

Equilibrium Analyses of Market-driven Agents

KWANG MONG SIM

Department of Information Engineering, Chinese University of Hong Kong,
Shatin, NT, Hong Kong. Email: kmsim@ie.cuhk.edu.hk

Market-driven agents (MDAs) are negotiation agents that make adjustable amounts of concession by considering factors such as competition, deadlines, and trading options. While previous work demonstrates the properties of *MDAs* by experiments, this paper provides mathematical analyses to show that (1) for a given market situation, an *MDA* makes *minimally sufficient concession*, and (2) the strategies adopted by *MDAs* are in equilibrium. The results show that *MDAs* are stable (stability is an essential evaluation criterion for negotiation agents) and they avoid making excessive or inadequate concession in a market situation.

Additional Key Words and Phrases: e-Negotiation, Market-driven Agent, Sequential Equilibrium.

1. INTRODUCTION

Although there are many extant negotiation agents designed for e-commerce ([1,2], just to name a few because of space limitation), strategies adopted by these agents do not take the dynamics of the market into consideration. This paper presents a *market-driven model* [3] for designing negotiation agents (section 2), and proves that for a given market situation, a market-driven agent (*MDA*) does not over-compromise or under-compromise (section 3) and the strategies adopted by *MDAs* are stable (section 4). In dynamic markets where products/services become available, and traders enter and leave a market, the condition for deliberation changes as new opportunities and threats are constantly being introduced. Additionally, negotiation is also bounded by time, since deadline puts negotiators under pressure [4]. Previous empirical results [5] show that in general, *MDAs* [3,5,6] outperform fixed strategy negotiation agents [1] in many situations. In [4], a more flexible negotiation model defines a range of strategies and tactics for generating (counter-)proposals based on time, resource and behaviors of negotiators. The major difference between this work and [4] is that other essential factors such as competition, trading alternatives and differences among negotiators' proposal are also considered. In multilateral negotiations, it is intuitive to think that a negotiator's bargaining power is affected by the number of competitors and trading alternatives. Good options give a negotiator more "power" because the negotiating party needs not pursue the negotiation with any sense of desperation [7, p157], and negotiations may break down because the parties cannot resolve their differences [8, p94]. However, this work does not compete with [1], [4] or the current literature, but rather complements them by modeling the behaviors of negotiation agents in changing market situations. In designing *MDAs* [5,6], factors such as competition, trading alternatives, differences among negotiators and deadline are taken into account when an agent determines the amount of concession at each negotiation round (section 2).

This research was supported by a direct research grant (project code: 2050288) from the Faculty of Engineering, Chinese University of Hong Kong.

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2003 ACM

2. A MARKET-DRIVEN MODEL FOR DESIGNING NEGOTIATION AGENTS

In making concession, an *MDA* assesses the probability p of successfully completing a deal in different market situations. One way to assess p is to consider the value of k – the *spread* (difference) between an agent's bid/offer and that of others [3]. Let $v^{a \rightarrow j}_i$ be the utility of an agent and $w^{j \rightarrow a}_i$ be the best utility among all its trading partners in the current trading cycle. The (best) spread in the current cycle i is $k_i = v^{a \rightarrow j}_i - w^{j \rightarrow a}_i$. A market-driven agent makes concession by determining the utility of its *next* bid/offer $v^{a \rightarrow j}_{i+1}$ using the sum of the spread in the *next* cycle k_{i+1} and $w^{j \rightarrow a}_i$. Hence, $v^{a \rightarrow j}_{i+1} = k_{i+1} + w^{j \rightarrow a}_i$. An agent strives to attain the highest possible utility while maintaining a reasonable level of p in the next trading cycle. It is more probable to complete a deal when k_{i+1} is small, and vice versa. k_{i+1} is determined by assessing the current market situation, taking into account factors such as *opportunity*, *competition* and *deadline*:

$$k_{i+1} = [O(n^a_{i, <w^{j \rightarrow a}_i, v^{a \rightarrow j}_i>}, C(m^a_{i, n^a_i}) T(t, \tau, \epsilon)] k_i \quad (1)$$

Trading opportunity: $O(n^a_{i, <w^{j \rightarrow a}_i, v^{a \rightarrow j}_i>})$ determines the amount of concession based on trading alternatives (number of trading partners) and differences in offers/bids [3,6]. Suppose agent a engages j in round i . a 's last proposal generates a payoff of $v^{a \rightarrow j}_i$ for itself, and j 's last counter-proposal generates a payoff of $w^{j \rightarrow a}_i$ for a . If a accepts j 's counter-proposal, it will obtain $w^{j \rightarrow a}_i$ with certainty. If a insists on its last proposal, and i) if j accepts it, a will obtain $v^{a \rightarrow j}_i$ and ii) if j does not accept it, a may be subjected to a conflict utility c^a . c^a is the worst possible utility for a (i.e. an agent's payoff in the absent of agreement [9]), hence, $w^{j \rightarrow a}_i \geq c^a$. If j does not accept a 's last proposal, a may ultimately have to settle with lower utilities (the lowest possible being the conflict utility), if there are changes in the market situation in subsequent cycles. For instance, a may face more competitions in subsequent cycles, and may have to ultimately accept a utility that is lower than $w^{j \rightarrow a}_i$ (possibly as low as c^a). If the subjective probability of a obtaining c^a is $p^{a \leftrightarrow j}_{c,i}$ (*conflict probability*) and the probability that a achieves $v^{a \rightarrow j}_i$ is $1 - p^{a \leftrightarrow j}_{c,i}$, then, according to Zeuthen's analysis [9], if a insists on holding its last proposal, a will obtain a payoff of $(1 - p^{a \leftrightarrow j}_{c,i})v^{a \rightarrow j}_i + p^{a \leftrightarrow j}_{c,i}c^a$. Hence, a will find that it is advantageous to insist on its last proposal only if $(1 - p^{a \leftrightarrow j}_{c,i})v^{a \rightarrow j}_i + p^{a \leftrightarrow j}_{c,i}c^a \geq w^{j \rightarrow a}_i$. Hence, $p^{a \leftrightarrow j}_{c,i} \leq (v^{a \rightarrow j}_i - w^{j \rightarrow a}_i) / (v^{a \rightarrow j}_i - c^a)$. Consequently, the maximum value of $p^{a \leftrightarrow j}_{c,i}$ is the *highest probability of a conflict* that a may encounter in round i , given as [3,6]:

$$p^{a \leftrightarrow j}_{c,i} = (v^{a \rightarrow j}_i - w^{j \rightarrow a}_i) / (v^{a \rightarrow j}_i - c^a) \quad (2)$$

$p^{a \leftrightarrow j}_{c,i}$ is a ratio of difference between two utilities. While $v^{a \rightarrow j}_i - w^{j \rightarrow a}_i$ measures the cost of *accepting the trading agent's last offer* (the spread k or difference between the (counter-)proposals of a and j), $v^{a \rightarrow j}_i - c^a$ measures the cost of *provoking a conflict*. Hence, the probability that a will obtain a utility $v^{a \rightarrow j}_i$, with at least one of its n^a_i trading partners is

$$O(n^a_{i, <w^{j \rightarrow a}_i, v^{a \rightarrow j}_i>}, v^{a \rightarrow j}_i) = 1 - \prod_{j=1}^{n^a_i} \frac{v^{a \rightarrow j}_i - w^{j \rightarrow a}_i}{(v^{a \rightarrow j}_i - c^a)} \quad (3)$$

With ample trading options, $O(n^a_{i, <w^{j \rightarrow a}_i, v^{a \rightarrow j}_i>}, v^{a \rightarrow j}_i)$ approaches 1, an *MDA* has more bargaining power, and hence, makes smaller compromises. Since there are a very large number of trading partners, the likelihood that some other agent's proposals is potentially closer to the *MDA*'s proposal is higher (proven in [3]). As $O(n^a_{i, <w^{j \rightarrow a}_i, v^{a \rightarrow j}_i>}, v^{a \rightarrow j}_i)$ reduces to 0, an *MDA*'s bargaining power decreases. It makes larger compromises. Even with a very large number of trading options, it would be difficult for an *MDA* to reach a consensus if there are large differences between itself and other negotiating parties (proven in [3]).

Competition: $C(m^a_{i, n^a_i})$ determines the probability that an agent a is ranked as the *most preferred* trading partner by at least one other agent in round i . If a has m^a_{i-1}

competitors, and one trading partner, the probability that a is not the most preferred trading partner is $(m_i^a - 1)/m_i^a$. If a has $m_i^a - 1$ competitors, and n_i^a trading partners, the probability that a is *not* the most preferred partner of *all* its trading partners is $[(m_i^a - 1)/m_i^a]^{n_i^a}$. Hence, in round i , the probability that a is considered the most preferred trading partner by at least one agent is:

$$C(m_i^a, n_i^a) = 1 - [(m_i^a - 1)/m_i^a]^{n_i^a} \quad (4)$$

If there is a very large number of trading partners n_i^a , then the chances of being considered the most preferred trading partner by some parties is higher. When the number of competitors m_i^a increases significantly, the likelihood of being considered the most preferred trading partner becomes much lower. These claims were proven in [3].

Deadline: $T(t, \tau, \lambda)$ is a time-dependent function given as:

$$T(t, \tau, \lambda) = 1 - (t/\tau)^\lambda \quad (5)$$

where t is current trading time, τ is the deadline, and λ is a nonnegative temporal sensitivity factor. At the start of negotiation ($t=0$), an *MDA* offers its initial price since $T(t, \tau, \lambda) = 1 - (0/\tau)^\lambda = 1$; when its deadline is reached, the *MDA* offers its reserve (final) price since $T(t, \tau, \lambda) = 1 - (\tau/\tau)^\lambda = 0$. Although agents can select from an infinitely many strategies with respect to remaining trading time (one for each value of λ , where $\lambda \geq 0$), they can be classified as follows [3]:

1) *Linear:* $\lambda=1$ and $T(t, \tau, \varepsilon) = 1 - (t/\tau)$. An *MDA* makes a *constant* rate of concession.

2) *Conservative:* $k_{i+1} = [1 - (t/\tau)^\lambda]k_i$ where $\lambda > 1$, an *MDA* makes smaller concession in early rounds and larger concessions in later rounds.

3) *Conciliatory:* $k_{i+1} = [1 - (t/\tau)^\lambda]k_i$ where $\lambda < 1$, an *MDA* makes larger concessions in the early trading rounds and smaller concessions at the later stage.

4) The extreme and special cases when $\lambda = 0$ (always giving in) and $\lambda = \infty$ (never give in) will not be considered because they represent the situations when no negotiation is needed and non-negotiable trading respectively.

In [10] *MDAs* are designed with a time-dependent function $T(t, \tau, \varepsilon) = 1 - (t/\tau)^{1/\varepsilon}$ that models its degree of sensitivity to approaching deadlines with respect to the *MDA*'s eagerness ε . ε represents the user's desire to complete the deal. Given 2 agents a_1 and a_2 with the same deadline $\tau_1 = \tau_2$, but different eagerness ε_1 and ε_2 with $\varepsilon_1 > \varepsilon_2$ then at time t , $T(t, \tau_1, \varepsilon_1) < T(t, \tau_2, \varepsilon_2)$; a_1 having higher eagerness makes more concessions than a_2 . Consequently, a_1 adopts a strategy that is more likely to complete a deal faster. This corresponds with the intuition that negotiators that are more eager to complete a deal are more likely to concede faster. In [10], *MDAs* are designed with only conservative (with $0 < \varepsilon < 1$) and linear strategies (with $\varepsilon = 1$).

3. MINIMALLY SUFFICIENT CONCESSION

Given the current market situation, an *MDA* strives to attain the highest possible utility while maintaining a reasonable probability p of reaching a consensus. However, both increasing an *MDA*'s utility and increasing p are incompatible objectives. One way to deal with incompatible objectives is to determine the most important objective, and express all other objectives as constraints. For *MDAs*, maximizing utility is the most important objective, but they are also designed to maintain a minimum p . This section shows that with respect to opportunity and competition (*propositions 1 and 2*), the amount of concession made by an *MDA* is *minimally sufficient* [11].

Definition 1: A *sufficient concession* is one that increases a market-driven agent's probability of reaching a consensus.

Definition 2: A concession Δk is *minimally sufficient* if it achieves the *highest possible* utility for a market-driven agent while maintaining a minimum probability p_{\min} of reaching a consensus.

Proposition 1: With respect to trade opportunity, the amount of concession Δk_{opp} made by a market-driven agent is minimally sufficient.

Proof: The proof proceeds by induction on the number of trading partners n .

Base case: When $n=1$. Let Δk_{opp} be the amount of concession made by an *MDA*. Since $p^{a \leftarrow j}_{c,i} \leq (v_i^{a \rightarrow j} - w_i^{j \rightarrow a}) / (v_i^{a \rightarrow j} - c^a)$ and $p = 1 - p^{a \leftarrow j}_{c,i}$, it follows that p is the minimum probability of reaching a consensus when $p^{a \leftarrow j}_{c,i} = (v_i^{a \rightarrow j} - w_i^{j \rightarrow a}) / (v_i^{a \rightarrow j} - c^a) = k_i / (v_i^{a \rightarrow j} - c^a)$. When $n=1$, with respect to trade opportunity, $k_{i+1} = O(1, \langle w_i^{l \rightarrow a}, v_i^{a \rightarrow l} \rangle) k_i$ and $\Delta k_{opp} = k_i - k_{i+1} = 1 - O(1, \langle w_i^{l \rightarrow a}, v_i^{a \rightarrow l} \rangle) = (v_i^{a \rightarrow l} - w_i^{l \rightarrow a}) / (v_i^{a \rightarrow l} - c^a) = k_i / (v_i^{a \rightarrow l} - c^a)$.

Inadequate concession: Let Δk_u be any amount of concession such that $\Delta k_u < \Delta k_{opp}$ and p_u be the corresponding probability of reaching a consensus if the *MDA* concedes with Δk_u . Although Δk_u generates a higher utility for the *MDA* because the corresponding spread $k_u > k_i$, it follows that $k_u / (v_i^{a \rightarrow l} - c^a) > p^{a \leftarrow l}_{c,i} = k_i / (v_i^{a \rightarrow l} - c^a)$, and $p_u < p = p_{\min}$, which contradicts definition 2.

Excessive concession: Let Δk_{ov} be any amount of concession such that $\Delta k_{ov} > \Delta k_{opp}$ and p_{ov} be the corresponding probability of reaching a consensus if the *MDA* concedes with Δk_{ov} . Since the corresponding spread $k_{ov} < k_i$, it follows that $p_{ov} > p = p_{\min}$. However, Δk_{ov} is sufficient (definition 1) but not minimally sufficient because it generates a lower utility for the *MDA*. Hence, Δk_{opp} is minimally sufficient.

Induction case: When $n > 1$, assume as an induction hypothesis that an *MDA* makes minimally sufficient concession when $n < N$. With respect to competition, when $n = N$, it follows that $k_{i+1} = O(N, \langle w_i^{j \rightarrow a}, v_i^{a \rightarrow j} \rangle) k_i$ and $\Delta k_{opp} = k_i - k_{i+1} = 1 - O(N, \langle w_i^{j \rightarrow a}, v_i^{a \rightarrow j} \rangle)$

$$= \prod_{j=1}^N \frac{v_i^{a \rightarrow j} - w_i^{j \rightarrow a}}{(v_i^{a \rightarrow j} - c^a)} = \frac{k_i^{a \rightarrow 1} \times k_i^{a \rightarrow 2} \times \dots \times k_i^{a \rightarrow N}}{(v_i^{a \rightarrow j} - c^a)^N} = \frac{k_i^{a \rightarrow 1} \times k_i^{a \rightarrow 2} \times \dots \times k_i^{a \rightarrow r-1}}{(v_i^{a \rightarrow j} - c^a)^{r-1}} \times \frac{k_i^{a \rightarrow r}}{(v_i^{a \rightarrow r} - c^a)} \times \frac{k_i^{a \rightarrow r+1} \times \dots \times k_i^{a \rightarrow N}}{(v_i^{a \rightarrow j} - c^a)^{N-r}}.$$

Let r be any one of the N trading partners of the *MDA*. Let Δk_r be the amount of concession that the *MDA* makes to r . Let $\Delta k_r = \frac{k_i^{a \rightarrow r}}{(v_i^{a \rightarrow r} - c^a)} = \frac{v_i^{a \rightarrow r} - w_i^{r \rightarrow a}}{(v_i^{a \rightarrow r} - c^a)}$. By the same

argument given above (the base case), Δk_r is also minimally sufficient.

Let $\{1, 2, \dots, r-1, r+1, \dots, N\}$ be the set of *all other* trading partners of the *MDA* other than r and $\Delta k_{N-1} = \frac{k_i^{a \rightarrow 1} \times k_i^{a \rightarrow 2} \times \dots \times k_i^{a \rightarrow r-1}}{(v_i^{a \rightarrow j} - c^a)^{r-1}} \times \frac{k_i^{a \rightarrow r+1} \times \dots \times k_i^{a \rightarrow N}}{(v_i^{a \rightarrow j} - c^a)^{N-r}}$ such that $\Delta k_{opp} = \Delta k_{N-1} \times \Delta k_r$. By

induction hypothesis, Δk_{N-1} is minimally sufficient. Since Δk_{N-1} and Δk_r both are minimally sufficient, it follows that Δk_{opp} is also minimally sufficient \blacklozenge

Proposition 2: With respect to competition in a given market size with m numbers of competitors and n numbers of trading partners, the amount of concession Δk_{com} made by a market-driven agent is minimally sufficient.

Proof: Since market-driven agents are utility maximizing agents, a market-driven agent B_l is *not* likely to reach a consensus with another market-driven agent S_i if B_l 's proposal is not ranked as the best proposal by S_i . Suppose B_l has n trading partners and m competitors. Let $E_1, E_2, \dots, E_{n-1}, E_n$ be the events that B_l 's proposal is not ranked the best proposal by *exactly* 1, 2, ..., $n-1$, n other trading partners respectively. The *worst case* scenario is E_n , when B_l 's proposal is not ranked the best proposal by *all* of its n trading partners, in which case, it is *most unlikely* to reach a consensus with any $S_i \in \mathbf{S}$. Let P_c be the probability of conflict with respect to competition. P_c is the largest possible when B_l 's proposal is not ranked as the best proposal by *all* $S_i \in \mathbf{S}$. Let p be the probability of reaching a consensus. Since $p=1-P_c$, with respect to competition, p is the smallest possible when E_n occurs. $\Delta k_{com} = k_t - k_{t+1} = 1 - C(m^a, n^a) = \left[(m^a - 1) / m^a \right]^{n^a}$ is determined based on the assumption that E_n occurs.

Inadequate-concession: By making a concession $\Delta k_u < \Delta k_{com}$, it will be *even less likely* that B_l 's proposal is ranked the best proposal by some other agent. Consequently, the corresponding probability $p_u < p = p_{min}$ violates definition 2.

Excessive-concession: Let Δk_{ov} be any amount of concession such that $\Delta k_{ov} > \Delta k_{com}$ and the corresponding probability of completing the deal be p_{ov} . Although $p_{ov} > p = p_{min}$, Δk_{ov} generates a lower utility for the *MDA* since the corresponding spread $k_{ov} < k$. Thus, Δk_{ov} is sufficient but not minimally sufficient. Hence, Δk_{com} is minimally sufficient ♦

4. SEQUENTIAL EQUILIBRIUM

The amount of concession that an *MDA* makes with respect to its deadline is $\Delta k_{DL} = 1 - T(t, \tau, \varepsilon) = (t / \tau)^{1/\varepsilon}$. Since Δk_{DL} is not determined by the number of trading partners, for the purpose of analyzing the time-dependent strategies of *MDAs*, a multi-lateral negotiation is treated as many threads of one-to-one negotiation. This section shows that (i) when an *MDA* adopts a conservative strategy, it requires more trading rounds to reach a consensus with its trading partner (*lemma 3.1*) than using either the conciliatory or linear strategy, (ii) the conservative strategy is the optimal strategy for an *MDA* regardless of the strategy adopted by its trading partner (*lemma 3.2*), and (iii) the strategies of *MDAs* are in equilibrium (*proposition 3*).

Lemma 3.1: A market-driven agent b adopting a strategy with $\lambda^b > 1$ requires more negotiation rounds to reach an agreement than when b adopts a strategy with $\lambda^b \leq 1$ regardless of the strategy λ^s adopted by its trading partner s .

Proof: Let T_1 (respectively, T_2) be the number of rounds that b adopting a strategy with $\lambda^b > 1$ (respectively, $\lambda^b \leq 1$) completes a deal with s . Let $k_{1,t}$ (respectively, $k_{2,t}$) be the spread between b adopting λ^b_1 (respectively, λ^b_2) and s in round t . Since market-driven strategies are designed for both buyer and seller agents, it can be assumed without loss of generality that b makes a concession first, hence, at round t it follows that:

$$k_{a,t} = \begin{cases} [T(1, \tau^b, \lambda^b_a) \times T(2, \tau^s, \lambda^s) \dots \times T(t, \tau^s, \lambda^s)] \times k_{a,0} & t \text{ is odd} \\ [T(1, \tau^b, \lambda^b_a) \times T(2, \tau^s, \lambda^s) \dots \times T(t, \tau^b, \lambda^b_a)] \times k_{a,0} & t \text{ is even} \end{cases}$$

where $a=1$ if b adopts λ^b_1 , otherwise b adopts λ^b_2 . Consequently, $k_{a,t}$ can be re-written as:

$$k_{a,t} = \begin{cases} \prod_{i=0}^I \left[1 - \left(\frac{2i+1}{\tau^b} \right)^{\lambda_a^b} \right] \times \left[1 - \left(\frac{2i}{\tau^s} \right)^{\lambda^s} \right] \times k_{a,0} & , t=2I+1 \\ \prod_{i=0}^I \left[1 - \left(\frac{2i+1}{\tau^b} \right)^{\lambda_a^b} \right] \times \left[1 - \left(\frac{2i+2}{\tau^s} \right)^{\lambda^s} \right] \times k_{a,0} & , t=2I+2 \end{cases}$$

where $a=1$ if \mathbf{b} adopts λ_1^b , otherwise \mathbf{b} adopts λ_2^b . Since T_1 (respectively, T_2) is the number of rounds for \mathbf{b} adopting λ_1^b (respectively, λ_2^b) to reach an agreement with \mathbf{s} , it follows that $t \rightarrow T_1$ iff $k_{1,t} \rightarrow 0$ and $t \rightarrow T_2$ iff $k_{2,t} \rightarrow 0$. Since the initial spreads for both strategies λ_1^b and λ_2^b is $k_{1,0} = k_{2,0}$, and for all i , $\left[1 - \left(\frac{2i+1}{\tau^b} \right)^{\lambda_1^b} \right] > \left[1 - \left(\frac{2i+1}{\tau^b} \right)^{\lambda_2^b} \right]$, it follows

$$\text{that:} \quad \text{For all } t > 0, k_{1,t} > k_{2,t} \quad (6)$$

Consequently, at round t , if $k_{1,t} \rightarrow 0$, then $t \rightarrow T_1$ and $k_{2,t} \rightarrow 0$. But, $k_{2,t} \rightarrow 0$ does not imply $k_{1,t} \rightarrow 0$. Hence, it follows that $T_1 > T_2$, and \mathbf{b} adopting a strategy with λ_1^b requires more negotiation rounds to reach an agreement with \mathbf{s} than when it adopts a strategy with λ_2^b .

Lemma 3.2: The conservative strategy ($\lambda_1^b > 1$) is the optimal strategy for a market-driven agent \mathbf{b} regardless of the strategy that its trading partner \mathbf{s} adopts.

Proof: Let T_1 (respectively, T_2) be the number of rounds that \mathbf{b} adopting $\lambda_1^b > 1$ (respectively, $\lambda_2^b \leq 1$) reaches an agreement with \mathbf{s} . Without loss of generality, assume that \mathbf{b} makes concession first. From lemma 3.1, it follows that $T_1 > T_2$, hence, at T_2 , there are two possible cases:

1) $T_2 = 2I+1$: For offer $l^s_{T_2}$ sent by \mathbf{s} in round T_2 :

$$\begin{aligned} l^{s \rightarrow b_1}_{T_2} &= l^s_0 - \sum_{i=1}^I (k_{1,2i}) \left(\frac{2i}{\tau^s} \right)^{\lambda^s} && \text{if } \mathbf{b} \text{ adopts } \lambda_1^b \\ l^{s \rightarrow b_2}_{T_2} &= l^s_0 - \sum_{i=1}^I (k_{2,2i}) \left(\frac{2i}{\tau^s} \right)^{\lambda^s} && \text{if } \mathbf{b} \text{ adopts } \lambda_2^b \end{aligned}$$

From (6) in lemma 3.1, for $t > 0$, $k_{1,t} > k_{2,t}$, hence, it follows that $k_{1,T_2} > k_{2,T_2}$ and $l^{s \rightarrow b_1}_{T_2} < l^{s \rightarrow b_2}_{T_2}$. Since at T_2 , \mathbf{b} adopting λ_2^b reaches an agreement with \mathbf{s} , $l^{s \rightarrow b_2}_{T_2} = l^{b_2 \rightarrow s}_{T_2}$. Since $T_1 > T_2$, \mathbf{s} and \mathbf{b} adopting λ_1^b will continue to make concession from $t > T_2$ until $t < T_1$. Hence, it follows that $l^{s \rightarrow b_1}_{T_1} < l^{s \rightarrow b_1}_{T_2}$. At T_1 when \mathbf{s} and \mathbf{b} adopting λ_1^b reach an agreement, $l^{b_1 \rightarrow s}_{T_1} = l^{s \rightarrow b_1}_{T_1}$. Consequently, it follows that $l^{b_1 \rightarrow s}_{T_1} = l^{s \rightarrow b_1}_{T_1} < l^{s \rightarrow b_1}_{T_2} < l^{s \rightarrow b_2}_{T_2} = l^{b_2 \rightarrow s}_{T_2}$. Hence, $v_b(l^{b_1 \rightarrow s}_{T_1}) > v_b(l^{b_2 \rightarrow s}_{T_2})$, where v_b is a function that generates the utilities of the proposals $l^{b_1 \rightarrow s}_{T_1}$ and $l^{b_2 \rightarrow s}_{T_2}$.

2) $T_2 = 2I+2$: For offer $l^s_{T_2}$ sent by \mathbf{s} in round T_2

$$\begin{aligned}
I^{s \rightarrow b}_{T_2} &= I^s_0 - \sum_{i=1}^{I+1} (k_{1,2i}) \left(\frac{2i}{\tau^s}\right)^x && \text{if } \mathbf{b} \text{ adopts } \lambda^{b_1} \\
I^{s \rightarrow b_2}_{T_2} &= I^s_0 - \sum_{i=1}^{I+1} (k_{2,2i}) \left(\frac{2i}{\tau^s}\right)^x && \text{if } \mathbf{b} \text{ adopts } \lambda^{b_2}
\end{aligned}$$

The proof that $v_b(I^{b_1 \rightarrow s}_{T_1}) > v_b(I^{b_2 \rightarrow s}_{T_2})$ for case (2) is analogous to case (1) and is omitted due to space limitation. Since $v_b(I^{b_1 \rightarrow s}_{T_1}) > v_b(I^{b_2 \rightarrow s}_{T_2})$ for cases (1) and (2), the *conservative strategy* is the optimal strategy for \mathbf{b} regardless of the strategy that \mathbf{s} chooses

♦

Sequential equilibrium: In analyzing the mutual strategic behaviors of market-driven agents, negotiation is viewed as a game Γ of incomplete information. In Γ , an *MDA* has imperfect information about the deadline and strategy of its trading partner. Since *MDAs* are imperfectly informed, the idea of *sequential equilibrium* [12, p95] is used to demonstrate that their strategies are in equilibrium [12, p41]. The notion of sequential equilibrium requires the specifications of: a *strategy profile* π and a *belief system* μ . Let $\pi = (S^b, S^s)$ such that S^b , and S^s are the strategies of buyer \mathbf{b} and seller \mathbf{s} . μ is a function that assigns a probability $\mu(h) \in [0, 1]$ to each decision node h in Γ such that $\sum_{h \in I} \mu(h) = 1$ for all

information sets I . For instance, at $t=1$ in Fig. 1, h_2, h_3 , and h_4 form the information set I_1 of \mathbf{s} , and $\mu(h_2) + \mu(h_3) + \mu(h_4) = 1$. μ represents an *MDA's* belief about the history of negotiation.

Definition 3: A sequential equilibrium of Γ is a pair (π, μ) that satisfies the following:

- (i) *Sequential rationality:* At each information set I of each agent a_i , the strategy of a_i is the best response to other agents' strategies (ie, optimal), given a_i 's beliefs in I .
- (ii) *Consistency:* At each information set I , a_i updates its beliefs based on strategy profile π using *Bayes rules*.

By Bayes rules, at I , the probability of reaching a decision node h_i according to π is

$$\frac{P(h_i, \pi)}{\sum_{h \in I} P(h, \pi)}$$

respectively according to π . For instance, at I_1 , if \mathbf{b} chooses B_1, B_2 , and B_3 with probabilities p_2, p_3 , and p_4 respectively, then \mathbf{s} must assign probabilities $\frac{p_2}{p_2 + p_3 + p_4}$,

$\frac{p_3}{p_2 + p_3 + p_4}$ and $\frac{p_4}{p_2 + p_3 + p_4}$ of reaching h_2, h_3 , and h_4 respectively.

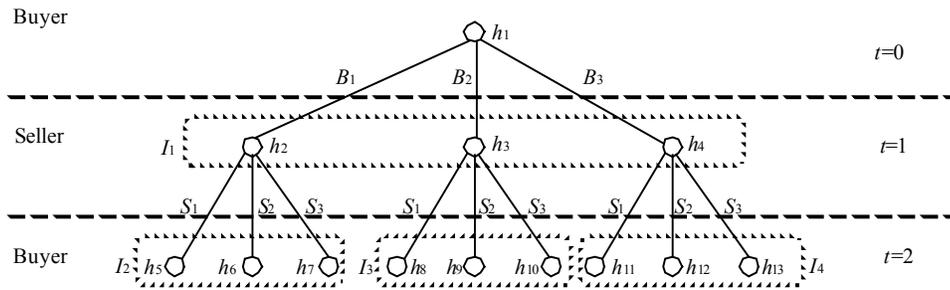
Proposition 3: There exists a unique sequential equilibrium of Γ .

Proof: In Γ , buyer \mathbf{b} and seller \mathbf{s} send proposals alternately. Assume without loss of generality that \mathbf{b} sends its proposal first. At $t=0$, buyer \mathbf{b} starts the negotiation by choosing its optimal strategy B_1 (lemma 3.2), and the game proceeds to $t=1$. At $t=1$, \mathbf{s} is unaware of the strategy that \mathbf{b} adopts, hence, at information set I_1 , \mathbf{s} does not know whether play reaches h_2, h_3 , or h_4 . However, regardless of \mathbf{s} 's belief about the history of negotiation, i.e. whether h_2, h_3 , or h_4 is reached, \mathbf{s} will choose S_1 because by lemma 3.2, S_1 is the optimal strategy for \mathbf{s} at I_1 regardless of the strategy that \mathbf{b} chooses at $t=0$. At $t=2$, even though \mathbf{b} does know which strategy \mathbf{s} adopts, \mathbf{b} knows that I_2 is reached with

probability of 1, because at $t=0$, \mathbf{b} chooses B_I . From lemma 3.2, B_I is also the optimal strategy for \mathbf{b} at I_2 . Iterating the arguments at $t=0$, $t=1$ and $t=2$, at every information set that is \mathbf{b} 's turn to move, the optimal strategy for \mathbf{b} is B_I , and at every information set that is \mathbf{s} 's turn to move, the optimal strategy for \mathbf{s} is S_I . Consequently, the strategy profile $\pi=(B_I, S_I)$ satisfies the requirement for sequential rationality. Furthermore, both \mathbf{b} and \mathbf{s} adhere to (B_I, S_I) regardless of their beliefs about the negotiation histories. Hence, (B_I, S_I) is a sequential equilibrium of Γ . In addition, the sequential equilibrium at (B_I, S_I) is also unique because B_I and S_I are the optimal strategies of \mathbf{b} and \mathbf{s} at every of their information sets, and there do not exist any other equilibrium where either \mathbf{b} or \mathbf{s} deviates from its optimal strategy ♦

5. DISCUSSION AND CONCLUSION

Propositions 1 and 2 further elaborate and support the claim previously made in [3] that *MDAs* are designed to avoid making excessive concessions in favorable markets or to make inadequate compromises in unfavorable markets. Consequently, with respect to market opportunity and competition, an *MDA* does not have the incentive of making any other amounts of concession other than Δk_{opp} and Δk_{com} . Proposition 3 shows that the time-dependent strategies of *MDAs* are in equilibrium. In [10], (τ, ϵ) forms the time-dependent strategy of an *MDA*. Hence, *MDAs* with different ϵ adopts different (conservative) strategies in making concession. From a negotiator's perspective, one resists time pressure by adjusting one's level of aspiration to what is believed to be optimal within the available time [13]. In [10], the amount of concession Δk_{DL} made by an *MDA* with respect to deadline is based on ϵ and regardless of the strategy adopted by its trading partner. Consequently, the strategies of *MDAs* are in equilibrium because with respect to opportunity, competition and deadline, every *MDA* does not find any incentive to deviate from making Δk_{opp} , Δk_{com} and Δk_{DL} respectively. Theoretical results from this work show that *MDAs* are stable (and are designed to behave in the desired manner). In addition to showing the desirable property of stability, which is one of the essential evaluation criteria for automated negotiation systems [11, p21], this work also supplements and complements the empirical results obtained from previous work on *MDAs* [6,10]. Both theoretical results from this work, as well as empirical results previously obtained from [6,10] collectively show that *MDAs* are stable, and achieve relatively high utility and success rates. While this work (in its present stage) assumes that all negotiating parties are *MDAs*, future work will explore analyses of negotiation between *MDAs* and agents adopting other strategies.



B_1, S_1 --Conservative strategy B_2, S_2 --Linear strategy B_3, S_3 --Conciliatory strategy
 Fig. 1. A Negotiation Game with Incomplete Information

REFERENCES

- [1] A. Chavez and P. Maes. *Kasbah: An agent marketplace for buying and selling goods*. 1st Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Tech., London, 1996.
- [2] D. Zeng and K. Sycara. *Bayesian Learning in Negotiation*. Int. J. Human-Computer Studies, Vol 48, pp. 125-141, 1998.
- [3] K. M. Sim. *A Market-driven Model for Designing Negotiation Agents*. In Computational Intelligence, Special issue in Agent Tech. for E-commerce, vol. 18, no. 4, 2002.
- [4] P. Faratin, C. Sierra, N. R. Jennings, *Negotiation Decision Functions for Autonomous Agents*. Int. J. Robotics and Autonomous System. Vol.24, No.3: 159-182
- [5] K. M. Sim and E. Wong. *Towards Market Driven Agents for E-auction*. IEEE Trans. Syst., Man Cybern, A, Vol.31. No.6, Nov.2001, pp 274-284.
- [6] K. M. Sim and C.Y. Choi. *Agents that React to Changing Market Situations*. IEEE Trans. Syst., Man Cybern B: Cybernetics, Vol. 33, No. 2, April 2003, pp 188-201.
- [7] R. Lewicki. *Negotiating Strategically*. In A. Cohn (ed.): The Portable MBA in Management, pp. 147-289, John-Wiley Publisher, 1993.
- [8] M. Bazerman and M. Neale. *Negotiating Rationally*. The Free Press, NY, 1992.
- [9] F. Zeuthen. *Problem of Monopoly and Economic Warfare*. London UK, Routledge and Kegan-paul.1967
- [10] K. M. Sim and S.Y. Wang. *Designing Flexible Negotiation Agent with Relaxed Decision Rules*. To appear in Proc. of the 2003 IEEE/WIC Int. Conf. on Intelligent Agent Tech.
- [11] J. Rosenschein., and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. Cambridge, Mass.: MIT Press, 1994.
- [12] Osborne M.J. and Rubinstein A. *Bargaining and Markets*. Academic Press, 1990.
- [13] Marsh P.D.V. *Contract Negotiation Handbook* (3rd Ed.), MPG Books Ltd., Bodmin, 2001.