Monotonic Mixing of Decision Strategies for Agent-based Bargaining

Jan Richter, Matthias Klusch, and Ryszard Kowalczyk

Swinburne University of Technology, Melbourne, Australia
German Research Center for Artificial Intelligence, Saarbrücken, Germany
{jrichter,rkowalczyk}@swin.edu.au,klusch@dfki.de

Abstract. In automated bargaining a common method to obtain complex concession behaviour is to mix individual tactics, or decision functions, by a linear weighted combination. In such systems, the negotiation process between agents using mixed strategies with imitative and non-imitative tactics is highly dynamic, and non-monotonicity in the sequence of utilities of proposed offers can emerge at any time even in cases of individual cooperative behaviour and static strategy settings of both agents. This can result in a number of undesirable effects, such as delayed agreements, significant variation of outcomes with lower utilities, or a partial loss of control over the strategy settings. We propose two alternatives of mixing to avoid these problems, one based on individual imitative negotiation threads and one based on single concessions of each tactic involved. We prove that both produce monotonic sequences of utilities over time for mixed multi-tactic strategies with static and dynamically changing weights thereby avoiding such dynamic effects, and show with a comparative evaluation that they can provide utility gains for each agent in many multi-issue negotiation scenarios.

1 Introduction

Automated negotiation between rational software agents is considered key to facilitate intelligent decision-making between two or more parties which are in conflict about their goals. In such environments, the agents acting on behalf of their users (buyers, sellers) have no or only uncertain knowledge about opponent's behaviours and can use a range of different strategies to conduct negotiation. In automated bargaining or bilateral negotiation, two rational agents negotiate by alternatively exchanging offers over issues of a service or product where each agent has the preference to achieve the highest possible utility from an outcome while the common interest is to find an agreement before the deadline. A common approach for the agents to propose offers is to use individual decision functions, also called tactics, and mix them based on linear weighted combinations to create complex concession behaviour in the form of negotiation strategies. For instance, the prominent service-oriented negotiation model by Faratin et al [6] proposes different types of tactics such as behaviour-, time- or resource-dependent that can be mixed together. In this paper, we demonstrate that non-monotonic behaviour in the form of non-monotonic offer curves and utility sequences can easily emerge at any time as a result of the *dynamic effects of an agent system in which the agents use mixed strategies* involving behaviour-dependent and -independent tactics. In other words, such a system created by negotiating agents using mixed strategies may generate non-monotonic behaviour even when strategy settings and mixing weights of both agents are *static* and all involved tactics are cooperative in that their individual concession behaviour is monotonic. As a result, the agent's own aggregated utilities over all negotiated issues can be non-monotonic as well, implying that it proposes offers increasing its own overall utility. These effects are often considered undesirable for automated single- and multi-issue negotiations [10] [6] as they may delay final agreements, have significantly varying outcomes with lower utilities, or result in a partial loss of the agents' control over their strategy due to the high sensitivity of parameters. In this paper, we therefore provide examples as well as an analysis and evaluation of the traditional mixing method with respect to the self-emergence of non-monotonic behaviour in the implied negotiation process. In particular, we propose two alternative mixing mechanisms based on linear weighted combination of tactics that solve this problem: the first using individual imitative negotiation threads, and the second combining individually proposed concessions of each tactic involved. We prove that both methods avoid these dynamic effects and ensure monotonic behaviour, the first for static and the second for dynamic weights. We further demonstrate by means of a comparative experimental evaluation that the proposed mechanisms provide utility gains for both parties in many negotiation settings.

In the next section, we briefly introduce the basic model as well as pure and mixed tactics for agent-based bargaining and discuss the emergence of non-monotonic behaviour in multi-tactic strategies. The two alternative mixing mechanisms are proposed in Section 3, while the results of an experimental evaluation are given in Section 4. Related work is presented in Section 5, and, finally we conclude in Section 6.

2 Mixing Negotiation Tactics

The negotiation model we consider in this paper has been introduced in [6] where two agents a and b exchange offers and counteroffers on a number of real-valued issues such as price or delivery time. The sequence of all offers exchanged between agents a and b until time t_n is termed a negotiation thread:

$$X^{t_n}_{a\leftrightarrow b} = (x^{t_i}_{a\leftrightarrow b})_{i=1,\ldots,n} = (x^{t_1}_{a\to b}, x^{t_2}_{b\to a}, x^{t_3}_{a\to b}, \ldots, x^{t_n}_{b\to a}). \tag{1}$$

The offer $x^{t_n}_{b\to a}$ at time t_n indicates the last element of the negotiation thread where t_i represent discrete time points where $t_{i+1} > t_i$ with $i = 1, 2, \ldots, n$ and $n \in \mathbb{N}$. The next counteroffer of agent a given the thread is then $x^{t_{n+1}}_{a\to b}$. We assume that each agent has a negotiation interval $D^a_j = [min^a_j, max^a_j]$ for each issue j where min^a_j and max^a_j are the initial and reservation values, respectively, if a is a buyer agent whereas the opposite holds for a seller, and D_j is the issue domain with $D^a_j, D^b_j \subseteq D_j$. Each agent has a utility function $U^a_j : D^a_j \to [0, 1]$ associated to each issue which assigns a score to the current value within its acceptable interval. We assume that utility functions are monotonically increasing or decreasing depending on the issue and the role of the agent. For example, for issue *price* the utility function is decreasing for a buyer and increasing for a seller. For each offer x of an agent a, the aggregated utility function $U^a(x) = \sum_j w_j^a U_j^a(x_j)$ determines the score for all issues j, where the weight w_{aj} represents the relative importance of issue j to agent a with $\sum_{1 \leq j \leq p} w_j^a = 1$. Agents may include discounts or negotiation costs, however, for simplicity we do not consider such a case here. The agents exchange offers alternately until one agent accepts or withdraws from the negotiation. An offer is accepted by agent a if the overall utility of agent b's last offer is equal or higher than a's next offer, such that $U^a(x_{b \to a}^{t_n}) \geq U^a(x_a^{t_{n+1}})$. An agent withdraws if it reaches its deadline t_{max}^a. Even though the utility structure may be more complex and of a different shape the functioning of the negotiation strategies described in this paper are best measured when using above linear utility function. Accordingly, it has been shown [8] that strategies well-suited for monotonic utility models do not cope well with non-monotonic utility spaces, so that we restrict to the former.

2.1 Negotiation Tactics and Strategies

A common method to generate offers is to use tactics or decision functions which utilize changes in the negotiation environment such as proposals from negotiation partners, or available resources such as time or the number of negotiating agents. In particular, a tactic τ_j^a is as a function mapping the mental state (about its environment) of an agent a to the issue domain D_j with $\tau_j^a : MS_a \to D_j$. Typical examples of such tactics are the time-, resource- or behaviour-dependent tactics proposed in [6]. A wide range of different negotiation strategies can be created by an agent through mixing of pure tactics. Faratin et al [6] introduces the concept of strategies where tactics are mixed based on a weight matrix

$$\Gamma_{a \to b}^{t_{n+1}} = \begin{pmatrix} \gamma_{11} & \gamma_{12} & \cdots & \gamma_{1m} \\ \gamma_{21} & \gamma_{22} & \cdots & \gamma_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{p1} & \gamma_{p2} & \cdots & \gamma_{pm} \end{pmatrix} \tag{2}$$

where $\gamma_{ji} \in [0,1]$ is the weight of tactic i for issue j. The weighted linear combination of tactics is then defined by the weighted sum of proposed offers of each tactic $x_{a \to b}^{t_{n+1}}[j] = \sum_{i=1}^m \gamma_{ji} \cdot \tau_{ji}$ where weights are normalized with $\sum_{i=1}^m \gamma_{ji} = 1$. The weighted counterproposal extends the negotiation thread by appending $x_{a \to b}^{t_{n+1}}$ whereby each row in the matrix represents a weighted linear combination of m tactics for one issue. Different types of negotiation behaviour can be obtained by weighting a given set of tactics in different ways. For example, the agent's mental state can change and generate a new weight matrix [4] depending on the current environment and belief of the agent. The above method of using pure or mixed tactics represent decision functions which an agent uses to make *concessions* such that $U^a(x_{a \to b}^{t_{n+1}}) < U^a(x_{a \to b}^{t_{n-1}})$. In multi-issue negotiations, an agent can also make trade-offs where the next offer has the same utility as its previous offer (both are on the same indifference curve) with $U^a(x^{t_{n+1}}_{a\to b}) = U^a(x^{t_{n-1}}_{a\to b})$. In this paper, we focus on the concession-making mechanisms as detailed above and refer to [11, 5] for well-discussed trade-off mechanisms.

2.2 Definition of Monotonic Tactics

To determine if a mixed strategy generates a monotonic offer sequence we distinguish between monotonic behaviour-dependent and -independent pure tactics:

Definition 1. *Given a negotiation between agents a and b, a* monotonic behaviour-independent tactic $\tau^a_j(t_k)$ *of agent a for issue j is a function generating offers at any times $t_k, t_i \in T_n$ such that $\tau^a_j(t_k) \geq \tau^a_j(t_i)$ if U^a is decreasing or $\tau^a_j(t_k) \leq \tau^a_j(t_i)$ if U^a is increasing under the condition that $k, i \in \{1, 2, \ldots, n\}$ and $k > i$.*

Definition 2. *Given a negotiation between agents a and b at time t_n, a* monotonic behaviour-dependent tactic $\tau^a_j(\widetilde{X}^{t_n}_{a\leftrightarrow b})$ *generates an offer using any sequence $\widetilde{X}^{t_n}_{a\leftrightarrow b} = (x^t_{a\leftrightarrow b})_{t\in\widetilde{T}_n}$ where $\widetilde{T}_n \neq \emptyset$ and $\widetilde{T}_n \subseteq T_n = \{t_1, \ldots, t_n\}$ under the conditions that there exists at least one offer $x^{t_i}_{b\to a} \in D^b_j$ of agent b in the sequence such that*

- *$\tau^a_j(\widetilde{X}^{t_n}_{a\leftrightarrow b}) \geq \tau^a_j(\widetilde{X}^{t_{n-2}}_{a\leftrightarrow b})$ if the sequence of opponent's offers $(x^t_{b\to a})_{t\in\widetilde{T}_n}$ and U^a is monotonic decreasing or*
- *$\tau^a_j(\widetilde{X}^{t_n}_{a\leftrightarrow b}) \leq \tau^a_j(\widetilde{X}^{t_{n-2}}_{b\leftrightarrow a})$ if the sequence of opponent's offers $(x^t_{b\to a})_{t\in\widetilde{T}_n}$ and U^a is monotonic increasing.*

Definition 1 typically represents tactics depending on a particular resource which state may change over time. Throughout the paper we denote this class of tactics with $\tau_{j,\mathrm{time}}$ for issue j. In the simplest case the tactic may depend on time or the number of negotiation rounds. For instance, the polynomial and exponential time-dependent decision functions proposed by Faratin et al [6] represent such tactics as they generate offers in a monotonically decreasing or increasing manner. In the case of a resource-dependent tactic, however, the resource may diminish and increase over time such that a monotonic sequence of offers is not guaranteed. An imitative tactic according to Definition 2 uses historical offers from the opponent to propose counteroffers by preserving a monotonic offer sequence as long as the opponent's sequence is monotonic as well. We refer to such imitative tactics as $\tau_{j,\mathrm{beh}}$. For instance, the imitative tit-for-tat tactics in [6] fulfil this definition. Once non-monotonicity is introduced by one partner it can in turn cause a non-monotonic offer sequence of the opponent depending on the degree of how much the concessions are copied. As a result, if monotonic tactics are mixed together, non-monotonic behaviour can emerge even when both agents apply monotonic tactics as we demonstrate in the next section.

2.3 Monotonicity of Mixed Strategies

It is often argued [10, 4] that the process of negotiation should be designed in a way that agents make concessions, or seek for joint improvements, i.e. in the form of trade-off proposals, in a negotiation. This implies monotonic behaviour: an agent makes proposals such that the aggregated utility of its next offer is equal (trade-off) or lower (concession) than the aggregated utility of its previous offer, such that $U^a(x^{t_{n+1}}_{a\to b}) \leq U^a(x^{t_{n+1}}_{a\to b})$. In the following, we say that agents have monotonic concession behaviour if they propose offers according to this principle. In single-issue negotiations agents typically have opposing utility structures such that a non-monotonic sequence of offers increases the risk of a withdrawal of the opponent. In that sense an agent a is acting rational in single-issue negotiations if it concedes towards the last offer of its opponent, thereby trying to increase the opponent's utility such that the sequence of its own utilities is monotonically decreasing. In multi-issue negotiations, however, an offer of an agent a with a higher aggregated utility for a as compared to its previous offer can not easily be detected by the opponent as the utility structures are unknown to each other. If, in turn, the opponent's utility for a's last offer is lower as a's previous offers, the opponent may assume that a made a trade-off proposal and can therefore not detect the cause of such non-monotonic behaviour. It is also argued that agents behaving non-monotonic under time-constraints can be advantageous and the question whether automated negotiation should be designed in a way that non-monotonic behaviour is ensured is widely discussed in the research literature [12]. However, non-monotonicity in the sequence of proposed offers and their respective aggregated utilities of an agent can easily emerge at any time as a result of the *dynamic effects of an agent system in which the agents use mixed strategies*. Intuitively, non-monotonic behaviour can occur when an agent changes its strategy, e.g. the mixing weights, during the encounter. However, automatic non-monotonic behaviour can also be observed when imitative and non-imitative tactics are mixed by a linear weighted combination without the agent changing its strategy, i.e. even in the case of *static strategy settings and mixing weights*. A simple example shall demonstrate this:

Example 1: Assume a negotiation between two agents a and b at time t_n where agent a applies a mixed strategy with static weight γ using one time-dependent tactic $\tau^a_{\text{time}}(t_{n+1})$ and one imitative tactic simply copying the concession of the partner (basic absolute tit-for-tat): $\tau^a_{\text{beh}}(x^{t_{n-2}}_{b\to a}, x^{t_{n-1}}_{a\to b}, x^{t_n}_{b\to a}) = x^{t_{n-2}}_{b\to a} - x^{t_n}_{b\to a} + x^{t_{n-1}}_{a\to b}$ such that agent a's next offer is $x^{t_{n+1}}_{a\to b} = \gamma \cdot \tau^a_{\text{time}}(t_{n+1}) + (1-\gamma) \cdot \tau^a_{\text{beh}}(x^{t_{n-2}}_{b\to a}, x^{t_{n-1}}_{a\to b}, x^{t_n}_{b\to a})$. Given the thread $(\ldots, x^{t_{n-2}}_{b\to a}, x^{t_{n-1}}_{a\to b}, x^{t_n}_{b\to a}) = (\ldots, 30, 10, 20)$, agent a's next time-dependent proposal $\tau_{\text{time}}(t_{n+1}) = 11$ and the mixing weight $\gamma = 0.5$, the next counteroffer is $x^{t_{n+1}}_{a\to b} = 0.5 \cdot 11 + 0.5 \cdot 20 = 15.5$. Now assume, agent b replies with a comparatively small concession $x^{t_{n+2}}_{b\to a} = 19$ and agent a's next time-dependent proposal is $\tau_{\text{time}}(t_{n+3}) = 12$, then agent a's response is lower than its previous offer and thus non-monotonic with $x^{t_{n+3}}_{a\to b} = 0.5 \cdot 12 + 0.5 \cdot 16.5 = 14.25$. In this example the non-monotonic behaviour emerges in static mixed strategies with imitative and non-imitative tactics even though the sequence of opponents' offers is monotonic and all involved tactics are monotonic as well. This is because the imitative tactic is *not independent* from the other tactics in the mix since it uses the last offer of the current negotiation thread which is different from the individually proposed one. In addition, if both

	Buyer Agent	Seller Agent
Issue 1	$min_1 = 10, max_1 = 25, w_1 = 0.7$	$min_1 = 15,\ max_1 = 30,\ w_1 = 0.5$
	Mixed Strategy ($\gamma_1 = 0.3$): $\tau_{1,\text{time}}$: polynomial, $\beta = 5$ $\tau_{1,\text{beh}}$: absolute tft, $\delta = 1$	Mixed Strategy ($\gamma_1 \in \{0.1, 0.12\}$): $\tau_{1,\text{time}}$: polynomial, $\beta = 1$ $\tau_{1,\text{beh}}$: absolute tft, $\delta = 1$
Issue 2	$min_2 = 20, max_2 = 40, w_2 = 0.3$	$min_2 = 30, max_2 = 50, w_2 = 0.5$
	Mixed Strategy ($\gamma_1 = 0.4$): $\tau_{2,\text{time}}$: polynomial, $\beta = 2$ $\tau_{2,\text{beh}}$: absolute tft, $\delta = 1$, $R = 0$	Mixed Strategy ($\gamma_1 = 0.2$): $\tau_{2,\text{time}}$: polynomial, $\beta = 0.3$ $\tau_{2,\text{beh}}$: relative tft, $\delta = 1$

(a) Example 2: Negotiation settings

Fig. 1: Example 2 settings and offer and utility curves using the Traditional (straight) or the negotiation thread-based mixing (dotted)

agents have imitative tactics in their mix a non-monotonic sequence of offers is copied to some degree and may thus reproduce the non-monotonicity in the sequence of opponent's offers and vice versa. If agents have opposing utility functions, such that a non-monotonic utility sequence of one agent then also causes a non-monotonic utility sequence of the partner's offers. This can result in a delay of agreements, varying outcomes as compared to mixing methods with monotonic offer sequences, and a high sensitivity in terms of the strategy parameters making it difficult for an agent to apply such strategies in real world scenarios. In such cases, the described dynamics of the system result in a partial loss of the agent's control over its strategy since small changes of parameters may change the offer curves to a large degree.

Example 2: The settings of the second example with multiple issues are shown in the table in figure 1(a), and figure 1(b-e) shows the non-monotonic offer and utility curves of both agents (straight) and how it is reproduced if the traditional mixing method is used (as a comparison the dotted curves show the thread-based mixing from next section). As a result, the utility curves of both agents are non-monotonic with a delayed agreement. In the case of the agents having different deadlines this behaviour might also result in no agreement. The example further demonstrates that in such scenarios it is difficult to find suitable strategy parameters since the outcome utility may change significantly for slightly different settings as shown in figure 1 for different γ settings (the difference in utility is 0.1 for both agents when the seller changes γ_1 from 0.1 to 0.12). The agent can avoid non-monotonic behaviour by applying a simple min- or max-*constraint* (if it is a buyer or seller, respectively) to the next offer proposal to ensure that the agent's own utility does not increase with the new offer. However, the offer curve then rapidly changes to linear and the agent may propose the same offer over a long time period which may also increase the risk of the opponent's withdrawal. In the next section, we therefore present two alternative mixing mechanisms producing monotonic offer and utility sequences thereby avoiding the dynamic effects described above.

3 Monotonic Mixing Mechanisms

3.1 Mixing based on Negotiation Threads

To calculate the imitative tactics in mixed strategies using the traditional mixing method the last offer in the current thread is used. The imitative part of the strategy does therefore not represent an individually applied behaviour-dependent tactic. Another intuitive method is to use the last offers of each imitative tactic involved in the mix. This can be interpreted as using individual negotiation threads $X^{t_n}_{a\leftrightarrow b}[j,k]$ where k denotes the k'th behaviour-dependent tactic $\tau_{jk}(\tilde{X}^{t_n}_{a\leftrightarrow b}[j,k])$ for issue j. As a result, offers from all imitative functions have to be stored in order to be used in the calculation of next proposals. Formally, the linear weighted combination of tactics can now be written as:

$$x^{t_{n+1}}_{a\to b}[j] = \sum_{i=1}^{l} \gamma_{ji} \cdot \tau_{ji}(t_{n+1}) + \sum_{k=l+1}^{m} \gamma_{jk} \cdot \tau_{jk}(\tilde{X}^{t_n}_{a\leftrightarrow b}[j,k]) \qquad (3)$$

where m and l denote the total number and the number of behaviour-independent tactics, respectively. Unlike the traditional mixing method in Section 2.1 this method can be regarded as a true linear weighted combination of tactics in which all involved tactics are independent from each other.

Theorem 1. *The mixing mechanism using individual negotiation threads for each behaviour-dependent tactic results in a monotonic offer curve if monotonic tactics from definitions 1 and 2 are used with static weights for all tactics.*

Proof Let $X^{t_n}_{a\leftrightarrow b}$ be the negotiation thread at time t_n with $x^{t_n}_{b\to a}$ being the last offer and $x^{t_{n+1}}_{a\to b}$ being the next counteroffer of agent a then according to Definition 1 and 2 $\gamma_k \cdot \tau_k(\tilde{X}^{t_n}_{a\leftrightarrow b}[k]) \geq \gamma_k \cdot \tau_k(\tilde{X}^{t_{n-2}}_{a\leftrightarrow b}[k])$ and $\gamma_i \cdot \tau_i(t_{n+1}) \geq \gamma_i \cdot \tau_i(t_{n-1})$ if U^a is decreasing and all $\gamma_i, \gamma_k \geq 0$. Since each term of the sum in (3) at t_n is larger than the corresponding term of the sum at t_{n-2} it follows that $x^{t_n+1}_{a\to b} \geq x^{t_n-1}_{a\to b}$. The same line of reasoning can be followed for increasing utility functions U^a. $\square$

Figure 1 shows the monotonic offers curves and the resulting monotonic utility sequence when both agents use this mixing mechanism for Example 2 (dotted). The outcome is changed in favour of the seller and agreement is reached earlier. Agents using this mechanism do not expose the dynamic effects as described in section 2.3. However, the mechanism does not force the agent to propose offers in a monotonic manner. For instance, if the opponent still proposes offers in a non-monotonic sequence, an imitative tactic in the mix may still copy it to some degree. The agent can choose to strictly ensure monotonicity by applying a constraint C to the imitative tactic: $C(\tau_{jk}(\tilde{X}^{t_n}_{a\leftrightarrow b}[j,k]), x^{t_{n-1}}_{a\to b}[j,k])$ where $C \equiv \min$ if U^a decreasing and $C \equiv \max$ if U^a increasing. The individual imitative thread used by this method does not represent the actual negotiation thread. This seems counter-intuitive as the offer curve and the outcome of the individually applied imitative tactics might indeed be different from the mixed strategy.

3.2 Mixing based on Single Concessions

This mixing type calculates individual next concessions for each tactic to mix behaviour-dependent and -independent tactics as defined in Section 2.2:

$$\begin{aligned} x_{a\to b}^{t_{n+1}}[j] = x_{a\to b}^{t_{n-1}}[j] &+ \sum_{i=1}^{l} \gamma_{ji} \cdot (\tau_{ji}(t_{n+1}) - \tau_{ji}(t_{n-1})) + \dots \\ &+ \sum_{k=l+1}^{m} \gamma_{jk} \cdot (\tau_{jk}(\tilde{X}_{a\leftrightarrow b}^{t_n}[j]) - x_{a\to b}^{t_{n-1}}[j]) \end{aligned} \tag{4}$$

with m and l denoting the total number and the number of behaviour-independent tactics respectively. In order to use concessions at least two offers of the opponent are necessary. Any of the former mechanisms can be used for initial offers as they propose the same offers in the first round. Concessions for behaviour-independent tactics are, since they do not depend on opponent' offers, the difference $\tau_{ji}(t_{n+1}) - \tau_{ji}(t_{n-1})$ between the calculated offer at t_{n+1} and the previous individual offer at t_{n-1}. For the imitative tactic we can not follow the same line of reasoning because, as described in the previous section, the last offer of the individually applied imitative tactic is unknown. However, suppose that the agent changed its strategy to the pure imitative tactic at time t_{n+1} the last offer is still be $x_{a\to b}^{t_{n-1}}$ and hence the next offer is given by $\tau_{jk}(\tilde{X}_{a\leftrightarrow b}^{t_n}[j])$. We can hence calculate the behaviour-dependent concession by the difference between the proposed imitative offer and the last offer of the agent. This approach provides monotonic offer curves similar to the negotiation thread-based mixing and also avoids non-monotonic aggregated utilities over time. The major advantage, however, is that a monotonic sequence of utilities is also never introduced if the agent changes weights for tactics dynamically.

Theorem 2. *The mixing mechanism based on single concessions of pure tactics results in a monotonic offer curve (and therefore preserves a monotonic sequence of utilities) if monotonic tactics from Definitions 1 and 2 are used.*

Proof Let $X_{a\leftrightarrow b}^{t_n}$ be the negotiation thread at time t_n with $x_{b\to a}^{t_n}$ being the last offer and $x_{a\to b}^{t_{n+1}}$ being the next counteroffer of agent a then according to Definition 1 the behaviour-independent concession $\tau_{ji}^{a}(t_{n+1}) - \tau_{ji}^{a}(t_{n-1})$ is always greater zero if U_a is increasing. The offer proposed by the pure behaviour-dependent tactics $\tau_{jk}^{a}(\tilde{X}_{a\leftrightarrow b}^{t_n}[j])$ for issue j is greater than the previous offer $x_{a\to b}^{t_{n-1}}[j]$ if monotonic tactics from Definition 2 are used and the opponent never introduces non-monotonicity. The behaviour-dependent concession $\tau_{jk}^{a}(\tilde{X}_{a\leftrightarrow b}^{t_n}[j]) - x_{a\to b}^{t_{n-1}}[j]$ is therefore always greater zero. For all weights $\gamma_i, \gamma_k \geq 0$ follows that each term of the sum in Eq. (4) is greater zero and hence $x_{a\to b}^{t_n+1} \geq x_{a\to b}^{t_n-1}$. The same line of reasoning can be followed for an increasing scoring function U_a. $\square$

Similar to the previous method the agent can strictly avoid imitating a non-monotonic sequence of opponent's offers by applying a constraint C to each imitative concession in (4) written as $C(\tau_{jk}(\tilde{X}_{a\leftrightarrow b}^{t_n}[j]) - x_{a\to b}^{t_{n-1}}[j], 0)$ where $C \equiv \min$ if

s / b	CaS	CaL	BaS	BaL
CaS	0/0	7/0.03	30/0.45	100/0.19
CaL	9/0.06	23/0.03	88/0.45	100/0.18
BaS	89/0.36	89/0.43	0/0	0/0
BaL	88/0.06	100/0.16	0/0	0/0

s / b	CrS	CrL	BrS	BrL
CrS	4/0.09	4/0.09	70/0.47	100/0.14
CrL	16/0.01	22/0.04	99/0.24	100/0.09
BrS	74/0.46	48/0.55	0/0	0/0
BrL	99/0.13	100/0.26	0/0	0/0

(a) Non-monotonicity in single-issue negotiation (rate in %/max. variation in utility)

Fig. 2: Buyer's (bottom) and seller's (top)

U^a decreasing or $C \equiv \max$ if U^a increasing. In contrast to the thread-based mixing this mechanism needs no separate negotiation threads and produces monotonic offer curves even for dynamically changing weights.

4 Evaluation

We present the results of a comparative evaluation of the mixing mechanisms with respect to their non-monotonic behaviour and its respective effects in different bilateral single- and multi-issue negotiation settings. As the number of possible mixes of tactics is infinite, we restrict the evaluation to a mix of two tactics from [4], one behaviour- and one time-dependent, for each agent with static weights throughout the encounter and the following settings:

- *Time-dependent (poly.)*: (C)onceder: $\beta \in \{3, 7\}$; (B)oulware: $\beta \in \{0.1, 0.3\}$
- *Behaviour-dependent*: (a)bsolute tft: $\delta = 1, R(M) = 0$; (r)elative tft: $\delta = 1$
- *Weights*: (S)mall: $\gamma \in \{0.1, 0.3\}$; (L)arge: $\gamma \in \{0.7, 0.9\}$

where, for example, 'CaS' denotes the strategy group containing conceder time-dependent and absolute tit-for-tat tactics mixed by small weights. Before considering a multi-issue scenario we are interested in when and to what degree non-monotonic behaviour emerges in static mixed strategies. For that reason, we consider first a simple single-issue scenario where two agents, a buyer (b) and a seller (s), negotiate about a issue 1 from example 1 with partially overlapping intervals and equal deadlines $t^s_{max} = t^b_{max} = 20$. The tables in figure 2 illustrates the rate (%) of negotiations with non-monotonic offer curves in the case of both agents applying the traditional linear weighted combination of tactics for a particular strategy group. Numbers below the rate are the maximum variation in terms of non-monotonicity occurred in utility for either the seller or buyer agent. As we can see, the dynamic occurrence of non-monotonic behaviour in static strategy settings is not a negligible side-effect. In such scenarios the variation is higher in the case of oppositional applied time-dependent tactics in the mix, such as conceder against boulware. In the second multi-issue scenario using issue 1 and 2 form example 2, we compare the different mixing mechanisms (cf. 2.1 to 3.2) applying the same strategy groups. The performance is measured using the aggregated linear utility (cf. 2.1). Due to the large number of possible strategy assignments we choose three scenarios, where the buyer applies a more cooperative (CaS/CaL), a more competitive strategies (BaS/Bal) for both issues, or a mixture of both (BaS/CaL), whereas the seller is cooperative for issue 1 (CaS) and applies different combinations for issue 2. Figures 2(b) and (c) show the buyer's and seller's aggregated utilities for the different multi-issue scenarios in which both agents use the same mixing mechanism. In each diagram a group of bars represent one strategy scenario, where the different bars depict the mixing mechanism from left (light) to right (dark): (1) Linear weighted combination (cf. 2.1), (2) Constrained linear weighted combination (cf. 2.3), (3) thread-based mechanism (cf. 3.1) and (4) concession-based mechanism (cf. 3.2). Both agents gain higher utilities in many strategy scenarios using the proposed mixing mechanisms, however, most significantly when the buyer applies the competitive (BaS/Bal) mixed strategy. The monotonic mixing may also shift utility from an agent that gained advantage through its non-monotonic utility sequence to the other agent with monotonic behaviour (see buyer strategy CaS/Cal). Both monotonic mechanisms perform similar since all pure tactics are treated independently in both methods. In general, we further observed the effect that the difference between traditional and the monotonic mixing mechanisms increases when the time-dependent tactics and the mixing weights are oppositional, i.e. one agent uses conceder with small mixing weights while the other agent employs boulware tactics with large mixing weights and vice versa. For instance, if both agents use similar strategies (both cooperative or both competitive) utilities are similar for all mixing mechanisms. These observations correspond to the results from the first experiment (cf. figure 2(a)) where oppositional concession behaviours exposed the highest rate of non-monotonicity.

5 Related Work

A large number of negotiation scenarios have been studied to provide effective negotiation mechanisms and strategies, while, however, many focus on single families of tactics [6], trade-offs mechanisms [4] or meta-strategies [11], but do not consider the dynamic effects in the negotiation process. For example, Fatima et al [7] investigate scenarios of single- and multi-issue negotiation where agents have only partial information about each other trying to find optimal strategies that most exploit the opponent. The work focus on the effect of time, information states and discounting factors on the outcome while comparisons are made to equilibrium solutions but are limited to time-dependent tactics. Evaluation results for pure, static and dynamic mixed strategies are presented in [4] with focus on the influence of long and short term deadlines, and initial offers. Matos et al [9] propose the application of genetic algorithms to determine most successful mixed strategies that evolve depending on the environment and strategy of the opponent. Both approaches demonstrate that mixed strategies perform better than pure tactics in terms of gained utility and negotiation cycles, but do not investigate the mechanism of their mixing with respect to the emergence of non-monotonic behaviour. Cardoso et al [3], and Brzostowski et al [2] [1] consider the mixing of different tactic families to evaluate adaptive strategies based on reinforcement learning, respectively, heuristic predictive methods or regression analysis with respect to their negotiation outcomes only. Sierra and Ros [11] propose to let an agent make concessions through single or mixed tactics whenever a deadlock occurs, i.e. the opponent's last offer does not improve the utility of the offer two steps before, otherwise a trade-off tactic is used. However, utilities of offers may also decrease when pure tactics are combined as shown in this paper. Our work is different in that it focuses on the analysis of the mixing mechanism itself, and proposes new mechanisms that, in contrast to the commonly used mixing of tactics, avoid the dynamic emergence of non-monotonic utility sequences during the process of negotiation, thereby also avoiding the drawbacks described in this paper.

6 Conclusions

We provided an investigation of (non-)monotonic behaviour of multi-tactic strategies created by different mechanisms for mixing pure tactics in bilateral single- and multi-issue negotiations. The traditional mixing based on linear weighted combination can undesirably expose non-monotonic utilities over time, even in cases of individual cooperative behaviour and static strategy settings of both agents, if behaviour-dependent and -independent tactics are used. As alternative, we proposed two mixing mechanisms that solve this problem by provably producing monotonic concession behaviour for static and dynamic weights: the first using imitative negotiation threads and the second single concessions for each involved tactic. A comparative evaluation showed provided evidence that both mechanisms yield higher utilities for both agents in many multi-issue negotiation scenarios as compared to traditional mixing when both agents use the same mixing mechanism.

References

1. J. Brzostowski and R. Kowalczyk. Adaptive negotiation with on-line prediction of opponent behaviour in agent-based negotiations. In *Proceedings of the IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology*, Washington, DC, USA, 2006. IEEE Computer Society.
2. J. Brzostowski and R. Kowalczyk. Predicting partner's behaviour in agent negotiation. In *Proceedings of the Fifth Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, New York, NY, USA, 2006. ACM.
3. H. L. Cardoso and E. Oliveira. Using and evaluating adaptive agents for electronic commerce negotiation. London, UK, 2000. Springer-Verlag.
4. P. Faratin. *Automated Service Negotiation Between Autonomous Computational Agents*. PhD thesis, University of London, 2000.
5. P. Faratin, M. Klein, H. Sayama, and Y. Bar-Yam. Simple negotiating agents in complex games: Emergent equilibria and dominance of strategies. In J.-J. Meyer and M. Tambe, editors, *Intelligent Agents VIII*, volume 2333 of *LNAI*, 2002.
6. P. Faratin, C. Sierra, and N. R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.

7. S. S. Fatima, M. Wooldridge, and N. R. Jennings. Bargaining with incomplete information. *Annals of Mathematics and Artificial Intelligence*, 44(3), 2005.
8. Ito, T., Hattori, H., Klein, M.: Multi-issue negotiation protocol for agents: exploring nonlinear utility spaces. In: 20th international joint conference on Artifical intelligence, pp. Morgan Kaufmann Publishers Inc, San Francisco, CA, USA (2007)
9. N. Matos, C. Sierra, and N. R. Jennings. Determining successful negotiation strategies: An evolutionary approach. In *3rd Int. Conf. on Multi-Agent Systems*, Washington, DC, USA, 1998. IEEE Computer Society.
10. H. Raiffa. *The art and science of negotiation*. Harvard University Press Cambridge, Mass. (1982)
11. R. Ros and C. Sierra. A negotiation meta strategy combining trade-off and concession moves. *Autonomous agent and multiagent systems*, 12, 2006.
12. P. Winoto, G. I. McCalla, and J. Vassileva. Non-monotonic-offers bargaining protocol. *Autonomous Agents and Multi-Agent Systems*, 11(1), June 2005.