$\frac{\partial \ln[a(\beta)]}{\partial \beta} = -\ln(-\beta) - 1 + \ln(-\beta) - \frac{(1-\beta)}{\beta} = -\frac{1}{\beta} > 0.$$

This implies that $a(\beta)$ is increasing in β. Hence, both $(1-2r)$ and $a(1-2r)$ are decreasing in r. Because $a(\beta) \in (0,1)$ and $(1-2r) < 0$, the ratio is decreasing in r. Since r is increasing in μ and decreasing in σ^2, this completes the proof. □

5.8 Proof of Theorem 4

Proof. In the case of decaying or volatile goods, consider an SPE π that has a threshold form in which the cutoff rule and the value function are both homogeneous of degree one; so that, $c(ks_t^{ij}) = kc(s_t^{ij})$ and $V_i(kc, ks_t^{ij}; \pi) = kV_i(c, s_t^{ij}; \pi)$ for any constant $k > 0$, where $j = 1$ if $i = A$ and $j = 2$ if $i = B$. Note that $c(s_t^{ij})$ represents the cost cutoff when the remaining amount of the good is s_t^{ij}.

The Bellman equation at state (c, s_t^{ij}) is given by:

$$\rho V_i(c, s_t^{ij}; \pi) = -\delta s_t^{ij} \frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial s_t^{ij}} + \frac{1}{2}\xi^2 (s_t^{ij})^2 \frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial (s_t^{ij})^2} + \mu c \frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial c} + \frac{1}{2}\sigma^2 c^2 \frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2}.$$

Because of the assumption that $kV_i(c, s_t^{ij}; \pi) = V_i(kc, ks_t^{ij}; \pi)$, differentiating with respect to k results in:

$$V_i(c, s_t^{ij}; \pi) = c\frac{\partial V_i(kc, ks_t^{ij}; \pi)}{\partial c} + s_t^{ij} \frac{\partial V_i(kc, ks_t^{ij}; \pi)}{\partial s_t^{ij}}. \tag{4}$$

Substituting $k = 1$ yields:

$$V_i(c, s_t^{ij}; \pi) = c\frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial c} + s_t^{ij}\frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial s_t^{ij}}. \tag{5}$$

Differentiating (4) with respect to k and setting $k = 1$, we have:

$$0 = c^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2} + 2cs_t^{ij}\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c \partial s_t^{ij}} + (s_t^{ij})^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial (s_t^{ij})^2}. \tag{6}$$

Now differentiate (4) with respect to c and set $k = 1$ to obtain:

$$\frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial c} = \frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial c} + c\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2} + s_t^{ij}\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial s_t^{ij} \partial c}, \quad \text{or}$$

$$0 = c\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2} + s_t^{ij}\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial s_t^{ij} \partial c}. \tag{7}$$

Substituting (7) into (6), we have:

$$0 = c^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2} - 2c^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2} + (s_t^{ij})^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial (s_t^{ij})^2}, \quad \text{or}$$

$$c^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2} = (s_t^{ij})^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial (s_t^{ij})^2}. \tag{8}$$

Substituting (5) and (8) into the original Bellman equation yields:

$$\rho V_i(c, s_t^{ij}; \pi) = -\delta\left(V_i(c, s_t^{ij}) - c\frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial c}\right) + \frac{1}{2}\xi^2 c^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2}$$

$$+\mu c\frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial c} + \frac{1}{2}\sigma^2 c^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2}, \quad \text{or}$$

$$(\rho + \delta)V_i(c, s_t^{ij}; \pi) = (\delta + \mu)c\frac{\partial V_i(c, s_t^{ij}; \pi)}{\partial c} + \frac{1}{2}(\sigma^2 + \xi^2)c^2\frac{\partial^2 V_i(c, s_t^{ij}; \pi)}{\partial c^2}.$$

□

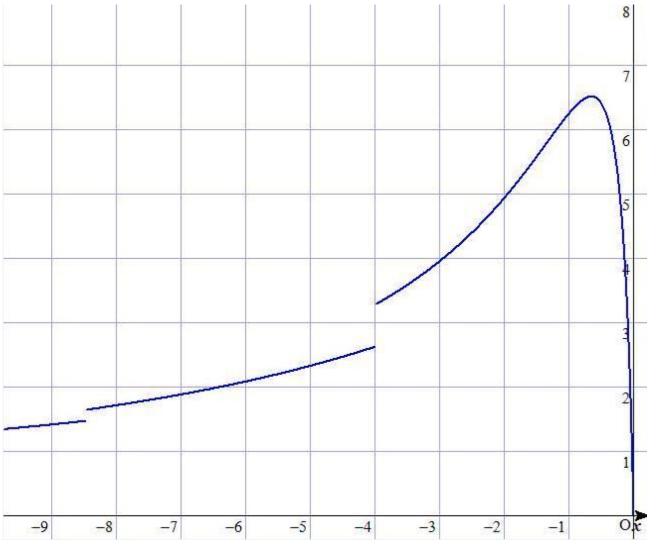

Figure 1: c*(K) vs. β for (K, s)=(10, 50).

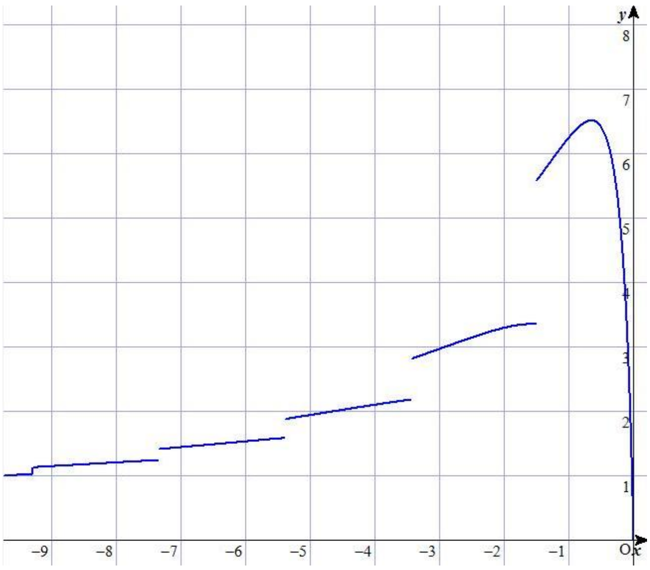

Figure 2: c*(K) vs. β for (K, s)=(20, 50).

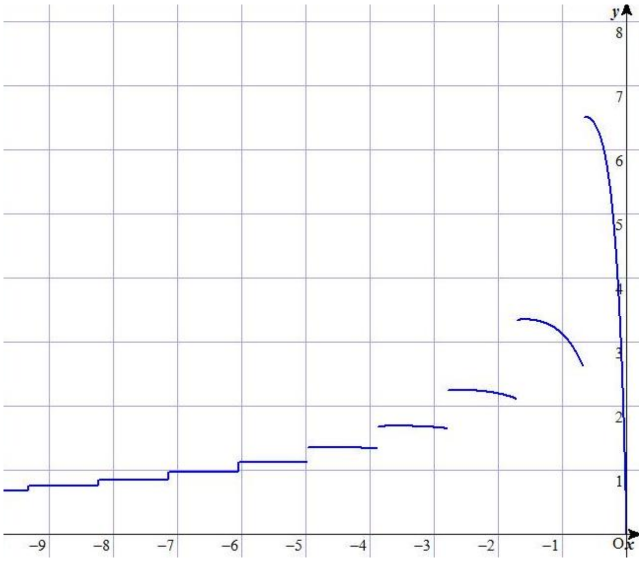

Figure 3: c*(K) vs. β for (K, s)=(30, 50).

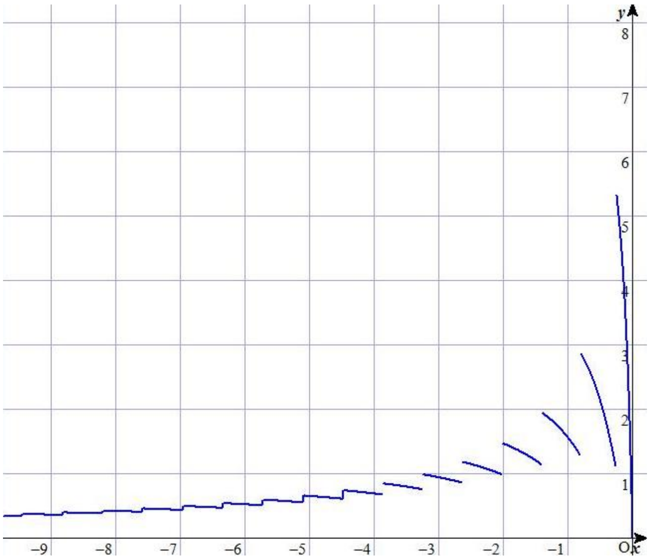

Figure 4: c*(K) vs. β for (K, s)=(40, 50).