Automated Negotiation with No Information about Partner Utility Functions Using the Tentative Acceptance Unique Offers Protocol

Yasser Mohammad

Preprint (submitted version). This preprint has not undergone peer review or any post-submission improvements or corrections. The Version of Record of this article is published in *Autonomous Agents and Multi-Agent Systems*, 40(1), 21 (2026), and is available online at `https://doi.org/10.1007/s10458-026-09745-9`.

Automated negotiation with no information about partner utility functions using the Tentative Acceptance Unique Offers Protocol

Yasser Mohammad[1,2,3*]

[1*]Data Science Laboratories, NEC CORPORATION, Shiba, Minato, 108-8001, Tokyo, Japan.
[2]National Institute of Advanced Industrial Science and Technology (AIST), Aomi, Minato, 135-0064, Tokyo, Japan.
[3]Assiut University, AUN Street, Assiut, 71515, Assiut, Egypt.

Corresponding author(s). E-mail(s): y.mohammad@nec.com;

Abstract

Automated Negotiation (AN) is an approach for reaching agreements between self-interested agents. Early work in automated negotiation focused on proving game-theoretic results for simplified negotiation scenarios (e.g. when all agents' preferences are common knowledge). More recent research - with some exceptions - focuses on developing novel mediated negotiation mechanisms or effective strategies for the Alternating Offers Protocol (AOP) and its variations. Recently, we proposed the Tentative Acceptance Unique offers (TAU) protocol as an alternative to AOP with promising empirical results [1]. In this paper, we develop a Perfect Bayesian Equilibrium (PBE) strategy for TAU in bilateral negotiations with no information about partner preferences. We analyze its theoretical properties showing that it renders TAU exactly complete, and Pareto-optimal for all bilateral discrete negotiation scenarios, and fair in the sense of a modified Egaliterian criterion for *almost* all bilateral discrete negotiation scenarios. Moreover, we empirically compare TAU combined with the proposed strategy against AOP with its state-of-the-art strategies and show empirically that it provides a higher expected advantage for all agents while achieving higher agreement rates, Pareto-optimality, fairness, and welfare on widely used negotiation scenarios.

Keywords: Automated Negotiation, Bargaining, Mechanism Design, Agreement Technologies

1 Introduction

The ubiquitous presence of negotiation in human society finds its most common expression in bargaining, a back-and-forth exchange of offers culminating in agreement or stalemate. Automated negotiation builds upon this established framework, injecting automated agents as potential participants. Fueled by the growing integration of AI into business operations [2], research in automated negotiation aims to equip these agents with the capability to negotiate agreements that maximize utility for their represented entities. Automated negotiation has been applied in a variety of application domains including IoT permission management, Wi-Fi channel allocation [3], agricultural supply chain support, broad-spectrum supply chain management [4], vehicle routing [5], path planning [6], and even educational tools used to hone student negotiation skills [7].

Automated negotiation's roots lie in games like Nash's Bargaining Game [8], a single-round game where two agents divide a fixed pie. Each announces their desired portion, and both receive their demand if the sum doesn't exceed one. Otherwise, each receives a known predefined *reserved* value [8]. Rubinstein 1982 [9] extended the game to an infinite number of rounds with exponential or linear asymmetric time-pressure on agents and found this game's unique Perfect Equilibrium in which it turns out that the negotiation ends in a single-round anyway. However, both models assume that preferences and reserved values are common knowledge. This situation rarely – if ever – encountered in real-world applications.

The Alternating Offers Protocol (AOP) is the most popular negotiation protocol for situations with private preferences. Its multilateral extension, the Stacked Alternating Offers Protocol (SAOP), further expands its applicability to situations with more than two agents [10].

Exploration beyond AOP has led to alternatives, often involving mediators to facilitate agreement. One such example is the Single Text (ST) protocol, where a third party proposes tentative agreements based on agent responses and refines them iteratively [11]. While these mediated protocols hold promise, real-world business applications face a crucial concern: the mediator's influence on the final agreement. In ST, for instance, the mediator generates all proposals leading to the final outcome, making them a strategic player in the negotiation. Given their self-interest in real-world scenarios, mediators could potentially introduce bias (e.g., favoring larger businesses) in their proposal generation.

Several strategies have been proposed for AOP and SAOP [12]. Despite decades of research, there is no known PBE for AOP on the set of discrete scenarios with private preferences. Moreover, most existing state-of-the-art (SOTA) strategies for AOP have no theoretical guarantees, with some notable exceptions including the Greedy Concession Algorithm [13] and its variants [14] that assumes knowledge of offer acceptance probability. Notably, MiCRO was recently proposed as a robust baseline for AOP, guaranteeing Pareto-optimal agreements for all-rational discrete *balanced* scenarios [15].

Recently, we introduced Generalized Bargaining Protocols (GBP) as a generalization of AOP keeping its unmediated nature, inherent rationality, and distributed

implementability while opening the door for new designs that can provide theoretical and empirical advantages [1]. We also proposed the Tentative Acceptance Unique Offers Protocol (TAU) as an example and showed empirically that a simple strategy called Conceding Accepting Better (CAB) can achieve higher expected agent advantage, and overall designer-score compared with top strategies for AOP. No theoretical analysis of TAU or CAB was provided and the evaluation considered only the case of all agents employing CAB.

This paper has three main contributions: We construct *a* PBE strategy for TAU called Wasteful offering Accepting Rational (WAR) for negotiations with no-information about partner's preferences and show that it weakly-dominates CAB whenever an agreement is reached. Secondly, we provide theoretical analysis of TAU when combined with either CAB or WAR in terms of completeness, Pareto-optimality and fairness of resulting negotiation outcomes. Thirdly, the paper proposes a simple method for adapting any AOP strategy to be used with TAU. Finally, we empirically evaluate CAB, the proposed strategy and adapter on multiple scenario datasets and show major improvements over using AOP with its state-of-the-art strategies in all evaluation criteria (i.e. agent advantage, completeness, Pareto-optimality, and fairness).

2 Automated Negotiation

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.

Scenario

A Negotiation Scenario λ is defined as a tuple $(\mathcal{N}^\lambda, \mathcal{D}^\lambda, \mathcal{I}^\lambda)$ where $\mathcal{N}^\lambda$ is the set of agents numbered from 1 to n_A, $\mathcal{D}^\lambda$ is the negotiation domain, and $\mathcal{I}^\lambda$ is a tuple defining all the information about the negotiation scenario available to each agent.

Domain, Outcome Space, and Extended Outcome Space

The negotiation domain $\mathcal{D}^\lambda \equiv \left(\Omega^\lambda, \mathcal{F}^\lambda\right)$ defines (1) The outcome space Ω^λ of size n_o containing all potential agreements (Called outcomes). The special outcome $\phi \notin \Omega^\lambda$ is always assumed to exist to represent disagreement, and we define the *extended outcome space* Ω^+ as $\Omega^\lambda \bigcup \{\phi\}$.

Preferences

Agent preferences $\mathcal{F}^\lambda$ define their *valuation* of different outcomes. It can either be ordinal ($\succcurlyeq$) defining an ordering over Ω^+ – per agent – or cardinal (u) defining a mapping per agent from Ω^+ to $\Re$ with higher values indicating better outcomes. $a \succcurlyeq b$ means that a is not worse than b. Symbols $\succ, \approx$ are defined accordingly. If the codomain of u is the range $[0, 1]$, the utility function is called normalized. We use $\succcurlyeq_i$ and u_i to refer to the preferences/utility function for i. We use $r_i(\omega)$ to represent the relative rank of ω in Ω^+ according to $\succcurlyeq_i$ normalized so that $r_i(\phi) = 0$ and $\max(r_i(\omega)) = 1$. We define the ω's *value* for i ($\hat{u}_i(\omega)$) to be $r_i(\omega)$ for ordinal preferences and $u_i(\omega)$ normalized the same way for cardinal preferences.

Information Tuple

The *Information Tuple* $\mathcal{I}^\lambda$ is a tuple of n_A information data-structures $(\{\mathrm{I}_i^\lambda : i \in \mathcal{N}^\lambda\})$ where I_i^λ represents all the information available to agent i about the preferences $\mathcal{F}^\lambda$ and strategies of all agents including itself. This can be modeled by a joint probability distribution over $\mathcal{F}^\lambda$ and the set of strategies available to agents (Π): $\mathrm{I}_i^\lambda \equiv \left(p_i^{\mathcal{F}^\lambda,\Pi}\right)$. For the rest of this paper, we assume that the agent knows only the outcome space and its own preferences (i.e. $\mathrm{I}_i^\lambda = \left(\Omega^\lambda, \mathcal{F}_i^\lambda\right)$).

Important Sets

We define the following sets for a scenario λ:

Rational Outcome Set $\overline{\Omega}^\lambda$*:* The subset of the outcome space that a rational outcome at least as good as disagreement for everyone: $\left\{\omega : \omega \in \Omega \wedge \succcurlyeq_i \phi \forall i \in \mathcal{N}^\lambda\right\}$.

Win-Win deals $\widehat{\Omega}^\lambda$*:* Rational outcomes strictly better than disagreement for at least one agent:
$\left\{\omega : \omega \in \overline{\Omega}^\lambda \wedge \exists i \in \mathcal{N}^\lambda \rightarrow \omega \succ_i \phi\right\}$.

Pareto Outcome Set $\mathbf{P}$*:* Rational outcomes that cannot be improved for one agent without making at least one other agent worst off:
$\left\{\psi : \psi \in \overline{\Omega}^\lambda \wedge \left[\neg\exists \left(\omega \in \Omega^\lambda \setminus \{\psi\}, b \in \mathcal{N}^\lambda\right) : \omega \succ_b \psi \wedge \left(\omega \succcurlyeq_i \psi \forall i \in \mathcal{N}^\lambda\right)\right]\right\}$.

2.1 Protocols, Mechanisms and Strategies

Negotiations are conducted according to a protocol ($\mathcal{P}$) dictating what can be done by every agent at any point of time. In this paper, we focus on *bargaining* protocols in which agents propose offers and their partners can either accept these offers, reject and counter them, or immediately end the negotiation (by sending the disagreement outcome ϕ).

Protocol

In this paper, protocols are formally defined using a simplified version of GBP [1] in which all offers are restricted to be a single outcome. This allows us to represent the protocols discussed in this paper using a common framework. A bargaining protocol consists of four components: a *breaking rule* (ζ) defining timing constraints on the negotiation (if any), an *activation rule* (χ) that selects an agent to offer or respond to an offer, an *offering constraint* (γ) that defines – at every point in the negotiation – the set of outcomes available to offer for the active agent and an *evaluation rule* (σ) that defines the conditions for breaking the negotiation.

Strategy

A negotiation strategy describes the behavior of an agent in accordance to a protocol $\mathcal{P}$. Its input is the tuple $\left(n_A, i, \Omega, \mathrm{I}_i, \mathcal{T}^k\right)$, where i is the agent controlled by the strategy, $\mathcal{T}^k$ is the *negotiation trace* recording all offers exchanged and all responses to them up to step k. We use $\mathbb{T}$ to represent the set of all possible traces. The notation $i \approx \pi$ is used to indicate that agent i uses strategy π.

For any *bargaining* protocol, a negotiation strategy π makes two types of decisions: it needs to decide what to offer when it is its turn, and it needs to decide whether to accept an offer from a negotiation partner. A negotiation strategy for agent i (π_i) is thus defined as a tuple (η_i, ρ_i) where:

Offering Policy $\eta_i : \mathbb{T} \times \mathbb{S}(\Omega) \to \Omega^+$*:* Generates a valid outcome (offer) ω from the valid offers set $\Omega_i^v \subseteq \Omega$ selected by the protocol's constraint γ, or ϕ to immediately end the negotiation.

Acceptance Policy $\rho_i : \mathbb{T} \times \Omega \to \{0, 1\}$*:* Accepts or rejects an offer (1 indicates acceptance).

Both policies have access to the agent's information ($\mathcal{D}_i$) and the negotiation trace $\mathcal{T}$.

SAR

In order to compare different negotiation protocols, we need a way to assign strategies to negotiation scenarios called a *Strategy Assignment Rule* (SAR) and represented by Υ hereafter. A SAR is *implementable* if it assigns a strategy to every agent based only on information available to that agent (i.e. can be implemented by the agent itself). The simplest type of SAR is a *pure* SAR which assigns the same strategy π to all agents, and we refer to it as Υ_π. A SAR that always assigns a strategy from a set Π is referred to as Υ_Π.

Mechanism

We define a mechanism $\mathcal{P}_\Upsilon$ as the tuple ($\mathcal{P}$, Υ). Given a set of scenarios Λ, a mechanism $\mathcal{P}_\Upsilon$ generates a probability distribution over the extended outcome space giving the probability of getting every outcome as the final negotiation outcome called the *negotiation outcome distribution* $\mathcal{P}_\Upsilon(\Lambda) : \Omega^+ \to [0, 1]$. If $\Lambda = \{\lambda\}$, we indicate the probability distribution by $p_{\omega_*^\lambda}$. We indicate executing a *single* negotiation using mechanism $\mathcal{P}_\Upsilon$ on λ leading to $\omega_*^\lambda \in \Omega^+$ as a negotiation outcome by $\mathcal{P}_\Upsilon(\lambda) \equiv \omega_*^\lambda$. If the SAR is pure, we will simplify this notation as $\mathcal{P}_\pi(\lambda)$.

Negotiation Session

Now we can define a negotiation session using a mechanism as an execution of Algorithm 1. The process starts by assigning strategies to agents (Line 1). The following process is then repeated until the breaking rule returns a nonzero value (Line 3), an agent leaves the negotiation (Line 13), or an agreement is reached (Line 14): The next agent to offer and to respond are selected by the activation rule (Line 4). The selected agent is then asked to offer (Line 6) and responses are collected (Line 7-10). The evaluation rule is then invoked (Line 12) to decide whether the negotiation ends in disagreement (Line 13), agreement (Line 14), or continues (Line 17).

2.2 AOP and its Strategies

The Alternating Offers Protocol (AOP) $\mathcal{P}_{ao}$ is the most widely used bilateral negotiation protocol. An agent starts the negotiation by offering an outcome to its partner. Whenever an agent receives an offer, it can (1) accept it leading to an agreement, (2)

Algorithm 1 Bargaining Negotiation Mechanism: $\mathcal{P}_\Upsilon(\lambda)$

Input: Scenario $\lambda \equiv \left(\mathcal{N}^\lambda, \left(\Omega, \mathcal{F}^\lambda\right), \mathcal{I}^\lambda \equiv \left\{\mathrm{I}_i^\lambda : i \in \mathcal{N}^\lambda\right\}\right)$, Mechanism $\mathcal{P}_\Upsilon = (\mathcal{P}, \Upsilon) : \mathcal{P} = (\zeta, \chi, \gamma, \sigma)$
Output: $\omega_* \in \Omega \cup \{\phi\}$

$(\eta_i, \rho_i) \forall i \in \mathcal{N}^\lambda \leftarrow \Upsilon(\lambda)$ ▷ Assign strategies
$k, t, t_0, \mathcal{T} \leftarrow 0, 0, clock(), \langle\rangle$ ▷ Init step, time and trace
while $\zeta(\mathcal{T}, k, t; \lambda) = 0$ **do** ▷ check the breaking rule
  $i, \mathcal{R} \leftarrow \chi(\mathcal{T}, k; \lambda)$ ▷ Next offerer and responder(s)
  $\Omega_i^v \leftarrow \gamma(\mathcal{T}, i; \lambda)$ ▷ Valid offering set
  $\omega_k \leftarrow \eta_i\left(\mathcal{T}, \Omega_i^v; \mathrm{I}_i^\lambda\right)$ ▷ Receive an offer
  $R_k \leftarrow \langle\rangle$ ▷ Initialize responses to the empty tuple
  for $y \in \mathcal{R}$ **do** ▷ Collect responses
    $R_k \leftarrow R_k + \left\langle y, \rho_y(\mathcal{T}, \omega_k; \mathrm{I}_y^\lambda)\right\rangle$ ▷ Response map
  end for
  $\mathcal{T} \leftarrow \mathcal{T} + \langle \omega_k, i, R_k\rangle$ ▷ Update the trace
  $z \leftarrow \sigma\left(\mathcal{T}, k; \lambda\right)$ ▷ Evaluation rule
  if $z = \phi$ **then return** ϕ ▷ End in **disagreement**
  else if $z \in \Omega$ **then return** ω_k ▷ End with **agreement**
  end if ▷ Otherwise continue to the next step
  $k, t \leftarrow k + 1, clock() - t_0$ ▷ Update step and time.
end while
return ϕ ▷ Disagreement due to breaking rule (timeout).

leave the negotiation leading to disagreement or (3) send a counter-offer to the partner. This process repeats until either a predefined number of rounds or seconds have passed which automatically ends the negotiation with a disagreement. This can be formally expressed as: $(\zeta_{ao}, \chi_{ao}, \gamma_{ao}, \sigma_{ao})$ where:

$$\zeta_{ao}(\mathcal{T}^k, t) = k > N \vee t > T$$
$$\chi_{ao}(k) = (k \mod n_A, \mathcal{N} \setminus \{k \mod n_A\})$$
$$\gamma_{ao}(\mathcal{T}^k, i) = \Omega$$
$$\sigma_{ao}(\mathcal{T}^k, k) = \begin{cases} \phi & \text{if } \omega_k = \phi \\ \omega_k & \text{if } \sum_{j=1}^{n_A - 1} r_j^k = n_A - 1 \\ \triangleright & otherwise \end{cases}$$

where ω_k is the last offer, r_j^k is the jth response received in step k, N (T) is the maximum allowed number of steps (seconds), t is the number of seconds since the beginning of the negotiation, and "▷" means "*continue the negotiation*" (i.e. check the breaking rule).

The Minimal Concession in Reply to New Offers (MiCRO) policy is a recently proposed offering policy for AOP that provides several attractive features [15]. It is only defined for scenarios with a finite number of outcomes Λ_D. MiCRO is an offering policy that works as follows: If the number of unique offers from the partner is higher than or equal to the number of unique offers it has proposed so far, and there are any

offers better than disagreement that it did not offer yet (Ω_{new}); it offers the best offer in Ω_{new}. Otherwise, it offers a random offer from the set it already offered before (Ω_{old}). MiCRO can be combined with any acceptance policy, but the one recommended by de Jonge 2024 [15] (and used in our paper) is to accept an offer if and only if it is better than the worst offer MiCRO is willing to offer in the next step. The MiCRO strategy is defined here as the MiCRO offering policy and this acceptance policy.

MiCRO achieves the slowest possible concession as long as the partner is changing its offers (interpreted as concession). If played against itself, it can be shown to be Pareto-efficient whenever an agreement is reached [15] but it is easy to show that sometimes a win-win deal is available but cannot be found (i.e. MiCRO is not complete). Consider two MiCRO agents A and B such that $\omega_1 \succ_A \phi \succ_A \omega_2 \succcurlyeq_A \omega_3$, and $\omega_3 \succ_B \omega_2 \succ_B \omega_1 \succ_B \phi$. Agent A will keep repeating ω_1 and B will offer ω_3 followed by ω_2 then keep sampling from $\{\omega_2, \omega_3\}$ which will lead to disagreement (ϕ) no matter how long the negotiation is allowed to run despite the fact that a win-win deal exists (ω_1). In summary $AOP(\Upsilon_{MiCRO})$ is exactly optimal and complete only when both agents can concede almost the same number of times towards the agreement [15]. This set is called **balanced** scenarios in [15] (Λ_B^B)[1].

2.3 Evaluating Negotiations

Evaluating negotiation mechanisms is multifaceted. Mohammad 2023 [16] 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[2]. **Fairness** is inversely proportional to the distance between the negotiation outcome and the *nearest* Bargaining Solution (i.e. Kalai [17], Nash [8], Kalai-Smorodinsky [18]) or an ordinal version of this solution that utilizes outcome ranks instead of utilities [16]. **Welfare** is the sum of advantages of all agents. All of these measures are normalized to have a maximum value of one. We call a mechanism *exactly* rational if its rationality is one. The same applies for all other performance measures.

These evaluation criteria are only defined for a given set of scenarios. In this paper, we consider mostly the set of scenarios with discrete outcomes spaces, full information about own preferences and no-information about partner(s) preferences (Λ_{DN}). An important subset of these are scenarios in which for any outcome in the Pareto Outcome Set, there are no non-pareto outcomes that have the same utility for any agent (Λ_U).

[1] In this paper, we consider a scenario *balanced* if it is balanced under *any* solution concept [15].

[2] 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

3 TAU Protocol

The Tentative Acceptance Unique Offers Protocol (TAU) applies the following simple modifications to standard AOP/SAOP [1].

1. Agents cannot repeat offers except their last one. If all agents are repeating their last offer, the negotiation ends with disagreement.
2. Once an agent starts repeating an offer, it must keep repeating it forever.
3. An agreement is reached if and only if an outcome is offered by every agent and accepted by every agent. By making all agreements tentative, TAU allows strategies to better explore the outcome space.
4. The negotiation terminates if an agreement is reached, an agent explicitly ends the negotiation (by offering ϕ or breaking the rules above) or all agents started repeating.

We need three more definitions to formally describe this protocol: Ω_i^k is the set of i's offers up to and including at step k, ω_i^{last} is the last offer from i (ϕ if none), and $\mathbb{A}_i^k$ is the set of all offers accepted by i up to and including at step k.

Let, k be the current step and n_o be the size of the outcome space, and Ω_i^k be the offer from agent i at step k, TAU can formally be defined as $(\zeta_\tau, \chi_\tau = \chi_{ao}, \gamma_\tau, \sigma_\tau)$

$$\zeta_\tau(\mathcal{T}^k, t) = 0$$

$$\gamma_\tau(\mathcal{T}^k, i) = \begin{cases} \left\{\omega_i^{last}\right\} \cup \left(\Omega \setminus \Omega_i^k\right), & \text{if } \Omega_i^k \neq \Omega_i^{k-1} \\ \left\{\omega_i^{last}\right\}, & \text{if } \Omega_i^k = \Omega_i^{k-1} \end{cases}$$

$$\sigma_\tau(\mathcal{T}^k, k) = \begin{cases} \omega, & \text{if } \forall i \in \mathcal{N} : \omega \in \mathbb{A}_i^k \cap \Omega_i^k \\ \phi, & \text{if } \forall i \in \mathcal{N} : \Omega_i^k = \Omega_i^{k-1} \vee \exists i \in \mathcal{N} : \phi \in \Omega_i^k \\ \triangleright, & otherwise. \end{cases}$$

3.1 Conceding Accepting Better: CAB

The Conceding Accepting Better (CAB) is the only known strategy for the TAU protocol [1]. CAB's offering policy progresses through two stages: Rational, and Repeat. In the first stage (Rational), it offers all its rational outcomes from best to worst breaking ties in some predefined shared lexical ordering. The ordering does not matter as long as it is shared, and no agent can benefit from breaking ties any other way. In the final Repeat stage – reached when there are no more rational outcomes in the outcome space to offer – CAB repeats its last offer which will always be the worst rational offers with ties broken according to the shared lexical ordering. CAB's acceptance policy (AcceptBetter hereafter) is quiet simple: accept if and only if the current offer is not worse than disagreement and is either the first offer received or is better than the last outcome accepted before.

4 Proposed Strategies

4.1 Wasteful offering Accepting Rational: WAR

We propose the Wasteful offering Accepting Rational (WAR) strategy for TAU which will be shown to be a PBE equilibrium for the widest set of scenarios considered in this paper (Λ_{DN}). WAR's offering policy progresses through three stages: Irrational, Rational, and Repeat. In the first stage (Irrational) WAR offers all its irrational outcomes (in any order). The remaining two stages are exactly the same as CAB but do not use the shared lexical ordering.

The main intuition behind the Irrational stage, is to first use all of our irrational outcomes as *time-wasters* to entice our partner to concede. We can do that safely under TAU (but not AOP) because we will later have the chance to reject these offers when they are offered to us. In other words, none of these irrational outcomes can become an agreement.

Let $\succ_l$ be a shared full ordering of the outcome space, WAR's offering policy can be formally written as:

$$\eta_i(k) = \begin{cases} \sim \underline{\Omega}_i^v(k), & \text{if } \underline{\Omega}_i^v(k) \neq \{\} \\ \sim \overline{\Omega}_i^v(k), & \text{if } \underline{\Omega}_i^v(k) = \{\} \wedge \overline{\Omega}_i^v(k) \neq \{\} \\ \omega_i^{k-1}, & \text{otherwise} \end{cases} \tag{1}$$

where $\sim X$ is any sample drawn from X, ω_i^{k-1} is the agent's last offer, $\underline{\Omega}_i^v(k)$ is the set of irrational outcomes that agent i can offer:

$$\underline{\Omega}_i^v(k) \equiv \left\{\omega : \omega \in \Omega_i^v(k) \cap (\Omega \setminus \overline{\Omega}_i) \wedge \omega \succcurlyeq_i \psi \forall \psi \in \Omega_i^v(k)\right\},$$

and $\overline{\Omega}_i^v(k)$is the set of rational outcomes that agent i can offer:

$$\overline{\Omega}_i^v(k) \equiv \left\{\omega : \omega \in \Omega_i^v(k) \cap \overline{\Omega}_i \wedge \omega \succcurlyeq_i \psi \quad \forall \psi \in \Omega_i^v(k)\right\},$$

and $\Omega_i^v(k) \equiv \Omega \setminus \Omega_i^{k-1}$, and Ω_i^k is i's offers up to step k.

WAR's acceptance policy simply accepts any rational outcome (AcceptRational hereafter):

$$\rho_i(\omega, k) = \mathbb{1}\left[\omega \succcurlyeq_i \phi\right] \tag{2}$$

4.2 Adapting AOP strategies to TAU

TAU changes the meaning of offers and acceptances making both of them *tentative* and adds new constraints on the offering strategy. Even though any existing AOP (or SAOP) strategy can be used with TAU without modification it will most likely lead to poor negotiation outcomes because most of these strategies do repeat offers frequently and have a conservative acceptance strategy which will result in almost all negotiations ending in disagreement.

To adapt any existing AOP (or SAOP) strategy (η_i, ρ_i) to TAU (as $(\eta_i^\tau, \rho_i^\tau)$), we apply the following two conservative changes:

Offering policy

Reproduce the same offer as the original offering strategy if it is valid, otherwise replace it with the valid offer with the nearest value for the agent limiting the search to rational outcomes if the original offer was rational. Let $\omega^a \equiv \eta_i\left(\Omega_i^v(i), \mathcal{T}^i\right)$, this process can be expressed as:

$$\eta_i^\tau(\Omega_i^v(i), \mathcal{T}^i) = \begin{cases} \underset{\omega \in \overline{\Omega}_i \bigcap \Omega_i^v(k)}{\operatorname{argmin}} \left|\hat{u}(\omega^a - \hat{u}(\omega)\right|, & \text{if } \omega^a \in \overline{\Omega}_i \\ \underset{\omega \in \Omega_i^v(k)}{\operatorname{argmin}} \left|\hat{u}(\omega^a - \hat{u}(\omega)\right|, & \text{otherwise} \end{cases}$$

Acceptance policy

Replace rejections from the acceptance policy with the AcceptRational policy. Acceptance decisions are passed unchanged: $\rho_i^\tau(\omega, \mathcal{T}) = \rho_i^{ao}(\omega, \mathcal{T}) \vee \mathbb{1}\left[\omega \succ_i \phi\right]$.

This adaptation procedure can lead to a change in the behavior of the adapted strategy depending on the density of outcomes when projected on negotiator utilities as well as the fraction of the outcome-space explored by the original strategy during the negotiation. Nevertheless, as will be shown empirically in Section 6, the adapted strategy leads in practice to improved negotiation outcomes in almost all cases.

5 Theoretical Guarantees

In this section, we provide a set of theorems to establish the theoretical properties of TAU when combined with CAB and WAR. Full proofs are provided in the appendix.

Theorem 1. *TAU always terminates in a finite number of steps T not greater than $n_o + 1$ for any strategy assignment rule on Λ_{DN}.*

$$T \leq n_o + 1 \tag{3}$$

Theorem 2. *TAU is exactly rational with any strategy assignment rule that always assigns strategies with the AcceptRational or AcceptBetter acceptance policies on Λ_{DN} including WAR and CAB.*

Theorem 3 (Optimality of CAB)**.** *TAU(Υ_{CAB}) is Exactly Optimal on Λ_{DN}.*

Proof Sketch. If there are no win-win deals, and using Proposition 2, optimality is trivially guaranteed. If there are any win-win deals, we prove that if we assume that the negotiation outcome ω_*^λ is not on the Pareto Outcome Set, another outcome ψ must exist which dominates it, and is offered and accepted by all agents before this negotiation outcome which is a contradiction as ψ (not ω_*^λ) would be the negotiation outcome. □

Theorem 4 (Optimality of WAR)**.** *TAU(Υ_{WAR}) is Exactly Optimal on Λ_U.*

Proof Sketch. We show that the proof for Theorem 3 works for WAR as long as there are no Pareto-outcome that has the same value for some agent as a non-Pareto-outcome. This condition is always guaranteed for any scenario in Λ_U by definition. □

Theorem 5 (Completeness)**.** *TAU(Υ_{WAR}) and TAU(Υ_{CAB}) are Exactly Complete on Λ_{DN}^B.*

Proof Sketch. If an agreement is reached, and using Theorem 3 and the fact that the Pareto Outcome Set is a subset of the win-win deals, we establish exact completeness. If no agreement is reached, we establish that both agents are repeating offers in the last step. We then show that this can only happen if the set of win-win deals is empty which establish exact completeness for this case as well. □

Theorem 6 (CAB Fairness)**.** *$TAU(\Upsilon_{CAB})$ is Exactly Fair on Λ^B_{DN} with no irrational outcomes. Specifically, it finds an Ordinal Kalai Bargaining Solution in that case:*

$$TAU_{CAB}(\lambda) \in \left\{\omega : \omega \in \underset{\omega \in \mathbf{P}^\lambda \bigcup \{\phi\}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} r_i(\omega)\right\} \forall \lambda \in \Lambda_{DN}$$

Proof Sketch. Each agent using CAB will rank its outcomes descendingly and then go over the list from top to bottom stopping at its worst rational outcome. Agreement will be reached whenever these two pointers intersect which happens at the Ordinal Kalai Bargaining Solution. □

Theorem 7 (WAR Fairness)**.** *$TAU(\Upsilon_{WAR})$ ($TAU(\Upsilon_{CAB})$) is Exactly Fair for on Λ^B_U (Λ^B_{DN} with no irrational outcomes). Specifically, it finds the Ordinal Kalai Bargaining Solution:* $TAU_{WAR}(\lambda) \in \left\{\omega : \omega \in \underset{\omega \in \mathbf{P}^\lambda \bigcup \{\phi\}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} r_i(\omega) - r_i(\phi)\right\}$

Proof Sketch. The proof follows the same steps as the previous one but it shows that the introduction of the irrational phase in WAR correctly changes the agreement point to match the Ordinal Kalai Bargaining Solution for negotiations with irrational outcomes. □

We say that a strategy **weakly-dominates** another for a set of scenarios, iff it never leads to an agreement with lower advantage for any agent ignoring negotiations leading to disagreement.

Theorem 8 (Weak Domination)**.** *WAR weakly dominates CAB for TAU on Λ^B_{DN}.*

Proof Sketch. We show that whenever Υ_{WAR} finds an outcome in the Pareto Outcome Set, this outcome must be either the same or better for every agent compared with what CAB finds. Moreover, whenever WAR finds an outcome that is not in the Pareto Outcome Set, this outcome must have the same utility for the agent running WAR as the one that would have been found if it was running CAB. Taken together, the theorem is proven. □

Theorem 9 (Equilibrium)**.** *Υ_{WAR} is a PBE of TAU for on Λ^B_{DN}.*

Proof Sketch. We start by defining the induced incomplete information game corresponding to TAU on Λ_{DN} and characterise its information sets (Section C.11.1). We then define the beliefs of agents in the game (Section C.11.2) and describes a consistent Bayesian belief update given that the partner is using WAR (Section C.11.3). We then show that assuming no information about the partner makes this process much simpler than if partial information about partner preferences (e.g. a linear aggregation utility function) is available. Finally, we show that the best response to WAR is itself by establishing that changing neither the acceptance policy nor the offering policy can lead to any improvement in expected value when facing WAR (Section C.11.4). □

All theoretical results presented in the paper can be summarized in Table 1. The following sets of scenarios are used in this table and throughout this document:

Λ_{DN} All scenarios with a discrete outcome space and no information about partner preferences.

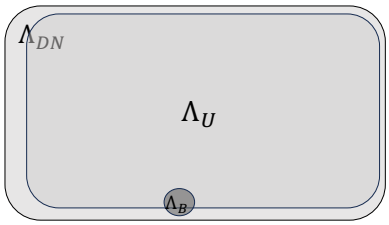

Fig. 1: Sets of scenarios used in this work and their relationship. Note that the size of Λ_B is vanishingly small compared to Λ_{DN} while Λ_U is almost equal in size to Λ_{DN}

Table 1: Theoretical Result Summary

Set	Strategy	Complete	Optimal	Fair	Equilibrium
Λ_{DN}	TAU-CAB	✓x	✓	x	x
	TAU-WAR	✓	x	?	✓
	AOP-MiCRO	x	x	x	?
Λ^{B}_{DN}	TAU-CAB	✓	✓	x	x
	TAU-WAR	✓	x	?	✓
	AOP-MiCRO	x	x	x	?
Λ^{RB}_{DN}	TAU-CAB	✓	✓	✓	✓
	TAU-WAR	✓	x	✓	x
	AOP-MiCRO	x	x	x	?
Λ^{B}_{U}	TAU-CAB	✓	✓	x	x
	TAU-WAR	✓	✓	✓	✓
	AOP-MiCRO	x	✓	✓	?
Λ^{RB}_{B}	TAU-CAB	✓	✓	✓	✓
	TAU-WAR	✓	✓	✓	x
	AOP-MiCRO	✓	✓	✓	✓

? Unknown ^{B} Bilateral ^{R} All Rational
✓x Conjecture (No Proof Provided) ✓/x Proven true/false.

Λ_U All scenarios in Λ_{DN} with the added requirement that no two outcomes have the same rank (or utility) for any agent.

Λ^{RB}_{B} Balanced scenarios as defined by de Jonge 2024 [15]. These are scenarios for which MiCRO is optimal and are characterised with the property that the number of outcomes better than the agreement is almost equal for all agents. Moreover, there are exactly two agents (bilateral negotiations) and all outcomes must be rational for both agents.

Fig. 1 shows these three sets of scenarios and their relationships. Because Λ_U only requires that no two outcomes have *exactly* the same utility value, any scenario Λ_{DN} can be converted to a scenario in Λ_U by adding a vanishingly small value to the utility of some of its outcomes. The constraint on Λ_B – on the other hand – is much stricter as to convert a domain not in Λ_B to one in Λ_B, the number of outcomes above reserved value must be changed which is a discrete change.

6 Empirical Evaluation

Table 2: The datasets used in our experiments.*Strategy* refers to the number of AOP strategies used. *Combination* refers to the total number of strategy combinations used per domain.

	Dataset	Strategy	Combination	Domain	Scenario	Negotiation	Type	Years
1	main	11	411	184	736	302,496	both	2010-22‡
2	2010	7	179	3	3	537	bilateral	2010
3	2010*	4	68	3	3	204	bilateral	2010
4	2011	8	228	8	8	1,824	bilateral	2011
5	2012	8	228	15	15	3,420	bilateral	2012
6	2013	7	179	14	14	2,506	bilateral	2013
7	2015	23	1,683	15	15	25,245	multilateral	2015
8	2015*	8	228	15	15	3,420	multilateral	2015
9	2016	16	836	15	15	12,540	multilateral	2016
10	2016*	10	344	15	15	5,160	multilateral	2016
11	2017	18	1,048	10	10	10,480	multilateral	2017
12	2017*	10	344	10	10	3,440	multilateral	2017
13	2018	20	1,284	4	4	5,136	multilateral	2018
14	2018*	13	563	4	4	2,252	multilateral	2018
15	2024	20	884	488	488	1,294,176	bilateral†	2024
16	2024*	10	244	488	488	119,072	bilateral†	2024
		138	8,751	672	1,308	1,658,360		2010-24‡

* Finalists only ‡ Except 2014 † Known partner utility up to reserved value.

The Automated Negotiating Agents Competition (ANAC) is the largest international automated negotiation competition running since 2010 in conjunction with AAMAS or IJCAI [19, 20]. We used ANAC as the source of all our scenarios and AOP strategies. In this section we evaluate TAU with CAB, the proposed WAR (Section 4.1) strategy, and the proposed strategy adaptation method (Section 4.2) empirically against AOP with baseline and SOTA strategies for AOP.[3] All experiments were conducted using a fork of NegMAS 0.10.12 [21] platform with its NegMAS-GENIUS bridge [22] for the official implementations of AOP strategies and the ANL-Agents official implementation for 2024's strategies, on machines with 64 Xenon CPUs and 700 GB of RAM.

For the main experiment, we used six baselines: Time-based Conceder, Linear and Boulware strategies, the behavioral Nice Tit for Tat (NTfT) strategy [23], and the recently proposed MiCRO [15]. We also use the following set of SOTA strategies for AOP: Caduceus [24] , Atlas3 [25](winners of multilateral competitions without uncertainty or learning in 2015 and 2016), AgentK by Kawaguchi et al. 2013 [26], Hardheaded by van Krimpen et al. 2013 [27], CUHK by Hao and Leung 2014 [28], and AgentGG [29] (most recent winners of bilateral competitions without knowledge

[3]Extra results and analysis are available in the technical appendix. The full source-code is available in the supporting webpage at: https://github.com/yasserfarouk/tau_paper

Table 3: Summary for empirical evaluation. *Adapted* refers to the AOP strategies adapted using the proposed adapter to run under TAU.

Dataset	Protocol	AOP		TAU					
	Variant	n_o	3min	Ada.†	Ada.‡	CAB†	CAB‡	WAR†	WAR‡
main	Advantage	0.41	0.45	0.62	0.60	0.69	0.66	**0.72**	**0.72**
	DesignerScore	0.50	0.54	0.81	0.81	**0.93**	**0.93**	0.90	0.90
	Speed	5.31	0.83	9.91	10.59	**33.19**	33.19	23.03	23.03
	Completeness	0.60	0.65	0.88	0.88	**1.00**	**1.00**	**1.00**	**1.00**
	Fairness	0.73	0.76	0.92	0.91	0.97	0.97	**0.99**	**0.99**
	Optimality	0.80	0.82	0.94	0.94	**1.00**	**1.00**	**1.00**	**1.00**
	Welfare	0.63	0.67	0.87	**0.96**	0.91	0.87	**0.96**	0.91
	Privacy	**0.95**	0.93	0.90	0.87	0.84	0.84	0.56	0.71
2010	Advantage	0.38	0.49	0.64	0.67	**0.80**	0.76	**0.80**	0.76
	DesignerScore	0.41	0.53	0.81	0.86	**0.99**	**0.99**	**0.99**	**0.99**
2010*	Advantage	0.37	0.70	**0.80**	**0.80**	**0.80**	**0.80**	**0.80**	**0.80**
	DesignerScore	0.42	0.79	**0.99**	**0.99**	**0.99**	**0.99**	**0.99**	**0.99**
2011	Advantage	0.55	0.64	0.67	0.68	**0.75**	0.74	**0.75**	**0.75**
	DesignerScore	0.60	0.68	0.79	0.82	**0.90**	**0.90**	**0.90**	**0.90**
2012	Advantage	0.14	0.58	0.62	0.64	**0.74**	**0.74**	**0.74**	**0.74**
	DesignerScore	0.15	0.66	0.79	0.85	**0.97**	**0.97**	**0.97**	**0.97**
2013	Advantage	0.47	0.47	0.78	0.78	**0.79**	**0.79**	**0.79**	**0.79**
	DesignerScore	0.49	0.48	0.94	0.94	**0.98**	**0.98**	**0.98**	**0.98**
2015	Advantage	0.33	0.23	0.56	0.57	0.74	0.74	**0.76**	0.73
	DesignerScore	0.31	0.22	0.65	0.68	0.81	0.81	**0.90**	**0.90**
2015*	Advantage	0.38	0.18	0.69	0.69	0.74	0.74	**0.76**	0.74
	DesignerScore	0.35	0.17	0.82	0.82	0.81	0.81	**0.90**	**0.90**
2016	Advantage	0.23	0.29	0.57	0.57	0.68	0.67	**0.69**	0.63
	DesignerScore	0.25	0.32	0.70	0.71	0.83	0.83	**0.91**	**0.91**
2016*	Advantage	0.25	0.38	0.57	0.57	0.68	0.66	**0.69**	0.63
	DesignerScore	0.28	0.43	0.73	0.74	0.83	0.83	**0.91**	**0.91**
2017	Advantage	0.37	0.47	0.72	0.72	0.74	0.76	**0.82**	0.79
	DesignerScore	0.36	0.45	0.79	0.79	0.78	0.78	**0.95**	**0.95**
2017*	Advantage	0.37	0.48	0.68	0.69	0.74	0.73	**0.82**	0.76
	DesignerScore	0.37	0.47	0.77	0.78	0.78	0.78	**0.95**	**0.95**
2018	Advantage	0.32	0.45	0.60	0.61	**0.76**	0.70	0.73	0.70
	DesignerScore	0.34	0.47	0.70	0.72	0.83	0.83	**0.92**	**0.92**
2018*	Advantage	0.38	0.51	0.62	0.63	**0.76**	0.73	0.73	0.73
	DesignerScore	0.41	0.52	0.72	0.75	0.83	0.83	**0.92**	**0.92**
2024	Advantage	0.33		0.26	0.25	0.34	0.30	**0.41**	0.33
	DesignerScore	0.70		0.59	0.58	0.87	0.87	**0.92**	**0.92**
2024*	Advantage	0.38		0.23	0.23	0.34	0.26	**0.41**	0.31
	DesignerScore	0.83		0.54	0.55	0.87	0.87	**0.92**	**0.92**

*: Finalist-only datasets. †: Against same type ‡: Against everyone

of partner preferences, uncertainty in own preferences or learning at ANAC in 2010, 2011, 2012, 2023). All 121 combinations of these strategies were used.

To avoid cherry picking, we use all scenarios of ANAC 2010 to 2022 except 2014 which used nonlinear utility functions not compatible with most AOP strategies tested (184 unique scenarios). For each scenario, we varied the reserved value of the second

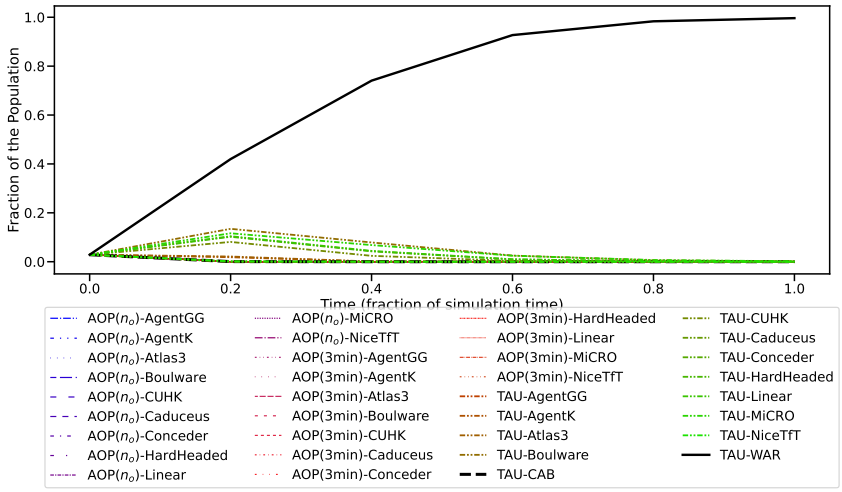

Fig. 2: Replicator Dynamics. AOP strategies go extinct almost immediately and WAR clearly outperforms all other options.

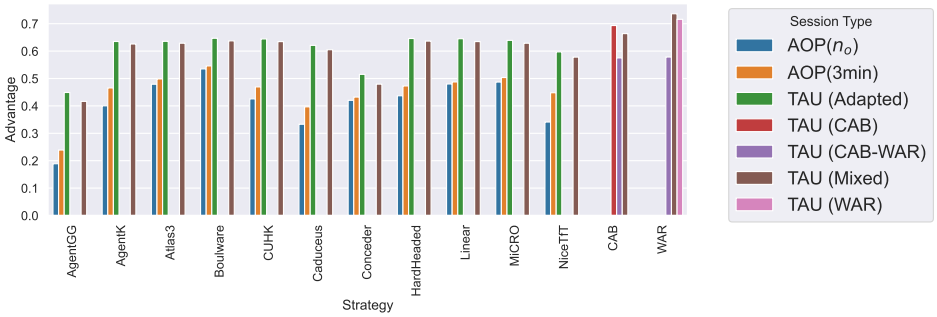

Fig. 3: Analysis of *advantage* for the main dataset. The *Adapter* condition consists of adapted strategies negotiating among themselves, The *Mixed* condition consists of adapted strategies negotiating against CAB or WAR. Using the proposed adapter under TAU improves the advantage of All AOP strategies. WAR outperforms all other conditions. CAB outperforms adapted strategies.

agent so that 100%, 90%, 50%, and 10% of the outcome space is rational to simulate outside options. For multilateral scenarios, all agents except the first always used the same strategy. For AOP we ran negotiations with either a 3 minutes time limit (standard for ANAC competitions) of $2n_o + 1$ steps (intrinsic limit of TAU). For TAU, we use WAR, CAB, and the adapter proposed in Section 4.2 on the same 11 strategies leading to 169 different strategy combinations. This leads to 302,496 different negotiations in total. All differences were checked using factorial dependent

t-test (and Wilcoxon's rank test) with Bonferroni's multiple-comparisons correction and differences reported hereafter are all statistically significant.

One possible issue with this experimental design is that most AOP strategies selected are *strong* which prevents agents from *exploiting* weak agents. Also, some of the AOP strategies were not designed to handle non-zero reserved values because they were rarely used in ANAC. To resolve these two issues, we repeated the same experiment using all scenarios and strategies of each year of ANAC between 2010 and 2018 with the exception of 2014 available through the Genius 10 platform (8 datasets) without varying reserved values. Some of these agents were too weak that they did not pass the initial qualification round of ANAC. We repeated the experiment again for all scenarios and only finalist agents for all years in which this happened (2010, 2015-2018) leading to 5 more datasets. We repeated the same process with scenarios and strategies from ANAC 2024 in which AOP strategies had access to partner utility functions with the proposed strategy and adapter lacking such information (repeating every run three times). In total, we used 16 datasets with $1,658,360$ negotiations. Table 2 describes the characteristics of all the datasets used in this paper. The rest of this section will focus on the main experiment.

6.1 Game Theoretic Analysis

Focusing on the 24 mechanisms we have in the main experiment (i.e. 11 strategies for AOP and 13 for TAU), we construct a general-sum simultaneous-action game. An *action* in this game corresponds to selecting a mechanism which means selecting a protocol and a strategy for that protocol. An action profile is the set of actions of both agents. The payoffs for any action profile are expected *advantage* over all scenarios for the two agents. Payoffs are set to reserved-values if the agents selected different protocols. The Lemke-Howson Algorithm [30] and explicit best-response checking of pure strategy profiles were used to find Nash Equilibria of this game. Only one dominant pure Nash-equilibrium was found corresponding to using TAU with WAR (Theorem 9).

We also applied replicator dynamics to discover which game-strategies are fixated in the population if we start with the same proportion for each of them. Fig. 2 shows the results of this experiment. AOP strategies disappear from the population relatively quickly. All TAU strategies increase their fractions in the beginning but only WAR keeps increasing its fraction until fixation.

These results were replicated for all of our 14 scenario/strategy datasets with no information about partner utility functions. In seven cases, some adapted TAU strategies appeared as undominated Nash equilibira. In all cases, TAU(WAR) appeared as a pure undominated Nash equilibrium. TAU(WAR) was always the condition with maximum frequency at the end of the replicator dynamics simulation for all *bilateral* negotiation datasets (except for 2013 in which CAB took this place). This repeated for the multilateral datasets except for years 2017 and 2018. In call cases, AOP strategies were *always* driven to extinction almost immediately. For 2024 datasets, in which agents knew partner utility functions up to reserved values, AgentRenting2024 drove all other AOP, adapted and TAU-pure strategies extinct. Nevertheless, WAR still provided an undominated Nash equilibrium of the game even under this condition. This shows that the proposed method is still empirically effective outside the domain for

which it was theoretically analyzed (e.g. multilateral negotiations and partial information about partner utility functions) but when — good enough — information about the partner is available, it may not be rational to switch to TAU as expected.

6.2 Advantage and Designer Score

Table 3 shows the results of applying the evaluation criteria proposed by Mohammad 2023 [16]. For TAU, we distinguish between six conditions: 1) Adapted-runs between AOP-adapted strategies and themselves, and 2) between an adapted strategy and all TAU strategies, CAB/WAR runs between each of these strategies and itself (3, 4), and CAB/WAR against all TAU strategies (5, 6). In all cases, we can see that TAU with any set of strategies considered outperforms AOP in both designer-score and advantage except for 2024 datasets. The price is a reduction in Privacy (5% for the proposed adapter, and 24% for the WAR strategy). As expected, WAR and CAB are exactly complete and optimal. Moreover, they were the *only* strategies that achieved this. Despite having multilateral negotiations, fairness is still 99% (97%) for WAR (CAB) which is a 13% (11%) improvement over AOP with strategies and 7% (5%) over adapted strategies.

For the main experiment, WAR (CAB) increases agent advantage by 27% (21%) over AOP, while increasing the designer-score by 36% (39%). The higher advantage of the proposed WAR compared with CAB explains the results found earlier (Fig. 2). The improvements are even *higher* for all other datasets. Checking individual factors of the designer-score reveals that the main reason CAB achieves a higher score is its higher welfare. This may be due to the fact that WAR agents tend to negotiate longer (due to the Irrational stage) leading to more concession from the partner which reduces the overall welfare. This is supported by the finding that on average CAB was 44% times faster than WAR. It is interesting to note that both were much faster than both AOP variants.

Fig. 3 shows the advantage for each strategy tested when run under each condition in Table 3. The following results can be gleaned from this figure. Firstly, for every AOP strategy, the adapted version under TAU leads to higher advantage. Secondly, WAR achieves the highest advantage (0.73) when negotiating against other adapted strategies, then when negotiating against itself (0.71) and finally against CAB (0.58). Even the minimum of these is 10% higher than the best advantage of any AOP strategy.

Comparing adapted strategies under TAU to AOP variants, an improvement of at least 14% in advantage and 25% in designer-score was achieved for every individual strategy. This is lower than what WAR promises in all tested datasets but is still a major improvement justifying switching from AOP to TAU with the proposed adapter except in the extremely unlikely case that the modest 5% reduction in privacy is considered more important than the 14% minimum increase in advantage (and the 86.6% speed improvement). Moreover, these improvements do not only hold on average. Comparing each of the $7,768$ individual strategy combinations over all datasets (excluding 2024), the proposed adapter always led to improvements in both advantage and designer-score with only two exceptions in multilateral negotiation cases both in dominated strategies. These findings explain the fact that AOP strategies were driven to extinction relatively quickly in the game-theoretic analysis above.

The proposed adapter achieves a 27% improvement over AOP in designer-score. Moreover, TAU-native strategies outperform adapted strategies if played against themselves with CAB outperforming WAR by only 3%. As expected, both WAR and CAB achieve 100% completeness and optimality with WAR slightly outperforming CAB in fairness by 2% (despite the fact that 75% of the scenarios did have irrational outcomes). WAR outperformed adapted strategies and AOP strategies in all criteria. TAU negotiations under the proposed adapter were 87% faster than AOP negotiations showing that the speed we gain by avoiding repetitions more than compensates for the price of the sort and search operations needed for the adapter. Moreover, TAU-native strategies outperformed adapted strategies with the proposed WAR being 44% slower than CAB because of the irrational stage (Section 4.1).

Taken together, these results suggest that the proposed adapter delivers higher advantage and designer-scores with a slight reduction in privacy while the proposed WAR strategy can provide an even higher advantage and designer score at the price of even higher information revelation.

7 Related Work

7.1 MiCRO Strategy

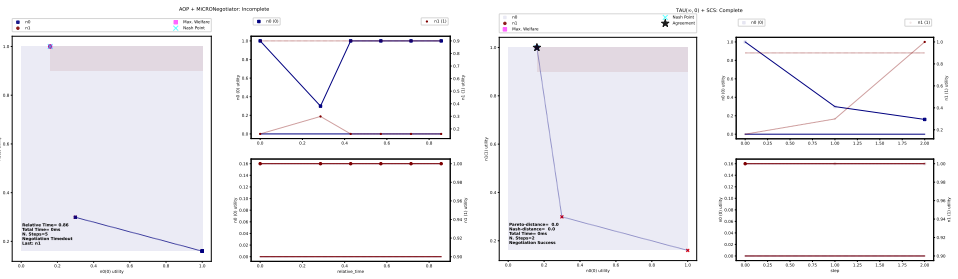

(a) MiCRO reaches no agreement (b) SCS reaches agreement

Fig. 4: A simple example of running MiCRO under AOP and CAB under TAU for the NiceOrDie ANAC scenario. MiCRO fails to reach agreement while CAB succeeded helped by the protocol

MiCRO [15] is the most related AOP strategy to CAB. When MiCRO plays against itself, its offering strategy (which is the main component of MiCRO) is *exactly* the same as CAB with only one crucial difference. MiCRO can run out of outcomes to offer if the scenario is not balanced. The paper presents an example. Fig. 4 provides another example taken from the ANAC domain NiceOrDie.

Given the design of AOP, MiCRO does not have the option of repeating offers at the end here because this can be exploited. A partner can just repeat its offers early and MiCRO will not know the difference. This is why the design of the protocol matters. In AOP, repeating an offer does not entail a penalty. In TAU, repeating an

offer will force the agent to never change it again. That is why TAU native strategies (like CAB and WAR) can assume that an agent repeating its offers has already run out of concessions to make and can concede themselves.

The offering policy of WAR, on the other hand, is different from MiCRO because of the added irrational phase. Again, MiCRO cannot employ such tactic as offering an irrational outcome under AOP can lead to immediate acceptance which is irrational for the agent to ever do (except if there is learning between negotiations which we do not consider in this work). TAU, by redefining agreement as requiring explicit offering and acceptance from *all* agents avoids this issues allowing agents to better explore the outcome space.

It can easily to be shown that MiCRO is not a best response to itself when the agent has some specific information about its partner. The general idea is to use some of the rounds early in the negotiation to offer outcomes that are known (or have high probability) of being worse than the next best outcome for MiCRO independent of their utility for MiCRO's opponent (distracting offers). These will be considered concessions and will force MiCRO to concede more than against itself.

In many practical situations, it will be easy to find outcomes that are — with high probability — not the next best for MiCRO. Consider a negotiation between a buyer using MiCRO and a seller knowing this fact (with the seller starting the negotiation). No matter how many other issues are involved in the negotiation, it is always the case that the buyer prefers lower prices. This means that the seller knows that any outcome that has any price less than the highest price will not be in the set of outcomes with highest utility for the buyer ($\Omega^{\top}_{buyer}$) and can confidently use the second highest price with any other value for the rest of the issues as a distracting first offer. It can then use MiCRO for the rest of the negotiation. This can either lead to a better or the same agreement for the seller but never a worse one.

It is instructive to see how much information we need to know about partner's preferences to find a better response for WAR. For MiCRO, we only need to know one outcome that is not on its set of best outcomes to find a better response. WAR starts the negotiation by offering all of its irrational outcomes which means that using *time-wasters* will not be effective as the opponent is already doing the same.

For the buyer-seller example discussed, the seller cannot just use any offer with a non-maximal price as its first offer (as it did with MiCRO) because this outcome will be accepted by WAR if it is better than disagreement and may become the negotiation outcome which means it should either be irrational for the buyer or a good outcome for the seller. Outcomes with low prices will likely not satisfy the first criterion and outcomes with high prices will likely not satisfy the second criterion. There is no way to know which outcomes are in the irrational set ($\underline{\Omega}_{WAR}$) before the negotiation without knowing a lot of information about its market position (and any outside options it may have).

In summary, it is, arguably, much easier to exploit information about partner preferences to find a better response for MiCRO compared with WAR.

7.2 MCP and Zeuthen

The most similar protocol to TAU is the Monotonic Concession Protocol (MCP) with its multilateral extension [31]. The Zeuthen [32] strategy is an almost stable strategy for MCP that involves monotonic concession based on calculating the *willingness to risk disagreement* for *both* agents introduced in 1930 by the Danish economist Frederik Zeuthen. Zeuthen involves slow concession in a way similar to MiCRO and CAB. MCP and Zeuthen are designed for cases of *complete information* about partners' preferences (to enable verifiability). Endriss 2006 [31] proposed the egocentric concession method for MCP which requires no direct access to the partner preferences but all agents still need to calculate their *willingness to risk disagreement* which *still* requires access to partner preferences. A major difference between TAU and MCP is that MCP defines concession by referring to the partner utility function which implies that agent preferences are common knowledge. No such assumption is implied by TAU. Moreover, MCP does not allow repeated offers with the same utility while TAU allows these repetitions as long as the outcome itself changes. MCP assumes that all offers are sent in parallel but TAU does not require this assumption.

8 Discussion, Limitations and Future Directions

Taken together, the empirical results reported in Section 6 suggest that rational agents should always choose to switch from AOP to TAU with the proposed adapter (as it leads to higher advantage for every evaluated strategy as long as they value their privacy less than 2.3 times their advantage). Moreover, theoretical and empirical results (Section 5 and 6) show that, Under TAU, they should switch to the proposed WAR strategy because it achieves the highest advantage. Fortunately, this actually leads to increased speed, completeness, optimality, fairness and welfare. Repeating the analysis on the remaining 15 datasets revealed the same patterns for all of them with the following exceptions: 1) In the smallest dataset (finalists of year 2010 with only 3 scenarios), no difference was detected in designer-score between WAR and adapted strategies. 2) A single strategy (Agent33) on the 2018 multilateral negotiation dataset did not benefit from switching to TAU. 3) Six agents in the 2024 dataset did not benefit from switching to TAU using the proposed adapter which is expected given the fact that agents had access to partner utility functions and were likely to only offer Pareto-optimal offers. The only surprise here is that the other 14 strategies still benefited form the switch. Moreover, the average advantage and designer-score is still highest for WAR despite having no access to partner utilities or even preferred outcome ordering. All others benefited.

Empirical Limitations

One limitation shared by WAR and CAB is that it is not clear how to extend them to continuous outcome spaces. Moreover, WAR reveals too much information about agent's preferences that may make it impractical for some applications. The current implementation of WAR, CAB, and the proposed adapter scale in $O(n_o \log n_o)$ which works fine for outcome spaces with up to few million outcomes but cannot scale to

larger outcome spaces. Using sampling instead of pre-ordering outcomes may resolve this issue [33].

Theoretical Limitations

Theorem 9 is established only for the case of bilateral negotiations on discrete outcome-spaces with absolutely no information about partner's preferences. In most practical situations, some information about partner's preferences are available which allows the agent to *learn* about the partner's preferences from its behavior which invalidates our proof for Theorem 9. This means that – in principle – it may be possible to find a better response for WAR when such information about partner's preferences are available. The resulting theoretical situation is interesting. A Subgame Perfect Equilibrium for automated bilateral negotiation with *full information* under AOP (with infinite rounds under time-pressure) is known since 1982 [9] and this paper introduces a Perfect Bayes Equilibrium for automated bilateral negotiation with *no information* under TAU for discrete outcomes spaces. Nevertheless, the most interesting practical situations are cases with *partial* information about partner preferences for which theoretical guarantees are not available. In the future, we will work on approximations of WAR (and variations of TAU) that reduce information revelation, improving scalability, and making use of partial information.

Ordinal Kalai Bargaining Solution

WAR and CAB were shown to find a *fair* agreement in the sense of finding the Ordinal Kalai Bargaining Solution. Can they guarantee finding the Kalai Bargaining Solutions itself? No. The reason is that these strategies have no access to the true utility value of any outcomes. They only use the partial ordering induced by the utility function if one is defined. As such, it is not possible for them to find the exact Kalai Bargaining solution. This limitation though is alleviated partially in practice as in most cases the Ordinal Kalai Bargaining Solution is close to the Kalai Bargaining Solution.

Effect of Privacy Valuation

As described in the paper, information revelation is not usually considered when evaluating automated negotiation strategies and protocols. Nevertheless, we believe that it is an important factor that needs to be taken into account in any such evaluation. In this work, the proposed approach is *guaranteed* to have a level of information revelation that is at least the same as the baselines and is expected — and found to be — to be higher. If we only consider the Advantage accrued to the agent, the proposed approach appears even better than in the results reported so far. In the paper, we used the same weight for Privacy and Advantage when calculating the agent score arguing that this is an upper limit on the expected valuation of privacy in realistic situations. Fig. 5 shows the effect of privacy weight on the improvement achieved by the proposed method in agent score for different dataset. For the largest main experiment, the agents need to value their privacy more than three times their utility/advantage/profit in order to prefer using AOP over the proposed approach. The smallest such cross-over was 1.2 for year 2015 and the largest was 8 for year 2010. In all cases agents valuing

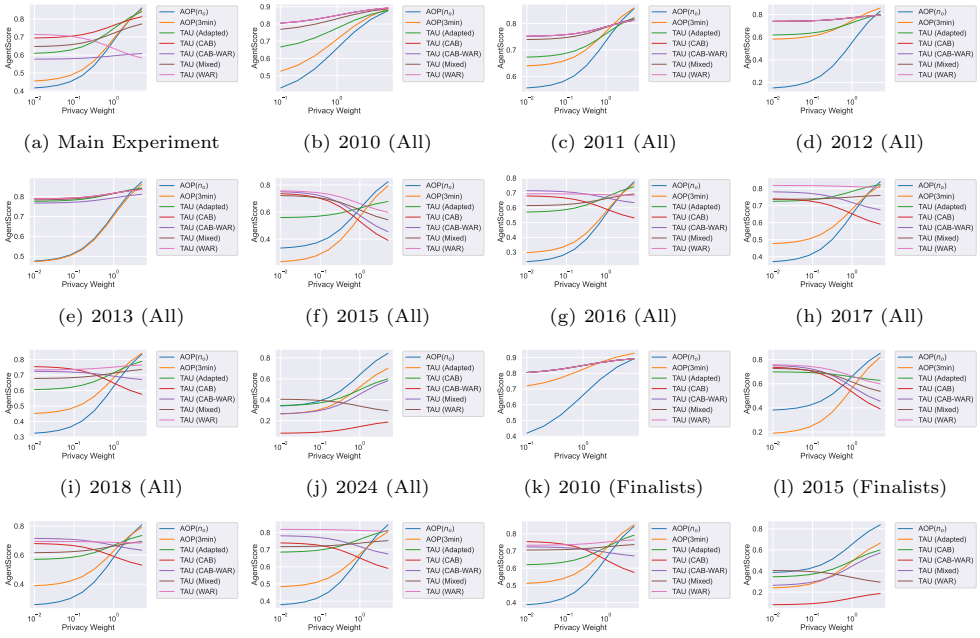

(a) Main Experiment (b) 2010 (All) (c) 2011 (All) (d) 2012 (All)

(e) 2013 (All) (f) 2015 (All) (g) 2016 (All) (h) 2017 (All)

(i) 2018 (All) (j) 2024 (All) (k) 2010 (Finalists) (l) 2015 (Finalists)

(m) 2016 (Finalists) (n) 2017 (Finalists) (o) 2018 (Finalists) (p) 2024 (Finalists)

Fig. 5: Effect of weight valuation on the improvement on agent score.

their advantage at least as much as their privacy will choose the proposed approach as stated in the paper.

Continuous Outcome Spaces

Continuous outcome-spaces provide a challenge to the proposed approach in two ways. Regarding the protocol, it is trivial to achieve the required constraint of non-repetition by simply changing offers by a small fraction that (assuming continuity of the utility function) does cause relatively small change in utility. The adaptation algorithm described in the paper needs to find the *nearest* outcome to the one received from the adapted strategy in utility. This is not even defined in continuous outcome spaces. An extension of TAU that can define some limit on similarity between offers may be able to extend the proposed method to continuous and hybrid outcome spaces but it is not clear whether this will still result in similar improvements to what we reported in this paper or not.

Repeating Negotiations

It is customary when running tournaments between negotiation strategies (e.g. in the official ANAC competition) to repeat each session multiple times (usually ten) to average over stochasticity in strategies' behavior. In this paper, we opted to run each session once due to the computational cost entailed by repetition given the huge number of sessions we ran with some of which taking hours because we did not limit

the time allowed in the rounds-limited sessions. Nevertheless, this should not affect our results as we are not comparing the underlying strategies themselves. Our goal is to compare the proposed protocol and adapter with AOP. This calls for having a larger number of negotiation scenarios (940 in our experiment) and larger combinations of strategies (2,529 in our experiment).

Learning between Negotiations

Strategies for ANAC 2017 and 2018 were designed to learn from repeated encounters with their partners. In our evaluations, each strategy encounters each partner once in the same scenario. This disables learning for these strategies. This means that any comparison between the underlying strategies for these years is unfair and we do not conduct any such comparison.

8.1 Scalability Analysis

TAU has an intrinsic time-limit which scales linearly with the size of the outcome-space $O(n_o)$. This means that it scales *exponentially* with the number of issues. Moreover the simplest implementation of the proposed adapter needs $O(n_o \log n_o)$ operations for sorting the outcome space. On the other hand, TAU's strategies have the advantage that they are never tempted to repeat offers unnecessarily (thanks to TAU disallowing it). Most AOP strategies — on the other hand — do waste a large fraction of negotiation time just repeating offers.

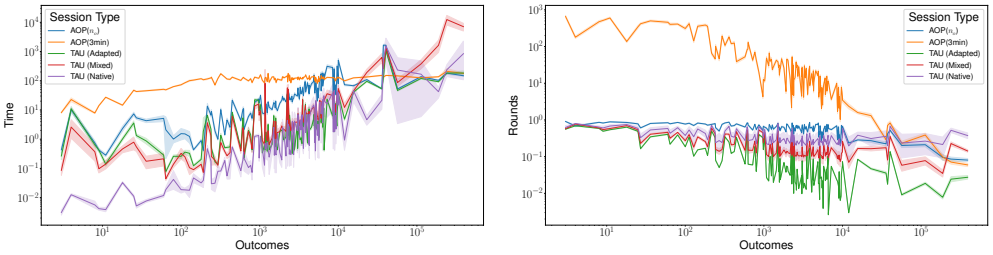

(a) Time

(b) Number of Rounds

Fig. 6: Scalability with outcome space size

Fig. 6 shows the effect of outcome-space size on the actual wall time (Fig. 6a) and number of rounds used relative to the outcome-space size (Fig. 6b). In terms, of real time, TAU outperformed both AOP methods for outcome spaces up to a tens of thousands of outcomes. This may be due to never repeating offers. For outcome spaces larger than this, and up to fifty thousand outcomes, both approaches start to require almost the same time. For outcome spaces larger than a hundred thousand, AOP outperforms TAU. This is most likely due to outcome space sorting implemented by the adapter. Fig. 6b shows the same comparison in terms of the number of rounds needed to end the negotiation. In this case, the proposed approach seems to almost always outperform AOP.

8.2 Notes about timing

The description of TAU is agnostic on the timing of offers as long as all responses to every offer are processed before another offer from the same agent is allowed.

If TAU is implemented serially using round-robin (as usually done for AOP), coming later in the ordering imposed entails a slight advantage as theagent may need to concede slightly less. Consider a bilateral zero-sum negotiation with two outcomes that are both rational to both negotiators. The agent that starts the second round will have to concede to its less favorite outcome which becomes the agreement. To avoid this asymmetry, we can either run all offering policies in parallel and serialize the evaluation of agreements after reception of responses randomly or just randomly assign an order of offering at every round.

It is interesting to note that AOP suffers from a bias toward agents that come earlier in the negotiation if they can determine the last round and have some information about the likelihood that each outcome is rational for their partner. In such cases, the first agent to offer in the last round becomes a de-facto dictator. The usual way to handle this issue is to run negotiation limited by real time (e.g. 3min) instead of the number of rounds but this solution reduces the risk instead of eliminating it. For TAU, the randomization proposed in this section completely eliminates any ordering effect.

9 Conclusion

This paper constructed a PBE strategy of TAU for bilateral negotiations with no information about partner preferences (WAR). We showed that WAR guarantees completeness for all discrete scenarios, and optimality and fairness for the subset of them with no repeated utility values. The paper also proposed a simple method for adapting existing AOP strategies to TAU. Empirical analysis showed that the proposed adapter and strategy can improve expected agent advantage, agreement rate, Pareto-optimality, fairness, and welfare at the expense of a mild increase in information revelation.

Supplementary information.

Appendix The supplementary file (appendix.pdf) provides detailed proofs of all the theorems described in this work as well as extended discussion and more detailed experimental results.

Data and Code The supplementary materials include the full source code as well as all raw data to reproduce all the results reported in this paper.

Declarations

- Funding: The author did not receive support from any organization for the submitted work.
- Conflict of interest: The author has no relevant financial or non-financial interests to disclose.
- Ethics approval and consent to participate: N/A

- Data availability: Experimental results as well as all negotiation scenarios used in this study are to be made available online on GitHub in a support page upon acceptance. All scenarios are currently available in the supplementary materials.
- Code availability: The source code for all the experiments reported in this paper are to be made available on GitHub in a support page upon acceptance. The source code is also included in the supplementary materials

If any of the sections are not relevant to your manuscript, please include the heading and write 'Not applicable' for that section.

Appendix A Notation

A.1 Used in the Paper

$\mathbb{N}$ Integers starting at zero.
$\mathbb{N}^+ \equiv \mathbb{N}$ Natural numbers (positive integers from 1).
$\mathbb{N}^+$ Natural numbers (positive integers starting at 1).
$\{0, 1\}$ The set of booleans (True, False).
$\Re$ The set of real numbers.
$\mathbb{S}(x)$ The power-set of set x (all possible subsets of x).
$\mathbb{S}^i(x)$ Sets of cardinality i belonging to $\mathbb{S}(x)$.
λ Negotiation scenario.
$\mathcal{P}$ Negotiation Protocol.
$\mathcal{N}^\lambda$ Negotiators for Scenario λ.
Π Strategies associated with negotiators for Scenario λ.
n_A The number of negotiators/agents/preferences in a negotiation.
$\mathcal{D}^\lambda$ Negotiation Domain for scenario λ (Outcome Space + Preferences).
Ω Outcome space (all possible agreements).
Ω_i^v The set of valid outcomes that can be offered by agent i.
n_o The cardinality of Ω (size of the outcome space).
ϕ Represents disagreement as an output of a negotiation, ending negotiation as an offer from an agent, and the empty set otherwise.
Ω^+ Outcome space or disagreement ϕ (all possible negotiation outcomes).
$\Omega_i^\top$ Best set of outcomes for negotiator i (none of these outcomes is better than any other in the set and any of them is better than any outcome not in the set).
$\Omega_i^\bot$ Worst set of outcomes for negotiator i (none of these outcomes is worse than any other in the set and any of them is worse than any outcome not in the set).
$\mathcal{F}^\lambda i$ Preferences of negotiator i in scenario λ.
u_i Utility function of negotiator i (cardinal).
r_i Relative rank of outcomes in Ω^+ for negotiator i (ordinal/cardinal).
$\hat{u}_i$ Value for agent i defined as normalized utility value (cardinal) or normalized relative rank (ordinal) for negotiator i.
$\omega_1 \succcurlyeq_i \omega_2$ Outcome ω_1 is not worse for negotiator i than outcome ω_2 given the preferences $\mathcal{F}^\lambda i$.
$\omega_1 \succ_i \omega_2$ Outcome ω_1 is better for negotiator i than outcome ω_2 given the preferences $\mathcal{F}^\lambda i$.

$\omega_1 \approx_i \omega_2$ Outcome ω_1 is neither better not worse for negotiator i compared with outcome ω_2 given the preferences $\mathcal{F}^\lambda i$.

$\mathcal{I}^\lambda$ Negotiator's Information Sets representing all information available to all negotiators in scenario λ.

$\mathrm{I}^\lambda i(\mathcal{F}^\lambda)$ Information available to negotiator i in scenario λ about preferences of all negotiators $\mathcal{F}^\lambda$ in that scenario.

$\mathbb{D}$ A set of domains.

$\mathcal{D}_i^\lambda$ All the information available to negotiator i in scenario λ.

$\overline{\Omega}_i \subseteq \Omega$*:* The subset of the outcome space Ω satisfying: $\omega \succcurlyeq_i \phi$.

$\overline{\Omega}^\lambda$*:* All possible rational agreements: $\bigcap_{i \in \mathcal{N}^\lambda} \overline{\Omega}_i$.

$\underline{\Omega}_i \subseteq \Omega$*:* The subset of the outcome space Ω satisfying: $\phi \succ_i \omega$.

$\underline{\Omega}^\lambda$*:* All irrational agreementst that should never occur defined as: $\bigcup_{i \in \mathcal{N}^\lambda} \underline{\Omega}_i$.

$\widehat{\Omega}^\lambda$*:* Win-win deals defined as the set of rational outcomes strictly better than disagreement:
$\left\{\omega \in \overline{\Omega}^\lambda : \forall i \in \mathcal{N}^\lambda \rightarrow \omega \succ_i \phi\right\}$.

$\mathbf{P}$*:* The pareto frontier defined as the set of rational outcomes that cannot be improved for one negotiator without making at least one other negotiator worst off: $\psi \in \mathbf{P} \subseteq \overline{\Omega}^\lambda \leftrightarrow \neg\exists \omega \in \Omega \setminus \{\psi\} : \exists i \in \mathcal{N}^\lambda : \omega \succ_i \psi \wedge \forall x \in \mathcal{N}^\lambda \setminus \{i\}\, \omega \succcurlyeq_i \psi$.

ω_*^λ The outcome of a negotiation using scenario λ.

$\lambda()$ The outcome of a negotiation using scenario λ (same as ω_*^λ).

$\mathcal{T}^k$ The negotiation trace showing all offers and responses (in order) in scenario λ up to round k (i.e. $\mathcal{T}^k = \Omega_i^k, \Omega_j^k, \mathbb{A}_i^k, \mathbb{A}_j^k$).

$\mathbb{T}^\lambda k$ The domain of all possible traces.

f The Filter used in representing a negotiation protocol.

σ The Evaluation Strategy used in representing a negotiation protocol.

$\triangleright$ Continue negotiation (returned by the evaluation policy to signal that the negotiation should go on).

π A negotiation strategy (consists of an offering policy and an acceptance policy).

η_i The offering policy of negotiator i.

ρ_i The Acceptance Policy of negotiator i.

$replace(\pi, \pi')$ Replace strategy π with strategy π'

Ω_x^j An ordered tuple of j most recent **unique** offers proposed by negotiator x so far ordered from earliest to latest.

$\mathbb{A}_x^j$ An ordered tuple of j most **unique** recent offers accepted by negotiator x so far ordered from earliest to latest.

$\Omega_i^v(k)$ The set of valid outcomes that can be offered by agent i at round k.

η_i^k The offer from negotiator i at round k.

A.2 Used only in this Appendix

a_{ji}^k The response from negotiator i at round k to the offer received from negotiator j.

$\hat{a}_i^k(\omega)$ The outcome ω will be accepted by negotiator i if it was offered at round k.

$\hat{a}_i(\omega)$ The outcome ω will be accepted by negotiator i at some time during the negotiation.

$J_i(\omega)$ The first round at which agent i would make offer ω. We set $J_i(\omega) = \infty$ if i will never make the offer during a negotiation.
$\Omega_i^v(k)$ The set of valid outcomes that can be offered by agent i at round k.
η_i^k The offer from agent i at round k.
a_{ji}^k The response from agent i at round k to the offer received from agent j.
$\hat{a}_i^k(\omega)$ The outcome ω will be accepted by agent i if it was offered at round k.
$\hat{a}_i(\omega)$ The outcome ω will be accepted by agent i at some time during the negotiation.
$J_i(\omega)$ The first round at which agent i would make offer ω. We set $J_i(\omega) = \infty$ if i will never make the offer during a negotiation.
T number of offers from the agent with maximum number of offers by the end of the negotiation (i.e. negotiation length).
$i \approx \pi$ agent i uses strategy π.
$\eta_i \approx \eta_\pi$ agent i uses π's offering policy.
$\rho_i \approx \rho_\pi$ agent i uses π's acceptance policy.

Appendix B Formal Definitions and Properties of Strategies

Let Λ_D be the set of all scenarios with discrete outcome spaces and let Λ_{DN} but the subset of Λ_D with no information about partner preferences (i.e. $\mathcal{D}^\lambda$ is discrete and $\mathrm{I}_i^\lambda = \{\mathcal{F}_i^\lambda\}$). Note that – crucially – this assumption implies that the agent cannot know which outcomes are rational for its partner. For the rest of this section, we always assume that we are considering Λ_{DN}.

Definition 1. *Let $\succ_l$ be a shared full ordering of the outcome space, CAB's offering policy can be formally written as:*

$$\eta_{CAB}(k) = \omega_i^k = \begin{cases} \sup_{\succ_l} \overline{\Omega}_i^v(k), & \text{if } \underline{\Omega}_i^v(k) = \{\} \wedge \overline{\Omega}_i^v(k) \neq \{\} \\ \omega_i^{k-1}, & \text{otherwise} \end{cases} \tag{B1}$$

where $\sup_{\succcurlyeq}(X)$ is the best outcome in the set X according to ordering $\succcurlyeq$, ω_i^{k-1} is the agent's last offer, and $\overline{\Omega}_i^v(k)$is the set of rational outcomes that agent i can offer:

$$\overline{\Omega}_i^v(k) \equiv \left\{\omega : \omega \in \Omega_i^v(k) \cap \overline{\Omega}_i \wedge \omega \succcurlyeq_i \psi \quad \forall \psi \in \Omega_i^v(k)\right\}$$

where $\Omega_i^v(k) \equiv \Omega \setminus \Omega_i^{k-1}$, and Ω_i^k is i's offers up to step but not including k.

This strategy offers outcomes in descending order of value for itself repeating when it cannot find any new outcomes to offer. Ties are broken using any predefined lexical ordering on the outcomes. It does not matter what this ordering is. What matters is that it is shared between all agents. Without this shared ordering, exact optimality is lost. It is important that agents not following this lexical ordering rule can never gain any utility because of it which means this rule is incentive compatible. We implicitly assume that $\omega_i^0 = \phi$ to handle the case with no rational outcomes.

Definition 2. *CAB's acceptance policy:*

$$\rho_{CAB}(\omega, k) = \mathbb{1}\left[\omega \succcurlyeq_i \phi \wedge \left(\left|\mathbb{A}_i^{k-1}\right| < 1 \vee \omega \succ_i A_i^{last}\right)\right] \tag{B2}$$

We will use the following properties of offers directly deducible from Eq. B1 and Eq. 1

- For CAB (but not WAR), a better offer is always offered before a worse offer:

$$\psi \succ_i \omega \wedge \omega_i^k = \omega \implies \psi \in \Omega_i^k \quad \forall \psi, \omega \in \Omega \tag{B3}$$

- For CAB and WAR, a better rational offer is always offered before a worse rational offer:

$$\psi \succ_i \omega \wedge \omega_i^k = \omega \implies \psi \in \Omega_i^k \quad \forall \psi, \omega \in \Omega \tag{B4}$$

- CAB and WAR never offer ϕ under TAU which means they never choose to explicitly end the negotiation.

$$\neg \exists i \in \mathbb{N} : \omega_i^i = \phi \tag{B5}$$

- CAB and WAR never repeat an offer before all offers not worse than it are tried:

$$\omega_i^k = \omega_i^{k-1} \implies \omega \in \Omega_i^{k-1} \quad \forall \omega \in \Omega : \omega \succ_i \omega_i^k \tag{B6}$$

- CAB and WAR never repeat an offer before trying all their rational outcomes:

$$\omega_i^k = \omega_i^{k-1} \implies \omega \in \Omega_i^{k-1} \qquad \forall \omega \in \overline{\Omega} \tag{B7}$$

Because win-win deals are a subset of rational outcomes:

$$\omega_i^k = \omega_i^{k-1} \implies \omega \in \Omega_i^{k-1} \qquad \forall \omega \in \widehat{\Omega} \tag{B8}$$

Considering the acceptance policies of CAB and WAR (Eq. B2, Eq. 2), we can directly deduce the following properties:

- If CAB or WAR are willing to accept an offer, they will accept any better offer anytime not later in the negotiation:

$$\omega \succ_i \psi \wedge \hat{a}_i^i(\psi) \implies \hat{a}_i^j(\omega) \qquad \forall j \leq i \tag{B9}$$

The reason that this holds is that CAB will only accept an offer if it is better than everything that was offered before and WAR accepts any rational offer. This means that both will accept any better offer at the same or any earlier round.

- If CAB (but not WAR) rejects an offer, it must have received an offer not worse for itself earlier:

$$\omega_j^k = \omega \wedge \neg a_{ji}^k \implies \exists \psi \in \Omega_j^{k-1} : \psi \succcurlyeq_i \omega \tag{B10}$$

- If CAB (but not WAR) actually accepts an offer at step k, it will reject any worse offer coming later:

$$\omega_j^k = \omega \wedge a_{ji}^k \implies \neg \hat{a}_i^j(\omega) \quad \forall j > i \forall \omega \in \Omega \tag{B11}$$

- If two outcomes belong to the set of outcomes accepted by CAB (but not WAR), they cannot be as good as each others for it (one of them must be strictly worse than the other):

$$\forall i \in \mathbb{N} \forall \omega, \psi \in \Omega \quad \omega \in i \wedge \psi \in i \implies \neg \omega \approx_i \psi \tag{B12}$$

We can also directly see the following properties of running TAU with CAB and WAR for all negotiators from the definitions above:

- If an outcome is offered by both negotiators and is acceptable to both negotiators by some time step k, the negotiation ends with this outcome as an agreement:

$$\exists \omega \in \Omega : \omega \in \Omega_i^k \wedge \hat{a}_i^j(\omega) \forall j < k \forall i \in \mathcal{N} \implies \omega_* = \omega \tag{B13}$$

- For an outcome to become the agreement of a negotiation, it must be offered by all negotiators and be acceptable by each of them before and at the round it received it:

$$\omega_* = \omega \implies \omega \in \Omega_i^{k_i} \wedge \hat{a}_i^j(\omega) \forall j < k_i \forall i \in \mathcal{N} \tag{B14}$$

Appendix C Full Proofs

C.1 TAU Always Terminates for Λ_{DN}

In the paper, Section "Tentative Acceptance Unique Offers Protocol", we claimed that TAU always terminates.

Theorem 1. *TAU always terminates in a finite number of steps T not greater than $n_o + 1$ for any strategy assignment rule on Λ_{DN}.*

$$T \leq n_o + 1$$

Proof. The size of valid set is the difference between the outcome space size and the number of unique offers for every negotiator:

$$n_v \equiv \min_{i \in \mathcal{N}} |\Omega_i^v| = \min_{i \in \mathcal{N}} n_o - |\Omega_i^\infty| \leq n_o - \max_{i \in \mathcal{N}} |\Omega_i^\infty| \tag{C15}$$

Let $n_v(k)$ be the size of the valid set of outcomes from all negotiators as defined above at round k.

For any round k if all negotiators repeat their last offer, the negotiation terminates:

$$\omega_i^k = \omega_i^k \forall i \in \mathcal{N} \Rightarrow terminates(TAU, k).$$

From this we conclude that if the negotiation is not terminated at round k, there is at least one negotiator not repeating at this round:

$$\neg terminates(TAU, k) \Rightarrow \exists y \in \mathcal{N} : \omega_j^k \neq \omega_j^k$$

$$\therefore \neg\text{terminates}(TAU, k) \Rightarrow \exists y \in \mathcal{N} : \quad |\omega_i^\infty|_k > |\omega_i^\infty|_{k-1} \tag{C16}$$

From Eq. C16 and Eq. C15,

$$\therefore \neg terminates(TAU, k) \Rightarrow n_v(k) < n_v(k-1)$$

Which means that if a negotiation did not end at round k, the size of the valid set of outcomes n_v will go down.

$$\therefore \lim_{k\to\infty} n_v(k) = 0$$

Negotiators will eventually run out of unique outcomes to offer and the negotiation will terminate.

$\therefore$ TAU will always terminate. □

Empirical support comes from the fact that in all our experiments TAU terminated.

C.2 TAU($\Upsilon_{WAR,CAB}$) is Exactly Rational on Λ_{DN}

Theorem 2. *TAU is Exactly Rational with any strategy assignment rule that always assigns strategies with the AcceptRational or AcceptBetter acceptance policies ($TAU(\Upsilon_{*B,*R})$) on Λ_{DN}.*

$$\neg\exists\lambda \in \Lambda_{DN} \exists i \in \mathcal{N}^\lambda : \lambda_{TAU}(\Upsilon_{*B,*R}) \prec_i \phi \tag{C17}$$

Proof. $\because$ From Eq. B2 and Eq. 2, an agent using AcceptBetter or AcceptRational will only accept an offer if it is better than disagreement.

$\because$ From Eq. B14, an agreement can only be reached if all agents accept it.

$\therefore$ An outcome that is worse than disagreement for any negotiator using AcceptBetter (Eq: B2) or AcceptRational (Eq: 2) cannot be the output of a negotiation.

$$\therefore \omega_*^\lambda \neq \phi \implies \omega_* \in \overline{\Omega} \tag{C18}$$

□

Lemma 1. *$TAU(\Upsilon_{CAB,WAR})$ is exactly rational.*

Proof. This follows directly from Theorem 2 the fact that CAB uses AcceptBetter and WAR uses AcceptRational. □

Empirical support comes from the fact that in all cases in all experiments, TAU never ended with an agreement with a utility lower than the reserved value of any negotiator involved. This can be seen in from the fact that Pareto Optimality was always exactly one implying that all outcomes are on the Pareto-frontier which contains only rational outcomes by definition.

C.3 TAU(Υ_{CAB}) is Exactly Optimal on Λ_{DN}

Proposition 1 (Acceptability). *In a bilateral negotiation of Λ_{DN} with TAU(Υ_{CAB}), an agent accepts an outcome iff it is Pareto-efficient or as good as a Pareto-efficient outcome for its partner.*

Proof. The acceptance policy of CAB (AcceptBetter) will only accept an outcome if it is better than everything it received so far (Eq. B9). Therefore, for an agent i running CAB:

$$\hat{a}_i^i(\omega) \implies \omega \succ_i \psi \qquad \forall \psi \in \Omega_j^i$$

Moreover, CAB offers outcomes in descending order. Therefore, for an agent j running CAB:

$$\omega_j^i = \omega \implies \omega \succcurlyeq_j \psi \qquad \forall \psi \in \Omega \setminus \Omega_j^i$$

Agent i will accept ω iff it was offered by j at some step i at which it was acceptable by i.

$$a_{ji}^k \implies \omega_j^k = \omega \wedge \hat{a}_i^j(\omega)$$

$$\therefore a_{ji}^k \implies \omega \succ_i \psi \forall \psi \in \Omega_j^i \wedge \omega \succcurlyeq_j \psi \forall \psi \in \Omega \setminus \Omega_j^i$$

The inequality case gives the Pareto frontier:

$$\psi \in \Omega_j^i \wedge \omega \succ_j \psi \forall \psi \in \Omega \setminus \Omega_j^i \implies \omega \in \mathbf{P}$$

The equality case, gives outcomes that are as good as a Pareto efficient outcome for agent j:

$$\psi \in \Omega_j^i \wedge \exists \psi \in \Omega \setminus \Omega_j^i \omega \approx_j \psi \implies \omega \in \{\alpha \in \Omega : \alpha \approx_j \omega\}$$

But if there is another outcome that is the same for j as ω, the Pereto-frontier must have at least one such outcome:

$$\psi \in \Omega_j^i \wedge \exists \psi \in \Omega \setminus \Omega_j^i \omega \approx_j \psi \implies \exists \alpha \in \mathbf{P} : \alpha \approx_j \omega$$

Taken together:

$$\therefore a_{ji}^k \implies \omega \in \mathbf{P} \vee \exists \alpha \in \mathbf{P} : \alpha \approx_j \omega$$

□

As a simple example illustrating a case where Pareto-inefficient outcome is accepted consider the following outcome space ($\Omega = \{\omega_1, \omega_2, \omega_3\}$) with the following preferences: $\omega_1 \approx_j \omega_2 \succ_j \omega_3 \succ_j \phi$ and $\omega_3 \succ_i \omega_2 \succ_i \omega_1 \succ_i \phi$. Assume that j started by offering ω_1, as this is the first offer for agent i, it will accept it. Nevertheless ω_1 is Pareto-inefficient because ω_2 weakly dominates it. Note that in this case, the other condition is satisfied because $\omega_1 \approx_j \omega_2$. [4]

[4] There is nothing special about ω_1 being the first offer here. We could have just added two other outcomes in either side of the ordering and got the same result.

Proposition 2 (Pareto →Accepted)**.** *In a bilateral negotiation of Λ_{DN} with TAU(Υ_{CAB}), every Pareto-efficient outcome is accepted by at least one agent or has the same value as the agreement for all agents.*

Proof. From Theorem 5 (Completeness), we know that if the negotiation ends in disagreement, there are NO rational outcomes, and no outcomes on the Pareto frontier and the proposition holds trivially. Therefore, we need only concern ourselves with proper agreements in finite time $T \leq n_o$:

$$\omega_* \in \Omega$$

Let's assume that there is another outcome ψ with unidential value for all negotiators on the Perto-frontier that is not accepted by one of the negotiators:

$$\psi \in \mathbf{P} \wedge \neg\psi \approx_z \omega_* \forall z \in \mathcal{N} \wedge \exists i \in \mathcal{N} \neg \exists k \leq T : a^k_{ji}$$

This means it is either never offered by j or was not acceptable by i when it was offered by j:

$$\psi \notin \Omega_j \vee \exists k \leq T : \omega^k_i = \psi \wedge \neg \hat{a}^i_k(\psi)$$

Case 1: ψ not offered $\psi \notin \Omega_j$

Because the negotiation ends with agreement ω_*, all outcomes better for any agent z must have been offered (descending order of offering):

$$\omega_* \in \Omega \implies \alpha \in \Omega_z \qquad \forall \alpha \succ_z \omega_* \forall z \in \mathcal{N}$$

But,

$$\because \psi \in \mathbf{P}$$
$$\therefore \exists z \in \mathcal{N} \psi \succ_z \omega_* \vee \psi \approx_z \omega_* \forall z \in \mathcal{N}$$

The first subcase implies that ψ must have already been offered by z before the agreement ω_*:

$$\exists z \in \mathcal{N} \psi \succ_z \omega_* \implies \exists k < T : \omega^k_z = \psi$$

which contradicts the assumption of Case 1. The second subcase implies that ψ must have the same value for all agents which directly contradicts our assumption about ψ.

$\therefore$ Case 1 leads to a contradiction.

Case 2: ψ was not acceptable $\exists k \leq T : \omega^k_i = \psi \wedge \neg \hat{a}^k_i(\psi)$

For ψ not to be acceptable by agent i at step k, a better outcome (for i) must have been offered and accepted by i before that:

$$\neg \hat{a}^k_i(\psi) \implies \exists \omega \in \Omega \exists z \in \mathcal{N} \exists j < k : \omega^j_z = \omega \wedge a^j_{zi} \wedge \omega \succ_i \psi$$

$$\therefore \neg \hat{a}^k_i(\psi) \implies \omega \succ_i \psi$$

But for ω to be offered by z before ψ, ω must be at least as good as ψ for z:

$$\omega^j_z = \omega \wedge \omega^i_z = \psi \implies \omega \succcurlyeq_i \psi$$

$$\therefore \omega \succcurlyeq_i \psi \wedge \omega \succ_i \psi$$
$$\therefore \psi \notin \mathbf{P}$$

Which contradicts our assumption.

$\therefore$ Case 2 leads to a contradiction.

$\therefore$ Both cases lead to a contradiction and there are no other possibilities.

$\therefore \psi$ cannot exist.

$$\therefore \neg\exists\psi : \psi \in \mathbf{P} \wedge \neg\psi \approx_z \omega_* \forall z \in \mathcal{N} \wedge \exists i \in \mathcal{N} \neg\exists k \leq T : a^k_{ji}$$

□

Theorem 3 (Optimal)**.** *TAU(Υ_{CAB}) is Exactly Optimal with any strategy assignment rule that uses only CAB on Λ_{DN}.*

$$\omega^\lambda_* \in \mathbf{P}^\lambda \cup \phi \qquad \forall \lambda \in \Lambda_{DN} \tag{C19}$$

Proof. We will assume that both negotiators are using CAB. We need to consider too cases.

Case 1 (No win-win deals): $\left|\widehat{\Omega}\right| = 0$.

In this case, there are no outcomes that strictly dominate disagreement.

$\because$ TAU(Υ_{CAB}) is Exactly Rational (Lemma 1)

$\therefore$ It will either lead to disagreement or to some agreement $\omega_* \succcurlyeq_i \phi \forall i \in \mathcal{N}$. This is the optimal result.

$$\left|\widehat{\Omega}\right| = 0 \implies optimal(TAU) \tag{C20}$$

Case 2 (some win-win deals): $\left|\widehat{\Omega}\right| > 0$.

There are some win-win outcomes which means that the Pareto-frontier is not empty. We now need to show that: $\left|\widehat{\Omega}\right| > 0 \implies \omega_* \in \mathbf{P}$

We will use a proof by contradiction. Assume that $\omega_* \notin \mathbf{P}$ and that agreement was reached at round T:

From Theorem 1, we know that ω_* must be rational: $\omega_* \in \overline{\Omega}$.

From the definition of pareto-optimality:

$$\therefore \exists\psi \neq \omega_* \exists z \in \mathcal{N} : \psi \succcurlyeq_i \omega_* \forall i \in \mathcal{N} \wedge \psi \succ_z \omega_*$$

$$\because \psi \succ_z \omega_* \tag{C21}$$

because CAB always offer outcomes from best to worst (Eq. B3), ψ must have been offered by z:

$$\therefore J_z(\psi) < J_z(\omega_*)$$

but ω_* was also offered by z to become an agreement in the first place (Eq. B14):

$$\omega_* \in \Omega_z \tag{C22}$$

$$\therefore J_z(\psi) < J_z(\omega_*) \leq T$$
$$\therefore \psi \in \Omega_z \tag{C23}$$

ψ was offered by z earlier. Now consider how would other agents respond to this offer. ω_* must be acceptable by all agents at the end of the negotiation ($\hat{a}_i^T(\omega_*)$):

$$\because \hat{a}_i^T(\omega_*) \wedge \psi \succcurlyeq_i \omega_* \quad \forall i \in \mathcal{N}$$

From Eq. B9:

$$\therefore \hat{a}_i^i(\psi) \quad \forall i < T \quad \forall i \in \mathcal{N} \setminus \{z\}$$

Because ψ is acceptable to all agents and is offered at step $k \equiv J_z(\psi)$ by z, it must have been accepted by everyone:

$$\because \omega_i^k = \psi \wedge \hat{a}_i^k(\psi) \forall i \in \mathcal{N} \setminus \{z\}$$

$$\therefore a_{zi}^k \forall i \in \mathcal{N} \setminus \{z\} \tag{C24}$$

Consider any agent $i \neq z$, we know that it accepted ψ and ω:

$$\because \psi \in \ \wedge \omega_* \in \ \wedge \psi \succcurlyeq_i \omega_* \tag{C25}$$

From Eq. B12[5]:

$$\therefore \psi \succ_i \omega_*$$

This is exactly the form of Eq. C21, and by the same steps (applied for all agents) we can arrive at:

$$\psi \in \Omega_i \forall i \in \mathcal{N} \tag{C26}$$

and

$$a_{iz}^k \forall i \in \mathcal{N} \forall z \in \mathcal{N} \setminus \{i\} \tag{C27}$$

Which means that ψ is offered and accepted by every negotiator and therefore must be the agreement:

$$\psi = \omega_*$$

This is a contradiction as we assumed $\omega_* \neq \psi$.

$$\therefore \left|\widehat{\Omega}^\lambda\right| > 0 \implies \omega_* \in \mathbf{P}$$

Combining the two cases above completes the proof. □

Empirical support comes from the fact that in all experiments, TAU(Υ_{CAB}) found agreements on the Pareto-frontier. This can be seen from the fact that Pareto Optimality was always exactly one which implies that all outcomes are on the Pareto-frontier.

[5] This is exactly where the proof breaks for WAR as this is not guaranteed in that case.

C.4 TAU(Υ_{WAR}) is NOT Exactly Optimal on Λ_{DN}

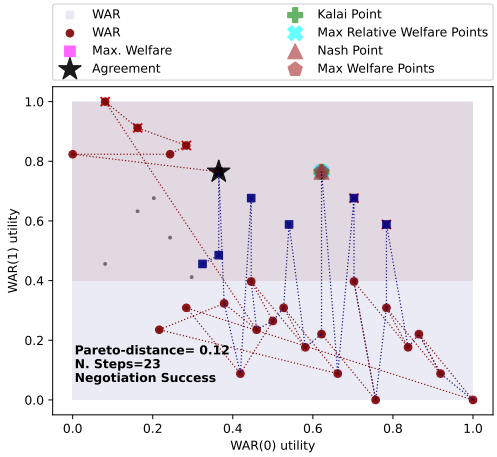

Fig. C1: Example from ANAC 2016 domains (Defensive Charms) showing that TAU(Υ_{WAR}) is not exactly optimal.

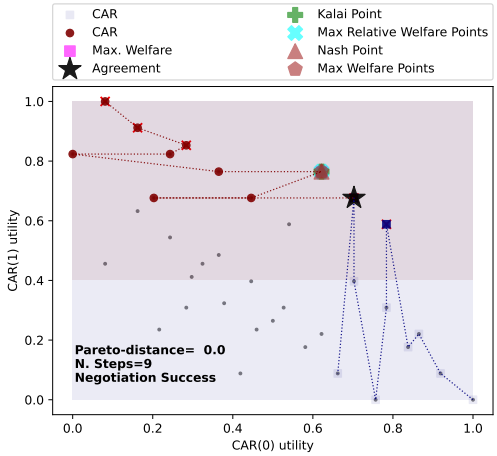

Fig. C2: Example from ANAC 2016 domains (Defensive Charms) showing that TAU(Υ_{CAB}) can – in this case – find an agreement on the Perto-frontier. Compare this to WAR (Fig. C2)

Proposition 3 (WAR is not exactly optimal)**.** *TAU(Υ_{WAR}) is NOT exactly optimal.*

Proof. We need only provide a single counter example to prove this proposition. Fig. C1 shows an example from the Defensive Charms domain (ANAC).[6] A smaller example can be constructed with only three outcomes $(\omega_1, \omega_2, \omega_3)$ with the following rankings $\omega_1 \succ_i \omega_2 \approx_i \omega_3 \succ_i \phi$ and $\omega_3 \succ_j \omega_2 \succ_j \omega_1 \succ_j \phi$ for agents i and j. Agent i can offer the outcomes in the order $\langle \omega_1, \omega_2, \omega_3 \rangle$ while Agent j is forced to offer them in the following order $\langle \omega_3, \omega_2, \omega_1 \rangle$. Let's assume that at the second round, agent j offered first. In this case, once agent i offers ω_2, it will become the agreement: $\omega_*^\lambda = \omega_2$ but $\omega_2 \notin \mathbf{P}$ because ω_3 weakly dominates it. □

The only case this can happen in a bilateral negotiation (the only point at which the previous proof failed) is when there are two outcomes ω, ψ such that $(\approx_i (\omega, \psi)$ and $\succ_j (\psi, \omega))$. Moreover, both outcomes are rational for both agents. If j have already offered ω (implying it also offered ψ) by the time i offers ω, and it happened that i offers ω before ψ (which is only possible because $\approx_i (\omega, \psi)$), then ω will be offered by both agents and accepted by both at that point and will become the agreement despite the fact that it is weakly dominated by ψ and cannot lie on the Pareto-frontier. It is important to note that the whenever TAU(Υ_{WAR} reaches a non-optimal agreement, it will be with no loss of utility for the agent that accepted it. The reason CAB is not subject to this problem is that it only accepts outcomes that are strictly better than what it accepted before so in our example above agent i will not accept ω_2 when offered to it because it already accepted ω_3 and $\omega_3 \approx_i \omega_2$.

C.5 TAU(Υ_{WAR}) is Exactly Optimal on Λ_U

Theorem 4. *TAU(Υ_{WAR}) is exactly optimal for Λ_U.*

Proof. The proof follows exactly the same as in the proof for Theorem 3 up until and including Eq. C25. Instead of using Eq. B12 to continue the proof, we use the fact that no two outcomes have the same utility value in any scenario that belongs to Λ_{DN} to show that:

$$\therefore \psi \succ_i \omega_*$$

The rest of the proof follows exactly as in Theorem 3. □

C.6 TAU(Υ_{CAB}) is Exactly Complete on Λ^B_{DN}

Proposition 4 (CAB Completeness)**.** *TAU(Υ_{CAB}) is Exactly Complete for bilateral negotiations on Λ_{DN}.*

$$\left|\widehat{\Omega}^\lambda\right| \neq 0 \rightarrow \omega_*^\lambda \in \widehat{\Omega}^\lambda \quad \forall \lambda \in \Lambda^B_{DN} \tag{C28}$$

[6] CAB still finds a pareto efficient outcome in this case as shown in Fig. C2.

Proof. We will assume that both negotiators are using CAB. For TAU(Υ_{CAB}) to be incomplete, we must have at least one win-win deal in a negotiation that ends either with disagreement or an agreement not in the set of win-win deals. Formally:

$$\neg complete(TAU) \implies \left|\widehat{\Omega}\right| > 0 \wedge \omega_* \notin \widehat{\Omega}$$

$\therefore$ we only need to consider negotiations with $\left|\widehat{\Omega}\right| > 0$ and show that in all such negotiations $\omega_* \neq \phi \wedge \omega_* \in \widehat{\Omega}$.

$\because$ The negotiation always terminates (Theorem 1) after a finite number of rounds T.

$\therefore$ Either an agreement is reached or the negotiation ends in disagreement because of repetition.

Case 1 (agreement): $\left|\widehat{\Omega}\right| > 0 \wedge \omega_* \neq \phi$

Let's consider first the case where a negotiation ends with an agreement: $\omega_* \neq \phi$ From C17, we know that the agreement is rational:

$$\omega_* \neq \phi \implies \omega_* \in \overline{\Omega}$$

Because TAU (with CAB) is Exactly Optimal (Theorem 3), we know:

$$\omega_* \neq \phi \implies \omega_* \in \mathbf{P}$$

From the definition of the Pareto-frontier:

$$\left|\widehat{\Omega}\right| > 0 \implies \mathbf{P} \subset \widehat{\Omega}$$

$$\therefore \left|\widehat{\Omega}\right| > 0 \wedge \omega_* \neq \phi \implies \omega_* \in \widehat{\Omega} \tag{C29}$$

Case 2 (disagreement) $\left|\widehat{\Omega}\right| > 0 \wedge \omega_* = \phi$**.**

Can this happen? From Eq B14, disagreement ($\omega_* = \phi$) can happen only in one of two ways:

1. One negotiator ends the negotiation: $\exists i \in \mathcal{N} : \omega_i^T = \phi$. This cannot happen because CAB never explicitly end a negotiation (See Eq. B5).
2. Both negotiators are repeating offers: $\forall i \in \mathcal{N} : \omega_i^T = \omega_i^{T-1}$. We need to show that this is also cannot happen when there are any win-win deals.

$\therefore$ We only need to show that it is never the case that $\left|\widehat{\Omega}\right| > 0 \wedge \exists T \in \mathbb{N} : \omega_i^T = \omega_i^{T-1} \quad \forall i \in \mathcal{N}$.

We will do a proof by contradiction again. To arrive at round T, all rational outcomes for i must have been offered by the agent (because CAB and WAR never repeat before trying all of their rational outcomes – Eq B7):

$$\omega_i^T = \omega_i^{T-1} \implies \omega \in \Omega_i^{:T-1} \quad \forall \omega \in \overline{\Omega}_i \tag{C30}$$

This is true for both negotiators which means:

$$\omega_i^T = \omega_i^{T-1} \forall i \in \mathcal{N} \implies \omega \in \Omega_i^{:T-1} \quad \forall \omega \in \overline{\Omega}_i \forall i \in \mathcal{N} \tag{C31}$$

This means that all possible rational outcomes for both negotiators have been offered.

$$\begin{gathered}\because \widehat{\Omega} \subseteq \overline{\Omega} \subseteq \overline{\Omega}_i \forall i \in \mathcal{N} \\ \therefore \omega \in \Omega_i^T \quad \forall \omega \in \widehat{\Omega} \quad \forall i \in \mathcal{N}\end{gathered} \tag{C32}$$

Now consider negotiator i that rejects win-win deal $\omega \in \widehat{\Omega}$. From Eq. B10, we know that a better or equivalent offer for i must have been proposed earlier by the other agent j:

$$\omega_j^k = \omega \wedge \omega \in \widehat{\Omega} \wedge \neg a_{ji}^k \implies \exists \psi \in \Omega_j^{:k-1} \cap \widehat{\Omega} : \psi \succcurlyeq_i \omega \tag{C33}$$

Because ψ was offered first, it is also not worse than ω for agent j:

$$\begin{gathered}\because J_j(\psi) < J_j(\omega) \\ \therefore \psi \succcurlyeq_j \omega \wedge \psi \succcurlyeq_i \omega \\ \therefore \psi \succcurlyeq \omega \\ \because \omega \in \widehat{\Omega} \\ \therefore \psi \in \widehat{\Omega}\end{gathered}$$

This means that a negotiator can reject a win-win deal only if it accepted one earlier. Therefore no negotiator rejects all win-win deals.

$$\forall i \in \mathcal{N} \exists \omega_i \in \widehat{\Omega} : \omega_i \in \mathbb{A}_i$$

Let $Q_i \subset \widehat{\Omega}$ be the set of win-win deals rejected by negotiator i and $q_i^\top$ ($q_i^\perp$) be the best (worst) such outcome in Q_i for i. By definition

$$q_i^\top \succcurlyeq_i q_i^\perp \wedge q_j^\top \succcurlyeq_j q_j^\perp \tag{C34}$$

Because a rejected offer must have a non-dominated offer accepted before it (Eq. B10), we have:

$$\exists \alpha_i \in \widehat{\Omega} - Q_i : \alpha_i \succcurlyeq_i q_i^\top \tag{C35}$$

Because the negotiation did not end in agreement until step T and all win-win deals were offered by both negotiators (Eq. C55), we must have no win-win deal that is accepted by both negotiators:

$$\left| \widehat{\Omega} \setminus (Q_i \cup Q_j) \right| = 0 \tag{C36}$$

From Eq. C35, Eq. C34, and E. C36 (and their symmetric counterparts), we get:

$$q_j^\top \succcurlyeq_i \alpha_i \succcurlyeq_i q_i^\top \tag{C37}$$

By symmetry

$$q_i^\top \succcurlyeq_j \alpha_j \succcurlyeq_j q_j^\top \tag{C38}$$

From Eq. C37 and Eq. C38 we get:

$$q_i^\top \approx_i q_j^\top \wedge q_i^\top \approx_j q_j^\top \tag{C39}$$

In summary:

$$\left|\widehat{\Omega} \setminus (Q_i \cup Q_j)\right| = 0 \implies q_i^\top \approx_i q_j^\top \wedge q_i^\top \approx_j q_j^\top \tag{C40}$$

Because of the lexical ordering rule in Eq. B1, we know that one of these two outcomes must have been received first by both negotiators. Moreover, by definition, they both have been rejected. Therefore, there is another outcome that must have been accepted by both negotiators before them which must be a win-win deal and cannot be a member of either Q_i or Q_j:

$$\exists \psi \in \widehat{\Omega} : \psi \notin Q_i \wedge Q_j \tag{C41}$$

From Eq. C40 and Eq. C41 we have:

$$\left|\widehat{\Omega} \setminus (Q_i \cup Q_j)\right| = 0 \implies \left|\widehat{\Omega} \setminus (Q_i \cup Q_j)\right| > 0 \tag{C42}$$

which is a contradiction disproving the third way of ending the negotiation with no agreement.

$$\therefore \left|\widehat{\Omega}\right| > 0 \wedge \omega_i^T = \omega_i^{T-1} \forall i \in \mathcal{N} \implies \omega_* \neq \phi$$

$$\therefore \left|\widehat{\Omega}\right| > 0 \wedge \omega_* = \phi \implies \omega_* \neq \phi \tag{C43}$$

which is a contradiction for Case 2. From Eq. C43 and Eq. C29, we get:

$$\therefore \left|\widehat{\Omega}\right| > 0 \implies \omega_* \in \widehat{\Omega} \tag{C44}$$

$\therefore$ TAU is Exactly Complete on Λ_{DN}^B for any strategy assignment rule that employ only CAB.

□

This is supported empirically in the paper by the fact that in all experiments, the agreement rate (completeness) of TAU with CAB when negotiating against itself was always 100%.

C.7 TAU(Υ_{WAR}) is Exactly Complete on Λ_{DN}^B

Proposition 5 (Completeness). *TAU(Υ_{WAR}) is Exactly Complete for bilateral negotiations on Λ_{DN}.*

$$\left|\widehat{\Omega}^\lambda\right| \neq 0 \rightarrow \omega_*^\lambda \in \widehat{\Omega}^\lambda \quad \forall \lambda \in \Lambda_{DN}^B \tag{C45}$$

Proof. We will assume that both negotiators are using WAR. For TAU(Υ_{WAR}) to be incomplete, we must have at least one win-win deal in a negotiation that ends either with disagreement or an agreement not in the set of win-win deals. Formally:

$$\neg complete(TAU) \implies \left|\widehat{\Omega}\right| > 0 \wedge \omega_* \notin \widehat{\Omega}$$

$\therefore$ we only need to consider negotiations with $\left|\widehat{\Omega}\right| > 0$ and show that in all such negotiations $\omega_* \neq \phi \wedge \omega_* \in \widehat{\Omega}$.

$\because$ The negotiation always terminates (Theorem 1) after a finite number of rounds T.

$\therefore$ Either an agreement is reached or the negotiation ends in disagreement because of repetition.

Case 1 (agreement): $\left|\widehat{\Omega}\right| > 0 \wedge \omega_* \neq \phi$

Let's consider first the case where a negotiation ends with an agreement: $\omega_* \neq \phi$

From C17, we know that the agreement is rational:

$$\omega_* \neq \phi \implies \omega_* \in \overline{\Omega}$$

From the agreement assumption and Eq. B14, the agreement was offered and was acceptable by all negotiators.

$$\omega_* \neq \phi \implies \omega_* \in \Omega_i^T \wedge \hat{a}_i^j(\omega_*^\lambda) \forall j < T \forall i \in \mathcal{N} \tag{C46}$$

Now consider a rational outcome that is not a win-win deal: $\psi \in \overline{\Omega} \setminus \widehat{\Omega}$

$$\because \psi \in \overline{\Omega} \setminus \widehat{\Omega}$$

$$\therefore \omega \succ_i \psi \qquad \forall i \in \mathcal{N} \forall \omega \in \widehat{\Omega} \tag{C47}$$

From Eq. B8, every outcome in $\widehat{\Omega}$ was offered before ψ by all negotiators.

From Eq. B9, assuming that ψ is acceptable at time T implies that every outcome in $\widehat{\Omega}$ will be acceptable if offered anytime before T. From Eq. B3, we then have:

$$\omega_* = \psi \implies \hat{a}_i^T(\psi) \longrightarrow \hat{a}_i^k(\omega) \forall k < T \tag{C48}$$

From Eq C47 and Eq C48 we have:

$$\therefore \omega_* = \psi \implies \exists i_i < T : \omega_i^{i_i} = \omega \wedge \hat{a}_j^{i_i}(\omega) \forall i, y \in \mathcal{N} \forall \omega \in \widehat{\Omega}$$

From Eq. B3, and Eq C47, we can conclude that all win-win deals were offered during the negotiation by all negotiators:

$$\omega_* \notin \widehat{\Omega} \implies \omega_* \in \Omega_i^T \quad \forall \omega \in \widehat{\Omega} \quad \forall i \in \mathcal{N} \tag{C49}$$

From Eq. B9, and Eq C47, we can conclude that all win-win deals were acceptable during the negotiation by all negotiators:

$$\omega_* \notin \widehat{\Omega} \implies \hat{a}_i^j(\omega) \quad \forall j < T \forall \omega \in \widehat{\Omega} \forall i \in \mathcal{N} \tag{C50}$$

Substituting Eq. C49, Eq. C50 into Eq. B13, we get:

$$\begin{aligned} \omega_* \notin \widehat{\Omega} &\implies \hat{a}_i^j(\omega) \wedge \omega_* \in \Omega_i^T \quad \forall j < T \forall \omega \in \widehat{\Omega} \forall i \in \mathcal{N} \\ &\implies \omega_* \in \widehat{\Omega} \end{aligned} \tag{C51}$$

We have proven a contradiction.

$$\therefore \left|\widehat{\Omega}\right| > 0 \wedge \omega_* \neq \phi \implies \omega_* \in \widehat{\Omega} \tag{C52}$$

Case 2 (disagreement) $\left|\widehat{\Omega}\right| > 0 \wedge \omega_* = \phi$**.**

Can this happen? From Eq B14, disagreement ($\omega_* = \phi$) can happen only in one of two ways:

1. One negotiator ends the negotiation: $\exists i \in \mathcal{N} : \omega_i^T = \phi$. This cannot happen because WAR never explicitly end a negotiation (See Eq. B5).
2. Both negotiators are repeating offers: $\forall i \in \mathcal{N} : \omega_i^T = \omega_i^{T-1}$. We need to show that this is also cannot happen when there are any win-win deals.

$\therefore$ We only need to show that it is never the case that $\left|\widehat{\Omega}\right| > 0 \wedge \exists T \in \mathbb{N} : \omega_i^T = \omega_i^{T-1} \quad \forall i \in \mathcal{N}$.

We will do a proof by contradiction again. To arrive at round T, all rational outcomes for i must have been offered by the agent (because WAR never repeats before trying all of its rational outcomes – Eq B7):

$$\omega_i^T = \omega_i^{T-1} \implies \omega \in \Omega_i^{:T-1} \quad \forall \omega \in \overline{\Omega}_i \tag{C53}$$

This is true for both negotiators which means:

$$\omega_i^T = \omega_i^{T-1} \forall i \in \mathcal{N} \implies \omega \in \Omega_i^{:T-1} \quad \forall \omega \in \overline{\Omega}_i \forall i \in \mathcal{N} \tag{C54}$$

This means that all possible rational outcomes for both negotiators have been offered.

$$\begin{aligned} &\because \widehat{\Omega} \subseteq \overline{\Omega} \subseteq \overline{\Omega}_i \forall i \in \mathcal{N} \\ &\therefore \omega \in \Omega_i^{:T} \quad \forall \omega \in \widehat{\Omega} \quad \forall i \in \mathcal{N} \end{aligned} \tag{C55}$$

WAR never rejects a win-win deal $\omega \in \widehat{\Omega}$ because it accepts all rational offers. From Eq. C55, we now know that all win-win deals have been offered and accepted by all agents. This means that at least one of them must be an agreement.

$$\begin{aligned} \therefore \left|\widehat{\Omega}\right| > 0 \wedge \omega_i^T = \omega_i^{T-1} \forall i \in \mathcal{N} &\implies \omega_* \neq \phi \\ \therefore \left|\widehat{\Omega}\right| > 0 \wedge \omega_* = \phi &\implies \omega_* \neq \phi \end{aligned} \tag{C56}$$

which is a contradiction for Case 2. From Eq. C56 and Eq. C52, we get:

$$\therefore \left|\widehat{\Omega}\right| > 0 \implies \omega_* \in \widehat{\Omega} \tag{C57}$$

$\therefore$ TAU is Exactly Complete on Λ^B_{DN} for any strategy assignment rule that employ only WAR.

□

This is supported empirically in the paper by the fact that in all experiments, the agreement rate (completeness) of TAU with WAR when negotiating against itself was always 100%.

Theorem 5 (Completeness)**.** *TAU(Υ_{WAR}) and TAU(Υ_{CAB}) are Exactly Complete on Λ^B_{DN}.*

Proof. The proof follows directly from combining Proposition 4 and 5. □

Lemma 2 ($\Upsilon_{CAB,WAR}$ Not Complete)**.** *TAU($\Upsilon_{CAB,WAR}$) (i.e. WAR negotiating against CAB) is not exactly complete on Λ_{DN}.*

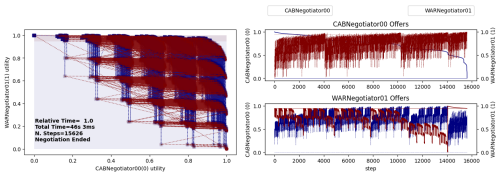

Fig. C3: CAB against WAR cannot reach agreement in the Car domain with 10% rational outcomes even though win-win deals exist.

What Theorem 5 and Theorem 5 established is that Υ_{WAR} and Υ_{CAB} are exactly complete on Λ_{DN} which means that WAR is exactly complete when negotiating against itself and CAB is exactly complete when negotiating against itself. What happens if CAB is negotiating against WAR. In some cases, the negotiation will end in disagreement even when win-win deals exist. The reason is that WAR will start with offering irrational outcomes that will tend to have high utility for its partner running CAB. The partner will accept these offers and as a result cannot accept anything worse for itself later (because CAB uses AcceptBetter). Moreover, CAB will never accept these irrational outcomes. This is why the negotiation may end in disagreement. An example of this happening from the experiment reported in the paper is on the domain Car with 10% of the outcomes being rational for the second agent and all of them being rational for the first agent in which running CAB against WAR leads to no agreement even though win-win deals exist. Fig. C3 shows this negotiation.

Conjecture 1 (Υ_{CAB} Not Complete on Multilateral)**.** *CAB is not exactly complete on Λ_{DN} (i.e. only on Λ^{B}_{DN}).*

The proof of Theorem 5 depended crucially on the fact that CAB will not get stuck unable to accept any offer after it accepted an offer from its partner that is then rejected by that partner. In a multilateral negotiation, it may be the case that the partner will accept that offer later but another partner will not which means that an agreement may not be reachable. This is not a proof (and that is why this is a conjecture). We did not find any counter example in the 184 scenarios evaluated in this paper.

C.8 TAU(Υ_{CAB}) is Exactly Fair on Λ^{RB}_{DN}

Theorem 6 (CAB Fairness)**.** *TAU(Υ_{CAB}) is Exactly Fair for bilateral negotiations on Λ_{DN} with no irrational outcomes. Specifically, it finds an Ordinal Kalai Bargaining Solution in that case:*

$$\left|\overline{\Omega}^{\lambda}\right| > 0 \rightarrow \lambda_{TAU}(\Upsilon_{CAB}) \in \Omega_k \quad \forall \lambda \in \Lambda_{DN} \tag{C58}$$

where $\Omega_k = \left\{ \omega : \omega \in \underset{\omega \in \mathbf{P}^{\lambda} \bigcup \{\phi\}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} r_i(\omega) \right\}$

Proof. From Theorem 5 (Completeness), we know that whenever there is disagreement, there are no Pareto-efficient outcomes and Eq. C58 holds trivially. We need to focus on the case when there is a proper agreement:

$$\omega_* \in \Omega$$

From Theorem 3 (Optimality), we know that — in this case — the agreement is Pereto-efficient.

$$\omega_* \in \mathbf{P}$$

The offering policy of CAB is to order rational outcomes in descending order of value and offer them in sequence repeating the last one. This assigns each outcome a unique integer giving the step it is offered (i.e. $R_i(\Omega_i^{\top}) \equiv R_i(\omega^T) = 0 \forall \omega^T \in \Omega_i^{\top}, R_i(\phi) = \left|\overline{\Omega}_i\right|$). We can model the fact that CAB never offer irrational outcomes by assigning all irrational outcomes an infinite value inf (From Theorem 1, none of these outcomes will ever be reached). Let's call this assigned integer the offering order $R_i(\omega)$.

The definition of the value function r can be related to this offering order by:

$$r_i(\omega) \approx \frac{R_i(\phi) - R_i(\omega)}{R_i(\phi) - R_i(\Omega_i^{\top})} \quad \forall \omega \in \overline{\Omega}_i$$

$$r_i(\omega) \approx \frac{|\Omega| - R_i(\omega)}{|\Omega| - 0} \quad \forall \omega \in \overline{\Omega}_i$$
$$r_i(\omega) \approx \frac{|\Omega| - R_i(\omega)}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i$$

The reason this is not a full equality is that outcomes with the same value for the agent have different consecutive offering orders but their value is the same. This will not affect the logic of this proof though. What really matters for our purposes is the following fact:

$$f\left(r_i(\omega)\right) = \frac{|\Omega| - R_i(\omega)}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i \tag{C59}$$

where $f(x)$ is a monotonically non-decreasing mapping. It assigns the same rank to consecutive outcomes with the same utility with the property

$$f(x_1) > f(x_2) \implies x_1 \geq x_2$$

For an agreement to be reached at step T in TAU, the same outcome must be offered by all agents and the agreement is reached at the step in which the last agent offers it:

$$\omega_* \in \Omega_i^T \qquad \forall i \in \mathcal{N}$$
$$\therefore T = \max_{i \in \mathcal{N}} R_i(\omega_*)$$
$$\therefore \omega_* \in \left\{ \omega : \omega \in \operatorname*{argmin}_{\omega \in \mathbf{P}} \max_{i \in \mathcal{N}} R_i(\omega) \right\}$$

Note that in the last step we had to take into account the fact that multiple outcomes may be the same for all agents.

$$\therefore \omega_* \in \left\{ \omega : \omega \in \operatorname*{argmin}_{\omega \in \mathbf{P}} \max_{i \in \mathcal{N}} c - c \times f\left(r_i(\omega)\right) \right\}$$

where $c = |\Omega|$. But c does not depend on ω or the agent index:

$$\therefore \omega_* \in \left\{ \omega : \omega \in \operatorname*{argmin}_{\omega \in \mathbf{P}} \max_{i \in \mathcal{N}} -f(r_i(\omega)) \right\}$$

$$\therefore \omega_* \in \left\{ \omega : \omega \in \operatorname*{argmin}_{\omega \in \mathbf{P}} \max_{i \in \mathcal{N}} -r_i(\omega) \right\}$$

because f is monotonically non-decreasing

$$\therefore \omega_* \in \left\{ \omega : \omega \in \operatorname*{argmin}_{\omega \in \mathbf{P}} - \min_{i \in \mathcal{N}} r_i(\omega) \right\}$$

$$\therefore \omega_* \in \left\{ \omega : \omega \in \operatorname*{argmax}_{\omega \in \mathbf{P}} \min_{i \in \mathcal{N}} r_i(\omega) \right\} = \Omega_k$$

□

The proof did not explicitly use the assumption that all outcomes are rational but this is implied in using the Kalai Bargaining Solution which is defined for this case [17]. The standard Kalai Bargaining solution for cardinal preferences is defined by:

$$\Omega_k^* = \left\{ \omega : \omega \in \underset{\omega \in \mathbf{P}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} u_i(\omega) - u_i(\phi) \right\}$$

which is not the same as the Ordinal Kalai Bargaining Solution we define here as:

$$\Omega_k = \left\{ \omega : \omega \in \underset{\omega \in \mathbf{P}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} r_i(\omega) \right\}$$

except for the case when all outcomes are rational.

C.9 TAU(Υ_{WAR}) is Exactly Fair on Λ_U^B

Theorem 7 (WAR Fairness)**.** *TAU(Υ_{CAB}) is Exactly Fair for bilateral negotiations on Λ_U. Specifically, it finds an Ordinal Kalai Bargaining Solution in that case:*

$$\left|\overline{\Omega}^{\lambda}\right| > 0 \rightarrow \lambda_{TAU}(\Upsilon_{CAB}) \in \Omega_k \quad \forall \lambda \in \Lambda_{DN} \tag{C60}$$

where $\Omega_k = \left\{ \omega : \omega \in \underset{\omega \in \mathbf{P}^{\lambda} \bigcup \{\phi\}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} r_i(\omega) - r_i(\phi) \right\}$

Proof. From Theorem 5 (Completeness), we know that whenever there is disagreement, there are no Pareto-efficient outcomes and Eq. C60 holds trivially. We need to focus on the case when there is a proper agreement:

$$\omega_* \in \Omega$$

From Theorem 4 (Optimality), we know that — in this case — the agreement is Pereto-efficient.

$$\omega_* \in \mathbf{P}$$

The offering policy of WAR is to first offer all irrational outcomes then to order rational outcomes in descending order of value and offer them in sequence repeating the last one. This assigns each outcome a unique integer giving the step it is offered (i.e.). We modeled the fact that WAR offers all irrational outcomes first by assigning adding the number of irrational outcomes to the step at which all rational outcomes are offered (From Theorem 1, none of these outcomes will ever be reached). Let's call this assigned integer the offering order $R_i(\omega)$.

This leads to the following two equations:

$$R_i(\Omega_i^{\top}) = R_i(\omega^T) = |\underline{\Omega}_i| \, \forall \omega^T \in \Omega_i^{\top}$$

$$R_i(\phi) = |\Omega|$$

The definition of the value function r can be related to this offering order by:

$$r_i(\omega) \approx \frac{R_i(\phi) - R_i(\omega) + n_{ir}}{R_i(\phi) - R_i(\Omega_i^\top) + n_{ir}} \quad \forall \omega \in \overline{\Omega}_i$$

where $n_{ir} = |\underline{\Omega}_i|$

$$\therefore r_i(\omega) - r_i(\phi) \approx \frac{|\Omega| - R_i(\omega) - |\underline{\Omega}_i|}{|\Omega| - |\underline{\Omega}_i| + |\underline{\Omega}_i|} - r_i(\phi) \quad \forall \omega \in \overline{\Omega}_i$$

$$\therefore r_i(\omega) - r_i(\phi) \approx \frac{|\overline{\Omega}_i| - R_i(\omega)}{|\Omega|} - \frac{|\overline{\Omega}_i| - |\Omega|}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i$$

$$\therefore r_i(\omega) - r_i(\phi) \approx \frac{|\overline{\Omega}_i| - R_i(\omega)}{|\Omega|} + \frac{|\underline{\Omega}_i|}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i$$

$$\therefore r_i(\omega) - r_i(\phi) \approx \frac{|\Omega| - R_i(\omega)}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i$$

$$\therefore r_i(\omega) - r_i(\phi) \approx 1 - \frac{R_i(\omega)}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i$$

$$\therefore f\left(r_i(\omega) - r_i(\phi)\right) = 1 - \frac{R_i(\omega)}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i$$

$$\therefore f\left(r_i^*(\omega)\right) = 1 - \frac{R_i(\omega)}{|\Omega|} \quad \forall \omega \in \overline{\Omega}_i \tag{C61}$$

where $r_i^*(\omega) \equiv r_i^*(\omega) - r_i(\phi)$. $f(x)$ is a monotonically non-decreasing mapping. It assigns the same rank to consecutive outcomes with the same utility with the property

$$f(x_1) > f(x_2) \implies x_1 \geq x_2$$

Eq. C61 has the same form as Eq. C59. Applying the same argument exactly using r^* instead of r, we get at:

$$\therefore \omega_* \in \left\{\omega : \omega \in \underset{\omega \in \mathbf{P}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} r_i^*(\omega)\right\}$$

$$\therefore \omega_* \in \left\{\omega : \omega \in \underset{\omega \in \mathbf{P}}{\operatorname{argmax}} \min_{i \in \mathcal{N}} \left(r_i(\omega) - r_i(\phi)\right)\right\} = \Omega_k$$

□

C.10 WAR weakly dominates CAB on Λ_{DN}^B

Theorem 8 (Weak Domination). *WAR weakly dominates CAB for TAU on* Λ_{DN}^B.

$$\begin{aligned} &\mathcal{P}_\Upsilon(i \approx WAR, ...) = \phi \\ &\vee \hat{u}_i(\mathcal{P}_\Upsilon(i \approx WAR, ...)) \geq \hat{u}_i(\mathcal{P}_\Upsilon(i \approx CAB, ...)) \\ &\forall i \in \mathcal{N}^\lambda \forall \lambda \in \Lambda_{DN}^B \end{aligned} \tag{C62}$$

Proof. Let's assume — without loss of generality — that agent j uses CAB and run two negotiations N^{λ}_{WAR} and N^{λ}_{CAB} where agent i is using WAR and CAB respectively. From now one, we mark a value v that belongs to N^{λ}_{σ} as $z\,|_{\sigma}$ (i.e. v given that agent i uses strategy σ). For example $\omega_*\,|_{CAB}$ is the agreement when agent i plays CAB against j. [7]

Case 1: No Irrational Outcomes for i

Consider first any scenario $\lambda \in \Lambda_{DN}$ with no irrational outcomes for i (i.e. $\overline{\Omega}_i = \Omega$). In this case CAB and WAR offer exactly in the same way. The only difference is that CAB accepts all rational outcomes while WAR accepts only outcomes better than the best it accepted so far.

Because WAR and CAB have exactly the same offering strategy in this case, the offers exchanged until the step $\min(T\,|_{WAR}, T\,|_{CAB})$ will be exactly the same. We now consider the three alternatives for the negotiation length:

1. $T\,|_{WAR} < T\,|_{CAB}$: The agreement cannot be worse for i because it offers its outcomes in descending order of their value (i.e. $\omega_*\,|_{WAR} \geq \omega^{\lambda}_*\,|_{CAB}$)
2. $T\,|_{WAR} = T\,|_{CAB}$: The agreement must be the same (i.e. $\omega_*\,|_{WAR} = \omega^{\lambda}_*\,|_{CAB}$)
3. $T\,|_{WAR} > T\,|_{CAB}$: The agreement worse for i in principle. This cannot happen though. For the negotiation to take longer in this case, ω_* must have been offered by both agents and must have been accepted by j (because its acceptance strategy does not change). This means that to reach $T\,|_{WAR}$, ω_* must have been rejected at step $T\,|_{optim}$ by WAR. Nevertheless, ω_* must be acceptable for CAB (otherwise it would have not been an outcome). But the acceptance strategy of WAR accepts any outcome acceptable by CAB which is a contradiction.

We have shown that in the first two subcases $\omega_*\,|_{WAR} \succcurlyeq_i \omega_*\,|_{CAB}$ and the third subcase cannot happen.

$$\text{Case 1} \implies \omega_*\,|_{WAR} \succcurlyeq_i \omega_*\,|_{CAB}$$

Case 2: Some Irrational Outcomes for i We will consider the same three subcases based on the relative length of the two negotiations assuming that there are n_i irrational outcomes for agent i:

1. $T\,|_{WAR} < T\,|_{CAB}$: The agreement cannot be worse for i because it offers its outcomes in descending order of their value (i.e. $\omega_*\,|_{WAR} \geq \omega^{\lambda}_*\,|_{CAB}$)
2. $T\,|_{WAR} = T\,|_{CAB}$: The agreement cannot be worse for i because it will be offering an outcome n_i steps higher in its ordered outcome set and j will be offering the same outcome it always offers at this step. (i.e. $\omega_*\,|_{WAR} \geq \omega^{\lambda}_*\,|_{CAB}$).
3. $T\,|_{WAR} > T\,|_{CAB}$: The agreement worse for i in principle. We will consider this case in more details.

From the three subcases above, we conclude that the only viable possibility to have a worse outcome for i in N_{WAR} is to have a longer negotiation:

$$\omega_*\,|_{WAR} \prec_i \omega_*\,|_{CAB} \implies T\,|_{WAR} > T\,|_{CAB}$$

We again have three subcases to consider based on the difference between these two times:

[7] j always plays CAB in this proof.

1. $T\ |_{WAR}\ -T\ |_{CAB} < n_i$: i will be offering an outcome $n_i - (T\ |_{WAR}\ -T\ |_{CAB})$ steps higher in its ordered outcome list and j will be offering an outcome worse for itself than $\omega_*\ |_{CAB}$. There are two possibilities depending on who offered the outcome that was accepted as the negotiation outcome:
   (a) Agent j: Because $\omega_*\ |_{WAR} \preceq_j\ \omega_*^\lambda\ |_{CAB}$, $\omega_*^\lambda\ |_{CAB}$ must have been offered earlier by j in N_{WAR}. Because, $\omega_*\ |_{CAB}$ is rational for i, it must have been accepted at that point by i. Assuming $\omega_*\ |_{WAR} \prec_i\ \omega_*\ |_{CAB}$, then $\omega_*\ |_{CAB}$ must also have been offered by i in N_{WAR}. Moreover, given that i offered n_i irrational outcomes first, the step at which it offered $\omega_*\ |_{CAB}$ in N_{WAR} must be larger than the step it offered that outcome in N_{CAB} which means that must have been accepted by j. We just showed that $\omega_*\ |_{CAB}$ must have been offered and accepted by both agents in N_{WAR}. This implies that $\omega_*\ |_{WAR} = \omega_*\ |_{CAB}$. Therefore, $\omega_*\ |_{WAR} \approx_i\ \omega_*\ |_{CAB}$.
   (b) Agent i: Because $\omega_*\ |_{WAR}$ is higher in the ordered outcome list than $\omega_*\ |_{CAB}$, it cannot be worse. Therefore, $\omega_*\ |_{WAR} \succcurlyeq_i\ \omega_*\ |_{CAB}$.

   Taken together this shows that $T\ |_{WAR}\ -T\ |_{CAB} < n_i \implies \omega_*\ |_{WAR} \approx_i\ \omega_*\ |_{CAB}$
2. $T\ |_{WAR}\ -T\ |_{CAB} = n_i$: i will be offering an outcome $n_i - (T\ |_{WAR}\ -T\ |_{CAB}) = 0$ steps higher in its ordered outcome list which means it is offering the same outcome (i.e. $\omega_*\ |_{WAR} = \omega_*^\lambda\ |_{CAB}$). Therefore, $\omega_*\ |_{WAR} \approx_i\ \omega_*\ |_{CAB}$.
3. $T\ |_{WAR}\ -T\ |_{CAB} > n_i$: i will be offering an outcome $(T\ |_{WAR}\ -T\ |_{CAB}) - n_i$ steps lower in its ordered outcome list and j will be offering an outcome worse for itself than $\omega_*\ |_{CAB}$. There are two possibilities depending on who offered the outcome that was accepted as the negotiation outcome:
   (a) Agent j: Following the same argument for this condition above we reach the same conclusion. Therefore, $\omega_*\ |_{WAR} \succcurlyeq_i\ \omega_*\ |_{CAB}$.[8]
   (b) Agent i: Because $\omega_*\ |_{WAR} \prec_i\ \omega_*\ |_{CAB}$, $\omega_*^\lambda\ |_{CAB}$ must have been offered earlier. The same reasoning as in the previous case applies from there. $\omega_*\ |_{WAR} \succcurlyeq_i\ \omega_*\ |_{CAB}$. But that is a contradiction which means that this case cannot happen.

We have shown that in all possible situations, $\omega_*\ |_{WAR} \succcurlyeq_i\ \omega_*\ |_{CAB}$. □

Why would anyone care about this notion of weak domination? In some practical situations, it is possible to repeat a failed negotiation. In such cases, knowing that s_1 weakly dominates s_2 — in the sense given above — suggests that we should start by using s_1 and if an agreement is not reached, we can switch to s_2. In our specific case, this suggests attempting WAR first then if and only if we could not reach agreement we can switch to CAB.

C.11 WAR is a PBE of TAU on Λ^B_{DN}

We proceed in the following steps: In Section C.11.1 we rigorously define the induced incomplete information game corresponding to TAU on a set of scenarios and define its information sets. Section C.11.2 defines the beliefs of agents in the game. Section C.11.3 describes *consistent* Bayesian belief update given that the partner is using WAR and

[8] The only difference is that we do not need to use the assumption that $\omega_*\ |_{WAR} \prec_i\ \omega_*\ |_{CAB}$ because it is guaranteed in this case.

shows that assuming no information about the partner makes this process much simpler than if partial information about partner preferences (e.g. a linear aggregation utility function) is available. Finally, Section C.11.4 completes the proof by showing that WAR is a Perfect Bayesian Equilibrium for the induced game making it — by definition — a negotiation equilibrium.

C.11.1 The Induced Game

Let n_s be the number of scenarios is a scenario set Λ, n_o^λ be the cardinality of the outcome space for $\lambda \in \Lambda$ and n_A^λ be the number of agents in that scenario. Moreover, let's define Λ^Ω as the subset of Λ that share the same outcome space Ω. n_o^Ω is defined correspondingly as the number of outcomes in Ω which is shared by all members of Λ^Ω.

In this section, we only consider the bilateral negotiations case (i.e. $n_A = 2$). For simplicity of exposition we assume that all scenarios have the same outcome size n_o and will generalize that back in the following proposition.

Given a set of scenarios Λ, Nature moves first and once by selecting a scenario from this set of scenarios based on some probability distribution. Nature never moves again. In this paper, we are focused on the case of a general negotiator (i.e. one that can negotiate in any scenario) with no information about the partner which means that all scenarios are sampled with equal probability.

Information sets correspond to decision points of agents. There is one information set for each offering and each acceptance decision by each agent. Because TAU always terminates in $n_o + 1$ rounds (Theorem 1), the longest path in the game tree is of length $4n_o + 4 + 1$ (1 for nature's move and a factor 4 for the two offers and two acceptance decisions per round). Each offering node has a branching factor of $n_o + 1$ and each acceptance node has a branching factor of 2 while the root node (Nature's) has a branching factor of n_s. If we consider the two levels corresponding to an offer and its response in a round together, we get a branching factor of $2n_o + 2$ with a tree of height $1 + 2(n_o + 1)$ leading to a total of $n_s \left(2n_o + 2\right)^{2(n_o+1)}$ nodes.

The only incomplete information in the game is which scenario are we playing which means that the nodes of each decision level in the game (i.e. each decision made by each agent) are distributed into information sets of size n_s each[9]. Each information set corresponds to one possible negotiation *trace* defined as the complete history of the decisions made in the negotiation (common knowledge).

We can summarize this information in the following proposition

Proposition 6 (Game). *TAU induces a game* G^Λ *for any set of bilateral scenarios* Λ. *The game* G^Λ *is an incomplete information game with at most* $\max_{\lambda\in\Lambda} 5 + 4n_o^\lambda$ *containing a total of* $\sum_{\lambda\in\Lambda} \left(2n_o^\lambda + 2\right)^{2\left(n_o^\lambda+1\right)}$. *These nodes decompose into information sets corresponding to specific negotiation traces. Each information set has exactly* n_s *nodes (all at in the same level).*

[9] We choose the common practice of having the largest possible information sets and leaving it to agent beliefs to assign zero probabilities to irrelevant nodes. For example, in the first node, we have $\sum_{\lambda in \Lambda} K_\lambda$ nodes, where K_λ are all possible preference assignments for scenario λ. Nevertheless, the agent about to move knows its own preferences and the outcome space which assigns zero probability to most of these nodes in its belief.

C.11.2 Modeling Beliefs in the Induced Game

The belief of an agent in the induced game at an information set is a probability distribution over all possible nodes in that set. This is a distribution over n_s values which is in general infinite in size. In this paper, we are only interested in analyzing games played using TAU against WAR. This will simplify this distribution a lot. WAR does not use any cardinal information in its decision (i.e. it is fully characterized by full ordering it uses for offering and the rational outcome set). Moreover, the negotiation outcome (payoffs) will be determined completely by the order in which offers were proopsed and the binary acceptance decisions. This means that all preference allocations for any given scenario that lead to the same ordering of outcomes for all agents will correspond to the same payoffs and may be collapsed together. Moreover, equality does not matter when WAR is playing itself as there is a full ordering imposed by the fact that if two outcomes are the same for an agent, they must be offered according to the predefined (yet arbitrary) common ordering defined by the mechanism.

Because the outcome space Ω^λ is common knowledge for all agents, the induced game decomposes into a set of independent games each sharing the same outcome space Ω (i.e. each game corresponds to on Λ^Ω set of scenarios). Beliefs over these games are disjoint and will be zero for all of these decomposed games except the one currently being played. From now on, we will drop the superscript Ω from Λ^Ω and n_o^Ω because all other members of the belief are zero all the time.

In summary, for the purposes of this paper, we can model an agent's belief as a distribution over all possible permutations of the extended outcome space of our partner at every negotiation. There is always $(n_o + 1)!$ such permutations which is always finite for Λ_{DN} that we are concerned with in this proof. We use Λ to represent the set of all possible permutations of *osplus* and λ to be one such permutation. Moreover, we define λ_ω as all permutations ending with ω, λ^ω as all permutations starting with ω.

Proposition 7 (Belief)**.** *The induced game* G^Λ *induced by the bilateral discrete scenarios with no information* Λ_{DN} *when agent* j *is playing WAR has information sets for which the number of nodes with nonzero probabilities is* $(n_o + 1)!$ *corresponding to all possible permutations of* Ω^+ *where* n_o *is the number of outcomes of the outcome space (known at the first round).*

$$\lambda \in \Lambda_{DN} \wedge j \approx WAR \implies \mathbb{B}_i \equiv \{p(\lambda) \forall \lambda \in \Lambda\}$$

C.11.3 Bayesian Belief Update

We are only considering belief update compatible with knowing that our partner is using WAR because that is the only update relevant for *checking* that WAR is a pure PBE for the induced game.

We also need the following Lemma about its three stages:

Lemma 3 (WAR's Repetition)**.** *WAR never ends a negotiation by sending ϕ and never repeats until round $n_o + 1$ where n_o is the size of the outcome space.*

$$j \approx WAR \wedge \omega_j^k = \omega_j^{k-1} \implies k = n_o$$

Proof. The first phase of WAR involves sending all irrational outcomes $\underline{\Omega}_i$ and the second phase all rational outcomes $\overline{\Omega}_i$. But these are all possible outcomes: $\Omega = \overline{\Omega}_i \cup \underline{\Omega}_i$.

$\therefore$ The second phase (Rational) ends only when all outcomes in the outcome space are offered and there are n_o of them.

$\therefore$ The third phase (Repeat) can only start at round $n_o + 1$. □

Consider a negotiation between i and j on a scenario from Λ_{DN} in which j uses WAR. We augment the belief of i with another data structure called its Partner's Offer Rationality Distribution $\mathbb{K}_i$ which is defined as a distribution over the rationality of every past or potential offer from j:[10]

$$\mathbb{K}_i^k \equiv p(\omega_j^k \succcurlyeq_j \phi) \qquad \forall k \in [1, n_o]$$

Three special properties of this distribution will have an important role for the proposed belief update rule:

Ignorance A partner rationality belief is in the Ignorance state iff $\mathbb{K}_i^k \notin \{0, 1\} \quad \forall k \in [1, n_o]$

RationalAfter(k) A partner rationality belief is in the Ignorance state iff $\mathbb{K}_i^k = 1 \quad \forall k \leq j \leq n_o$

IrrationalBefore(k) A partner rationality belief is in the Ignorance state iff $\mathbb{K}_i^k = 0 \quad \forall 1 \leq j \leq k$

The final component of the belief state for i is the Learned Inequality Set $\mathcal{O}_i$ consisting of a set of inequalities that it learns from the negotiation trace.

The full belief state of i at step k consists of the tuple $\langle \mathbb{B}_i, \mathbb{K}_i \mathcal{O}_i \rangle$. Hereafter, we use p_k to indicate a probability as in i's belief state at step k.

It is important to note at this point that WAR's behavior does not allow i to ever have a probability that an equality holds other than zero or one. That is because the only information available to it is the order at which offers are given which can in some cases be used to infer inequality and acceptance decisions which can be used to infer inequality between an outcome and disagreement. Notice also that elements of $\mathcal{O}_i$ are always strict inequalities because they are intended to model the order in which offers are given by j rather than the underlying preferences that may have some equalities. Moreover, it is not possible to infer equality of any two outcomes from the behavior of WAR because $\omega_1 \approx_j \omega_2$ can only be inferred from knowing that $\omega_1 \succcurlyeq_j \omega_2$ and $\omega_1 \preccurlyeq_j \omega_2$, but TAU does not allow repetition which make it impossible to learn both facts.

[10] We ignore round $n_o + 1$ because it does not matter as the negotiation will end with disagreement anyway once this round is reached.

Initial Belief

The belief of i is initialized by assigning a probability value for every permutation of the extended outcome space as described above. We are only concerned with the case with no information ($\mathrm{I}^{\lambda} i = \mathcal{F}^{\lambda} i$). In this case, all permutations are equally likely and they all get a probability of $\frac{1}{(n_o+1)!}$. Moreover:

$$p_0(R_j(\omega_1) > R_j(\omega_2)) = p_0(R_j(\omega_2) > R_j(\omega_1)) = \frac{1}{2}$$
$$\forall \omega_1, \omega_2 \in \Omega^+$$

where $R_j(\omega)$ is the round at which j will offer ω (See the proof for Theorem 8).

Therefore, i must assign a probability $\frac{n_o!}{(n_o+1)!} = \frac{1}{n_o}$ to the proposition that agent j's offers will be rational because this only happens if all outcomes are rational which means that ϕ comes last in the permutation of Ω^+. This implies that i starts its Partner's Offer Rationality Distribution in the Ignorance state. Moreover, the $|\mathcal{O}| = 0$ as all permutations are equally likely implying by symmetry that there are no known inequalities.

$$\begin{aligned} \mathbb{K}_i^0 &= \left\{ \frac{1}{n_o} \forall \omega \in \Omega \right\} \\ \mathcal{O}_i^0 &= \{\} \\ \mathbb{B}_i^0 &= \left\{ \frac{1}{n_o!} \forall \lambda \in \Lambda \right\} \end{aligned} \tag{C63}$$

Update Process

When j acts, i uses its current belief as the prior and then updates the probability distribution over permutations in the following step.

Firstly, it determines its the new partner rationality belief: (1) If j offers an outcome it already rejected earlier, then we know that it is in the Irrational phase. (2) If j offers an outcome that it accepted earlier then we know that it is in the Rational phase. (3) If j accepts an offer that it offered earlier, then we know that it is in the Rational phase at least since it made that offer. (4) If j rejects an offer that it offered earlier, we know that it was in the Irrational phase up to the point it made that offer. Using these rules, i determines its current state. These are the four main if-statements in Algorithm 2

Secondly, if the state switches to Rational, i examines the negotiation history and assigns all outcomes that were offered at the point the switch is determined to have happened to the set of rational outcomes. If the state switches to Irrational, all offers from i so far are assigned to the set of irrational outcomes. This is the first operation in every if-statement in Algorithm 2.

Thirdly, the Learned Inequality Set is updated given the information we have so far. (1) if an offer is rejected by j, it must be irrational, and (2) if it is accepted, it must be rational. These correspond to the first two if-statements in Algorithm 3. Finally, (3) if i knows a step j at which j is in its rational phase (i.e. Rational(j)), then it can add an equality for any two subsequent offers of j stating that the earlier is better than the later. This is the third if-statement in Algorithm 3.

Algorithm 2 Partner's Offer Rationality Distribution Update: UpdateOR

Input: $\omega_i^k, \omega_j^k, a_{ij}^k, \mathcal{T}^k$
Output: $\mathbb{K}_i$

if $\exists j < k : \omega_i^j = \omega_j^k \wedge \neg a_{ji}^j$ **then** ▷ Offers rejected
  Set IrrationalBefore(k)
  $p_k(\omega_j^k \succcurlyeq_j \phi) \leftarrow 0 \qquad \forall 1 \leq k \leq k$
end if
if $\exists j < k : \omega_i^j = \omega_j^k \wedge a_{ji}^j$ **then** ▷ Offers accepted
  Set RationalAfter(k)
  $p_k(\omega_j^k \succcurlyeq_j \phi) \leftarrow 1 \qquad \forall n_o \geq k \geq k$
end if
if $a_{kj} \wedge j < k : \omega_j^j = \omega_i^k$ **then** ▷ Accepts offered
  Set RationalAfter(j)
  $p_k(\omega_j^k \succcurlyeq_j \phi) \leftarrow 1 \qquad \forall n_o \geq k \geq j$
end if
if $\neg a_{kj} \wedge j < k : \omega_j^j = \omega_i^k$ **then** ▷ Rejects offered
  Set IrrationalBefore(j)
  $p_k(\omega_j^k \succcurlyeq_j \phi) \leftarrow 0 \qquad \forall 1 \leq k \leq j$
end if
Renormalize $\mathbb{K}_i^k$: $(p_i(a) \leftarrow \frac{p_i(a)}{\sum_a p_i(a)})$
return $\mathbb{K}_i^k, \mathbb{B}_i^k$

Algorithm 3 Learned Inequality Set Update: UpdateInequalities

Input: $\mathbb{K}_i^k, \mathcal{O}_i^{k-1}, \omega_i^k, \omega_j^k, a_{ij}^k, \mathcal{T}^k$
Output: $\mathcal{O}_i$

$\mathcal{O}_i^k \leftarrow \mathcal{O}_i^{k-1}$
if a_{kj} **then** ▷ Accepts
  $\mathcal{O}_i^k \leftarrow \mathcal{O}_i^k \bigcup \{\langle \omega_i^k, \phi \rangle\}$
end if
if $\neg a_{kj}$ **then** ▷ Rejects
  $\mathcal{O}_i^k \leftarrow \mathcal{O}_i^k \bigcup \{\langle \phi, \omega_i^k \rangle\}$
end if
if $\exists j \in [1, k] : \text{Rational}(j)$ **then**
  $\mathcal{O}_i^k \leftarrow \mathcal{O}_i^k \bigcup \{\langle \omega_i^k, \phi \rangle\}$
  $j \leftarrow \min_{1 \leq k \leq n_o} \text{Rational}(k)$ ▷ Earliest rational offer
  for $j \leq k \leq k$ **do**
    $\mathcal{O}_i^k \leftarrow \mathcal{O}_i^k \bigcup \{\langle \omega_j^k, \omega_j^l \rangle \quad \forall k < l \leq k\}$ ▷ Subsequent offers are worse
  end for
end if
return $\mathcal{O}_i^k$

Algorithm 4 Information Set Belief Update: UpdateISBelief

Input: **Input:** $\mathbb{K}_i^k, \mathcal{O}_i^k, \mathbb{B}_i^{k-1}$
Output: $\mathbb{B}_i^k$

for $\omega_1, \omega_2 \in \mathcal{O}_i^k$ **do**
  for $\lambda \in \Lambda$ **do** ▷ Remove incompatible permutations
    if $R_j(\omega_1) \mid_\lambda > R_j(\omega_2) \mid_\lambda$ **then**
      $p_i(\lambda) \leftarrow 0$
    end if
  end for
end for
Renormalize $\mathbb{B}_i^k$: $(p_i(a) \leftarrow \frac{p_i(a)}{\sum_a p_i(a)})$
return $\mathbb{B}_i^k$

Algorithm 5 Belief Update Process: UpdateBelief

Input: $\mathbb{K}_i^{k-1}, \mathbb{B}_i^{k-1}, \mathcal{O}_i^{k-1}, \omega_i^k, \omega_j^k, a_{ij}^k, \mathcal{T}^k$
Output: $\mathbb{K}_i^k, \mathbb{B}_i^k, \mathcal{O}_i^k$

$\mathbb{K}_i^k, \mathbb{B}_i^k \leftarrow$ UpdateOS$(\omega_i^k, \omega_j^k, a_{ij}^k, \mathcal{T}^k)$
$\mathcal{O}_i^k \leftarrow$ UpdateInequalities$(\mathbb{K}_i^k, \mathcal{O}_i^{k-1}, \omega_i^k, \omega_j^k, a_{ij}^k, \mathcal{T}^k)$
$\mathbb{B}_i^k \leftarrow$ UpdateISBelief$(\mathbb{K}_i^k, \mathcal{O}_i^k, \mathbb{B}_i^k)$
return $\mathbb{K}_i^k, \mathbb{B}_i^k, \mathcal{O}_i^k$

Finally, the distribution over permutations is updated according to Algorithm 4 by assigning zero probability to all permutations that are incompatible with the current set of rational and irrational outcomes.

The overall process is deccribed in Algorithm 5 The process and arguments described above uses all the information available to i from the behavior of j under the assumption that j is using WAR which leads to the following proposition

Proposition 8 (Consistency)**.** *The belief update process described by Algorithm 5 form a consistent belief with any strategy profile that has agent j playing WAR on the Λ_{DN} scenario set.*

C.11.4 WAR is a PBE for the Induced Game

First of all, we need the following lemma about WAR's decision making in information sets.

Lemma 4 (WAR cannot be affected)**.** *WAR's offers and responses (i.e. all its decisions) are independent of the negotiation trace given its preferences. In other words, there is no way to influence the behavior of WAR when negotiating with it.*

Proof. Offers of WAR are purely determined by the relative ordering of outcomes in its preferences and the common ordering defined for the negotiation. Both are independent of any events that happen during the negotiation.

Responses of WAR depend only on the rationality of the outcome which is — again — independent of any events that happen during the negotiation. □

Let agent j be using WAR in a bilateral negotiation on a scenario from Λ_{DN} with the TAU protocol:

Lemma 5 (Uniform Belief)**.** *At any step of the negotiation, agent i's belief is uniform over all nodes in the active information set that have nonzero probabilities.*

$$p_k(\lambda) \in \left\{0, \frac{1}{n_p^k}\right\} \qquad \forall \lambda \in \Lambda$$

where $n_p^k = |\{\lambda : p_k(\lambda) \neq 0\}|$

Proof. Agent i's belief is initialized according to Eq C63. In this state, the lemma follows trivially.

$$\therefore p_0(\lambda) \in \left\{0, \frac{1}{n_p^k}\right\} \qquad \forall \lambda \in \Lambda$$

After every update of the belief state using Algorithm 4, each permutation will either have been assigned a zero probability or be divided by the sum of all probabilities that are nonzero (The re normalization step).

$$\because p_{k-1}(\lambda) \in \left\{0, \frac{1}{n_p^{k-1}}\right\} \qquad \forall \lambda \in \Lambda$$

For permutations assigned zero probability, the Lemma follows trivially again. For those not assigned a zero probability, their probability will be updated only by the re normalization step.

$$\therefore p_k(\lambda) = p_{k-1}(\lambda) \times \frac{n_p^{k-1}}{n_p^k}$$

$$\therefore p_k(\lambda) \in \left\{0, \frac{1}{n_p^k}\right\} \qquad \forall \lambda \in \Lambda$$

□

Lemma 6 (Unknown Ranking of New Outcomes)**.** *Let $\omega \in \Omega$ be an outcome that is not yet offered by i:*

$$\omega \notin \Omega_i^k \implies p_i(\omega \succ_j \phi) = p_i(\phi \succ_j \omega). \tag{C64}$$

Let $\omega_1, \omega_2 \in \Omega$ be two members of the outcome space that are not yet offered by j:

$$\omega_1, \omega_2 \notin \Omega_i^k \implies p_i(\omega_1 \succ_j \omega_2) = p_i(\omega_1 \succ_j \omega_2). \tag{C65}$$

Proof. Agent i's belief is initialized according to Eq C63. In this state, the lemma follows trivially. From Lemma 5, all permutations that have nonzero probability are equally probable.

Let $\omega \in \Omega$ be an outcome that is not yet offered by i. The only cases in which an inequality involving ϕ is added are in lines 3, 6, 9 of Algorithm 3 corresponding in each case to an offer from i. Crucially, WAR never offers ϕ so line 12 in the algorithm will not add an inequality with ϕ on either side. Moreover, all cases in which the belief (p_i) changes in Algorithm 2 (lines 3, 7, 11, 15) involve in the condition for running them an offer from i. There are no other places in which p_i changes in Algorithm 5 (except for the re normalization step which has no effect for our purposes). This proves Eq. C64.

Let $\omega_1, \omega_2 \in \Omega$ be two members of the outcome space that are not yet offered by j. By the same argument in the previous paragraph, there are no places in which p_i changes in Algorithm 5 for these outcomes. This proves Eq. C65.

□

Theorem 9 (Equilibrium)**.** *Υ_{WAR} is a negotiation equilibrium of TAU for bilateral negotiations on Λ_{DN}.*

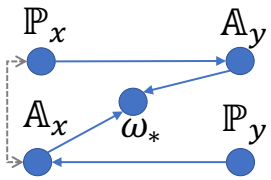

Fig. C4: The final agreement is fully determined by the sequence of acceptances of all partners that are in turn fully determining by the sequence of offers from each partner.

Proof. We assume that agent j uses WAR and consider its partner i. Our goal is to show that the best response of i is to use WAR as well. Given that we established the consistent belief update rule earlier, this concludes the proof.

Firstly we define the expected utility of i [11] using strategy π (represented as $i \approx \pi$):

$$\mathrm{EU}(\pi) = \sum_{\lambda \in \Lambda^{DN}} \sum_{\omega \in \Omega^+_\lambda} p\left(\omega \mid i \approx \pi\right) \hat{u}_i(\omega), \tag{C66}$$

where $p\left(\omega \mid i \approx \pi\right) \equiv p[\omega^\lambda_* = \omega : i \approx \pi]$ is the probability that the scenario λ will end with the outcome ω if agent i uses strategy π and $\Omega^+_\lambda = \Omega^\lambda \cup \{\phi\}$.

$$\therefore \mathcal{B}(WAR) = \operatorname*{argmax}_\pi \mathrm{EU}(\pi)$$

[11] If preferences are given as a partial ordering instead of a utility function. We can define the utility of an outcome $\hat{u}_i(\omega)$ as its rank in that ordering with $\hat{u}_i(\omega) = 0 \forall \omega \in \Omega^+_i$ and $\hat{u}_i(\omega) \leq n_o \forall \omega \in \Omega^\top_i$.

We know that WAR offers and accepts only rational outcomes. This means that no matter what π is:

$$p(\omega \mid i \approx \pi) = 0 \quad \forall \omega \notin \overline{\Omega}_y$$

$$\therefore \mathrm{EU}(\pi) = \sum_{\lambda \in \Lambda^{DN}} \sum_{\omega \in \overline{\Omega}_y^\lambda} p(\omega \mid i \approx \pi)\, \hat{u}_i(\omega), \tag{C67}$$

$$\therefore \mathcal{B}(WAR) = \operatorname*{argmax}_s \sum_{\lambda \in \Lambda^{DN}} \sum_{\omega \in \overline{\Omega}_y^\lambda} p(\omega \mid i \approx WAR)\, \hat{u}_i(\omega) \tag{C68}$$

$$\therefore \mathcal{B}(WAR) = \operatorname*{argmax}_s \sum_{\lambda \in \Lambda^{DN}} \mathrm{EU}^\lambda(s) \tag{C69}$$

where $\mathrm{EU}^\lambda(\pi) = \sum_{\omega \in \overline{\Omega}_y^\lambda} p(\omega \mid i \approx \pi)\, \hat{u}_i(\omega)$

Also we know from Theorem 1 that the negotiation for scenario λ will end at some finite time step $T_s \in \mathbb{N}^+$ for any strategy s. This means that if we can show that there is no strategy that can outperform WAR against itself for any scenario, we have shown that WAR is a best response.

We will now focus on a single negotiation scenario $\lambda \in \Lambda^{DN}$ in agent j uses WAR and agent i uses $\hat{s}$.

Case 1: Changing the acceptance policy ($\hat{s} = (\eta_{\hat{s}}, \rho_{\hat{s}} \neq \rho_{WAR})$):

Let's consider a strategy $\hat{s} = (\eta_{\hat{s}}, \rho_{\hat{s}} \neq \rho_{WAR})$. From Lemma 4, we can conclude that the only way for $\hat{s}$ to influence the final agreement ω_*^λ is by changing $\mathbb{A}_i$ (because it cannot influence Ω_j or $\mathbb{A}_j$ and Ω_i is independent of the acceptance policy given the offering policy which we will consider later).

Because $\rho_{\hat{s}} \neq \rho_{WAR}$, an outcome ω must exist that is either accepted by $\hat{s}$ and rejected by WAR or the other way around. Let ψ be the best outcome for i that it accepted so far:

$$\psi = \operatorname*{argmax}_{\omega \in \Omega_j^*(k)} \hat{u}_i(\omega)$$

The only way this decision can affect the expected utility is if either $\hat{\omega}$ or ψ ends up being the agreement:

$$\mathrm{EU}^\lambda(\hat{s}) \neq \mathrm{EU}^\lambda(WAR) \implies \omega_*^\lambda \in \{\hat{\omega}, \psi\}, \tag{C70}$$

1-1) $\hat{\omega}$ is rejected by $\hat{s}$ and accepted by WAR when offered by j at step k.

The only way this can affect expected utility is by preventing $\hat{\omega}$ from becoming an agreement (i.e. $p(\hat{\omega} \mid i \approx \hat{s}) = 0$ and $p(\hat{\omega} \mid i \approx WAR) = 1$). Most importantly, this decision does not affect $p(\omega \mid i \approx \hat{s})$ for any outcomes except $\hat{\omega}$ (Lemma 4).

From Eq. C68 and the fact that all utilities and probabilities are positive:

$$\therefore \text{Case 1-1} \implies \mathrm{EU}^\lambda(\hat{s}) \leq \mathrm{EU}^\lambda(WAR) \tag{C71}$$

1-2) $\hat{\omega}$ is accepted by $\hat{s}$ and rejected by WAR when offered by j at step k.

Because WAR only rejects irrational outcomes,

$$\therefore \hat{u}_i(\hat{\omega}) \leq \hat{u}_i(\phi) \tag{C72}$$

If the negotiation ends with ψ, there is no change in expected utility and if it ends by $\hat{\omega}$, the expected utility cannot increase because WAR against itself will never lead to an irrational outcome (Theorem 2).

$$\therefore \text{Case 1-2} \implies \text{EU}^\lambda(\hat{s}) \leq \text{EU}^\lambda(WAR) \tag{C73}$$

From Eq C71 and Eq C73 (and the fact that these are all the possible options for Case 1):

$$\therefore \hat{s} = (\eta_{\hat{s}}, \rho_{\hat{s}} \neq \rho_{WAR}) \implies \text{EU}^\lambda(\hat{s}) \leq \text{EU}^\lambda(WAR) \tag{C74}$$

$\therefore$ No matter what offering policy is used, it is not possible to find a better response to WAR that has a different acceptance policy. Now we need to show that fixing the acceptance policy as ρ_{WAR}, no offering policy can improve upon η_{WAR}.

Case 2: Changing the offering policy only: $\hat{s} = (\eta_{\hat{s}}, \rho_{\hat{s}} = \rho_{WAR})$

Let's consider a strategy $\hat{s} = (\eta_{\hat{s}} \neq \eta_{WAR}, \rho_{WAR})$. $\hat{s}$ must offer some outcome $\hat{\omega}$ at round k while WAR would have offered $\omega^* \mid_{WAR}$. The situation is more complicated than when we changed acceptance policy because the offering policy of i can affect the final agreement in two ways (Fig C4):

2-A) by affecting the acceptance policy changing $\mathbb{A}_j^{T_{\hat{s}}}$ for the partner.

2-B) by affecting the offering policy changing $\Omega_j^{T_{\hat{s}}}$ for the partner. This can only happen if the negotiation length $T_{\hat{s}}$ is greater than what it would have been if WAR was used instead of $\hat{s}$ (T_{WAR}) because Ω_j is independent of the behavior of i (Lemma 4).

Because WAR offers all rational outcomes in descending order of their value for itself, $\hat{\omega}$ cannot be better than any outcome offered by WAR up to step k

$$\therefore \hat{u}_i(\hat{\omega}) \leq \hat{u}_i(\omega) \quad \forall \omega \in \Omega_i^{:k}$$

Moreover, it cannot be better than $\omega^* \mid_{WAR}$

$$\therefore \hat{u}_i(\hat{\omega}) \leq \hat{u}_i(\omega^* \mid_{WAR}) \tag{C75}$$

We need to consider three sub-cases depending on when did the negotiation end:

2-1) If the negotiation ends before k then $\hat{s}$ and WAR behave in the same way:

$$T_{\hat{s}} < k \implies \text{EU}^\lambda(\hat{s}) = \text{EU}^\lambda(WAR) \wedge T_{WAR} = T_{\hat{s}}$$

2-2) If the negotiation ends at step k when $\hat{s}$ is used by i (i.e. $T_{\hat{s}} = i$), then we have three possibilities:

2-2-1) The negotiation ended with the outcome i accepted. In this case, it will be the same for $\hat{s}$ and WAR because both use ρ_{WAR}:

$$T_{\hat{s}} = k \wedge \omega_* = \omega_j^k \wedge a_{ii} \implies \text{EU}^\lambda(\hat{s}) = \text{EU}^\lambda(WAR)$$

2-2-2) $\hat{\omega}$ is the agreement which means that WAR accepted $\hat{s}$ but this implies that it would have accepted any outcome $\omega \succcurlyeq_i \omega_{\hat{s}}$ including $\omega^* \mid_{WAR}$ (Eq. B9). From Eq. C75 and Eq. C68, we therefore conclude that:

$$T = k \wedge \omega_* = \hat{\omega} \implies \text{EU}^\lambda(\hat{s}) \leq \text{EU}^\lambda(WAR)$$

2-2-3) The negotiation ends with disagreement. Because η_i is the only policy different from the case when i uses WAR, this can only change the situation if i ends the negotiation by sending ϕ or by repeating its last offer. In both cases, i receives its reserved value. Given that TAU with WAR is Exactly Rational (Theorem 2), the expected utility cannot go up:

$$T = k \wedge \omega_* = \phi \implies \text{EU}^\lambda(\hat{s}) \leq \text{EU}^\lambda(WAR)$$

Combing cases 2-2-1, 2-2-2, 2-2-3, we have:

$$T = k \implies \text{EU}^\lambda(\hat{s}) \leq \text{EU}^\lambda(WAR)$$

2-3) The negotiation ends after round k.

If the negotiation ends after round k, we have two possibilities based on the time at which the negotiation ends ($T_{\hat{s}}$) relative to what would have happened if η_{WAR} was used (T_{WAR}) instead of $\eta_{\hat{s}}$:

2-3-1) Shorter or same length negotiation: $k < T_{\hat{s}} \leq T_{WAR}$.

In this case, offering $\hat{\omega}$ can only change the final agreement by changing the set of outcomes accepted by partner j:

$$\therefore T_{\hat{s}} \leq T_{WAR} \wedge \text{EU}^\lambda(\hat{s}) \neq \text{EU}^\lambda(WAR) \implies \mathbb{A}_j^{T_{\hat{s}}} \neq \mathbb{A}_j^{T_{WAR}}$$

because $\eta_{\hat{s}}$ has no effect on the order of outcomes in Ω_j.

Moreover, the only way $\mathbb{A}_j^{T_{\hat{s}}} \neq \mathbb{A}_j^{T_{WAR}}$ can affect $\text{EU}^\lambda(\hat{s})$ is by changing the final agreement which implies that WAR accepts some outcome offered by i at step $j > i$ that it would not have accepted if i did not offer $\hat{\omega}$ at step i but we know that $\hat{\omega}$ is rejected by j because the negotiation did not end at step i.

$$\begin{aligned}
\therefore \mathbb{A}_j^{T_{\hat{s}}} \neq \mathbb{A}_j^{T_{\hat{s}}} &\implies \exists i < j \leq T_{\hat{s}} : \hat{a}_j^{T_{\hat{s}}}(\omega_i^j) \\
&\wedge \omega_* = \omega_i^j \\
&\implies \omega^* \mid_{WAR} \succ_i \omega_i^j \\
&\implies \hat{u}_i(\omega^* \mid_{WAR}) > \hat{u}_i(\omega_i^j) \\
&\implies \hat{a}_j^{T_{\hat{s}}}(\omega_i^j) \\
&\implies \hat{a}_j^{T_{\hat{s}}}(\omega^* \mid_{WAR})
\end{aligned}$$

where we used Eq. B9 in the last step. But this implies that if agent i did offer $\omega^* \mid_{WAR}$ at any point during the negotiation, it would have been accepted changing the agreement from ω_i^j to $\omega^* \mid_{WAR}$ which we have already shown to have a higher utility.

$$\begin{aligned}
\therefore T_{\hat{s}} \leq T_{WAR} \wedge \text{EU}^\lambda(\hat{s}) \neq \text{EU}^\lambda(WAR) \\
\implies \text{EU}^\lambda(\hat{s}) < \text{EU}^\lambda(WAR) \\
\therefore k < T_{\hat{s}} \leq T_{WAR} \implies \text{EU}^\lambda(\hat{s}) \leq \text{EU}^\lambda(WAR)
\end{aligned} \tag{C76}$$

2-3-2) Longer negotiation: $k < T_{\hat{s}} \wedge T_{\hat{s}} > T_{WAR}$.

To force negotiation extension, i (employing $\hat{s}$) cannot offer the agreement WAR would have reached (because the acceptance policies of both agents and the offering policy of j did not change):

$$\omega_* \mid_{WAR} \notin \Omega_i^{T_{WAR}} \mid_{\hat{s}}$$
$$\omega_* \mid_{WAR} \in \Omega_i^{T_{WAR}} \mid_{WAR}$$

In this case, the change in the agreement must come through a change in a sequence of acceptances (Fig. C4):

$$\begin{aligned} \therefore T_{\hat{s}} > T_{WAR} \wedge \mathrm{EU}^{\lambda}(\hat{s}) \neq \mathrm{EU}^{\lambda}(WAR) \\ \implies \exists z \in \mathcal{N} : \mathbb{A}_z^{T_{\hat{s}}} \neq \mathbb{A}_z^{T_{WAR}} \end{aligned} \tag{C77}$$

If $z \neq i$, we have the same case as in 2-3-1) and we have shown that this cannot lead to an improved expected utility which implies that the change must come from agent i:

$$\begin{aligned} \therefore T_{\hat{s}} > T_{WAR} \wedge \mathrm{EU}^{\lambda}(\hat{s}) \neq \mathrm{EU}^{\lambda}(WAR) &\implies \mathbb{A}_i^{T_{\hat{s}}} \neq \mathbb{A}_i^{T_{WAR}} \\ &\implies \Omega_j^{T_{\hat{s}}} \neq \Omega_j^{T_{WAR}} \end{aligned}$$

From Lemma 4, we know that the first T_{WAR} outcomes in $\Omega_j^{T_{\hat{s}}}$ cannot be changed.

$$\begin{aligned} \therefore T_{\hat{s}} > T_{WAR} \wedge \mathrm{EU}^{\lambda}(\hat{s}) \neq \mathrm{EU}^{\lambda}(WAR) \\ \implies \exists j \in \mathbb{N}^+ : T_{WAR} < j \leq T_{\hat{s}} \wedge \omega_* = \omega_j^j \end{aligned} \tag{C78}$$

Intuitively, this means that the only way for i to improve its utility using $\hat{s}$ over what it could have achieved using WAR is to extend the negotiation forcing its partners to propose more outcomes and guarantee that the negotiation ends with agreement on one of these newly proposed outcomes.

In summary, i needs to have ω_i^i be a *distracting outcome*. This outcome must be dominated by the final agreement for WAR (i.e. $\omega_i^i \prec_i \omega_* \mid_{WAR}$) before T_{WAR} (i.e. $i < T_{WAR}$). From the first requirement we know that $\omega_i^i \notin \Omega_i^i$ and from the second we know that $\omega_* \mid_{WAR} \notin \Omega_i^i$:

$$\therefore \omega_i^i, \omega_* \mid_{WAR} \notin \Omega_i^i$$

From Lemma 6, we conclude that

$$\therefore p_i(\omega_i^i \succ_j \omega_* \mid_{WAR}) = p_i(\omega_* \mid_{WAR} \succ_j \omega_i^i)$$

Therefore, it is impossible to find $\omega_i^i \prec_j \omega_*$ and in this case also there is no way to improve the expected utility of the agent

$$\therefore k < T_{\hat{s}} \wedge T_{\hat{s}} > T_{WAR} \implies \mathrm{EU}^{\lambda}(\hat{s}) \leq \mathrm{EU}^{\lambda}(WAR)$$

Which concludes the proof for Cass 2-3

$$\therefore T > i \implies \mathrm{EU}^{\lambda}(\hat{s}) \leq \mathrm{EU}^{\lambda}(WAR)$$

From cases 2-1, 2-2, 2-3, varying the offering policy also does not lead to a better response to WAR.

$$\therefore WAR \in \mathcal{B}(WAR)$$

□

References

[1] Mohammad, Y.: Generalized bargaining protocols. In: Australasian Joint Conference on Artificial Intelligence, pp. 261–273 (2023). Springer

[2] PRNewswire: Digital Process Automation Market by Component, Business Function, Deployment Type, Organization Size, Industry Vertical And Region - Global Forecast to 2023. Accessed: May 19th 2021. http://tiny.cc/573o9y

[3] De La Hoz, E., Marsa-Maestre, I., Gimenez-Guzman, J.M., Orden, D., Klein, M.: Multi-agent nonlinear negotiation for wi-fi channel assignment. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, pp. 1035–1043 (2017). International Foundation for Autonomous Agents and Multiagent Systems

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

[5] Jonge, D., Bistaffa, F., Levy, J.: Multi-objective vehicle routing with automated negotiation. Applied Intelligence, 1–24 (2022)

[6] Inotsume, H., Aggarwal, A., Higa, R., Nakadai, S.: Path negotiation for self-interested multirobot vehicles in shared space. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 11587–11594 (2020). IEEE

[7] Johnson, E., Gratch, J., default, D.: Towards an autonomous agent that provides automated feedback on students' negotiation skills. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, pp. 410–418 (2017). International Foundation for Autonomous Agents and Multiagent Systems

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

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

[10] Aydoğan, R., Festen, D., Hindriks, K.V., Jonker, C.M.: Alternating offers protocols for multilateral negotiation. Modern approaches to agent-based complex automated negotiation, 153–167 (2017)

[11] Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y.: Protocols for negotiating complex contracts. IEEE Intelligent Systems **18**(6), 32–38 (2003)

[12] Baarslag, T., Hindriks, K., Hendrikx, M., Dirkzwager, A., Jonker, C.: Decoupling negotiating agents to explore the space of negotiation strategies. In: Marsa-Maestre, I., Lopez-Carmona, M.A., Ito, T., Zhang, M., Bai, Q., Fujita, K. (eds.) Novel Insights in Agent-based Complex Automated Negotiation, pp. 61–83. Springer, Berlin, Heidelberg (2014)

[13] Baarslag, T., Gerding, E.H., Aydogan, R., Schraefel, M.: 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, pp. 190–197 (2015). IEEE

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

[15] Jonge, D.: Theoretical properties of the micro negotiation strategy. Autonomous Agents and Multi-Agent Systems **38**(46), 1–43 (2024)

[16] Mohammad, Y.: Evaluating automated negotiations. In: IEEE International Conference on Agents (ICA) (2023). IEEE

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

[18] Kalai, E., Smorodinsky, M.: Other solutions to nash's bargaining problem. Econometrica: Journal of the Econometric Society, 513–518 (1975)

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

[20] Aydoğan, R., Baarslag, T., Fujita, K., Hoos, H.H., Jonker, C.M., Mohammad, Y., Renting, B.M.: The 13th international automated negotiating agent competition challenges and results. In: International Joint Conference on Artificial Intelligence, pp. 87–101 (2022). Springer

[21] Mohammad, Y., Greenwald, A., Nakadai, S.: Negmas: A platform for situated negotiations. In: Twelfth International Workshop on Agent-based Complex Automated Negotiations (ACAN2019) in Conjunction with IJCAI (2019)

[22] Lin, R., Kraus, S., Baarslag, T., Tykhonov, D., Hindriks, K., Jonker, C.M.: Genius: An integrated environment for supporting the design of generic automated negotiators. Computational Intelligence **30**(1), 48–70 (2014) https://doi.org/10.1111/j.1467-8640.2012.00463.x

[23] Baarslag, T., Hindriks, K., Jonker, C.: A tit for tat negotiation strategy for real-time bilateral negotiations. In: Complex Automated Negotiations: Theories, Models, and Software Competitions, pp. 229–233. Springer, Berlin, Heidelberg (2013)

[24] Gunes, T.D., Arditi, E., Aydogan, R.: Collective voice of experts in multilateral negotiation. In: PRIMA 2017: Principles and Practice of Multi-Agent Systems: 20th International Conference, Nice, France, October 30–November 3, 2017, Proceedings 20, pp. 450–458 (2017). Springer

[25] Mori, A., Ito, T.: Atlas3: a negotiating agent based on expecting lower limit of concession function. In: Modern Approaches to Agent-based Complex Automated Negotiation, pp. 169–173. Springer, Berlin, Heidelberg (2017)

[26] Kawaguchi, S., Fujita, K., Ito, T.: In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) AgentK2: Compromising Strategy Based on Estimated Maximum Utility for Automated Negotiating Agents, pp. 235–241. Springer, Berlin, Heidelberg (2013)

[27] Krimpen, T., Looije, D., Hajizadeh, S.: In: Ito, T., Zhang, M., Robu, V., Matsuo, T. (eds.) HardHeaded, pp. 223–227. Springer, Berlin, Heidelberg (2013)

[28] Hao, J., Leung, H.-f.: In: Marsa-Maestre, I., Lopez-Carmona, M.A., Ito, T., Zhang, M., Bai, Q., Fujita, K. (eds.) CUHKAgent: An Adaptive Negotiation Strategy for Bilateral Negotiations over Multiple Items, pp. 171–179. Springer, Berlin, Heidelberg (2014)

[29] Aydoğan, R., Baarslag, T., Fujita, K., Mell, J., Gratch, J., De Jonge, D., Mohammad, Y., Nakadai, S., Morinaga, S., Osawa, H., *et al.*: Challenges and main results of the automated negotiating agents competition (anac) 2019. In: Multi-Agent Systems and Agreement Technologies: 17th European Conference, EUMAS 2020, and 7th International Conference, AT 2020, Thessaloniki, Greece, September 14-15, 2020, Revised Selected Papers 17, pp. 366–381 (2020). Springer

[30] Lemke, C.E., Howson, J.T. Jr: Equilibrium points of bimatrix games. Journal of the Society for industrial and Applied Mathematics **12**(2), 413–423 (1964)

[31] Endriss, U.: Monotonic concession protocols for multilateral negotiation. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 392–399 (2006)

[32] Fatima, S., Kraus, S., Wooldridge, M.: Principles of Automated Negotiation.

Cambridge University Press, USA (2014)

[33] Koça, T., Jonge, D., Baarslag, T.: Search algorithms for automated negotiation in large domains. Annals of Mathematics and Artificial Intelligence, 1–22 (2023)