Relative Performance of Bilateral Multiattribute Negotiation Strategies in Open Markets

Jacob Brue, Joe Shymanski, Selim Karaoglu, Sandip Sen
University of Tulsa
800 S. Tucker Dr., Tulsa, OK 74104
{jacob-brue, joe-shymanski, sek6301, sandip-sen}@utulsa.edu

Abstract

The long-running Automated Negotiating Agents Competition (ANAC) is comprised of various agent-agent and human-agent negotiation leagues. One such competition is the Automated Negotiation League (ANL) which involves repeated, bilateral negotiation over multiple issues. Researchers have investigated a tournament setting for this scenario involving a small, fixed number of agents. We are interested in automated agents participating in large and open marketplaces containing many instances of well-known agent types of varying sophistication. We experiment with four representative negotiation behaviors as agent types: Hardliner, Boulware, Conceder, and Tit-for-Tat. We simulate open markets with varying negotiation domain sizes, agent type distributions, and negotiation time available to evaluate the relative performances of different negotiation strategies. We analyze and report relative performances of the strategies on relevant performance metrics. We also extend this analysis using a head-to-head matrix.

Introduction

Real-world applications of intelligent agents include the retail, e-commerce, legal, business, and industrial sectors where there is an increased demand for automated systems that can accurately represent a party in a negotiation to secure better deals for both parties (Ransbotham et al. 2017; Tung 2019). There is increasing interest in agents representing human users in negotiating deals with other human and autonomous agents (Baarslag et al. 2017; Lin et al. 2014; Peled, Gal, and Kraus 2013; Rosenfeld et al. 2015).

Researchers investigating variations of negotiation scenarios and the corresponding negotiation strategies have created benchmark domains and platforms that allow the evaluation of competing approaches. Various agent-agent and human-agent negotiation environments have been instituted, providing testing grounds and insight into the design and deployment of effective automated negotiation approaches. One such environment is the Automated Negotiation League (ANL), held as part of the Automated Negotiating Agents Competition (ANAC) and in association with major international conferences, which involves bilateral negotiation between two agents over multiple issues. Variants of this environment have been experimented with over the past decade (Aydoğan et al. 2019; Aydoğan et al. 2020; Baarslag et al. 2012; Mell et al. 2018) and researchers have analyzed the fielded agent behaviors and tournament outcomes from these competitions (Baarslag et al. 2013; de Jong 2022). Although these tournaments and corresponding analysis shed light on effective strategies within such structured environments, it is unclear how these results generalize to more open and unstructured environments.

We investigate the relative performance of representative agent strategies in open environments, such as large marketplaces that involved a diverse population of agents in flux. Individual agents in such marketplaces, representing known strategy types, are likely to differ in the choice of parameters or other aspects of their negotiation strategies that affect outcome and performance differences. It is unlikely that any pair of agents interact multiple times. Hence, the use of approaches that learn specific opponent models over multiple interactions is not helpful. When selecting representative negotiation strategies to experiment with in open markets, therefore, we exclude learning strategies based on:

- As argued by (de Jong 2022), simpler, nonadaptive strategies can be competitive with more complex strategies.
- We believe that large and open marketplaces, with low entry barriers and demographically diverse participants, will mostly contain easy-to-implement, relatively simple strategies rather than more complex approaches requiring a scholarly appreciation of bilateral, multi-issue negotiation theory.
- Interactions between specific individuals in such marketplaces are likely sporadic, even anonymous, and may not provide sufficient opportunities to develop an accurate model of an opponent's behavior from early interactions to be leveraged in later sustained interactions.

Representative strategies included in our experiments on a simulated open market were based on a review of the ANL competition entries over the past few years. Most of the participants are variants of agents with time-varying "aspiration levels", where they are more willing to accept lower utility deals as the negotiation deadline approaches. The concession rate of aspirations over the negotiation period can be parameterized. The extreme case of this is a *Hardliner* agent that never lowers its aspiration level. Based on the shape of the curve representing a change in aspiration level, we identify two additional agent types when compared to a linear decrease with time: the *Boulware* and *Conceder* agents, with aspiration levels that decrease slower and faster than linear, respectively. We also include an adaptive, *tit-for-tat* strategy that responds to the negotiation partner's lowering of aspiration levels (Mirzayi, Taghiyareh, and Nassiri-Mofakham 2022).

To simulate open market scenarios, we ran experiments that contained many bilateral negotiations. For each negotiation, two agent strategies are sampled from fixed distributions over the four identified strategy types. Individual sampled agents are not uniquely identified, and learning agents, which build and leverage opponent models from past interactions with their opponent, cannot use that information in this market. This process simulates a large population of agents, each using a variant of a particular parameterized negotiation strategy. Utilities of all instantiated agents of a type are aggregated to report the overall performance of that type. We identify dominant strategies for representative agent distributions. We also analyze the effect of other system parameters on the relative performance of agent strategy types.

The goal of this research is to understand how market demographics and negotiation domains favor certain strategies. These results can be used to select the preferred strategy to adopt when entering a marketplace where market demographics and relevant parameters can be estimated.

Related Work

Although previous research on negotiation has focused more on developing smart negotiation strategies, there has been some work on architectures that specify the underlying components and reasoning modules. We use a general purpose negotiation architecture (Xu et al. 2020) to implement the agents used in the simulations of the open marketplace.

Significant effort has been put into developing adaptive strategies that learn opponent models and strategies (Bagga, Paoletti, and Stathis 2022; Baarslag et al. 2016; Mirzayi, Taghiyareh, and Nassiri-Mofakham 2022; Sengupta, Mohammad, and Nakadai 2021). We do not include adaptive strategies in the mix of agents we test in our simulated marketplace as it is unlikely that the same opponent will be seen again in a large, open market environment.

We do not utilize an evolutionary framework where agents adopt the strategies of higher-performing agents in the population as we believe that agents may not have access to global information about relative performance of agent strategy types. Therefore, the analysis and identification of emergent behavior from multiagent learning (Bloembergen et al. 2015) and study of equilibria conditions involving evolutionary stable strategies (Hines 1987; Morales et al. 2018; Mori and Ito 2016) are not applicable to this work.

ANAC has received attention from multiagent researchers investigating interesting negotiation scenarios and strategies (Aydoğan et al. 2020; Baarslag et al. 2012; Jonge et al. 2018). To run our experiments, we adapted the GeniusWebPython software platform, used in the ANAC ANL competitions to run tournaments between a small, fixed number of competitors. The four major strategy types we use for experimentation are inspired by common agent behaviors in ANL competitions that have also been observed in human negotiations (Maaravi, Idan, and Hochman 2019; Zohar 2015) including hard-lining, strategic conceding, ultimatums, positive reinforcement, splitting the difference, etc.

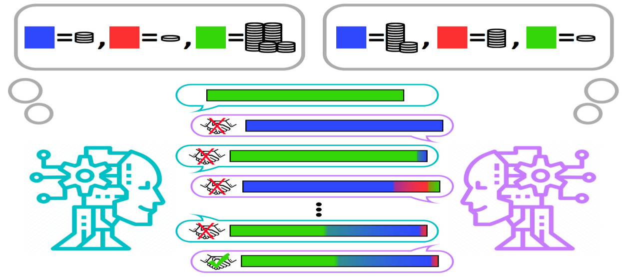

Figure 1: A bilateral negotiation scenario.

Negotiation in an Open Marketplace

The 2022 ANL (Automated Negotiation League) competition used Alternating Offers Protocol, which requires agents to offer a bid, receive an offer from the opponent, then either accept their offer or generate a new one. The utility of an agent is based on their preference profile and the accepted bid. Both agents receive zero utility if no agreement is reached. A negotiation domain is the set of issues (represented with colors in Figure 1) and the collection of values (represented with coin stacks in Figure 1) over which agents negotiate. Each issue has an associated set of values. An offer consists of a chosen value for every issue. The utility function of an agent is represented by a linear additive preference profile. The preference profiles for the two agents in a negotiation comprise a scenario (see Figure 1).

A bilateral negotiation takes place between a host agent, a, and a partner agent, p. Each negotiator has an associated profile P consisting of a vector of issue weights and a utility for each possible value for each issue. Let $\mathcal{I}$ be the set of issues being negotiated. $\forall i \in \mathcal{I}$, w_i^a (w_i^p) is the issue weight for issue i for the agent (partner). u_{ij}^a (u_{ij}^p) is the utility for the agent (partner) for the jth value of the ith issue. An offer O is a vector of issues of length $|\mathcal{I}|$, which contains a value for each issue being negotiated. O_i represents the value of the ith issue of the offer. If both agents agree on an offer O, the utility to the agent is $U_a(O) = \sum_{k \in \mathcal{I}} w_k^a u_{kO_k}^a$ and that to the partner is $U_p(O) = \sum_{k \in \mathcal{I}} w_k^p u_{kO_k}^p$. If two agents fail to arrive at an agreement by the end of the negotiation deadline, d, they receive a default utility of $U_{conflict}$, which is zero for our experiments. We compute the average utility received by an agent over all of its negotiations, $\bar{U}_a$.

Each negotiation round includes a time limit. To control for factors such as the speed of the machine running the marketplace simulation and the speed of the agent implementations, we used simulated instead of real time. Simulated time advances by a time step sampled from a normal distribution.

Negotiation Strategy Types

Monotonic Concession with Time (MCT)

MCT strategies open with selfish offers and concede later.
Initial Offer: Offer with highest utility, $\arg\max_O U_a(O)$.

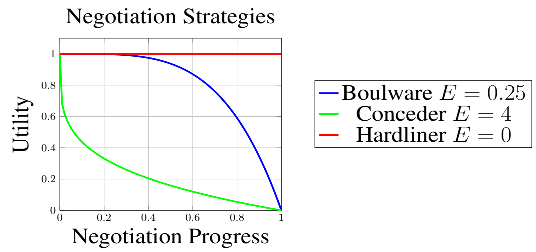

Figure 2: Negotiation strategies for different agent classes.

Offer response: Offer with utility closest to, but not lower than the target utility. The target utility decreases as the negotiation progresses: $T(p) = 1 - p^{1/E}$, where $0 \leq p \leq 1$ varies from the start (0) to the end (1) of the negotiation.

Variations on MCT: The parameter E controls how quickly the agent drops its expectation from its maximum (1) to its minimum acceptable utility (0). We used three agents for testing (Mccalley et al. 2020) (see Figure 2):

Hardliner agent: A hard-balling agent with $E = 0$ that offers only its best bids the whole round and never concedes.

Boulware agent: A stubborn agent with $0 < E < 1$ that offers its best bids for much of the round and concedes only late in the round.

Conceder agent: An eager, concession-happy agent with $E > 1$ that drops its expected utility early in the round.

Tit-for-Tat Strategy

A *tit-for-tat* negotiation strategy mimics the most recent move of the partner. We chose the MiCRO agent (de Jong 2022), which concedes only if it detects that its partner has conceded, for our implementation of the tit-for-tat agent.

Initial Offer: Best offer in its profile.

Offer Response: Only accepts offers it has previously proposed. If the partner has previously made the current offer, the MiCRO agent will counter with one of the offers it has previously sent. Otherwise, the agent will send the highest offer that it has not yet proposed.

Observations on the MiCRO Agent: The MiCRO agent always makes the minimal possible concession and does not attempt to match the intensity of concession displayed by its negotiating partner. The effectiveness of this strategy could depend heavily on the negotiation time allotted.

Experiment Design

Experiment Timelines

The simulated time starts at 0. After both agents produce a response, if no offer has been accepted, the simulated time is increased by a value sampled from a Normal distribution with a mean of 1 and a standard deviation of 0.2. If no offer is accepted before $t = 500$, negotiation is terminated with zero utility for both agents. We expected that each agent would receive a higher average utility when there was more time available to negotiate because more mutually favorable offers can be explored. We predicted that agents that eventually concede heavily (Boulware agents and Conceder agents) would be especially more likely to achieve mutually beneficial offers. We tested several marketplaces with different negotiation time limits of 10, 250, 500, and 1000.

Agent Distributions

We evaluated the relative performances of the identified strategies on some likely and interesting marketplaces:

Even Distribution: The four strategy types have an equal chance of being selected for each negotiation.

Majority Type: In each “majority” market, one of the four strategy types is instantiated with a probability of 0.55, while the other three are selected with a probability of 0.15 each. These configurations help us understand how being in the majority or the minority affects the average utility of different strategies. Later, we suggest a model to generalize varying the proportion of each strategy.

Experiment Domains

Each marketplace simulation sampled uniformly from 50 scenarios, with domains generated as in ANAC 2022. Issue weights are selected using an even Dirichlet distribution. The value utilities are selected using an even Dirichlet distribution scaled linearly to the range of [0,1].

Domains can have different bid space sizes. We were interested in the effect that the bid space size has on negotiations. We expected that the smaller domains would result in a lower utility for the Boulware, Conceder, and MiCRO agents, since there are fewer good compromises available. The default target for the size of the bid space is chosen uniformly in the range of [200, 10000], which corresponds to the distribution of the domains used in the ANAC 2022 competition. We developed two other sets of 50 scenarios, a small one with bid spaces in a target range of [50, 200], and a large one with a target range of [10000, 30000]. A domain is generated with a bid space size within $\pm 10\%$ of the target.

As default profiles are generated identically, no profile has an advantage. This is often not true for real-life marketplaces. Hence, we designed a new collection of lopsided domains. The utilities of one of the profiles are modified so that nonzero value utilities are improved using the formula $u_{new} = 0.5 + 0.5 * u_{old}$. We expect that the agents that concede heavily will improve the most in these domains.

Experiment Parameters

For each marketplace scenario, 3000 negotiations were run. We sample Boulware E values from a gamma distribution with the shape parameter $k = 4$ and the scale parameter $\theta = 0.05$. We sample Conceder E values from a gamma distribution with shape parameter $k = 4$ and scale parameter $\theta = 0.5$, and then add one to ensure $E > 1$.

Experimental Results and Discussion

Time

The first marketplace measures the effect of allotted negotiation time on average agent utilities. We predicted that allowing agents more time during a negotiation would lead to more offers being exchanged and thus a higher agreement rate. Experiments results, shown in Figure 3A, confirmed our suspicions. Figure 3B shows a consistent increase in average utility directly related to an increase in negotiation agreement rate. For the remainder of our experiments, we

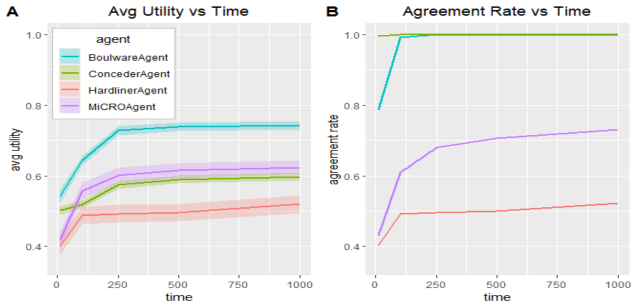

Figure 3: Average utility and agreement rate as a function of negotiation time (default domain, even class distribution).

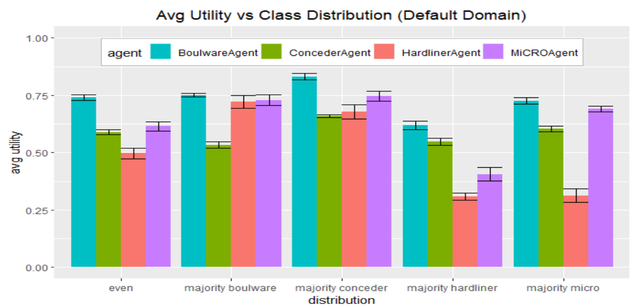

Figure 4: Average utility as a function of different class distributions using only the default domain.

set the time parameter to 500.0 to balance simulation speed and agent performance.

Market Configurations

Agent Distribution

Even: When deciding how to analyze the marketplace, we started with default parameters and minimal variation to start. Thus, we initially simulated a marketplace with an even distribution of the four agent classes on the default domain; this way, we could see which agent or agents would stand out on an even playing field.

We compared negotiation strategies in an even distributed domain space with respect to their utility scores. ANOVA test results show the difference between average class utility is statistically significant ($F = 117.8$ and $p < 2.10^{-16}$). Tukey's HSD test shows that the only non-significant difference is between MiCRO and Conceder ($p = 0.22$).

Majority Type: We next bolstered a single agent type such that they held a majority share of the marketplace, keeping the rest of the agents equally represented. We chose a proportion of 0.55 for the majority type and 0.15 for the minority type. The average utility results for both the equal and majority distributions are shown in Figure 4. Boulwarism significantly outperformed all other strategies in every scenario *except* when Boulware agents are in the majority.

We conducted ANOVA tests for these five class distributions and observed they are statistically significant in all different domains (Figure 4). The p value for these distributions is less than 2.10^{-16}, and the F values for majority Boulware, Conceder, Hardliner, and MiCRO distributions are 141.9, 93.81, 188.9, and 340.9, respectively. Figure 4 shows that the Boulware class has the highest average utility

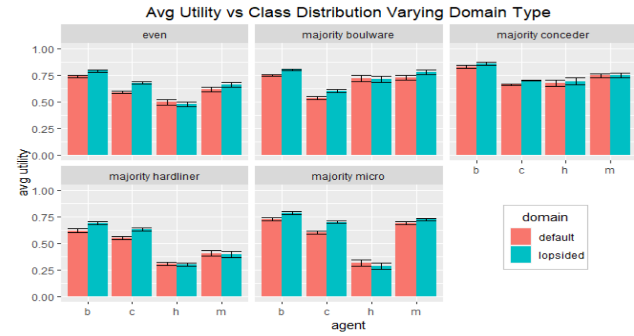

Figure 5: Average utility as a function of class distribution and lopsidedness of the domain.

for all the distributions. To see the statistical significance of the class differences, we compared each pair of negotiation classes for different distributions with Tukey's HSD test. For the majority Boulware distribution, the only class pairs without a significant difference are Boulware & MiCRO ($p = 0.2$) and Hardliner & MiCRO ($p = 0.95$). When the dominant class is Conceder, only the Conceder and Hardliner classes are not significantly different ($p = 0.36$). All other differences are statistically significant.

Domain Distribution

Lopsidedness: We were interested in the effect of asymmetric scenarios where one profile has a significant advantage. We used the set of 50 lopsided scenarios with one profile having increased utility values. With this experiment, we hoped to show which agents perform better in an asymmetric negotiation environment.

In Figure 5, we compare the results of the experiment using the lopsided domains to the experiment using the default domains. These tests include the even distribution and the majority distributions. In every scenario, our expectations that the Boulware and Conceder agents would achieve improved utility were met. This is because these agents concede significantly from their best bid and the only bids that are improved in a lopsided domain are the suboptimal bids. The Hardliner and MiCRO agents did not see significant change. The MiCRO agent concedes only a small amount and the Hardliner agent does not concede at all. The lopsided domains only benefit agents which concede. Since none of these agents is more likely to eventually reach an agreement with the Hardliner based upon the fact that their available utilities are more favorable, the Hardliner does not benefit from the lopsided domains. The MiCRO agent may be more likely to reach a better early agreement, but it appears that this difference is not always significant.

Domain sizes: Agents may perform differently based on the size (number of bids) of the domain. In Figure 6, we compare the average utility of the agents in the even distribution and the majority distributions across the three domain sizes: small, default, and large. The significant differences in the Boulware and Conceder agents suggest that these agents receive less utility in large negotiations. Hardliner and MiCRO agents do not experience any significant changes associated with the size of the domain. This result is unexpected

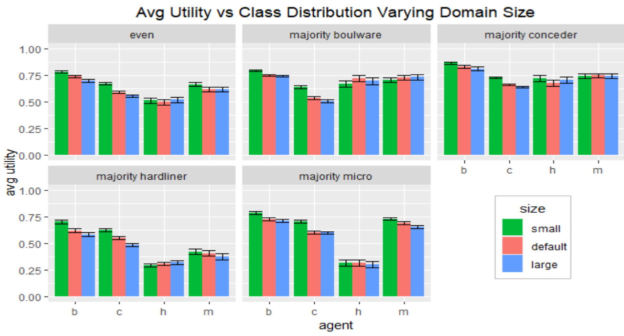

Figure 6: Average utility as a function of class distribution and domain size.

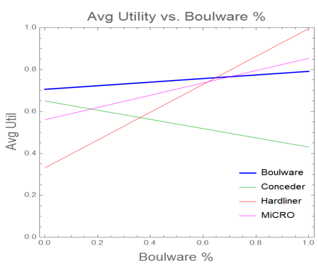

(a) Varying Boulware %

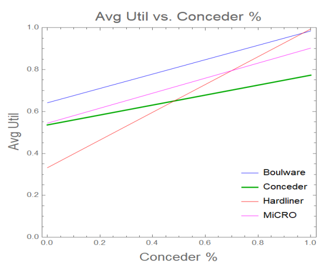

(b) Varying Conceder %

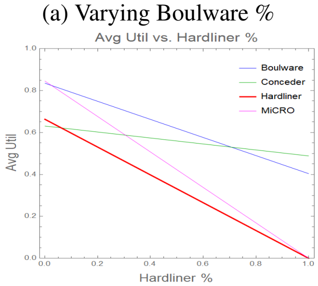

(c) Varying Hardliner %

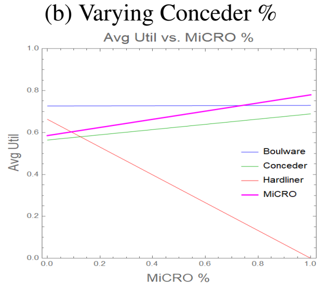

(d) Varying MiCRO %

Figure 7: Effects on average utility when varying one class proportion and keeping all others even.

since we predicted that large domains would provide more opportunities for mutually beneficial outcomes. One likely cause for this trend is that larger domains have fewer or more sparse strong offers that provide high utility for both agents. This means that only two compromising agents can benefit from the potentially strong compromises, while the uncompromising agents are likely to reach worse outcomes.

Head-To-Head Matrix Model

Even after many hours of collecting these simulated data, these findings are still limited in scope, and we had to make some arbitrary decisions about the distributions chosen. Hence, we decided to do a formal analysis to support and extend our findings: using a head-to-head matrix. This matrix will contain the average performance of each strategy against every other strategy, including itself. The values for the matrix were found by simulating 100 head-to-head negotiations for each agent pair using the default domain. To find a class performance given a specific marketplace distribution, the head-to-head matrix will be multiplied by a column

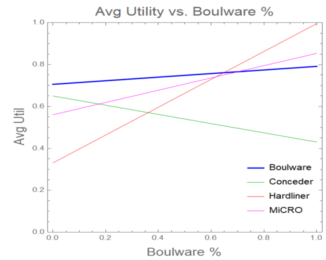

(a) All classes

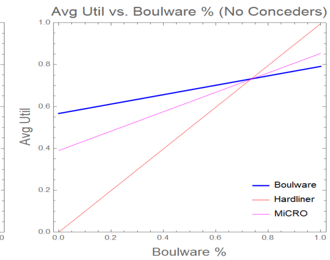

(b) No Conceders

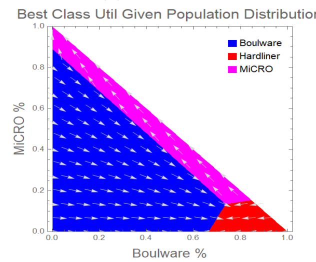

(c) Population shifts

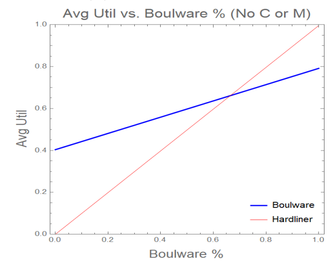

(d) No Conceders, MiCROs

Figure 8: Analysis of potential Boulware equilibrium.

vector containing the four class proportions. The resulting row vector will contain the average utility of each of the four classes against the specified marketplace distribution.

We start by varying one class while keeping all others even, as presented in Figure 7, where the bold line represents the varied class. The data in Figure 7 is validated by matching the data in Figure 4 in the simulated distributions.

According to Figure 7a, Boulwarism should increase its share in the market until around 60-65%, where a potential equilibrium point may occur. However, this does refute the possibility of an all-Boulware market equilibrium, as both Hardliners and MiCROs entering an all-Boulware market will perform better. Figure 7b shows that the Conceder class never reaches domination and hence it would theoretically fall out of contention in this kind of marketplace. However, as the Conceder proportion increases, so too does overall utility in the market, which has positive social welfare implications. Unsurprisingly, an increase in the Hardliner population leads to a decrease in overall utility in the market, according to Figure 7c. Also, Hardliners perform the worst as their share of the market increases, which means they will likely be eliminated from the population. Hardliners act like a “predator” class, so they starve as the “prey” population declines. Thus, an all-Hardliner marketplace is untenable under these conditions. Figure 7d reveals that MiCRO is the dominant strategy when it saturates the market. This implies that a pure MiCRO strategy market should remain stable as long as its share of the population remains sufficiently large. This explains why the strategy, which is proposed to be “optimal” by its creator, performed in the middle of the pack in the 2022 ANL competition: none of the other agents employed the strategy.

With this preliminary observation, we now consider the two potential equilibria scenarios: Boulwarism with Hardliners and MiCROs and a purely MiCRO marketplace.

First, we look at Boulwarism. When we remove the under-

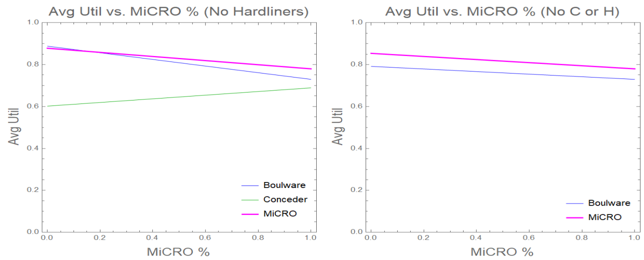

(a) No Hardliners (b) No Conceders, Hardliners

Figure 9: Analysis of potential MiCRO equilibrium.

performing Conceders from the population, as in Figure 8b, we are left with a potential equilibrium point of roughly 74% Boulware, 13% Hardliner, and 13% MiCRO. This also supports the simulated investigation, but this may not be stable.

The top left, bottom left, and bottom right corners of the ternary graph in Figure 8c represent marketplaces that contain only the MiCRO, Hardliner, or Boulware strategy, respectively. The arrows suggest the direction in which the marketplace distribution will shift over time as underperforming strategies leave the market and are replaced with better-performing strategies.

As we can see, the best class for the majority of marketplace distributions is Boulware (the blue region), so the population makeup will shift to the bottom right of the graph toward pure Boulwarism. However, too much Boulwarism will cause Hardliners to reign supreme (the red region), shifting the population back to the left. If at any time the population distribution enters the magenta region, the MiCRO strategy will become the top performer, moving the marketplace up and the to the left. This is still within the magenta region, so a purely MiCRO marketplace is indeed a stable equilibrium. A distribution where the three colored regions meet appears to be an equilibrium, suggested by Figure 8c. However, it is not stable, as any movement upward will lead to a MiCRO-dominated market, and any movement downward will cause an equilibrium between Boulwares and Hardliners. If this bottom equilibrium is reached, Figure 8d shows the point at which it will occur: about two-thirds Boulware and one-third Hardliner. This implies that MiCRO is pushed out of the market completely.

To investigate a possible MiCRO equilibrium suggested by Figures 7d and 8c, we have removed the Hardliner strategy that performs poorly in a market densely populated by MiCRO users. As such, Figure 9a shows that MiCRO is now the top performer in nearly all of the marketplace distributions. We then remove the Conceders, which always take the last place in Figure 9a. The market we are left with in Figure 9b is completely dominated by the MiCRO strategy.

Thus, our analysis shows that a marketplace could either reach a stable class distribution of roughly 67% Boulwarism and 33% Hardlining or it can be dominated by the MiCRO strategy entirely. This second outcome would not have been discovered using only the simulated data, so it is crucial that we were able to verify our findings and extend the analysis using the head-to-head matrix. We must note that we have not assumed these marketplaces can evolve in the traditional sense. Therefore, these "equilibria" only represent *likely* marketplace distributions resulting from many population shifts where agents freely enter and exit based on how well their negotiation strategy worked for them.

Conclusions & Future Work

We identified four negotiation strategy types commonly used in structured, closed tournaments involving repeated bilateral, multi-issue negotiations between only a few players. We seek to understand the relative efficacy of these algorithms in large, open markets where one is unlikely to meet the same partner again. This constraint negates any advantage of learning approaches that first build partner strategy models from interaction history and then leverage that to gain a competitive advantage. We assumed that agents cannot access or adopt partner behaviors and thus evolutionary dynamics of agent strategy populations or *evolutionary* equilibria were not the focus of our investigation. Rather, we wanted to see the relative performance of widely used negotiation strategy types in representative market configurations to provide prescriptions of negotiation strategies to adopt when deploying agents in such open markets.

We derive these conclusions from experiments simulating distributions of likely strategies in large, open markets:

- The Boulware strategy outperforms the other three common strategies for a large variety of market configurations.
- For default domains, the MiCRO strategy obtains the highest utility in homogeneous populations and also dominates Boulwarism in a mixed population of two strategies. However, MiCRO agents perform poorly when other types of agent are present in the market (Hardliners in particular). This explains MiCRO's lackluster results in ANL 2022.
- A surprising result is that, in niche situations, the Conceder agents outperformed the Hardliner agents. This suggests that while it might be tempting to use a Hardliner strategy (likely the easiest strategy to implement) in the marketplace, playing nice, i.e., conceding, may be a preferable option!

It would be useful to study negotiation approaches that learn over a negotiation round. A Boulware agent, for example, can learn that the partner is conceding and hence slow its concession rate. Such adaptation occurs during a single negotiation and is feasible in a large and open market. It would be instructive to evaluate the relative performance of "smart" agents in a mix with the other strategy distributions.

It can also be useful to develop a "population-adaptive" MCT agent that adapts its E value using its interaction history-based estimate of the strategy type distribution in the market. Such an adaptive agent should be tested to compare its performance with that of the most competitive fixed agent strategy across a diverse set of market configurations.

Our simulations use variable E values to represent realistic marketplaces where participants choose that constant based on local bias, knowledge, and preferences. Theoretical analysis of distributions of strategies with fixed behaviors (E values) and domain types is feasible. Comparisons of the approximations of such theoretical analyses with our experimental results can provide coarse-grained recommendations over a larger space of agent distributions but for specific domain types.

References

[Aydoğan et al. 2019] Aydoğan, R.; Fujita, K.; Baarslag, T.; Jonker, C. M.; and Ito, T. 2019. Anac 2018: Repeated multilateral negotiation league. In *Annual Conference of the Japanese Society for Artificial Intelligence*, 77–89. Springer.

[Aydoğan et al. 2020] Aydoğan, R.; Baarslag, T.; Fujita, K.; Mell, J.; Gratch, J.; Jonge, D. d.; Mohammad, Y.; Nakadai, S.; Morinaga, S.; Osawa, H.; et al. 2020. Challenges and main results of the automated negotiating agents competition (anac) 2019. In *Multi-agent systems and agreement technologies*. Springer. 366–381.

[Baarslag et al. 2012] Baarslag, T.; Hindriks, K.; Jonker, C.; Kraus, S.; and Lin, R. 2012. The first automated negotiating agents competition (anac 2010). In *New Trends in agent-based complex automated negotiations*. Springer. 113–135.

[Baarslag et al. 2013] Baarslag, T.; Fujita, K.; Gerding, E. H.; Hindriks, K.; Ito, T.; Jennings, N. R.; Jonker, C.; Kraus, S.; Lin, R.; Robu, V.; et al. 2013. Evaluating practical negotiating agents: Results and analysis of the 2011 international competition. *Artificial Intelligence* 198:73–103.

[Baarslag et al. 2016] Baarslag, T.; Hendrikx, M. J.; Hindriks, K. V.; and Jonker, C. M. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems* 30(5):849–898.

[Baarslag et al. 2017] Baarslag, T.; Kaisers, M.; Gerding, E.; Jonker, C. M.; and Gratch, J. 2017. When will negotiation agents be able to represent us? the challenges and opportunities for autonomous negotiators. International Joint Conferences on Artificial Intelligence.

[Bagga, Paoletti, and Stathis 2022] Bagga, P.; Paoletti, N.; and Stathis, K. 2022. Deep learnable strategy templates for multi-issue bilateral negotiation. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 1533–1535.

[Bloembergen et al. 2015] Bloembergen, D.; Tuyls, K.; Hennes, D.; and Kaisers, M. 2015. Evolutionary dynamics of multi-agent learning: A survey. *Journal of Artificial Intelligence Research* 53:659–697.

[de Jong 2022] de Jong, D. 2022. An analysis of the linear bilateral anac domains using the micro benchmark strategy. 223–229. International Joint Conferences on Artificial Intelligence.

[Hines 1987] Hines, W. 1987. Evolutionary stable strategies: a review of basic theory. *Theoretical Population Biology* 31(2):195–272.

[Jonge et al. 2018] Jonge, D. d.; Baarslag, T.; Aydoğan, R.; Jonker, C.; Fujita, K.; and Ito, T. 2018. The challenge of negotiation in the game of diplomacy. In *International Conference on Agreement Technologies*, 100–114. Springer.

[Lin et al. 2014] Lin, R.; Gal, Y.; Kraus, S.; and Mazliah, Y. 2014. Training with automated agents improves peoples behavior in negotiation and coordination tasks. *Decision Support Systems (DSS)* 60(1–9).

[Maaravi, Idan, and Hochman 2019] Maaravi, Y.; Idan, O.; and Hochman, G. 2019. And sympathy is what we need my friend—polite requests improve negotiation results. *PLOS ONE* 14(3):1–22.

[Mccalley et al. 2020] Mccalley, J.; Zhang, Z.; Vishwanathan, V.; and Honavar, V. 2020. Chapter 4: Multiagent negotiation models for power system applications.

[Mell et al. 2018] Mell, J.; Gratch, J.; Baarslag, T.; Aydoğan, R.; and Jonker, C. M. 2018. Results of the first annual human-agent league of the automated negotiating agents competition. In *Proceedings of the 18th International Conference on Intelligent Virtual Agents*, 23–28.

[Mirzayi, Taghiyareh, and Nassiri-Mofakham 2022] Mirzayi, S.; Taghiyareh, F.; and Nassiri-Mofakham, F. 2022. An opponent-adaptive strategy to increase utility and fairness in agents' negotiation. *Applied Intelligence* 52(4):3587–3603.

[Morales et al. 2018] Morales, J.; Wooldridge, M.; Rodríguez-Aguilar, J. A.; and López-Sánchez, M. 2018. Off-line synthesis of evolutionarily stable normative systems. *Autonomous agents and multi-agent systems* 32(5):635–671.

[Mori and Ito 2016] Mori, A., and Ito, T. 2016. Automated negotiating agent based on evolutionary stable strategies. In *Computer and Information Science 2015*. Springer. 27–41.

[Peled, Gal, and Kraus 2013] Peled, N.; Gal, Y. K.; and Kraus, S. 2013. An agent design for repeated negotiation and information revelation with people. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, 789–795. Association for the Advancement of Artificial Intelligence.

[Ransbotham et al. 2017] Ransbotham, S.; Kiron, D.; Gerbert, P.; and Reeves, M. 2017. Reshaping business with artificial intelligence: Closing the gap between ambition and action. *MIT Sloan Management Review* 59(1).

[Rosenfeld et al. 2015] Rosenfeld, A.; Zuckerman, I.; Segal-Halevi, E.; Drein, O.; and Kraus, S. 2015. Negochat-a: a chat-based negotiation agent with bounded rationality. *Autonomous Agents and Multi-Agent Systems* 30(1):60–81.

[Sengupta, Mohammad, and Nakadai 2021] Sengupta, A.; Mohammad, Y.; and Nakadai, S. 2021. An autonomous negotiating agent framework with reinforcement learning based strategies and adaptive strategy switching mechanism. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 1163–1172.

[Tung 2019] Tung, K. 2019. Ai, the internet of legal things, and lawyers. *Journal of Management Analytics* 6(4):390–403.

[Xu et al. 2020] Xu, B.; Hale, J. A.; Pritchard, S.; and Sen, S. 2020. An application of the infinite framework in a human-agent negotiation competition. In *Proceedings of the 8th International Conference on Human-Agent Interaction*, 32–40.

[Zohar 2015] Zohar, I. 2015. "the art of negotiation" leadership skills required for negotiation in time of crisis. *Procedia - Social and Behavioral Sciences* 209:540 – 548. The 3rd International Conference.