

Optimal Non-adaptive Concession Strategies with Incomplete Information

Tim Baarslag, Rafik Hadfi, Koen Hindriks, Takayuki Ito and Catholijn Jonker

Abstract When two parties conduct a negotiation, they must be willing to make concessions to achieve a mutually acceptable deal, or face the consequences of no agreement. Therefore, negotiators normally make larger concessions as the deadline is closing in. Many time-based concession strategies have already been proposed, but they are typically heuristic in nature, and therefore, it is still unclear what is the right way to concede toward the opponent. Our aim is to construct *optimal* concession strategies against specific classes of acceptance strategies. We apply sequential decision techniques to find analytical solutions that optimize the expected utility of the bidder, given certain strategy sets of the opponent. Our solutions turn out to significantly outperform current state of the art approaches in terms of obtained utility. Our results open the way for a new and general concession strategy that can be combined with various existing learning and accepting techniques to yield a fully-fledged negotiation strategy for the alternating offers setting.

Tim Baarslag

Interactive Intelligence Group, Delft University of Technology, Mekelweg 4, Delft, The Netherlands. e-mail: T.Baarslag@tudelft.nl

Rafik Hadfi

Department of Computer Science and Engineering, Graduate School of Engineering, Nagoya Institute of Technology, Gokiso, Showa-ku, Nagoya 466-8555, Japan. e-mail: rafik@itolab.nitech.ac.jp

Koen Hindriks

Interactive Intelligence Group, Delft University of Technology, Mekelweg 4, Delft, The Netherlands. e-mail: K.V.Hindriks@tudelft.nl

Takayuki Ito

Department of Computer Science and Engineering, Graduate School of Engineering, Nagoya Institute of Technology, Gokiso, Showa-ku, Nagoya 466-8555, Japan. e-mail: ito.takayuki@nitech.ac.jp

Catholijn Jonker

Interactive Intelligence Group, Delft University of Technology, Mekelweg 4, Delft, The Netherlands. e-mail: C.M.Jonker@tudelft.nl

1 Introduction

A key insight of negotiation research is that making concessions is crucial to conducting a successful negotiation. There are important reasons to make concessions during the negotiation [25]: it is often used to elicit cooperation from the other, in the hope that the other will reciprocate in kind. Second, it conveys information to the opponent, both about the negotiator’s preferences and about the perceptions of the opponent. But most importantly, it is the time pressure of the negotiation itself (typically in the form a deadline or a perceived maximum number of bidding rounds) that operates as a force on the parties to concede [8]. An approaching deadline puts important pressure on the parties to reduce their aspirations, especially when the time pressure heightens, which is referred to as the “eleventh hour effect”.

Given the paramount importance of time in bargaining, it is not surprising that many negotiating agents adjust their level of aspiration based on the time that is left in the negotiation. There is a clear rationale behind the design of such agents, given their aim to maximize the chance of reaching an agreement in a limited amount of time. For example, well-known time dependent tactics (TDT’s) [10, 11], such as *Boulware* and *Conceder*, are characterized by the fact that they consistently concede throughout the negotiation process as a function of time. Time-based concession curves can also be observed in practice in the Automated Negotiating Agents Competition (ANAC) [3, 32]. However, in the TDT’s, as well as in some very effective agents, such *Agent K* [16] (winner of ANAC 2010) and *HardHeaded* [30] (winner of ANAC 2011), the specific concession curve is selected rather arbitrarily, and is not informed by any other insights; therefore, they make largely unfounded choices on how much to concede at each time interval.

Alternatively, behavior dependent tactics (e.g. reciprocating the opponent’s concessions by tit for tat [5, 10]) base their decision to make concessions on the actions of the other negotiating party. However, such *adaptive* approaches do not give us any information on how to concede based on time alone.

Work that presents optimal choices of how much to concede includes game theoretic work (e.g. [26]) and single-shot bargaining, also known as the *ultimatum game* [23]. However, these approaches usually assume a complete information setting, or a game where the deal is struck immediately, which we cannot apply to a typical concession-based negotiation. Furthermore, this type of work typically revolves around equilibrium strategies, which assumes full rationality on the part of both agents. We are more interested in optimal solutions for one negotiating party, playing against various classes of acceptance strategies.

This paper aims to find out how time pressure alone, in the form of a deadline, should influence the concession behavior of a negotiator against specific opponent classes. To do so, we employ methods from sequential decision theory to devise negotiation strategies that make *optimal* concessions

at each negotiation round. Finally, we show that an agent making these optimal concessions performs better than any other in our experimental setup.

We begin with an example in Section 2 that sets the stage for our time-based concession model in Section 3. We apply our methods to find optimal concessions against opponents that accept according to acceptance thresholds in Section 4 and 5. We subsequently compare the optimal bidding strategy with state of the art bidding strategies in a series of tests (Section 6). We conclude our paper with a discussion of related work (Section 7) and the contributions of this paper, as well as its implications (Section 8).

2 An Example

The following example serves as an illustration of the basic insights that have motivated our approach.

Suppose agent B is negotiating the purchase of a house with a buyer, agent A . As in the rest of this paper, we will only focus on one of the two parties; in this case B .

B has set the opening price at \$300,000, but is (secretly) willing to go down to \$250,000 if necessary. B has to strike a balance between the probability that A buys the house, and getting as much utility out of the agreement as possible (i.e., not conceding too much). Our main question is: what offers or concessions should B make towards A in order to perform optimally and maximize his own outcome? That is, what is the right way to lower the price of the house depending on the remaining time and A 's acceptance policy?

Let us consider the easiest case where B has only *one* offer to make to A , which A can either accept or reject. If B makes an offer $x \in [\$250,000; \$300,000]$ to A , then this will yield him the following utility:

$$U(x) = \begin{cases} \frac{x - \$250,000}{\$50,000}, & \text{if } A \text{ decides to accept,} \\ 0, & \text{when } A \text{ rejects the offer.} \end{cases}$$

Now suppose A is prepared to pay up to \$280,000 for the house. Of course, B does not know this, but instead presumes that A 's *acceptance threshold* could be anywhere between \$200,000 and \$300,000 with equal probability. It follows that B believes the chance that A accepts decreases linearly in terms of the price offered. Then, B 's strategy is simple: he should set the price to

$$\arg \max_x U(x) \cdot P(A \text{ accepts } x), \quad (1)$$

which is equal to

$$\arg \max_x \left(\frac{x - \$250,000}{\$50,000} \right) \cdot \left(1 - \frac{x - \$200,000}{\$100,000} \right). \quad (2)$$

We can readily compute that the maximum is reached for $x = \$275,000$, and so with only one offer to make, B should pick this price in order to optimize his utility. Note how B 's offer falls exactly in between his reservation price and his opening price. Fortunately for both, the price x is also actually lower than the maximum A was willing to pay, and therefore, she will buy the house.

This simple case serves as a good intuition to proceed to the more general setting we consider below.

3 Negotiation Model

Our negotiation model builds upon the alternating offers protocol [26]: two agents A and B exchange offers in turns on a fixed negotiation domain Ω that contains all possible negotiation outcomes. A and B have utility functions $U_A : \Omega \rightarrow [0, 1]$ and $U_B : \Omega \rightarrow [0, 1]$ and reservation values $rv_A, rv_B \in [0, 1]$ respectively. The agents will only propose offers that they deem acceptable, namely bids with higher utility than their reservation value. Likewise, the reservation values acts as the lowest utility they are willing to accept.

The process is illustrated in Figure 1. With n rounds remaining, B starts by making an offer $B_n \in \Omega$, which A can either respond to with an accept or a counteroffer. If A accepts, the negotiation ends in agreement, and both sides obtain their respective utility. If A rejects the offer, she replies with a counteroffer A_n , the number of remaining rounds decreases by one, and B can make a second offer, B_{n-1} . This process continues until B makes his last offer B_1 . If A also turns down this last offer, A and B receive their respective reservation values rv_A and rv_B .




Fig. 1: Sequential diagram of the negotiation process. B starts the negotiation with bid B_n , and after every proposal, the other responds with a counteroffer or an ACCEPT.

Throughout this work, we are concerned with finding optimal concessions for one of the parties, given certain strategies used by the opponent. We focus on agent B for his *Bidding* role, and on A for her *Accepting* role. In particular, we do not focus on what offers A should generate, or what acceptance policy B should employ.

There is an important class of strategies that allows us to make some further simplifications. In line with our aim to find concession curves that depend only on the time remaining, we will assume B 's strategy is not adaptive (i.e., B does not change his behavior according to the bids that have been exchanged). This means B can completely ignore A 's bids, and we can code them by ACCEPT or REJECT instead. Other examples of such agents include the family of TDT's [10], the *sit-and-wait* agent [2], and IAMcrazyHaggler [31]. Note that this then effectively defines an one-sided bidding protocol in which B submits n bids to A in a sequential manner, which is equivalent to a repeated ultimatum game with reservation values [28].

The fact that B does not adapt to his opponent puts him in a significantly more difficult epistemic position than in a typical negotiation, since B could normally gather valuable information from A 's counter offers. This holds in a very fundamental sense: apart from the incoming offers, B can only distinguish n different possible states, namely the number of rejects that he received. This means that we can already compute the appropriate *concession curve* before the start of the negotiation; i.e., the bids or utilities that need to be sent out at every round.

Also note that as time moves *forward*, the indexing of the offers by B runs *backward*. This has the advantage that it simplifies the calculations of certain recurrence relations we will encounter in this paper, as the expected utility with $j + 1$ rounds remaining depends on the expected utility in the future (i.e., with j rounds remaining). Secondly, it allows us to define our model without a preset deadline, as we do not need to specify the maximum number of rounds beforehand.

4 Conceding and Accepting

The aim of this paper is to find the right concession behavior for B given that he only has n offers to try out. B will only propose offers that he himself deems acceptable, namely bids in the following set:

$$O_B = \{\omega \in \Omega \mid U_B(\omega) \geq rv_B\}. \quad (3)$$

B can make his bids in many ways (see also our discussion of related work in Section 7).

A well-known method using concession curves are the TDT's such as *Boulware* or *Conceder*. When there are j rounds remaining (out of a total of

n rounds), this strategy makes a bid with utility closest to

$$P_{\min} + (P_{\max} - P_{\min}) \cdot \left(1 - k - (1 - k) \cdot \left(1 - \frac{j}{n} \right)^{\frac{1}{e}} \right),$$

for certain choices of k , e , P_{\min} and P_{\max} , which control the concession rate and the minimum and maximum target utilities. These tactics form the basis of many successful negotiation strategies [3, 11, 16, 30], in which variants of the above formula are used to decide the appropriate concessions, combined with advanced techniques such as preference modeling and strategy prediction. In what follows, we will propose a different kind of concession curve that is also able to provide the groundwork for designing an advanced negotiation agent.

B 's optimal strategy of course depends on how A chooses to accept or reject. While B makes his offers, A has to solve some kind of stopping problem to decide when the offer is sufficient.

Player A will have to either accept or reject a bid $\omega \in \Omega$ in every round j . A may update her beliefs in round j based on ω and the bidding history of the opponent, which is an element h_j from the set of offers $\mathcal{H}_j = \Omega^{n-j}$ made by B . Note that we only need to consider the bids made by B , since the rejects of A are already implicitly represented. For example, for $n = 8$ and $j = 5$, the information set of A is a history of rejected offers $h_5 = (B_6, B_7, B_8) \in \mathcal{H}_5$, adjoined with the bid B_5 of the current round.

Hence, we can represent A 's acceptance strategy in round j as a function

$$\alpha : \mathcal{H}_j \times \Omega \rightarrow \{\text{ACCEPT}, \text{REJECT}\}. \quad (4)$$

In this paper, we will focus on a specific set of strategies that accept according to *acceptance thresholds*. That is, A 's acceptance strategy in round j is specified by a utility constant α_j (possibly dependent on rv_A) such that

$$\alpha(h_j, \omega) = \begin{cases} \text{ACCEPT} & \text{if } U_A(\omega) \geq \alpha_j, \\ \text{REJECT} & \text{otherwise.} \end{cases} \quad (5)$$

We believe acceptance thresholds are a natural set of acceptance strategies to consider, since it is reasonable to assume that if A finds a bid acceptable, then so is any bid with higher utility.

One of the simplest acceptance strategies for A is *satisficing*: accepting any offer with a utility above her reservation value by setting all α_j equal to rv_A . This is done, for example, when a negotiator is more concerned with getting any deal at all than reaching the best possible deal (e.g., [22]). Needless to say, this is a very simple acceptance strategy, as A normally wants to get as much out of the negotiation as possible. In other applications, illustrated in Figure 2, A 's threshold would be higher at the beginning of the negotiation, and would slowly decrease towards rv_A . However, as we shall see in the next

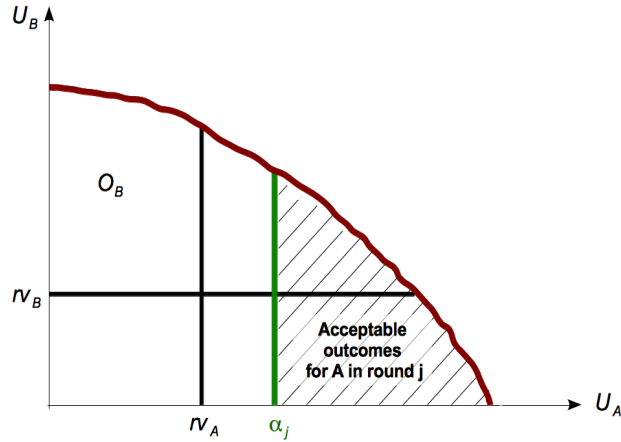


Fig. 2: A 's acceptable offers for round j consist of all bids ω with $U_A(\omega) \geq rv_A$, while B 's possible offers O_B consist of all bids ω such that $U_B(\omega) \geq rv_B$.

section, the case of a satisficing acceptor already requires highly effective conceding behavior by B .

A might also employ a more fundamental approach by trying to optimize her own utility, taking into account the number of rounds remaining. Optimal stopping theory provides optimal solutions for A for the case that A has incomplete information about B 's offers. From [7, 19, 33], we know that the optimal solution against a bidder B with completely unknown utility goals is as follows:

Proposition 1. *When B makes random bids of utility uniformly distributed in $[0, 1]$, and with j offers still to be observed, A 's optimal acceptance strategy is to accept an offer of utility x exactly when $x \geq v_j$, where v_j satisfies the following equation:*

$$\begin{cases} v_0 = rv_A, \\ v_j = \frac{1}{2} + \frac{1}{2}v_{j-1}^2. \end{cases} \quad (6)$$

As it will turn out, this strategy is closely related to the strategy B should follow against satisfying acceptors.

5 Making Optimal Offers

We will now outline our general method to make optimal offers, which extends our running example from Section 2 to a general domain with an arbitrary number of remaining rounds.

Suppose A uses an acceptance strategy based on acceptance thresholds $(\alpha_0, \alpha_1, \dots)$, with $\alpha_j \geq rv_A$; that is, A accepts an offer $\omega \in \Omega$ with j remaining

rounds exactly when $U_A(\omega) \geq \alpha_j$. Suppose we have $j + 1$ rounds to go, and B has to decide the optimal bid B_{j+1} to make.

The general formula for the expected utility $U_{j+1}(\omega)$ for B of offering $\omega \in O_B$ is as follows:

$$U_{j+1}(\omega) = \begin{cases} U_B(\omega), & \text{if } A \text{ accepts,} \\ \text{the expected utility} \\ \text{with } j \text{ remaining rounds, if } A \text{ rejects.} \end{cases}$$

When there are no more rounds remaining, B will get rv_B , hence $U_0(\omega) = rv_B$. The recursive nature of this equation allows us to employ techniques from sequential decision theory to formulate the optimal way to make concessions.

We will write U_{j+1} for the highest expected utility B can obtain in round $j + 1$, which is thus given by the following equations:

$$\begin{aligned} U_0 &= rv_B, \quad \text{and} \\ U_{j+1} &= \max_{\omega \in O_B} U_B(\omega) \cdot P(U_A(\omega) \geq \alpha_{j+1}) + U_j \cdot P(U_A(\omega) < \alpha_{j+1}) \\ &= U_j + \max_{\omega \in O_B} (U_B(\omega) - U_j) \cdot P(U_A(\omega) \geq \alpha_{j+1}). \end{aligned}$$

The corresponding optimal *bid* that B should make is given by the ω that maximizes the above equation, and therefore,

$$B_{j+1} = \arg \max_{\omega \in O_B} (U_B(\omega) - U_j) \cdot P(U_A(\omega) \geq \alpha_{j+1}).$$

Solving this equation would be straightforward, if it were not for the fact that B in general does not have full knowledge of a number of aspects in this equation: U_A is of course unknown to B , and so are the acceptance thresholds α_j . To make matters worse, the acceptance thresholds generally depend on the reservation value of A , which is also unknown to B .

B will therefore have to make some assumptions about U_A and her concession thresholds. Assume that B has an estimation of the reservation value of A , which does not change according to A 's behavior. As in [20], the estimation is characterized by a probability distribution $F_j(x)$ for every remaining round j , where $F_j(x)$ denotes the probability that A 's reservation value is no greater than x .

We will now analytically solve the specific case of an opponent with a satisficing accepting strategy. To get concrete examples of concession curves, we will consider a classical buyer-seller scenario as in [10, 11] to evaluate our solutions. We study the negotiation scenario of *Split the Pie* [26, 29], where two players have to reach an agreement $x \in [0, 1]$ on the partition of a pie of size 1. The pie will be partitioned only after the players reach an agreement. In this setting, we instantiate $\Omega = [0, 1]$ to represent a pie of size 1, with A and B having opposing preferences on them: $U_B(x) = x$ and $U_A(x) = 1 - x$.




Fig. 3: Graph of B_j and U_j for remaining rounds $j \in \{0, 100\}$.

We assume that A 's acceptance strategy is limited to the class where she accepts any offer that is better than her reservation value; i.e., $\forall_j \alpha_j = rv_A$, where we assume rv_A is uniformly distributed. The probability that A accepts in round j is now:

$$P(U_A(\omega) \geq \alpha_j) = P(U_A(\omega) \geq rv_A) = F_j(U_A(\omega)). \quad (7)$$

The general formula now simplifies to

$$U_{j+1} = U_j + \max_{x \geq rv_B} (x - U_j) \cdot (1 - x). \quad (8)$$

Note that the following holds, even for the general setting:

Proposition 2. *When B has more rounds remaining, B can expect to get more utility out of our negotiation. That is, for every remaining round j , we have $U_{j+1} \geq U_j$.*

Proof. B can always make the bid $\omega_{\max} = \arg \max_{\omega} U_B(\omega)$ with $j + 1$ rounds remaining, to get at least as much utility as in the next round.

With the help of Prop. 2, we can show the maximum in equation (8) is attained for $x = B_j$, where B_j satisfies the following relationship:

$$\begin{cases} B_1 = \frac{rv_B + 1}{2}, \\ B_{j+1} = \frac{1 + U_j}{2}. \end{cases} \quad (9)$$




Fig. 4: Expected utility U_j for B with $j \in \{1, 2, 3, 4\}$ rounds to go, depending on B 's reservation value rv_B .

This yields the following recurrence relation for U_j :

$$\begin{cases} U_0 = rv_B, \\ U_{j+1} = \frac{1}{4}(U_j + 1)^2. \end{cases} \quad (10)$$

With these equations, we can now compute the expected value of the optimal bids B has to make in terms of his own reservation value (see Figure 4). For example:

$$\begin{aligned} U_1 &= \frac{1}{4}(1 + rv_B)^2, \\ U_2 &= \frac{1}{64}(5 + rv_B(2 + rv_B))^2, \\ &\vdots \end{aligned}$$

The corresponding optimal concessions by B are as follows:

$$\begin{aligned} B_1 &= \frac{1 + rv_B}{2}, \\ B_2 &= \frac{1}{8} \cdot (5 + rv_B \cdot (rv_B + 2)), \\ &\vdots \end{aligned}$$

Note that Eq. 10 is related to the logistic map: if we substitute $U_j = 1 - 4x_j$ in Eq. 10 we get the equivalent relation of the logistic map $x_j = x_{j-1}(1 - x_{j-1})$ at $r = 1$, so we cannot expect to solve the recurrence relation. Note, however, that both equations (9) and (10) are expressed in terms of the value U_j . We obtain a much more elegant formulation when we express them in terms of B_j :

Proposition 3. *For all $j \geq 1$,*

$$B_{j+1} = \frac{1}{2} + \frac{1}{2}B_j^2,$$

and

$$U_j = B_j^2.$$

Proof. Both statements follow from rewriting equations (9) and (10).

Corollary 1. *For $rv_B = 0$, there exists the following connection between B_j , U_j , and the optimal stopping cut-off values v_j from equation (6):*

$$B_j = v_j,$$

and

$$U_j = v_j^2.$$

Proof. When $rv_B = 0$, then $B_1 = \frac{1}{2}$, and consequently, the B_j sequence has the same starting value and definition as the optimal stopping sequence v_j .

Corollary 1 shows that optimal *bidding* against a satisficing acceptor with unknown reservation value is the mirrored version of optimal *accepting* against a bidder that makes unknown offers. In both cases, the idea is the same, but with switched roles: both the optimal bidder and the optimal stopper aim to pick the optimal utility threshold that simultaneously maximizes the expected utility of an agreement and the chance of acceptance, given stochastic behavior by the opponent.

We conclude this section with an example that relates our results to the housing example of Section 2.

Example 1. Our assumptions in this section are consistent with our housing example when we scale the pie Ω to $[\$250,000; \$300,000]$, and when $rv_B = 0$ corresponds to the utility of selling the house for $\$250,000$. Our results indicate that the seller of the house should act as follows: we see that indeed, $B_1 = \frac{1}{2}$, so with one bid remaining, B should make a bid halfway between $\$250,000$ and $\$300,000$, as we had calculated before, with an expected utility of $U_1 = \frac{1}{4}$.

Also, $B_2 = \frac{5}{8}$, which means with *two* bids remaining, B should offer $\$250,000 + \frac{5}{8} \cdot \$50,000 = \$281,250$. See Table 1 for other entries. Note that since B does not learn anything about A , B actually overestimates the expected utility, and more so when the number of rounds increases (as can be

j	Offer	Utility
1	\$275,000	0.25
2	\$281,250	0.39
3	\$284,766	0.48
4	\$287,086	0.55
5	\$288,754	0.60
10	\$293,055	0.74
100	\$299,060	0.96

Table 1: The optimal offers to make and their expected utility in the housing example, given the remaining amount of j bidding rounds.

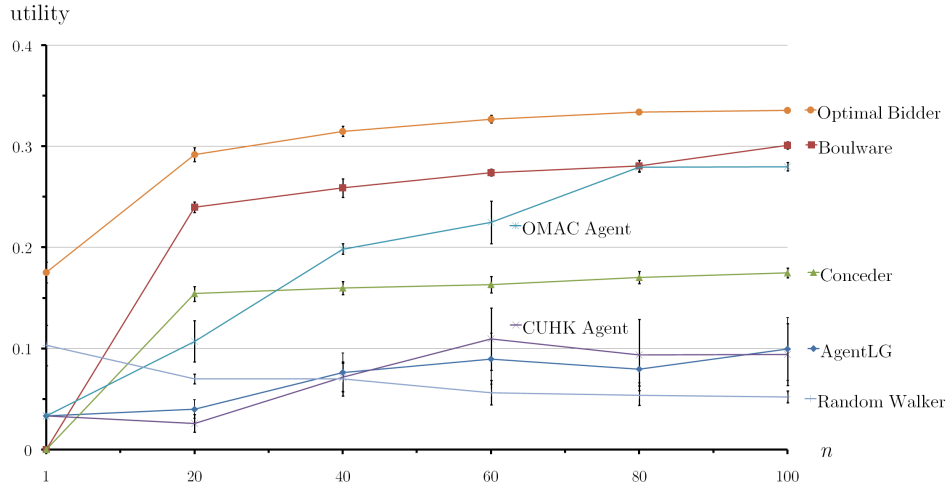


Fig. 5: The utility obtained by the bidding strategies in our experiments for different values of the total number of rounds n . The vertical bars indicate one standard deviation from the mean.

seen in Table 1). This is because B is not able to deduce, from A 's rejects, that high utility values for B are simply not attainable.

6 Experiments

In order to test the efficacy of the optimal bidding technique given by Prop. 3, we integrated it into a fully functional negotiating agent. Given j remaining rounds it makes an offer with utility target B_j as defined by equation (9). It does not accept any offers and it does not model the opponent in any way.

For the opponents (side A), we selected various well-known negotiation agents available for our setting, including the top three agents of ANAC 2012 [32], namely *CUHK Agent* [15], *AgentLG*, and *OMAC Agent* [9]. We also included the time dependent tactics (TDT’s) *Boulware* (with concession factor $e = 0.2$), and *Conceder* ($e = 2$) taken from [10]. As a baseline, we included the *Random Walker* strategy [14], which generates random bids. We then compared the optimal bidder’s performance with the same set of strategies on side B . To analyze the performance of different agents, we employed GENIUS [21], which is an environment to design and evaluate automated negotiators’ strategies and their components [4].

Note that some of the agents in our setup are originally designed to work with the alternating offers protocol, while we essentially employ a one-sided bidding protocol; however, since our model is a simplified alternating offers protocol, it is easy to adjust the agents to work in our setting: for the opponents, we ignore any bids that are sent out; i.e., when side A accepts, the negotiation ends in agreement, while A ’s counter-offers count as rejects and are ignored. In effect, this means only A ’s acceptance mechanisms [6] are used.

For our negotiation scenario we use a discretized version of *Split the Pie*. We set $rv_B = 0$, and we selected varying reservation values for A : $rv_A \in \{0, 0.1, 0.2, \dots, 0.9\}$. We ran our experiments using varying total number of rounds $n \in \{1, 20, 40, \dots, 100\}$, and we repeated every negotiation session 5 times for statistical significance.

The results of our experiment are plotted in Figure 5, with the total amount of rounds n varying between 1 and 100. As is evident from the results, the optimal bidder significantly outperforms all agents in all cases (*one-tailed t-test*, $p < 0.01$ for every n). The good relative performance of the optimal bidder is even more pronounced for $n = 1$, and it is easy to see why. Many of the other strategies will persist in aiming for a high utility, even with only few bids to send out. This results in many break-offs, while optimal bidder will settle for a much lower value in the last rounds with more potential agreements. For example, for $n = 1$, optimal bidder sets his utility target halfway between his reservation value and the maximum attainable; i.e., $B_1 = \frac{rv_B + 1}{2}$. On the other hand, with more bids remaining, optimal bidder acts like an extreme *Boulware* strategy, trying to get as much out of the negotiation as possible. Indeed, optimal bidder and *Boulware* tend to act more similar as n increases. *CUHK Agent* (the winner of ANAC 2012) and *AgentLG* obtain particularly low scores. The main reason for this is that these strategies are very behavior-dependent and do not take into account the remaining time as much as they need to.

Note that almost all agents obtain higher utilities with more negotiation rounds. This is to be expected, as more time allows for a more fine-grained search of what is acceptable for the opponent. The only exception is *Random Walker*, who only increases the chances to make a disadvantageous bid for itself with more available time.

7 Related Work

Concession making is an important process as it might critically affect the outcomes of the negotiation. Despite its importance, the literature lacks a fully comprehensible theory where the uncertainty settings as well as the optimal concession strategies are formally represented.

Our interest lie in investigating the optimal concession strategies against particular classes of acceptance strategies. The concessions we are referring to, relate to the subjective type of concessions as proposed in the categorization of [17]. However, our focus is on the type of concessions that are mainly driven by time dependence for the offers generation, such as the time dependent tactics (*Boulware*, *Conceder*) found in [10, 11, 12, 13], or the *monotonic concession protocol* [24], where each agent first starts with the best offer for itself and then iteratively makes compromises to the other agent.

The bargaining game we take as main setting could be seen as an instance of the ultimatum game, in which a player proposes a deal that the other player may only accept or refuse [28]. Another work, [27], explores a similar bargaining model as well; that is, models with one-sided incomplete information and one sided offers. It investigates the role of confrontation in negotiations and uses optimal stopping is to decide whether or not to invoke conflict. Our setting can also be found in [1], which presents an alternating offer protocol for bilateral bargaining with imperfect information and deadline constraints. Here too, the authors use backward induction, in order to provide a way to find all sequential equilibria.

To the best of our knowledge, this is the first work that makes usage of the optimal stopping rule to *generate* offers in an incomplete information setting and compares it to other concession techniques. Most of the previous work makes use of optimal stopping theory in a negotiation setting to formulate acceptance strategies [7, 18, 19, 33]; for instance, to find the stopping time as to maximize the acceptor's utility. In this work, the acceptor (typically represented by the buyer) can either stop the negotiation by accepting the seller's offer or continue the negotiation by rejecting. Other such problems concern the decision to accept sequential job offers while trying to maximize the sum of the payments of all accepted jobs [33]. Similar to our model, the agent is only incentivized by the limited time resource and the expectation of high utility in the future.

The major difference with our work and the above approaches is that we use the optimal stopping rule for concessions, instead of focusing on the optimal time to accept. Our work is defined more as the *complimentary* version of these approaches, in the sense that our formulation of optimal bidding rules happen to resemble optimal acceptance rules. Another key point is that we do not assume that the players' strategies are fixed, which allows us to formulate optimal bidding strategies against certain types of accepting strategies.

8 Conclusion

This paper presents a theoretical model to formulate optimal concession curves against strategies that decide when to accept using acceptance thresholds. We calculated these optimal concessions by employing sequential decision methods, and we showed that they significantly outperform state of the art concession strategies, even against a much wider set of acceptance strategies. As far as the authors are aware, this is the first time such informed time-based concession strategies have been formulated and tested in practice.

Our results demonstrate that our optimal bidding mechanism is an effective way for a negotiating agent to take into account the passing of time. We believe even more effective concession curves can be computed by studying more general negotiation settings (e.g., dynamic, multi-lateral, and concurrent negotiation environments) and broader ranges of acceptance strategies than we did in this paper. Further improvement could be made by introducing a form of learning to the optimal bidder. Eventually, we envision a design of an automated negotiator that incorporates our optimal concession curve with regard to time-related concessions, while other types of concessions (e.g., to elicit cooperation or to convey information) are handled separately by other concession modules.

References

1. Bo An, Nicola Gatti, and Victor Lesser. Bilateral bargaining with one-sided uncertain reserve prices. *Autonomous Agents and Multi-Agent Systems*, 26(3):420–455, 2013.
2. Bo An, Kwang Mong Sim, Liang Gui Tang, Chun Yan Miao, Zhi Qi Shen, and Dai Jie Cheng. Negotiation agents’ decision making using markov chains. In Takayuki Ito, Hiromitsu Hattori, Minjie Zhang, and Tokuro Matsuo, editors, *Rational, Robust, and Secure Negotiations in Multi-Agent Systems*, volume 89 of *Studies in Computational Intelligence*, pages 3–23. Springer Berlin Heidelberg, 2008.
3. Tim Baarslag, Katsuhide Fujita, Enrico H. Gerding, Koen Hindriks, Takayuki Ito, Nicholas R. Jennings, Catholijn Jonker, Sarit Kraus, Raz Lin, Valentin Robu, and Colin R. Williams. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence*, 198(0):73 – 103, 2013.
4. Tim Baarslag, Koen Hindriks, Mark Hendrikx, Alex Dirkzwager, and Catholijn Jonker. Decoupling negotiating agents to explore the space of negotiation strategies. In Ivan Marsa-Maestre, Miguel A. Lopez-Carmona, Takayuki Ito, Minjie Zhang, Quan Bai, and Katsuhide Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 61–83. Springer, Japan, 2014.
5. Tim Baarslag, Koen Hindriks, and Catholijn Jonker. A tit for tat negotiation strategy for real-time bilateral negotiations. In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories*,

- Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 229–233. Springer Berlin Heidelberg, 2013.
6. Tim Baarslag, Koen Hindriks, and Catholijn Jonker. Effective acceptance conditions in real-time automated negotiation. *Decision Support Systems*, 60:68–77, April 2014.
 7. Tim Baarslag and Koen V. Hindriks. Accepting optimally in automated negotiation with incomplete information. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 715–722, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems.
 8. Peter J.D. Carnevale and Edward J. Lawler. Time pressure and the development of integrative agreements in bilateral negotiations. *The Journal of Conflict Resolution*, 30(4):pp. 636–659, 1986.
 9. S. Chen and G. Weiss. OMAC: a discrete wavelet transformation based negotiation agent. In I. Marsa-Maestre, M.A. Lopez-Carmona, T. Ito, M. Zhang, Q. Bai, and K. Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 187–196. Springer, Japan, 2014.
 10. Peyman Faratin, Carles Sierra, and Nick R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159 – 182, 1998. Multi-Agent Rationality.
 11. S. Shaheen Fatima, Michael Wooldridge, and Nicholas R. Jennings. Optimal negotiation strategies for agents with incomplete information. In *Revised Papers from the 8th International Workshop on Intelligent Agents VIII*, ATAL '01, pages 377–392, London, UK, UK, 2002. Springer-Verlag.
 12. Shaheen Fatima, Michael Wooldridge, and Nicholas R. Jennings. Optimal negotiation of multiple issues in incomplete information settings. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 3*, AAMAS '04, pages 1080–1087, Washington, DC, USA, 2004. IEEE Computer Society.
 13. Shaheen S. Fatima, Michael Wooldridge, and Nicholas R. Jennings. Bargaining with incomplete information. *Annals of Mathematics and Artificial Intelligence*, 44(3):207–232, July 2005.
 14. Dhananjay K. Gode and Shyam Sunder. Allocative efficiency in markets with zero intelligence (ZI) traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
 15. Jianye Hao and Ho-Fung Leung. ABiNeS: An adaptive bilateral negotiating strategy over multiple items. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology - Volume 02*, WI-IAT '12, pages 95–102, Washington, DC, USA, 2012. IEEE Computer Society.
 16. Shogo Kawaguchi, Katsuhide Fujita, and Takayuki Ito. Compromising strategy based on estimated maximum utility for automated negotiating agents. In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, pages 137–144, Berlin, Heidelberg, 2012. Springer-Verlag.
 17. Gregory E. Kersten, Rustam Vahidov, and Dmitry Gimon. Concession-making in multi-attribute auctions and multi-bilateral negotiations: Theory and experiments. *Electronic Commerce Research and Applications*, 12(3):166 – 180, 2013. Negotiation and E-Commerce.
 18. K. Kolomvatsos, C. Anagnostopoulos, and S. Hadjiefthymiades. Determining the optimal stopping time for automated negotiations. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 99:1–1, September 2013.

19. Bjorn Leonardz. *To stop or not to stop. Some elementary optimal stopping problems with economic interpretations*. Almqvist & Wiksell, Stockholm, 1973.
20. Cuihong Li, J. Giampapa, and K. Sycara. Bilateral negotiation decisions with uncertain dynamic outside options. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 36(1):31–44, Jan 2006.
21. Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. Genius: An integrated environment for supporting the design of generic automated negotiators. *Computational Intelligence*, 30(1):48–70, 2014.
22. Christoph Niemann and Florian Lang. Assess your opponent: A bayesian process for preference observation in multi-attribute negotiations. In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *Advances in Agent-Based Complex Automated Negotiations*, volume 233 of *Studies in Computational Intelligence*, pages 119–137. Springer Berlin / Heidelberg, 2009.
23. H. Raiffa. *The Art and Science of Negotiation*. Belknap Press of Harvard University Press, 1982.
24. Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press, Cambridge, MA, USA, 1994.
25. Jeffrey Z. Rubin and Bert R. Brown. *The social psychology of bargaining and negotiation*. Academic Press, New York, 1975.
26. Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.
27. Santiago Sanchez-Pages. The use of conflict as a bargaining tool against unsophisticated opponents. ESE Discussion Papers 99, Edinburgh School of Economics, University of Edinburgh, March 2004.
28. Tilman Slembeck. Reputations and fairness in bargaining - experimental evidence from a repeated ultimatum game with fixed opponents. Experimental, EconWPA, 1999.
29. Ingolf Stahl. *Bargaining theory*. Economic Research Institute, Stockholm, 1972.
30. Thijs van Krimpen, Daphne Looije, and Siamak Hajizadeh. Hardheaded. In Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo, editors, *Complex Automated Negotiations: Theories, Models, and Software Competitions*, volume 435 of *Studies in Computational Intelligence*, pages 223–227. Springer Berlin Heidelberg, 2013.
31. Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nicholas R. Jennings. Iamhaggler: A negotiation agent for complex environments. In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, pages 151–158, Berlin, Heidelberg, 2012. Springer-Verlag.
32. C.R. Williams, V. Robu, E.H. Gerding, and N.R. Jennings. An overview of the results and insights from the third automated negotiating agents competition (ANAC 2012). In I. Marsa-Maestre, M.A. Lopez-Carmona, T. Ito, M. Zhang, Q. Bai, and K. Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 151–162. Springer, Japan, 2014.
33. Mengxiao Wu, Mathijs Weerd, and Han Poutré. Acceptance strategies for maximizing agent profits in online scheduling. In Esther David, Valentin Robu, Onn Shehory, Sebastian Stein, and Andreas Symeonidis, editors, *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, volume 119 of *Lecture Notes in Business Information Processing*, pages 115–128. Springer Berlin Heidelberg, 2013.