Tackling the Protocol Problem in Automated Negotiation

Blue Sky Ideas Track

Yasser Mohammad
NEC CORPORATION
Tokyo, Japan
y.mohammad@nec.com

ABSTRACT

Automated Negotiation (AN) is a research field with roots extending back to the mid-twentieth century. There are two dominant AN research directions pursued by the AAMAS community in recent years: (1) designing new heuristic or ML/RL/MARL-based strategies for the simplest bargaining mechanism called the Alternating Offers Protocol (AOP) and its extensions and (2) defining new mediated mechanisms that require a trusted third party. Intelligence lies in the strategy in the first direction and the mechanism in the latter. Either way, evaluation is almost always conducted in terms of empirical evaluation in some chosen set of negotiation scenarios. This paper argues for more efforts towards tackling the problem of *protocol design* in automated negotiation more rigorously by integrating ideas from mechanism-design literature. This requires, as a first step, a common language for expressing different negotiation protocols and strategies. We provide such a language which can represent a wide variety of negotiation protocols (both mediated and unmediated). We briefly outline our early effort in using this approach to provide a novel protocol with a provable Perfect Bayesian Equilibrium strategy that is also empirically effective.

KEYWORDS

Automated Negotiation, Bargaining, Mechanism Design

ACM Reference Format:
Yasser Mohammad. 2025. Tackling the Protocol Problem in Automated Negotiation: Blue Sky Ideas Track. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 5 pages.

1 INTRODUCTION

Automated negotiation is a process by which *self-interested* agents try to achieve an *agreement* that is at least as good as no agreement for all agents involved. Automated negotiation is receiving more interest from the research community recently due to the accelerated pace at which businesses employ intelligent agents to manage different production and business processes and the need for these agents to work together as self-interested agents representing their institutions. Recent applications of automated negotiation include permission management in IoT systems, Wi-Fi channel assignment [11, 26], agriculture supply chain support, supply chain management [35], vehicle routing [9], path planning [14], and providing feedback for student negotiation skills [16].

Automated negotiation has a long history as a research discipline dating back to Nash's solution of the bargaining game [38]. Nash's Bargaining game was a single round protocol. It was later extended by finding a perfect equilibrium for a more realistic time-limited bargaining protocol [39]. The preferences of all players was assumed to be common knowledge in that work.

One of the most commonly used automated negotiation protocols that assume private preferences is the Alternating Offers Protocol (AOP) with its multilateral extension (the Stacked Alternating Offers Protocol — SAOP by Aydoğan et al. [1]). Several strategies have been proposed for negotiation under this mechanism [5]. These methods can be classified into exact solutions of simplified problems [4] and heuristic solutions to the most general case [18]. Most recently, MiCRO was proposed for AOP guaranteeing that agreements – if found – are Pareto-optimal [8].

Alternatives to the AOP mechanism have been proposed that are mostly mediated. A commonly used example is the Single Text (ST) negotiation mechanism in which a third-party mediates the interaction between agents by offering a tentative agreement for them and updating it based on their response [22, 23]. Several alternatives for the mediation algorithm have been proposed [15, 21, 22, 28]. In this work, we are mostly concerned with unmediated negotiation protocols because they do not rely on a trusted-third party. We will comment on the extension to mediated protocols at the end.

This paper argues for a research program that tries to develop automated negotiation protocols with known (and hopefully simple) dominant strategies and equilibria, shifting the intelligence to the protocol while providing rigorous tools for analyzing these protocols and their strategies. This extends recently published works that attempt to design an equilibrium strategy for AOP for some discrete outcome spaces [8] provide metrics for evaluating negotiation protocols [30] and generalizing AOP to a larger set of unmediated protocols [34].

2 THE NEGOTIATION PROBLEM

A negotiation is an interaction between two or more self-interested agents trying to reach an agreement beneficial to at least one of them. As a simple example, consider a procurement manager in a company trying to negotiate the quantity, price and delivery dates of some material with one supplier. In this case, the scenario is defined by the possible set of agreements (called the *outcome space*) which is all the possible combinations of values for the three negotiation *issues* (e.g. 10 items tomorrow at 1$ each is an *outcome*),

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA.

and the information the procurement manager and the supplier know about their own and their partner's preferences regarding different possible agreements. For example, the procurement manager knows that she prefers lower prices while the supplier prefers higher prices. Moreover, she may know that this specific supplier (based on earlier interactions or knowledge of their production capacity) prefers higher quantities if the delivery date is late but lower quantities otherwise. Note that the procurement manager may care about increasing the supplier's profit to ensure future supplies if that is achievable at no or low cost for itself. Such factors must be included in the preferences of the procurement manager. Other information that may be available to the negotiator (procurement manager in our case) is information about the partner(s) strategy. For example, she may know that the supplier is stubborn and does not change their first offer that much. This information needs to be considered when the procurement manager decides what to do (i.e. the negotiation *strategy*).

The *protocol* in our example is the set of rules concerning the actions allowed by the procurement manager and the supplier during their interaction. Examples may include time pressure (agreement must be reached within two days), timing constraints (no party can change its standing offer except after receiving a reply from her partner), modalities (email but not phone), etc. The protocol simply defines who can do what when.

Most research in unmediated automated negotiation targets the problem of designing an effective *strategy* for a predefined negotiation *protocol* on a possible set of negotiation *scenarios*. In our running example, how should the procurement manager (or supplier or both) behave during the negotiation [3, 13, 24, 33]. Since Nash's pioneering work [38], the negotiation protocol in most of this research is a variation of the simplest possible human negotiation protocol: *haggling*. One negotiator starts by offering an outcome which is then either accepted or rejected by the partner who can now counter it with her own offer. Any negotiator can leave the negotiation at any time. Variations include extension to multilateral negotiation [1], restricting offers to be non-repeating [4, 33], allowing any-time offering [10], extension to multiple concurrent (or interleaved) negotiations [2, 29, 40], allowing for user's preferences during the negotiation [7, 37], etc. Even in these cases, the research focused on the negotiation *strategy* under one of these variants with little work trying to compare these variants or discover new unmediated protocols with desirable features. In summary, most research in automated negotiation is focusing in *strategy design*. There is a research gap in the area of *protocol design*.

The appeal of the alternating offers protocol (AOP) is a contributing factor to this situation. Firstly, the protocol is easy to understand and follows closely human's negotiation behavior. Secondly, it is trivial to know when the negotiation session ends and whether an agreement is reached or not. The protocol is memoryless in this regard as it requires no information about previous actions given the last action to determine agreement of failure. This may have contributed in making haggling an effective negotiation strategy for people. Thirdly, there is a wealth of existing research and strategies against which new proposals can be evaluated and contrasted. Finally, there are no known strategies that provide the crucial combination of completeness, Pareto-optimality, fairness (e.g. Nash, Kalai, or Kalai-Smorodonisky) while constituting an equilibrium of the protocol [31]. One of the few exception is the MiCRO strategy [8] which was recently shown to guarantee completeness, Pareto-optimality and a sense of fairness for *balanced* negotiation scenarios (i.e. ones allowing for agreement through almost the same number of concessions) [17].

The last *appealing* point for AOP for the researcher is its main weakness as a protocol from the social and business point of view. It leads to the current situation of having hundreds of negotiation strategies (e.g. the strategies submitted to the Automated Negotiating Agents Competition [6] which is running yearly since 2010 and available from platforms like GENIUS [25] and NegMAS [36]) but no *guarantees* on performance. AOP has some structural problems as well. For example, it encourages time-wasting with some evaluations citing more than 75% offer repetition rate [32]. This is not surprising giving that an agent that could successfully waste negotiation time will likely to get more concessions from its partner. Moreover, the decision to offer an outcome comes with the possibility of it becoming the agreement which may lead to insufficient outcome-space exploration. Allowing the offerer of the agreement-to-be a last-moment veto [40] can reduce this problem but does not eliminate it (at least we do not no of any way to rigorously show that it does).

Nevertheless, we know from research in auctions (a related agreement technology) that the *best* protocol may not be the obvious one. For example, second-price auctions are known to be Pareto-optimal and incentive compatible while most auction formats traditionally used by people like ascending auctions and descending auctions are not. We argue that it is time for the automated negotiation research community to search for a similar kind of an unmediated negotiation protocol that transfers the intelligence from the strategy to the protocol.

The first step towards achieving this goal is to have a common mathematical formalism that can represent a wide variety of unmediated negotiation protocols. We can then systematically search this space for negotiation protocols that can achieve the designers objectives (e.g. fairness, welfare maximization, Pareto-optimality, completeness, etc).

3 FORMAL DEFINITION OF A NEGOTIATION

Formally, a negotiation *scenario* s is defined as a tuple $(\mathcal{A}^s, \mathcal{D}^s, \mathcal{I}^s)$ where $\mathcal{A}^s$ is the set of agents numbered from 1 to n_A, $\mathcal{D}^s$ is the negotiation domain and $\mathcal{I}^s$ is a tuple defining information available to each agent. The negotiation domain $\mathcal{D}^s \equiv (\Omega^s, \mathcal{F}^s)$ defines (1) The outcome space (Ω^s of size n_o) comprising all possible agreements. A special outcome $\phi \notin \Omega^s$ is always assumed to exist to represent disagreement and we define the *extended outcome space* Ω^+ as the $\Omega^s \bigcup \{\phi\}$. (2) Agent preferences $\mathcal{F}^s$ which can either be ordinal[1] ($\succcurlyeq$) defining an ordering[2] over Ω^+ – per agent – or cardinal (u) defining a mapping per agent from Ω^+ to $\Re$ with higher values indicating better outcomes[3]. We use $\succcurlyeq_x$ and u_x to refer to the preferences at index x. The *Information Set Tuple* $\mathcal{I}^s$ is a tuple of n_A information sets $\left(\left\{\mathrm{I}^s_x(\mathcal{F}^s) : x \in \mathcal{A}^s\right\}\right)$ where $\mathrm{I}^s_x(\mathcal{F}^s)$ represents

[1]We assume standard transitivity on preferences.
[2]$a \succcurlyeq b$ means that a is not worse than b. Symbols $\succ, \approx$ are defined accordingly.
[3]If the codomain of u within the range $[0, 1]$, the utility function is called normalized. Time-pressure can be modeled here by a discounting factor as in Rubinstein [39].

all the information available to agent x about the preferences $\mathcal{F}^s$ of all agents including itself

A negotiation *strategy* describes the behavior of an agent in accordance of a protocol $\mathcal{P}$ for the set of scenarios it was designed to handle. Its input is the tuple $(n_A, x, \Omega, \mathrm{I}_x)$, where x is the agent index controlled by the strategy.

Given a negotiation scenario s and a strategy profile π, the result of all agents $\mathcal{A}^s$ following their assigned strategies selects a member of the extended outcome space called the *negotiation outcome* $\omega_*^s \in \Omega^+$. We indicate executing the negotiation on s using strategy profile π resulting on ω_*^s as a negotiation outcome by $s(\pi) = \omega_*^s$.

Because the strategy running agent x has access only to (Ω, I_x^s), while the full game induced by the scenario (G^s) is defined by $(\Omega, \mathcal{F}^s)$, it faces a Game with Incomplete Information (GII) in which nature moves first assigning agent and partner(s) types (defined as the Ω, and preferences $\mathcal{F}^s$ sampled according to $\mathcal{I}^s$). Nature never plays again. The size of this GII is huge (and in most cases infinite). No wonder we only have heuristic approaches except for the case with complete information ($\mathrm{I}_x^s = \mathcal{F}_x^s$) [12, 38, 39]. Most recent research in automated negotiation assumes that I_x^s is simply the agent's own preferences($\mathcal{F}_x^s$) or complete ($\mathcal{F}^s$) but in most realistic situations, agents have some partial information about partner preferences.

3.1 Evaluating Negotiations

Evaluating negotiation protocols is multifaceted. [30] proposed using a two dimensional criteria: Designer and Agent scores. The agent score is defined as a linear combination of the expected **advantage** (received utility - reserved value) of the agent and **privacy** which is inversely proportional to the amount of information about agent's preferences revealed to the partner(s) by the end of the negotiation. The designer score is defined as the product of rationality, completeness, optimality, fairness and welfare. **Rationality** is the fraction of negotiations ending in rational outcomes including disagreement. **Completeness** is the fraction of negotiations with non-empty win-win deals ending in agreement. **Optimality** is inversely proportional to the distance between the negotiation outcome and the Pareto Outcome Set on the utility space[4]. **Fairness** is inversely proportional to the distance between the negotiation outcome and the *nearest* Bargaining Solution (i.e. Kalai [19], Nash [38], Kalai-Smorodinsky [20]) or an ordinal version of this solution that utilizes outcome ranks instead of utilities [31]. **Welfare** is the sum of all agent advantages. All measures are normalized to have a maximum value of one. A protocol is *exactly* rational if its rationality is one. The same applies for all other performance measures.

4 PROPOSED RESEARCH PROGRAM

Now that we know how to evaluate negotiation protocols, we can define the proposed research program. Firstly we need to find a mathematical framework that can be used to define generalizations of AOP that keep its core advantages (e.g. no trusted third party, easy and efficient implementation, distributed execution, clear definition of termination) while allowing the designer to generate new protocols and examine them mathematically.

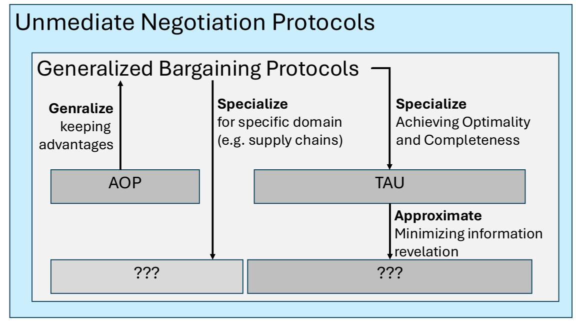

Figure 1: The proposed research program

One possible goal of this research program is to discover a *single* negotiation protocol with a unique Perfect Bayesian Equilibrium (PBE) strategy profile that maximizes all designer metrics while maximizing all agent metrics in all negotiation scenarios. The existence of such a protocol is highly unlikely but proving this is an interesting research question in itself.

A more realistic goal is to find a set of protocols for sets of negotiation scenarios for which a known strategy profile provides is a PBE that maximizes the designer score. We call a protocol that achieves this for a given set of scenarios a Outcome-Perfect Negotiation Mechanism for these scenarios. The outline of this program is shown in Fig. 1. Even in 2014, such a goal was recognized as an important direction for automated negotiation research [27] but was never achieved. We believe that with the recent advances in protocol evaluation and representation outlined in this paper, the research community can achieve this goal providing a major breakthrough for the field and opening the door for new innovations.

Looking again at mechanism design research in auction theory, we can see that a crucial factor in the success of this discipline is the availability of a common mathematical framework for expressing auction mechanisms. An auction strategy is fully specified by the bidding strategy of the agent which is a function mapping value (utility) to a bid. An auction mechanism is fully specified by two rules: 1) an *allocation rule* which determines who gets the item(s) being auctioned. It is a function that maps the bids of the participants to a decision on the winner(s). 2) a *payment rule* that determines how much each participant pays. It maps the bids of the participants to a payment amount for each bidder, including the winner(s) and losers.

For the proposed program of protocol design, we rely on the recently proposed Generalized Bargaining Protocols (GBP) [32].

Two common features of all GBPs are 1) All offers originate from the agents. 2) agents can leave the negotiation anytime preventing agreement (ensuring exact rationality).

A GBP negotiation between n_A agents is carried out in n_T negotiation threads. Each negotiation thread is assigned to exactly one agent called its owner. One or more agents are assigned as responders on each thread to respond to offers from the thread

[4]The utility space is an N dimensional space for a negotiation with N agents in which each dimension represents the value (utility) for one agent upon which all outcomes are projected as points

owner. Associated with each thread is an ordered tuple of outcomes called the Tentative Agreement Tuple for this thread (Ω_x^*) which is initialized by the empty tuple $\langle\rangle$ and is used to represent potential agreements offered by the thread owner.

A GBP is defined as a tuple $(N_\pi, N_\rho, \alpha, \rho, \chi, f, \gamma, \sigma)$:

Offer, Response Cardinality $N_\pi, N_\rho \in \mathbb{Z}^+ \cup \infty$ The maximum allowed cardinality of any offer Ω_x (response Ω_{yx}). These are usually set to 1.

Assignment Rule $\alpha : \mathbb{T} \to \mathcal{A}$ Defines the negotiation threads ($\mathbb{T}$) and selects exactly one agents to own each thread.

Responders Rule $\rho : \mathbb{T} \to \mathbb{P}(\mathcal{A})$ Selects one or more agents to respond to offers in every thread.

Activation Rule $\chi : \mathbb{T} \to \mathbb{P}(\mathbb{T})$ Selects one or more threads to activate (i.e. run the owner's offering policy).

Filtering Rule $f : \Omega \to \mathbb{P}(\Omega)^{n_T}$ Filters out invalid outcomes for the current state of the mechanism from the outcome space leading to the Valid Outcome Space for each agent Ω_x^v for thread x.

Update Rule γ Reads all offers and responses and updates the Tentative Agreement Tuple (TAT) for each negotiation thread x (Ω_x^*). When $\left|\Omega_x^*\right| \leq 1$, we use ω_x^* to indicate its single element or ϕ if Ω_x^* is empty.

Evaluation Rule $\sigma : (\mathbb{P}(\Omega))^{n_T} \to \mathbb{P}(\Omega) \cup \{\triangleright\}$ Reads all Tentative Agreement Tuples and returns one of two decisions: 1) an outcome-set ending the negotiation or 2) The value $\triangleright$ which activates the Activation Rule to select the next set of Offering Policies to activate.

The first two rules (assignment rule, responders rule) define a graph connecting agents allowing them to exchange offers. The activation rule defines the timing of the protocol (which agent is allowed to send offers at any point of time).

Despite its apparent complexity (two parameters and seven rules), in most cases the mechanism designer can focus almost completely on the filtering, update, and evaluation rules because the remaining three rules and parameters are mostly fixed by the negotiation context. Nevertheless, having the expressiveness of the full GBP framework is important in modeling different negotiation contexts (e.g. negotiations designed to reduce the search space rather than reaching an agreement, concurrent negotiations).

A GBP agent strategy is thus defined by two components (very similar to AOP strategies):

Offering policy $\pi_x : \mathbb{P}(\Omega) \to \mathbb{P}^{N_\pi}(\Omega)$ Receives a set of outcomes (the Valid Outcome Space) $\Omega_x^v \subseteq \Omega$ from which it selects a sub-set Ω_x (called an offer) to send through a thread it owns. Offering the empty set {} is interpreted as ending the negotiation immediately leading to disagreement. The cardinality of any offer must be less than the N_π parameter of the mechanism.

Selection policy $\rho_x : \mathbb{P}^{N_\pi}(\Omega) \to \mathbb{P}^{N_\rho}(\Omega)$ Receives an offer (Ω_y) on a given thread y and returns a subset of it as a response $\Omega_{yx} \subseteq \Omega_y$. The empty set {} is interpreted as rejection. The cardinality of any response must be less than the N_ρ parameter of the mechanism.

An attractive feature of the GBP formulation is that all rules and policies are defined as simple set operations. This simplifies the process of analysis.

It is possible to extend the GBP framework to single-text mediated protocols (e.g. [22, 23]) by simply moving the offering policies *inside* the mediator and define the agent strategy by just the selection policy. Extension to other unmediated protocols is possible by extending the offering and selection rules to assign a value from a predefined domain to each member of the outcome-space. These extensions are beyond the scope of this paper.

4.1 Tentative Agreements Unique Offers (TAU)

TAU (TAU) is a GBP proposed in [32]. that modifies the filtering rule, update rule and evaluation rule of AOP. Agents can offer in parallel or a round-robin fashion. We provide this as an example of a protocol design research within the proposed framework. We then describe a PBE strategy for this protocol that achieves perfect rationality, completeness, Pareto-optimality, fairness (in an ordinal sense of the Kalai [20]) on a large subset of discrete negotiation strategies with no-information about partner preferences. As the goal of this paper is to introduce the research program and argue for its timeliness and importance, this section provides only a sketch of this protocol and strategy to showcase he possibilities.

In TAU Agents are allowed to repeat offers but once an agent starts repeating an offer, it is forced to continue repeating the same offer forever. Moreover, an agreement is reached only if the same offer was offered by every agent and was selected by every agent. This is a memory-full protocol as evaluating whether an agreement is reached or not requires checking the complete negotiation trace to confirm that the above condition is satisfied. The Wasting Accepting Rational (WAR) is a simple but unobvious strategy for TAU with discrete outcome-spaces. It starts by offering all irrational outcomes (i.e. outcomes worse than disagreement) in random order. Once no more irrational outcomes are available, WAR offers all rational outcomes from best to worst (breaking ties randomly). If there are no more rational outcomes, WAR repeats its last offer. WAR's selection policy extremely simple: select any rational outcome.

It can be shown that TAU is an Outcome-Perfect Negotiation Mechanism (using WAR) for the set bilateral negotiation scenarios of discrete outcome spaces for which no outcome has the exact same utility value as the final agreement for one but not all negotiators. Moreover, TAU(WAR) can be shown to be fairer and faster in practice than state-of-the-art strategies for AOP with the only disadvantage of increased information revelation.

5 CONCLUSION

Our main goal in this work is to advocate for a new research program in automated negotiation mechanism design with the aim of finding new protocols with desired properties for different negotiation conditions. We believe that GBP provides a useful framework in this direction. It is still bargaining in the sense that agents exchange offers, but it is free from some constraints of AOP that are mostly there for humans not automated agents. Even in 2014, such a goal was recognized as an important direction for automated negotiation research [27] but was never achieved. We believe that with the recent advances in protocol evaluation and representation outlined in this paper, the research community can achieve this goal providing a major breakthrough for the field and opening the door for new types of innovation.

REFERENCES

[1] Reyhan Aydoğan, David Festen, Koen V Hindriks, and Catholijn M Jonker. 2017. Alternating offers protocols for multilateral negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, 153–167.

[2] Tim Baarslag. 2024. Multi-deal Negotiation. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*. 2668–2673.

[3] Tim Baarslag and Enrico H Gerding. 2015. Optimal Incremental Preference Elicitation during Negotiation.. In *IJCAI*. 3–9.

[4] Tim Baarslag, Enrico H Gerding, Reyhan Aydogan, and MC Schraefel. 2015. Optimal negotiation decision functions in time-sensitive domains. In *2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, Vol. 2. IEEE, 190–197.

[5] Tim Baarslag, Koen Hindriks, Mark Hendrikx, Alexander Dirkzwager, and Catholijn Jonker. 2014. Decoupling negotiating agents to explore the space of negotiation strategies. In *Novel Insights in Agent-based Complex Automated Negotiation*. Springer, 61–83.

[6] Tim Baarslag, Koen Hindriks, Catholijn Jonker, Sarit Kraus, and Raz Lin. 2012. The first automated negotiating agents competition (ANAC 2010). In *New Trends in agent-based complex automated negotiations*. Springer, 113–135.

[7] Tim Baarslag and Michael Kaisers. 2017. The value of information in automated negotiation: A decision model for eliciting user preferences. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 391–400.

[8] Dave de Jonge. 2022. An Analysis of the Linear Bilateral ANAC Domains Using the MiCRO Benchmark Strategy. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI)*.

[9] Dave de Jonge, Filippo Bistaffa, and Jordi Levy. 2022. Multi-objective vehicle routing with automated negotiation. *Applied Intelligence* (2022), 1–24.

[10] Dave De Jonge and Carles Sierra. 2015. NB 3: a multilateral negotiation algorithm for large, non-linear agreement spaces with limited time. *Autonomous Agents and Multi-Agent Systems* 29, 5 (2015), 896–942.

[11] Enrique De La Hoz, Ivan Marsa-Maestre, Jose Manuel Gimenez-Guzman, David Orden, and Mark Klein. 2017. Multi-agent nonlinear negotiation for Wi-Fi channel assignment. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1035–1043.

[12] Francesco Di Giunta and Nicola Gatti. 2006. Bargaining over multiple issues in finite horizon alternating-offers protocol. *Annals of Mathematics and Artificial Intelligence* 47, 3 (2006), 251–271.

[13] Walaa H. El-Ashmawi, Diaa Salama Abd Elminaam, Ayman M. Nabil, and Esraa Eldesouky. 2020. A chaotic owl search algorithm based bilateral negotiation model. *Ain Shams Engineering Journal* 11, 4 (2020), 1163–1178. https://doi.org/10.1016/j.asej.2020.01.005

[14] Hiroaki Inotsume, Aayush Aggarwal, Ryota Higa, and Shinji Nakadai. 2020. Path negotiation for self-interested multirobot vehicles in shared space. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 11587–11594.

[15] Takayuki Ito, Mark Klein, and Hiromitsu Hattori. 2008. A multi-issue negotiation protocol among agents with nonlinear utility functions. *Multiagent and Grid Systems* 4, 1 (2008), 67–83.

[16] Emmanuel Johnson, Jonathan Gratch, and David DeVault. 2017. Towards An Autonomous Agent that Provides Automated Feedback on Students' Negotiation Skills. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 410–418.

[17] Dave de Jonge. 2024. Theoretical properties of the MiCRO negotiation strategy. *Autonomous Agents and Multi-Agent Systems* 38, 2 (2024), 46.

[18] Catholijn M Jonker, Reyhan Aydogan, Tim Baarslag, Katsuhide Fujita, Takayuki Ito, and Koen V Hindriks. 2017. Automated Negotiating Agents Competition (ANAC).. In *AAAI*. 5070–5072.

[19] Ehud Kalai. 1977. Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica: Journal of the Econometric Society* (1977), 1623–1630.

[20] Ehud Kalai and Meir Smorodinsky. 1975. Other solutions to Nash's bargaining problem. *Econometrica: Journal of the Econometric Society* (1975), 513–518.

[21] Ryohei Kawata, Takaki Yasui, Mark Klein Yuta Hosokawa, and Katsuhide Fujita. 2019. *A Mediation for Multi-Issue Negotiation with Genetic Algorithm*. Technical Report 119–330. IEICE - AI19.

[22] Mark Klein, Peyman Faratin, Hiroki Sayama, and Yaneer Bar-Yam. 2003. Negotiating complex contracts. *Group Decision and Negotiation* 12 (2003), 111–125.

[23] Mark Klein, Peyman Faratin, Hiroki Sayama, and Yaneer Bar-Yam. 2003. Protocols for negotiating complex contracts. *IEEE Intelligent Systems* 18, 6 (2003), 32–38.

[24] Guoming Lai and Katia Sycara. 2009. A generic framework for automated multi-attribute negotiation. *Group Decision and Negotiation* 18 (2009), 169–187.

[25] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. 2014. Genius: An Integrated Environment for Supporting the Design of Generic Automated Negotiators. *Computational Intelligence* 30, 1 (2014), 48–70. https://doi.org/10.1111/j.1467-8640.2012.00463.x

[26] Ivan Marsa-Maestre, Enrique de la Hoz, Jose Manuel Gimenez-Guzman, David Orden, and Mark Klein. 2019. Nonlinear negotiation approaches for complex-network optimization: a study inspired by Wi-Fi channel assignment. *Group Decision and Negotiation* 28, 1 (2019), 175–196.

[27] Ivan Marsa-Maestre, Mark Klein, Catholijn M Jonker, and Reyhan Aydo ğan. 2014. From problems to protocols: Towards a negotiation handbook. *Decision Support Systems* 60 (2014), 39–54.

[28] Ivan Marsa-Maestre, Miguel A Lopez-Carmona, Juan R Velasco, Takayuki Ito, Mark Klein, and Katsuhide Fujita. 2009. Balancing utility and deal probability for auction-based negotiations in highly nonlinear utility spaces. In *Twenty-first international joint conference on artificial intelligence*.

[29] Yasser Mohammad. 2021. Concurrent local negotiations with a global utility function: a greedy approach. *Autonomous Agents and Multi-Agent Systems* 35, 2 (2021), 1–31.

[30] Yasser Mohammad. 2023. Evaluating Automated Negotiations. In *2023 IEEE International Conference on Agents (ICA)*. 77–82. https://doi.org/10.1109/ICA58824.2023.00022

[31] Yasser Mohammad. 2023. Evaluating Automated Negotiations. In *2023 IEEE International Conference on Agents (ICA)*. 77–82. https://doi.org/10.1109/ICA58824.2023.00022

[32] Yasser Mohammad. 2023. Generalized Bargaining Protocols. In *Australasian Joint Conference on Artificial Intelligence*. Springer, 261–273.

[33] Yasser Mohammad. 2023. Optimal time-based strategy for automated negotiation. *Applied Intelligence* 53, 6 (2023), 6710–6735.

[34] Yasser Mohammad. 2024. Generalized Bargaining Protocols. In *AI 2023: Advances in Artificial Intelligence*, Tongliang Liu, Geoff Webb, Lin Yue, and Dadong Wang (Eds.). Springer Nature Singapore, Singapore, 261–273.

[35] Y. Mohammad, K. Fujita, A. Greenwald, M. Klein, S. Morinaga, and S. Nakadai. 2019. ANAC 2019 SCML. http://tiny.cc/f8sv9y

[36] Yasser Mohammad, Amy Greenwald, and Shinji Nakadai. 2019. NegMAS: A platform for situated negotiations. In *Twelfth International Workshop on Agent-based Complex Automated Negotiations (ACAN2019) in conjunction with IJCAI* (Macau, China).

[37] Yasser Mohammad and Shinji Nakadai. 2018. FastVOI: Efficient Utility Elicitation During Negotiations. In *International Conference on Principles and Practice of Multi-Agent Systems (PRIMA)*. Springer, 560–567.

[38] John F Nash Jr. 1950. The bargaining problem. *Econometrica: Journal of the Econometric Society* (1950), 155–162.

[39] Ariel Rubinstein. 1982. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society* (1982), 97–109.

[40] Colin R Williams, Valentin Robu, Enrico H Gerding, and Nicholas R Jennings. 2012. Negotiating concurrently with unknown opponents in complex, real-time domains. In *Proc. of the Twentieth European Conference on Artificial Intelligence*. 834—839.