Correlation Analysis of Social Welfare and the Gini Coefficient in Multi-Party Automated Negotiations

Shin Sakakibara[1], Shun Okuhara[1]

[1]Department of Information Engineering, Mie University, Tsu, Japan

Abstract

Social Welfare is widely used to evaluate the efficiency of automated negotiation outcomes, but aggregate utility alone does not capture how benefits are distributed among participants. The Gini coefficient has been introduced to measure this inequality in multi-party negotiation, yet it remains unclear whether its relationship with Social Welfare reflects negotiation or the structure of the bounded utility space. This study examines the relationship in trilateral negotiations on the Jupiter platform with four combinations of time-dependent concession strategies. Preference profiles are shared across all scenarios and a random baseline, and all bids of each profile are enumerated so that every agreement can be compared with its alternatives. Social Welfare and the Gini coefficient are negatively correlated in every scenario and also across all enumerated bids, indicating that the correlation is largely a property of the utility space. Conditional on Social Welfare, agreements between homogeneous strategies are more equal than expected, whereas agreements in the heterogeneous scenario are less equal, with the Boulware agent obtaining the highest utility. These results show that the Gini coefficient reveals distributional differences that Social Welfare alone does not.

Keywords: Automated Negotiation, Social Welfare, Gini Coefficient

1. Introduction

1.1. Background of the Study

Automated negotiation is a technology in which multiple autonomous agents (AI) with conflicting interests act as proxies for humans, autonomously proposing terms and making concessions to one another in an effort to reach an agreement. Automated negotiation has rapidly gained attention in recent years as an AI-powered autonomous decision-support technology, and negotiation problems involving the distribution of benefits among multiple agents are highly anticipated to find applications in various real-world fields, such as traffic flow management, electricity trading in smart grids, and supply chain coordination in logistics [1]–[3].

In such automated negotiations, "fairness" is considered one of the most critical factors when evaluating the proposed agreements. Traditionally, "social welfare" has been widely used as a representative evaluation metric in automated negotiation. While social welfare is a metric that measures the total utility gained by all agents, there is a challenge in that pursuing only the maximization of this total may fail to detect situations where specific individual agents suffer significant disadvantages or dissatisfaction. A biased agreement outcome that lacks fairness may fail to gain the acceptance of the humans they represent and could significantly undermine trust in the agent system. Therefore, when considering the real-world application of automated negotiation and the maintenance of sustainable cooperative relationships, it is important to go beyond mere global optimization and quantitatively evaluate how utility is distributed among participants.

Several measures have been proposed to evaluate the fairness of negotiation outcomes. In multiagent resource allocation, negotiated allocations have been analyzed under notions of social welfare other than the utilitarian sum, such as egalitarian welfare and Lorenz optimality [4]. For negotiation outcomes, conflict degree [5] measures the geometric distance from the axis of symmetry in utility space. Its formulation, based on a two-dimensional utility space, is primarily applicable to bilateral negotiations. To address multi-party settings, Hara [6] introduced the Gini coefficient, a measure of distributional inequality commonly used in economics, and compared its characteristics with those of the Theil index [7]. Hara's study provided an approach to evaluating inequality among multiple negotiating agents but did not examine the relationship between Social Welfare and the Gini coefficient. Moreover, the evaluation was conducted in a simplified, self-developed simulation environment, leaving this relationship unexamined in widely used automated negotiation environments.

1.2. Research Purpose

The purpose of this study is to clarify how the negative correlation between Social Welfare and the Gini coefficient in multi-party automated negotiation is shaped by two factors: the mathematical constraints of the bounded utility space, and the selection of particular outcomes through negotiation. Using the automated negotiation platform "Jupiter" [8], we conduct trilateral negotiations and compare the agreements with the full set of possible outcomes under the same preference profiles. In particular, we ask whether negotiated agreements are more or less equal than other outcomes with the same Social Welfare.

The contribution of this study is empirical: it separates the correlation produced by the bounded utility space from the differences in how negotiated agreements are distributed within that space. Because the domain is small, all outcomes of each profile can be enumerated, and agreements can be compared

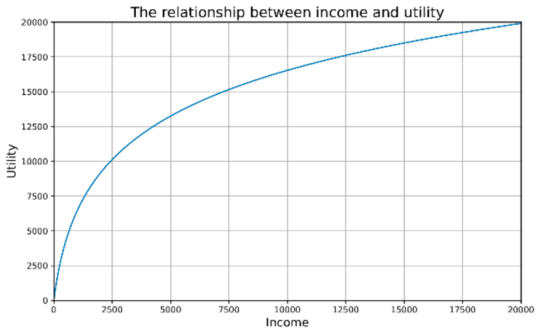


Fig. 1: The relationship between income and utility.

exactly with every alternative under the same preferences. We therefore position this study as a detailed case study of this separation.

2. Performance Metrics

2.1. Social Welfare

Let $A = \{1, 2, \ldots, n\}$ denote the set of participating agents in the negotiation. When an agreement ω_{agree} is reached, the utility achieved by each agent i is denoted by $U_i(\omega_{\text{agree}}) \in [0.0, 1.0]$. Social Welfare ($SW$), which measures collective efficiency, is defined as the arithmetic sum of the utilities attained by all agents [3]:

$$SW = \sum_{i=1}^{n} U_i(\omega_{\text{agree}}). \tag{1}$$

Because individual utilities are normalized to $[0.0, 1.0]$, the domain of Social Welfare in a trilateral negotiation ($n = 3$) is bounded within $[0.0, 3.0]$. If a negotiation breaks down and all agents receive zero utility, $SW = 0.0$; conversely, in an ideal agreement where all agents simultaneously achieve complete satisfaction, SW reaches its theoretical maximum of 3.0.

2.2. The Gini Coefficient

To evaluate the inequality of utilities in an agreement, we adopt the Gini coefficient formulation introduced by Hara et al. In this paper, the Gini coefficient is treated as a measure of utility inequality among agents, not as a comprehensive measure of fairness; fairness in a broader normative sense also depends on reservation utilities, bargaining power, and other contextual factors. The Gini coefficient originated in welfare economics to evaluate inequality in wealth and income distributions [9], [10]. The model by Hara et al. incorporates the Law of Diminishing Marginal Utility into agent evaluations. The Law of Diminishing Marginal Utility states that as an individual's income or wealth increases, the marginal increase in utility obtained from an additional unit of income progressively diminishes. In other words, a small gain provides substantial satisfaction in a state of deprivation, whereas the same incremental gain yields only marginal psychological satisfaction in an already affluent state (Fig. 1).

Following the design of Hara et al., the trilateral Gini coefficient G is computed based on transformed utility evaluations that incorporate diminishing marginal utility [6]:

$$G = \frac{\sum_{i=1}^{3}\sum_{j=1}^{3}|Y_i - Y_j|}{2 \cdot 3^2 \cdot \overline{Y}}, \qquad Y_i = \ln(1 + \alpha U_i) \quad (\alpha = 100). \tag{2}$$

Following the formulation in Hara [6], each agent's utility U_i is mapped to an underlying evaluation proxy Y_i via a concave logarithmic transformation, where $\alpha > 0$ is a scaling parameter that governs inequality sensitivity. Here, Y_i is the transformed evaluation value of agent i, and $\overline{Y} = \frac{1}{3}\sum_{i=1}^{3} Y_i$ represents the mean evaluation across all participants. (G is undefined if all utilities are zero, which cannot occur here: all sessions reached agreement, and every bid gives each agent a utility of at least 0.1.) A value of $G = 0.0$ signifies complete equality (all agents obtain identical payoffs). In a population of n agents, G is bounded above by $(n-1)/n$; for $n = 3$, values approaching 0.667 indicate that one agent obtains almost all of the surplus.

3. Experimental Configuration

We implemented the negotiation environment using the Jupiter platform under the Stacked Alternating Offers Protocol (SAOP), which is the standard protocol for multilateral negotiations [8], [11].

3.1. Negotiation Strategy Models

We deployed three representative time-dependent concession strategies (Conceder, Linear, and Boulware). At negotiation time $t \in [0, T_{\max}]$, an agent's target utility $U(t)$ is defined as [12]:

$$U(t) = U_{\max} - (U_{\max} - U_{\min}) \left(\frac{t}{T_{\max}}\right)^{p}, \tag{3}$$

where we set $U_{\max} = 1.0$ and $U_{\min} = 0.0$ across all agents.

The concession exponent p dictates the pace of concession (Fig. 2):

- Conceder tactic ($p = 0.2$): A cooperative strategy that concedes rapidly in the early stages of negotiation.
- Linear tactic ($p = 1.0$): A neutral strategy that concedes at a constant, steady rate over time.
- Boulware tactic ($p = 5.0$): A stubborn strategy that maintains high aspiration levels until near the deadline and concedes abruptly at the end.

3.2. Scenarios and Experimental Settings

Four scenarios were configured:

- Scenario A (Heterogeneous): Conceder, Linear, and Boulware agents, proposing in this order.
- Scenario B (Conceder $\times$ 3): Three Conceder agents.
- Scenario C (Linear $\times$ 3): Three Linear agents.
- Scenario D (Boulware $\times$ 3): Three Boulware agents.

The experimental parameters are as follows:

- Protocol: SAOP with a fixed round-robin proposal order (Agent 1 $\rightarrow$ Agent 2 $\rightarrow$ Agent 3).

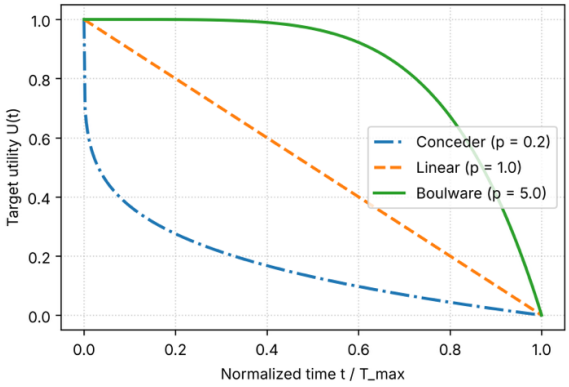


Fig. 2: Target utility dynamics across concession tactics.

- Domain: Atlas3 / triangularFight (2 issues, 3 values each, 9 bids).
- Preferences: Issue weights are drawn from a Dirichlet(1, 1) distribution, and value evaluations from $\mathcal{U}(0.1, 1.0)$, normalized so that the best value of each issue is 1. Utilities are linear additive without discounting, and the reservation value is 0.
- Deadline: 100 rounds, in which every agent acts once per round.
- Trials: 1,000 preference profiles. Each profile is generated once and used in all four scenarios and in the random baseline, so that all comparisons are paired by profile.
- Random baseline: For each profile, one of the 9 bids is selected uniformly at random as the agreement.

An agent i accepts the offer ω it last received at time t if and only if it satisfies its current target:

$$\text{Accept } \omega \iff U_i(\omega) \geq U_i(t). \tag{4}$$

Otherwise, the agent proposes a bid with $U_i(\omega) \geq U_i(t)$ generated by the standard bid generator of Jupiter. Following SAOP, an offer becomes an agreement when all the other agents accept it. All 4,000 negotiation sessions reached agreement before the deadline.

3.3. Analysis

Because the domain has only nine bids, we enumerated all bids of every profile ($9 \times 1000 = 9000$ bids) and computed their SW and G. This enumerated set represents the outcome space without negotiation and serves as the reference for the following analyses.

1) Gini at fixed Social Welfare within a profile: We fit a model with profile fixed effects to the enumerated bids,

$$G_{tk} = \mu_t + \beta_1 SW_{tk} + \beta_2 SW_{tk}^2 + \beta_3 SW_{tk}^3 + \varepsilon_{tk}, \tag{5}$$

where t is the profile and k the bid. The Gini residual of an agreement ω in profile t is $e(\omega) = G(\omega) - \hat{\mu}_t - \hat{f}(SW(\omega))$, where $\hat{f}$ is the fitted cubic term. A negative mean residual indicates that agreements exhibit lower Gini coefficients than predicted by the fitted model, conditional on SW and profile fixed effects. As a direct check, we also compare each agreement with the other bids of the same profile whose SW differs by at most 0.1; such bids exist in 52%–72% of sessions.

2) SW-matched reference correlation: We resample the enumerated bids so that their SW histogram (bin width 0.05) matches that of each scenario and compute Pearson r (500 repetitions). Matching equalizes the distribution of SW but not the distribution of G at a given SW. A difference from the matched value therefore shows that agreements are located differently from the outcome space at given SW levels; it does not by itself measure the effect of a strategy.

3) Reference bids: For each profile, we identify the bid maximizing SW, the Nash bid maximizing $\prod_i U_i$ [13], and the bid minimizing G, and record whether the agreement coincides with each.

4) Utilities by strategy: In Scenario A, we compare the utilities obtained by the three strategies within each session.

Confidence intervals are obtained by bootstrap over profiles (4,000 resamples). Because all scenarios share the same profiles, comparisons between scenarios use per-profile differences.

4. Results

4.1. Correlation between Social Welfare and the Gini Coefficient

Table 1 summarizes the 1,000 sessions of each scenario and the random baseline. In the random baseline, the mean Social Welfare was 2.1293, the mean Gini coefficient was 0.0359, and the correlation was $r = -0.5767$ (Spearman $\rho = -0.6202$). In all negotiated scenarios, the correlation was also negative ($r = -0.6272$ to -0.8836), with the strongest value in Scenario A ($r = -0.8836$, $\rho = -0.9235$). Compared with the baseline, the mean Gini coefficient was lower by 22.5% in Scenario A and by 55.4% to 66.8% in Scenarios B–D, and the Pareto optimality rate rose from 54.1% to 92.6%–99.7%.

In the scatter plots, baseline outcomes spread over $SW \in [1.0, 3.0]$ with Gini values up to 0.122 (Fig. 3), Scenario A ranges over $SW \in [1.55, 3.0]$ with Gini values up to 0.101 (Fig. 4), and agreements in Scenarios B–D concentrate in the region $SW \geq 2.2$ and $G \leq 0.03$ (Figs. 5 to 7).

In Fig. 8, the maximum Gini coefficient in Scenarios B–D is 0.061, and most of their agreements lie below $G = 0.015$.

In all scatter plots, no point lies above a downward-convex curve that reaches $(3.0, 0.0)$. Because individual utilities cannot exceed 1.0, the room for unequal allocations shrinks as SW approaches 3.0, and the concave transformation in (2) further compresses differences among high utilities [6]. The upper bound of G therefore decreases to zero along this curve.

The negative correlation is also a property of the outcome space itself. Across all 9,000 enumerated bids, $r = -0.5945$, and within individual profiles the correlation over the nine bids was negative in 93.4% of profiles (median $r = -0.626$).

Sensitivity to α: We recomputed the correlation using raw utilities and $\alpha = 1$, 10, and 100. In Scenario A, r was -0.9413, -0.9245, -0.8945, and -0.8836, respectively, and

Table 1: Statistical summary of trilateral negotiations ($N = 1000$ per scenario). All sessions reached agreement.

Scenario	Mean SW	Mean Gini	Pearson r	r 95% CI	p	Spearman ρ	Pareto rate	SW / max SW	Nash bid rate	Mean agree time
Random Baseline	2.1293	0.0359	-0.5767	$[-0.617, -0.535]$	1.00×10^{-89}	-0.6202	0.5410	0.7997	0.0970	–
A (Heterogeneous)	2.4735	0.0278	-0.8836	$[-0.897, -0.869]$	$< 10^{-300}$	-0.9235	0.9970	0.9251	0.3230	0.2903
B (Conceder × 3)	2.5506	0.0160	-0.6272	$[-0.679, -0.573]$	1.95×10^{-110}	-0.6836	0.9260	0.9544	0.4690	0.0177
C (Linear × 3)	2.6296	0.0119	-0.6773	$[-0.721, -0.631]$	3.34×10^{-135}	-0.7444	0.9930	0.9830	0.7270	0.2415
D (Boulware × 3)	2.6251	0.0121	-0.6910	$[-0.738, -0.642]$	7.90×10^{-143}	-0.7534	0.9860	0.9811	0.7230	0.7220

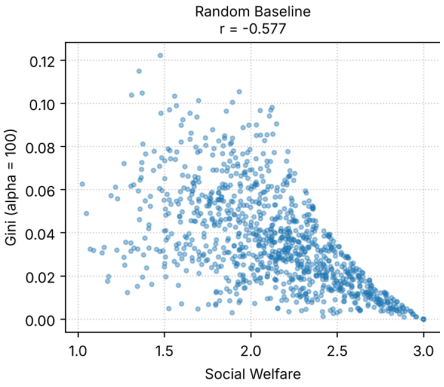


Fig. 3: Scatter plot of Social Welfare vs. Gini coefficient in the Random Baseline.

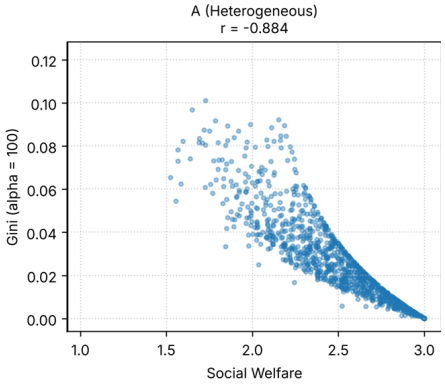


Fig. 4: Scatter plot of Social Welfare vs. Gini in Heterogeneous Trilateral Negotiation (Scenario A).

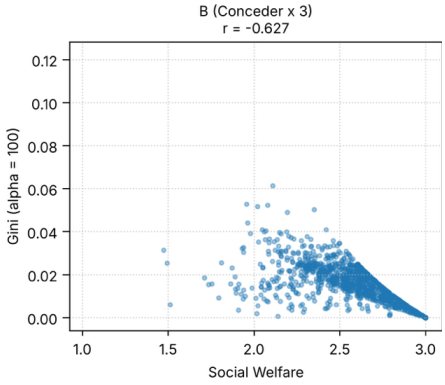


Fig. 5: Scatter plot of Social Welfare vs. Gini in Conceder × 3 (Scenario B).

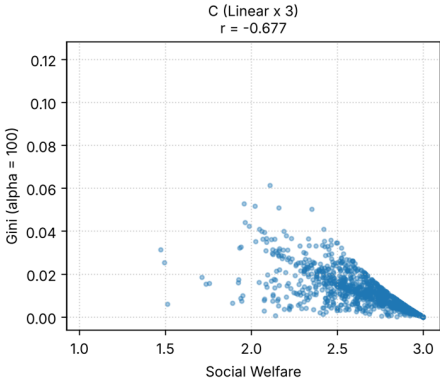


Fig. 6: Scatter plot of Social Welfare vs. Gini in Linear × 3 (Scenario C).

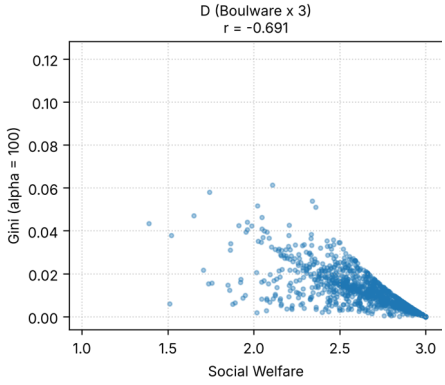


Fig. 7: Scatter plot of Social Welfare vs. Gini in Boulware × 3 (Scenario D).

Table 2: Comparison with the enumerated outcome space. Residuals are in units of 10^{-3}.

	Residual e		Near-	Pearson r		
	Mean	95% CI	SW	Obs.	Matched	95% CI
Random	0.4	$[-0.6, 1.3]$	1.0	-0.577	-0.589	$[-0.622, -0.550]$
A	4.4	$[3.7, 5.1]$	5.3	-0.884	-0.750	$[-0.777, -0.720]$
B	-4.3	$[-5.1, -3.6]$	-7.4	-0.627	-0.758	$[-0.784, -0.730]$
C	-4.9	$[-5.6, -4.2]$	-11.2	-0.677	-0.797	$[-0.825, -0.767]$
D	-4.9	$[-5.5, -4.2]$	-11.2	-0.691	-0.800	$[-0.828, -0.771]$

in Scenarios B–D it ranged from -0.627 to -0.706. The negative correlation was therefore maintained across the tested values of α.

4.2. Comparison with the Enumerated Outcome Space

Table 2 reports the analyses of Section 3.3. The random baseline serves as a check: its residuals are indistinguishable from zero and its correlation equals the matched value.

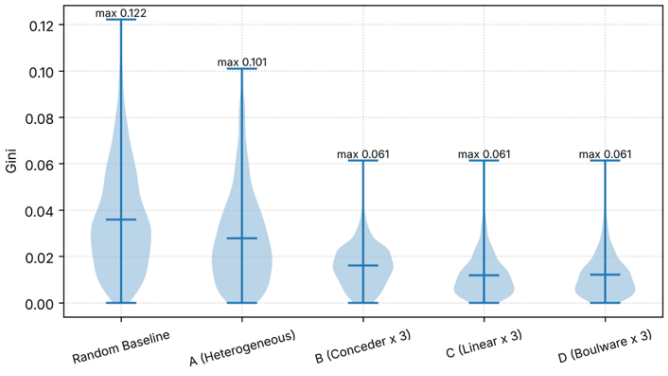


Fig. 8: Distribution of the Gini coefficient across scenarios (violin plot).

Gini at fixed Social Welfare: Conditional on SW and profile fixed effects, agreements in Scenarios B–D had lower Gini coefficients than predicted (mean e from -4.9×10^{-3} to -4.3×10^{-3}), and agreements in Scenario A had higher Gini coefficients (4.4×10^{-3}). The direct comparison with bids of similar SW in the same profile ("Near-SW" in Table 2) gave the same signs. Paired by profile, the residual in Scenario A exceeded that in Scenarios B, C, and D by 8.7, 9.3, and 9.2×10^{-3}, respectively (all 95% CIs above zero). The signs were unchanged for raw utilities and for $\alpha = 1$ and 10 (Scenario A: 17.0, 13.7, and 7.8×10^{-3}; Scenarios B–D: -17.3 to -7.6×10^{-3}, all CIs excluding zero), and also when the reference was pooled across profiles without fixed effects.

To give a sense of scale, at $SW = 2.62$ the utility vectors $(1.00, 0.82, 0.80)$ and $(1.00, 0.90, 0.72)$ have $G = 0.011$ and 0.016. A difference of 0.005 in G thus corresponds to moving about 0.08 of utility from the worst-off agent to another agent. Consistent with this, the mean utility of the worst-off agent was 0.774 in Scenario C, 0.617 in Scenario A, and 0.493 in the baseline.

Correlation: When the enumerated bids were resampled to the SW distribution of each scenario, the matched correlation was -0.750 to -0.800. The observed correlation in Scenarios B–D (-0.627 to -0.691) was weaker than this value, while that in Scenario A (-0.884) was stronger. Paired by profile, the difference from the baseline correlation was $\Delta r = -0.307$ (95% CI $[-0.353, -0.265]$) in Scenario A, -0.050 ($[-0.113, 0.015]$) in B, -0.101 ($[-0.163, -0.041]$) in C, and -0.114 ($[-0.174, -0.056]$) in D.

Reference bids: The max-SW bid coincided with the min-G bid in only 49.6% of profiles, whereas it coincided with the Nash bid in 94.0%. Agreements matched the min-G bid in 20.0% (A) to about 66% (C, D) of sessions.

4.3. Utilities by Strategy in Scenario A

In Scenario A, the mean utility was 0.993 (95% CI $[0.992, 0.995]$, median 1.000) for the Boulware agent, 0.758 ($[0.746, 0.770]$) for the Linear agent, and 0.722 ($[0.708, 0.736]$) for the Conceder agent. The Boulware agent obtained a higher utility than the Conceder agent in 86.6% of sessions and a lower utility in 3.1% (mean difference 0.271, Wilcoxon $p < 10^{-140}$). To gauge the effect of the fixed proposal order, we compared agents at positions 1 and 3 in the homogeneous scenarios: the mean difference was 0.054 in Scenario B and below 0.003 in Scenarios C and D.

5. Discussion

5.1. Mathematical Constraints and Outcome Selection

The results separate two sources of the negative correlation.

1) The bounded utility space: Because utilities are bounded by 1.0, an agreement with $SW = 3.0$ must give every agent utility 1.0 and thus $G = 0.0$, and the upper bound of G decreases as SW increases. This constraint explains the correlation in the random baseline.

2) Selection of outcomes through negotiation: Negotiation moved agreements toward high SW, where the range of possible Gini values is narrow. Most of the stronger correlation in the negotiated scenarios follows from this shift: once the enumerated bids are matched to the same SW distribution, they show an even stronger correlation than Scenarios B–D. The change in r relative to the baseline should therefore not be read as evidence that negotiation itself strengthens the relationship. Because matching does not equalize the distribution of G at a given SW, the remaining gap describes where agreements lie in the outcome space, not the effect of a strategy. Homogeneous scenarios selected bids that are more equal than expected given their SW, which weakens the dependence of G on SW. Scenario A selected bids that are less equal, particularly at lower SW, which strengthens it.

5.2. Strategic Heterogeneity and Distributional Inequality

Scenario A shows a higher mean Gini coefficient and lower mean Social Welfare than the three homogeneous scenarios, and its agreements are less equal than expected given their SW and profile. The per-agent utilities show where this inequality lies: under the order Conceder $\rightarrow$ Linear $\rightarrow$ Boulware, the Boulware agent obtained nearly its maximum utility in most sessions, while the Conceder agent obtained 0.27 less on average.

Previous research has shown that the order of negotiation affects outcomes in multilateral negotiation and has proposed a protocol that adjusts the order through pre-negotiation [14]. In the present study, the proposal order is fixed; therefore, the effects of strategy composition and proposal position cannot be fully separated. The position effect in the homogeneous scenarios (at most 0.054) is much smaller than the gap of 0.271, which suggests that position alone does not produce it. This comparison, however, places different strategies at each position. Separating the two effects requires running all six assignments of strategies to positions on the same profiles.

In Scenarios B–D, agreements were more equal than expected given their SW, and the mean Gini coefficient was less than half of the baseline value. When all agents follow the same concession curve, they tend to agree on the bid that maximizes the minimum utility (Section 5.3). In Scenarios C and D, the agreements coincided with the Nash bid in about 72% of sessions (Table 1).

5.3. Temporal Characteristics of the Homogeneous Scenarios

The three homogeneous scenarios yield similar outcomes but differ in agreement time. In Scenario B, agreement was reached almost immediately (mean normalized agreement time $t_{\mathrm{agree}}/T_{\max} = 0.0177$, round $\sim$2), with a Nash bid rate of 46.9% and a Pareto rate of 92.6%. In Scenario C, agreement was reached at $t_{\mathrm{agree}} = 0.2415$ (round $\sim$24), with the highest SW (2.6296, 98.3% of the maximum), the lowest Gini (0.0119), and a Nash bid rate of 72.7%. In Scenario D, agreement was reached at $t_{\mathrm{agree}} = 0.7220$ (round $\sim$72) with almost the same outcome as Scenario C ($SW = 2.6251$, Nash bid rate 72.3%); the two scenarios reached the same bid in 86.4% of profiles.

Thus, Linear $\times$ 3 and Boulware $\times$ 3 produced outcomes of similar quality, but Boulware $\times$ 3 needed about three times as long, which is a cost in time-sensitive applications. In Scenarios C and D, the agreement time is largely determined by the concession rule: with identical concession curves, agreement occurs approximately when the common target $U(t)$ falls to $\max_\omega \min_i U_i(\omega)$, the highest utility that any bid guarantees to the worst-off agent. The predicted time $(1 - \max_\omega \min_i U_i(\omega))^{1/p}$ matched the observed agreement time with a mean absolute error of 0.016 in both scenarios, and the agreement coincided with the bid attaining this maximum in 91.3% (C) and 88.6% (D) of sessions.

5.4. Complementary Information Provided by the Gini Coefficient

Social Welfare and the Gini coefficient covary strongly, but the results show three situations in which the Gini coefficient adds information.

1) Distribution at the same Social Welfare: Agreements with the same SW can differ in inequality. In Scenario A, agreements around $SW \approx 2.1$ range from $G \approx 0.02$ to $G \geq 0.08$ (Fig. 4). The Gini residual in Table 2 shows that this is systematic: conditional on SW within each profile, Scenario A is less equal and Scenarios B–D are more equal than expected.

2) Efficiency and equality within a profile: The max-SW bid was also the min-G bid in only half of the profiles. Choosing an agreement by SW alone, or by the Nash product, which coincided with the max-SW bid in 94% of profiles, does not guarantee the most equal outcome.

3) Who bears the inequality: Aggregate SW and G do not identify which agent is disadvantaged. In Scenario A, the per-agent utilities revealed that the Conceder agent, which proposes first, bears the loss. The Gini coefficient is therefore most informative when combined with per-agent utilities.

6. Conclusion

This study investigated the relationship between Social Welfare and the Gini coefficient in trilateral automated negotiations. The two were negatively correlated in all scenarios and across all enumerated bids, so the correlation is primarily a property of the bounded utility space. Negotiated agreements had higher Social Welfare than the baseline, and most of the stronger correlation in the negotiated scenarios is explained by this shift toward high Social Welfare. Conditional on Social Welfare within each profile, homogeneous strategy combinations reached agreements that were more equal than expected, whereas the heterogeneous combination reached less equal agreements; these signs held for all tested values of α. Under the fixed proposal order of the heterogeneous scenario, the Boulware agent obtained the highest mean utility, whereas the Conceder agent obtained the lowest. These differences are not visible from Social Welfare alone, and the Gini coefficient, together with per-agent utilities, captures them.

This study is a case study of trilateral negotiation in a nine-bid domain, where full enumeration allows exact comparison with all alternatives. In this domain, the max-SW bid coincided with the Nash bid in 94% of profiles, a rate that may differ in larger domains. Future work will run all six assignments of strategies to proposal positions to separate strategy and position effects, and will extend the analysis to larger domains, continuous utility spaces, and adaptive agents. The within-profile Gini residual can also serve as an evaluation criterion for negotiation protocols and mediator agents that consider both efficiency and inequality.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP22H00533.

References

[1] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, M. J. Wooldridge, and C. Sierra, "Automated negotiation: prospects, methods and challenges," *Group Decision and Negotiation*, vol. 10, no. 2, pp. 199–215, 2001.

[2] S. Kraus, *Strategic Negotiation in Multiagent Environments*, ser. Intelligent Robots and Autonomous Agents. MIT Press, 2001.

[3] K. Fujita, "Research trends in automated negotiation in multi-agent systems," *Systems, Control and Information*, vol. 67, no. 1, pp. 25–32, 2023.

[4] U. Endriss, N. Maudet, F. Sadri, and F. Toni, "Negotiating socially optimal allocations of resources," *Journal of Artificial Intelligence Research*, vol. 25, pp. 315–348, 2006.

[5] T. Toyama and T. Ito, "A metric of opposition level between the utility information based on game theoretical fairness for multi-issue negotiation," in *Proceedings of the 31st Annual Conference of the Japanese Society for Artificial Intelligence*. The Japanese Society for Artificial Intelligence, 2017, pp. 2F2–5.

[6] A. Hara, "Proposal and comparison of fairness metrics in multi-party automated negotiation," Bachelor's thesis, Department of Information Engineering, Faculty of Engineering, Mie University, 2024.

[7] H. Theil, *Economics and Information Theory*. North-Holland, 1967.

[8] T. Fukui and T. Ito, "A proposal of automatic negotiation simulator Jupiter for negotiation agents using machine learning," in *The 32nd Annual Conference of the Japanese Society for Artificial Intelligence*, 2018.

[9] H. Suzuki, "Investigation of factor decomposition methods of Gini coefficient and application to regional wage disparities," *Journal of Regional Policy Research (Chiiki Seisaku Kenkyu)*, vol. 19, pp. 1–78, 2006.

[10] A. B. Atkinson, "On the measurement of inequality," *Journal of Economic Theory*, vol. 2, no. 3, 1970.

[11] R. Aydoğan, D. Festen, K. V. Hindriks, and C. M. Jonker, "Alternating offers protocols for multilateral negotiation," in *Modern Approaches to Agent-based Complex Automated Negotiation*, ser. Studies in Computational Intelligence. Springer, 2017, vol. 674, pp. 153–167.

[12] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation decision functions for autonomous agents," *Robotics and Autonomous Systems*, vol. 24, no. 3–4, 1998.

[13] J. F. Nash, "The bargaining problem," *Econometrica*, vol. 18, no. 2, pp. 155–162, 1950.

[14] S. Okuhara and T. Ito, "An automated negotiating protocol that adjusts the negotiation order based on pre-negotiations," *Transactions of the Japanese Society for Artificial Intelligence*, vol. 39, no. 1, pp. B-M24_1–11, 2024, (in Japanese).