



# A Negotiator's Backup Plan: Optimal Concessions with a Reservation Value

## Extended Abstract

Tamara C.P. Florijn  
Centrum Wiskunde & Informatica  
Amsterdam, The Netherlands  
Utrecht University  
Utrecht, The Netherlands  
Tamara.Florijn@cw.nl

Pinar Yolum  
Utrecht University  
Utrecht, The Netherlands  
P.Yolum@uu.nl

Tim Baarslag  
Centrum Wiskunde & Informatica  
Amsterdam, The Netherlands  
Utrecht University  
Utrecht, The Netherlands  
T.Baarslag@uu.nl

## ABSTRACT

Automated negotiation is a well-known mechanism for autonomous agents to reach agreements. To realize beneficial agreements quickly, it is key to employ a good bidding strategy. When a negotiating agent has a good back-up plan, i.e., a high reservation value, failing to reach an agreement is not necessarily disadvantageous. Thus, the agent can adopt a risk-seeking strategy, aiming for outcomes with a higher utilities.

Accordingly, this paper develops an optimal bidding strategy called MIA-RVelow for bilateral negotiations with private reservation values. The proposed greedy algorithm finds the optimal bid sequence given the agent's beliefs about the opponent in  $O(n^2D)$  time, with  $D$  the maximum number of rounds and  $n$  the number of outcomes. The results obtained here can pave the way to realizing effective concurrent negotiations, given that concurrent negotiations can serve as a (probabilistic) backup plan.

## KEYWORDS

Reservation value; Bidding strategy; Concessions; Negotiation; One-to-one negotiation; Optimal offers; Automated negotiation

### ACM Reference Format:

Tamara C.P. Florijn, Pinar Yolum, and Tim Baarslag. 2024. A Negotiator's Backup Plan: Optimal Concessions with a Reservation Value: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

## 1 INTRODUCTION

Negotiation is often described as making concessions toward a mutually agreeable outcome [15]. Deciding on the right concessions is a challenge: being too stubborn can lead to non-agreements, while conceding too fast can result in bad agreements. An important factor in choosing the right negotiation strategy is that of a *backup plan*, which serves as a possible outcome that a negotiator can safely fall back on in case the negotiation ends in disagreement. Accordingly, this paper studies how the negotiator's backup plan can

be used to design an optimal bidding strategy and how it influences a negotiator's concession behavior.

We consider the utility of this backup plan conceptually as a *reservation value* [1, 2, 8, 10, 11, 19], interpreted as the value an agent receives if the negotiation fails [4, 13, 21]. Various works consider how a reservation value can be used in a negotiation. A reservation value often functions as a lower bound of a bidder's utility target, signifying what can be obtained in other concurrent negotiation threads [9, 16] or as a minimum acceptance threshold [12]. Alternatively, reservation values are used to define heuristics in a multi-threaded negotiation context [7, 20, 22]. However, these works focus mostly on heuristic strategies and do not provide a theoretical proof of optimality.

In this work, we are interested in finding the optimal bidding curve given a private reservation value. The most closely related work is the Greedy Concession Algorithm (GCA) [3], which finds the theoretically optimal solution for time-sensitive domains with a static acceptance model, but without a reservation value. Mohammad [14] has proven GCA can handle repeating offers. They have also introduced a faster GCA (QGCA) and extend to negotiations with a nonzero reservation value, but without direct proof for this setting.

Our main contribution is a bidding strategy called a Marginal Improvement Algorithm for Reservation Values (MIA-RVelow), which maximizes the marginal improvement of expected utility in a greedy way to find the optimal bid sequence given the reservation value. The algorithm finds an optimal sequence for static acceptance models with a time complexity of  $O(n^2D)$ .

## 2 PROBLEM SETTING

We focus on the bidding strategy (not the acceptance strategy or opponent model [5, 6, 17]) of a general agent for bilateral negotiation, where agents  $A$  and  $B$  take turns making bids, in accordance with the *alternating offers protocol* [18]. If no agreement is reached before the maximum number of rounds, i.e. a deadline  $D$ , then the agent receives the reservation value  $rv$ , which is considered private information and hence remains unknown to the opponent.

All the possible agreements together are called the *negotiation domain*  $\Omega = \{\omega_1, \omega_2, \dots, \omega_n\}$ . Both agents have preferences over the domain  $\Omega$ ; we adopt the perspective of agent  $B$  and examine their bidding strategy. Agent  $B$  associates with  $\omega_i \in \Omega$  a utility between 0 and 1, which we denote by  $u(\omega_i)$ .  $B$  has access to their own utility function, but  $A$ 's utility function is unknown to  $B$ . Therefore,  $B$



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).*

uses an opponent model, with  $a(\omega_i)$  the acceptance probability of bid  $\omega_i$ .

EXAMPLE 1. Imagine Bo ( $B$ ) would like to go for dinner together with their friend Ace ( $A$ ), and they both have to agree on a restaurant before dinner time. To come to the best dinner plan, Bo has to decide what bids to make and in what order. Suppose there are three options to consider: Italian, Sushi, and Fast food. In Table 1, Bo’s utility for each bid is shown, along with the corresponding (estimated) probability that Ace accepts the bid. Suppose Bo only has time for one dinner proposal out of the three options. Using Bo’s belief regarding Ace’s acceptance probabilities, Bo’s first (and only) bid should maximize the expected utility, so the best choice is proposing the Fast food place, which yields an expected utility of 0.27.

**Table 1: Possible bids and their expected utilities for one bid.**

Bid	$B$ ’s utility	$A$ ’s acceptance probability	Expected utility
Italian	0.5	0.4	0.20
Sushi	0.9	0.2	0.18
Fast food	0.3	0.9	0.27

The example above illustrates one bid only; if we consider multiple bids strategically, the agent (Bo) can plan multiple bids ahead and decide on the whole bid sequence  $\pi = (x_1, x_2, \dots, x_k)$ , with  $x_1, \dots, x_k \in \Omega$  before the negotiation start given a static acceptance model. Note that early termination is irrational without bidding costs.  $B$ ’s expected utility of making bids,  $EU_{\mathbf{rv}}$  of  $\pi$ , is the expected utility of the first bid, or, if it is rejected, the expected utility of the second bid, and so on, ending with the reservation value if no agreement is reached:

$$EU_{\mathbf{rv}}(\pi) = \sum_{i=1}^k u_i \cdot a_i \prod_{j=1}^{i-1} (1 - a_j) + \mathbf{rv} \cdot \prod_{j=1}^k (1 - a_j).$$

The overall aim is to select the optimal sequence of  $D$  bids  $\pi^* \in \Omega^D$  such that the expected utility  $EU_{\mathbf{rv}}$  is maximized:

$$\pi^* = \operatorname{argmax}_{\pi \in \Omega^D} EU_{\mathbf{rv}}(\pi). \quad (1)$$

A naive strategy to find this optimal sequence  $\pi^*$  would be to evaluate all sequences from  $\Omega^D$  and pick the sequence with the highest expected utility. As the agreement space grows larger, this quickly becomes infeasible. Therefore, we develop a novel strategy called MIA-RVeloous that finds  $\pi^*$  in  $O(n^2D)$  time.

### 3 CONCESSION BEHAVIOR

The MIA-RVeloous algorithm, a variation of [3], is presented in Algorithm 1. Intuitively, MIA-RVeloous works by first adding the reservation value as a special bid, and then greedily adding the bid with the best marginal improvement to utility in expectation.

THEOREM 3.1. *MIA-RVeloous selects the optimal sequence of bids  $\pi^* = \operatorname{argmax}_{\pi \in \Omega^D} EU_{\mathbf{rv}}(\pi)$  given deadline  $D$ .*

#### Algorithm 1 Marginal Improvement Algorithm with $\mathbf{rv}$ (MIA-RVeloous)

```

1: for  $k \in \{1, \dots, D\}$  do
2:    $x_k \leftarrow \operatorname{argmax}_{\omega \in \Omega \setminus \{x_1, \dots, x_{k-1}\}} EU_{\mathbf{rv}}(\operatorname{sort}(\{x_1, \dots, x_{k-1}, \omega\}))$ 
3:  $\pi^* \leftarrow \operatorname{sort}(\{x_1, \dots, x_k\})$  //sorted by decreasing utility

```

We can prove the optimality as formulated in Theorem 3.1. To illustrate further, we extend our dinner example from Section 2.

EXAMPLE 2. Bo and Ace negotiate about where to go for dinner. But this time, if the negotiation fails, Bo would eat a meal from the freezer, which yields a reservation value of 0.2. Given the three options, what bids should Bo optimally propose? As shown before, the Fast food place should be chosen first without a backup plan. However, with a reservation value, it may not be immediately clear. If a bid’s utility is lower than the reservation value, we can safely ignore the bid. Otherwise, we can use MIA-RVeloous to determine what bid should be chosen with one available bid and a reservation value of 0.2, and apply Algorithm 1, as shown below.

**Table 2: Calculations of the expected utility of Example 1.**

First bid	If rejected:	Expected utility
Italian	$\mathbf{rv}$	$0.5 \cdot 0.4 + (1 - 0.4) \cdot 0.2 = 0.32.$
Sushi	$\mathbf{rv}$	$0.9 \cdot 0.2 + (1 - 0.2) \cdot 0.2 = \boxed{0.34}.$
Fast food	$\mathbf{rv}$	$0.3 \cdot 0.9 + (1 - 0.1) \cdot 0.2 = 0.29.$

Since the Sushi bid yields the highest expected utility, the first bid should be Sushi, which is more risky than Fast food meaning that it has a higher utility and a lower acceptance probability. This is an example of how the presence of a reservation value stimulates the agent to seek the riskier option.

### 4 DISCUSSION

We can show that our proposed bidding strategy, MIA-RVeloous, finds the optimal bid sequence for bilateral negotiations with a reservation value with a time complexity of  $O(n^2D)$ .

Our work leads to further research directions. In this paper, we consider a given reservation value, which we can extend to a scenario where the agent receives the reservation value with a certain probability. Probabilistic reservation values can serve their function in concurrent negotiations, as previously proposed [7]: an agent can fall back on ongoing concurrent negotiations after a negotiation thread failed, which results in probabilistic backup plans. Therefore, an important future direction is to extend our work on the optimal concession curve with a reservation value to one-to-many and many-to-many negotiations. The results obtained here can pave the way to realizing effective concurrent negotiations.

### ACKNOWLEDGMENTS

This publication is part of the Vidi project COMBINE (VIdi.203.044) which is (partly) financed by the Dutch Research Council (NWO).

## REFERENCES

- [1] Reyhan Aydoğan, Katsuhide Fujita, Tim Baarslag, Catholijn M. Jonker, and Takayuki Ito. 2020. ANAC 2018: Repeated Multilateral Negotiation League. In *Advances in Artificial Intelligence (Advances in Intelligent Systems and Computing, Vol. 1128)*, Yukio Ohsawa, Katsutoshi Yada, Takayuki Ito, Yasufumi Takama, Eri Sato-Shimokawara, Akinori Abe, Junichiro Mori, and Naohiro Matsumura (Eds.). Springer, Cham, 77–89. [https://doi.org/10.1007/978-3-030-39878-1\\_8](https://doi.org/10.1007/978-3-030-39878-1_8)
- [2] Reyhan Aydoğan, Katsuhide Fujita, Tim Baarslag, Catholijn M. Jonker, and Takayuki Ito. 2021. ANAC 2017: Repeated Multilateral Negotiation League. In *Advances in Automated Negotiations (Studies in Computational Intelligence, Vol. 905)*, Takayuki Ito, Minjie Zhang, and Reyhan Aydoğan (Eds.). Springer, Singapore, 101–115. [https://doi.org/10.1007/978-981-15-5869-6\\_7](https://doi.org/10.1007/978-981-15-5869-6_7)
- [3] Tim Baarslag, Enrico H. Gerding, Reyhan Aydoğan, and M. C. Schraefel. 2015. Optimal Negotiation Decision Functions in Time-Sensitive Domains. In *2015 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, Vol. 2. IEEE, Singapore, 190–197. <https://doi.org/10.1109/WI-IAT.2015.161>
- [4] Tim Baarslag, Rafik Hadfi, Koen Hindriks, Takayuki Ito, and Catholijn M. Jonker. 2016. Optimal Non-adaptive Concession Strategies with Incomplete Information. In *Recent Advances in Agent-based Complex Automated Negotiation*, Naoki Fukuta, Takayuki Ito, Minjie Zhang, Katsuhide Fujita, and Valentin Robu (Eds.). Studies in Computational Intelligence, Vol. 638. Springer, Cham, 39–54. [https://doi.org/10.1007/978-3-319-30307-9\\_3](https://doi.org/10.1007/978-3-319-30307-9_3) Series Title: Studies in Computational Intelligence.
- [5] Tim Baarslag, Mark J. C. Hendriks, Koen V. Hindriks, and Catholijn M. Jonker. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems* 30, 5 (Sept. 2016), 849–898. <https://doi.org/10.1007/s10458-015-9309-1>
- [6] Tim Baarslag, Koen Hindriks, Mark Hendriks, Alexander Dirkzwager, and Catholijn M. Jonker. 2014. Decoupling Negotiating Agents to Explore the Space of Negotiation Strategies. In *Novel Insights in Agent-based Complex Automated Negotiation*, Ivan Marsa-Maestre, Miguel A. Lopez-Carmona, Takayuki Ito, Minjie Zhang, Quan Bai, and Katsuhide Fujita (Eds.). Vol. 535. Springer Japan, Tokyo, 61–83. [https://doi.org/10.1007/978-4-431-54758-7\\_4](https://doi.org/10.1007/978-4-431-54758-7_4) Series Title: Studies in Computational Intelligence.
- [7] Cuihong Li, K. Sycara, and J. Giampapa. 2005. Dynamic Outside Options in Alternating-Offers Negotiations. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*. IEEE, Big Island, HI, USA, 35b–35b. <https://doi.org/10.1109/HICSS.2005.212>
- [8] Itza Tlaloc Quetzalcoatl Curiel-Cabral, Sonia Di Giannatale, and Giselle Labrador-Badia. 2024. Risk Aversion, Reservation Utility and Bargaining Power: An Evolutionary Algorithm Approximation of Incentive Contracts. *Computational Economics* 63, 2 (Feb. 2024), 477–511. <https://doi.org/10.1007/s10614-022-10349-0>
- [9] Peyman Faratin, Carles Sierra, and Nick R. Jennings. 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems* 24, 3–4 (Sept. 1998), 159–182. [https://doi.org/10.1016/S0921-8890\(98\)00029-3](https://doi.org/10.1016/S0921-8890(98)00029-3)
- [10] Deepak Malhotra and Max Bazerman. 2007. *Negotiation Genius: How to Overcome Obstacles and Achieve Brilliant Results at the Bargaining Table and Beyond*. Bantam Books, New York.
- [11] Khalid Mansour and Ryszard Kowalczyk. 2015. An Approach to One-to-Many Concurrent Negotiation. *Group Decision and Negotiation* 24, 1 (Jan. 2015), 45–66. <https://doi.org/10.1007/s10726-014-9379-5>
- [12] Sahar Mirzayi, Fattaneh Taghiyareh, and Faria Nassiri-Mofakham. 2022. An opponent-adaptive strategy to increase utility and fairness in agents' negotiation. *Applied Intelligence* 52, 4 (2022), 3587–3603. <https://doi.org/10.1007/s10489-021-02638-2>
- [13] Yasser Mohammad. 2021. Concurrent local negotiations with a global utility function: a greedy approach. *Autonomous Agents and Multi-Agent Systems* 35, 2 (June 2021), 28. <https://doi.org/10.1007/s10458-021-09512-y>
- [14] Yasser Mohammad. 2023. Optimal time-based strategy for automated negotiation. *Applied Intelligence* 53, 6 (March 2023), 6710–6735. <https://doi.org/10.1007/s10489-022-03662-6>
- [15] Hobart Peyton Young (Ed.). 1991. *Negotiation analysis*. University of Michigan Press.
- [16] Howard Raiffa. 1982. *The art and science of negotiation*. Harvard University Press, Cambridge, MA.
- [17] Yousef Razeghi, Ozan Yavuz, and Reyhan Aydoğan. 2020. Deep reinforcement learning for acceptance strategy in bilateral negotiations. *Turkish Journal of Electrical Engineering & Computer Sciences* 28, 4 (2020), 1824–1840. <https://doi.org/10.3906/elk-1907-215>
- [18] Ariel Rubinstein. 1982. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society* 50, 1 (1982), 97–109. <https://doi.org/10.2307/1912531> Publisher: JSTOR.
- [19] Noah G. Susskind. 2011. Wiggle room: Rethinking reservation values in negotiation. *Ohio State Journal on Dispute Resolution* 26 (2011), 79–118.
- [20] Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nicholas R. Jennings. 2012. Negotiating Concurrently with Unknown Opponents in omplex, Real-Time Domains. In *ECAI 2012. Frontiers in Artificial Intelligence and Applications*, Vol. 242. IOS Press, Amsterdam, 834–839. <https://doi.org/10.3233/978-1-61499-098-7-834>
- [21] Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nicholas R. Jennings. 2014. An Overview of the Results and Insights from the Third Automated Negotiating Agents Competition (ANAC2012). In *Novel Insights in Agent-based Complex Automated Negotiation*, Ivan Marsa-Maestre, Miguel A. Lopez-Carmona, Takayuki Ito, Minjie Zhang, Quan Bai, and Katsuhide Fujita (Eds.). Studies in Computational Intelligence, Vol. 535. Springer, Tokyo, 151–162. [https://doi.org/10.1007/978-4-431-54758-7\\_9](https://doi.org/10.1007/978-4-431-54758-7_9)
- [22] Celal O. B. Yavuz, Çağır Süslü, and Reyhan Aydoğan. 2020. Taking Inventory Changes into Account While Negotiating in Supply Chain Management. In *Proceedings of the 12th International Conference on Agents and Artificial Intelligence*, Vol. 1. SciTePress - Science and Technology Publications, Malta, 94–103. <https://doi.org/10.5220/0008976900940103>