

The Value of Temptation

Claudia V. Goldman^{1*}

clag@cs.umass.edu

Sarit Kraus^{2,3}

sarit@cs.biu.ac.il

¹Department of Computer Science

University of Massachusetts

Amherst, MA 01003

ph.: 1-413-5451985

²Department of Computer Science

Bar-Ilan University, Ramat-Gan 52900, Israel

³Institute for Advanced Computer Studies, University of Maryland

College Park, MD 20742

*This work was performed while the author was a post-doctoral researcher at Bar Ilan University. This material is based upon work supported in part by NSF under grant number IIS9907482 and IIS0222914. The address for correspondence is: Claudia V. Goldman - 140 Governor's Dr. - Department of Computer Science - University of Massachusetts - Amherst, MA 01003 - ph.: 413-5451985

Abstract

There is an implicit assumption in electronic commerce that induces the buyers to believe that their deals will be handled appropriately. However, after a seller has already committed to a buyer, he may be tempted by several requests though he will not be able to supply them all. We analyze markets in which a finite set of automated buyers interacts repeatedly with a finite set of automated sellers. These sellers can satisfy one buyer at a time, and they can be tempted to break a commitment they already have. We have found the perfect equilibria that exist in markets with a finite horizon, and with an unrestricted horizon. A significant result stemming from our study reveals that sellers are almost always tempted to breach their commitments. However, we also show that if markets' designers implement an external mechanism that restricts the automated buyers actions, then sellers will keep their commitments.

Keywords: E-commerce, Perfect equilibrium, software agents.

Marketplaces, either traditional or electronic, are sites in which sellers and buyers meet in order to trade. The trade could be for goods or for services. In both cases, the demand may exceed the supply and a seller that has already committed to a buyer, may be tempted by several profitable requests of other buyers that he will not be able to supply if he keeps his promises. The seller, that would like to maximize his benefits faces the strategic question of whether to keep his promise or to break it and to increase his short term benefits. In markets where trade continues repetitively, buyers should take into consideration the long term effects of trading with sellers who may break their promises. The buyers that would also like to maximize their benefits need to decide how to react when a seller breaks a promise.

Such situations occur in B2C electronic markets where usually the consumers are unable to verify the actual stock of the business with whom they are trading. In addition, in such electronic markets, there is an implicit assumption which induces the buyers to believe that whenever a deal is made between a buyer and a seller, it will be performed as agreed. However, usually there is no binding contract between the parties. After a buyer has successfully placed his order, he gets the feeling that the seller to whom he has approached is, although implicitly, *committed* to performing the request as submitted. Nevertheless, it may be the case that after a seller of a service or a product has already committed to a buyer, he may be tempted by several requests though he will not be able to supply them all. Usually, inventories are not checked before the buyer places his submission, and he still receives a successful message regarding his request. Later on, the same buyer may realize that he will not actually receive the order as it has originally been placed.

We interviewed several customers of on-line stores and collected many examples including the following. A customer placed an electronic order through a web-site of a pharmacy that advertised its products on sale. However, the number of products he actually received was smaller than those that he had successfully ordered. Another example, is of an order with an extremely large number of books with the same book-title that was successfully placed (by mistake) in one of the electronic book stores. The customer did not receive any comment on the availability of the number of books. Although, the store may still have been able to process such a large number of books, it seems that its inventory was not checked while the customer was placing his order. There is a gap between the belief of the customer that his order will be processed completely and successfully, and the actual result. This gap results from the *implicit* contracts existing in B2C transactions, and the fact that in electronic markets the customers are not able to *see* what is actually happening on the sellers' side.

Similar situations occur also in traditional markets. For example, a plumber may have committed himself to serve a certain customer. However, he may be tempted by a larger customer and may change his plans at the last moment. Thus he may not keep his promise to the original customer by not coming on the agreed day and time.

Other related scenarios occur in physical stores when the merchandise may be kept in a stock room and thus it is not visible to the customers. During our interviews we heard the following example. A salesperson of a cellular phone company promised a new model of a phone to a customer. However the delivered appliance was not of the promised model and the claim was that they do not have it in their stock. Still, this behavior is more common in electronic markets where sellers are uncertain regarding the number of buyers that will approach their site and the volume of their orders. It results from liquidity and dynamism of

the activity on the web, where sites are easily accessible by buyers from diverse geographical locations, and the cost of switching from one site to another is very low.

When the sellers receive multiple orders for the same product which they may be unable to supply to all of the customers, it is not clear what strategies these sellers should use to solve these conflicts. One option is to have sellers respect the customers' order, and implement the strategy *first in first served* which is the most commonly used in physical stores due to the fair feeling it gives the customers. But when the customers themselves can not verify what was the order of placing their orders, it is still not clear whether the sellers should benefit from this strategy. Since there is no explicit contract that binds the seller to actually fulfilling an electronic order, the question of which are the best strategies to use in such markets should not be considered trivial.

Intention reconciliation was studied in the framework of teams [Sullivan et al., 1999] which is not the case of B2C interactions. The notion of reputation was studied in [Castelfranchi et al., 1998] assuming that the reputation values are communicated among the interacting parties. This is not the case that we are studying, where competing buyers and sellers rarely communicate between themselves. We focus on the strategies that will result in equilibrium of electronic marketplaces that do not allow for additional payments beyond those that the buyers are requested to pay for the products or services purchased. We also assumed that the price of similar products is the same for all the buyers. That is, we did not analyze buyers with different types (see [Goldman et al., 2001]) and we do not allow buyers to pay higher prices for the same products in order to guarantee the delivery of the product to themselves. We consider situations where there is no social law that will bind the sellers to keep their commitments. In particular, we do not consider markets where there is a mechanism to enforce sellers to pay penalties for breaching [Sandholm et al., 1999].

We are interested in studying the effect of taking into account the reputation of the sellers at the design stage, i.e., how we should design buyers' and sellers' strategies in order to implement them to act on behalf of human users in B2C interactions where each user would like to maximize his own benefits. Economic studies on B2C (e.g., [van der Heijden et al., 2000, Schmitz, 2002] and others (as reported in section 5)) study certain markets that do not fit the problem we handle here. In particular the question of how to design strategies for automated buyers and sellers where the interactions may not be trusted has not been studied to the best of our knowledge.

Therefore, we are particularly interested in studying the equilibria that exist in such B2C situations. Are these markets self-stable? Are the sellers trustful? Based on the information that resides in the web-site where a buyer places his order, can this buyer be certain that if he receives a message describing the successful status of his order submission he will obtain the expected result from this order? The answer to the latter question is, unfortunately, negative, as we will show in this paper. In other words, current electronic sellers on the Internet have the potential not to be trustful. That is, following their equilibrium strategies, the sellers will indeed breach their implicit contracts with their customers if they happen to receive better deals, even after a submission for an order has already been successfully processed. However, we show that this trust problem can be avoided by adding an external regulator to the market. By a regulator we mean a mechanism that regulates which seller a buyer can approach.

Given an external mechanism as presented in this paper, we show that the equilibrium strategies will induce the sellers *not to breach* their implicit commitments. The external mechanism transforms the sellers' breaches from beneficial to non-beneficial, and therefore the sellers will prefer to remain reliable.

We formally specified conditions for which equilibria exist in markets composed of finite sets of sellers and buyers who interact repeatedly. The seller must decide whether to keep its promise that may lead to some losses, or whether to break the promise and increase its short term gains, but irritate one of its clients. A buyer that is irritated by a seller is motivated not to approach the irritating seller in succeeding encounters. This behavior is considered a punishment for not providing the expected supply or service.

In principle, there could be two possible equilibria:

- The buyers will not punish sellers that do not keep their promises. In this case the sellers will most likely be tempted by a larger deal, and will therefore break any commitment they have.
- The buyers will punish sellers that do not keep their promises. The sellers may refrain from breaking commitments if the loss of a buyer's deal in the future does offset the increment in the immediate gain the seller obtains when he is tempted.

Our main result is that in most of the cases the first equilibrium exists. We showed this for any probability with which a tempting deal can occur, assuming the characteristics of the market in question. This is true for a general class of buyers that we have characterized. There are other classes of buyers, for whom being constantly rejected is very costly. In this case, buyers may choose not to return to sellers that do not keep their promises. However, such a market is not stable, and an equilibrium does not exist. For maintaining stability in such markets, or to maintain stability while the sellers keep their promises we propose an external regulator that will be implemented as a mechanism that binds the potential buyers' threats not to return to an unreliable seller. Practically, this means that the market's manager can prevent a certain buyer from approaching a certain seller who has broken a commitment to this buyer in the past. Markets can store log files with the history of interactions maintained among the buyers and sellers that trade in these markets. We are assuming that binding the sellers not to breach their commitments is not possible. We assume that there is no social law that can be imposed on an electronic competitive market that will tell a seller not to sell to a certain buyer. In this paper we deal with a market in which the sellers may behave unreliably.

The paper is organized as follows: The model is first presented in section 1. The strategies studied shed light on the sellers' decision as to whether or not it is worthwhile for them to be tempted to leave a deal they are already committed to. The profiles of strategies in perfect equilibrium that were found are detailed in section 3 following the main theorem presented in section 2. In section 4, we propose the implementation of an external mechanism that results in a stable market with reliable sellers and present formal results for a simple market. Work conducted on trust, reconciliation and reputation functions in the framework of teams and economies is surveyed in section 5. Finally, we provide our conclusion in section 6. The formal detailed proofs of the theorems can be found in the appendixes.

1 The Model and its Dynamics

The formal model was simplified as much as possible specifically to concentrate on the factors that are needed for studying the temptation problem. There is a finite set \mathcal{B} of buyers B^i and a finite set \mathcal{S} of sellers S^j that interact repeatedly. Each buyer in \mathcal{B} would like to buy a *single product* that the sellers in \mathcal{S} can sell (we refer to it as a product, but it can also represent a service the seller can provide; e.g., assisting an Internet user from a help-desk) for a fixed price p . Each seller has a non-restrictive amount of the product, but can serve only one customer at a time. This limitation is the cause of the problem handled in this paper.¹

In each time period a buyer would like to buy a given number of units of the product and will send a request to one of the sellers. $\beta(B^i, \tau, tc)$ will denote a request of buyer B^i to buy τ units at the unit price p at trading cycle tc . When both parties agree for every cycle upon a request, we will refer to it as a deal. In this paper we assume that a request is valid for only one time cycle and its values do not change during its time cycle. We also assume that the size of a request, τ , is bounded and denoted $\bar{\tau}$.

1.1 The Dynamics of the Trade

Each buyer could approach any seller for a given request. If a seller, who was approached by a buyer, agrees to sell², then the buyer will indeed buy from this seller.

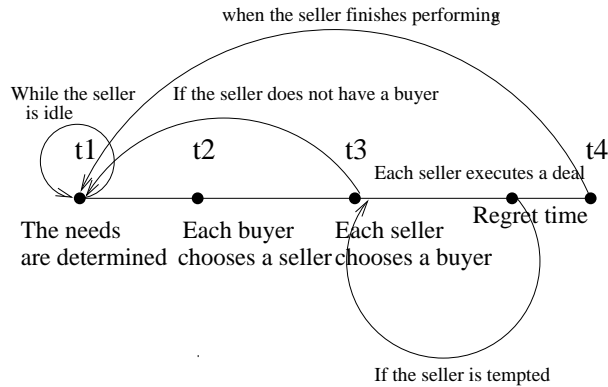


Figure 1: A Trading Cycle tc .

Figure 1 shows the steps followed by the trading agents in one trading cycle. At time t_1 , the needs of each buyer are determined, i.e., $\beta(B^i, \tau, tc)$. We assume a positive β exists for every buyer for every trading cycle tc . By time t_2 , each buyer chooses a seller to whom he will send his request. At t_3 , each seller sends each of his possible buyers an answer either agreeing to perform β or rejecting the request.

When a buyer receives a negative answer, he may approach another seller. This can be done at time $t_{3+\epsilon}$ that will be called the *regret time*. We use $\hat{\gamma}$ to denote such a request asked

¹This is one of our simplification assumptions. In general, even in cases in which sellers can serve several buyers simultaneously, such a problem arises.

²The sellers decide according to their utility that is based on the size of the buyers' requests.

Notation	Description
tc	The trading cycle. tc_i will denote the i^{th} trading cycle,
K	The number of trading cycles in a market with a finite horizon
K'	$K' = K - tc$, i.e., K' denotes the number of cycles that are left until the end of the market.
$\mathcal{B} = \{B^i\}, 1 \leq i < \infty$	The set of buyers.
$\mathcal{S} = \{S^j\}, 1 \leq j < \infty$	The set of sellers. The problem is relevant when $ \mathcal{B} > \mathcal{S} $.
$\beta = \beta(B^i, \tau, tc)$	B^i requests a deal β composed of τ units of the product at trading cycle tc .
$\bar{\tau}$	The maximal number of units in a request.
$Val(\beta)$	The value of β is $Val(\beta) = \tau * p$.
$\hat{\gamma}$	The deal offered to a seller at regret time. $Val(\hat{\gamma}) > Val(\beta)$.
Unreliable seller	A seller is unreliable (with respect to a buyer to whom he was committed to perform this buyer's β deal) if he accepts $\hat{\gamma}$ at regret time.
$U^{S^j}(\beta)$	The seller S^j 's utility function. $U^{S^j}(\beta) = Val(\beta) = \tau * p$.
$U^{B^i}(\beta)$	The buyer B^i 's utility function.
δ	The time discount factor.
$Pr(\hat{\gamma})$	The probability that a seller will be tempted to break its promise in a given time period.

Table 1: A Summary of all the Notations Defined in this Section.

at regret time. A seller may prefer $\hat{\gamma}$ over the original β he is committed to. The dilemma occurs only when the benefits of the seller from its original β are lower than that of $\hat{\gamma}$.

Up to time t_4 , each seller fulfills the deal to which he has committed himself, either at t_3 or at $t_{3+\epsilon}$. A buyer can decide to leave the market at time t_1 of any trading cycle. This may happen when the buyer is rejected by all the sellers in the market, and when the buyer does not return to such sellers.

In this paper we consider the case in which only one regret time exists between t_3 and t_4 . We will consider situations of finite horizon where the number of trading cycles, K , is known to all the buyers and to all the sellers. We will also consider situations of unrestricted horizon where there is no limit on the number of trading cycles. We will denote the i 's cycle tc_i .

Table 1 presents a summary of our model's notations.

1.2 Utility functions

Each seller and each buyer has a utility function denoted U^{S^j} and U^{B^i} , respectively. The agents have a time constant discount rate $0 \leq \delta \leq 1$. When considering finite horizon situations, we will assume that $\delta = 1$.

A seller's utility from a request $\beta(B^i, \tau, tc)$ is linear in $\tau * p$. Thus, we will refer to $\tau * p$ as the "value" of β for the sellers and we will denote it $Val(\beta)$.

The maximal utility a buyer can attain is when his request is satisfied. In addition he prefers not receiving a commitment at all over reaching a deal with a seller that will break its promise. Hence, the buyers' utility satisfies the following inequalities: $U^{B^i}(\beta \text{ is performed by } S^j) >$

U^{B^i} (β is not performed but no commitment is made) $>$
 U^{B^i} (a seller promised to do β but later broke its commitment).

The utility a buyer attains when he is rejected by a seller that was committed to him is denoted as U_{reg} and his utility when he leaves the market is denoted U_{exit} .³

For simplicity, we assume that a buyer's utility function is as follows:

$$U^{B^i}(\beta) = \begin{cases} 1 & \text{if } \beta \text{ was performed during the trading cycle.} \\ 0 & \text{if } B^i \text{ was not chosen by the seller he approached either at } t_3 \text{ or at } t_{3+\epsilon}. \\ U_{reg} & \text{if } B^i \text{ was rejected by the seller at regret time.} \\ U_{exit} & \text{if } B^i \text{ opts out of the market.} \end{cases}$$

For example, U_{reg} may be $-\frac{1}{2}$ and U_{exit} may be equal to -1 .

1.3 Strategies and Equilibrium

A seller's strategies differ as to whether or not he keeps his promise to the buyer to whom he has committed himself in the current trading cycle. That is, a seller performs β or breaks his promise and performs the request $\hat{\gamma}$, it obtained at regret time. This question is relevant when $Val(\hat{\gamma}) > Val(\beta)$ and thus we will discuss the sellers' strategies only in such situations.

On the other hand, the buyers' strategies differ as to whether or not the buyer returns to a seller that has broken his promises. Given that the buyers do not come back to a seller who has behaved unreliably, we distinguish among three types of buyers: not restricted buyers, exclusive buyers, and loyal buyers.

Definition 1 Not restricted, exclusive, and loyal buyers:

- *A buyer is not restricted from approaching a seller if he may choose this seller at time t_2 . The buyer's strategy does not prevent him from approaching this seller.*
- *A buyer is exclusive to a seller j , if his strategy prevents him from approaching any seller i , $i \neq j$.*
- *A buyer is loyal to a seller if he keeps choosing the same seller after the seller has chosen him and performed the requested β . If the seller obtains equal valued deals, he would prefer the deal of his loyal buyer.*

In the following sections we consider the problem of how a rational agent chooses its strategy. A useful notion is the Nash equilibrium [Nash, 1953] that requires that if all the agents use the strategies specified for them in the strategy profile of the Nash equilibrium, then no agent is motivated to deviate and implement another strategy. However, in a market with multiple cycles such as the one considered in this paper some absurd Nash equilibrium may exist: an agent may use a threat that would not be carried out if the agent were put in the position to do so, since the threat move would give the agent a lower payoff than it would get by not doing the threatened action [Fudenberg and Tirole, 1992]. This happens since Nash

³A buyer may need to leave the market if he does not return to unreliable sellers, and all the sellers he has approached have broken their commitments to this buyer throughout the trade period.

equilibrium strategies may be in equilibrium only at the beginning of the market, but may be unstable in intermediate stages. Thus, we will also use the concept of the sub-game perfect equilibrium (SPE) [Osborne and Rubinstein, 1994, Osborne and Rubinstein, 1990] which is a stronger concept. SPE requires that in any cycle of the market, no matter what the history is, no agent has a motivation to deviate and follow another strategy different from that defined in the strategy profile.

2 Buyers Return to Unreliable Sellers

It seems that punishing an unreliable seller by not returning to him after he breaks his promise is a good strategy for a buyer. However, when the number of sellers is finite, as considered in this paper, this strategy may cause the buyer to leave the market if all the sellers have irritated him. In this section we characterize the situations in which it is beneficial for a buyer to return to unreliable sellers. The characterization depends mainly on the utility function of the buyers, but also on the probability that a seller will be tempted to break its promise. This probability depends on the possible values of the size of the requests, i.e., the possible values of τ , and the way the needs of the buyers are determined at the beginning of a trading cycle. This is because the origin of the temptation at the regret time is caused by the buyers that were not chosen by the seller they approached initially. We denote this probability $Pr(\hat{\gamma})$ and demonstrate it in the following example.

Example 1 Consider a market of three buyers and two sellers, i.e., $\mathcal{B} = \{B^1, B^2, B^3\}$ and $\mathcal{S} = \{S^1, S^2\}$. The possible values of τ are 1, 2, 3, and 4. Suppose also, that the needs of each buyer at the beginning of a trading cycle are uniformly and independently distributed over $\{1, 2, 3, 4\}$ and that a buyer chooses randomly between the two sellers if he is not restricted from approaching either of them. We will show that $Pr(\hat{\gamma}) = 7/32$.

There are 64 combinations of 3 requests of the three buyers over $\{1, \dots, 4\}$. A seller will be tempted by a buyer if (i) originally only one buyer approached him and two buyers have approached the other seller; (ii) the request of the buyer that approached him was lower than the request of the minimum between the requests received by its opponent. For example, if he obtained a request of 1 and its opponent received requests of 3 and 4. There are 14 such cases that we enumerate in the following list (the first element in the pair, is the request obtained by the seller and the second element in the pair is an ordered list of the requests sent to its opponent) $[\{1\}\{2,2\}]; [\{1\}\{2,3\}]; [\{1\}\{2,4\}]; [\{1\}\{3,2\}]; [\{1\}\{3,3\}]; [\{1\}\{3,4\}]; [\{1\}\{4,2\}]; [\{1\}\{4,3\}]; [\{1\}\{4,4\}]; [\{2\}\{3,3\}]; [\{2\}\{3,4\}]; [\{2\}\{4,3\}]; [\{2\}\{4,4\}]; [\{3\}\{4,4\}]$.

Thus, in 14 cases out of the 64 possible cases a seller will be tempted, i.e., $Pr(\hat{\gamma}) = 7/32$.

The next theorem specifies the conditions in which buyers will return to unreliable sellers. It serves as the basis for the perfect equilibria that we identify in the next section. Remember that in a market with a finite horizon, the trade lasts for K trading cycles and in a market of unrestricted horizon the buyers' and sellers' utility function have a time constant discount rate denoted δ .

Theorem 1 (Sellers can be Unreliable)

1. **The buyers will return to unreliable sellers** — *If the market satisfies the following conditions:*

- **Finite Horizon:** $U_{reg} * (1 + Pr(\hat{\gamma})(K - 2)) > U_{exit} + (K - 2)$
- **Unrestricted Horizon:** $\frac{1}{1-\delta} * [U_{reg} * (Pr(\hat{\gamma}) - 1) - 1] > U_{exit}$

thus if all the sellers in the market are unreliable; i.e., they accept any $\hat{\gamma}$ deal at regret time, then a buyer B will benefit more from returning to an unreliable seller, than from not returning, regardless of the strategies of the other buyers.

2. **The sellers are unreliable with respect to returning buyers** — *If the buyers always keep themselves not restricted from approaching any seller in the market, then a seller's expected utility increases when he accepts a $\hat{\gamma}$ deal at regret time, regardless of what the strategy of the other sellers are. This holds when the horizon is finite as well as when it is unrestricted.*

Sketch of Proof.

1. **The Buyers** — We show that when U_{reg} and U_{exit} follow the relation stated in the theorem, the expected utility of a buyer that returns to an unreliable seller is larger than the expected utility it could have attained, had he not returned to unreliable sellers. The full details of these computations appear in Appendix A.1 for a market with a finite horizon and in Appendix A.2 for a market with an unrestricted horizon.
2. **The Sellers** — Since the behavior of the buyers does not vary with respect to the behavior of the sellers, and since $Val(\hat{\gamma}) > Val(\beta)$, it is clear that the expected utility of a seller who accepts a $\hat{\gamma}$ deal at regret time is larger than the expected utility he obtains if he rejects such a deal. Since the buyers keep themselves not restricted from approaching any seller, this theorem is true, regardless of what the other sellers do.

□

If the buyers' utility function does not follow the inequalities stated in the previous theorem, the market will not remain stable. Rational buyers will not respect their own threats of not returning to unreliable sellers. We demonstrate the theorem in the next two examples.

Example 2 *We return to the market of example 1 composed of three buyers, two sellers and request values uniformly distributed over $\{1, 2, 3, 4\}$. Suppose the market lasts for 20 trading cycles, i.e., $K = 20$. In this case if $U_{reg} = -1$ and $U_{exit} = -23$ then buyers will benefit from returning to unreliable sellers. If the horizon of the market is unrestricted, such buyers will also be induced to return to unreliable sellers when $\delta < 0.99$.*

In the next example we consider a situation with more buyers.

Example 3 Consider a market of four buyers and two sellers where the requests are uniformly and independently distributed over $\{1, 2, 3\}$. In this case, the probability of a seller facing a $\hat{\gamma}$ deal is $Pr(\hat{\gamma}) = 34/81$. A seller is tempted in situations where he is first approached by one buyer and the other seller is originally approached by three buyers, two of which eventually will be rejected. In addition the value of its request is lower than at least one of the values of the requests of the buyers that were rejected by the other seller. Another situation in which the seller may be tempted is when he is originally approached by two buyers, and the tempting deal is offered by the buyer that has been rejected by the second seller. Examples of such situations include: $[\{1, 1\}\{2, 2\}]$; $[\{1\}\{2, 1, 2\}]$; $[\{3, 3\}\{2, 1\}]$; etc.

Since the sellers of this market are tempted more often than these of the previous example, buyers with the same utility function (i.e., $U_{reg} = -1$ and $U_{exit} = -23$) return to unreliable sellers even for markets that last for a smaller number of time periods, i.e., when $K \leq 17$. Similarly, they will return even if δ is smaller, i.e., $\delta < 0.98$.

Notice that theorem 1 holds for markets where $Pr(\hat{\gamma})$ (the probability of a seller being offered a larger deal at regret time) follows the next inequalities for the finite and the unrestricted cases, respectively:

- **Finite Horizon:** $Pr(\hat{\gamma}) > \frac{U_{exit} + (K-2) - U_{reg}}{U_{reg} * (K-2)}$
- **Unrestricted Horizon:** $Pr(\hat{\gamma}) > \frac{U_{exit} * (1-\delta) + U_{reg} + 1}{U_{reg}}$

Assuming that $U_{exit} < U_{reg} < 0$ and $K > 1$ (K is the number of trading steps in the market), in the finite horizon case we find that the denominator of $Pr(\hat{\gamma})$ is always negative, and the numerator is positive leading to $Pr(\hat{\gamma})$ being as small as we want ($Pr(\hat{\gamma})$ should be positive). Therefore the result presented in theorem 1 holds for markets where $\hat{\gamma}$ may not happen so often, even rarely. Hence, our results reveal the asymmetry that exists in marketplaces where sellers may behave unreliably and the buyers cannot punish them by not returning to them for a whole range of markets which may include rare large deals or even very frequent large deals that will cause the sellers to breach their commitment.

For the unrestricted horizon case, solving $Pr(\hat{\gamma})$ results in the condition that $Pr(\hat{\gamma}) > \frac{U_{exit} * (1-\delta) + U_{reg} + 1}{U_{reg}}$. Assuming again that $U_{exit} < U_{reg} < 0$, and the fact that $0 < \delta < 1$, we find that the denominator is always negative, and the numerator is positive when $U_{exit} * (1-\delta) + U_{reg} > -1$. Since U_{exit} is smaller than U_{reg} we can write $U_{exit} = U_{reg} - x$ for some positive x . Then, the numerator will be positive when $(U_{reg} - x) * (1-\delta) + U_{reg} > -1$, leading to $U_{reg} > \frac{-1-x(\delta-1)}{(2-\delta)}$.

3 The Power of Unreliable Sellers

Based on Theorem 1, a perfect equilibrium can be identified for the markets characterized in that theorem. This finding is disappointing since sellers do not keep their promises when following the strategies of the perfect equilibrium.

Theorem 2 (Perfect Equilibrium in a Market with Unreliable Sellers) *If the market satisfies the following conditions:*

- **Finite Horizon:** $U_{reg} * (1 + Pr(\hat{\gamma})(K - 2)) > U_{exit} + (K - 2)$
- **Unrestricted Horizon:** $\frac{1}{1-\delta} * [U_{reg} * (Pr(\hat{\gamma}) - 1) - 1] > U_{exit}$

then the following strategies are in perfect equilibrium:

- **The Buyers:** *The buyers remain not restricted from approaching the sellers in the market even if the sellers have behaved unreliably.*
- **The Sellers:** *The sellers accept any $\hat{\gamma}$ deal at regret time.*

When the horizon is unrestricted this profile is the unique perfect equilibrium for any $0 \leq \delta \leq 1$.

Sketch of Proof.

Following theorem 1 (The Sellers), a seller will benefit most from accepting any $\hat{\gamma}$ deal at regret time, if all the buyers return eventually to the unreliable seller. This was proven for a general market and its proof did not assume any market in particular. In addition, theorem 1(The Buyers) stated that if the sellers behave unreliably then the buyers will be induced to remain not restricted from buying from these sellers. \square

One may assume that even though the sellers do not keep their promises when following the equilibrium strategies, buyers still may not lose since an agent may play two roles in the market: sometimes he is the rejected buyer and sometimes he is the one whose requests are satisfied at regret time. However, since a buyer loses when he is rejected, and in such situations he may be rejected quite often, his average benefit is lower when the sellers do not keep their promises than in the case that they do keep their promises.

We have formally found additional Nash equilibria for a simple market with an unrestricted horizon composed of three buyers and two sellers (see theorem 3 and the corollary of theorem 5 in section 4). These profiles include cases where the sellers avoid breaking commitments. Since these profiles are Nash equilibrium but they are not in perfect equilibrium, it is unreasonable to recommend them to designers of automated agents. Such Nash equilibria enable buyers to threaten sellers with threats that the buyers themselves cannot carry out. These equilibria are known as absurd Nash equilibria[Fudenberg and Tirole, 1992]. The reason for this is that Nash equilibrium strategies are stable only at the beginning of the trade. However, agents may benefit from deviating from these strategies at intermediate stages of the trade.

4 Reliable Sellers

The results obtained so far reveal that in an electronic market with no external intervention, sellers have incentives to break commitments. Buyers are compelled to return to such unreliable sellers. Even though there are certain cases for which sellers will remain reliable, these cases comprise only Nash equilibrium (see theorem 3 and the corollary of theorem 5) which the agents themselves will not be able to follow given the perfect equilibrium shown in the former section.

To overcome such undesirable situations we propose an external mechanism that can be implemented in the electronic marketplace. This mechanism will prevent the buyers from returning to unreliable sellers. Markets' managers can store the log files containing information about the transactions performed by the buyers and sellers trading in the market. These managers can impose constraints on the buyers' behavior with respect to the sellers they are allowed to approach. Thus, a buyer may not be allowed to actually approach a seller who has behaved unreliably in a previous trade round. In this paper we are dealing with a market in which the sellers may behave unreliably.

We have analytically studied the effect of implementing such a mechanism in a simple market composed of three buyers and two sellers, when the horizon is finite (theorem 4) and when the horizon is unrestricted (theorem 5).

In particular, we study in detail the minimal market in which the temptation is relevant. This market is composed of three buyers $\{B^1, B^2, B^3\}$ and two sellers $\{S^1, S^2\}$. If there were more sellers than buyers, then, the buyer who was rejected by an unreliable seller at regret time, would have had an idle seller who could supply his β . When there is a single seller in the market, the regret time is irrelevant.

We assume that $p = 1$. As in example 1 in the simple market we analyzed, τ can be assigned the values of one, two, or three with a probability of $\frac{1}{3}$. Therefore the difference between the value of a deal $\hat{\gamma}$ offered to a seller at regret time and a deal β to which the seller is already committed can be either one or two (these $\hat{\gamma}$ deals will be denoted $\hat{\gamma}_1$ and $\hat{\gamma}_2$, respectively).

Since there are more buyers than sellers in the market, it is in the buyers' interest to be able to approach both sellers along the trading cycles. That is, even if a seller breaks a commitment to a buyer, this buyer may nevertheless decide to come back to this seller during a later trading cycle.

On the other hand, if the sellers knew that the buyers would come back to them, even if they had behaved unreliably, they would not have an incentive not to break commitments and would therefore always be tempted. Thus, the buyers would benefit from not coming back to a seller who has broken a commitment to them if this behavior could cause the sellers to keep their promises. Even though this seems to be a simple market, deciding what action will yield the maximal expected utility is not a trivial task.

4.1 The Possible States of the Simple Market

The dynamics of the simple market lead it through six different states. The tree constructed from these states and the transitions between them is depicted in Figure 2. An exclusive buyer is denoted in the figure by, “ , after the buyer's name. The buyers' names can be interchanged in the figure. The market is initially in a state in which all buyers are not restricted from approaching any seller. In all of the strategies considered in this paper, we assume that, during the first trading cycle, each buyer chooses each one of the sellers with an equal probability of $\frac{1}{2}$. Each buyer keeps choosing the seller who has also chosen him during the last cycle, and keeps choosing one seller with a probability of one half, when he has not been chosen by a seller and also has not been rejected by any seller.

The agents remain in the same state in two cases: 1) as long as the sellers do not accept

any $\hat{\gamma}$ deals proposed at regret time, or 2) as long as the buyers keep returning to a seller who accepted a $\hat{\gamma}$ deal at regret time. As a consequence, the movement of the agents between the states occurs when the sellers accept $\hat{\gamma}$ deals requested at regret time, and when the buyers do not come back to these sellers. The expected utility for the buyers and the sellers can be computed by evaluating the possible paths that these agents follow. These computations are illustrated in Appendix B.

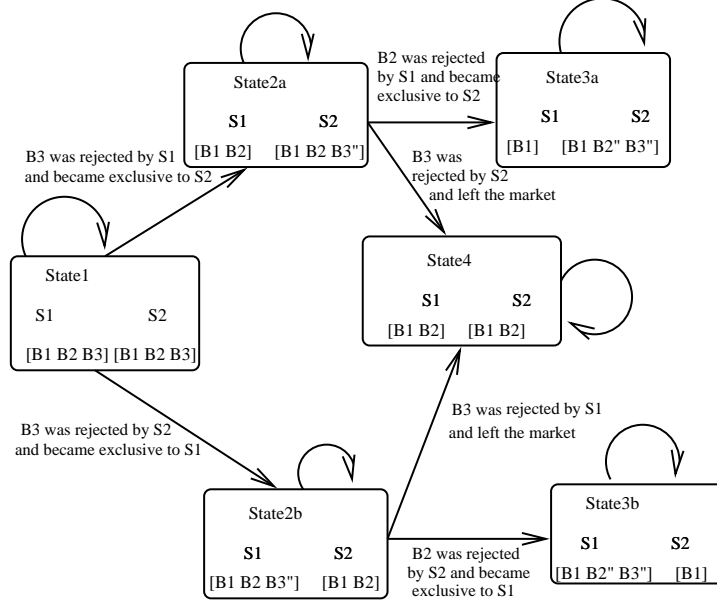


Figure 2: State Transitions in a Market with Three Buyers and Two Sellers.

Each one of the states is described in more detail hereafter:

1. **State1** - This is the initial state of the market. The three buyers are not restricted from approaching any of the sellers.
2. **State2a** - Two buyers are not restricted from approaching any specific seller; the third buyer is exclusive to seller S^2 . This state is reached when, for example, B^3 was chosen by S^1 and B^2 was preferred by S^2 over B^1 . B^1 approaches S^1 at regret time; S^1 decides to accept B^1 's deal, rejecting B^3 , making him exclusive to S^2 . Other arrangements of buyers (by changing the name of the buyers in this example) are possible and constitute other ways in which the market can reach State2a. **State2b** is symmetric to State2a where the result is that a buyer becomes exclusive to seller S^1 instead of S^2 .
3. **State3a** - One buyer is not restricted from approaching any specific seller; the other two buyers are exclusive to seller S^2 . This state is reached after the market is in State2a; for example B^2 is loyal to S^1 , and B^3 is loyal and exclusive to S^2 . B^1 lost to B^3 at t_3 and approaches S^1 at regret time. S^1 accepts B^1 's deal, thus making B^2 exclusive to S^2 . **State3b** is symmetric to State3a, where the exclusive buyers are exclusive to S^1 instead of S^2 .

4. **State4** - There are two loyal buyers (each one is loyal to a different seller), and the third buyer has opted out of the market. State4 can be reached from State2a or State2b. For example, if the market is in State2a, then one of the buyers (e.g., B^3) is exclusive to S^2 . In the next trading cycle, B^3 becomes loyal to S^2 . In the following cycle, S^2 makes a better deal at regret time, e.g., request made by B^1 . If S^2 accepts any deal at regret time that is larger than the one he is committed to, and a buyer does not return to an unreliable seller, then during this trading cycle, B^3 will be rejected by S^2 at regret time. Notice that B^3 was already exclusive to S^2 , and since there are only two sellers in the market, the buyer is compelled to leave the market.

4.2 An Equilibrium Example with Reliable Sellers

We have found a certain range of values for δ in which the sellers will refrain from breaking certain commitments they have towards the buyers. These strategies were found to be in Nash equilibrium as shown in the next theorem. Nevertheless, as explained in section 3, there is a perfect equilibrium that instructs the sellers to break their commitments given the opportunity. This will lead the buyers to return to these unreliable sellers. We bring the following theorem as an example of the possible existence of additional Nash equilibrium. However for practical reasons, these strategies will not be implemented since perfect equilibrium should be preferred.

Theorem 3 (Nash Equilibrium - Unrestricted Horizon) *When $0.92 < \delta < 1$, the following strategies are in Nash equilibrium:*

- **The Buyers:** *The buyers remain not restricted from approaching an unreliable seller that has accepted a $\hat{\gamma}_2$ deal, at regret time. The buyers will not return to an unreliable seller if the seller has accepted a $\hat{\gamma}_1$ deal, at regret time.*
- **The Sellers:** *The sellers accept a $\hat{\gamma}_2$ deal at regret time, and they do not accept a $\hat{\gamma}_1$ deal.*

Sketch of Proof. The complete computations showing that neither the sellers nor the buyers will benefit from deviating from the strategies in this theorem appear in appendix C.

□

In the next sections, we find the strategies that sellers will follow given an external mechanism which prevents buyers from returning to unreliable sellers. The strategies computed for markets with finite and unrestricted horizons maximize the sellers' expected utility. It will be shown that sellers will eventually behave reliably if the external mechanism is imposed on the buyers' selections.

4.3 The Sellers' Optimal Strategies in a Simple Market

In the previous section (section 3), we showed that if the buyers cannot be forced not to return to a seller who broke a commitment to them, then the sellers will break any commitments they have when they are offered larger deals at regret time. The following theorem states

the strategy that will maximize the seller's expected utility, if the buyers *can* be forced not to return.

Theorem 4 (Binding Threats - Finite Horizon) *If a buyer can be obliged not to come back to an unreliable seller, for example by an external mechanism, the only way for the seller to maximize his expected utility is to follow the next strategy:*

State1: *if the market is in State1:*

- *If $K' \leq 5$, the seller accepts any $\hat{\gamma}$ deal at regret time,⁴*
- *If $5 < K' \leq 12$, the seller accepts any $\hat{\gamma}_2$ deal at regret time,⁵*
- *If $K' \geq 13$, the seller rejects any $\hat{\gamma}$ deal offered to him at regret time,*

State2a: *if the market is in State2a:⁶*

- *If $K' \leq 3$, the seller accepts any $\hat{\gamma}$ deal at regret time,*
- *If $3 < K' \leq 7$, S^1 accepts any $\hat{\gamma}$ deal at regret time, and seller S^2 accepts any $\hat{\gamma}_2$ deal at regret time,*
- *If $K' \geq 8$ the seller rejects any $\hat{\gamma}$ deal offered to him at regret time.*

Sketch of Proof. We have recursively computed the expected utility functions for the sellers, assuming that the buyers will not return to them if they behave unreliably. The computation was performed in a backwards manner, starting from State4 and State3, and moving back to States 2a, 2b, and 1. For any possible $Val(\hat{\gamma}) - Val(\beta)$, we checked for which K' , the difference between the expected utility when performing β and performing $\hat{\gamma}$ is smaller than $Val(\hat{\gamma}) - Val(\beta)$. The full computations of the required utility functions appear in Appendix D.1. \square

As we see from the above theorem the buyers can indeed increase their expected benefits by utilizing an external mechanism that prevents them from returning to a seller who broke a commitment to them for at least 12 cycles. That is, the above strategies are not in Nash equilibrium. Given that sellers follow the strategies of the theorem, a buyer can increase its expected utility by returning to unreliable sellers for $K' \leq 12$. However, this will lead the sellers to break their promises also for $K' \geq 13$; this in turn will reduce the overall expected utility of the buyers.

The following theorem presents the strategy that maximizes the sellers expected utility when the buyers' threats are binding and when the horizon is unrestricted.

Theorem 5 (Binding Threats - Unrestricted Horizon) *If the buyers are forced not to return to unreliable sellers, each seller will maximize his own expected utility by using the following strategy:*

- I. *If $0 \leq \delta \leq 0.907$, each seller accepts any $\hat{\gamma}$ deal at regret time.*

⁴In this work K' denotes the number of cycles that are left until the end of the market; i.e., $K' = K - tc$.

⁵ $\hat{\gamma}_2$ is a $\hat{\gamma}$ deal offered at regret time, such that $Val(\hat{\gamma}_2) - Val(\beta) = 2$.

⁶ S^1 and S^2 have to be exchanged to obtain the sellers' strategies profile for State2b.

II. If $0.92 \leq \delta < 1$, each seller accepts any $\hat{\gamma}_2$ deal at regret time and rejects any $\hat{\gamma}_1$ deal at regret time.

III. If $\delta \rightarrow 1$, each seller rejects any $\hat{\gamma}$ deal at regret time.

For $0.907 < \delta < 0.92$, the strategy that maximizes one seller's expected utility depends on the strategy that maximizes the other seller's expected utility. In particular, if one seller accepts any $\hat{\gamma}$ deal at regret time, the other seller will benefit by accepting only $\hat{\gamma}_2$ deals and *visa versa*.

Sketch of Proof.

- **I.** $0 \leq \delta \leq 0.907$ The buyers never return to an unreliable seller in the binding case. Assuming that one seller follows the strategy stated in the profile proposed in the theorem, the other seller can deviate by rejecting any deal at regret time or by accepting only $\hat{\gamma}_2$ deals. In Appendix D.2, section D.2.1, we show that the expected utility of a seller who deviates is less than his utility if he would follow the strategies as stated in the profile presented in the theorem.
- **II.** $0.92 \leq \delta < 1$ The buyers do not return to an unreliable seller in the binding case. According to the computations in Appendix D.2, section D.2.2, no seller can benefit from deviating from the strategies specified in the profile presented in the theorem; i.e., both sellers will accept $\hat{\gamma}_2$ deals at regret time.
- **III.** $\delta \rightarrow 1$

Suppose that the buyers do not return to unreliable sellers, and that one seller (e.g., S^2) does not accept any deal at regret time, as stated in the strategies' profile of the theorem. S^1 can deviate by accepting only $\hat{\gamma}_2$ deals. In this case, eventually the market will move to State2a and remain there. In this state S^1 will obtain an average of 2.18 per cycle.⁷

If S^1 deviates by accepting any deal at regret time, the market will eventually move to State3a, in which S^1 will obtain 2, on the average, because he will be approached by a single buyer.

However, if S^1 does not deviate from the strategy stated in the theorem's profile, the market will remain in State1 (refer to Figure 3). S^1 will obtain 2, with a probability of $\frac{1}{2}$, and with the same probability it will obtain 2.44. Therefore, on average, S^1 will obtain 2.2 in each iteration. By comparing the two possible deviations in which the expected utilities are 2.18 or 2, it is clear that it is better to accept 2.2 per cycle by following the strategy of the equilibrium.

Since the buyers' threats are binding, the sellers' strategy specified in the profile is unique in the sense that this is the strategy that maximizes the sellers' expected utility.

⁷ S^1 reaches a type I node at the leftmost subtree of Figure 4 with a probability of $\frac{4}{27}$ and obtains a constant of 2.444. S^1 reaches a type I node at the rightmost subtree of Figure 4 with a probability of $\frac{1}{3}$ and obtains a constant of 2. Finally, S^1 will reach the other nodes with a probability of $1 - (\frac{4}{27} + \frac{1}{3})$ and will obtain a constant of 2.222. Therefore, at State2a, S^1 will obtain on average: $\frac{4}{27} * 2.444 + \frac{1}{3} * 2 + (1 - (\frac{4}{27} + \frac{1}{3})) * 2.222 = 2.18$.

In the range $0.907 < \delta < 0.92$, if one seller accepts any $\hat{\gamma}$ deal at regret time, the other seller will benefit by accepting only $\hat{\gamma}_2$ deals. That is, the second seller gains more by having exclusive buyers who were rejected by the first seller. \square

In the above theorem the case in which $\delta \rightarrow 1$ is special since the sellers always keep their promises. This leads to the following corollary.

Corollary 1 *If $\delta \rightarrow 1$, then the following strategies profile is a Nash equilibrium:*

- **The Buyers:** *A buyer does not return to an unreliable seller.*
- **The Sellers:** *A seller rejects any $\hat{\gamma}$ deal at regret time.*

In order to extend the above results to markets different from the simple market assumed in this section⁸, there is a need to compute the seller's expected long term loss from breaking a promise due to losing a buyer. This loss depends on the number of buyers in the market and the possible values and distribution of their needs. For a specific market, it is possible to follow the same steps as in the proof of the above theorems as presented in the appendix and identify the appropriate strategies.

We can show that the strategy that maximizes the sellers' expected utility is monotonic in the number of remaining trading cycles, K' (see lemma 1). This result can be beneficial in order to prune the space of the sellers' strategies. This space is searched when looking for the strategy that maximizes the sellers' utility given that the buyers are forced not to return to unreliable sellers.

The data structure used to describe a seller's strategy is a list of pairs: $[(1 y_1) \dots (\bar{\tau} - 1 y_{\bar{\tau}-1})]$. The first number in each pair denotes the difference that a certain $\hat{\gamma}$ deal will make to the immediate utility that the seller will obtain, if the seller accepts this $\hat{\gamma}$ deal at regret time. $\bar{\tau}$ is the maximal value that a deal β can obtain. If the number of the remaining trading cycles is less or equal to y_i (the second number in each pair), the seller will accept a $\hat{\gamma}$ deal at regret time, such that $Val(\hat{\gamma}) - Val(\beta) = i$. For example, $[(1 5) (2 12)]$ represents a seller's strategy in which the seller accepts $\hat{\gamma}_1$ deals when there are less than 5 remaining trading cycles, and he accepts $\hat{\gamma}_2$ deals when there are less than 12 remaining trading cycles.

Lemma 1 *Assume that the buyers' threat not to return to unreliable sellers is binding. If $[(1 y_1) \dots (\bar{\tau} - 1 y_{\bar{\tau}-1})]$ is the strategy that maximizes the expected utility of a seller with $|\mathcal{B}|$ buyers, $|\mathcal{S}|$ sellers and deals that range from 1 to $\bar{\tau}$, then $y_1 \leq y_2 \leq \dots \leq y_{\bar{\tau}-1}$.*

Proof. Let $\hat{\gamma}_i$ and $\hat{\gamma}_j$ be two possible $\hat{\gamma}$ deals in a general market such that $Val(\hat{\gamma}_i) - Val(\beta) = i$ and $Val(\hat{\gamma}_j) - Val(\beta) = j$. Assume that $i < j$ and, in contradiction to the lemma, that $y_i > y_j$. This means that for some period of time, such that there are less than y_i remaining trading cycles, but there are more than y_j trading cycles, the seller will accept $\hat{\gamma}_i$ deals at regret time, but will reject a larger deal such as $\hat{\gamma}_j$.

This is unreasonable, because if the strategy that maximizes the seller's expected utility allows him to accept a deal such as $\hat{\gamma}_j$ when only y_j trading cycles remain; then it is clear

⁸A simple market consists of three buyers, two sellers, and τ values that could be one, two or three with a probability of $\frac{1}{3}$.

that while there are more remaining trading cycles (like y_i) this seller cannot be affected by accepting $\hat{\gamma}_j$ deals if he can accept $\hat{\gamma}_i$ deals at the same time.

As long as there are less remaining trading cycles, the seller can lose less from accepting larger deals at regret time. This is true since as long as there are less trading cycles remaining, the gain of a seller from a tempting deal has more probability of being larger than the loss he may incur from losing one buyer. This also means that as long as there are more trading cycles remaining, the seller is more cautious about the tempting deals he may accept. Therefore as long as there are more trading cycles remaining he will risk the loss of a buyer if the tempting deal is worth it (i.e., is large enough). In our case, $\hat{\gamma}_j$ is larger than $\hat{\gamma}_i$, therefore $y_i > y_j$ would not be possible. \square

5 Related Work

Sullivan et. al. [Sullivan et al., 1999] have investigated strategies for intention reconciliation in a team framework. Self-interested but collaborative agents share a goal. Their intentions toward team-related actions may conflict with their individual intentions and as a consequence they may consider breaking a commitment to the team. A task may be given to another agent who can replace the agent whose self-directed intentions conflict with those of the team. Otherwise the task will not be done. When one agent does not keep his commitment, the whole team incurs a cost. In our case, a seller and a buyer do not comprise a team. Buyers and sellers are self-interested. They may choose their partners according to reputation measures and expected utility calculations, instead of considering themselves committed to a team.

An extensive overview of the concept of trust can be found in [Marsh, 1994]. Here, we point out the necessity to consider trust and reputation in the design of agents' behavior, in E-markets in particular. Castelfranchi et al. [Castelfranchi et al., 1998] study the control that agents may have over others in order to comply with norms by communicating reputation values. Here, we concentrate on the individual buyer and the individual seller, and how each one of them decides on his strategy of behavior when the possibility of breaking a commitment exists. In E-commerce, the direct communication between buyers is rare, and global information about sellers may be found if some external mechanism exists that objectively inspects the service given by the sellers. The direction we have chosen to study in this paper is the influence that reputation has on markets by considering its possible effects at the design stage and by describing the equilibria that exist for different markets of sellers and buyers (i.e., all buyers and all sellers are aware of the temptations the sellers may be faced with, and they are also aware of the possibility that the buyers may react to the breaking of a commitment by not returning to a specific seller).

We assume that there are only implicit contracts between the buyers and the sellers, and as such we cannot assume that these contracts are binding. The work presented in [Sandholm et al., 1999] allows agents to break a commitment by paying a predetermined penalty. In our scenario, the implicit contracts can be broken unilaterally by the sellers. There are no predetermined penalties, but the buyers can punish an unreliable seller by refraining from approaching him.

Economic studies have been done on infinite markets in which an agent whose contract has been violated will unavoidably look for another partner [Fafchamps, 1998]. Agents have types in such markets, and a screening cost is incurred in the process of finding new partners. Relational contracting that is based on trust and reputation is shown to prevent the breaking of commitments when agents have long-lives and the screening costs are high. In a market like the one studied in our work, each encounter between a seller and a buyer has its own value, so that we do not have a general payoff function for all these encounters. The role of reputation has also been analyzed in the context of bargaining [Abreu and Gul, 2000].

Reputation systems can maintain records of complaints or can provide numerical ratings of participants in E-markets (e.g., www.ebay.com and www.amazon.com). These reputation values are left for the subjective judgment of the human user, as to how to understand them and how to behave accordingly. In our framework, the sellers' reputation is taken into account at the design stage. The reputation is not considered a function, as in [Zacharia, 1999]. The agents do not have types attached to them [Fudenberg and Tirole, 1992] that lead them to act in different ways. Our interest is in the notion of reputation and how it influences the interacting agents' strategies and hence the design of equilibrium strategies. It is easy to show that there is no additional benefit to letting the agents change the reputation values during trade in light of the results we have already presented.

6 Conclusions and Future Directions

E-markets are mostly based on implicit contracts between buyers and sellers in B2C interactions. In this paper, we have studied the strategic behavior of the buyers and sellers when the sellers can be tempted by other deals after having committed themselves to satisfy a certain buyer.

The results we obtained for the set of buyers characterized in theorem 1 lead to the perfect equilibria found for markets with finite and unrestricted horizons (section 3). When the equilibrium strategies are implemented, the sellers will always break commitments to buyers when they are offered larger deals at regret time. The buyers, on the other hand, will behave loyally by returning to a seller that has supplied their request or will need to return to an unreliable sellers to avoid leaving the market. If the conditions we have assumed for the buyers' utility function do not hold, the market will not remain stable. There is no equilibrium with pure strategies for such markets. Even if part of the buyers do not return to unreliable sellers, any rational buyer will benefit most by returning.

These results show that in electronic marketplaces, where no external mechanism can regulate the buyers behavior, as in most of the B2C interactions today, the sellers have more power than the buyers. This asymmetry expresses itself in the unfair treatment which a buyer can naively receive. Specifically, a buyer who *successfully* places his order and who believes that eventually this order will be performed as expected, can indeed be deceived by not being fully serviced as expected. This change in the service expected can be due to more beneficial requests the seller is faced with at the same time, and which cannot be fulfilled given its stock at that time. These situations arise in current electronic trade, as we exemplified in the Introduction, when for example one buyer is trying to purchase a large amount of products that are on sale. The seller may receive at the same time many other

orders for the same products, and may prefer to satisfy other customers instead of one single buyer. Given the equilibrium strategies found in this paper, and assuming that the buyers are not restricted in their future purchases as to which sellers they can approach, these buyers will not receive the full order as expected because these sellers will always prefer to fulfill larger deals even after they have committed themselves to some deal (e.g., by sending a successful message by E-mail to the first customer).

We have also shown that a remedy to such asymmetry is the imposition of an external mechanism for the B2C trade. Such a regulator system can actually prevent buyers from approaching unreliable sellers. For example by tracking the log histories of the buyers, a market's manager can know when to prevent a buyer from approaching a seller that has behaved unreliably to this buyer in the past. However, threats were shown not to be effective in a market without such an external regulator. But, given that such a mechanism *actually* prevents the buyers from approaching unreliable sellers, these sellers will not breach their commitments for their own benefit. The implementation of such a regulator is our main recommendation for B2C electronic marketplaces. Then, buyers whose orders were successfully processed can be certain that these orders will be fulfilled as expected. Testing the inventories before accepting an order seems a simple solution, but then again if the sellers themselves will be responsible for that, there is no guarantee that the buyers can trust this test. A seller that accepts an order after having tested his inventory may be tempted by a larger customer later on (and before the first order was actually fulfilled). Thus, if the sellers are aware of this regulating system then it is not beneficial for them to lose a customer for a short term larger deal. This regulating system will indeed lead the sellers to satisfy the requests received in a first-in first-out order, resulting in buyers remaining loyal to the sellers, and preventing any mistrust in the market.

Another alternative for solving the problem posed in this paper is an external mechanism that will regulate the sellers behavior. For example, monetary penalties can be imposed on such sellers. The main problems with such penalties is who will collect them and how will they be used. On the other hand, since the long-term benefits of the buyers increase when there is a mechanism that regulates their behavior, it is less problematic to implement such mechanism. Therefore, we analyzed markets with no external mechanism that will prevent the sellers from breaking their promises, but we focused on how the buyers' strategies and an external regulator will cause the sellers to behave reliably.

Thus, the main conclusion of this work is that E-markets that are based on implicit contracts, encompass the potential of sellers that cannot be trusted. Nevertheless, we have shown that if electronic markets' designers impose an external regulating system the sellers will benefit from not breaking their promises.

There are many open interesting directions to explore. It is possible to study the buyers' and sellers' strategies equilibria in more complex marketplaces when some of our assumptions are relaxed. For example, the market in consideration can be dynamic, i.e., the set of buyers and sellers are not static, and the software agents that represent the buyers and the sellers may have learning capabilities.

Another assumption that can be relaxed is the assumption of trading with fixed prices. There are two cases that should be distinguished considering the buyers' in markets with negotiable prices. One consists of markets where the buyers can pay any price requested

by the sellers, but rational buyers will prefer to pay the least. The second case consists of markets where buyers may have limited resources and therefore may not be able to actually buy at any price.

In the first case, even though the buyers would prefer to trade with the sellers that ask for the lowest prices, whenever a seller breaches a commitment with a buyer, and assuming that this buyer does not return to unreliable sellers, then the buyer will have to buy from a more expensive seller. In the second case, when buyers are limited in the prices that they can pay, then the buyers that do not return to unreliable sellers will find themselves in even worse situations than in the first case. So, from the buyers' perspective it seems that similar results to the ones presented in this paper will be obtained. Buyers will return to unreliable sellers if there is not any external mechanism that will bind the buyers' threats not to return to unreliable sellers.

From the perspective of the sellers, relaxing the constraint of maintaining equal and fixed prices for the products can lead to a more interesting market where sellers adapt their prices to the demand of the market, trying to supply as many buyers as possible (in other words, trying to minimize the number of buyers being rejected). This topic is related to the price wars that evolve in competitive markets. A special market where information is traded was studied in [Kephart et al., 2000] and [Yarom et al., 2002, Yarom et al., 2003]. The dynamics of the prices in the market was explicitly studied. However, the problem of temptation and mistrust which we considered in this paper was not dealt with. An interesting future direction would be to combine these two approaches and analyze the dynamics of pricing in the framework of markets where breaching of commitments can take place.

References

- [Abreu and Gul, 2000] Abreu, D. and Gul, F. (2000). Bargaining and reputation. *Econometrica*.
- [Castelfranchi et al., 1998] Castelfranchi, C., Conte, R., and Paolucci, M. (1998). Normative reputation and the cost of compliance. *Journal of Artificial Societies and Social Simulation*, 1(3). <http://www.soc.surrey.ac.uk/JASSS/1/3/3.html>.
- [Fafchamps, 1998] Fafchamps, M. (1998). Market emergence, trust and reputation. *unpublished manuscript*. <http://www.economics.ox.ac.uk/faculty/members/fafchamps/>.
- [Fudenberg and Tirole, 1992] Fudenberg, D. and Tirole, J. (1992). *Game Theory*. MIT Press.
- [Goldman et al., 2001] Goldman, C. V., Kraus, S., and Shehory, O. (2001). Equilibria strategies for selecting sellers and satisfying buyers. *Cooperative Information Agents V*, LNAI 2182:166—177.
- [Kephart et al., 2000] Kephart, J. O., Hanson, J. E., and Greenwald, A. R. (2000). Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752.
- [Marsh, 1994] Marsh, S. P. (1994). *Formalising Trust as a Computational Concept*. PhD thesis, University of Stirling.
- [Nash, 1953] Nash, J. F. (1953). Two-person cooperative games. *Econometrica*, 21:128–140.
- [Osborne and Rubinstein, 1990] Osborne, M. J. and Rubinstein, A. (1990). *Bargaining and Markets*. Academic Press Inc., San Diego, California.
- [Osborne and Rubinstein, 1994] Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts.
- [Sandholm et al., 1999] Sandholm, T., Sikka, S., and Norden, S. (1999). Algorithms for optimizing leveled commitment contracts. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 542–547, Stockholm, Sweden.
- [Schmitz, 2002] Schmitz, S. W. (2002). The effects of e-commerce on the structure of intermediation. Technical Report 2, Austrian Academy of Sciences - ICE. To appear in C. Steinfield (ed.), *New Directions in Research on E-Commerce*.
- [Sullivan et al., 1999] Sullivan, D. G., Glass, A., Grosz, B. J., and Kraus, S. (1999). Intention reconciliation in the context of teamwork: an initial empirical investigation. In Klusch, M., Shehory, O., and Weiss, G., editors, *Cooperative Information Agents III*, Lecture Notes in Artificial Intelligence, Vol. 1652, pages 138–151. Springer-Verlag, Berlin.
- [van der Heijden et al., 2000] van der Heijden, H., Verhagen, T., and Creemers, M. (2000). Predicting online purchase behavior: replications and tests of competing models. Serie Research Memoranda 0016, Free University Amsterdam.

- [Yarom et al., 2002] Yarom, I., Goldman, C. V., and Rosenschein, J. S. (2002). The impact of infocenters on e-marketplaces. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy.
- [Yarom et al., 2003] Yarom, I., Goldman, C. V., and Rosenschein, J. S. (2003). Software middlemen benefit traders in information e-markets. *In preparation*.
- [Zacharia, 1999] Zacharia, G. (1999). Collaborative reputation mechanisms for online communities. Master's thesis, MIT.

A Sellers will be unreliable

A.1 Proof of Theorem 1 - Markets with Finite Horizon

We need to prove that if the market satisfies the following condition:

- $U_{reg} * (1 + Pr(\hat{\gamma})(K - 2)) > U_{exit} + (K - 2)$

then if all the sellers in the market are unreliable; i.e., they accept any $\hat{\gamma}$ deal at regret time, then a buyer B will benefit more from returning to an unreliable seller, than from not returning, regardless of the strategies of the other buyers.

Proof. We compare the expected utility a buyer obtains when

1. he returns to an unreliable seller, and when
2. he does not return to an unreliable seller.

We divide the K trading cycles into intervals of trading cycles. tc_1 denotes the beginning of the market. tc_{K_j} $j \geq 1$ will denote the trading cycles in which a buyer B is rejected at regret time in case (1). tc'_{K_j} $j \geq 1$ will denote the trading cycles in which a buyer B is rejected at regret time in case (2).

tc_{K_1} is the first trading cycle in which buyer B is rejected at regret time and it is equal to tc'_{K_1} . However, the other time intervals may be different. We will denote by $tc_{K_{|\mathcal{S}|}}$ the $|\mathcal{S}|^{th}$ time that B was rejected by a seller at regret time, in case (1). $tc'_{K_{|\mathcal{S}|}}$ denotes the last trading cycle when B was rejected at regret time in a market in case (2). At this trading cycle B opts out of the market.

An interval of trading cycles that occurs between any two time limits tc_i and tc_j is denoted $[tc_i, tc_j]$. The expected utility of B during such a period when it does return to an unreliable seller is denoted $EU(B, [tc_i, tc_j])$. The total expected utility of B in this case along the entire market is given by the sum of $EU(B) = EU(B, [tc_1, tc_{K_1}]) + EU(B, [tc_{K_1}, tc_{K_2}]) + \dots + EU(B, [tc_{K_{|\mathcal{S}|}}, tc_K])$.

Similarly, when B does not return to an unreliable seller, $EU'(B) = EU'(B, [tc_1, tc_{K_1}]) + EU'(B, [tc_{K_1}, tc'_{K_2}]) + \dots + EU'(B, [tc'_{K_{|\mathcal{S}|}}, tc_K])$.

Up to tc_{K_1} , B has not been rejected at any regret time, and therefore $EU(B, [tc_1, tc_{K_1}])$ and $EU'(B, [tc_1, tc_{K_1}])$ are equal.

In the worst case, when B returns to unreliable sellers, B will be rejected $|\mathcal{S}|$ number of times and will still remain in the market for another $tc_K - tc_{K_{|\mathcal{S}|}}$ trading cycles, in which, with a probability of $Pr(\hat{\gamma})$ he will obtain U_{reg} . During the intervals $[tc_{K_i}, tc_{K_j}]$ where $1 \leq i \leq |\mathcal{S}|-1$ and $2 \leq j \leq |\mathcal{S}|$, B will obtain a utility of zero. This refers to the intervals during which B is not rejected, so the worst case is that it is not chosen by any seller, leading to a utility value of zero. Therefore $EU(B) \geq EU(B, [tc_1, tc_{K_1}]) + U_{reg} * |\mathcal{S}| + Pr(\hat{\gamma}) * U_{reg} * (tc_K - tc_{K_{|\mathcal{S}|}})$.

When B does not return to unreliable sellers, he can be rejected at most $|\mathcal{S}|$ number of times. In $|\mathcal{S}|-1$ times, B will receive U_{reg} and it will attain U_{exit} in the $|\mathcal{S}|^{th}$ time, and will opt out of the market. During the intervals in which B is not rejected, in the best case, B will always have his deal performed, leading to a utility of $1 * (tc'_{K_{|\mathcal{S}|}} - tc_{K_1} - 1)$. Therefore,

the upper bound for B's expected utility is given by: $EU'(B, [tc_1, tc_{K_1}]) + U_{reg} * (|\mathcal{S}|-1) + (tc_K - tc_{|\mathcal{S}|}) * 1 + U_{exit}$.

To prove the theorem we need to find the conditions for which $EU(B)$ is larger than $EU'(B)$.

$EU(B) > EU'(B) \iff U_{reg} + U_{reg} * Pr(\hat{\gamma}) * (tc_K - tc_{K_{|\mathcal{S}|}}) > U_{exit} + (tc_K - tc_{K_{|\mathcal{S}|}})$. The markets for which the problem studied in this paper is relevant, include at least 2 sellers. Substituting $tc_{K_{|\mathcal{S}|}}$ with two, leads us to the characteristic of the utility function we assumed the buyers have.

A.2 Proof of Theorem 1 - Markets with Unrestricted Horizon

We need to prove that if the market satisfies the following condition:

- $\frac{1}{1-\delta} * [U_{reg} * (Pr(\hat{\gamma}) - 1) - 1] > U_{exit}$

then if all the sellers in the market are unreliable; i.e., they accept any $\hat{\gamma}$ deal at regret time, then a buyer B will benefit more from returning to an unreliable seller, than from not returning, regardless of the strategies of the other buyers.

Proof. Denote by case (1) the case, in which a buyer B always returns to an unreliable seller. $EU(B)$ stands for the expected utility a buyer B attains in case (1). Denote by tc_{K_1} the first time when B is rejected at regret time. In the worst case, B is going to be rejected $|\mathcal{S}|$ times, each time obtaining a utility of U_{reg} discounted by the corresponding exponent of the δ discount factor. For the worst case, we can assume that these rejections will result in a value very close to zero due to the discount rate with which they are computed. After being rejected for $|\mathcal{S}|$ times, B still remains in the market and attains an expected utility that is at least $U_{reg} * Pr(\hat{\gamma}) * (\frac{1}{1-\delta})$.⁹

Denote by case (2) the case, in which a buyer B never returns to an unreliable seller. $EU'(B)$ stands for the expected utility obtained by a buyer B in case (2). B is rejected for the first time at the same tc_{K_1} as in case (1). Therefore, there is no difference in the expected utility values attained by B in both cases, (1) and (2), in the interval $[tc_1, tc_{K_1}]$. In case (2), B is going to be rejected $|\mathcal{S}|$ times, in $|\mathcal{S}|-1$ of them B obtains an expected utility of U_{reg} weighted by a finite sum of discounted values (although this series is finite, its sum is bounded by $\frac{1}{1-\delta}$). In the best of the cases of type (2), B may have obtained a utility of one for the deals he asked for before he was lead to opt out the market and when he was not rejected. Therefore the expected utility of B in case (2) is bounded from above by $EU'(B) \leq EU'(B, [tc_1, tc_{K_1}]) + 1 * \frac{1}{1-\delta} + U_{reg} * (|\mathcal{S}|-1) * \frac{1}{1-\delta} + U_{exit}$.

The conditions for which $EU(B)$ is larger than $EU'(B)$ are such that $U_{reg} * Pr(\hat{\gamma}) * \frac{1}{1-\delta} - U_{reg} * (|\mathcal{S}|-1) * \frac{1}{1-\delta} - \frac{1}{1-\delta} > U_{exit}$. The size of the sellers' set in the market should be at least 2, for our problem to be relevant. This concludes the proof of the theorem.

If any of the other buyers do not return to an unreliable seller, then the probability of B being rejected decreases, and this can only improve his situation when he still comes back to unreliable sellers. If none of the other buyers return to unreliable sellers, then it is clear

⁹We assume that the market has unrestricted horizon, therefore the values of U_{reg} will be accumulated with a weight of the infinite series $\delta^{c_1} + \dots + \dots = \frac{1}{1-\delta}$.

that B will prefer to come back, so that he becomes the loyal, exclusive, and sole buyer of one seller who will always satisfy him. \square

B The Computation of Expected Utilities in a Simple Market

We exemplified the method of computing the expected utilities of the sellers in a simple market. This market is composed of three buyers, and two sellers. The trade is performed over deals which can take values of $\{1, 2, 3\}$ with a probability of $1/3$. In State1, the market can start from four possible initial scenarios. These scenarios, with their possible continuations, are shown in Figure 3. It is assumed that the buyers do not return to unreliable sellers.

The immediate utility that each one of the sellers will obtain at each node is specified using the format $S^i \leftarrow x$ when $i \in \{1, 2\}$ and x is the utility. The probabilities of reaching each one of the nodes appear on the edges.

There are eight possibilities for the three buyers to approach two sellers. In two cases, the three buyers approach the same seller (either S^1 or S^2). In the remaining six possibilities, two buyers approach the same seller, and the third buyer approaches the other seller. When the three buyers approach S^1 (see the leftmost subtree in Figure 3), the expected utility of S^1 is $\frac{8}{3}$,¹⁰ resulting from the maximal value obtained from all of the combinations of triplets of $\{1, 2, 3\}$. Consequently, S^2 's utility will be the maximal value he can obtain from the remaining two buyers' deals (i.e., S^2 's immediate expected utility is 2). The leftmost subtree in Figure 3 is symmetrical to the rightmost subtree in the figure. The subtrees in the middle differ in whether or not S^1 or S^2 is faced with a tempting deal $\hat{\gamma}$ at regret time. From State1-2, there are 22 cases in which there is no possible deal at regret time whose value is larger than the β deal to whom S^1 is already committed. In these 22 possibilities, the immediate expected gain of S^1 is $24/11$. Then, with an equal probability the market remains in State1-2 or in State1-3, depending on whether or not the buyer that is not restricted (and is not loyal) will choose to approach S^2 or S^1 . From State1-2, there are five cases in which a $\hat{\gamma}$ deal exists.¹¹ If S^1 accepts the tempting $\hat{\gamma}$ deal, and the buyers do not return to unreliable sellers, the market moves to State2a.

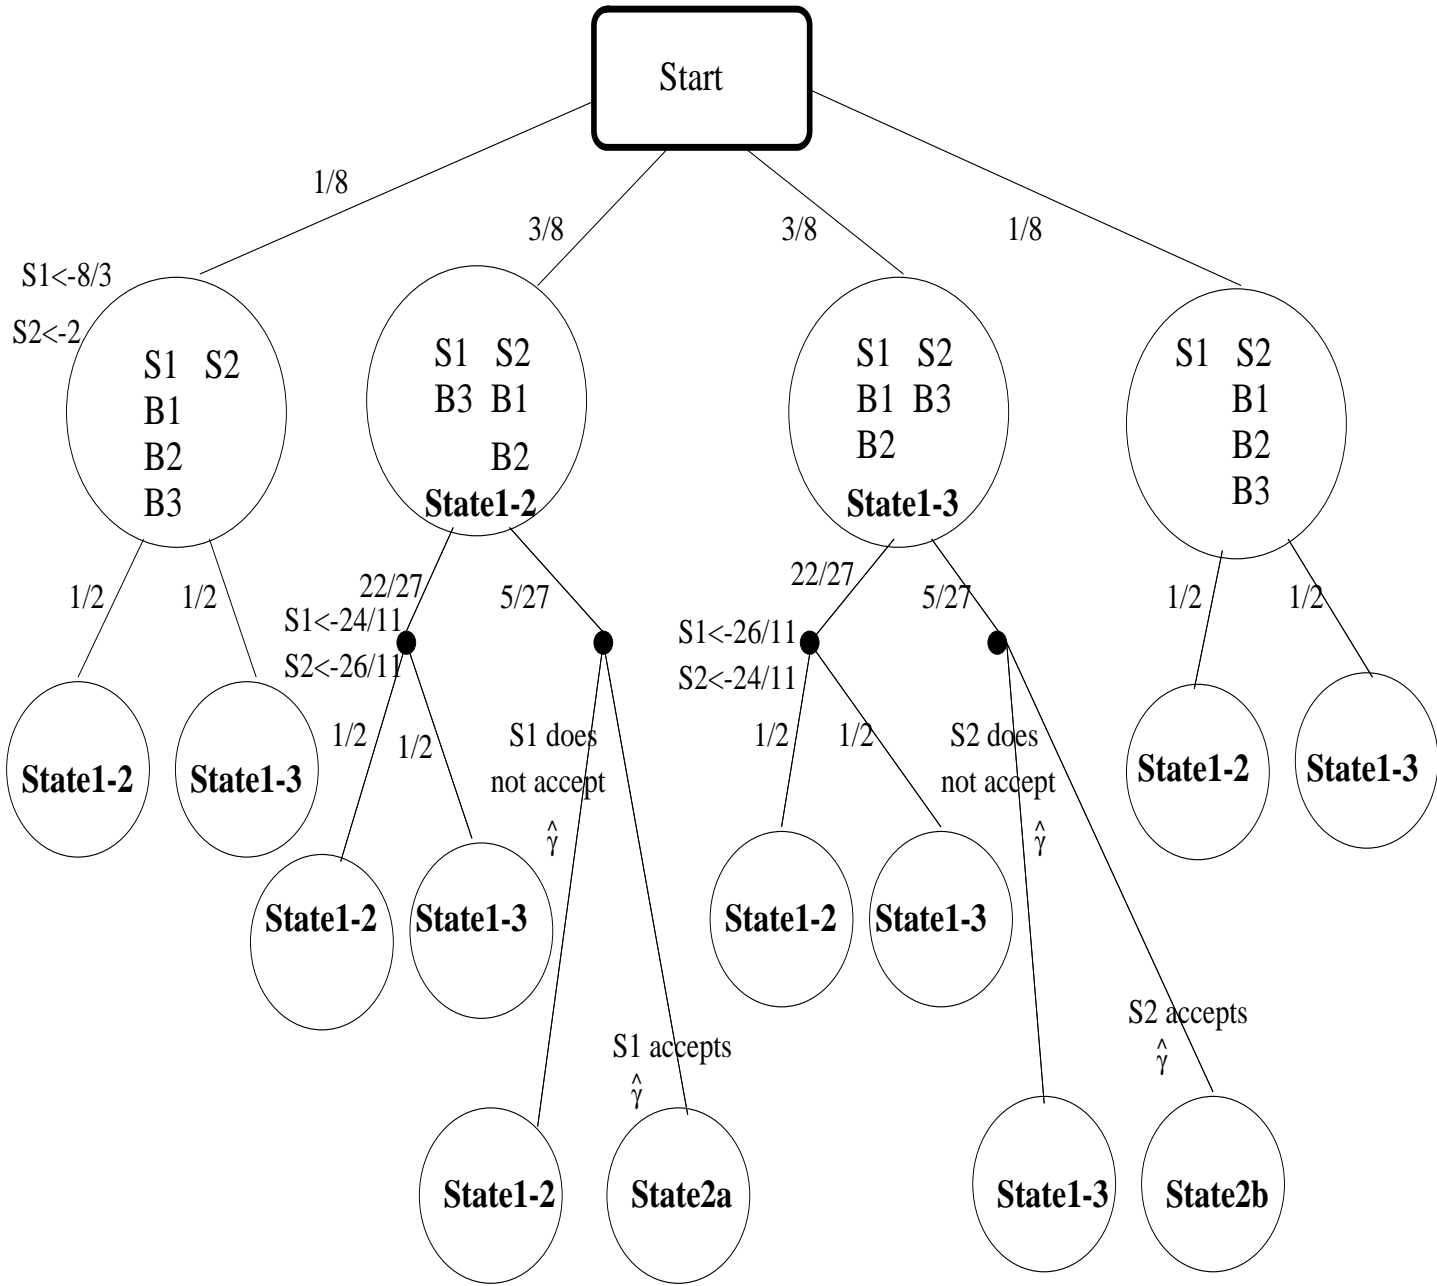
Similarly, Figure 4 describes the possible scenarios when the market is in State2a and the buyers do not return to unreliable sellers.

There are two special leaves that denote the possible transfer of the market either to State3a or to State4. Three types of nodes are distinguished in the tree corresponding to State2a:

¹⁰ $(1 * 1 + 2 * 7 + 3 * 19)/27 = 8/3$; the value 1 is accepted by the seller with a probability of $1/3$ once, when all of the buyers offered a deal with value 1, the value 2 is the maximal value in 7 cases, the value 3 is chosen by S (because it is the maximal value) in 19 cases when all of the buyers approach the same seller with $\beta = 1, 2, 3$.

¹¹These cases correspond to when: 1) S^1 is offered $\{1\}$, and S^2 is offered $\{2, 2\}$; 2) S^1 is offered $\{1\}$, and S^2 is offered $\{3, 3\}$; 3) S^1 is offered $\{1\}$, and S^2 is offered $\{2, 3\}$; 4) S^1 is offered $\{1\}$, and S^2 is offered $\{3, 2\}$; 5) S^1 is offered $\{2\}$, and S^2 is offered $\{3, 3\}$. When a seller has two offers, the first one represents the value of the deal offered by the loyal buyer.

Figure 3: State1 - The Sellers' Utility.



- **Type I Nodes** — B^2 becomes loyal¹²
- **Type II Nodes** — B^2 remains not restricted
- **Type III Nodes** — A deal is proposed at regret time; the seller can move the market to State3a or State4.¹³

Computing the expected utilities of the sellers in other markets can be done in the same way as shown here, when the appropriate trees are developed to capture the numerical values resulting from the market that is being analyzed.

C Proof of Theorem 3 - Unrestricted Horizon

We need to show that when $0.92 < \delta < 1$, the following strategies are in Nash equilibrium:

- **The Buyers:** The buyers remain not restricted from approaching an unreliable seller that has accepted a $\hat{\gamma}_2$ deal, at regret time. The buyers will not return to an unreliable seller if the seller accepted a $\hat{\gamma}_1$ deal, at regret time.
- **The Sellers:** The sellers accept a $\hat{\gamma}_2$ deal at regret time, and they do not accept a $\hat{\gamma}_1$ deal.

Sketch of Proof.

1. **The Sellers** — Assume that the buyers and one seller behave according to the strategies stated in the profile in the theorem. The other seller may deviate by rejecting all the deals requested from him at regret time, or by accepting any deal that is requested from him at regret time. Based on the following computations,¹⁴ the seller will be better off by not deviating from the strategies specified in the profile proposed in the theorem for $0.92 < \delta < 1$.

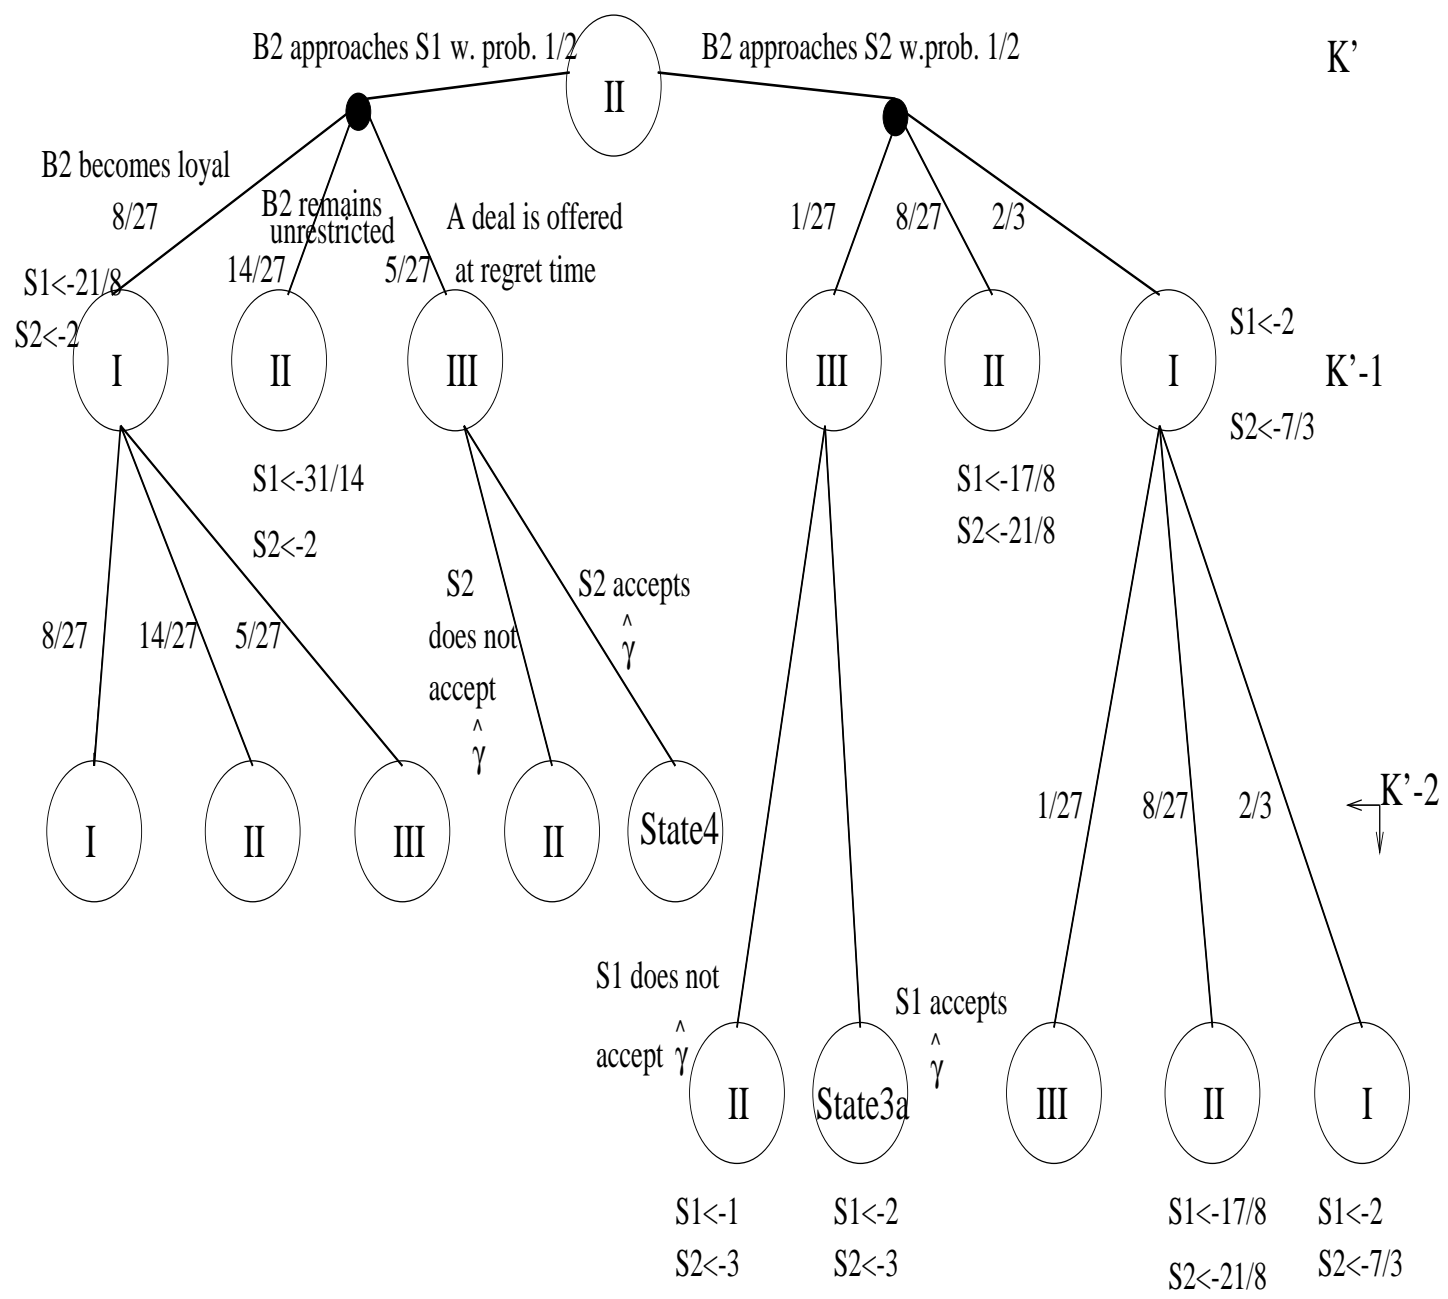
We assume that the buyers return to unreliable sellers only after these sellers have accepted a $\hat{\gamma}_2$ deal offered to them at regret time. One of the two sellers follows the strategy specified in the profile presented in the theorem for $0.92 < \delta < 1$; i.e., he will accept only $\hat{\gamma}_2$ deals at regret time and will reject any $\hat{\gamma}_1$ deal offered to him at regret time.

¹² B^2 was chosen for the presentation, although any other two buyers can be in B^2 's place.

¹³The market can move to State3a with a probability of $1/27$ when S^1 is offered $\{1\}$, and S^2 is offered $\{2, 3\}$. The market can move to State4 with a probability of $5/27$, when 1) S^1 is offered $\{2, 2\}$, and S^2 is offered $\{1\}$; 2) S^1 is offered $\{2, 3\}$, and S^2 is offered $\{1\}$; 3) S^1 is offered $\{3, 2\}$, and S^2 is offered $\{1\}$; 4) S^1 is offered $\{3, 3\}$, and S^2 is offered $\{1\}$; and 5) S^1 is offered $\{3, 2\}$, and S^2 is offered $\{1\}$. If a seller obtains two deals, the first one represents the value of the loyal buyer.

¹⁴The functions were developed based on the trees in Figures 3 and 4, with the relevant corrections based on the case we were testing. The expected utility for each seller is computed by summing up all of the possible combinations of paths through which the market could be moved with respect to a certain seller. The utility a seller gains in each one of these paths is given by multiplying the probability of reaching the situation given by the node (denoted on the edges) by the utility a seller will obtain at that node (denoted next to each node). A summary of the notations is presented in Appendix E.

Figure 4: State2a - The Sellers' Utility.



We first computed the expected utility of the second seller, when he follows the same strategy as stated in the profile of the other seller. Then we also computed the expected utility of the deviating seller.

- (a) *Both sellers accept only $\hat{\gamma}_2$ deals at regret time.* The market will remain in State1 forever, because the buyers will return to the sellers after they have behaved unreliably. The expected utility of S^1 is given by $z2_{\delta,S^1}^4 = \frac{2.074+0.407*\delta*z3_{\delta,S^1}^4}{1-0.593*\delta}$, and $z3_{\delta,S^1}^4 = \frac{2.444+0.407*\delta*z2_{\delta,S^1}^4}{1-0.593*\delta}$, resulting in

$$EU_{\delta,S^1}^4 = 2.278 + 0.5 * \delta * (z2_{\delta,S^1}^4 + z3_{\delta,S^1}^4).$$

- (b) *S^1 accepts any deal offered to him at regret time, S^2 accepts only $\hat{\gamma}_2$ deals at regret time.*

The expected utility of S^1 in State2a is composed of:

$$f_{\delta,S^1}^5 = \frac{2.241 + 0.148 * \delta * h_{\delta,S^1}^5 + 0.333 * \delta * g_{\delta,S^1}^5 + \frac{0.037}{1-\delta}}{1 - 0.5 * \delta}$$

$$g_{\delta,S^1}^5 = \frac{2.037 + 0.296 * \delta * f_{\delta,S^1}^5 + \frac{0.074}{1-\delta}}{1 - 0.667 * \delta}$$

$$h_{\delta,S^1}^5 = \frac{2.444 + 0.704 * \delta * f_{\delta,S^1}^5}{1 - 0.296 * \delta}.$$

In State1, S^1 's expected utility is built of $z2_{\delta,S^1}^5 = \frac{2.222+0.407*\delta*z3_{\delta,S^1}^5+0.148*\delta*g_{\delta,S^1}^5}{1-0.444*\delta}$, and $z3_{\delta,S^1}^5 = \frac{2.444+0.407*\delta*z2_{\delta,S^1}^5}{1-0.593*\delta}$, resulting in:

$$EU_{\delta,S^1}^5 = 2.333 + 0.444 * \delta * z2_{\delta,S^1}^5 + 0.5 * \delta * z3_{\delta,S^1}^5 + 0.056 * \delta * g_{\delta,S^1}^5.$$

Comparing both cases, we found that for $\delta > 0.92$, S^1 will benefit most by following the strategy specified in the profile; i.e., he will accept only $\hat{\gamma}_2$ deals, at regret time.

Another possibility is for one seller to deviate by rejecting any deal at regret time. We assumed that the buyers return to unreliable sellers. Based on theorem 1(Sellers), no seller can benefit from deviating by rejecting any $\hat{\gamma}$ deal at regret time.

2. **The Buyers** — Assume that both sellers and two buyers follow the strategies stated in the profile proposed in the theorem. One buyer can deviate from the strategy specified in this profile by not returning to an unreliable seller or by always returning to him. If a buyer deviates by always returning to an unreliable seller, he also will be returning to an unreliable seller when the seller accepts $\hat{\gamma}_1$ deals. This is an irrelevant case, since, as stated in the profile, no seller will accept such a deal at regret time. Therefore the deviating buyer could not profit from the deviation.

If a buyer deviates by never returning to an unreliable seller, he will decrease his utility. From theorem 1, we already know that, regardless of what the other buyers strategies are, one buyer will gain more by returning to an unreliable seller than by not returning

to him. This is true for all the buyers. That is, the deviating buyer cannot benefit from not returning to an unreliable seller.

Therefore, this buyer would be better off by abiding by the strategy stated in the profile proposed in the theorem. Consequently the strategies as described in the theorem are in equilibrium for $0.92 < \delta < 1$.

□

D The advantages of an external binding mechanism

D.1 Proof of Theorem 4 - Finite Horizon

If a buyer can be obliged not to come back to an unreliable seller, e.g., by an external mechanism, the only way for the seller to maximize his expected utility is to follow the next strategy:

State1: if the market is in State1 then

- If $K' \leq 5$, the seller accepts any deal $\hat{\gamma}$ at regret time.¹⁵
- If $5 < K' \leq 12$, the seller accepts any deal $\hat{\gamma}_2$ at regret time.
- If $K' \geq 13$ the seller rejects any $\hat{\gamma}$ deal offered to him at regret time.

State2a: If the market is in State2a then¹⁶:

- If $K' \leq 3$, the seller accepts any deal $\hat{\gamma}$ at regret time.
- If $3 < K' \leq 7$, S^1 accepts any deal $\hat{\gamma}$ at regret time, and seller S^2 accepts any deal $\hat{\gamma}_2$ at regret time.
- If $K' \geq 8$ the seller rejects any $\hat{\gamma}$ deal offered to him at regret time.

Proof.

The computations performed in this proof are based on the values explained in appendix B.

- $1 \leq K' \leq 5$:

We need to show that the expected utility of S^1 when performing any $\hat{\gamma}$ deal at regret time, is higher than when keeping the promise to perform β if the buyers' threats not to return are binding.

¹⁵ $K' = K - tc.$

¹⁶ S^1 and S^2 has to be exchanged to obtain the sellers' strategies profile for State2b.

In State1, S^1 's expected utility when he accepts any $\hat{\gamma}$ deal at regret time is given, following Figure 3, by:

$$\begin{aligned} EU_{K',S^1}^1 &= \frac{1}{8} \left(\frac{8}{3} + \frac{1}{2} z_{2_{K'-1},S^1}^1 + \frac{1}{2} z_{3_{K'-1},S^1}^1 \right) + \frac{1}{8} \left(2 + \frac{1}{2} z_{2_{K'-1},S^1}^1 + \frac{1}{2} z_{3_{K'-1},S^1}^1 \right) + \\ &+ \frac{3}{8} \left(\frac{22}{27} \left(\frac{24}{11} + \frac{1}{2} z_{2_{K'-1},S^1}^1 + \frac{1}{2} z_{3_{K'-1},S^1}^1 \right) + \frac{5}{27} \left(\frac{1}{5} * 3(2 + g_{K'-1,S^1}^1) + \frac{1}{5} * 2(3 + g_{K'-1,S^1}^1) \right) \right) + \\ &+ \frac{3}{8} \left(\frac{22}{27} \left(\frac{26}{11} + \frac{1}{2} z_{2_{K'-1},S^1}^1 + \frac{1}{2} z_{3_{K'-1},S^1}^1 \right) + \frac{5}{27} \left(\frac{1}{5} (2 + gb_{K'-1,S^1}^1 + \frac{1}{5} * 4(3 + gb_{K'-1,S^1}^1)) \right) \right). \end{aligned}$$

where $z_{2_{K'-1},S^1}^1$ and $z_{3_{K'-1},S^1}^1$ are the expected utility of S^1 at State1-2 and State1-3, respectively.

Applying algebraic manipulations, we obtain:

$$EU_{K',S^1}^1 = 2.333 + 0.431(z_{2_{K'-1},S^1}^1 + z_{3_{K'-1},S^1}^1) + 0.069(g_{K'-1,S^1}^1 + gb_{K'-1,S^1}^1).$$

Based on the notation explained in Appendix E, g_{K',S^1}^1 is the expected utility S^1 will obtain at State2a after a deal is accepted at regret time. We need to compute it in order to find the expected utility at State1. f_{K',S^1}^{1or2} is the expected utility that S^1 will attain at type II node. $h_{K'-1,S^1}^{1or2}$ is the expected utility that S^1 will attain in the state represented by type I node at the leftmost subtree (i.e., when B^2 becomes loyal to S^1); and $g_{K'-1,S^1}^{1or2}$ is the expected utility that S^1 will attain in the state represented by type I node at the rightmost subtree (i.e., when B^2 becomes loyal to S^2).

In State2a, the only $\hat{\gamma}$ deal that S^1 can be facing has a value of 2 when he has already committed himself to a β deal with value of 1. If S^1 accepts this $\hat{\gamma}_1$ deal, he will move the market to State3a, in which he will obtain a utility of 2 during the K' remaining trading cycles. Seller S^2 in State2a may face two possible $\hat{\gamma}$ deals with values of 2 or 3 when he himself is already committed to a β deal with a value of 1 or 2. If S^2 accepts any of these deals at regret time, he will move the market to State4, in which he will receive an expected utility of 2 for the K' remaining trading cycles.

The following functions were found after we checked for which K' , and for which difference $Val(\hat{\gamma}) - Val(\beta)$, it is worthwhile for a seller to accept the deal $\hat{\gamma}$ requested from him at regret time. In this case, such a seller will lose one buyer if this buyer does not return to an unreliable seller.

If $1 \leq K' \leq 3$ (both sellers accept any $\hat{\gamma}$ deal at regret time),

$$f_{K',S^1}^1 = \frac{121}{54} + \frac{4}{27} h_{K'-1,S^1}^1 + \frac{11}{27} f_{K'-1,S^1}^1 + \frac{1}{3} g_{K'-1,S^1}^1 + \frac{6}{27} * (K' - 1)$$

$$g_{K',S^1}^1 = \frac{55}{27} + \frac{2}{3} g_{K'-1,S^1}^1 + \frac{8}{27} f_{K'-1,S^1}^1 + \frac{2}{27} * (K' - 1)$$

$$h_{K',S^1}^1 = \frac{22}{9} + \frac{8}{27} h_{K'-1,S^1}^1 + \frac{14}{27} f_{K'-1,S^1}^1 + \frac{10}{27} * (K' - 1).$$

If $4 \leq K' \leq 7$ (S^1 accepts the $\hat{\gamma}_1$ deal at regret time, and S^2 accepts the $\hat{\gamma}_2$ deal at regret time),

$$f_{K',S^1}^2 = \frac{119}{54} + \frac{4}{27}h_{K'-1,S^1}^2 + \frac{26}{54}f_{K'-1,S^1}^2 + \frac{1}{3}g_{K'-1,S^1}^2 + \frac{2}{54} * (K' - 1)$$

$$g_{K',S^1}^2 = \frac{55}{27} + \frac{2}{3}g_{K'-1,S^1}^2 + \frac{8}{27}f_{K'-1,S^1}^2 + \frac{2}{27} * (K' - 1)$$

$$h_{K',S^1}^2 = \frac{22}{9} + \frac{8}{27}h_{K'-1,S^1}^2 + \frac{18}{27}f_{K'-1,S^1}^2 + \frac{2}{27} * (K' - 1).$$

We evaluated the following inequalities:

If $(1 + f_{K',S^1}^1 > 2 + 2 * K')$, then S^1 will not accept the deal at regret time and will keep the market in State2a instead of moving it to State3a.

If $(1 + f_{K',S^2}^1 > 2 + 2 * K')$, then S^2 will reject a $\hat{\gamma}_1$ deal. If $(1 + f_{K',S^2}^1 > 3 + 2 * K')$, then S^2 will reject a $\hat{\gamma}_2$ deal. If $(2 + f_{K',S^2}^1 > 3 + 2 * K')$, then S^2 will reject a $\hat{\gamma}_1$ deal. The results we obtained show that:

- if there are at most 7 remaining trading cycles, S^1 will accept the $\hat{\gamma}_1$ offered to him, and S^2 will accept any $\hat{\gamma}$ deal proposed to him at regret time.
- if there are at most 3 remaining trading cycles, S^2 will only accept the $\hat{\gamma}_1$ deals.

These results show that while $K' \leq 5$, even at State2a, S^1 will keep accepting a deal offered to him at regret time.

The corresponding inequalities for S^1 at State1¹⁷ also show that for $K' \leq 5$, S^1 will benefit most by accepting any $\hat{\gamma}$ deal at regret time. Remember that as short the market is, the sellers gain more from the difference between the $\hat{\gamma}$ deal and the β deal, than from the loss they incur by remaining with a single buyer (that will give the seller a secure gain of only 2 until the end of the market).

• $6 \leq K' \leq 12$

S^1 benefits most by accepting only those $\hat{\gamma}_2$ deals offered to him at regret time. The expected utility of S^1 while $6 \leq K' \leq 12$ is given by:

If $6 \leq K' \leq 7$,

$$EU_{K',S^1}^3 = \frac{1}{8} \left(\frac{8}{3} + \frac{1}{2}z2_{K'-1,S^1}^3 + \frac{1}{2}z3_{K'-1,S^1}^3 \right) + \frac{1}{8} \left(2 + \frac{1}{2}z2_{K'-1,S^1}^3 + \frac{1}{2}z3_{K'-1,S^1}^3 \right) + \frac{3}{8} \left(\frac{22}{27} \left(\frac{24}{11} + \frac{1}{2}z2_{K'-1,S^1}^3 + \frac{1}{2}z3_{K'-1,S^1}^3 \right) + \frac{5}{27} \left(\frac{4}{5} \left(\frac{5}{4} + z2_{K'-1,S^1}^3 \right) + \frac{1}{5} (3 + g_{K'-1,S^1}^2) \right) \right) +$$

¹⁷If $(1 + \frac{1}{2}z2_{K'-1,S^1}^1 + \frac{1}{2}z3_{K'-1,S^1}^1 < 2 + g_{K'-1,S^1}^1)$ and if $(2 + \frac{1}{2}z2_{K'-1,S^1}^1 + \frac{1}{2}z3_{K'-1,S^1}^1 < 3 + g_{K'-1,S^1}^1)$, and if $(1 + \frac{1}{2}z2_{K'-1,S^1}^1 + \frac{1}{2}z3_{K'-1,S^1}^1 < 3 + g_{K'-1,S^1}^1)$, then S^1 will accept the $\hat{\gamma}$ deal at regret time, and the market will move to State2a.

$$+ \frac{3}{8} \left(\frac{22}{27} \left(\frac{26}{11} + \frac{1}{2} z 2_{K'-1, S^1}^3 + \frac{1}{2} z 3_{K'-1, S^1}^3 \right) + \frac{5}{27} \left(\frac{4}{5} \left(\frac{11}{4} + z 3_{K'-1, S^1}^3 \right) + \frac{1}{5} (3 + g b_{K'-1, S^1}^2) \right) \right).$$

The final expression is obtained after simple algebraic manipulations:

$$EU_{K', S^1}^3 = 2.278 + 0.486(z 2_{K'-1, S^1}^3 + z 3_{K'-1, S^1}^3) + 0.014(g_{K'-1, S^1}^2 + g b_{K'-1, S^1}^2),$$

where $z 2_{K', S^1}^3$, and $z 3_{K', S^1}^3$ are S^1 's expected utility at State1-2, and State1-3, respectively. g_{K', S^1}^2 is the expected utility when the market moves to State2a due to a $\hat{\gamma}_2$ deal, and $g b_{K', S^1}^2$ is the expected utility of S^1 when the market moves to State2b. In State2a, S^1 will accept any $\hat{\gamma}_1$ deal offered to him at regret time.

If $8 \leq K' \leq 12$, S^1 will refrain from accepting any further deals at regret time at State2a. The expression for the expected utility of S^1 at State1 (when there are $8 \leq K' \leq 12$ remaining trading cycles) is given by

$$EU_{K', S^1}^3 = 2.278 + 0.486(z 2_{K'-1, S^1}^3 + z 3_{K'-1, S^1}^3) + 0.014(g_{K'-1, S^1}^4 + g b_{K'-1, S^1}^4),$$

where g_{K', S^1}^4 , and $g b_{K', S^1}^4$ are the expected utility S^1 will attain in State2a and State2b respectively.

We compared for which K' , and for which difference between possible $\hat{\gamma}$ deals and β deals, the seller will benefit most from accepting them at regret time. The only inequality that will be true here is: for $6 \leq K' \leq 7$, if S^1 accepts a $\hat{\gamma}$ deal with a value of 3, instead of a β with a value of 1, and moves the market to State2a, in which S^1 will accept any $\hat{\gamma}_1$ deal offered at regret time (in State2a, S^1 will only face $\hat{\gamma}_1$ deals at regret time). For $8 \leq K' \leq 12$, the inequality will still be true, when S^1 accepts a $\hat{\gamma}$ deal with a value of 3, instead of a β deal with a value of 1 and moves the market to State2a, in which he will refrain from accepting any deals at regret time.

- $K' \geq 13$

We show that if more than 13 trading cycles remain in the market, a seller will benefit most by rejecting any $\hat{\gamma}$ deal offered to him at regret time. We show this formally for S^1 , though it can be similarly computed for S^2 . In such a case, the market will remain in State1 (see figure 3). Following the notation in Appendix E, $z 2_{K', S^i}^5$ stands for the expected utility the seller S^i will obtain at State1-2 when there are K' remaining trading cycles. $z 3_{K', S^i}^5$ expresses the expected utility a seller S^i will obtain at State1-3, correspondingly (when the seller refrains from accepting any $\hat{\gamma}$ deal at regret time).

$$z 2_{K', S^1}^5 = \frac{22}{27} * \left(\frac{24}{11} + \frac{1}{2} * z 2_{K'-1, S^1}^5 + \frac{1}{2} * z 3_{K'-1, S^1}^5 \right) + \frac{5}{27} * \left(\frac{6}{5} + z 2_{K'-1, S^1}^5 \right)$$

$$z 3_{K', S^1}^5 = \frac{22}{27} * \left(\frac{26}{11} + \frac{1}{2} * z 2_{K'-1, S^1}^5 + \frac{1}{2} * z 3_{K'-1, S^1}^5 \right) + \frac{5}{27} * \left(\frac{14}{5} + z 3_{K'-1, S^1}^5 \right).$$

The expected utility for S^1 when the market remains in State1 is given by:

$$EU_{K', S^1}^5 = \frac{1}{8} \left(\frac{8}{3} + \frac{1}{2} z 2_{K'-1, S^1}^5 + \frac{1}{2} z 3_{K'-1, S^1}^5 \right) + \frac{1}{8} \left(2 + \frac{1}{2} z 2_{K'-1, S^1}^5 + \frac{1}{2} z 3_{K'-1, S^1}^5 \right) +$$

$$\begin{aligned}
& + \frac{3}{8} \left(\frac{22}{27} \left(\frac{24}{11} + \frac{1}{2} z_{K'-1, S^1}^5 + \frac{1}{2} z_{3^{K'-1, S^1}}^5 \right) + \frac{5}{27} \left(\frac{6}{5} + z_{2^{K'-1, S^1}}^5 \right) \right) + \\
& + \frac{3}{8} \left(\frac{22}{27} \left(\frac{26}{11} + \frac{1}{2} z_{K'-1, S^1}^5 + \frac{1}{2} z_{3^{K'-1, S^1}}^5 \right) + \frac{5}{27} \left(\frac{14}{5} + z_{3^{K'-1, S^1}}^5 \right) \right).
\end{aligned}$$

Applying algebraic operations, the expression becomes:

$$EU_{K', S^1}^5 = 2.25 + 0.5 * (z_{K'-1, S^1}^5 + z_{3^{K'-1, S^1}}^5).$$

We have compared the utility S^1 will obtain in the following cases: if S^1 accepts a deal at regret time (i.e., S^1 gains $Val(\hat{\gamma})$ and continues from State2a where S^1 will either accept any $\hat{\gamma}$ deal or will not accept any of them). The second case is when S^1 remains committed to the β deal accepted at t_3 and may continue in State1 either by rejecting any deal at regret time or by accepting any deal at regret time. The result shows that for $K' \geq 13$, it is always beneficial for S^1 to reject any $\hat{\gamma}$ deal offered to him. There are still too many trading cycles for the sellers to take the risk of losing a buyer.

D.2 Proof of Theorem 5 - Unrestricted Horizon

If the buyers are forced not to return to unreliable sellers, each seller will maximize his own expected utility by following the strategies corresponding to the profiles:

- I. If $0 \leq \delta \leq 0.907$, each seller accepts any $\hat{\gamma}$ deal at regret time.
- II. If $0.92 \leq \delta < 1$, each seller accepts any $\hat{\gamma}_2$ deal at regret time and rejects any $\hat{\gamma}_1$ deal at regret time.
- III. If $\delta \rightarrow 1$, each seller rejects any $\hat{\gamma}$ deal at regret time.

D.2.1 Case I. $0 \leq \delta < 0.907$

Buyers are forced not to come back to unreliable sellers. We denote by S^1 the seller who follows the strategy specified in the profile presented in the theorem; that is, he accepts any $\hat{\gamma}$ deal offered to him at regret time. S^2 can deviate from the strategy specified in the profile by accepting only $\hat{\gamma}_2$ deals at regret time, or by rejecting any $\hat{\gamma}$ deal at regret time.

We show later in this section D.2.3 that if the buyers do not return to unreliable sellers and one seller accepts deals at regret time, the other seller will not benefit from rejecting deals offered to him at regret time.

The other possibility is for S^2 to deviate by accepting $\hat{\gamma}_2$ deals at regret time. The expected utility of S^2 in State2a is based on the following functions:

$$\begin{aligned}
f_{\delta, S^2}^6 &= \frac{2.185 + 0.148 * \delta * h_{\delta, S^2}^6 + 0.333 * \delta * g_{\delta, S^2}^6 + \frac{0.082}{1-\delta}}{1 - 0.481 * \delta} \\
g_{\delta, S^2}^6 &= \frac{2.444 + 0.296 * \delta * f_{\delta, S^2}^6 + \frac{0.091}{1-\delta}}{1 - 0.667 * \delta} \\
h_{\delta, S^2}^6 &= \frac{2.037 + 0.667 * \delta * f_{\delta, S^2}^6 + \frac{0.074}{1-\delta}}{1 - 0.296 * \delta}.
\end{aligned}$$

The expected utility of S^2 in State2b is based on the following functions:

$$fb_{\delta,S^2}^6 = \frac{2.426 + 0.148 * \delta * hb_{\delta,S^2}^6 + 0.333 * \delta * gb_{\delta,S^2}^6 + \frac{0.074}{1-\delta}}{1 - 0.481 * \delta}$$

$$gb_{\delta,S^2}^6 = \frac{2.074 + 0.296 * \delta * fb_{\delta,S^2}^6 + \frac{0.074}{1-\delta}}{1 - 0.667 * \delta}$$

$$hb_{\delta,S^2}^6 = \frac{2.778 + 0.667 * \delta * fb_{\delta,S^2}^6 + \frac{0.074}{1-\delta}}{1 - 0.296 * \delta}.$$

S^2 's expected utility at State1 is based on $z2_{\delta,S^2}^6 = \frac{2.444+0.407*\delta*z3_{\delta,S^2}^6+0.185*\delta*g_{\delta,S^2}^6}{1-0.407*\delta}$, and $z3_{\delta,S^2}^6 = \frac{2.074+0.407*\delta*z2_{\delta,S^2}^6+0.037*\delta*gb_{\delta,S^2}^6}{1-0.556*\delta}$, leading to

$$EU_{\delta,S^2}^6 = 2.278 + 0.431 * \delta * z2_{\delta,S^2}^6 + 0.486 * \delta * z3_{\delta,S^2}^6 + 0.014 * \delta * gb_{\delta,S^2}^6 + 0.069 * \delta * g_{\delta,S^2}^6.$$

However, if both sellers abide by the strategies in the profile presented in the theorem, i.e., S^2 also accepts any $\hat{\gamma}$ deal at regret time, S^2 's expected utility will be as follows:

$$f_{\delta,S^2}^7 = \frac{2.259 + 0.148 * \delta * h_{\delta,S^2}^7 + 0.333 * \delta * g_{\delta,S^2}^7 + \frac{0.23}{1-\delta}}{1 - 0.407 * \delta}$$

$$g_{\delta,S^2}^7 = \frac{2.444 + 0.296 * \delta * f_{\delta,S^2}^7 + \frac{0.091}{1-\delta}}{1 - 0.667 * \delta}$$

$$h_{\delta,S^2}^7 = \frac{2.074 + 0.519 * \delta * f_{\delta,S^2}^7 + \frac{0.37}{1-\delta}}{1 - 0.296 * \delta}.$$

The expected utility for S^2 in State1 is based on $z2_{\delta,S^2}^7 = \frac{2.444+0.407*\delta*z3_{\delta,S^2}^7+0.185*\delta*g_{\delta,S^2}^7}{1-0.407*\delta}$, and $z3_{\delta,S^2}^7 = \frac{2.222+0.407*\delta*z2_{\delta,S^2}^7+0.185*\delta*gb_{\delta,S^2}^7}{1-0.407*\delta}$, resulting in

$$EU_{\delta,S^2}^7 = 2.333 + 0.431 * \delta * (z2_{\delta,S^2}^7 + z3_{\delta,S^2}^7) + 0.069 * \delta * (gb_{\delta,S^2}^7 + g_{\delta,S^2}^7).$$

After we computed all of these formulas, we found that for $0 \leq \delta \leq 0.907$, S^2 will benefit most by following the strategy specified in the profile stated in the theorem; i.e., he will also accept any $\hat{\gamma}$ deal offered to him at regret time.

D.2.2 Case II. $0.92 < \delta < 1$

The buyers are forced not to come back to unreliable sellers. One seller (e.g., S^2) follows the strategy specified in the profile, as stated in the theorem for $0.92 < \delta < 1$; i.e., he accepts only $\hat{\gamma}_2$ deals at regret time.

S^1 can deviate from the sellers' strategy as specified in the profile, by accepting any $\hat{\gamma}$ deal offered to him at regret time, or by rejecting any $\hat{\gamma}$ deal offered to him at regret time. We show in section D.2.3 that if the buyers do not return to unreliable sellers and one seller accepts deals at regret time, the other seller will not benefit from rejecting deals offered to

him at regret time. The proof in the appendix was presented for any value of $0 \leq \delta < 1$, and, in particular, the result is true for $0.92 < \delta < 1$.

So it remains to check which utility will be greater for S^1 , the one in which he abides by the strategy specified in the profile as described in the theorem or the one he deviates from.

If S^1 accepts any $\hat{\gamma}$ deal offered to him at regret time, his expected utility in State2a is computed as follows:

$$\begin{aligned} f_{\delta,S^1}^6 &= \frac{2.241 + 0.148 * \delta * h_{\delta,S^1}^6 + 0.333 * \delta * g_{\delta,S^1}^6 + \frac{0.074}{1-\delta}}{1 - 0.481 * \delta} \\ g_{\delta,S^1}^6 &= \frac{2.037 + 0.296 * \delta * f_{\delta,S^1}^6 + \frac{0.074}{1-\delta}}{1 - 0.667 * \delta} \\ h_{\delta,S^1}^6 &= \frac{2.444 + 0.667 * \delta * f_{\delta,S^1}^6 + \frac{0.074}{1-\delta}}{1 - 0.296 * \delta}. \end{aligned}$$

The expected utility of S^1 in State2b is based on the following:

$$\begin{aligned} fb_{\delta,S^1}^6 &= \frac{2.259 + 0.148 * \delta * hb_{\delta,S^1}^6 + 0.333 * \delta * gb_{\delta,S^1}^6 + \frac{0.185}{1-\delta}}{1 - 0.426 * \delta} \\ gb_{\delta,S^1}^6 &= \frac{2.444 + 0.333 * \delta * fb_{\delta,S^1}^6}{1 - 0.667 * \delta} \\ hb_{\delta,S^1}^6 &= \frac{2.074 + 0.519 * \delta * fb_{\delta,S^1}^6 + \frac{0.37}{1-\delta}}{1 - 0.296 * \delta}. \end{aligned}$$

In State1, S^1 's expected utility is given by $z_{\delta,S^1}^6 = \frac{2.222+0.407*\delta*z_{\delta,S^1}^6+0.185*\delta*g_{\delta,S^1}^6}{1-0.407*\delta}$, and $z_{\delta,S^1}^6 = \frac{2.444+0.407*\delta*z_{\delta,S^1}^6+0.037*\delta*gb_{\delta,S^1}^6}{1-0.556*\delta}$, comprising the final function:

$$EU_{\delta,S^1}^6 = 2.333 + 0.431 * \delta * z_{\delta,S^1}^6 + 0.486 * \delta * z_{\delta,S^1}^6 + 0.069 * \delta * g_{\delta,S^1}^6 + 0.014 * \delta * gb_{\delta,S^1}^6.$$

If both sellers follow the strategies in the profile of the theorem, they accept only $\hat{\gamma}_2$ deals at regret time, S^1 's expected utility is based on the functions that follow:

In State2a, S^1 's expected utility is composed of:

$$\begin{aligned} f_{\delta,S^1}^8 &= \frac{2.222 + 0.148 * \delta * h_{\delta,S^1}^8 + 0.333 * \delta * g_{\delta,S^1}^8 + \frac{0.037}{1-\delta}}{1 - 0.5 * \delta} \\ g_{\delta,S^1}^8 &= \frac{2 + 0.333 * \delta * f_{\delta,S^1}^8}{1 - 0.667 * \delta} \\ h_{\delta,S^1}^8 &= \frac{2.444 + 0.667 * \delta * f_{\delta,S^1}^8 + \frac{0.074}{1-\delta}}{1 - 0.296 * \delta}. \end{aligned}$$

In State1, S^1 's expected utility is composed of the following two functions: $z_{\delta,S^1}^8 = \frac{2.074+0.407*\delta*z_{\delta,S^1}^8+0.037*\delta*g_{\delta,S^1}^8}{1-0.556*\delta}$, and

$$z_{\delta,S^1}^8 = \frac{2.444+0.407*\delta*z_{\delta,S^1}^8+0.037*\delta*gb_{\delta,S^1}^8}{1-0.556*\delta},$$

$$EU_{\delta,S^1}^8 = 2.278 + 0.486 * \delta * (z_{\delta,S^1}^8 + z_{\delta,S^1}^8) + 0.014 * \delta * (g_{\delta,S^1}^8 + gb_{\delta,S^1}^8).$$

The results we obtained for $0.92 < \delta < 1$ show that S^1 will benefit most by accepting only $\hat{\gamma}_2$ deals at regret time.

D.2.3 Both Sellers will be Unreliable

In this section we show that the sellers obtain their maximal expected utility when the buyers do not return to unreliable sellers, and both sellers are unreliable. That is, all the other combinations of sellers will not lead to stable strategies at regret time.

Assume that the buyers will not return to unreliable sellers, and that one seller (e.g., S^2) rejects any $\hat{\gamma}$ deal proposed to him at regret time. We computed the expected utility S^1 will attain when 1) he also rejects any $\hat{\gamma}$ deal offered at regret time; 2) he accepts any $\hat{\gamma}$ deal offered to him at regret time; and 3) when he accepts only $\hat{\gamma}_2$ deals at regret time.

The functions were developed based on the trees in Figures 3 and 4.

1. *Both sellers reject any $\hat{\gamma}$ deal.* The market remains in State1 forever. The expected utility of S^1 is given by $z2_{\delta,S^1}^1 = \frac{2+0.407*\delta*z3_{\delta,S^1}^1}{1-0.593*\delta}$, and $z3_{\delta,S^1}^1 = \frac{2.444+0.407*\delta*z2_{\delta,S^1}^1}{1-0.593*\delta}$, giving

$$EU_{\delta,S^1}^1 = 2.25 + 0.5 * \delta * (z2_{\delta,S^1}^1 + z3_{\delta,S^1}^1).$$

2. *S^1 accepts any $\hat{\gamma}$ deal offered to him at regret time, and S^2 rejects any deal offered to him at regret time.* The corresponding utility functions for S^1 at State2a are the following:

$$f_{\delta,S^1}^2 = \frac{4.482 + 0.296 * \delta * h_{\delta,S^1}^2 + 0.667 * \delta * g_{\delta,S^1}^2 + \frac{0.074}{1-\delta}}{2 - \delta}$$

$$g_{\delta,S^1}^2 = \frac{6.111 + 0.889 * \delta * f_{\delta,S^1}^2 + \frac{0.222}{1-\delta}}{3 - 2 * \delta}$$

$$h_{\delta,S^1}^2 = \frac{2.444 + 0.704 * \delta * f_{\delta,S^1}^2}{1 - 0.296 * \delta}.$$

The expected utility for S^1 at State1 is built of $z2_{\delta,S^1}^2 = \frac{2.222+0.407*\delta*z3_{\delta,S^1}^2+0.185*\delta*g_{\delta,S^1}^2}{1-0.407*\delta}$, and $z3_{\delta,S^1}^2 = \frac{2.444+0.407*\delta*z2_{\delta,S^1}^2}{1-0.593*\delta}$, leading to

$$EU_{\delta,S^1}^2 = 2.333 + 0.431 * \delta * z2_{\delta,S^1}^2 + 0.5 * \delta * z3_{\delta,S^1}^2 + 0.069 * \delta * g_{\delta,S^1}^2.$$

3. *S^1 accepts only $\hat{\gamma}_2$ deals at regret time, while S^2 rejects any deal offered at regret time.* The expected utility of S^1 in State2a is based on the following functions:

$$f_{\delta,S^1}^3 = \frac{2.222 + 0.148 * \delta * h_{\delta,S^1}^3 + 0.333 * \delta * g_{\delta,S^1}^3}{1 - 0.519 * \delta}$$

$$g_{\delta,S^1}^3 = \frac{2 + 0.333 * \delta * f_{\delta,S^1}^3}{1 - 0.667 * \delta}$$

$$h_{\delta,S^1}^3 = \frac{2.444 + 0.704 * \delta * f_{\delta,S^1}^3}{1 - 0.296 * \delta}.$$

In State1, S^1 's expected utility is based on $z2_{\delta,S^1}^3 = \frac{2.074+0.407*\delta*z3_{\delta,S^1}^3+0.037*\delta*g_{\delta,S^1}^3}{1-0.556*\delta}$, and $z3_{\delta,S^1}^3 = \frac{2.444+0.407*\delta*z2_{\delta,S^1}^3}{1-0.593*\delta}$, resulting in

$$EU_{\delta,S^1}^3 = 2.278 + 0.486 * \delta * z2_{\delta,S^1}^3 + 0.5 * \delta * z3_{\delta,S^1}^3 + 0.014 * \delta * g_{\delta,S^1}^3.$$

S^1 gains the largest expected utility by deviating from S^2 's strategy, by accepting any deal offered to him at regret time for *any* value of δ .

E The Seller's Expected Utility Notation

The expected utility of a seller in the market is given by the following five parameters:

1. *The horizon of the market:* \mathbf{K} ' indicates the number of remaining trading cycles for which the market lasts, and δ indicates that the length of the market is uncertain.
2. *The identity of the seller:* either \mathbf{S}^1 or \mathbf{S}^2 .
3. *The state in which the market is.* This is given by the roles of the buyers relative to the sellers (i.e., which buyers are not restricted, which buyers are loyal, which buyers are in the market and have not opted out). This parameter is refined even more by the state represented by the nodes that have been identified in Figures 3 and 4 (as shown in Appendix B). The labels used for the expected utility notation that a seller obtains in **State1** is EU, and z2 and z3 are used for the two special states encountered when a seller is already in State1, i.e., **State1-2**, and **State1-3**. For State2, we use the labels, f,g,h for the special states **State2a-II**, **State2a-I-right**, **State2a-I-left**, respectively; and fb, gb, and hb for the special states **State2b-II**, **State2b-I-right**, **State2b-I-left**.
4. *The buyers' profile.* These values will differ when the buyers come back to unreliable sellers (denoted \mathbf{BR}), when the buyers do not come back to unreliable sellers (denoted $\mathbf{\bar{B}R}$), when the buyers come back to an unreliable seller only if the seller has broken his commitment to the buyer due to a $\hat{\gamma}_2$ deal (i.e., \mathbf{BR}_2) or due to a $\hat{\gamma}_1$ deal (i.e., \mathbf{BR}_1).
5. *The last two parameters correspond to the sellers' profile.* This parameter is a pair of values, in which the first one determines the $\hat{\gamma}$ deals that the corresponding seller is willing to accept at regret time if he is offered such a deal, and the second one corresponds to the second seller. The values for this parameter consist of any combination of the following: $\hat{\gamma}$, $\hat{\gamma}_1$, $\hat{\gamma}_2$, $\bar{\hat{\gamma}}$ (when the seller rejects any $\hat{\gamma}$ deal offered at regret time).

In tables 2, 3 and 4, we explain the notation used to describe each one of the cases that were relevant to some of the proofs presented in this paper.

Notation	Description
$z2_{K',S^i}^1$	$EU(K', S^i, \text{State1-2}, BR, (\hat{\gamma}, \hat{\gamma}))$
$z3_{K',S^i}^1$	$EU(K', S^i, \text{State1-3}, BR, (\hat{\gamma}, \hat{\gamma}))$
EU_{K',S^i}^1	$EU(K', S^i, \text{State1}, BR, (\hat{\gamma}, \hat{\gamma}))$
f_{K',S^i}^1	$EU(K', S^i, \text{State2a-II}, BR, (\hat{\gamma}, \hat{\gamma}))$
g_{K',S^i}^1	$EU(K', S^i, \text{State2a-I-right}, BR, (\hat{\gamma}, \hat{\gamma}))$
h_{K',S^i}^1	$EU(K', S^i, \text{State2a-I-left}, BR, (\hat{\gamma}, \hat{\gamma}))$
gb_{K',S^i}^1	$EU(K', S^i, \text{State2b-I-right}, BR, (\hat{\gamma}, \hat{\gamma}))$
f_{K',S^i}^2	$EU(K', S^i, \text{State2a-II}, BR, (\hat{\gamma}_1, \hat{\gamma}_2))$
g_{K',S^i}^2	$EU(K', S^i, \text{State2a-I-right}, BR, (\hat{\gamma}_1, \hat{\gamma}_2))$
h_{K',S^i}^2	$EU(K', S^i, \text{State2a-I-left}, BR, (\hat{\gamma}_1, \hat{\gamma}_2))$
gb_{K',S^i}^2	$EU(K', S^i, \text{State2b-I-right}, BR, (\hat{\gamma}_1, \hat{\gamma}_2))$
$z2_{K',S^i}^3$	$EU(K', S^i, \text{State1-2}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
$z3_{K',S^i}^3$	$EU(K', S^i, \text{State1-3}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
EU_{K',S^i}^3	$EU(K', S^i, \text{State1}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
f_{K',S^i}^4	$EU(K', S^i, \text{State2a-II}, BR, (\hat{\gamma}, \hat{\gamma}))$
g_{K',S^i}^4	$EU(K', S^i, \text{State2a-I-right}, BR, (\hat{\gamma}, \hat{\gamma}))$
h_{K',S^i}^4	$EU(K', S^i, \text{State2a-I-left}, BR, (\hat{\gamma}, \hat{\gamma}))$
gb_{K',S^i}^4	$EU(K', S^i, \text{State2b-I-right}, BR, (\hat{\gamma}, \hat{\gamma}))$
$z2_{K',S^i}^5$	$EU(K', S^i, \text{State1-2}, BR, (\hat{\gamma}, \hat{\gamma}))$
$z3_{K',S^i}^5$	$EU(K', S^i, \text{State1-3}, BR, (\hat{\gamma}, \hat{\gamma}))$
EU_{K',S^i}^5	$EU(K', S^i, \text{State1}, BR, (\hat{\gamma}, \hat{\gamma}))$

Table 2: A Summary of the Notations Used to Present the Expected Utility of a Seller when the Horizon is Finite.

Notation	Description
z_{δ,S^i}^1	$EU(\delta, S^i, \text{State1-2}, BR, (\hat{\gamma}, \hat{\gamma}))$
z_{δ,S^i}^3	$EU(\delta, S^i, \text{State1-3}, BR, (\hat{\gamma}, \hat{\gamma}))$
EU_{δ,S^i}^1	$EU(\delta, S^i, \text{State1}, BR, (\hat{\gamma}, \hat{\gamma}))$
f_{δ,S^i}^2	$EU(\delta, S^i, \text{State2a-II}, BR, (\hat{\gamma}, \hat{\gamma}))$
g_{δ,S^i}^2	$EU(\delta, S^i, \text{State2a-I-right}, BR, (\hat{\gamma}, \hat{\gamma}))$
h_{δ,S^i}^2	$EU(\delta, S^i, \text{State2a-I-left}, BR, (\hat{\gamma}, \hat{\gamma}))$
z_{δ,S^i}^2	$EU(\delta, S^i, \text{State1-2}, BR, (\hat{\gamma}, \hat{\gamma}))$
z_{δ,S^i}^3	$EU(\delta, S^i, \text{State1-3}, BR, (\hat{\gamma}, \hat{\gamma}))$
EU_{δ,S^i}^2	$EU(\delta, S^i, \text{State1}, BR, (\hat{\gamma}, \hat{\gamma}))$
f_{δ,S^i}^3	$EU(\delta, S^i, \text{State2a-II}, BR, (\hat{\gamma}_2, \hat{\gamma}))$
g_{δ,S^i}^3	$EU(\delta, S^i, \text{State2a-I-right}, BR, (\hat{\gamma}_2, \hat{\gamma}))$
h_{δ,S^i}^3	$EU(\delta, S^i, \text{State2a-I-left}, BR, (\hat{\gamma}_2, \hat{\gamma}))$
z_{δ,S^i}^3	$EU(\delta, S^i, \text{State1-2}, BR, (\hat{\gamma}_2, \hat{\gamma}))$
z_{δ,S^i}^3	$EU(\delta, S^i, \text{State1-3}, BR, (\hat{\gamma}_2, \hat{\gamma}))$
EU_{δ,S^i}^3	$EU(\delta, S^i, \text{State1}, BR, (\hat{\gamma}_2, \hat{\gamma}))$
z_{δ,S^i}^4	$EU(\delta, S^i, \text{State1-2}, BR_2, (\hat{\gamma}_2, \hat{\gamma}_2))$
z_{δ,S^i}^4	$EU(\delta, S^i, \text{State1-3}, BR_2, (\hat{\gamma}_2, \hat{\gamma}_2))$
EU_{δ,S^i}^4	$EU(\delta, S^i, \text{State1}, BR_2, (\hat{\gamma}_2, \hat{\gamma}_2))$
f_{δ,S^i}^5	$EU(\delta, S^i, \text{State2a-II}, BR_2, (\hat{\gamma}, \hat{\gamma}_2))$
g_{δ,S^i}^5	$EU(\delta, S^i, \text{State2a-I-right}, BR_2, (\hat{\gamma}, \hat{\gamma}_2))$
h_{δ,S^i}^5	$EU(\delta, S^i, \text{State2a-I-left}, BR_2, (\hat{\gamma}, \hat{\gamma}_2))$
z_{δ,S^i}^5	$EU(\delta, S^i, \text{State1-2}, BR_2, (\hat{\gamma}, \hat{\gamma}_2))$
z_{δ,S^i}^5	$EU(\delta, S^i, \text{State1-3}, BR_2, (\hat{\gamma}, \hat{\gamma}_2))$
EU_{δ,S^i}^5	$EU(\delta, S^i, \text{State1}, BR_2, (\hat{\gamma}, \hat{\gamma}_2))$

Table 3: A Summary of the Notations Used to Present the Expected Utility of a Seller when the Horizon is Unrestricted.

Notation	Description
f_{δ,S^i}^6	$EU(\delta, S^i, \text{State2a-II}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
g_{δ,S^i}^6	$EU(\delta, S^i, \text{State2a-I-right}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
h_{δ,S^i}^6	$EU(\delta, S^i, \text{State2a-I-left}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
fb_{δ,S^i}^6	$EU(\delta, S^i, \text{State2b-II}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
gb_{δ,S^i}^6	$EU(\delta, S^i, \text{State2b-I-right}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
hb_{δ,S^i}^6	$EU(\delta, S^i, \text{State2b-I-left}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
z_{δ,S^i}^6	$EU(\delta, S^i, \text{State1-2}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
$z3_{\delta,S^i}^6$	$EU(\delta, S^i, \text{State1-3}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
EU_{δ,S^i}^6	$EU(\delta, S^i, \text{State1}, BR, (\hat{\gamma}, \hat{\gamma}_2))$
f_{δ,S^i}^7	$EU(\delta, S^i, \text{State2a-II}, BR, (\hat{\gamma}, \hat{\gamma}))$
g_{δ,S^i}^7	$EU(\delta, S^i, \text{State2a-I-right}, BR, (\hat{\gamma}, \hat{\gamma}))$
h_{δ,S^i}^7	$EU(\delta, S^i, \text{State2a-I-left}, BR, (\hat{\gamma}, \hat{\gamma}))$
gb_{δ,S^i}^7	$EU(\delta, S^i, \text{State2b}, BR, (\hat{\gamma}, \hat{\gamma}))$
z_{δ,S^i}^7	$EU(\delta, S^i, \text{State1-2}, BR, (\hat{\gamma}, \hat{\gamma}))$
$z3_{\delta,S^i}^7$	$EU(\delta, S^i, \text{State1-3}, BR, (\hat{\gamma}, \hat{\gamma}))$
EU_{δ,S^i}^7	$EU(\delta, S^i, \text{State1}, BR, (\hat{\gamma}, \hat{\gamma}))$
f_{δ,S^i}^8	$EU(\delta, S^i, \text{State2a-II}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
g_{δ,S^i}^8	$EU(\delta, S^i, \text{State2a-I-right}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
h_{δ,S^i}^8	$EU(\delta, S^i, \text{State2a-I-left}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
gb_{δ,S^i}^8	$EU(\delta, S^i, \text{State2b}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
z_{δ,S^i}^8	$EU(\delta, S^i, \text{State1-2}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
$z3_{\delta,S^i}^8$	$EU(\delta, S^i, \text{State1-3}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$
EU_{δ,S^i}^8	$EU(\delta, S^i, \text{State1}, BR, (\hat{\gamma}_2, \hat{\gamma}_2))$

Table 4: A Summary of the Notations Used to Present the Expected Utility of a Seller when the Horizon is Unrestricted - continued.