Autonomous Agents and Multi-Agent Systems, XX, XXX–XXX, 2005

Non-Monotonic-Offers Bargaining Protocol

PINATA WINOTO piw410@mail.usask.ca
Department of Computer Science, University of Saskatchewan, Saskatoon, SK S7N 5A9, Canada

GORDON I. McCALLA mccalla@cs.usask.ca
Department of Computer Science, University of Saskatchewan, Saskatoon, SK S7N 5A9, Canada

JULITA VASSILEVA jiv@cs.usask.ca
Department of Computer Science, University of Saskatchewan, Saskatoon, SK S7N 5A9, Canada

Abstract. This paper concerns the strengths and weaknesses of non-monotonic-offers in alternating-offer bargaining protocols. It is commonly assumed that bargainers submit monotonic offers over time corresponding to their belief revisions. However, through formal analysis and simulations, we are able to show that a non-monotonic-offers protocol can generate higher average surplus and a lower breakdown rate compared to a monotonic-offers protocol.

Keywords: automated negotiation, alternating offer protocol, multiagent systems, belief revision.

1. Introduction

Negotiation among artificial agents (automated negotiation) is different from that among humans, because humans have the creativity to establish a negotiation protocol, to find its weaknesses, and to use their richer communication channels to exploit or loosely follow it. For example, shilling by the auctioneer and collusion among bidders are commonly observed in auctions, and strategic delay and persuasion are commonly observed in bargaining. Up to now, artificial agents (or 'agents', for short) are still much less creative compared to humans. However, the advantages of automated negotiation are in finding resolution quickly with minimal human intervention, which reduces human workload. Those advantages have been utilized by sniping agents in online auctions, who mostly bid in the last 10 seconds before the closing time of the auctions [2]. The other advantages of agents are their less emotional and more consistent behaviors in making decisions and their more precise calculation and information processing in achieving their goals.

In this paper, we focus on the alternating-offers negotiation between a buyer and a seller agent [10]. The design of bargaining mechanisms has been well explored both in multi-agent systems and economics [5, 8, 9]. One of the common properties in alternating-offers bargaining is the (weak) monotonic (counter-) offers by bargainers, i.e. buyers/sellers may only insist on their previous offers or raise/reduce their offers monotonically until an agreement is reached. Existing bargaining tactics, for example time dependent tactics [3, 4], behavior dependent tactics [3], and market-driven strategies [11], all follow the monotonic-offers property. In [14], we have shown that the monotonic-offers property in the bargaining between two rational and self-interested agents is the consequence of the belief revision mechanism of those agents. Intuitively, if a buyer asks for $150 to buy an item, but the seller insists on $200, then the buyer may reduce its belief that the seller is willing to sell for $150, which forces the buyer to revise its offer to a higher price. However, this intuition is only correct if the buyer's valuation (reservation price or the highest price that it is willing to pay) is more than $150 and will not change until the end of the bargaining, and/or the buyer believes that the seller's valuation will not change either. If the buyer's valuation decreases over time, for example becomes $100, then the buyer must reduce its offer from $150 to a price lower than $100, if allowed by the protocol. Conversely, if the seller's valuation changes to $0 and the item must be sold at any price on the next round, then the buyer's best strategy is to offer the minimum price, say $10. Note here, a buyer that reduces its offers may not be considered as a serious or sincere buyer; thus, this is usually not acceptable in human bargaining. However, in this paper we will show that imposing a non-monotonic-offers protocol (N-protocol, for short) in bargaining may be better than imposing a monotonic-offers protocol (M-protocol). Let us first look at an illustrative example in Figure 1.

Scenario 1. Suppose that a buyer wants to buy a service, and it needs this at a specified time. However, if it cannot get the service during that period, its utility from the service decreases over time, and becomes zero if it gets the service after a time deadline. Thus, the buyer's valuation of the service will be high at the start, and will decrease until the deadline (as shown as a downward line in Figure 1). Suppose that the seller's valuation is constant over time (the horizontal line in Figure 1). Since both parties' valuations are private, the initial spread (the difference between initial offers by buyer and seller) will be relatively big. The spread decreases as the bargaining progresses (by sending offers and counter-offers, illustrated by dotted lines in Figure 1). However, under the M-protocol, the bargaining may be stuck even if both parties repeat the bargaining several times, as illustrated in Figure 1(a). The failures are caused by the buyer who cannot resume the negotiation after its offer approaches its valuation, because its next offer will be higher than its next valuation. However, under the N-protocol, the failure can be remedied as illustrated in Figure 1(b).

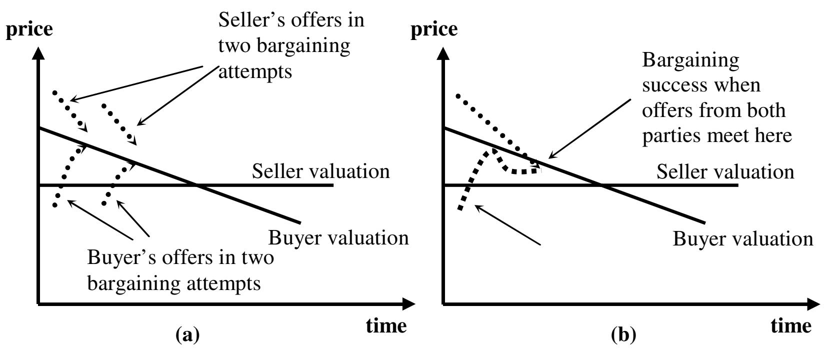

Figure 1. An example of the bargaining under (a) M-protocol and (b) N-protocol.

The illustration above only shows a partial scenario, which favors the N-protocol. The goal of the N-protocol is to allow agents to bargain as flexibly as possible so that they can always maximize their utility or increase the success rate, or both, depending on their goals. Indeed, some bargainers' decision processes under the N-protocol may suppress the convergence of spread in reaching a concession, which will be explained later. However, assuming that both bargainers are rational, slower convergence rates due to non-monotonic offers will not significantly reduce the overall performance of the bargaining in terms of agents' expected surplus.

2. Formal analysis

In this section, we will describe the bargaining model and prove the advantage of the N-protocol when all agents are bounded rational, self-interested and utility maximizers. The case of decision-theoretic or irrational agents will be analyzed empirically using simulation in Section 3.

2.1. Basic model

Suppose there are only two agents $i \in \{b, s\}$ bargaining over a single attribute item (e.g. item price), where b is the buyer agent and s is the seller agent, and they use an alternating-offers protocol. Without loss of generality, let each time period consist of a pair of decisions by the buyer and the seller and is labeled by $t = 1, 2, 3$, etc., where the buyer moves first. Also, suppose both parties do not know any information about their opponents (e.g. time deadline, valuation, or bargaining strategy), and the buyer shows a decreasing valuation over time, until a time deadline, T_{d}. Then, we define an alternative region X as all possible points in the bargaining space of the attribute (e.g. price from \$1 until \$10,000). Both agents have their private valuation $\{B, S\} \in X$ and preference order $\succsim_i$ over X (e.g. the seller prefers to sell for \$500 than \$499). For agent b, if an alternative $x \succ_b B$, then x generates a positive surplus, denoted by $\mathrm{Sur}_b(x) \geq 0$ or $\mathrm{Sur}_b(x)^+$; similarly for the seller we have $\mathrm{Sur}_S(x)^+$ if $x \succ_S S$. Moreover, a set $\mathrm{Acc}_i \subseteq X$ such that $\forall x \in \mathrm{Acc}_i \Rightarrow \mathrm{Sur}_i(x)^+$ is called an *acceptable set* of agent i. The intersection of buyer's and seller's acceptable set is called the *feasible set F*, which is a compact set (closed and bounded). Any alternative outside the feasible set is called a *disagreement set*, i.e. $D = X \backslash F$.

For example, if a buyer willing to buy for \$500 or less, and a seller is willing to sell for \$400 or more, then the feasible set is the set of prices between \$400 and \$500 inclusive. Moreover, all prices above \$500 and below \$400 are in the disagreement set. The existence of a feasible set does not guarantee the existence of a bargaining solution. The existence of a bargaining solution also depends on the agents' strategies, i.e. how their offers converge to a solution. Theoretical analysis usually assumes the agent's valuation is known by the agent, and mostly fixed until the deadline. However, in the real human negotiation world the valuation is not necessarily fixed or known in advance, because new information is added over time, the market situation changes during the bargaining, and the utility of the bargained item also changes.

Since agents are used to represent humans in automated negotiations, agent ability to perceive complex information is required in order to increase their effectiveness in the bargaining. For example, the agent may be able to check the market price from multiple fixed-price sellers, to assess the quality of the item and the seller's reputation, or to wait for other sellers with better deals, etc. Under uncertainty, the agent's rationality is defined as the behavior of an agent who only takes the choice yielding the highest expected gain. Since agents are bounded rational, they may consider different levels of future outcomes in their decisions. For instance, at time t the most myopic agent (lets say myopic-0) will only consider all possible outcomes at time t regardless of whether or not the negotiation proceeds to the next round $t+1$. On the other hand, myopic-1 agents will consider all possible outcomes at time t and $t+1$ (one round in the future), which may happen if the negotiation at time t did not conclude with any result.

Moreover, a myopic-K agent will take into account the situation when the negotiation does not yield any result until the K-th round. In reality, both human buyers and artificial agents are bounded rational and myopic to different degrees, never assuming that the seller will be the same type as them. Thus, they usually will not ask for a high price at the beginning of the negotiation because they always consider that they can get a better deal in the future. In our analysis, we will consider both myopic-0 and myopic-1 agents. We will not analyze myopic-K agents in this paper, but leave it for the future work. For simplicity, we will only focus on the buyer's perspective, which can be applied symmetrically for the seller; thus, we will drop the subscript b wherever it is not necessary.

2.1.1. Myopic-0 agents. As a rational agent, the buyer will never offer a price that generates negative surplus. Suppose that $\mathrm{Sur}(x_t)^+ = B_t - x_t$, where B_t is the buyer valuation at time t. Since the buyer only considers the immediate situation (myopic-0) and ignores the past and future situation, then its optimal decision depends only on the expected gain $\mathsf{E}G_t$:

$$\mathsf{E}G_t \equiv (1 - q_t)p(x_t)(B_t - x_t) + q_t B_\phi \qquad [1']$$

Here q_t is the likelihood of failure (breakdown) caused by the seller at time t which is independent of x_t, e.g. the seller has approached its deadline or has made a deal with another buyer, etc. B_ϕ is the buyer's valuation if the negotiation breaks down (we use ϕ to denote the time of breakdown), and $p(x_t)$ refers to its belief (subjective probability) that price x will be accepted by the seller at time t. If the negotiation does not break down, which may happen with probability $(1-q_t)$, then two possible states may happen. The first state is that the bargaining succeeds immediately (with probability $p(x_t)$), which gives the buyer positive surplus $(B_t - x_t)$. The other state is that the bargaining proceeds to the next period, which is ignored by the buyer since it is myopic. Let us assume that no surplus is generated from a breakdown ($B_\phi = 0$), and due to the independence of x_t from q_t, then the optimization problem to maximize $\mathsf{E}G_t$ in equation (1′) becomes

$$\mathrm{Max}_x\, \mathsf{E}G_t = \mathrm{Max}_x\, [p(x_t)(B_t - x_t)] \qquad [1]$$

Assumption 1. *A myopic-0 buyer will offer a price that yields the highest expected surplus at the present time t, i.e.,* $Max_x\ [p(x_t) \times Sur(x_t)^+]$.

2.1.2. *Myopic-1 agents.* A myopic-1 buyer will offer a price that yields the highest weighted sum of *expected* surplus at the present time t and the next round $t+1$. The buyer's problem can be stated as

$$\text{Max}_x \mathsf{E}G_t$$

where

$$\mathsf{E}G_t \equiv (1 - q_t)[p(x_t)(B_t - x_t) + \gamma(1 - p(x_t))\mathsf{E}G'_{t+1}] + q_t B_\phi \qquad [2']$$

In equation (2′) γ is the weight put for the expected gain in the next period, and $\mathsf{E}G'_{t+1}$ is the *estimation* of the expected gain made in the next period. Intuitively, as time goes by, the likelihood of breakdown increases, because both parties approach their time deadlines, or $q_{t+1} > q_t$. Assuming $p(x_t)$ does not change over time, and since $B_{t+1} < B_t$ (decreasing valuation), then the buyer may expect that $\mathsf{E}G'_{t+1} < \mathsf{E}G_t$. It is important here to differentiate between $\mathsf{E}G'_{t+1}$ and the true value $\mathsf{E}G_{t+1}$, because the latter may only be known at time $t+1$, i.e. after the buyer revises its belief. We may also assume that $\gamma < 1$. Similar to the analysis of myopic-0 buyers, if we assume that no surplus is generated from a breakdown and x_t is independent of q_t, then equation (2′) becomes

$$\text{Max}_x \mathsf{E}G_t = \text{Max}_x[p(x_t)(B_t - x_t) + \gamma(1 - p(x_t))\mathsf{E}G'_{t+1}] \qquad [2]$$

Assumption 2. *A myopic-1 buyer will offer a price that yields the highest expected gain according to equation (2).*

2.1.3. *Agents' opportunity to counter-offer.* Given the expected gain in equations (1′) and (2′), we can derive the next proposition regarding the agent's preference over an additional offer to a series of consecutive offers that ends up with a breakdown.

Proposition 1. *(Extending bargaining opportunity). A series of consecutive offers* $<x_1, x_2, \ldots, x_t, \phi>$ *is preferred to* $<x_1, x_2, \ldots, x_{t-1}, \phi>$ *for* $0 \leq x_t < B_t$.

Proof. For both myopic-0 and myopic-1 agents, the probability of bargaining success in the first period is $(1-q_1)p(x_1)$, and the probability of proceeding to the second period is $(1-q_1)(1-p(x_1))$. And, the probability of success in the second period is $(1-q_1)(1-p(x_1))(1-q_2)p(x_2)$. Hence, the probability that the bargaining will succeed in the i-th period is

$$(1 - q_i)p(x_i)\prod_{j=1}^{i-1}(1 - q_j)(1 - p(x_j))$$

Suppose the total probability that any offer in $X1 = < x_1, \ldots, x_{t-1}, \phi >$ is accepted by the seller is f_{t-1}. By adding $x_t > 0$, then the total probability that any offer in $X2 = < x_1, \ldots, x_t, \phi >$ is accepted by the seller is $f_t \geq f_{t-1}$. And since the last period is a breakdown, then both myopic-0 and myopic-1 agents will have the same expected gain over *X1*:

$$\mathsf{E}G_{X1} = \mathrm{Sur}(x_1)^+(1-q_1)p(x_1) + \sum_{i=2}^{t-1}\left(\mathrm{Sur}(x_i)^+(1-q_i)p(x_i)\prod_{j=1}^{i-1}(1-q_j)(1-p(x_j))\right),$$

and also over *X2*:

$$EG_{X2} = EG_{X1} + \mathrm{Sur}(x_t)^+(1-q_t)p(x_t)\prod_{j=1}^{t-1}(1-q_j)(1-p(x_j))$$

or, $\mathsf{E}G_{X2} \geq \mathsf{E}G_{X1}$. Since both the probability to find a concession and the expected surplus of *X2* is greater than that in *X1*, then we can conclude that *X2* is preferred to *X1*. □

Proposition 1 basically states that an additional opportunity to submit an offer is always preferred to ending up with a breakdown. In fact, this proposition applies not only for myopic-0 and myopic-1 agents, but also for myopic-*K* agents. From our previous example in Figure 1(a), a buyer's offers may be stuck on its valuation if its opponent does not accept its offer, which results in a breakdown. Thus, under the M-protocol, there is always a risk of early breakdown if $x_t > B_{Td}$, where T_d is the time deadline. Suppose there are two types of buyers, risk-seeking buyers, who may offer $x_t > B_{Td}$, and risk-averse buyers, who always offer $x_t \leq B_{Td}$. If we allow agents to repeat the bargaining after a breakdown as long as it is conducted before T_d, and there is no cost of doing it, then it can easily be shown that risk-seeking buyers outperform risk-averse buyers in terms of the success rate. However, if the repetition is restricted, e.g., due to a time delay on establishing a new bargaining session after a breakdown, which is a more realistic situation, then risk-averse buyers may outperform risk-seeking buyers in terms of the success rate. Intuitively, if the time delay needed for establishing a new bargaining session is relatively high when compared to T_d, then a risk-seeking agent (who tries to maximize expected gain) may be trapped in re-establishment of bargaining; resulting in a lower success rate. However, this situation can be avoided in the N-protocol, which is equivalent to a frictionless and repeatable bargaining protocol (without any cost or delay).

In the next two sections, further analysis that considers the sellers' acceptance criteria will be provided. Two types of acceptance criteria (evaluation function) are considered here: type-I and type-II agents.

2.2. *Type-I agent*

A type-I agent uses a typical evaluation function in accepting an offer by its opponent, i.e., accept an offer if it is at least as good as the counteroffer that will be sent by the agent in the next period. For instance, if a buyer plans to ask for \$50 in the next period, while the seller's current offer is \$40, then the buyer will accept the seller's current offer instead of asking for \$50. The evaluation criteria of type-I agents can be formulated in the following definition.

Definition 1. *A type-I agent uses the following evaluation function in making its decision:*

$$I_t = \begin{cases} \textit{Withdraw} & \textit{iff } t > T_d \text{ or } \mathrm{Sur}(x_{t+1})^{-} \text{for all allowed}^{1} \boldsymbol{x_{t+1}} \\ \textit{Accept} & \textit{iff } \mathrm{Sur}(y_t)^{+} \geq \mathrm{Sur}(x_{t+1})^{+} \text{and } t \leq T_d \\ \textit{Counter offer} & \textit{otherwise} \end{cases} \quad [3]$$

where I_t is the agent decision at time t, y_t is the offer by the agent's opponent at time t, x_{t+1} is the agent offer that will be proposed at time $t+1$, and $\mathrm{Sur}(x_{t+1})^{-}$ means negative surplus generated from x_{t+1}.

The evaluation function in (3) is also used in the automated negotiation literature [3, 4]. Recently, Sim and Wang [12] have adopted fuzzy rules to modify the acceptance criteria of type-I agents. Under their criterion, an agent will accept an offer even though it is slightly worse than its next counteroffer. Next, we will explain how the evaluation function affects the buyer's belief.

2.2.1. Agent's belief. From both equations (1) and (2), the maximization problem depends on the value of $p(x_t)$. Thus, we need to understand how to determine the value of $p(x_t)$. Since an agent can be programmed to have various belief functions $p(x_t)$, we may only restrict its characteristics according to its evaluation function. Suppose that the buyer knows that the seller uses the evaluation function in Definition 1. Then the buyer knows that the seller will accept its offer *iff* its offer is greater than or the same as the seller's next offer. However, offering a higher price will have a higher chance of exceeding the seller's next offer. Thus, $p(x_t)$ is an increasing function of x_t.

Proposition 2. *The buyer's belief that its offer x_t be accepted by a rational seller, $p(x_t)$ is an increasing function of x_t.*

Another characteristic is $p(x_t) = 1$ for $x_t \geq y_t$, where y_t is the seller's current offer; and $p(x_t) \to 0$ for $x_t < z$, where z is the buyer's estimation of the seller's valuation. Thus, the shape of $p(x_t)$ will be from the family of either step functions, linear functions, or S-shape functions, depending on the initial program, the value of y_t and z, and the updating method (belief revision). Intuitively, a buyer will reduce $p(x_t)$ if its offer, x_t has not been accepted by the seller, or increase $p(x_t)$ if the seller's counteroffer, y_t is very close to x_t.

Assumption 3. *(i) If the buyer's offer x_t is rejected, then the buyer will reduce $p(x')$ for all $x' \in$ Acc and $x' \leq x_t$, where a price smaller than x_t is reduced faster; thus yielding a steeper $p(x')$ at $x' \leq x_t$. (ii) If the seller reduces its previous offer x_{t-1}^S to $x_t^S < x_{t-1}^S$, then the buyer will increase its belief $p(x') = 1$ for $x' \geq x_t^S$. (iii) If the seller's offer is raised to $x_t^S > x_{t-1}^S$, then the buyer will decrease all belief of $x' < x_t^S$, i.e. $p(x')$ to a steeper one.*

Figure 2(a) shows an example of a buyer's belief at time t and $t+1$, i.e., after the seller drops its offer from \$200 to \$190. Suppose the buyer's valuation is \$180 and its previous offer x_t = \$160, with p(\$160) = 0.5. If the seller reduces its offer from \$200 to \$190, then the buyer may change its belief from $p(x_t)$ to $p(x_{t+1})$. The new beliefs are higher for prices above \$175, and lower for prices below \$175 (the intersection of $p(x_t)$ and $p(x_{t+1})$). Thus, the buyer now believes that the acceptance rate of offering \$160 is 0.3, which may force the buyer to raise its offer. However, if the seller insists on its current price, the buyer's belief will shift to the right, because by now, the buyer may think that the seller will not be able to concede (Figure 2(b)). More seriously, if the seller increases the price rather than conceding, then the buyer may think that the seller cannot even insist on its previous offer; thus, the buyer's belief will reduce more and becomes steeper (Figure 2(c)).

2.2.2. Analysis. From the seller's perspective, knowing that the buyer will update its belief as shown in Figure 2, the myopic-K seller has an incentive to deliberately stick on its current offer or even increase it in order to deceive the buyer to concede more (which may increase the seller's expected gain in the future). This may happen if the seller over-estimates its future expected gain $(\mathsf{E}G'_{t+1})_S$ or has a high value of γ; thus, making it willing to wait until the next round when the buyer concedes more. This strategy is known as "*strategic delay*" [1], discussion of which is beyond the scope of this paper. However, in the next section, we will show that a type-II seller also has incentive to increase its offers because it believes that doing so will increase the chance of early resolution. Since the intention of such a strategy is not to delay the negotiation, it is not strategic delay. To prevent agents from using strategic delay, we assume that all agents will not over-estimate and over-weight their future expected gain, have no intention of getting in a war-of-attrition (insisting until the last minute, or Boulwarism [3, 8]), and always choose an offer that maximizes their expected gain (equations (1) & (2)). This assumption also applies to the analysis of type-II agents in Section 2.3.

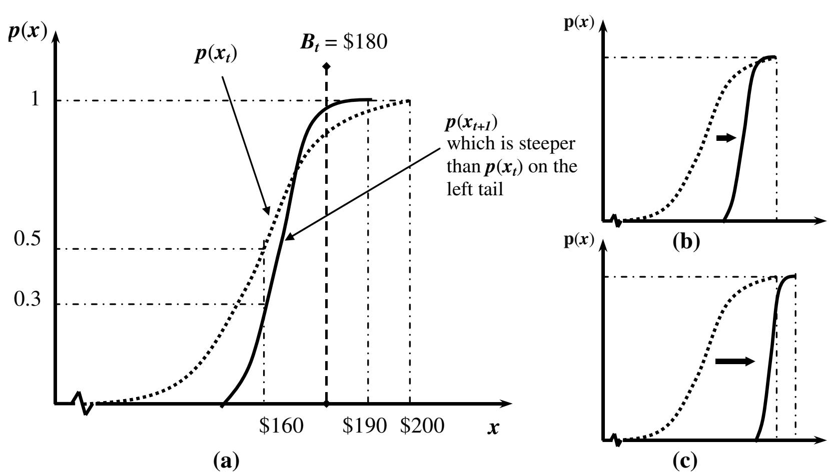

Figure 2. An example of the shifting of a buyer's belief towards seller's acceptance rate after (a) seller concedes, (b) seller insists, and (c) seller increases its price.

Given the characteristics of $p(x_t)$ in Proposition 2 and Assumption 3 above, a buyer agent is able to maximize its expected gain by offering x_t*, either in a continuous or discrete domain. In a continuous domain, we can derive x_t* by taking the first order derivative $d\mathsf{E}G_t/d\,x_t = 0$, which yields the optimal condition:

$$\text{(a) } \frac{dp(x_t^*)}{dx_t^*} = \frac{p(x_t^*)}{B_t - x_t^*} \text{ for myopic-0 buyers, and} \qquad [4a']$$

$$\text{(b) } \frac{dp(x_t^*)}{dx_t^*} = \frac{p(x_t^*)}{B_t - x_t^* - \gamma EG'_{t+1}} \text{ for myopic-1 buyers.} \qquad [4b']$$

Since most of the bargaining attributes (e.g. price) are in a discrete domain, we only focus on the optimal offer in the discrete domain (argmax $\mathsf{E}G_t$). In a discrete domain, the conditions are slightly different, i.e.

$$\text{(a) } \frac{\Delta p}{\Delta x} \approx \frac{p(x_t^a)}{B_t - x_t^*} \text{ for myopic-0 buyers, and} \qquad [4a]$$

$$\text{(b) } \frac{\Delta p}{\Delta x} \approx \frac{p(x_t^a)}{B_t - x_t^* - \gamma EG'_{t+1}} \text{ for myopic-1 buyers.} \qquad [4b]$$

Here x_t* is the optimal offer at time t, $x_t^a = x_t^* - \Delta x$ is the price just lower than x_t*, $\Delta p = p(x_t^*) - p(x_t^a)$ and Δx is the smallest interval of x (e.g. \$1). Moreover, the equality sign "≈" should be read "closest from the right to", which means that at x_t* the difference of LHS and RHS is minimized, but LHS ≥ RHS. From the optimal-offer condition above, we can prove that a buyer needs to adjust its offer in either an increasing or a decreasing way; which can only be realized under the N-protocol.

Proposition 3. *Suppose all agents under the N-protocol are type-I agents and this is common knowledge. If x_t* is an optimal offer at time t, then in order to maximize its expected gain at time $t+1$:*

(a) a myopic-0 buyer will monotonically increase its offers if $B_{t+1} - x_t^* \gg 0$, and decrease its offers if $B_{t+1} - x_t^* \to 0$ or $B_{t+1} - x_t^* < 0$;

(b) a myopic-1 buyer will monotonically increase its offers if $B_{t+1} - x_t^* \gg \gamma\mathsf{E}G'_{t+2}$, and decrease its offers if $B_{t+1} - x_t^* \to \gamma\mathsf{E}G'_{t+2}$ or $B_{t+1} - x_t^* < \gamma\mathsf{E}G'_{t+2}$.

Proof. (a) Suppose the buyer has already offered an optimal price at time t, i.e., x_t*. If this offer is not accepted by the seller, then by Assumption 3(i) the buyer will reduce both $p(x_t^*)$ and $p(x_t^a)$ at time $t+1$, where Δp becomes larger than before. Thus, from equation (4a), the LHS increases and the nominator of RHS decreases, which changes the equation such that LHS > RHS. Since the buyer's valuation decreases over time from B_t to B_{t+1}, say by ΔB, then the denominator of RHS decreases which may increase the RHS. Firstly, if ΔB is relatively small compared to $(B_t - x_t^*)$, i.e. when $(B_t - x_t^*) \gg \Delta B$ or $(B_{t+1} - x_t^*) \gg 0$, then the drop of B_t to B_{t+1}

is not enough to increase the RHS of (4a) to be equal to the LHS of (4a). Thus, the buyer needs to increase x_t* to a new optimal offer $x_{t+1}*$, which in turn increases $p(x_{t+1}{}^a)$ and decreases $(B_{t+1} - x_{t+1}*)$ (therefore, the RHS of (4a) increases more), and may also decrease Δp (therefore, the LHS of (4a) decreases) such that the equality of equation (4a) holds again. Secondly, if ΔB is relatively large compared to $(B_t - x_t*)$, i.e. when $(B_t - x_t*)$ is almost equal to ΔB or $B_{t+1} - x_t* \to 0$, then the drop of B_t to B_{t+1} increases the RHS of (4a) even more than the LHS of (4a). Consequently, the buyer needs to decrease x_t* to a new optimal offer $x_{t+1}*$, which in turn decreases $p(x_{t+1}{}^a)$ and increases $(B_{t+1} - x_{t+1}*)$ (thus, the RHS of (4a) decreases), and may increase Δp (thus, the LHS of (4a) may increase) such that the equality of equation (4a) holds again. Lastly, if $B_{t+1} - x_t* < 0$, then a negative surplus is generated at x_t*; thus, the buyer must offer a lower price.

(b) The proof is similar to the proof of (a), except with an additional term $\gamma \mathsf{E}G'_{t+1}$ that decreases over time as explained in Section 2.1.2. The decrease of this term, say by $\Delta\mathsf{E}G$, will increase the denominator of the RHS of (4b). However, if $\Delta B - \Delta\mathsf{E}G$ is relatively small compared to $(B_t - x_t* - \gamma\mathsf{E}G'_{t+1})$, i.e. when $(B_t - x_t* - \gamma\mathsf{E}G'_{t+1}) \gg \Delta B - \Delta\mathsf{E}G$, or $(B_{t+1} - x_t*) \gg \gamma\mathsf{E}G'_{t+2}$, then the buyer needs to increase x_t* to a new optimal offer $x_{t+1}*$ as in the proof of (a). The proofs of the case of $B_{t+1} - x_t* \to \gamma\mathsf{E}G'_{t+2}$ and $B_{t+1} - x_t* < \gamma\mathsf{E}G'_{t+2}$ are also similar to the steps used in the proof of (a). □

Proposition 3 describes the self-adjustment of the buyer's optimal offer in order to maximize its expected gain. When the surplus is large, the buyer tends to increase its offer. But when the surplus is small or negative, then it will reduce its offer. The convergence of a buyer's offers is guaranteed in Proposition 4.

Proposition 4. *Under the N-protocol, if all agents are type-I agents and this is common knowledge, then* $x* \to B$ *over time.*

Proof. First, we can prove the case of myopic-0 buyers using proposition 3(a). If $B_{t+1} - x_t* \gg 0$, then the buyer will increase its offer to $x_{t+1}* > x_t*$. Since the buyer's valuation decreases over time, then $B_{t+1} - x_{t+1}* < B_t - x_t*$, or $x*$ converges to B. If $B_{t+1} - x_t* \to 0$ or $B_{t+1} - x_t* < 0$, then the buyer will decrease its optimal offer; therefore, $p(x_{t+1}{}^a) < p(x_t{}^a)$ and Δp may increase. By alternating equation (4a), we have $B_t - x_t* \approx p(x_t{}^a)\, \Delta x/\Delta p$. Since, $p(x_t{}^a)$ decreases and Δp increases, then $B_t - x_t*$ *should* also decrease for the equality to hold; thus, the optimal $x*$ also converges to B. Similarly, to prove the case of myopic-1 buyers we use Proposition 3(b) and an alteration of equation (4b), $B_t - x_t* \approx p(x_t{}^b)\, \Delta x/\Delta p + \gamma\mathsf{E}G'_{t+1}$. When the buyer increases its offer, its optimal offer also converges to its valuation. When the buyer decreases its offer, by the fact that $p(x_t{}^b)$ decreases, Δp increases, and $\gamma\mathsf{E}G'_{t+1}$ decreases over time, then $B_t - x_t*$ also decreases over time; or, the optimal $x*$ also converges to B. □

From Proposition 3 and 4, we can derive several properties of type-I buyers as follows.

– The likelihood of breakdown does not affect myopic-0 buyers but does affect myopic-1 buyers. If the likelihood increases ($\gamma\mathsf{E}G'_{t+1}$ decreases), then a myopic-1 buyer's optimal offers will converge quickly to its valuation.

– If the buyer's valuation decreases sharply (ΔB is large), then its optimal offer may decrease too.

In our previous analysis, the agents were not concerned more about the failure or breakdown. However, in a real situation, humans may be concerned about a possible failure more than about the surplus. In such a situation, agents would be required to find concession as soon as possible, even with a zero surplus. Intuitively, agents in the N-protocol will outperform agents in the M-protocol in performing such tasks, as stated in Proposition 5 below.

Proposition 5. *If type-I agents are only concerned about the success rate, then the N-protocol is preferred over the M-protocol.*

Proof. The success rate depends on $p(x_t)$ and the probability that the seller's current offer y_t is lower than the buyer's next counteroffer, or $p(y_t \leq x_{t+1})$. Suppose that the buyer is only concerned about the success rate. Then, under the N-protocol the best strategy of the buyer is to offer its valuation, because both $p(B_t)$ and $p(y_t \leq B_{t+1})$ are the highest possible values. If the seller accepts B_t or counteroffers $y_t \leq B_{t+1}$, then the bargaining succeeds. If it does not, then the buyer will decrease its offer along with its valuation; thus, by applying Assumption 3(iii) to the seller, the seller will decrease its belief faster than that when the buyer increases its offer. Or, the seller will concede faster which increases $p(y_{t+1} \leq x_{t+2})$ at time $t+1$. Therefore, the success rate is higher when the buyer offers its valuation. But under the M-protocol the buyer cannot offer its valuation, because it is a decreasing function. Thus, the success rate is higher in the N-protocol. □

Combining Propositions 1 through 5, we can show that the N-protocol is better than the M-protocol for type-I agents.

Theorem 1. *The N-protocol is at least as good as the M-protocol for type-I agents.*

Proof (*sketch*). A protocol is preferred by agents if it can help the agents to achieve their goals better than another protocol. First, if the goal of agents is to maximize their expected gain, then it is shown by Proposition 3 that the N-protocol guarantees a way for the agents to offer a price that maximizes their expected gain; while agents in the M-protocol cannot because they are restricted to increase their offers only. This is true not only for a decreasing valuation, but also for an increasing one, because it is possible that $B_{t+1} - \gamma \mathsf{E}G'_{t+2} > B_t - \gamma \mathsf{E}G'_{t+1}$ (cf. equation (4b)), or the buyer has incentive to decrease its offer at time $t+1$ to maximize its expected gain. Second, if the goal of agents is to find a resolution as soon as possible, then it is shown by Proposition 5 that the N-protocol is preferred to the M-protocol. Finally, if the agents cannot optimize their offers but submit them according to a pre-defined strategy (a pre-defined sequence of offers), then buyers in the M-protocol may get stuck on their valuation (in case of decreasing valuation), which incurs some cost in order to repeat the bargaining. Buyers under the N-protocol will never get stuck, and by Proposition 1, this is preferred over getting stuck. While this does not apply in the case of increasing valuation, still the N-protocol incurs no extra cost. □

2.3. Type-II agent

In addition to the evaluation criterion used by a type-I agent, a type-II agent uses an additional evaluation function in accepting an offer by its opponent, i.e., accept an offer if it is *perceived* to generate an *optimal* surplus. For instance, if a buyer's current valuation and next offer are \$100 and \$80, respectively, while the seller's current offer is \$90, then the buyer will accept the offer if it predicts that the seller will not offer any better price until the end of the bargaining. The rationale is that some agents may be myopic and risk-averse towards breakdown in the future. Therefore, they may take a secure-yet-perceived-as-optimal offer even when a better one may exist in the future. While they also behave as type-I agents, this additional evaluation criterion makes type-II agents more 'vulnerable' to their opponent's strategy in the N-protocol. For instance, a seller may increase its offer because it believes that the buyer will be more likely to accept it. If the buyer thinks that the offer is acceptable and perceives it as the 'best' offer, then it will accept it, even if the seller may offer a better one in the future. By accepting it the buyer is not really maximizing its expected gain. Definition 2 below formalizes the evaluation criterion of type-II agents.

Definition 2. *A type-II agent uses the following evaluation function in making its decision*:

$$I_t = \begin{cases} \textit{Withdraw} & \textit{iff}\ t > T_d \text{ or } \mathrm{Sur}(x_{t+1})^- \text{ for all allowed } x_{t+1} \\ \textit{Accept} & \textit{iff}(\mathrm{Sur}(y_t)^+ \geq \mathrm{Sur}(x_{t+1})^+ \vee \wedge_j(\mathrm{Sur}^e(y_j) \\ & \leq (\mathrm{Sur}(y_t)^+)) \text{ and } t \leq T_d \\ \textit{Counter offer} & \textit{otherwise} \end{cases} \quad [5]$$

where $\mathrm{Sur}^e(y_j)$ *is the predicted value of* $\mathrm{Sur}(y_j)$ *and* $j = \{t+1, t+2, \ldots, T_d\}$.

2.3.1. Agent's belief. To avoid confusion, we will use $p(x_t)$ to represent the buyer's belief under a type-I evaluation function and $\pi(x_t)$ to represent its belief under a type-II evaluation function. Suppose at time t the buyer has a monotonically increasing belief function $\pi(x_t)$ as shown in Figure 3(a) and offers a price x_t*. Suppose x_t* is rejected and the bargaining goes to the next period $t+1$. Knowing that the seller uses the evaluation function in Definition 2, a buyer realizes that the seller may accept the same or a lower price if the price falls in the feasible set and is perceived by the seller as the best offer from the buyer; thus, $\pi(x_{t+1})$ may not be an increasing function. Instead, $\pi(x_{t+1})$ is the sum of the revised $\pi(x_t)$ according to Assumption 3, namely $p(x_{t+1})$ (dashed curve in Figure 3(b)), and a belief $k(x_{t+1})$ that represents the likelihood of x_{t+1} being accepted by the seller because it is perceived as the best offer from the buyer (top curve in Fig 3(b)). Or, $\pi(x_{t+1}) = min(p(x_{t+1}) + k(x_{t+1}), 1)$.

The highest value of $k(x_{t+1})$ is around x_t* for the following reasons:

- If the buyer offers a price $x_{t+1} \gg x_t$*, then the seller will learn that the buyer may concede more in the future; thus, for the seller x_{t+1} is not the best offer from the

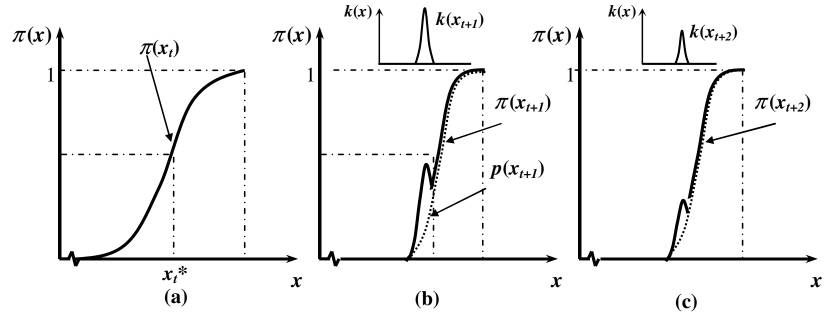

Figure 3. Belief revision of Type-II agents (a) before rejection, (b) after first rejection, (c) after second rejection.

buyer, or $k(x_{t+1}) = 0$. This is one of the reasons why human bargainers rarely concede quickly.

– If the buyer offers a price x_{t+1} slightly higher or lower than x_t* (say a price within a range $[x_t^*-\delta,\ x_t^*+\delta]$), then the seller will learn that the buyer has reached its limit of conceding; thus, x_{t+1} may be perceived as the best offer from the buyer.
– If the buyer offers a price $x_{t+1} \ll x_t^*$, then the seller realizes that the buyer may not offer a higher price in the future, or x_{t+1} is the best offer, but x_{t+1} may not be in the feasible set that can be accepted by the seller, or $k(x_{t+1}) = 0$.

If $k(x_{t+1})$ is high enough, the buyer has incentive to offer x_{t+1}^* which is slightly lower than x_t^* (which could be considered as an abuse of the N-protocol because it misleads the seller to accept a price that favors the buyer). However, if the seller does not accept x_{t+1}^*, then at time $t+2$ the buyer will reduce $k(x_{t+2})$ (the top of Figure 3(c)), because it realizes that x_{t+1}^* is less likely within the feasible set. If a new offer x_{t+2}^* which is close to x_{t+1}^* is submitted and rejected, then $k(x_{t+2})$ will be reduced more. This process may continue until a new $k(x)$ is generated if the buyer increases its offer to a price greater than $x_t^*+\delta$. If $k(x)$ is reduced several times (say n times), eventually it will become equal to zero, which leaves $\pi(x_{t+n}) = p(x_{t+n})$; thus, $\pi(x_{t+n})$ is a monotonically increasing function again.

Assumption 4. *The buyer agent will revise $k(x)$ according to its experiences. If a sequence of its offers $< x_t, x_{t+1}, x_{t+2}, \ldots, x_{t+n} >$ that are within $[x_t^*-\delta,\ x_t^*+\delta]$ is rejected, then $k(x_t^*-\delta)$ to $k(x_t^*+\delta)$ will be updated to zero, or $\pi(x_{t+n}) = p(x_{t+n})$.*

Depending on the value of $k(x)$ and n (the speed of reducing $k(x)$), we can characterize a buyer as a greedy or benevolent. If both $k(x)$ and n are large, then the buyer is greedy because it attempts to reduce its offer for several periods waiting for the seller to accept it. If $k(x) = 0$ for any x, then the buyer is benevolent and never reduces its offer for deception purpose, but reduces it according to Proposition 3 (only if facing negative/small surplus).

Similar to type-I agents, $\pi(x_t) = 1$ for $x_t \geq y_t$, where y_t is the seller's current offer; and $\pi(x_t) \to 0$ for $x_t < z$, where z is the buyer's estimation of the seller's valuation. Since the buyer does not know whether or not the seller is type-II, $k(x)$ depends on both x and the estimated ratio of sellers who perceived it as the signal of the best

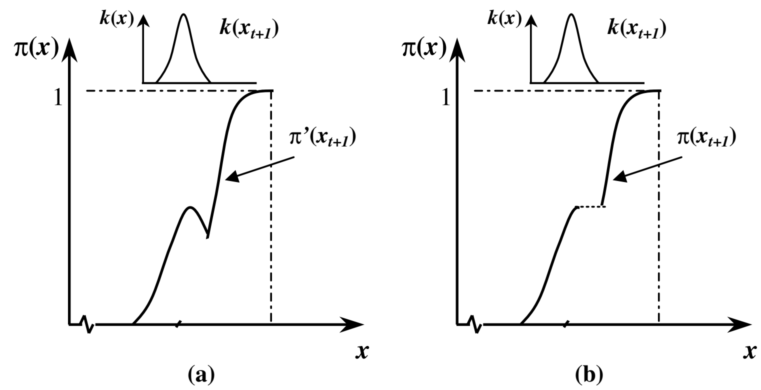

Figure 4. Belief of Type-II agents that is (a) stored and (b) actually considered.

offer from the buyer. However, the buyer does not need to consider all prices in order to find an optimal offer. Intuitively, if there are several prices with the same value π, then the lowest price will be selected because it will generate higher expected gain. Thus, the search space of the buyer is always an increasing $\pi(x_t)$, as shown in Figure 4(b).

2.3.2. Analysis. Let $p(x_t)$ be the agents' belief under a type-I evaluation function without considering $k(x_t)$. Then, we may modify equations (1) and (2) into

$$argmax_x \mathsf{E}G_t = argmax_x[\pi(x_t)(B_t - x_t)] \quad [6a]$$

$$argmax_x \mathsf{E}G_t = argmax_x[\pi(x_t)(B_t - x_t) + \gamma(1 - \pi(x_t))\mathsf{E}G'_{t+1}] \quad [6b]$$

where $\pi(x_t) = p(x_t) + k(x_t)$ if x_t within $[x_{t-1}*-\delta,\ x_{t-1}*+\delta]$, and $\pi(x_t) = p(x_t)$ otherwise.

Equations (6a) and (6b) can only be solved numerically. However, under certain conditions, the optimal offer may increase or decrease. For instance, x_t* will increase at a moment $\pi(x_t) = p(x_t)$ (by Assumption 4) and $B_{t+1} - x_t* \gg 0$ (myopic-0) or $B_{t+1} - x_t* \gg \gamma\mathsf{E}G'_{t+2}$ (myopic-1). And x_t* will decrease if $B_{t+1} - x_t* < 0$ (myopic-0) or $B_{t+1} - x_t* < \gamma\mathsf{E}G'_{t+2}$ (myopic-1). Thus, similar to the situation in type-I agents, we may conclude that type-II buyers will benefit from the N-protocol, because they can maximize their expected gain either by increasing or decreasing their offers. But, the benefit for type-II agents may not be as high as for type-I agents due to the possibility of abuses of the N-protocol by greedy agents, which reduces the convergence speed of the bargaining spread and may reduce the success rate. However, we can prove that the buyers' offers will eventually approach their valuations, and if they are only concerned about the success rate, then the N-protocol is better than the M-protocol.

Proposition 6. *Under the N-protocol, if all agents are type-II agents and this is common knowledge, then* $x_\infty* \to B_\infty$.

Proof. If the offers by a buyer decrease or increase to a value beyond the region $[x_t^*-\delta, x_t^*+\delta]$, then equations (6a) and (6b) become equations (1) and (2), respectively. Thus, we can apply the maximization analysis done for type-I agents; therefore, Proposition 4 holds and x_t^* converges to B. Consequently, a divergence appears only if the buyer increases and/or decreases the price within the region $[x_t^*-\delta, x_t^*+\delta]$. Let the buyer always offer prices within the region. However, under Assumption 4, within interval n the buyer will have $\pi(x_{t+n}) = p(x_{t+n})$. And n rounds after that, the buyer will have $\pi(x_{t+2n}) = p(x_{t+2n})$, where $p(x_{t+2n})$ is the updated version of $p(x_{t+n})$. If x_{t+2n}^* and x_{t+n}^* are the optimal offers at time $t+2n$ and $t+n$, respectively, and we ignore the situation between time $t+2n$ and $t+n$, then from proposition 4 we get $x_{t+2n}^* \rightarrow B_{t+2n}$, or x_{t+2n}^* relatively converges to B_{t+2n} compared to x_{t+n}^* from B_{t+n}. Iteratively, x_{t+3n}^* also relatively converges to B_{t+3n} compared to x_{t+2n}^* from B_{t+2n}. Therefore, $x_\infty^* \rightarrow B_\infty$. □

Proposition 7. *If type-II agents are only concerned about the success rate, then the N-protocol is preferred to the M-protocol.*

The proof of Proposition 7 is similar to the proof of Proposition 5, with an additional condition that the seller may concede faster because a decreasing of the buyer's offers along its valuation may be perceived by the seller as the 'best' offer from the buyer. And a faster concession by the seller will result in higher probability of finding a resolution; thus, the N-protocol is preferred to the M-protocol.

As for type-I agents, the N-protocol is better than the M-protocol for type-II agents if they are only concerned about the success rate (Proposition 7). Moreover, the N-protocol is better than the M-protocol for type-II agents also because it gives them flexibility to offer an optimal price. However, if neither agent is concerned about the success rate, then the success rate of the N-protocol may not be as high as the success rate of the M-protocol, especially when most agents are relatively greedy, i.e. n is relatively big compared to T_d. This becomes the serious weakness of the N-protocol.

Proposition 8. *Suppose the deadline of type-I agents is the same as the deadline of type-II agents. Then the success rate of the N-protocol for type-II agents is as low as* $(1/\tilde{n})$*-th of the success rate for type-I agents, where,* $\tilde{T_d}$ *and* $\tilde{n}$ *are the average value of* T_d *and* n *of type-II agents.*

Proof. (*sketch*). The worst case is that all type-II agents always choose an offer from the region $[x_t^*-\delta, x_t^*+\delta]$ which never yields a success. Since the effect of $k(x)$ disappears after on average $\tilde{n}$ periods, then on average, we can find $\tilde{T_d}/\tilde{n}$ periods where type-II agents are under the same condition as type-I agents. If the success rate of type-I agents under the N-protocol is m, then the success rate of type-II agents will be as low as $m/\tilde{n}$. □

Theorem 2. *The N-protocol is at least as good as the M-protocol for type-II agents if agents are benevolent.*

Proof. The same as the proof of theorem 1, except that agents are benevolent.

According to Theorem 1 and 2, we conclude that the N-protocol is better than the M-protocol in terms of success rate and expected gain for buyers if agents are benevolent. However, one of the objections to formal analysis lies in the strict assumption of agents' rationality as utility maximizers, which is not realistic in the real world. For example, what is the best belief revision procedure for the buyers in our analysis above? Can a buyer really maximize its expected gain? If the formal analysis is thus limited, then the only alternative analysis is through experimentation with artificial and/or human agents.

3. Simulation

3.1. Experimental design

First, 100 pairs of upward valuations (where both buyers and sellers have non-increasing valuations over time) are generated randomly under a pre-specified range. Then, two main parts of the experiment are designed, based on the protocol and the strategies used by agents:

– Agents use random strategies in bargaining;
– Agents use reactive (behavior-dependent) strategies in bargaining [3].

Furthermore, the first part of the experiment can be divided into four groups according to the strategies used by sellers:

– Risk-averse seller (R-averse): a seller who offers a monotonic-decreasing price but will not offer any price below the maximum valuation (in this case $100);
– Risk-seeking seller (R-seeking): a seller who offers a monotonic-decreasing price above its valuation, which may get stuck on its valuation;
– Non-monotonic-offer seller (N-seller): a seller who offers a monotonic-decreasing price if it is above its valuation unless it is stuck on its valuation (benevolent seller);
– Non-monotonic with random change (NR-seller): a seller who is similar to an N-seller, except that it may increase its price randomly (with probability equal to 0.1) in order to mislead type-II buyers (greedy seller).

Each group in the experiment is subdivided into four treatments based on the strategies used by sellers and buyers as shown in Table 1. Thus, a total of 16 groups of experiments are conducted in part 1, and each group is repeated 300 times for each pair of valuations, resulting in 480,000 trials in part 1. In this experiment, agents only follow random strategies, i.e. sellers (buyers) randomly increase (decrease) their

Table 1. Four treatments based on agents' evaluation.

		Seller	
		Type-I	Type-II
Buyer	Type-I	S1B1	S2B1
	Type-II	S1B2	S2B2

offers until the valuations are reached. Depending on the agents' characteristics, some may raise/drop their offers faster than others, but they never take their opponent's offers/behavior into consideration. The type-I agent uses the following strategy:

- if my opponent's current offer generates higher positive surplus than my offer which will be sent in the next period, then accept my opponent's current offer.

The type-II agent uses the same strategy as a type-I agent with two additional criteria:

- if my opponent's current offer generates higher positive surplus than the previous one, do nothing;
- if my opponent's current offer generates lower positive surplus than the previous one, accept it.

The experiment in part 2 is almost the same as in part 1, except that most agents consider their opponent's offers, and react accordingly. Four reactive strategies are considered:

- Tit-for-tat: the proponent's move is always the same as the opponent's previous move (I will concede if you conceded on the previous move, but I will divert if you diverted);
- Tit-for-2tat: the proponent's move is always the same as the opponent's previous two moves (I will concede/divert if you conceded/diverted on the previous two moves);
- Tat-for-tit: the proponent's move is the reverse of the opponent's previous move (I will concede/divert if you diverted/conceded on the previous move);
- Spread-driven: the proponent's move tries to reduce the spread of negotiation by a constant fraction (always concede unless it gets stuck on its valuation).

These reactive strategies are slightly different from those used in [3, 11]. Each experiment in part 2 consists of agents with the four reactive strategies above plus agents with random strategies. Eight different treatments are conducted in part 2: the 4 different treatments given in Table 1, and 2 different protocols for each of them (N-protocol and M-protocol). For the statistical analysis, each possible combination

Table 2. Parameters used in the experiment.

Parameters	Values
Maximum bargaining periods	99
Sellers' and buyers' initial valuation	\$50–\$85
Sellers' and buyers' final valuation	\$100
Increment of valuation	\$5
Range of sellers' initial offers	\$100–\$120
Range of buyers' initial offers	\$30–\$50
Min. increment/decrement of offers	\$1
Max. increment/decrement of offers	\$5

is repeated more than 30 times. In total, we have conducted 720,000 trials in part 2. Since two reactive agents may stand on their offers (e.g. two tit-for-tat agents will always use the same strategies), then we use a ‘‘tie-breaker’’ mechanism such that an agent will not consecutively diverting (in N-protocol) or insisting on their offers (in M-protocol) for more than 3 periods.

Table 2 shows general parameters used in both part 1 and part 2 of our experiments. We assume a very high cost in repeating a bargaining session. Thus, if an agent gets stuck on their valuation in M-protocol, then we will consider it as a breakdown. However, if an N- or NR-Seller gets stuck on its valuation, it will increase its offer to be \$5–\$10 higher than its valuation before monotonically decreasing it again.

3.1.1. Agents' valuation. Figure 5 shows four representative pairs of agents' valuations, which are generated randomly. The vertical axis represents the price, and the horizontal axis represents the time line (in periods). The thick line represents the buyer's valuation in each period, and the thin line represents the seller's valuation. The transaction may be made within the area where the thick line is on the top of the thin line, or when both of them are on the same horizontal line. Near the end of the bargaining period, both lines always overlap at \$100. This gives us a higher assurance of a success in the bargaining if the bargaining prolongs to the deadline, which may favor risk-averse sellers in part 1 of our experiment.

3.1.2. Evaluation criteria. In both part 1 and part 2, two main variables are recorded for evaluation purposes:

- total surplus generated from each group experiment, in terms of the sum of surplus for both buyers and sellers;
- number of breakdowns/successes.

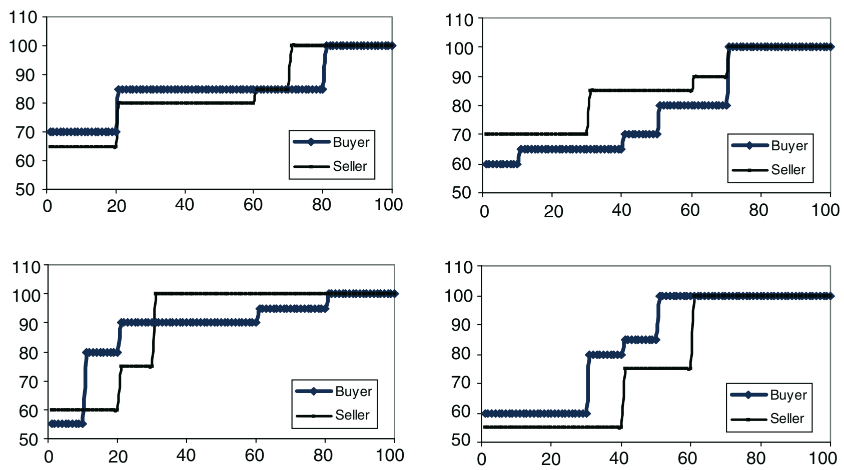

Figure 5. Examples of four pairs (out of 100) of upward valuations used in the experiments.

Table 3. Results of experiment in part 1.

	Ave. surplus	Ave. surplus/ transaction	Success rate
R-averse	6.657	6.967	94.903
R-seeking	7.098	8.773	80.399
N-seller	10.053	10.272	97.901
NR-seller	10.109	10.820	93.323

Based on those two values, three measures are computed: average surplus, average surplus per successful transaction, and success rate.

3.2. *Results*

It is shown in Table 3 that agents in the N-protocol (the two bottom rows) generate higher surplus compared to agents in the M-protocol (the two upper rows). The success rates of the N-protocol are also higher compared to the setting where sellers are risk-taking in the M-protocol (2nd row). These results justify our theoretical analysis that the N-protocol favors negotiation under sellers' upward valuations.

If we compare the results of each group of strategies used by buyers and sellers (four possible combinations of buyer/seller acceptance strategies), we find out that the effect of various strategies used by sellers and buyers is not significant in the N-protocol (Figure 6(a)). Smaller average surpluses are only observed in the case when sellers are type-II and buyers are type-I (shown as S2B1). But when we check the effect on the M-protocol, much smaller average surpluses are only observed when sellers are type-I (see Figure 6(b), in case S1B1 and S1B2). The result suggests that type-I sellers may reduce their overall surplus under the M-protocol, i.e. type-II sellers outperform type-I sellers. Moreover, introducing NR-sellers to the N-protocol reduces the success rate as predicted. The effort to increase the price in order to convince the buyer to make concessions earlier will prolong the bargaining, thus increasing the risk of breakdown. However, as expected, it generates higher total surplus, since more concessions are made and surplus is generated for both parties.

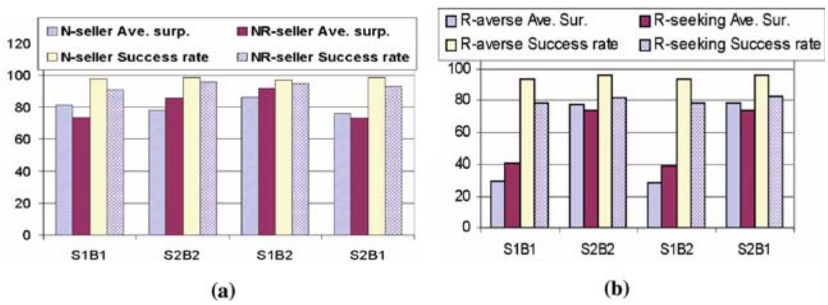

Figure 6. Average surpluses and success rates of four treatments in Table 1 for (a) N-seller and NR-seller, as well as (b) R-averse and R-seeking (average surplus is normalized).

Table 4. Result of experiment in part 2.

	Average surplus	Success rate
M-protocol	8.37	65
N-protocol	9.62	100

In part 2, the success rate and the average surplus generated are higher in the N-protocol, as shown in Table 4. 100% success rate is gained under our experiment in the N-protocol, which is much higher than the 65% in the M-protocol. All the results shown are statistically significant under the *t*-test for confidence level >95%. This result justifies that the N-protocol is better than the M-protocol in terms of average surplus gained by bargainers and their success rate, when both parties have increasing valuations.

4. Discussion, related and future work

As mentioned earlier in this paper, an alternative solution for an agent who has diminishing valuation includes re-opening negotiation with the same opponent or a new one every time it gets stuck in its valuation. However, this situation is not simulated in our experiment here because it is less realistic. Intuitively, if two bargainers have met before, then they may restart their initial offers closer than if they do not know each other in the first place, which may speed up the convergence to the concession. However, there are several weaknesses of this mechanism:

1. If the opponent is programmed to restart the bargaining session using its initial offer, then the chance of breakdown increases, because of the diminishing valuation of the proponent.
2. If there are other parties queuing for the bargaining, then a termination within two agents may reduce the chance for both to encounter each other again, which increases the chance of both to start their bargaining by offering their initial offer (to other new parties).
3. There may exist costs to restart the bargaining, as indicated previously.

In addition, if two bargainers can repeat their bargaining frictionless in the M-protocol, then the efficiency (in term of expected gain and success rate) gained by them is at most as high as the efficiency gained in the N-protocol, which becomes the upper bound of the efficiency of the M-protocol.

Finally, in the N-protocol, agents who use a simple bargaining strategy, such as random strategy or tit-for-tat, can gain higher efficiency (in term of success rate and surplus) compared to agents in the M-protocol. This suggests that the N-protocol may work better than the M-protocol for most type of agents, including those able to maximize expected gain, as shown in the theoretical analysis. However, this may not be true if we have the following types of agents in the bargaining:

– Irrational agents, who increase or decrease their offers arbitrarily; this can be avoided by the restrictions of the M-protocol.
– Misinformed agents, who cannot accept a non-monotonic offer, and therefore retreat from the bargaining immediately. These agents may perceive a non-monotonic offer as a sign of lack of seriousness of their opponent, or of a prolonged bargaining, or of a higher likelihood of breakdown. Thus, they may retreat from the bargaining, and try to find a new opponent.
– Nasty/greedy agents, who use non-monotonic offers to threaten their opponents, delay the bargaining, or mislead their opponents' belief. These agents can be better handled if agents have less pressure from time deadline or breakdown.

More issues may arise as a consequence of the N-protocol, but as shown in both the theoretical analysis and simulation, the N-protocol is better than the M-protocol under certain not uncommon conditions.

4.1. Related work

Research in bargaining has addressed several issues such as the deadline of the bargaining [7], the strategies under incomplete information [3, 11], agents with limited resources [6], and learning by buyer and seller [16]. Most studies on bargaining simplify the model by imposing several assumptions, such as common knowledge of deadline, valuation, risk-neutral attitude, etc. [4, 6, 7, 10]. Even if the exact value is unknown, most research assumes that the probability of a value is known (incomplete information). Under these assumptions, most bargaining problems are solvable using game-theoretic analysis, but at the cost of being less applicable in a real-world situation [5]. In competitive negotiation bargainers will withhold their private information, since revealing private information will subject them to exploitation by their opponents. Consequently, it leads to a more complex model, which is not solvable using game theory. The alternative solution is to use a heuristic approach in the design of negotiating agents [3, 11–13, 16]. However, adopting the N-protocol may alter the results of some analyses in game theory. For instance, let both the buyer and the seller know that the buyer's deadline is earlier than the seller's, and both valuations are decreasing over time and privately known. Then the best strategy for the seller in the M-protocol is to wait until the buyer's deadline and accept the buyer's offer as long as it generates positive surplus. This constitutes the sub-game perfect equilibrium, because there is no incentive for either the seller or the buyer to offer before the buyer's deadline, and no incentive for the seller to accept any offer before the deadline. On the contrary, under the N-protocol, the buyer may offer a higher price at the beginning of negotiation and reduce it along its valuation. Under this strategy, the best response of the seller is to accept the price that generates the highest surplus, which depends on the slope of the seller's valuation. Thus, the equilibrium point will be different for both the M-protocol and the N-protocol.

4.2. Future work

Through theoretical analysis (as mentioned in Section 2), we have identified several reasons for agents to use strategic delay. We plan to conduct further analysis of strategic delay in negotiation. Moreover, we also plan to analyze higher order myopic-K agents. Intuitively, myopic-K agents can be reduced into myopic-1 agents, but this intuition needs to be further explored. As to simulation, we also want to extend our experiment by using more sophisticated agent strategies, such as incorporating learning mechanisms to predict the opponent's valuation. Furthermore, the current simulation is conducted for the case of increasing valuation for both buyer and seller. It remains to study the other cases, such as decreasing valuation for buyers and increasing valuation for sellers, or decreasing valuations for both buyers and sellers but with different slope and deadline, etc. Finally, we also plan to extend the simulation to study the effects of the N-protocol with the presence of various nasty/greedy/irrational agents.

Acknowledgments

This research is supported by the Canadian Natural Sciences and Engineering Research Council (NSERC). This paper is based on [15], which was presented at AAMAS-04.

Notes

1. Not all prices can be offered in M-protocol, for example, a price lower than the buyer's previous offer is not allowed.

References

1. A. R. Admati and M. Perry, "Strategic delay in bargaining," *Rev. Econ. Stud.*, vol. 54, pp. 345–364, 1987.
2. R. Bapna, "When snipers become predators: Can mechanism design save online auctions?" *Commun. ACM*, vol. 46, no. 12, pp. 152–158, 2003.
3. P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," *Int. J. Robot. Auton. Syst.*, vol. 24, no. 3–4, pp. 159–182, 1998.
4. S. S. Fatima, M. Wooldridge, and N. R. Jennings, "Optimal negotiation strategies for agents with incomplete information," in *Proceedings of the ATAL-01*, Seattle, Washington, pp. 53–68, 2001.
5. N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge, "Automated negotiation: prospects, methods, and challenges," *J. Group Decision Negotiation*, vol. 10, no. 2, pp. 199–215, 2001.
6. K. Larson and T. W. Sandholm, "An alternating offers bargaining model for computationally limited agents," in *Proceedings of the AAMAS-02*, Bologna, Italy, pp. 135–142, 2003.
7. S. Kraus, J. Wilkenfeld, and G. Zlotkin, "Multiagent negotiation under time constraints," *Artifi. Intell.*, vol. 75, pp. 297–345, 1995.
8. H. Raiffa, *The Art and Science of Negotiation*, Harvard University Press: Cambridge, USA, 1982.
9. J. S. Rosenschein and G. Zlotkin, *Rules of Ecounter*, MIT Press, 1994.

10. A. Rubinstein, "A bargaining model with incomplete information about time preferences," *Econometrica*, vol. 53, pp. 1151–1172, 1985.
11. K. M. Sim, "A market-driven model for designing negotiation agents," *Comput. Intell.*, vol. 18, no. 4, pp. 618–637, 2002.
12. K. M. Sim, and S. Y. Wang, "Designing flexible negotiation agent with relaxed decision rules," in *Proceedings of the IAT-03*, Halifax, Canada, pp. 282–289, 2003.
13. K. Somefun, E. Gerding, S. Bohte, and H. Poutré, "Automated negotiation and bundling of information goods," in *Proceedings of the AMEC V*, Melbourne, Australia, pp. 1–17, 2003.
14. P. Winoto, G. McCalla, and J. Vassileva, "An extended alternating-offers bargaining protocol for automated negotiation in multi-agent systems," in *Proceedings of the CoopIS* 2002, Irvine, CA, pp. 179–194, 2002.
15. P. Winoto, G. McCalla, and J. Vassileva, "Non-monotonic-offers bargaining protocol," in *Proceedings of the AAMAS-04*, New York, NY, pp. 1072–1079, 2004.
16. D. Zeng, and K. Sycara, "Bayesian learning in negotiation," *Int. J. Hum.-Comput. Studi.*, vol. 48, pp. 125–141, 1998.