

# MINiDM: Multi-Issue Negotiation in Decentralised Data Marketplaces

Soulmaz Gheisari

s.gheisari@soton.ac.uk

University of Southampton  
Southampton, United Kingdom

Semih Yumusak

semih.yumusak@soton.ac.uk

University of Southampton  
Southampton, United Kingdom

Jaime Osvaldo Salas

j.o.salas@soton.ac.uk

University of Southampton  
Southampton, United Kingdom

George Konstantinidis

g.konstantinidis@soton.ac.uk

University of Southampton  
Southampton, United Kingdom

## ABSTRACT

This paper presents MINiDM, a novel negotiation framework designed for decentralised data marketplaces. MINiDM enables AI-driven agents to negotiate complex, multi-issue data-sharing agreements while ensuring compliance with legal and ethical standards such as GDPR. The framework combines game-theoretic strategies, formal policy languages and vocabularies (ODRL and DPV), and a lightweight multi-agent commitment protocol to automate fair transactions without relying on central authorities or blockchain-based contracts. Experimental results demonstrate that MINiDM outperforms existing protocols in terms of agreement rate, fairness, and negotiation efficiency, offering a promising direction for secure and autonomous data trading in Web3 ecosystems.

## KEYWORDS

Data Consumer; Data Provider; Data-Sharing; Game Theory; Negotiation

### PVLDB Workshop Format:

Soulmaz Gheisari, Jaime Osvaldo Salas, Semih Yumusak, and George Konstantinidis. MINiDM: Multi-Issue Negotiation in Decentralised Data Marketplaces. VLDB 2025 Workshop: 3rd Data Economy Workshop (DEC).

## 1 INTRODUCTION

Decentralised data marketplaces (DDMs), enable secure, privacy-preserving, and autonomous data exchanges without relying on central intermediaries. These marketplaces can establish trust, transparency, and immutability in transactions, as demonstrated by platforms such as Ocean Protocol [40], Gaia-X [10], and IPFS [7].

Despite their potential, DDMs face significant obstacles. First, ensuring trust and fairness is challenging, as there is no overarching entity to enforce reliability or verify commitments. Second, data providers often hesitate to share valuable information due to privacy concerns, exacerbated by opaque and manual negotiation

processes. Third, scalability remains a critical issue. Current negotiation models struggle to efficiently handle high-volume, multi-party interactions involving multiple interdependent attributes.

Moreover, regulatory compliance further complicates decentralised data trading. Legal frameworks such as the GDPR [1], AI Act [3], and Data Governance Act [2] impose strict requirements for privacy, fairness, and accountability. Formal policy languages and vocabularies such as ODRL (Open Digital Rights Language) [24], and DPV (Data Privacy Vocabulary) [28] can help enforce these requirements. However, ensuring real-time, automated enforcement of these policies in decentralised settings is challenging due to the lack of centralised oversight and the complexity of machine-readable contracts.

Traditional automated negotiation systems [35], although effective in centralised environments, are ill-suited for decentralised settings due to their reliance on predefined trust assumptions and limited policy enforcement capabilities. Protocols such as the Take-It-Or-Leave-It (TILI)[34] approach and the Monotonic Concession Protocol (MCP)[17] fail to adequately address the complexity of multi-issue negotiations and dynamic trade-offs.

To address these challenges, we propose MINiDM (Multi-Issue Negotiation in Decentralised Data Marketplaces), a novel game-theoretic framework for autonomous, policy-aware data-sharing agreements in decentralized marketplaces. MINiDM integrates multi-agent negotiation models, utility-based game-theoretic decision-making, and semantic web ontologies (ODRL and DPV) to enable AI-driven agents to generate offers, counteroffers, and resolve conflicts. Crucially, rather than relying on blockchain-based smart contracts, MINiDM employs a multi-agent commitment protocol that leverages digital signatures and a tamper-evident commitment registry to enforce and verify agreements. This approach ensures that negotiated terms are both final and verifiable, while reducing computational overhead.

The key contributions of this paper are:

- A scalable AI-driven negotiation framework for decentralised data marketplaces that employs multi-agent systems for autonomous negotiation, utility-based decision-making for fairness and efficiency, and privacy-preserving techniques for secure, policy-aware transactions.
- The integration of Semantic Web Ontologies (ODRL, DPV) for automated regulatory compliance, enabling AI agents

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

to dynamically enforce data policies and encode regulatory terms in machine-readable contracts.

- A robust game-theoretic model based on Multi-Attribute Utility Theory, Bayesian game theory, and Coalition Game Theory to handle multi-issue negotiations.
- Experimental validation demonstrating higher agreement rates, better utility balance, and improved negotiation efficiency compared to traditional approaches.

The remainder of this paper is structured as follows: Section 2 presents the framework overview, Section 3 describes implementation and evaluation, related works are studied in Section 4 and finally Section 6 concludes with future research directions.

## 2 FRAMEWORK OVERVIEW

Our framework models negotiation as a sequential, multi-issue game, where agents iteratively propose and refine offers based on their preferences, constraints, and regulatory requirements. The negotiation process follows the IDSA Contract Negotiation Protocol (CNP)<sup>1</sup>, incorporating utility modelling, ontology-driven compliance enforcement, and automated agreement execution.

Unlike traditional blockchain-based smart contracts, which often introduce high computational costs and rigidity, MINiDM adopts a multi-agent commitment protocol that leverages digital signatures and a tamper-evident commitment registry. This approach ensures verifiable, enforceable, and dispute-resistant agreements while maintaining computational efficiency.

The key components of MINiDM are the following:

**Multi-Agent Negotiation Model** Implements autonomous decision-making, enabling data providers and consumers to evaluate trade-offs dynamically.

**Utility-Based Optimisation** Uses a multiplicative utility function to capture complex interdependencies between negotiation attributes.

**Ontology-Driven Policy Enforcement** Integrates ODRL and DPV for automated regulatory compliance.

**Multi-Agent Commitment Protocol** Leverages digital signatures without relying on blockchain-based smart contracts.

The subsequent subsections detail utility modeling, negotiation protocol design, ontology integration, and algorithm execution.

### 2.1 Negotiation as a Multi-Issue Game

The negotiation process in MINiDM is modeled as a sequential multi-issue game, where AI-driven agents iteratively generate offers, counteroffers, and trade-offs based on dynamic preferences, constraints, and legal requirements. The framework supports multi-issue bargaining across the following attributes:

- **Data subset (DB):** Defines whether the consumer requests the full dataset or a subset selection.
- **Actor (TP):** Specifies who will receive access to the data (e.g., researchers vs. commercial entities).
- **Duration (D):** The length of data access rights (e.g., one-time, subscription-based).

- **Purpose (PU):** The intended use of the data (e.g., academic research vs. targeted marketing).
- **Action (A):** Defines permissible actions (e.g., read, process, share).
- **Price (P):** Represents the compensation model, which can be fixed, dynamic, or performance-based.

These attributes can be sorted into three categories, based on the types of their values: *nominal attributes* such as free text attributes, *numerical attributes* (price, duration and data subset), and *hierarchical attributes* (actors, purposes and actions that come from an ontology). Numerical attributes can be easily normalised to ensure comparability. Though the data subset attribute does not neatly fall into this category, we will model it as a numerical value corresponding to the proportion of data requested to the full data. Hierarchical attributes are weighted based on their position in a predefined hierarchy. For example, purpose value “research and development” is assigned a larger weight than “commercial research”, because it is more general, and thus higher in the hierarchy of purposes. To capture attribute interdependencies, MINiDM employs a multiplicative utility function. This approach would allow the negotiation protocol to better account for cases where one attribute (like price) can dramatically affect the utility of other attributes (such as duration or data access).

$$U(X) = \prod_{i=1}^n f_i(x_i, w_i) \quad (1)$$

Equation (1) shows our utility function  $U(X)$ , where  $X = (DB, TP, D, PU, A, P)$  is a set of negotiation attributes;  $x_i$  is the normalised value of attribute  $i$ ; and  $w_i$  denotes the relative weight of attribute  $i$ . While one use arbitrary functions, in this paper we use  $f_i(x_i, w_i) = x_i^{w_i}$  to ensure non-linearity, making utility highly sensitive to critical negotiation factors.

This models realistic risk aversion, preventing parties from accepting suboptimal trade-offs where a single issue is overly compromised. Furthermore, this approach enables dynamic trade-offs, allowing, for instance, consumers to accept a higher price for stronger privacy guarantees, and providers to reduce costs in exchange for shorter data-sharing durations. The degree of preference overlap between negotiation parties directly affects convergence speed and agreement likelihood: high preference alignment (e.g., similar price and privacy expectations) results in faster agreements; and low overlap leads to longer negotiations and lower agreement rates due to conflicting priorities.

Finally, both the provider and consumer assign weights to the dataset’s attributes based on their priorities to reflect their preferences and evaluate the target data source in utility function. In this case, weights assigned by the provider reflect the privacy sensitivity of attributes, and weights assigned by the consumer reflect how the attributes are crucial to their goals; attributes that the provider is not willing to share can be tagged as non-negotiable.

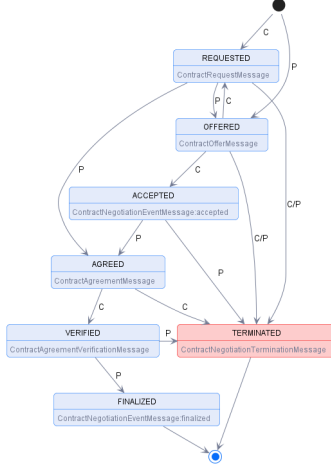
### 2.2 Negotiation Protocol

MINiDM follows the IDSA Contract Negotiation Protocol (CNP)<sup>2</sup> which is illustrated in Figure 1. Data providers list their datasets

<sup>1</sup><https://github.com/International-Data-Spaces-Association/ids-specification/>

<sup>2</sup><https://github.com/International-Data-Spaces-Association/ids-specification/blob/main/negotiation/contract.negotiation.protocol.md>

on decentralised data marketplaces, optionally specifying key attributes such as pricing, access conditions, and usage restrictions. When a consumer browses the marketplace and finds a dataset, they send a request to the provider.



**Figure 1: State Machine for the IDSA Contract Negotiation Protocol.**

**2.2.1 Negotiation Initialisation.** The data consumer’s agent initiates the negotiation by sending a request to the data provider, detailing their negotiation attributes, and containing any applicable legal or policy restrictions as regulatory constraints. Said request is structured using RDF to ensure machine-readable policy enforcement and compliance verification. Following this, the data provider receives the request and can evaluate it to either (1) agree the offer’s terms, ending the negotiation in a match, or (2) reject the offer and propose a counteroffer that is more suited to their goals by modifying any of the (negotiable) attributes.

**2.2.2 Iterative Adjustment.** Agents refine proposals to optimise their utility and get close to the opponent’s desired utility while ensuring regulatory compliance using Bounded Best-Response Dynamics (BBRD) [4]. BBRD is a negotiation strategy based on game theory, where players make counteroffers that: improve their utility compared to the opponent’s last offer; are worse than their previous offer (showing a controlled concession); and stay within their predefined preference bounds. This strategy ensures that the players gradually move toward an agreement while avoiding irrational concessions. This iterative process continues until an agreement is reached (contract generation) or the negotiation fails due to irreconcilable differences. Once both agents have reached a consensus, the agreement is formalised as an initial contract.

## 2.3 Ontology Integration

MINiDM integrates computational policies and semantic ontologies to enforce machine-readable, automated compliance with data-sharing regulations. The two key technologies are the Open Digital Rights Language (ODRL) and the Data Privacy Vocabulary (DPV)[43]. Together, they play a pivotal role in structuring negotiation terms, enabling precise alignment of the provider’s privacy

preferences with the consumer’s goals. ODRL allows for the definition of permissions, prohibitions or obligations on actions (e.g., use, aggregate, distribute, display), actors and assets, as well as define a hierarchical structure on said elements. On the other hand, DPV provides an ontology of terms related to the processing of data and relevant technologies, which includes a hierarchy of purposes to guide negotiation strategies. This layered ontology-driven approach ensures that agreements are semantically precise, reducing ambiguity and facilitating automation. Through semantic reasoning, the framework resolves conflicts that arise when a consumer requests actions or purposes outside the provider’s permissions. ODRL’s dependency structures allow for predefined conflict resolution strategies, such as prioritising permissions, prohibitions, or invalidating policies altogether. In our implementation we have used an extension of the formal semantics and rule-based reasoning framework introduced in [48] to handle conflict resolution.

In the ontology-driven conflict checking, AI agents automatically verify whether an offer adheres to predefined policies; if a conflict arises (e.g., an action is prohibited), the system suggests alternative offers applying semantic reasoning. Additionally, if multiple constraints exist, the system prioritises legal obligations over negotiable terms through a hierarchical prioritisation. This approach ensures regulatory adherence, interoperability across decentralised marketplaces, and reduced negotiation failure rates due to policy conflicts.

**2.3.1 Key GDPR Considerations in MINiDM.** To ensure compliance with the General Data Protection Regulation (GDPR), MINiDM has integrated privacy-preserving policies through ODRL and DPV ontologies. The key GDPR principles incorporated into the negotiation framework include:

- **Data Minimisation:** Data providers’ main strategy is to share subsets of the datasets; therefore, if a data consumer requests excessive information, the system enforces automatic rejection or counteroffers with reduced datasets.
- **Purpose Limitation:** Each dataset is tagged with an allowed usage purpose. If the data consumer attempts to negotiate for unauthorised usage (e.g., "marketing"), the negotiation process blocks the request.
- **Legal Basis Verification:** Every negotiation agreement must explicitly define a valid legal basis for data processing, such as user consent or contractual necessity. Any offer lacking a legal basis is automatically rejected. MINiDM ensures this by integrating an automated policy verification system that checks each offer against predefined legal basis categories within the ODRL-DPV framework. If an offer does not meet the necessary requirements, it is flagged for modification or rejected.
- **User Consent and Revocation:** Data providers can specify consent conditions that require consumers to support revocation mechanisms. Agreements must define procedures to withdraw consent and ensure compliance with data deletion requests.
- **Automated GDPR Compliance Checks:** MINiDM employs a structured compliance check mechanism to ensure that negotiations adhere to GDPR principles. This mechanism is implemented using a combination of rule-based

validation and constraint evaluation through ODRL-DPV ontologies.

By embedding GDPR compliance rules in the negotiation process, MINiDM ensures that agreements align with legal and ethical standards while maintaining fair and efficient negotiation outcomes.

## 2.4 Algorithm Design

In a non-Markovian process, at each iteration,  $t$ , the consumer's agent defines and resolves the conflicts between the last existing offer and the consumer's preferences and generates a request in order to minimise the distance to the provider's best offer in the previous rounds, from  $1$  to  $t-1$ . The provider's agent also evaluates the received request, resolves conflicts, and proposes a counteroffer with almost the same strategy to become closer to the best received request since the first iteration of negotiation. Figure 2 represents the automated negotiation architecture.




Figure 2: Automated negotiation architecture.

**2.4.1 Game Convergence Analysis.** Our negotiation process is modelled as a multi-agent, multi-issue game, where rational agents iteratively adjust their strategies to maximise their utility while ensuring regulatory compliance and mutual satisfaction. In this section, we establish:

- Convergence to a Nash Equilibrium (NE) – Ensuring that no agent has an incentive to unilaterally deviate from its strategy.
- Convergence to a Pareto-Optimal Solution – Ensuring that no further improvements can be made without reducing another agent's utility.

To formalise the convergence analysis, we define the strategic game model, establish the existence of Nash equilibria, and prove convergence to Pareto-optimality using first-order conditions and Zangwill's Theorem [53].

We define the negotiation game as a repeated multi-issue bargaining game between two agents: a data provider, and a data consumer. Each agent has a strategy space  $S_i$ , where  $s_i \in S_i$  represents an agent's offer in the negotiation. The joint strategy profile is given by:  $S = S_1 \times S_2$ .

As mentioned before, each agent has a utility function  $U_i : S \rightarrow \mathbb{R}$  defined as:

$$U_i(X) = \prod_{j=1}^n x_{ij}^{w_{ij}}, \quad (2)$$

where  $x_{ij}$  is the normalised value of attribute  $j$  for agent  $i$ , with  $0 \leq x_{ij} \leq 1$ .  $w_{ij}$  is the weight assigned to attribute  $j$  for agent  $i$ , satisfying:  $\sum_{j=1}^n w_{ij} = 1$ . The strategy update rule follows Bounded

Best-Response Dynamics (BBRD), meaning that each agent selects the best response within a feasible bound to optimise its utility at each step.

**Convergence to a Nash Equilibrium:** A Nash Equilibrium (NE) is a strategy profile  $(s_1^*, s_2^*)$  where no agent can improve its utility by unilaterally deviating:

$$U_1(s_1^*, s_2^*) \geq U_1(s_1, s_2^*) \forall s_1 \in S_1$$

$$U_2(s_1^*, s_2^*) \geq U_2(s_1^*, s_2) \forall s_2 \in S_2$$

By the Fan-Glicksberg Fixed-Point Theorem [23], a Nash Equilibrium exists if: the strategy space  $S$  is nonempty, compact, and convex; and the utility function  $U_i(X)$  is quasi-concave and upper-hemicontinuous.

Since each attribute is drawn from a convex space,  $x_{ij} \in [0, 1]$ , the strategy space  $S$  is convex.

The utility function  $U_i(X)$  is log-concave since:

$$\log U_i(X) = \sum_{j=1}^n w_{ij} \log x_{ij}, \quad (3)$$

which is a weighted sum of concave functions. Hence,  $U_i(X)$  is quasi-concave, ensuring the existence of an NE.

If an agent updates its offer  $X_i^{(t+1)}$  using a best-response function  $f$ ,  $X_i^{(t+1)} = f(X_i^t)$ . If there exists  $0 \leq \alpha < 1$  such that:

$$\|X_i^{(t+1)} - X_i^*\| \leq \alpha \|X_i^{(t)} - X_i^*\|, \quad (4)$$

Then by Banach's Fixed-Point Theorem [6, 13],  $X^t$  must converge to  $X^*$ . Since the negotiation space is bounded and updates are restricted,  $f$  exhibits contraction. Thus, BBRD ensures convergence to a Nash Equilibrium.

**Convergence to Pareto-Optimality:** A solution is Pareto-optimal if no agent can improve its utility without reducing the other agent's utility,  $\Delta U_1(X^*) = \lambda \Delta U_2(X^*)$  for some  $\lambda > 0$ .

Since  $U_i(X)$  is multiplicative,

$$\frac{\partial U_i}{\partial x_j} = w_j \prod_{k=1}^n x_k^{w_k - \delta_{jk}}. \quad (5)$$

Solving the Lagrange condition in Eq (5), we obtain:

$$\frac{w_1}{x_1^*} = \frac{w_2}{x_2^*} = \dots = \frac{w_n}{x_n^*}. \quad (6)$$

This condition ensures that no attribute can be improved without reducing another agent's utility. Now, we show that distance to the Pareto Frontier is monotonic. The distance of the current offer  $X^{(t)}$  to the Pareto frontier is  $d^{(t)}$  and is defined as:

$$d^{(t)} = \min_{X^* \in \mathcal{P}} \|X^{(t)} - X^*\|, \quad (7)$$

where  $\mathcal{P}$  represents the Pareto-optimal set. Each agent adjusts its proposal at round  $t + 1$  to maximise utility, therefore, each iteration reduces this distance, which implies:

$$d^{(t+1)} \leq d^{(t)}. \quad (8)$$

Since  $d^{(t)}$  is bounded below by zero, it converges. By Zangwill's Theorem [53], a sequence  $X^{(t)}$  converges to a fixed point if: (1)  $X^{(t)}$  is contained in a compact set; (2) the utility function is monotonically improving; (3) the decision rule ensures that the sequence

moves toward an optimal solution. Since all conditions hold, the final agreement lies on the Pareto frontier.

**Rate of Convergence to Pareto Optimality:** When preference alignment is high, convergence is fast due to rapid trade-offs. When preferences diverge significantly, convergence slows down due to iterative conflict resolution. The granularity of updates determines whether convergence is logarithmic or polynomial.

### 3 IMPLEMENTATION AND EVALUATION

To assess the performance of MINiDM, we developed an experimental platform simulating the negotiation process between data providers and consumers. The evaluation compares MINiDM against two established negotiation protocols: Take-It-Or-Leave-It (TILI), a rigid, one-shot negotiation strategy where the offer must be accepted or rejected without further interaction; and Monotonic Concession Protocol (MCP), a sequential approach where each party gradually concedes on negotiation attributes to reach an agreement. Since each concession is based only on the current negotiation state, standard MCP follows a Markovian process.

We conduct experiments using the *Bank Marketing Dataset* from the UCI repository<sup>3</sup>. The data is about direct marketing campaigns (phone calls) of a Portuguese banking institution [41]. Agents negotiate multiple attributes, including data subset, privacy constraints, and price.

#### 3.1 Experimental Setup

Experiments were conducted on a MacBook with an Apple M1 CPU and 8GB RAM. Each experiment was run for 100 iterations under varying conditions to evaluate the framework’s robustness. Table 1 summarises the experimental parameters.

**Table 1: Experimental Setup Parameters**

Parameter	Description	Values Tested
Number of Agents	Data provider and consumer	2
Utility Weights	Defined by main provider and consumer	Fixed
Privacy and Usage Control	Managed using ODRL and DPV ontologies	Enabled
GDPR Compliance	Enforced using ODRL-DPV extensions	Integrated
Negotiation Rounds	Maximum iterations	50, 100, 200

#### 3.2 Evaluation Metrics

We evaluate MINiDM based on the following key metrics:

- **Utility Values**
- **Fairness:** Fairness measures how balanced the final outcomes are between the provider and consumer:

$$\text{Fairness} = 1 - \left| \frac{U_{\text{provider}} - U_{\text{consumer}}}{\max(U)} \right| \quad (9)$$

- **Agreement Rate:** Percentage of negotiations that successfully reach an agreement.

$$\text{Agreement Rate} = \frac{\text{Number of Successful Negotiations}}{\text{Total Number of Simulations}} \times 100 \quad (10)$$

- **Time to Agreement:** Measures how execution time scales with preferences overlap and increasing negotiation attributes.

### 3.3 Results and Discussion

**3.3.1 Agreement Rate.** Figure 3 compares agreement rates under different preferences overlaps.

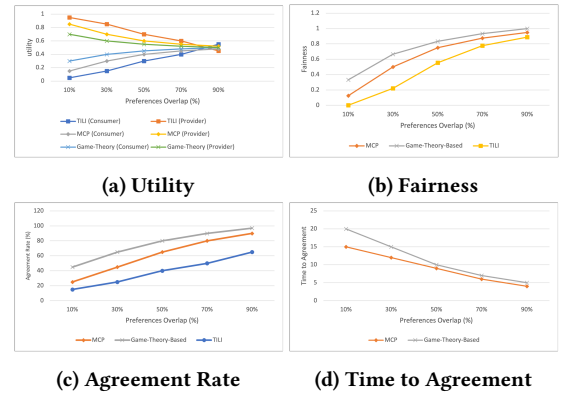
MINiDM achieves the highest agreement rates when preferences align (70%-90%), because the multiplicative utility function captures attribute dependencies, which become an advantage in scenarios with well-aligned preferences. However, agreement rates drop sharply in low-overlap scenarios because extreme interdependencies penalise offers with significant attribute mismatches.

MINiDM does dynamically reasoning and resolves some conflicts using hierarchical structures of attributes such as purposes and actions, and prioritisation rules. It uses dependency model, ODRL’s ConflictTerm, to determine whether permission, prohibition, or invalidation takes precedence. This capability is integral to its high agreement rates. Even in low-overlap scenarios, the protocol iteratively resolves discrepancies by offering alternative options (e.g., less critical purposes or actions), ensuring continued negotiation even when initial requests are misaligned.

TILI shows the lowest agreement rates, particularly in low-to mid-overlap scenarios (10%-50%), because it is rigid and non-iterative, meaning that agreements are only possible if the initial offer happens to align with the opponent’s preferences. However, its simplicity and speed make it suitable for high-overlap scenarios or where time to agreement is critical.

MCP falls between the other two protocols in agreement rates. It handles mid- to high-overlap scenarios relatively well (50%-90% overlap), since it facilitates iterative negotiations, allowing for gradual convergence. However, it lacks the nuanced adaptability of game-theory-based frameworks with advanced utility models.

**3.3.2 Utility Balance and Fairness.** Figures 3a shows the utility for both the consumer and provider, and 3b illustrates the fairness across protocols. MINiDM maintains superior fairness by dynamically adjusting offers, whereas TILI heavily favours the provider.



**Figure 3: Comparisons of Evaluation Metrics over Preferences Overlap Percentage.**

MINiDM offers the best utility balance as overlap increases, approaching a near-equal distribution (50/50) at 90% overlap, since the multiplicative utility function enforces strong inter-dependencies between attributes, encouraging both agents to make balanced concessions to reach agreement. However, in low-overlap scenarios,

<sup>3</sup><https://archive.ics.uci.edu/dataset/222/bank+marketing>

the provider often gains a disproportionate share due to stricter constraints in the utility model.

TILI shows the least balanced outcomes, heavily favoring the provider, because the consumer has limited negotiation leverage under this rigid framework, often accepting suboptimal offers or walking away. MCP falls between the other two protocols, balancing flexibility and fairness. As overlap increases, MCP gradually aligns the utilities but rarely achieves complete parity.

**3.3.3 Computational Complexity and Scalability.** Figure 3d shows the number of iteration requires to reach an agreement, for both MINiDM and MCP; TILI is not compared because it is a one-step negotiation protocol. As the results show, The MCP converges faster than MINiDM due to its streamlined and predictable concession-based approach, which narrows the negotiation space in a stepwise manner. MCP adjusts one or few attributes monotonically, avoiding iterative conflict resolution and complex utility recalculations, leading to linear and quicker convergence. However, this simplicity comes at the expense of reduced flexibility and fairness, making MCP less suitable for resolving conflicts in low-overlap scenarios or achieving balanced multi-issue agreements. In contrast, MINiDM’s advanced adaptability and conflict resolution extend negotiation time but ensure more equitable and optimal outcomes.

MINiDM’s non-linear complexity is justified by superior agreement rates and fairness. However, its computational complexity grows non-linearly as the number of attributes increases, primarily due to conflict resolution and utility recalculations for each added attribute. Since TILI sets a single, fixed offer, the number of attributes has minimal impact on its execution time. This makes it the fastest protocol in terms of scalability. However, the simplicity of TILI results in poor utility balance and low agreement rates in scenarios with many attributes, as the fixed offer may not align well with consumer preferences. Finally, MCP can handle multi-issue negotiations moderately well, making incremental concessions across attributes to narrow the negotiation space. It becomes increasingly slow as the number of attributes grows, due to the stepwise adjustment of each attribute in sequence.

## 4 RELATED WORKS

In this section we briefly review related works from to perspective: first, data sharing and its major concerns and then, automated negotiation and different approaches to it.

### 4.1 Data Sharing

Given the importance of data sharing as indicated by large research initiatives such as Gaia-X [30] and Catena-X<sup>4</sup> [16], lots of research has been doing on data sharing for decades; their focus has mainly been on privacy policies [27] and usage control [42, 44, 46, 50] in various dimensions and applications [9, 22, 37, 45, 47]. Some researches have also studied data sharing agreements from legal and technical perspective [11, 14, 39, 49]. In this study we study data sharing while considering privacy, usage control, and monetisation.

<sup>4</sup><https://catena-x.net/en/>

### 4.2 Automated Negotiation

Automated negotiation has been one of the research areas of interests for several decades, and a variety of negotiation frameworks have been proposed [5, 18–20, 25, 26, 29, 32, 38, 52] provided brief state of the art overviews of agent based automated negotiation. Automated negotiation, also powered by game theory and artificial intelligence, [15, 21]. Game theory is a sub-field of microeconomics which is widely acknowledged to provide the best available set of tools for the design of multi-agent architectures [8]. Game theory basically provides mathematical framework for analysing strategic decision-making, is particularly suited to situations where the actions of one participant affect the outcomes of others [31, 36]. By modeling data-sharing negotiations as sequential games, interactions between entities can be analyzed to predict behavior, identify potential conflicts, and design mechanisms that promote cooperation. The Vickrey-Clarke-Groves (VCG) mechanism in [12] encourages truthful information sharing, reducing asymmetry and facilitating optimal agreements. Game-theory-based models also adapt to evolving conditions, making them suitable for dynamic environments like smart grids, healthcare systems [51], and water transfer [33]. Our framework models the negotiation as a two-players game in which agents negotiate based on the strategies and preferences that are defined by the data provider and consumer to share data. When agents reach an agreement, it will be sent for the provider and consumer and they will finalise it as a contract.

## 5 DISCUSSION AND CONCLUSION

This paper introduced MINiDM, a framework for multi-issue negotiation in decentralised data marketplaces. By integrating formal policy languages and vocabularies (ODRL, DPV), AI agents, and a lightweight commitment protocol, MINiDM facilitates fair, GDPR-compliant negotiations without relying on blockchain or central authorities.

Our experiments show that MINiDM enables high agreement rates, improves negotiation fairness, and reduces delays and outperforms the TILI approach and MCP, which serve as the benchmarks for negotiation scenarios. The framework demonstrates that agent-based negotiation, coupled with policy reasoning and compliance checking, can enhance trust and autonomy in Web3 data trading.

Future work will focus on extending compliance logic, incorporating learning-based strategies, and exploring negotiation dynamics with incomplete or asymmetric information.

## ACKNOWLEDGMENTS

This research was partially funded by the Horizon Europe projects DATAPACT (No. 101189771), RAISE (No. 101058479) and UPGAST (No. 101093216).

## REFERENCES

- [1] 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). OJ L 119/1, 4/5/16.
- [2] 2022. Regulation (EU) 2022/868 of the European Parliament and of the Council of 30 May 2022 on European Data Governance and amending Regulation (EU) 2018/1724 (Data Governance Act). OJ L 152/1, 3/6/22.
- [3] 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and



- amending various Regulations and Directives (Artificial Intelligence Act). <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>. Official Journal of the European Union, L, Vol. 2024/1689.
- [4] Baruch Awerbuch, Yossi Azar, and Rohit Khandekar. 2008. Fast load balancing via bounded best response. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*. 314–322.
  - [5] Reyhan Aydoğan, David Festen, Koen V Hindriks, and Catholijn M Jonker. 2017. Alternating offers protocols for multilateral negotiation. *Modern approaches to agent-based complex automated negotiation* (2017), 153–167.
  - [6] Stefan Banach. 1922. Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta mathematicae* 3, 1 (1922), 133–181.
  - [7] Christian Bieri. 2021. An overview into the InterPlanetary File System (IPFS): use cases, advantages, and drawbacks. *Communication Systems XIV* 28 (2021).
  - [8] Ken Binmore and Nir Vulkan. 1999. Applying game theory to automated negotiation. *Netnomics* 1 (1999), 1–9.
  - [9] Luca Bonomi, Yingxiang Huang, and Lucila Ohno-Machado. 2020. Privacy challenges and research opportunities for genomic data sharing. *Nature genetics* 52, 7 (2020), 646–654.
  - [10] Arnaud Braud, Gaël Fromentoux, Benoit Radier, and Olivier Le Grand. 2021. The road to European digital sovereignty with Gaia-X and IDSA. *IEEE network* 35, 2 (2021), 4–5.
  - [11] Claudio Caimi, Carmela Gambardella, Mirko Manea, Marinella Petrocchi, and Debora Stella. 2015. Legal and technical perspectives in data sharing agreements definition. In *Annual Privacy Forum*. Springer, 178–192.
  - [12] M Caminati, M Kerber, C Lange, and C Rowat. 2015. Vickrey-clark-groves (vcg) auctions. *College Inf. Sci. Technol., Pennsylvania State Univ., State College, PA, USA* (2015).
  - [13] Krzysztof Ciesielski. 2007. On Stefan Banach and some of his results. *Banach Journal of Mathematical Analysis* 1, 1 (2007), 1–10.
  - [14] Gianpiero Costantino, Fabio Martinelli, Ilaria Matteucci, and Marinella Petrocchi. 2017. Analysis of data sharing agreements. In *International Conference on Information Systems Security and Privacy*, Vol. 2. SCITEPRESS, 167–178.
  - [15] Tinglong Dai, Katia Sycara, and Ronghuo Zheng. 2021. Agent reasoning in AI-powered negotiation. *Handbook of Group Decision and Negotiation* (2021), 1187–1211.
  - [16] Leon Ehlers. 2023. *The impact of Catena-X on digital business models in the automotive industry*. Ph.D. Dissertation.
  - [17] Ulle Endriss. 2006. Monotonic concession protocols for multilateral negotiation. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. 392–399.
  - [18] Shaheen Fatima, Sarit Kraus, and Michael Wooldridge. 2014. *Principles of automated negotiation*. Cambridge University Press.
  - [19] Katsuhide Fujita, Quan Bai, Takayuki Ito, Minjie Zhang, Fenghui Ren, Reyhan Aydoğan, and Rafik Hadfi. 2017. *Modern approaches to agent-based complex automated negotiation*. Springer.
  - [20] Naoki Fukuta, Takayuki Ito, Minjie Zhang, Katsuhide Fujita, and Valentin Robu. 2016. *Recent advances in agent-based complex automated negotiation*. Vol. 638. Springer.
  - [21] Enrico H Gerding, David Dominique Bernard van Bragt, and Johannes Antonius La Poutre. 2000. *Scientific approaches and techniques for negotiation: a game theoretic and artificial intelligence perspective*. Citeseer.
  - [22] Jin Gerlach, Thomas Widjaja, and Peter Buxmann. 2015. Handle with care: How online social network providers’ privacy policies impact users’ information sharing behavior. *The Journal of Strategic Information Systems* 24, 1 (2015), 33–43.
  - [23] Irving L Glicksberg. 1952. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proc. Amer. Math. Soc.* 3, 1 (1952), 170–174.
  - [24] Renato Iannella. 2007. Open digital rights language (ODRL). *Open Content Licensing: Cultivating the Creative Commons* (2007).
  - [25] Takayuki Ito. 2012. *New trends in agent-based complex automated negotiations*/[edited by] Takayuki Ito...[et al.]. Springer.
  - [26] Takayuki Ito, Minjie Zhang, Valentin Robu, and Tokuro Matsuo. 2013. *Complex automated negotiations: Theories, models, and software competitions*. Springer.
  - [27] Johnson Iyilade and Julita Vassileva. 2014. P2U: a privacy policy specification language for secondary data sharing and usage. In *2014 IEEE Security and Privacy Workshops*. IEEE, 18–22.
  - [28] Harshvardhan J. Pandit, Beatriz Esteves, Georg P. Krog, Paul Ryan, Delaram Golpayegani, and Julian Flake. 2024. Data Privacy Vocabulary (DPV)–Version 2.0. In *International Semantic Web Conference*. Springer, 171–193.
  - [29] Nicholas R Jennings, Peyman Faratin, Alessio R Lomuscio, Simon Parsons, Carlos Sierra, and Michael Wooldridge. 2001. Automated negotiation: prospects, methods and challenges. *International Journal of Group Decision and Negotiation* 10, 2 (2001), 199–215.
  - [30] Ilka Jussen, Frederik Möller, Julia Schweihoff, Anna Gieß, Giulia Giussani, and Boris Otto. 2024. Issues in inter-organizational data sharing: Findings from practice and research challenges. *Data & Knowledge Engineering* (2024), 102280.
  - [31] Anthony Kelly. 2003. *Decision making using game theory: an introduction for managers*. Cambridge University Press.
  - [32] Liu Kexing. 2011. A survey of agent based automated negotiation. In *2011 International Conference on Network Computing and Information Security*, Vol. 2. IEEE, 24–27.
  - [33] Chenming Li, Wei Huang, Zhenli Ma, Fengzhou Wang, Xiaodong Wei, and Lizhong Xu. 2014. An automated negotiation protocol based on game theory. *Computational Water, Energy, and Environmental Engineering* 3, 3 (2014), 111–118.
  - [34] Katrina Ligett and Aaron Roth. 2012. Take it or leave it: Running a survey when privacy comes at a cost. In *International workshop on internet and network economics*. Springer, 378–391.
  - [35] Xudong Luo, Yanling Li, Qiaojuan Huang, and Jieyu Zhan. 2024. A survey of automated negotiation: Human factor, learning, and application. *Computer Science Review* 54 (2024), 100683.
  - [36] Kaveh Madani and Jay R Lund. 2011. A Monte-Carlo game theoretic approach for multi-criteria decision making under uncertainty. *Advances in water resources* 34, 5 (2011), 607–616.
  - [37] Bradley Malin, David Karp, and Richard H Scheuermann. 2010. Technical and policy approaches to balancing patient privacy and data sharing in clinical and translational research. *Journal of Investigative Medicine* 58, 1 (2010), 11–18.
  - [38] Ivan Marsa-Maestre, Miguel A Lopez-Carmona, Takayuki Ito, Minjie Zhang, Quan Bai, and Katsuhide Fujita. 2014. *Novel insights in agent-based complex automated negotiation*. Vol. 535. Springer.
  - [39] Ilaria Matteucci, Marinella Petrocchi, Marco Luca Sbodio, and Luca Wiegand. 2011. A design phase for data sharing agreements. In *International Workshop on Data Privacy Management*. Springer, 25–41.
  - [40] Trent McConaghy. 2022. Ocean Protocol: tools for the Web3 data economy. In *Handbook on Blockchain*. Springer, 505–539.
  - [41] Sérgio Moro, Paulo Cortez, and Paulo Rita. 2014. A data-driven approach to predict the success of bank telemarketing. *Decision Support Systems* 62 (2014), 22–31.
  - [42] Andres Munoz-Arcenales, Sonsoles López-Pernas, Alejandro Pozo, Álvaro Alonso, Joaquín Salvachúa, and Gabriel Huecas. 2019. An architecture for providing data usage and access control in data sharing ecosystems. *Procedia Computer Science* 160 (2019), 590–597.
  - [43] Harshvardhan J Pandit and Beatriz Esteves. 2024. Enhancing Data Use Ontology (DUO) for health-data sharing by extending it with ODRL and DPV. *Semantic Web* 15, 4 (2024), 1473–1498.
  - [44] Jaehong Park and Ravi Sandhu. 2004. The UCONABC usage control model. *ACM transactions on information and system security (TISSEC)* 7, 1 (2004), 128–174.
  - [45] Thomas FJ-M Pasquier, Jatinder Singh, David Eysers, and Jean Bacon. 2015. CamFlow: Managed data-sharing for cloud services. *IEEE Transactions on Cloud Computing* 5, 3 (2015), 472–484.
  - [46] Alexander Pretschner, Manuel Hilty, and David Basin. 2006. Distributed usage control. *Commun. ACM* 49, 9 (2006), 39–44.
  - [47] Na Shi, Bo Tang, Ravi Sandhu, and Qi Li. 2021. DUCE: Distributed usage control enforcement for private data sharing in Internet of Things. In *Data and Applications Security and Privacy XXXV: 35th Annual IFIP WG 11.3 Conference, DBSec 2021, Calgary, Canada, July 19–20, 2021, Proceedings* 35. Springer, 278–290.
  - [48] Simon Steyskal and Axel Polleres. 2015. Towards formal semantics for ODRL policies. In *International Symposium on Rules and Rule Markup Languages for the Semantic Web*. Springer, 360–375.
  - [49] V Swamp, Len Seligman, and Arnon Rosenthal. 2006. Specifying data sharing agreements. In *Seventh IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY’06)*. IEEE, 4–pp.
  - [50] Carol Tenopir, Natalie M Rice, Suzie Allard, Lynn Baird, Josh Borycz, Lisa Christian, Bruce Grant, Robert Olendorf, and Robert J Sandusky. 2020. Data sharing, management, use, and reuse: Practices and perceptions of scientists worldwide. *PLoS one* 15, 3 (2020), e0229003.
  - [51] Wayes Tushar, Chau Yuen, Tapan K Saha, Sohrab Nizami, Mollah Rezaul Alam, David B Smith, and H Vincent Poor. 2023. A survey of cyber-physical systems from a game-theoretic perspective. *IEEE Access* 11 (2023), 9799–9834.
  - [52] Sam Vente, Angelika Kimmig, Alun Preece, and Federico Cerutti. 2020. The current state of automated negotiation theory: a literature review. *arXiv preprint arXiv:2004.02614* (2020).
  - [53] Willard I Zangwill. 1969. Convergence conditions for nonlinear programming algorithms. *Management Science* 16, 1 (1969), 1–13.