Sharing Personal Data via Incentive-based Negotiation
Preference Modeling and Empirical Analysis

Kuru, Ahmet; Aydogan, Reyhan; Ozturk, Pinar; Razeghi, Yousef

DOI
10.1145/3770751

Publication date
2025

Document Version
Final published version

Published in
ACM Transactions on Internet Technology

Sharing Personal Data via Incentive-based Negotiation: Preference Modeling and Empirical Analysis

AHMET KURU, Computer Science, Ozyegin Universitesi, Istanbul, Turkey
REYHAN AYDOGAN, Artificial Intelligence and Data Engineering, Ozyegin Universitesi, Istanbul, Turkey and Intelligent Systems, Delft University of Technology, Delft, Netherlands
PINAR OZTURK, Computer Science, Norwegian University of Science and Technology, Trondheim, Norway
YOUSEF RAZEGHI, Computer Science, Ozyegin Universitesi, Istanbul, Turkey

In an age where data is a pivotal asset for businesses, the ethical acquisition and use of personal information has become increasingly more significant. Empowering data providers with greater autonomy over their personal data is more important than ever. To address this, we propose a novel negotiation-based information-sharing framework that empowers individuals to actively negotiate the terms of their data sharing, addressing privacy concerns and ethical data usage. The framework enables users to determine what personal information they share and under what conditions, fostering a more balanced and transparent data exchange process. Our system allows data consumer agents to negotiate with their human users and can operate fully automatically, with agents representing data providers negotiating based on elicited preferences and needs. We propose novel preference modeling approaches and a negotiation framework to facilitate the bilateral sharing of information and incentives between data consumers and providers. User experiments demonstrate the efficacy of our negotiation approach and the effectiveness of the proposed preference models. Empirical results validate the benefits of the proposed framework.

CCS Concepts: • **Human-centered computing** → *Human computer interaction (HCI)*; • **Security and privacy** → **Usability in security and privacy;**

Additional Key Words and Phrases: Preference elicitation, automated negotiation

ACM Reference Format:
Ahmet Kuru, Reyhan Aydogan, Pinar Ozturk, and Yousef Razeghi. 2025. Sharing Personal Data via Incentive-based Negotiation: Preference Modeling and Empirical Analysis. *ACM Trans. Internet Technol.* 25, 4, Article 25 (November 2025), 26 pages. https://doi.org/10.1145/3770751

1 Introduction

Internet of things, e-services, and social media constitute an excellent resource of data for understanding the consumer behaviour. Such data has already been a major factor for companies

Authors' Contact Information: Ahmet Kuru (corresponding author), Computer Science, Ozyegin Universitesi, Istanbul, Turkey; e-mail: emre.kuru@ozu.edu.tr; Reyhan Aydogan, Artificial Intelligence and Data Engineering, Ozyegin Universitesi, Istanbul, Istanbul, Turkey, Intelligent Systems and Delft University of Technology, Delft, ZH, Netherlands; e-mail: reyhan.aydogan@ozyegin.edu.tr; Pinar Ozturk, Computer Science, Norwegian University of Science and Technology, Trondheim, Trndelag, Norway; e-mail: pinar@ntnu.no; Yousef Razeghi, Computer Science, Ozyegin Universitesi, Istanbul, Istanbul, Turkey; e-mail: yousef.razeghi66@gmail.com.

ACM 1533-5399/2025/11-ART25
https://doi.org/10.1145/3770751

to transform their business and achieve significant competitive advantage. As a consequence of companies' current practice of using customers' data for developing tailored and innovative services, people started to view their own personal information as a commodity that can be exchanged, either in return of specific services or even for money. However, consumer behaviour research reports that people are reluctant to share their personal data due to fear of possible illegal and unethical usage of their data and information [6, 9]. GDPR has also been effective in increasing people's awareness and concerns about sharing data/info. Putting these together, companies need personal data, people have data, and there is a need to regulate and facilitate how people can control the terms and conditions of sharing their data requested by the companies. The problem with marketing ethics has been studied for a long time [20] and some duties for the marketers were recognized such as justice, beneficence, noninjury, and so on [15].

We extend our data market framework [3] that revolves around an automated negotiation process between the data consumer and the data provider which gives the data provider an active role and power in deciding the terms of the data sharing, and facilitates the preservation of privacy in compliance with the individual data provider's privacy preferences. In this article, we propose a framework oriented around incentives that goes beyond mere discussion of the data content being shared. It also considers the terms under which the offer is made, specifically the selection of parties to share the data with, the time period for sharing, and the incentives that the data consumer should offer (e.g., usage is limited to 1 year, or sharing with third parties is not allowed). This approach ensures that the focus is not solely on the data itself but also on the broader context of its exchange.

In order to operate effectively in this negotiation framework, users should be able to express their privacy preferences in a quantifiable and structured way. This requirement introduces a significant challenge, as privacy preferences are often implicit, subjective, and difficult to articulate. One of the core motivations behind our work is to help users reflect on these concerns and better understand what kinds of data they value, and under which conditions they are willing to share them. To address this, we designed a preference elicitation process that not only captures user preferences effectively but also facilitates self-assessment and awareness during the process. On this end, we extend our preference model from [3] to more comprehensively capture individuals' sensitivity and their stance regarding various categories of personal data they generate, as well as their valuation of this information. Moreover, in our former work, preferential interdependency among data types were not considered for simplicity. However, sharing a particular data type may not seem very critical, but sharing another data type along with the other may be risky. That is, sharing only your neighborhood is not as risky as sharing both your neighborhood and email addresses. Therefore, instead of considering the risk of each single data type separately and aggregating their overall risk, it is desired to consider the interrelated data types as a single bundle and estimate the risk over bundles to capture the underlying preferential interdependency. Accordingly, we introduce a novel model to estimate user privacy violations.

Furthermore, this proposed framework has been assessed through human-agent experiments where the companies are represented as autonomous agents in the negotiation system. The proposed approach could also be used in an automated manner. In order to fully automate this process, the agents would need to elicit the preferences of the data providers to represent them. In order to evaluate our proposed approach, we conducted an extensive user experiment where the companies were represented as autonomous agents, and the users negotiated for themselves. We first evaluate our preference model by asking users to rate the offers they receive and give. The results demonstrate a close correlation between the preferences elicited from individuals and the given ratings. Moreover, we evaluate our proposed framework by considering the agreement rate, the acquired utilities for both parties and the number of interactions needed to reach an agreement. The results confirmed our framework's effectiveness.

In the following sections, we overview the related studies regarding privacy negotiation frameworks in Section 2. We provide a detailed background on Automated Negotiations in Section 3, particularly the underlying negotiation protocol, bidding strategy, and acceptance we utilized in our framework. Following this, Section 4 introduces our proposed framework, detailing its design and functionalities. We then present our Preference Elicitation Tool in Section 5.2. Section 6 discusses our experimental setup and analyzes the results, providing evidence of our framework's effectiveness. Lastly, we conclude with our future research direction in Section 7.
The primary contributions of this study are:

(1) **Preference Model**: We developed a computational model of preferences capable of detecting interdependent privacy risks in data, which enhances the management and understanding of privacy complexities.
(2) **Negotiation Framework**: We introduced a negotiation framework that facilitates interactions between data providers and consumers, improving outcomes for both parties.
(3) **Reasoning Mechanisms**: We implemented sophisticated reasoning mechanisms for both data providers and data consumers, enabling more informed decision-making during negotiations.
(4) **Comparative User Study**: We conducted a comprehensive user study to compare the efficacy of our negotiation-based approach with that of a baseline method, providing empirical evidence of its benefits. Furthermore, we analyzed the results elaborately from different perspectives, such as privacy awareness, personality, and negotiation attitudes.

2 Related Work

Numerous studies have explored the concerns of service consumers regarding the privacy of their personal data, which must be shared with service providers for processes like invoicing and shipment [12, 23, 24]. These studies highlight the limitations of the traditional approaches adopted by service providers, characterized as "take it or leave it"(where consumers must provide their personal data to access services) and "one-size-fits-all" (where all consumers are treated uniformly without considering the varying degrees of concern for their personal information). Such practices are shown to adversely affect user satisfaction. Consequently, these works advocate for a more adaptable strategy that incorporates privacy negotiation.

In this context, El-Khatib presents a privacy negotiation protocol where the data consumer and data provider negotiate their privacy policy [12]. In this protocol, the data consumer initiates the negotiation with an offer. The data provider can either accept this offer or reject it. In the case of rejection the provider will also provide an explanation regarding rejection. In this framework offers contain how the data provider's information will be used with an accompanying discount as an incentive. This pioneering work established the foundation for subsequent studies on privacy negotiation approaches. However, it operates under the assumption that the consumer agent has access to user preferences and employs a rule-based algorithm for decision-making. While this work introduces the concept of privacy policy negotiation, it does not proceed to evaluate the proposed system's effectiveness. Our research builds upon this foundation, advocating for negotiation while emphasizing the modeling of preferences and the reasoning processes of agents based on these preferences. Additionally, we conduct an evaluation of our system through experiments involving negotiation between users and agents.

Moreover, Preibusch frames this process as a dynamic back-and-forth discussion [24]. In that study, data providers are categorized into four types based on their data privacy concerns: those who are extremely cautious about any data use, those sensitive only to their financial and health data, those concerned exclusively with personal data (such as social security numbers, email addresses,

and bank account details), and those indifferent to data sharing. The service provider engages in negotiations with a service consumer only if the consumer does not fall into the first category, implying negotiations occur only when there is a potential for compromise. Unlike previous models, the negotiation process here is not concluded upon the service consumer's acceptance of an offer; it persists until either party decides to terminate the discussions. While both studies advocate for a more adaptable approach to privacy policy than traditional methods, they portray the service consumer as less empowered compared with the service provider, with consumers limited to either accepting or rejecting offers. In contrast, our framework equips service consumers with equal bargaining power, allowing for a more balanced negotiation process.

Furthermore, Baarslag et al. propose a negotiation mechanism for managing permissions, wherein the service consumer (i.e., the data provider) initiates the negotiation by making a partial offer concerning the data to be shared, and then invites the service provider to complete this offer with additional terms, such as a price discount [5]. In the proposed framework, requesting the service provider to complete the offer incurs a cost to the service consumer, thus preventing the service provider from disclosing its entire cost structure and precluding the agent from discovering all possible offer combinations. Unlike earlier studies, this model grants the service consumer greater bargaining power, though not equivalent to that of the service provider, as the latter always makes complete offers. Consequently, the service consumer is restricted from seeking further discounts or incentives and may hesitate to make more partial offers due to the associated costs. Our research diverges from Baarslag et al.'s study by addressing the limitation that service consumers can only make partial offers, a constraint that undermines the parity between service providers and consumers. We employ a different approach to evaluate complete offers, taking into consideration the potential interdependencies between various negotiation issues, contrary to the assumption in Baarslag et al.'s study that preferences are represented using additive utility functions without preferential interdependencies. However, we believe that such interdependencies may exist. For instance, the assessment of data usage could vary based on the type of data. Therefore, our work takes into account several factors to evaluate a bid, including the secrecy level of the information, the risk associated with sharing it, and the profit derived from the offered incentives. While our focus is primarily on the type of information and the incentives offered, Baarslag et al. also explore additional dimensions, such as the intended use of the data. This opens an avenue for future research to further explore the complexities of negotiation in permission management, including the nuanced interplay between data usage and other relevant factors. For example, the evaluation of data usage may differ depending on the type of data. Consequently, our work considers multiple criterion's when assessing a bid, including the data's privacy level, the possible risks involved if this information becomes public, and the benefits gained from the provided incentives. While our focus is mainly on the type of information and the incentives offered, Baarslag et al. examine additional aspects, such as the intended purpose of the data. This highlights an opportunity for future research to delve deeper into the complexities of negotiation in permission management, including the subtle interactions between data usage and other pertinent factors.

Recent research aims to automate the negotiation process to minimize human effort. Filipczuk et al. introduce a multi-issue negotiation model, where agents exchange partial and complete offers, negotiating over various bundled issues [14]. In this approach, the proposer (i.e., a data provider in our context) submits partial offers with some undetermined issues. The responder then completes the offer based on the partial proposal. Subsequently, the proposer can either accept or decline the complete offer. This system allows for human participants to provide feedback during the negotiation process, which is utilized solely for the purpose of eliciting user preferences without directly influencing the content of the offers. This method effectively reduces the cognitive burden on human users. On the other hand, we advocate that human users should retain the ultimate

authority to make decisions a critical aspect given the sensitive nature of the negotiations involved. Following, another study presents a context-aware agent-based framework that derives contexts from previous user experiences to model the current situation and makes privacy decisions on behalf of the user, even if the user has not explicitly been in a similar situation before [19]. In this framework, the agent can decide to accept, reject, or consult the user if the current context is ambiguous and thus requires the human to make the final decision. Furthermore, Kekulluoglu et al. introduce a methodology for a multi-agent management system for privacy in online social networks [17]. Each user is represented by an agent that aids in preserving their privacy through a hybrid negotiation architecture. This architecture integrates semantic representations of privacy rules and domains with agent-based decision-making using utility functions.

Another utilization of automated negotiations is big-data sharing, amplifying the risk of re-identification and privacy breaches, especially in the domains with diverse data sources ranging from IoT devices and social media with varying sensitivities and risks. Recent works have addressed these challenges by proposing frameworks that integrate privacy-preserving mechanisms with scalable negotiation strategies. For instance, Jung et al. tackle the challenge of balancing privacy risks with economic incentives in data sharing [16]. Their framework employs differential privacy to mediate negotiations, allowing data providers to adjust privacy parameters in exchange for compensation dynamically. This addresses the need for individualized privacy solutions in large-scale data markets, ensuring both privacy and fairness. Similarly, Ferradi et al. focus on the scalability and diversity of big data operations [13]. They provide a comprehensive survey of security and privacy-preserving techniques, emphasizing the importance of frameworks capable of managing heterogeneous data sources and multi-party negotiations. By highlighting the complexities of big data sharing, such as the interplay between privacy policies across different domains, their work underscores the necessity for adaptable negotiation mechanisms that address these interdependencies.

While existing methods show promise in terms of accuracy, they lack the capability to elucidate the rationale behind specific decisions to end-users. In response, Aycı et al. have introduced a methodology designed to generate explanations for the decision-making processes of privacy assistants [1]. Their method involves associating items with various topics to understand and communicate why an item is deemed private or public. This process employs topic modeling and machine learning classification techniques to identify and improve the topic descriptions. The critical aspect of their approach is using these identified topics to create intuitive, user-friendly explanations through textual or visual means, thereby enhancing the interpretability and transparency of the system's privacy decisions. There are also other works focusing on agreement over co-owned data's sharing policies amongst peers [11, 22] rather than negotiation between data consumers and providers. These works are complementary to our work. After the peers have reached an agreement, our framework could be used to negotiate with data consumers according to the groups agreed upon privacy policies.

Another crucial avenue of complementary research focuses on ensuring that data, once shared, is not subsequently re-shared with unauthorized third parties. This concern is especially critical in real-world deployments where data usage is difficult to monitor post-negotiation. Recent approaches have addressed this limitation by introducing enforceable data-sharing mechanisms. Urovi et al. propose LUCE, a blockchain-based platform for license accountability and compliance, which offers transparency and auditability regarding how shared data is used [27]. Similarly, Desai et al. explore the use of smart contracts to automatically enforce sharing agreements and penalize violations, thereby restricting data consumers from re-distributing data [21]. Furthermore, Braghin et al. present DLPFS, a framework for data-leakage prevention that employs access control and secure file handling mechanisms [7]. These systems are highly complementary to our work and could be integrated into our negotiation framework to provide end-to-end guarantees for data governance.

3 Agent-Based Negotiation

Automated negotiations structure the negotiation process between two or more agents over a finite set of n issues $\mathcal{I} = \{1, 2, \ldots, n\}$. Each issue $i \in \mathcal{I}$ has a range $\mathcal{D}_i$ of possible values. An outcome, $o \in \Omega$, is a complete assignment of values to all issues where Ω is the Cartesian Product of the values for each issue i in their $\mathcal{D}_i$. The set of all possible outcomes is defined as $\Omega = \mathcal{D}_1$ X $\mathcal{D}_2$ X... X $\mathcal{D}_n$.

The outcomes serve as offers, which agents evaluate using their utility functions that map each negotiation outcome to a real number in the range [0, 1], indicating the desirability of that outcome. Thus, a utility function mathematically represents the agent's preferences. Additive utility functions are the most commonly used utility functions in automated negotiations. Equation (1) illustrates such a utility function, where w_i represents the weight or importance of issue I_i for the agent, o_i represents the value for issue i from $\mathcal{D}i$ in the offer o, and Vi is the valuation function for issue i, which shows the desirability of the issue value o_i. An issue i's value o_i is more preferred when V_i returns a higher value. Each negotiating agent has its own utility function and uses it to determine what to offer and when to accept an offer. It is most commonly assumed that $\sum_{i \in n} w_i = 1$ and the domain of V_i is (0, 1) for any i. Agents generally do not know their opponent's preferences or *bidding strategies*.

$$\mathcal{U}(o) = \sum_{i=1}^{n} w_i \times V_i(o_i). \tag{1}$$

3.1 Negotiation Protocol

The ***Stacked Alternating Offers Protocol*** (**SAOP**) is one of the most commonly used turn-taking protocols, providing rules for interactions among agents. It determines the actions that can be taken during a turn and when to stop the negotiation [2]. The interaction begins with one agent making an offer. The receiving agent can then choose one of three actions: (i) *accept* the current offer, (ii) make a *counter offer*, or (iii) *end* the negotiation without an agreement. The first two actions, accepting the offer or making a counter offer, are based on the agent's utility function. This turn-taking interaction continues until an agreement is reached or the deadline is met.

3.2 Bidding Strategy

Any utility-based bidding strategy can be employed by the agent. According to such bidding strategies, agents calculate a target utility based on time or opponent's behavior and pick an offer that satisfies this target. In this study, as the bidding strategy, we adopted a bidding strategy from the literature [18] called the "Hybrid Negotiation" strategy, which combines time-based and behavior-based concession strategies in agent negotiation. This innovative approach ensures a dynamic adoptation to both the temporal constraints of the negotiation and the opponent's behavioral patterns. Traditionally, utility-based bidding strategies in the literature focus either on time or the opponent's behavior to calculate a target utility. In contrast, the "Hybrid Negotiation" strategy integrates these two factors.

By using this hybrid approach, our agents can make their decisions adaptable to their opponent's behavior as well as remaining time to complete the negotiation. Equation (2) shows the calculation of the target utility in this strategy.

$$TU_{Hybrid} = t^2 \times TU_{Time} + (1 - t^2) \times TU_{Behavior}. \tag{2}$$

TU_{Hybrid} is a weighted average of the time-dependent target utility (TU_{Time}) and the behavior-dependent target utility ($TU_{Behavior}$), with the weighting based on the square of the remaining time factor t. In the beginning of the negotiation, the agent takes into account its opponent's behavior more. While the deadline is approaching, it considers the target utility estimated by the time-based

strategy more.

$$TU_{Time}(t) = (1-t)^2 \times P_0 + [2 \times (1-t) \times t \times P_1] + t^2 \times P_2. \quad (3)$$

TU_{Time} represents the time-based component, which decreases exponentially as time advances, encouraging the agent to make concessions as the deadline approaches.

$$TU_{Behavior} = U(O_j^{t-1}) - \mu \times \Delta U. \quad (4)$$

$TU_{Behavior}$ takes into account the opponent's recent behavior, particularly how their utility has changed, allowing the agent to respond in kind.

$$\Delta U = \sum_{i=1}^{n} \left[W_i \times \left(U(O_h^{t-i}) - U(O_h^{t-i-1}) \right) \right]. \quad (5)$$

Equation (5) aggregates the changes in the opponent's utility over a window of n past bids, weighted by W_i, to inform the agent's behavior-based strategy.

$$\mu = P_3 + t \times P_3. \quad (6)$$

Finally, μ is a time-dependent parameter that influences the extent to which the agent mimics the opponent's behavior, increasing as the negotiation progresses.

3.3 Acceptance Strategy

A negotiation lasts until reaching an agreement or the deadline. When there is no agreement, both sides get a zero utility score or the reservation value if available when the deadline is up. The most common and successful strategy is the next utility acceptance strategy (AC_{next}). According to this strategy, the agent accepts its opponent's offer when the received utility of the opponent's current offer is equal to or greater than the agent's target utility estimated for the current round.

4 Negotiation-Based Data Sharing Framework

Building on our earlier work [3], we propose a platform for data sharing where two autonomous agents, a *data provider* and a *data consumer*, engage in a negotiation to decide together the terms and conditions for sharing data. The proposed data sharing approach is based on two key constructs. First, a data consumer (e.g., a company) must have specific some *business goals* that require the data owned by the data provider. Second, a data provider (e.g., a customer) must have some sort of **motivation** (i.e., incentive) for sharing their personal/private data. Figure 1 illustrates the architecture of the negotiation-based platform for data sharing. The negotiation between the data consumer and the provider is governed by the Alternating Offers Protocol [4, 25]. The data consumer initiates the negotiation with an offer, and the data provider either accepts the offer or makes a counter offer. Negotiation proceeds in a turn-based manner until the termination condition is satisfied, either through reaching a mutual agreement or upon hitting the deadline.

Particularly, after the enforcement of GDPR, individuals take privacy concerns more seriously and value their personal data accordingly. The suggested platform provides a negotiation mechanism for privacy-preserving data sharing. The data to be shared consists of a bundle of data pieces which may be of different data types (e.g., age, gender, GPS, etc.)[1] Certain data types may create a reason for privacy (or secrecy) concerns. In our framework, agents negotiate *(i)* the content of the data bundle (i.e. which data types to be shared), *(ii)* sharing policies (i.e., with whom to share, how long to share) and *(iii)* incentives to be given by data consumer.

As illustrated in Figure 1, a shared ontology is utilized to create a common understanding among agents. It formally defines data types (T) that can be exchanged, potential sharing durations (e.g.,

[1]Note that, It is important to distinguish between data and data types: data refers to the actual information, while data types indicate the category of that information. For example, '31' is data, whereas 'age' is the corresponding data type.

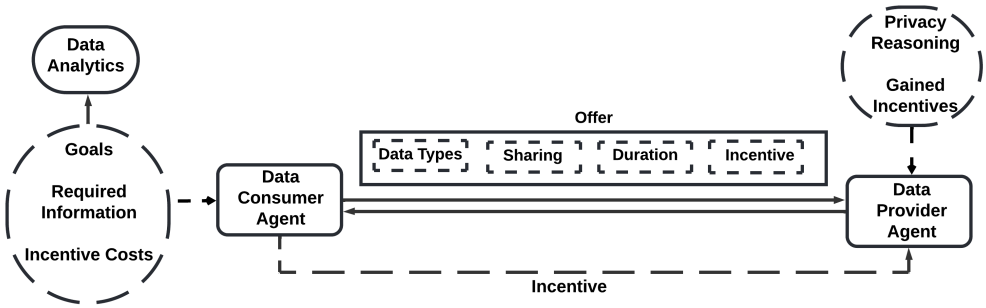

Fig. 1. Agent-based data sharing architecture.

six months, one year, etc.), possible sharing options such as "company only" or "company and third party" and various types of incentives (e.g., Call Minutes in telecommunications). Thus, the ontology's content is specific to the domain. In the proposed framework, both agents elicit stakeholders' preferences and beliefs and build a private knowledge base to be utilized for the reasoning process during the negotiation.

In this framework, an offer[2] can be formalized as follows: $o < d, t, s, p >$ where

- d represents a subset of data types under negotiation ($d \subseteq D$), with D being all possible data types defined in the ontology.
- t signifies the contract duration (i.e., the period during which the data will be shared once the offer is mutually accepted by both agents).
- s indicates the sharing policy, specifically detailing with whom the data may be shared.
- p represents the incentives offered in exchange for the shared data.

For example a possible offer in this context may look like, (o_1 =< {Occupation, Education Level}, "2 years", "share with the company and third parties) represented by the data consumer agent","3-month free phone call" > where *occupation and education level* constitute the data bundle, the duration of contract is 2 years, and the proposed incentive is *3-month free phone call*. In the following sections, we describe how the data consumer and provider agents assess offers that they receive on behalf of their stakeholders (i.e., company and individual customer).

4.1 Data Consumer Agent's Reasoning

To determine whether the data provider's offer is acceptable and what counteroffer to propose, the data consumer agent assesses the offers according to the expected utility function defined below. For this assessment, they require a knowledge base representing various characteristics of their owners. The consumer agent's knowledge base includes the possible *business goals of their owners*, the mapping of these goals to the necessary data types, and the importance of each goal to the data consumer agent. Additionally, the knowledge base contains the costs of incentives offered to the data provider to encourage the disclosure of the required data types. The agent's decision to accept or reject an offer is based on the calculated utility of the offer.

We propose a utility function for the data consumer agent, as shown in Equation (7), where the value of offers are determined by two factors: *the total utility of the included data types under specific conditions (i.e., duration and sharing policy)* $Value_{sharing}(d, t, s)$ and *the total cost of the incentives* $Value_{cost}(p)$. It is important to note that the ranges of these value functions should be

[2]In this work, the terms offer and bid are used interchangeably

consistent.

$$U(o < d, t, s, p >) = Value_{sharing}(d, t, s) - Value_{cost}(p). \quad (7)$$

When preparing its offer, the data consumer agent ($Agent_C$) makes certain that the data type bundle, d, adequately fulfills its desired objectives. Similarly, when assessing the data provider's counteroffer, the agent first examines whether the proposed bundle aligns with its requirements. This step is essential, as the provider ($Agent_P$) might have left out specific data types previously requested by the consumer due to privacy considerations.

$Agent_C$'s business goals are represented as $G = g_1, g_2, \ldots, g_k$, where k denotes the total number of goals. Each goal $g \in G$ requires specific data types to be acquired. For example, $d_1, d_2, d_3 \in D$ are needed to achieve g_1, while g_2 might only require d_3. $Agent_C$ seeks to acquire this data from $Agent_P$ through negotiation. Conversely, if some data types have high privacy value for $Agent_P$, they may avoid providing this data or demand more incentives. Typically, each goal has a different importance level for $Agent_C$. Therefore, a weight value, m_i, is associated with each goal $g_i \in G$, indicating its importance to $Agent_C$. The sum of these weights add up to one, i.e., $\sum_i^k m_i = 1$. In this framework, a goal is deemed to be *satisfiable* if $Agent_C$ possesses all the required data to achieve it.

We define the value of a data type bundle under specific sharing conditions $Value_{sharing}(d, t, s)$ as an additive function, as presented in Equation (8). Here, U_D represents the utility of the data type bundle, U_T represents the utility of the duration (the length of time the data will be shared with the company), and U_S represents the utility of the sharing policy (i.e., with which parties the data will be shared). The utility values range from zero to one. The importance of the data type bundle, duration, and sharing policy are denoted by w_d, w_t, and w_s respectively, with the sum of these weights equal to 1. It is assumed that there are no preferential dependencies among these factors. In other words, the data consumer would consistently favor a longer duration and a more extensive sharing policy, regardless of the data type bundle's content. However, the importance of these utilities may fluctuate based on the specific data types included in the bundle. For example, the utility of permitting the company to share the data bundle with a third party might be higher for certain data types due to particular business requirements. In such scenarios, generalized additive utility functions can be employed, where the utility of each factor is calculated for every possible data combination.

$$Value_{sharing}(d, t, s) = U_D(d) \times w_d + U_T(t) \times w_t + U_S(s) \times w_s. \quad (8)$$

The utility of a data type bundle, denoted as $U_D(d)$ in Equation (9), is determined by summing the weights of the goals that can be satisfied by d. As discussed, here, $Satisfiable(g_i, d)$ is equal to 1 if d includes all the data types needed for goal g_i; otherwise, $Satisfiable(g_i, d)$ is set to 0. It is crucial to highlight that this approach assigns a value to the bundle as a whole, rather than assessing each individual data item separately. The rationale behind this is that certain data items may hold little value on their own but become significant when combined with others. For example, if $Agent_C$ requires both d_1 and d_2 to fulfill g_1, the absence of either d_1 or d_2 would hinder the achievement of g_1.

$$U_D(d) = \sum_i^k m_i \times Satisfiable(g_i, d). \quad (9)$$

Remember that an offer is composed of four key components: the data type bundle, the duration, the sharing policy, and the incentives provided to the data provider. These incentives come with a cost, which $Agent_C$ seeks to minimize. The cost value of the incentives, denoted as $Value_{cost}(p)$, is a function that translates the cost of incentives for the obtained data types d into a real number between zero and one [0, 1]. A higher value suggests that the selected incentive is particularly expensive in relation to the data bundle. It's important to note that the cost value will be greater

than zero if the incentive cost for $Agent_C$ is lower in value compared with the worth of the data provided by $Agent_P$.

4.2 Data Provider's Reasoning

In this section, we describe how the data provider agent assesses an offer based on privacy concerns and the benefits of the incentives provided by the data consumer agent. We define privacy using two main components. The first is the **desire for secrecy**, which reflects the data owner's reluctance to share certain information simply because they prefer to keep it private. For instance, an individual may not want others to see a video of them singing off-key at a karaoke event. While sharing this video may not cause harm, it could cause discomfort. The second, is the **risk/fear of harm** and uncertainty regarding potential unethical or improper use of the data by others without consent. These components together determine the **privacy value** that the data provider assigns to specific data content. Consequently, the knowledge base of the data provider agent includes their preferences and beliefs about the secrecy and risk values associated with the information in their knowledge base.

People exhibit significant variations in their levels of secrecy and risk perception when sharing certain data types, as well as in how they value their private information. Some data types, such as social security numbers, are universally considered private. However, there is considerable variation among individuals (or individual companies) regarding the secrecy of other data types. Similarly, some people perceive certain data as highly susceptible to privacy breaches, while others may feel unconcerned about the same data. Thus, individuals value their data differently. In summary, an effective data sharing approach should be attuned to these individual differences.

During the negotiation, the data provider agent evaluates both the utility of the incentive proposed by $Agent_C$ and the extent of privacy intrusion involved in sharing the personal or private information requested. Equation (10) shows how $Agent_P$ assesses the utility of an offer, where c_p, c_t, and c_s serve as the coefficients for the expected tradeoff between incentive and privacy violation, duration, and sharing policy, respectively. It's important to note that these coefficients are normalized to sum to one. We assume that $Agent_P$ employs a utility function, $Value_{Incentive}(p)$, which translates each potential incentive into a real value between zero and one, reflecting the user's needs or preferences. Similarly, the provider agent uses utility functions for duration $U_T(t)$ and sharing policies $U_S(s)$, mapping the given values for these factors to a utility value within the zero to one range. Higher utility values correspond to more favorable conditions. For example, a sharing policy that limits data sharing to the company alone is preferred over one that includes sharing with third parties.

$$\begin{aligned} U(o < d, t, s, p >) = c_p \times [Value_{Incentive}(p) - Value_{privacy}(d)] \\ + c_t \times U_T(t) \\ + c_s \times U_S(s) \end{aligned} . \tag{10}$$

When calculating $Value_{privacy}$, which represents the cost of a privacy breach, $Agent_P$ assesses both the information's level of secrecy and the risks associated with sharing the requested data (i.e., potential negative consequences). Secrecy has a psychological aspect, reflecting an individual's desire to keep certain personal information private, regardless of whether it could be used against them. As such, secrecy is a deeply personal matter. In the proposed framework, the data provider will determine the secrecy level for each type of information. For instance, a data provider might be more reluctant to share their Social Security number than their phone number, so that $SL(SocialSecurityNumber) \succ SL(PhoneNumber)$. Higher secrecy levels indicate greater unwillingness to share the corresponding information. The secrecy scale ranges from zero to one.

Thus, the provider agent defines $Risk(x)$ to indicate the level of risk associated with sharing each data item $x \in I$ from their point of view. Although assessing the risk of individual data types may seem straightforward, the overall risk can be influenced by the combination of data types being shared. Sharing a single piece of data might not be risky on its own, but when combined with other data, it could reveal sensitive information. For example, a provider might assign a low risk to "occupation" and a medium risk to "GPS", but a high risk when these two are shared together, as it could disclose the company they work for. Consequently, providers may assess the risk of individual data types and their combinations differently. To address this, the framework allows providers to specify the risk associated with subsets of data types (i.e., $Risk_{DP}(Y)$ where $Y \subset D$) as well as the risk for each data type individually (i.e., $Risk(x)\ \forall x \in D$).

Accordingly, Equation (11) illustrates how the value of privacy is estimated in our framework. Here, $Risk(x)$ indicates the risk associated with sharing the requested item x from the data provider's perspective, and $Risk_{DP}(Y)$ represents the risk of sharing a subset of data types, with DP denoting data dependencies (i.e., sets of interdependent data types), where $Y \subset D$. SL signifies the normalized secrecy level of the given information type. For example, if a data provider believes that sharing both "GPS" and "occupation" together is riskier than sharing them individually, they define a dependency such as $GPS, occupation$. In our formulation, we opt to consider the maximum value of privacy violation rather than the average privacy violation of each piece of information in the bundle. This approach is chosen because a bundle may contain data types with either very high or very low privacy violations, and taking the average may not accurately reflect the significance of the violation.

Consequently, Equation (11) demonstrates how privacy value is calculated within our framework. In this context, $Risk(x)$ reflects the risk involved in sharing x from the perspective of the data provider, while $Risk_{DP}(Y)$ represents the risk linked to sharing a subset of data types, where DP stands for data dependencies (i.e., groups of interrelated data types) and $Y \subset D$. SL denotes the normalized secrecy level for the specific type of information. For instance, if a data provider considers that sharing both "age" and "gender" together poses a higher risk than sharing them separately, they would define a dependency such as $age, gender$. In our approach, we choose to prioritize the maximum value of privacy violation over the average violation for each data item in the bundle. This decision is based on the idea that a bundle may include data types with either extremely high or low privacy risks, and averaging might not fully capture the severity of the violation.

$$Value_{privacy}(I, d, s) = \max\Big(\max_{x \in D}(SL(x) * Risk(x)), \max_{Y \subseteq D \wedge Y \in DP \wedge z \in Y}(SL(z) * Risk_{DP}(Y))\Big) \tag{11}$$

5 Preference Elicitation and Negotiation Tool

The proposed negotiation approach and framework are versatile and can be applied to various data sharing scenarios. To illustrate and assess our approach, we examine a use case from the telecommunications sector. In this context, a telecommunications company aims to conduct data analytics on its customers' data. According to regulations, they must obtain customers' consent to keep and handle their private information. Rather than imposing a specific data sharing policy, the company engages with customers to establish a mutually acceptable policy using the proposed approach. For this purpose, the data consumer initiates a two-way negotiation process with data providers to determine what data will be shared and what incentives will encourage sharing. For example, the company might offer promotions such as "2GB's", "200 minutes", and so on, to gain customers' consent for the data sharing policy.

Table 1. Incentives

Incentive Name	Duration (Days)	SMS	Minutes	GB's
Incentive 1	7	100	60	1 GB
Incentive 2	30	50	30	1 GB
Incentive 3	30	100	60	-
Incentive 4	30	-	30	2 GB
Incentive 5	30	500	-	-
Incentive 6	30	-	100	-
Incentive 7	30	-	-	4 GB
Incentive 8	90	-	30	2 GB
Incentive 9	90	500	-	-
Incentive 10	90	-	100	-
Incentive 11	90	100	60	-
Incentive 12	90	-	-	4 GB

To develop such a mechanism, first we must identify the relevant data types, the kinds of incentives the company might provide in return for the requested data, and the sharing conditions (i.e., duration and intended recipients). Those issues and their possible values form the negotiation domain, which will be available for both sides as shown in Figure 1. For general usage purpose, we developed a domain elicitation tool where negotiation issues and their possible values are acquired from the data consumer side (Section 5.1). Other part of this tool is used for eliciting preferences of data provider and consumer separately (Section 5.2). Lastly, we present a human-agent negotiation tool where the data consumer agent on behalf of the company negotiates with a human negotiator representing the customer on the given domain (Section 5.3).

5.1 Shared Ontology for Negotiation Domain

We design a straightforward interface to elicit domain knowledge. In our scenario, the negotiation domain includes four key issues: the collection of all potential data types being negotiated, and the incentives proposed by the data consumer, the duration for which the data will be shared, and the intended recipients. The following *data types* are defined in our scenario: GPS information, marital status, education level, occupation, age, neighborhood, email, daily call duration, daily SMS usage, daily Internet usage, call log (i.e., log of all incoming and outgoing calls), application usage statistics, and daily social media activity (general usage duration, not content). Regarding *duration and sharing policies*, the contract duration specifies how long the access rights are granted, and the sharing policy indicates with whom the data will be shared. In our case, two sharing policy options are defined: (i) sharing only with the company, and (ii) sharing with the company and third parties. There are four sharing duration options: (i) six months, (ii) one year, (iii) three years, and (iv) five years. For *incentives*, 12 incentive packages are defined, as shown in Table 1.

5.2 Preference Elicitation

In this section, we provide the preference elicitation phase for both data consumer (i.e. telecommunication company) and data provider (i.e., customer).

5.2.1 Data Consumer. The data consumer begins by defining the company's business goals and linking the necessary data types to these goals, along with their respective importance to the company. Subsequently, the company provides cost details for the incentive packages listed in Table 1. The utility of an offer for the company is determined by how well it meets the company's

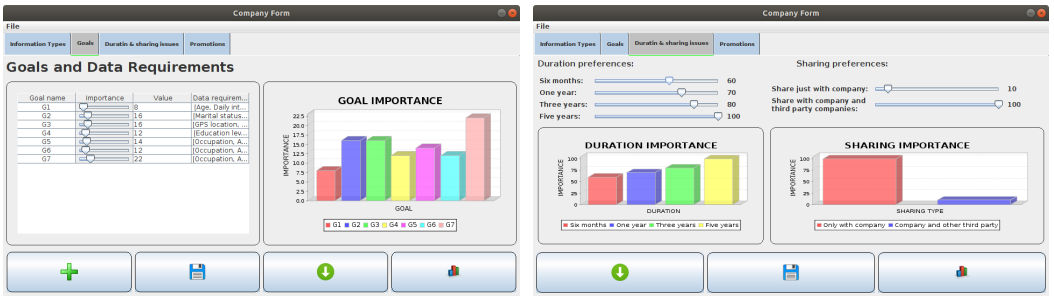

(a) Data Consumer's Goal Specification (b) Data Consumer's Preferences for Sharing Policy

Fig. 2. Comparison of goal specification and sharing policy preferences.

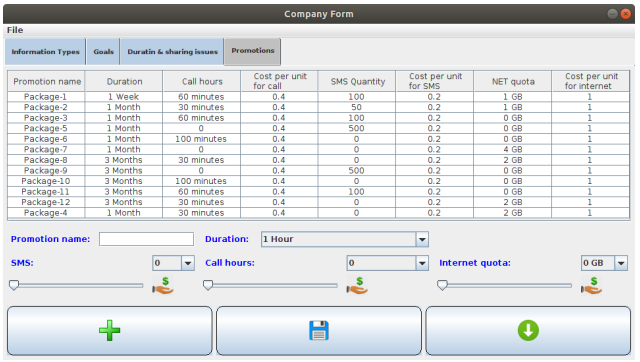

Fig. 3. Data consumer's incentive—promotion specification.

goals. Therefore, the data consumer agent assesses which goals can be satisfied by the data types included in the offer. Figure 2(a) illustrates the dialogue interfaces used to gather these inputs, with weight values ranging from 0 to 100. Figure 2(b) illustrates the process of eliciting the company's preferences related to sharing policy and duration. As the cost of promotion packages is a key factor in calculating the expected utility for the company agent, the company representative provides the cost details for each component of the packages. Figure 3 showcases the interface where both the content and the cost of the promotion packages are specified.

5.2.2 Data Provider. The data provider (i.e., in our framework, customers), can specify their preferences in three steps: (i) preferences regarding the sharing risk and secrecy level of each predefined data type, (ii) sharing policy (which parties to share with), and (iii) duration, as shown in Figure 4. As mentioned, some data types may pose more risk when shared together. Such combinations are referred to as *data dependencies* in our framework. Assessing data dependencies is subjective; thus, our tool elicits these combinations along with a quantified risk of sharing them together. Finally, the data provider should specify their preferences for promotion packages as shown in Figure 5, assigning a value between 0 and 100 to each package. Higher values indicate greater preference.

5.3 Negotiation Tool

We developed a negotiation tool that enables the company agent and customers to interact and reach a consensus on data type sharing. In our framework, the data consumer is represented by an

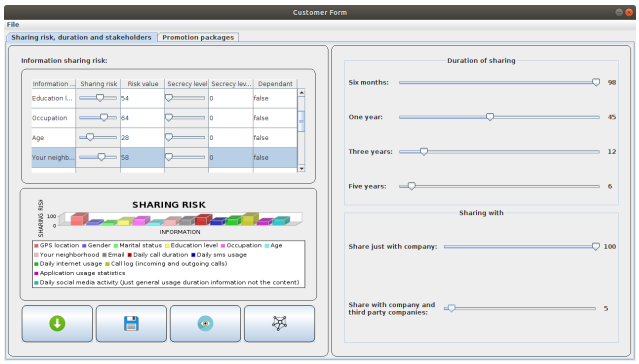

Fig. 4. Data provider's preferences regarding sharing risk, secrecy level, duration, and type of sharing.

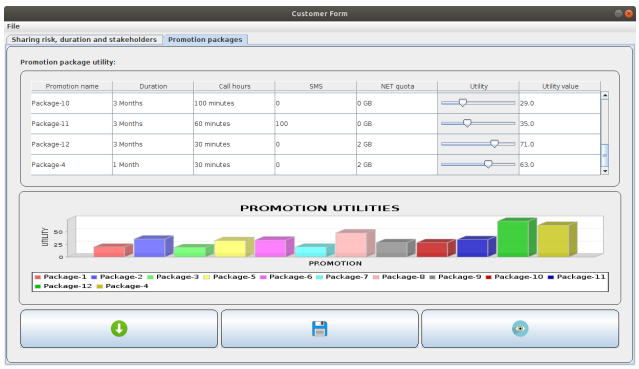

Fig. 5. Data provider's preferences on promotion packages.

intelligent agent, while the data providers are human participants. During our experiments, human participants specify their preferences related to the data sharing policy using the elicitation tool described above. These preferences are kept private by the system.

The negotiation starts with an offer from the company agent, who initially proposes a bid that maximizes its own utility. When the company agent's offers are displayed, the user rates them on a 0-10 Likert scale, as shown in Figure 6. If the user accepts the offer, the negotiation concludes with an agreement. Otherwise, the user makes a counter-offer by selecting values for each negotiation issue, as depicted in Figure 7. This interaction continues in a turn-taking manner.

6 Experimental Evaluation

We conducted a user experiment to evaluate our proposed framework. We recruited 60 participants in our experiment, ranging from bachelor's and master's students to office employees. Participants' ages ranged from 18 to 60, with a mean of 27.5, and there were 37 male and 23 female participants. In the experiments, each user played the role of the data provider, specifying their preferences using our preference elicitation tool and interacting with the company agent in two different sessions. In one session we adopt the baseline setting where participants could only accept or reject the given offers iteratively until reaching an agreement or deadline (i.e., take-it or leave-it in the literature), a commonly baseline in privacy negotiation literature [5, 12, 24]. In the other session (*our negotiation framework*), they actively negotiated with the consumer agent according to the proposed setting. To minimize the learning effect on negotiation results, we used a randomization technique (i.e.,

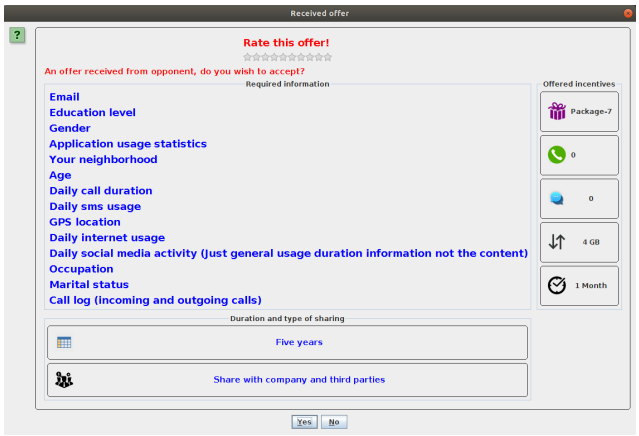

Fig. 6. Company agent's offer.

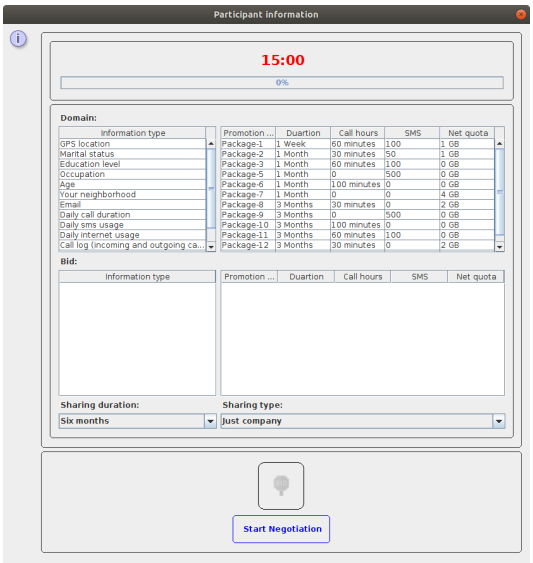

Fig. 7. Making counter offer.

counterbalancing). Half of the participants started with the baseline approach and then experienced our negotiation framework after a short break, while the remaining participants did the reverse.

To go into further detail, Figure 8 outlines the steps of our experiments. In the first step, users specify their preferences using the preference elicitation tool shown in Figure 4. They provide values for the secrecy level of each data type (SL), risk values of sharing each data type (Risk), risk dependencies ($Risk_{DP}$), and utility functions for duration ($U_T(t)$) and sharing policies ($U_S(s)$). Note that the coefficients for the expected tradeoff between incentive and privacy violation, duration, and sharing policy (c_p, c_t, and c_s, respectively) are set empirically. In Step 2, if the session requires negotiation, the system initiates the negotiation session after loading the preference profiles. Otherwise, the baseline approach of accepting/rejecting the given offers is used. During the session (Step 3), the exchanged offers, their subjective evaluations by the data provider (i.e., human participant),

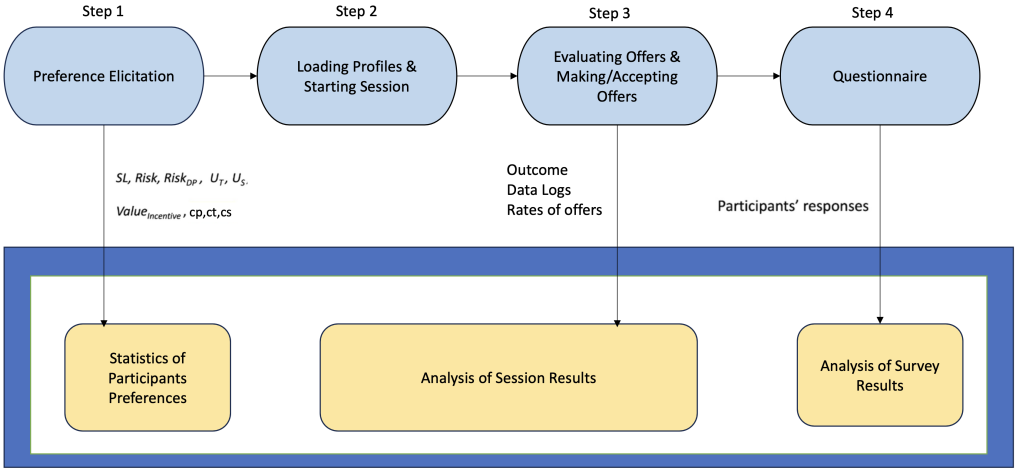

Fig. 8. Experimental setup.

and the agreements are recorded for session analysis. Finally, in Step 4, participants complete a questionnaire with 15 questions to assess their satisfaction level.

For this experiment, we developed a profile for the company agent. In our scenario, the company has identified seven specific goals as follows:

— Goal 1 with an importance of 8: Daily Internet Usage, Age
— Goal 2 with an importance of 16: Marital Status, Age, Daily Call Duration, Occupation
— Goal 3 with an importance of 16: Neighborhood, GPS, Occupation
— Goal 4 with an importance of 12: Daily Social Media Activity, Email, Age, Education Level, Occupation
— Goal 5 with an importance of 14: Gender, Application Usage Statistics, Age, Occupation
— Goal 6 with an importance of 12: Gender, Call List, Age, Occupation
— Goal 7 with an importance of 22: Application Usage Statistics, Age, Daily Call Duration, Daily Internet Usage, Daily SMS Usage, Occupation

In our preliminary research, we found that participants were more comfortable expressing their preferences on a scale from zero to 100, rather than using real numbers between 0 and 1. Consequently, we implemented this scale for all aspects of the expected utility, meaning that the agent's expected utility now ranges from zero to 100. Based on this, we established the utility values for other elements of the offers as follows. In summary, the company shows a preference for a 5-year duration over 3 years, 3 years over 1 year, and 1 year over 6 months. Additionally, it favors the sharing policy that includes sharing with third parties over sharing only with the company. Therefore, their utility values are set as follows:

— Duration (5 years) = 100
— Duration (3 years) = 87
— Duration (1 year)= 75
— Duration (6 months)= 60
— Sharing policy (company only) = 80
— Sharing(company and with third company) = 100

To estimate the overall expected utility, the weights are distributed as follows: 0.7 for goal satisfaction, 0.2 for duration, and 0.1 for sharing policy. The weight values guide the calculation,

Table 2. Cost of Incentives for the Company

Incentive	Cost	Incentive	Cost	Incentive	Cost	Incentive	Cost
Incentive 1	45.0	Incentive 2	23.0	Incentive 3	44.0	Incentive 4	14.0
Incentive 5	100.0	Incentive 6	45.0	Incentive 7	4.0	Incentive 8	14.0
Incentive 9	100.0	Incentive 10	40.0	Incentive 11	44.0	Incentive 12	14.0

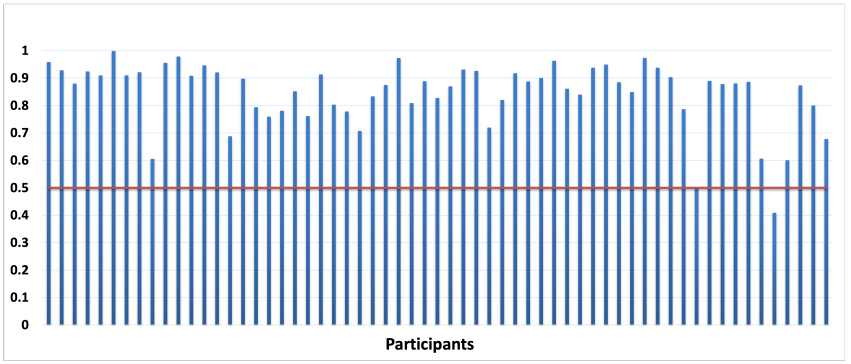

Fig. 9. Correlations between estimated utilities and participants' ratings.

beginning with determining the cost of each promotion (see Table 2), which is then normalized using the formula $\frac{cost-min(c)}{max(c)-min(c)} \times 100$, where $min(c)$ and $max(c)$ represent the lowest and highest costs among the packages.

In the following sections, we first investigate to what extent the proposed consumer preference model can capture participant's preferences (Section 6.1),analyze the session outcomes (Section 6.2), and provide an elaborate analysis of the outcomes with respect to the participants personality traits and privacy needs (Section 6.3). Finally, we provide a summary of our survey results (Section 6.4).

6.1 Assessment of Consumer Preference Model

The main components of the proposed preference model for the data provider are asked before the sessions (preference elicitation phase). It is worth noting that the system (i.e., data consumer agent) does not use this information while making its decisions. During the sessions, participants are asked to rate each offer made during the sessions out of 10. These ratings are also not used by the consumer agent. To assess the effectiveness and expressiveness of the proposed preference model for the data consumers, we examine the correlation between the given ratings by the participants and our model's estimation for the utility of offers.

Figure 9 shows the correlations for each participant's rankings with the estimated preference model where the x-axis and the y-axis denote the participant's corresponding correlation values, respectively. The results show that 96% of sessions have a correlation above 0.5. None of them involves a negative correlation. While 73% of sessions have a correlation above 80%, indicating that the elicited preferences and the users' ratings are highly correlated. Here, the blue bars show the correlation of the session, and the red line depicts the threshold of 0.5 correlation. These results support the idea that the proposed data provider preference model can capture the preferences of the participants in general. That means we can even automate the negotiation process. The data provider agent can negotiate on behalf of the human participants if it elicits their preferences.

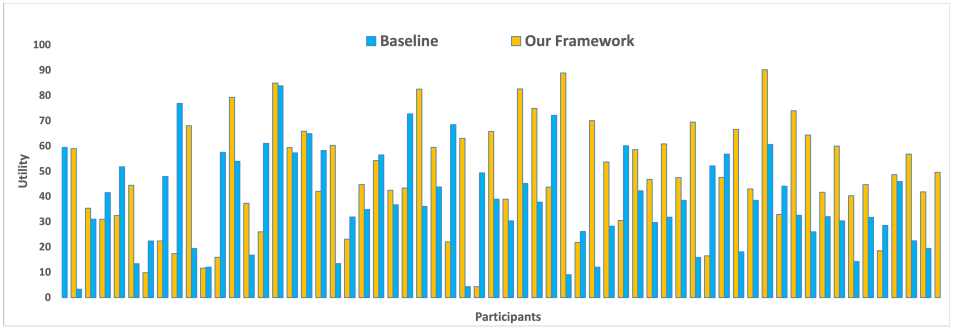

Fig. 10. Estimated user utility for both settings.

6.2 Analysis of Session Outcomes

Given that our experimental design utilizes a within-subject framework, we chose to conduct a two-tailed dependent sample t-test. To verify the suitability of this test, we initially conducted the Kolmogorov-Smirnov test for normality and Levene's test for homogeneity of variance. If the data distribution met the criteria of these tests, the dependent sample t-test was applied; otherwise, we employed a non-parametric statistical method, namely the Wilcoxon-Signed Rank test. All statistical results are presented with a 95 confidence interval.

We observe that 58 out of the 60 negotiation sessions have ended up with an agreement in our framework. While the baseline strategy has resulted in 49 agreements out of 60 negotiation sessions. In terms of the agreements, our framework outperformed the classic accept/reject baseline information-sharing system. Figure 10 shows the estimated utilities that participants received in both scenarios, where the x-axis and y-axis depict the participant number and the corresponding utility values that they received according to the estimated preference model, respectively. We observe that 28% of the participants received higher utility in the baseline while 69% of the users received higher utility in our framework. Overall, only 3% of the users received the same score. On average, the participants received around the utility of 38.56 in the baseline system, whereas they gained about the utility of 48.08 as seen in Figure 11(a), with the orange and blue bars representing the scores for the baseline and our framework, respectively. Since the data met the criteria for both the Kolmogorov-Smirnov test and Levene's Test, we utilized a two-tailed dependent sample t-test. The analysis revealed a statistically significant difference in the estimated participant utility at a 95 confidence level (t=2.59738, p=.0010569). In conclusion, data providers benefited more when using our framework to share their information with data consumers, potentially leading to higher user satisfaction.

On the other hand, we did not find any statistically significant difference in the agent's utility under a 95% confidence interval (z=-0.1708,p= 0.86502.), where we can observe the average utilities gained by the data consumer agent (68.2 versus 69.3 on average) in Figure 11(b). In order to evaluate the social welfare, we consider *Normalized Social Welfare* [10] (i.e., a product of both sides' utilities) as seen in Figure 11(c). When we applied the two-tailed dependent sample t-test to Normalized Social Welfare, we found a statistically significant effect with a confidence interval of 95% (t=2.76383, p=0.006614), with respective averages of 0.24 vs 0.30. As a result, social welfare and data provider's utilities are higher on the proposed approach while not significantly influencing the data consumer agent's utility.

Lastly, we analyzed the efficiency of the consensus-building process. As shown in Figure 12, our framework significantly reduced the number of rounds required to reach a consensus compared

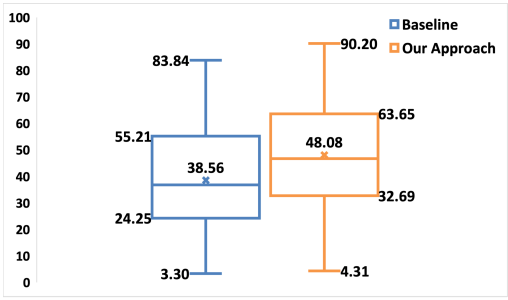

(a) User Utility

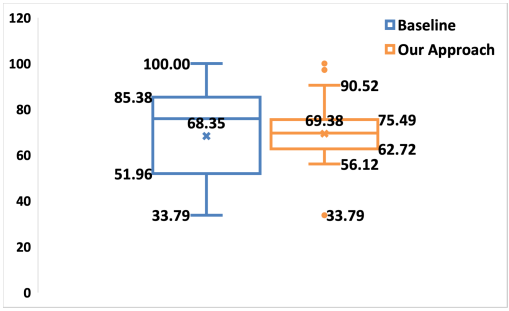

(b) Agent Utility

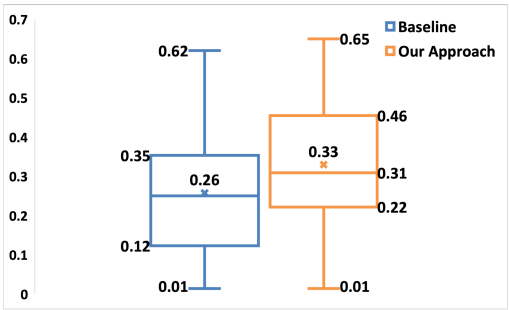

(c) Social Welfare

Fig. 11. Comparison of user utility, agent utility, and social welfare across experiments.

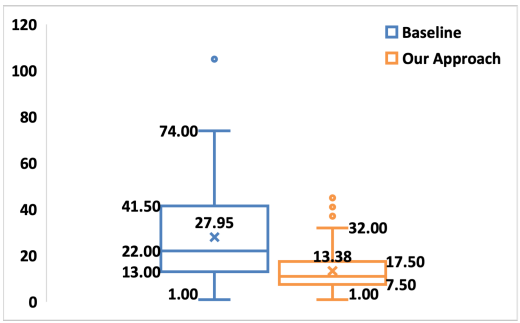

Fig. 12. Total number of rounds.

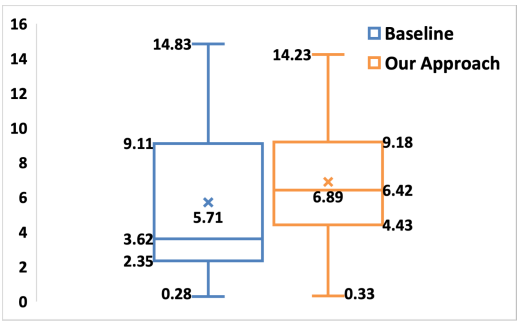

Fig. 13. Agreement time.

with the baseline (averages of 13.2 vs 27.9, respectively). This effect was statistically significant based on a Wilcoxon Signed-Rank test ($z = -5.3471$, $p < 0.0001$). However as depicted in Figure 13, this improvement came at the cost of longer session durations, where the baseline proved to be faster (6.89 vs 5.71, respectively). This reveals a fundamental trade-off: our framework promotes fewer, more substantive interactions to reach a decision, whereas the baseline relies on a larger quantity of quicker, less impactful exchanges. These results support that our framework successfully reduces the number of decision-making steps for the user, which is crucial in human-in-the-loop systems.

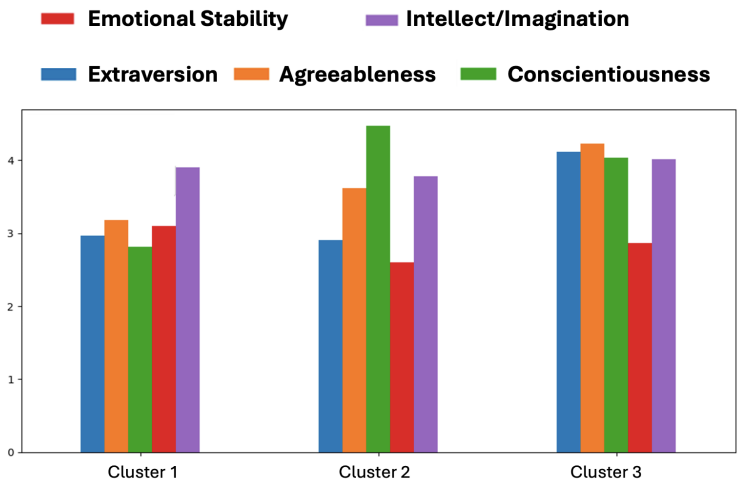

Fig. 14. Big five personality scores.

6.3 Further Analysis of the Outcomes

To gain deeper insights into our participants, we conducted a pre-survey, which consisted of the Big Five Personality test (Section 6.3.1), a negotiation self-assessment survey (Section 6.3.2), and a survey on how concerned they are about their privacy taking inspiration from the 'Westin's Privacy Index' (Section 6.3.3). We analyze the participants' outcomes in both setting by clustering them with respect to their personality, negotiation skills and privacy concerns.

6.3.1 Big Five Personality. We incorporated the Big Five Personality Test into our study to leverage its well-established framework for assessing fundamental personality traits. This test consists of several questions to assess five personality traits: *Extraversion*, *Agreeableness*, *Conscientiousness*, *Emotional Stability*, and *Intellect* as the key predictors of human behavior. We choose the IPIP-BFM-20 for its effectiveness as our survey of choice [26]. This survey helps us assess a score for each category so that a participant' personality can be represented as a vector of those scores. Accordingly, we clustered participants' personality data and observed three clusters regarding the Elbow analysis. Figure 14 shows the average scores of each cluster for five personality categories.

For each cluster, we analyzed the results of the outcomes illustrated in Figure 15. When we investigated the results, we noticed that most of the failures across both systems (i.e., not reaching an agreement) occurred in the first cluster and parties received relatively lower utility, where the agreeableness of the participants is the lowest according to the Personality tests. The third cluster received the highest utility on average, where the participants' agreeableness and extraversion scores were the highest compared with participants in other clusters.

6.3.2 Negotiation Self Assessment. To take participants' negotiation attitude into account, taking inspiration from a negotiation assessment survey from the literature, we asked them to rank the following items with respect to their priorities: decreasing opponent's utility, decreasing agreement time, increasing their own utility and getting the best deal for both sides where the highest and lowest priority are denoted by 4 and 1 points respectively [8]. Representing each participant as a vector of those ranks, we clustered and found out four clusters of participants as illustrated in Figure 16.

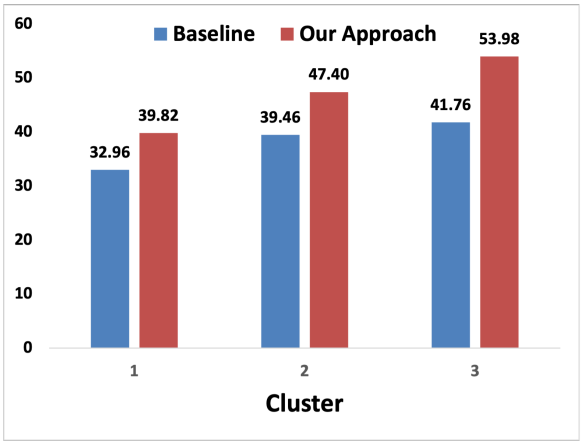

Fig. 15. User utilities for personality clusters.

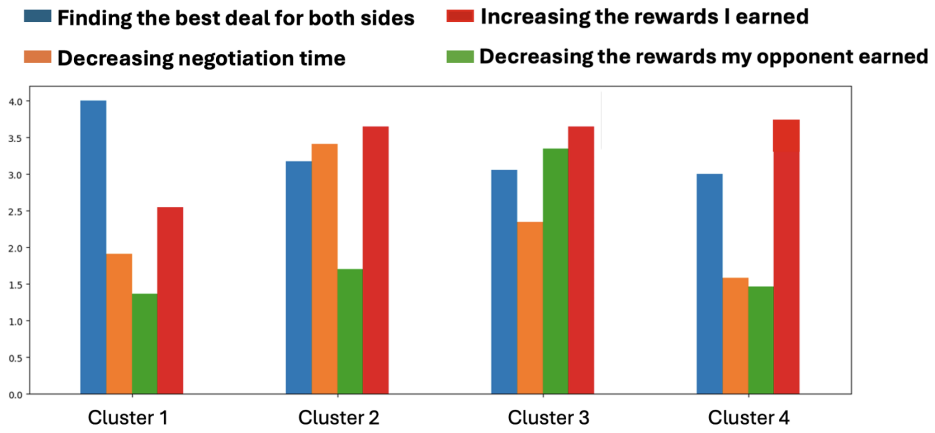

Fig. 16. Negotiation assessment scores.

For each cluster, we analyzed the results of the outcomes illustrated in Figure 17. When we investigated the results, we noticed that participants that were more cooperative gained the highest utility on average (see Cluster 1) while participants that were looking to decrease the opponents reward ended up with the lowest utility on average (see Cluster 3). Notably, both of the failed sessions in the proposed system, came from users that were in the fourth negotiation cluster (the cluster with the lowest interest in finding the best deal for both sides and decreasing negotiation time). In both cases, over 80%, both sessions included more than 80 percent of silent moves (i.e., repeated offers that were identical or nearly identical to the previous one, resulting in negligible or no change in utility for either party). Since there were not enough exchanged offers to find a suitable agreement for both sides, the negotiations timed out.

6.3.3 Privacy Self Assessment. Moreover, when we clustered our participants according to the answers from the Privacy Assessment Survey. Taking inspiration from Westin's Privacy Segmentation Index survey we cluster our participants on three groups [28]. We follow the convention of grouping our participants on the following categories similar to [14].

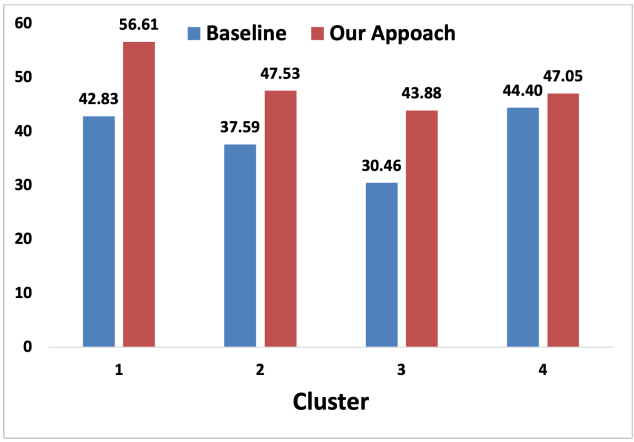

Fig. 17. User utilities for negotiation clusters.

(1) **Privacy Fundamentalists**: These are individuals who are highly concerned about their privacy. They are hesitant to share personal information and are skeptical of organizations using their data.
(2) **Privacy Unconcerned**: These people are not particularly worried about their privacy. They are more willing to share personal information and are less skeptical of organizations using their data.
(3) **Privacy Pragmatists**: These individuals will balance the benefits and risks of sharing personal information. They are willing to share their data if they believe the benefits outweigh the risks and if there are adequate protections.

When we analyze the average responses of the groups depicted in Figure 18. We can observe that group 1 aligns well with Privacy Fundamentalists, they believe it is absolutely essential to keep their personal information private (average score of 4.56 for Q4) and they are not willing to share their information no matter the benefits (average score of 3.84 for Q5). While group 3 aligns well with Privacy Pragmatists, they are just as concerned about their privacy (average score of 4.56 for Q4), however they tend to weigh the gained benefits against the potential risk of sharing their information (average score of 4.81 for Q5) . Lastly, group 2 aligns well with Privacy Unconcerned, they believe existing laws and organizational practices provide a reasonable level of protection for data providers (average score of 2.71 for Q3) and they don't believe it is absolutely essential to keep keep their personal information private hidden (average score of 2.48 for Q4).

For each cluster, we analyzed the results of the outcomes illustrated in Figure 19. When we investigated the results, we noticed that Privacy Pragmatists and Fundamentalists gained the highest utility on average, while Privacy Unconcerned ended up with the lowest utility on average. It can be easily seen that participants who are unconcerned about their privacy, gain little by adopting the proposed approach compared with the baseline approach (on average 37.74 versus 38.94). On the other hand, participants who care about their privacy (clusters 1 and 3) gained significantly more when they used our negotiation approach in contrast to the baseline case.

6.4 Post Survey

Alongside measuring objective performance indicators like agent and participant utility and the number of rounds, we also carried out a subjective assessment of the system through a questionnaire

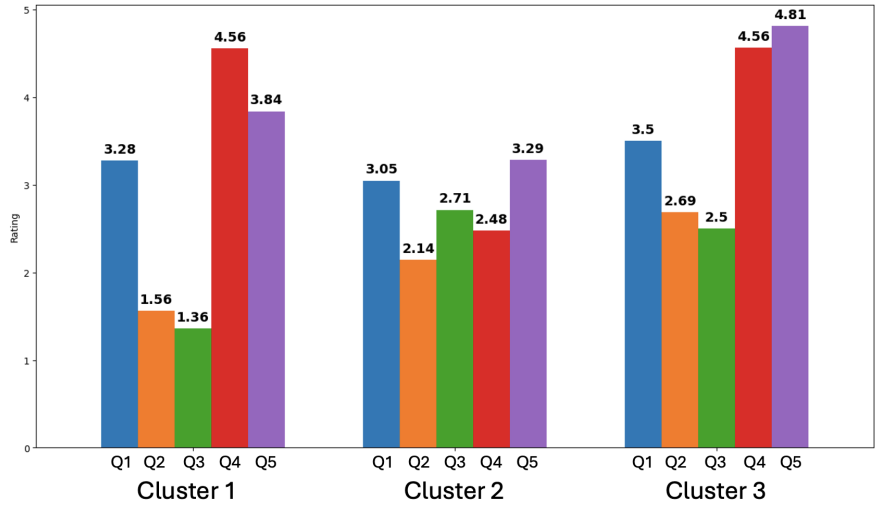

- **Q-1:** Consumers have lost all control over how personal information is collected and used by companies.
- **Q-2:** Most businesses handle the personal information they collect about consumers in a proper and confidential way.
- **Q-3:** Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today.
- **Q-4:** I believe it is absolutely essential to keep my personal information private. I do not trust companies to handle my data responsibly.
- **Q-5:** When deciding whether to share my personal information, I weigh the benefits (like personalized services, discounts, etc.) against the potential risks.

Fig. 18. Privacy personality centroids and survey questions.

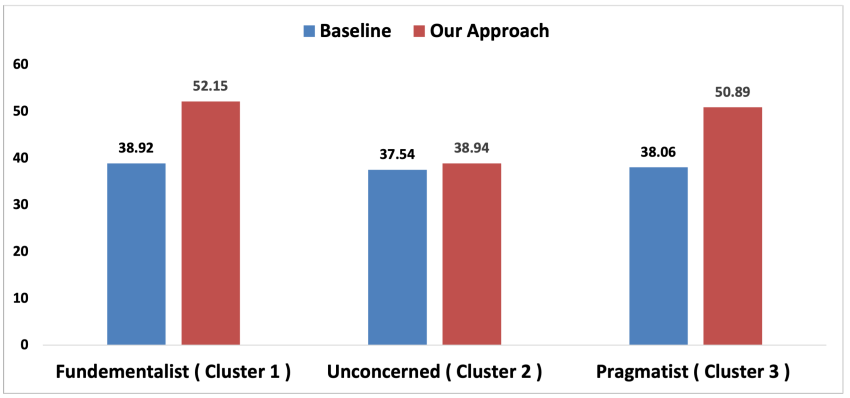

Fig. 19. User utilities for privacy clusters.

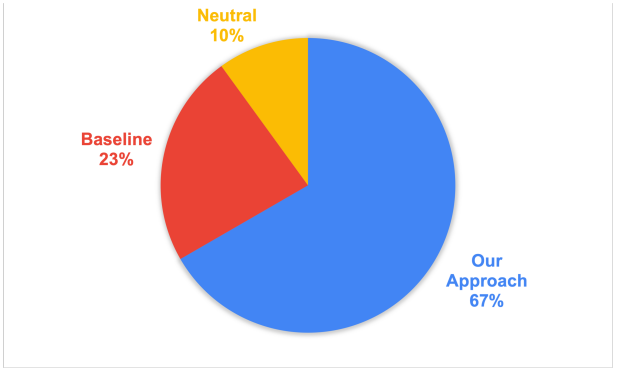

Fig. 20. Which system they prefer pie chart.

completed by participants at the conclusion of the experiment. The Likert scale questions were structured on a five-point scale, where 1 indicates strong disagreement, 3 represents neutrality, and 5 signifies strong agreement. When we analyze the responses given in our after-experiment survey. The results show that the majority of the participants prefer to use our negotiation-based information-sharing system. The distribution of which system the participants preferred is depicted in Figure 20. We can also observe from the responses that the participants agreed on the fact that the preference tool increased their awareness about the issues surrounding their privacy (average score of 3.93 ± 0.74 for Q3). It can also be seen that the participants have appreciated the ability to create their own offers and negotiate with the data consumer (average scores of 4.33 ± 1.08, $4.55 + 0.57$ for Q5 and Q6, respectively). The average responses given to each survey questions is depicted in Figure 21.

7 Conclusion

In conclusion, this study contributes to the ongoing conversation about incorporating negotiation mechanisms into data-sharing processes. As we aim at improving data-sharing frameworks, it is evident that data providers value the ability to make their own offers and engage in negotiations with data consumers.

This can be seen in the results of our empirical analysis as well. In this work, we also explore a preference negotiation tool to assess how well we can capture the preferences of data providers to present them as autonomous agents in automated negotiations. Our findings demonstrate that the preferences recorded are closely aligned with the actual preferences of the data providers. Suggesting that our proposed tool holds strong potential for representing data providers in automated negotiations. We plan to explore the effectiveness of fully automated negotiation mechanism for information sharing in future works.

Moreover, for future work, we plan to extend the application of the negotiation framework to different domains, such as healthcare, finance, and public services. These areas offer distinct challenges regarding data privacy and negotiation dynamics. Exploring these diverse contexts will enhance our understanding of the versatility and adaptability of our framework and contribute to its evolution into a robust tool for data-sharing capable of handling a diverse set of domains. Additionally, it would be interesting to conduct a larger-scale study to evaluate the framework's effectiveness across broader user populations and domains. To enable this, the framework can be deployed as a mobile application and a public competition can be organized to support large-scale, real-world participation. Lastly, while this study focused on modeling, negotiation, and preference elicitation, we recognize that formal security guarantees and enforcement mechanisms are

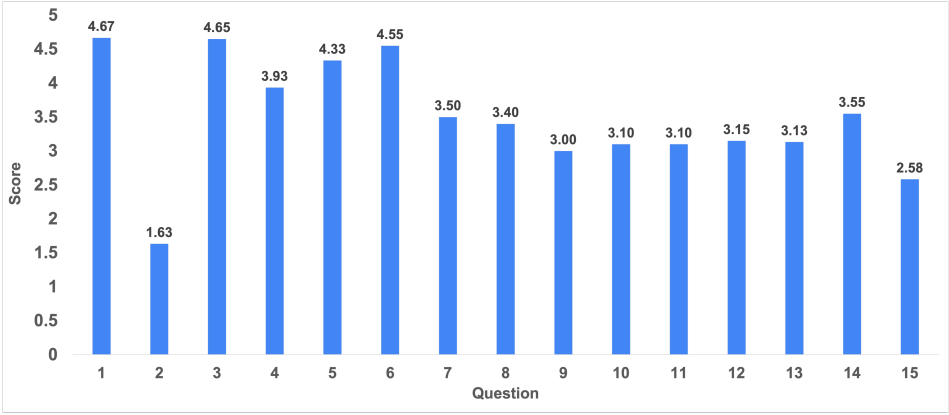

(1) The instructions provided to me for the experimental negotiation were clear.
(2) It was not clear to me how I should use the preference elicitation tool in which you specified your preferences on information sharing such as secrecy level, risk level, and so on.
(3) It was clear to me how I make my offers in the given negotiation tool.
(4) Specifying my preferences in the given tool increased my awareness of privacy.
(5) I like the idea of negotiating about the information sharing policy (i.e., types, duration, etc.) and incentives/promotion packages.
(6) Traditional approaches only allow to accept or reject a suggested offer. In the experiment, I appreciate being able to propose my own offer(s).
(7) Consider that a software agent negotiates with a company on my behalf after eliciting my privacy preferences. I will be confident that an agent can act on my behalf and help me in the negotiation process.
(8) Assessing my privacy preferences was a more challenging process than I thought.
(9) It does not make sense to me to negotiate on information sharing policies and incentives.
(10) My preferences on sharing information policy would be the same for any context. It does not matter whether it is telecommunications or healthcare or education.
(11) I would not let a software agent negotiate about information sharing policy on my behalf.
(12) My opponent made reasonable offers during the negotiation.
(13) My opponent took my privacy concerns into account.
(14) My opponent takes my previous offers into account while making its current offers.
(15) My opponent was not collaborative at all in finding a mutual agreement.

Fig. 21. The questions and average responses of post experiment survey.

essential for real-world deployment. In particular, future work should ensure that data users cannot re-share personal data with unauthorized third parties and incorporate system-level primitives and formal verification to protect the integrity and enforceability of negotiated agreements.

References

[1] Gönül Aycı, Arzucan Özgür, Murat Şensoy, and Pınar Yolum. 2023. Can we explain privacy? *IEEE Internet Computing* 27, 4 (2023), 75–80.
[2] Reyhan Aydoğan, David Festen, Koen V. Hindriks, and Catholijn M. Jonker. 2017. Alternating offers protocols for multilateral negotiation. In *Modern Approaches to Agent-Based Complex Automated Negotiation.* K. Fujita, Q. Bai, T. Ito, M. Zhang, F. Ren, R. Aydoğan, and R. Hadfi (Eds.). Springer, 153–167.

[3] Reyhan Aydoğan, Pinar Øzturk, and Yousef Razeghi. 2017. Negotiation for incentive driven privacy-preserving information sharing. In *International Conference on Principles and Practice of Multi-Agent Systems*. Springer, 486–494.
[4] Reyhan Aydoğan, David Festen, Koen V. Hindriks, and Catholijn M. Jonker. 2017. Alternating offers protocols for multilateral negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*. Springer, 153–167.
[5] Tim Baarslag, Alper T. Alan, Richard Gomer, Muddasser Alam, Charith Perera, Enrico H. Gerding, et al. 2017. An automated negotiation agent for permission management. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. 380–390.
[6] Markus C. Becker and ThorbjØrn Knudsen. 2005. The role of routines in reducing pervasive uncertainty. *Journal of Business Research* 58, 6 (2005), 746 – 757.
[7] Stefano Braghin, Marco Simioni, and Mathieu Sinn. 2022. DLPFS: The data leakage prevention filesystem. In *International Conference on Applied Cryptography and Network Security*. Springer, 380–397.
[8] Umut Çakan, M. Onur Keskin, and Reyhan Aydoğan. 2023. Effects of agent's embodiment in human-agent negotiations. In *Proceedings of the 23rd ACM International Conference on Intelligent Virtual Agents*. 1–8.
[9] Mary J. Culnan and Pamela K. Armstrong. 1999. Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science* 10, 1 (1999), 104–115.
[10] Enrique de la Hoz, Miguel Angel Lopez-Carmona, Mark Klein, and Ivan Marsa-Maestre. 2014. Alternative social welfare definitions for multiparty negotiation protocols. *Novel Insights in Agent-Based Complex Automated Negotiation* 535 (2014), 23–41.
[11] Daan Di Scala and Pınar Yolum. 2023. PACCART: Reinforcing trust in multiuser privacy agreement systems. arXiv:2302.13650. Retrieved from https://arxiv.org/abs//:2302.13650
[12] Khalil El-Khatib. 2003. A privacy negotiation protocol for web services. In *Workshop on Collaboration Agents: Autonomous Agents for Collaborative Environments*. Halifax, 85–92.
[13] Houda Ferradi, Jiannong Cao, Shan Jiang, Yinfeng Cao, and Divya Saxena. 2022. Security and privacy in big data sharing: State-of-the-art and research directions. arXiv:2210.09230. Retrieved from https://arxiv.org/abs//:2210.09230
[14] Dorota Filipczuk, Tim Baarslag, Enrico H. Gerding, and M. C. Schraefel. 2022. Automated privacy negotiations with preference uncertainty. *Autonomous Agents and Multi-Agent Systems* 36, 2 (2022), 49.
[15] Shelby Hunt and Scott Vitell. 2006. The general theory of marketing ethics: A revision and three questions. *Journal of Macromarketing* 26, 2 (2006), 143–153.
[16] Kangsoo Jung and Seog Park. 2019. Privacy bargaining with fairness: Privacy-price negotiation system for applying differential privacy in data market environments. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 1389–1394.
[17] Dilara Keküllüoğlu, Nadin Kökciyan, and Pınar Yolum. 2016. Strategies for privacy negotiation in online social networks. In *Proceedings of the 1st International Workshop on AI for Privacy and Security*. 1–8.
[18] Mehmet Onur Keskin, Umut Çakan, and Reyhan Aydoğan. 2021. Solver agent: Towards emotional and opponent-aware agent for human-robot negotiation. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*. 1557–1559.
[19] Nadin Kökciyan, Pınar Yolum, et al. 2022. Taking situation-based privacy decisions: Privacy assistants working with humans. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. 703–709.
[20] Gene R. Laczniak. 1983. Framework for analyzing marketing ethics. *Journal of Macromarketing* Spring (1983), 7–18.
[21] Kevin Liu, Harsh Desai, Lalana Kagal, and Murat Kantarcioglu. 2018. Enforceable data sharing agreements using smart contracts. arXiv:1804.10645. Retrieved from https://arxiv.org/abs//:1804.10645
[22] Gideon Ogunniye and Nadin Kökciyan. 2023. Contextual integrity for argumentation-based privacy reasoning. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 2253–2261.
[23] Gideon Ogunniye and Nadin Kokciyan. 2023. A survey on understanding and representing privacy requirements in the internet-of-things. *Journal of Artificial Intelligence Research* 76 (2023), 163–192.
[24] Sören Preibusch. 2005. Implementing privacy negotiation techniques in e-commerce. In *E-Commerce Technology, 2005. CEC 2005. Seventh IEEE International Conference on*. IEEE, 387–390.
[25] Ariel Rubinstein. 1982. Perfect equilibrium in a bargaining model. *Econometrica* 50, 1 (1982), 97–109.
[26] Ewa Topolewska, Ewa Skimina, Włodzimierz Strus, Jan Cieciuch, and Tomasz Rowiński. 2019. The short IPIP-BFM-20 questionnaire for measuring the big five. *Roczniki Psychologiczne* 17, 2 (2019), 385–402.
[27] Visara Urovi, Vikas Jaiman, Arno Angerer, and Michel Dumontier. 2022. Luce: A blockchain-based data sharing platform for monitoring data license accountability and compliance. *Blockchain: Research and Applications* 3, 4 (2022), 100102.
[28] Alan F. Westin and Danielle Maurici. 1998. *E-Commerce & Privacy: What Net Users Want*. Privacy & American Business Hackensack, NJ.

Received 8 January 2025; revised 31 May 2025; accepted 15 September 2025