

An Automated Negotiation Agent for Permission Management

Tim Baarslag

Centrum Wiskunde & Informatica
1098 XG Amsterdam
t.baarslag@cwi.nl

Alper T. Alan, Richard Gomer

University of Southampton
Southampton, SO17 1BJ
{a.t.alan,r.gomer}@soton.ac.uk

Muddasser Alam

University of Oxford
Oxford, OX1 2JD
moody@robots.ox.ac.uk

Charith Perera

The Open University
Milton Keynes, MK7 6AA
charith.perera@open.ac.uk

Enrico H. Gerdin,

m.c. schraefel

University of Southampton
Southampton, SO17 1BJ
{eg,mc}@ecs.soton.ac.uk

ABSTRACT

The digital economy is based on data sharing yet citizens have little control about how their personal data is being used. While data management during web and app-based use is already a challenge, as the Internet of Things (IoT) scales up, the number of devices accessing and requiring personal data will go beyond what a person can manually assess in terms of data access requests. Therefore, new approaches are needed for managing privacy preferences at scale and providing active consent around data sharing that can improve fidelity of operation in alignment with user intent. To address this challenge, we introduce a novel agent-based approach to negotiate the permission to exchange private data between users and services. Our agent negotiates based on learned preferences from actual users. To evaluate our agent-based approach, we developed an experimental tool to run on people's own smartphones, where users were asked to share their private, real data (e.g. photos, contacts, etc) under various conditions. The agent autonomously negotiates potential agreements for the user, which they can refine by manually continuing the negotiation. The agent learns from these interactions and updates the user model in subsequent interactions. We find that the agent is able to effectively capture the preferences and negotiate on the user's behalf but, surprisingly, does not reduce user engagement with the system. Understanding how interaction interplays with agent-based automation is a key component to successful deployment of negotiating agents in real-life settings and within the IoT context in particular.

Keywords

Automated negotiation; Negotiation agent; Privacy; Permissions; Mobile apps; Negotiation cost; Partial offers; Preference learning

1. INTRODUCTION

In the digital economy, our private information is increasingly being collected, shared and sold to third parties with little user control: users are asked to consent to personal data sharing transac-

tions by accepting privacy policies, which are almost never read, opaque, and lack any flexibility [2]. In recent years, an improved permission model has been introduced in smartphone apps, where users are able to permit access to certain types of data. A key challenge for this widely-adopted model, however, is a persistent lack of finely-tunable permission controls and clarity about the privacy trade-offs involved [48]. Even though individual permissions can be disabled, it is not clear how this affects the service if at all.

Multi-agent systems have been proposed for automating and negotiating privacy sharing decisions to make meaningful decisions on a user's behalf whilst minimizing the user burden (see Section 2 for a literature review). To date, this opportunity space has not been well-explored: there have been very few studies which propose practical automated negotiation solutions and none of these have been evaluated with real users using their real data. To this end, we address this gap by proposing a novel agent-based approach for negotiation of privacy permission management and by testing this approach with human participants using their actual data. This work sits within the wider agenda of privacy management that has received renewed momentum with the introduction of novel privacy laws (such as the EU's general data protection regulation, GDPR [15]), requiring greater transparency and user empowerment, and with opportunities for multi-agent systems to provide technological solutions.

In more detail, we design a novel negotiation strategy that makes optimal offers on behalf of the user with respect to the user's inferred utility function. A specific contribution here is the way in which the user's utility function is derived. Specifically, we first assess a user's privacy type by posing three key questions on privacy-awareness and then model the decisions made by this type (of users) from a separate baseline study. Finally, we establish a user preference model by computing a solution bounded by a set of linear inequalities that leads to a minimal number of constraint violations.

In addition, we introduce a novel alternating-offers, multi-issue negotiation protocol with costly quoting, in which users can propose partial offers (specifying values for some issues), and the service provider responds by proposing complete counter offers. We argue that such a protocol: 1) is more suitable for this domain, 2) allows for more collaborative exploration of the negotiation space to find mutually beneficial agreements, and 3) avoids distributive negotiation on single issues such as price. Using this protocol, the agent employs its strategy to negotiate autonomously on the user's

behalf. In a subsequent phase, the agent presents a potential agreement which the user can accept or manually override and continue the negotiation in order to further improve the offer.

To evaluate our approach, we develop an experimental tool which is installed on a user’s smartphone, allowing them to negotiate the sharing of data from their own mobile phone (e.g. contacts, messages, etc). In the scenario studied here, all of the value to the user from granting access to data is represented as a monetary reward, as an abstraction of the trade-off (usually in the form of a price discount) that accurately reflects real-world data exchanges. To be clear, the tool itself provides no service other than this reward, which aligns with our purpose to separate the issue of sharing information for the purpose of functionality, and for other uses such as targeted advertising. This allows us to study agent-supported negotiation with a greater degree of precision compared to previous studies, since our setup allows a clear trade-off between monetary reward and data stored on a personal phone.

The results of this agent-based negotiation approach are compared to a setting in which the user manually engages in a negotiation. From the user studies we find that the agent can accurately negotiate privacy decisions on behalf of the user without adversely affecting other key measures, such as privacy violations and user satisfaction. Surprisingly, however, the agent does not reduce the user bother. We discuss the implications of this for designing agent based systems for privacy negotiations.

The remainder of the paper is structured as follows. In Section 2 we discuss related approaches. In Section 3 we introduce the negotiation framework and Section 4 presents the agent. Then follows the design of our experimental tool and how the agent learns from previous interactions is discussed in Section 5. The experimental setup and results are presented in Sections 6 and 7. Finally, Section 8 concludes and discusses avenues for future work.

2. RELATED WORK

Various techniques have been used in privacy and permission negotiations, including rules based reasoning [47], game theory [59], question answer based profiling [20], direct user interaction [5], and learning through historical negotiations [25]. Table 1 summarizes the main privacy and trust negotiation approaches. These negotiations are conducted between consumers (e.g. web users) and service providers (e.g. web services). Ontologies are also commonly used to model privacy requirements, especially in more recent work [23, 24, 47] and to measure the privacy sensitivity of different pieces of data based on relationships to each other [23].

A notable milestone in online user privacy modeling and management is the Platform for Privacy Preferences (P3P), developed and recommended by the W3C. P3P is an XML mark-up language that allows websites to declare their intended use of information they collect about web browser users. The P3P Preference Exchange Language (APPEL) allows users to express their preferences, which can then be used by an agent to make semi-automated decisions regarding the acceptability of machine-readable privacy policies from P3P enabled Web sites. However, P3P/APPEL support did not get implemented widely due to lack of interest and use. Moreover, P3P is designed as a way to express privacy preferences but not as a negotiation framework. To address this, Bennicelli and Langendorfer [11] extended P3P/APPEL to include negotiation capabilities. In the negotiation process, they identify two main categories of user’s privacy requirements: *optional* and *mandatory* demands, which mirrors our negotiation protocol with *partial* and *complete* offers. Similarly, a classification of four such categories has been proposed by Kalyani and Adams [25], which can be assigned different weights so they can be effectively used in any

preference-based negotiation process such as ours, using a linear additive utility model.

To address the lack of negotiability of privacy, researchers have proposed different types of semi- and fully-automated privacy negotiation approaches. The inherent properties of multi-agent systems, such as intelligence and autonomy, make them an ideal intermediates for privacy negotiation between self-interested individuals. Existing works in this area utilize multi-agent systems to explore various aspects of privacy negotiation, ranging from information disclosure, secure data transfer to storage and trusted third party computation [49]. An agent can reveal its information based on policies, social relationships (disclosure based on trust and intimacy) and privacy-utility tradeoffs. Examples of such works using multi-agent systems include [28] and [50], which propose negotiation mechanisms to resolve policy conflicts among agents. Similarly, [30] and [60] propose mechanisms where users define their modular expected utility cost as a function of their private information. Our work falls into the privacy-utility tradeoff category, where an explicit model of the tradeoff governs each agent’s interaction with other agents in a given multi-agent system.

Our model of negotiation is an extension of the alternating offers protocol [36, 44] with partial offers (consistent with other negotiation models of partial offers, e.g. [43]) and, on top of that, costly completions. One of the main advantages of this approach is that an abundance of agents have been formulated for the alternating offers protocol [7, 13, 17, 18, 21, 26, 29, 58] which could be easily adapted to our model. Another approach is the model in [11], which consists of a negotiation process that includes five possible variants of settling an agreement. In most negotiation processes, these steps can be seen with minor differences [14, 25, 54, 59]. For example, Preibusch [42] has also proposed a privacy negotiation mechanism based on P3P, which allows multiple simultaneous offers to select from, instead of an alternating negotiation process.

A different strand of relevant work concerns the classification of privacy preferences. One classification identifies four types of users based on their privacy expectations [1, 46]: *privacy fundamentalists*, *profile averse*, *identity concerned*, and *marginally concerned*. Westin [57] has identified three similar main user types based on their attitudes and concerns about privacy: *fundamentalists*, *pragmatists*, and *unconcerned* (from most to least protective of their privacy). In this research, we have adopted Westin’s classification as it can be determined by an efficient three-question survey.

3. NEGOTIATION MODEL

The interaction between the agent and the service provider is governed by a negotiation protocol, which we first motivate and then introduce below.

3.1 A Permission Negotiation Protocol

We propose a novel variant of the well-known multi-issue alternating offer protocol [36, 44], in which offers are exchanged between two parties that specify values for each of the negotiable issues (in this case permissions). In our protocol, partial offers are submitted by the proposer (in this case the user or agent) specifying the requirements for only a subset of important issues. The responder (in this case the service provider) is then able to submit a complete counter offer based on the proposed offer. The proposer can either accept the complete offer or submit a new partial one. It can also break off the negotiation. Our premise is that, before the negotiation takes place, the service provider prescribes which additional issues make up a complete offer. This makes sure essential issues are under the service provider’s control (e.g. price, core functionality), and the service provider can thus always ‘price out’ any undesired

Model	Highlights of negotiation aspect
P3P/APPEL-based	Introduces a data model (extension of P3P) to describe privacy preferences and contracts [10] (2003), [35] (2005). Proposes automated negotiations algorithm for guaranteed termination and Pareto optimal results [32] (2004). Uses an extended P3P to describe alternative offers for the other party to choose from [41] (2005), [42] (2006). Proposes ‘Or Best Offer’-style privacy policy negotiation protocol with guaranteed termination [54] (2008).
Ontology-based	Uses ontologies to model common knowledge related to privacy and intellectual property laws [24] (2008). Proposes a data modeling technique to identify sensitivity levels of each data item [23] (2010).
Other	Uses user preferences and historic negotiation records to perform the negotiations [25] (2006). Uses XACML [37] as a policy description language [14] (2007). Derives a quantified privacy risk for each data item and uses it to determine an expected return [61, 59] (2008). Models privacy policies using rules. Users use weights to express the importance of each rule [51] (2011). Provides recommendations on which data items to share by learning sharing habits of similar users [20] (2015). Uses P2U [22] to define privacy policies. Proposes a bargaining model (as inspired by [33, 12]) [45] (2015). Describes Internet of Things protocols for enabling negotiable data access for future data marketplaces [39] (2016).

Table 1: Summary of related work.

partial offers. Such negotiations often occur in practice in a number of settings not necessarily limited to permissions management. For example, when negotiating insurance policies, buyers often specify certain conditions for the extent of cover, for which the seller completes the possible contract by proposing a price. Other examples include negotiating mortgages and broadband packages.

An advantage of this approach is that it prevents competitive, zero-sum negotiations on isolated issues (such as price), and instead promotes mutually beneficial deals. Moreover, the protocol allows users to focus on issues that are important to them and leave out less relevant issues, or issues for which users find it difficult to determine a precise value and are more naturally determined by the counter party. This is especially important in negotiating privacy permissions, because the benefits of privacy protection are often uncertain and intangible [3] and, as a result, users find it difficult to express the exact willingness to pay for revealing certain information. It is easier for consumers to decide, through *relative* comparisons, which of the complete offers they prefer in order to assess the value of protecting their privacy [52]. In this way, users can easily explore the set of possible agreements, while the service provider, provided with information about monetizing the data (e.g. through advertising), has the ability to exercise the final say and to calculate the value of a given combination of permissions.

3.2 Formal Negotiation Process

Fig. 1 shows an overview of the negotiation process. Formally, the negotiation domain is specified by m issues $I = \{1, \dots, m\}$ and a corresponding sets of possible values V_1, \dots, V_m . Then $\Omega = \prod_{i \in I} V_i$ is the set of all possible agreements/complete offers. For example, if the negotiable issues $I = \{1, \dots, 4\}$ correspond to *<Shared data, Purpose of sharing, Retention policy, Price discount>*, then an example agreement in Ω is *<GPS location, Targeted ads, Shared with third parties only, \$0,20>*.

The user or agent can propose a *partial* offer o over a subset $S \subseteq I$ of the negotiable issues; i.e., a value assignment $o \in \Omega|_S = \prod_{i \in S} V_i$. Upon receiving a partial offer o over the subset S , the service provider can complete o to a full offer $F(o) \in \Omega$ by supplying the values to the missing issues which we call a *quote*. That is, if we write $S = \{s_1, \dots, s_k\}$, then the service returns a complete quote $F(o)$, such that $o_i = F(o)_{s_i}$ for all $1 \leq i \leq k$. Continuing our example, the agent might not be aware of all privacy policy possibilities provided by the service and expect the service to supply values for *Retention policy* and *Price discount* by not having these issues included in S .

We assume that past quotes remain valid. Therefore, as the negotiation proceeds, the quotes received generate a growing set $Q \subseteq \Omega$

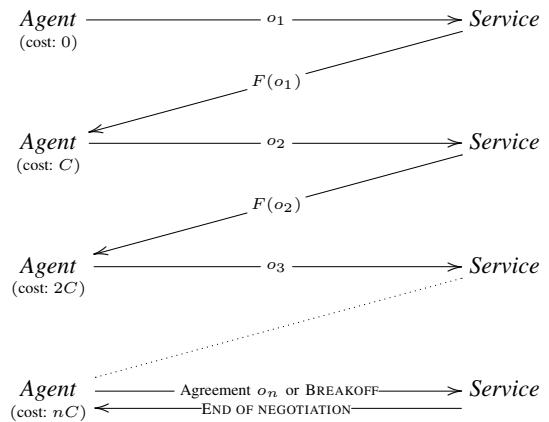


Figure 1: A depiction of the privacy negotiation process. Agent makes partial offers $o_i \in \Omega|_S$, and after every proposal, Service responds with a completed offer $F(o_i) \in \Omega$, incurring cost C . When the agent is satisfied with one of the possible agreements o_n , it is communicated to the service, and a deal is reached.

of possible outcomes that the agent can agree to. The negotiation ends when the agent either accepts a deal $q \in Q$ or actively ends the negotiation by signaling a break off.

In addition, we assume that every time the agent requests a quote, this incurs a cost C . This represents the service costs levied on the agent to process the request. The reasons for introducing these costs are twofold. First, this ensures that the agent does not simply explore all possible combinations, which is not realistic in many applications where time and quoting is costly. Second, and related to the first issue, it allows the service provider to gain surplus from the transaction by not revealing its complete cost structure [53].

4. THE NEGOTIATING AGENT

During negotiation the agent needs to determine, at any point in time, which of the partial offers, if any, to propose to the service, or which of the received quotes to accept. This is further complicated by the incrementally costly nature of the negotiation protocol: whether or not to make a new offer depends on the quoting cost, the expectation about future offers, and the attractiveness of the quotes received so far. To this end, we formulate a negotiation strategy that is optimal (with respect to expected utility according to the

agent's preference model), and which can engage with the service and specify when to end the negotiation process.

4.1 Calculating the Utility of an Offer

The agent's preferences are encoded by a value function $v : \Omega \rightarrow [0, 1]$ that maps every outcome to a utility value in the interval $[0, 1]$. Note that this function is only defined for *full* offers. Given this, if negotiation ends after requesting n quotes, the resulting utility is given by:

$$U(q, n) = \begin{cases} v(q) - nC & \text{if } q \in Q \text{ is accepted,} \\ r - nC & \text{if no agreement is reached.} \end{cases} \quad (1)$$

where $r \in [0, 1]$ is the reserve value of not reaching an agreement.

To simplify the learning task of the agent, in our experiments we will make the common assumption that the agent's utility function v is *linearly additive* [27]. That is, the utility $v(\omega)$ of an outcome $\omega = \langle \omega_1, \dots, \omega_m \rangle \in \Omega$ can be computed as a weighted sum from evaluation functions $e_i(\omega_i)$ as follows:

$$v(\omega) = \sum_{i=1}^m w_i \cdot e_i(\omega_i), \quad (2)$$

where $e_i(\omega_i)$ determines the value of the offer for the particular issue i , and w_i is the corresponding weight. The weights normalized such that $\sum w_i = 1$. In Section 5.2 we introduce an approach for deriving these weights based on the user's previous interactions.

4.2 Opponent Model

Now, to evaluate *partial* offers $o \in \Omega|_S$, we assume that, in addition to the user's utility function, the agent has a model of the likelihood of the completed counter offer it receives from the service. This depends on the negotiation strategy used by the opponent. In this paper we abstract away from modeling the actual strategy and, instead, assume that the probability of a complete offer given a partial offer o is given by a stochastic variable X_o with a cumulative distribution function (c.d.f.) $G_o(x)$, and we assume this function is known by the agent.

Using the previous example, if the agent makes a partial offer *<GPS location, Targeted ads, Shared with third parties only>*, then $G_o(x)$ would be the c.d.f. of the *Price discount*. Such a model could be constructed from prior knowledge or previous interactions with the service and can be based on the relative likelihood of the counter-proposal from the service (for an overview, see [8, 9]).

From this we can then derive the expected value of a partial offer. Specifically, we let $Y_o = v(X_o) \sim [0, 1]$ denote the stochastic variable describing the valuation, with corresponding c.d.f. $H_o(y)$. Given this, the expected value is $E[v|o] = \int_0^1 y H_o(y) dy$.

4.3 An Optimal Quoting Strategy

The agent's aim is to propose quotes with a high probability of a satisfactory completion without incurring high quoting costs. More formally, the aim is to find an optimal series of quotes that maximize the expectation over utility $U(q, n)$ w.r.t. G , using a minimal number of quoting requests, n . We use the approach from [6] to find the optimal quoting strategy which, as shown in [6], can be mapped onto a variant of the so-called Pandora's problem [56] (a search problem involving boxes that contain a stochastic reward).

Specifically, each partial offer o in $S|_\Omega$ can be regarded as a *closed* box with stochastic reward Y_o that can be opened at cost C , while every quote $q \in Q$ together with the break-off value can be represented by an *open* box with known reward (i.e., $v(q)$ and r respectively). As a consequence of Pandora's Rule [56] we can, for

every partial offer $o \in S|_\Omega$, assign an index z_o , satisfying

$$\int_{z_o}^{\infty} (y - z) dH_o(y) = C \quad (3)$$

that fully specifies when to propose it: namely, when it has the highest index in $S|_\Omega$ and exceeds the value of the quotes in Q so far. This procedure provides the agent with the following strategy:

Optimal Quoting Strategy. *Propose a partial offer $o \in \Omega|_S$ with the highest index z_o (as defined by eq. (3)) if it is higher than*

$$\max_{q \in Q} (v(q), r).$$

Otherwise, if $v(q)$ has the highest value for some $q \in Q$, offer q as the agreement. Otherwise, r is the highest value, and the agent should break off the negotiation.

It is shown in [56] that such a strategy is optimal in terms of maximized expected utility minus cumulative costs. In practice, the effectiveness of the optimal quoting strategy depends on the accuracy of the utility model v and the service model X_o ; however, with a faithful model, the agent's quoting strategy is optimal in a non-myopic sense: it will generate quotes taking into account not only the quoting costs, but also the incremental effect of any subsequent quoting taking place.

By employing this quoting strategy, the agent can make optimal trade-offs between the agreement utility and the sum of incurred quoting costs. We illustrate and test our strategy in the next section, by applying it to the domain of mobile permissions.

5. NEGOTIATING PERMISSIONS

To evaluate the negotiation protocol and the agent with real users and their actual personal data, we created a tool in the form of an Android app. Using this tool, both the agent and the user can negotiate permissions to access data on the user's mobile phone, in return for a monetary reward. We deliberately omit any other app functionality to avoid any personal preference towards the service and to focus on privacy preferences. Hence, the app's main function is to negotiate and collect user data. In what follows we first describe the tool and how the app interacts with the user, and then we explain how the agent derives the user's preferences (utility function) from previous interactions.

5.1 Interaction Design

The agent initially negotiates on the user's behalf using the approach detailed in Section 4. Once this is completed (i.e. Pandora's algorithm terminates), the user is presented with the best offer through a default setting screen (see Figure 2a), which shows the permissions and corresponding reward in the form of points (which directly translate into a monetary reward as explained in Section 6). The user can then accept this offer, or continue to negotiate manually by changing the settings and pressing *Quote*. This will send the partial offer to the service provider, who will complete the offer by responding with a specific number of points. Users can request multiple quotes and a quoting cost is levied for each request. In addition, a user can freely access previously received offers by pressing *Prev* and *Next*. Users can also choose not to share any data by setting all permissions to *Don't Share* in which case the reward is set to zero which can be accepted without pressing *Quote*.

Once a user accepts, the app randomly collects some data points from each shared data type. The user is then presented with a *review screen* (Figure 2b), which shows exactly what data is shared and asks the user to retrospectively express whether they are *Happy*

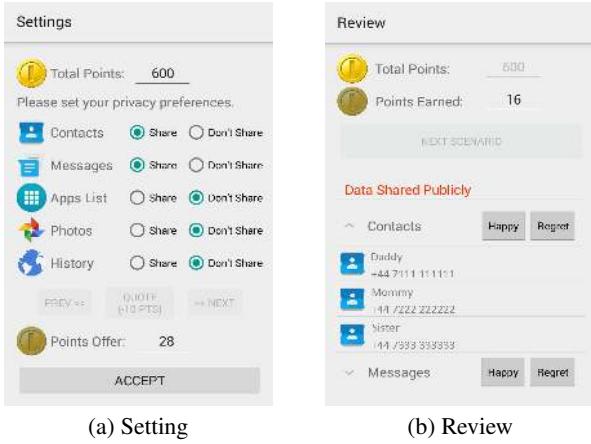


Figure 2: The design of our tool for data control negotiation: (a) where users can modify permissions and are offered different point quotes, and (b) where users can revise their personal data shared, and express their feeling of sharing each data type.

about or *Regret* sharing a specific data type. The purpose of this screen is to make users *aware* of their decisions, and also to *evaluate* the outcome of the negotiation.

5.2 Learning the Utility Function

In order for an agent to faithfully represent the user, it is important that the agent’s value function v (and thereby, the overall utility function U) has to be aligned with the user’s preferences. For this, we introduce a novel approach for deriving the valuation from past user interactions when they negotiate permissions manually (using the same tool but setting the initial/default screen to random). In more detail, we divide users into three categories based on their actual sharing behavior (see Section 6.5 for details): *Fundamentalists*, *Pragmatists*, and *Unconcerned*. Then, assuming a linear valuation function v (Equation 2), for each user category, we find the weights w_i for each issue by considering the rejected and accepted offers by users in that category as described in the following.

In more detail, the set of issues are represented by:

$$I = \langle x_1, \dots, x_m, r \rangle,$$

where every $x_i \in \{0, 1\}$ is a permission with binary values 0 (*Don’t Share*) or 1 (*Share*) and r is the *Reward* which is a continuous issue representing the points received. If we then normalize reward r such that $r \in [0, 1]$, the valuation is given by:

$$v(o) = r - x_1 w_1 - \dots - x_m w_m. \quad (4)$$

Given this, if a user rejects an offer o (including the setting of sharing nothing) in favor of accepting o' , we can assume that $v(o') \geq v(o)$. This can be written as:

$$(x'_1 - x_1)w_1 + \dots + (x'_m - x_m)w_m \leq r' - r.$$

If we do this for all rejected and accepted offers by all users in the particular category, we obtain a set of inequalities from which we can deduce the most appropriate overall weights w_i . To this end, note that this procedure transposes the problem into a set of linear inequalities of the form:

$$Aw \leq b,$$

where the entries of A and b correspond, for each equation, to the values of $x'_i - x_i \in \{-1, 0, 1\}$ and $b = r' - r$ respectively.

However, as this data stems from human interaction, and since the problem quickly becomes over-constrained, we find that these inequalities are typically not consistent. Therefore, we cannot simply use standard linear constraint solvers. To address this, instead, we find a solution that best satisfies the constraints following the techniques described in [40]. Specifically, we determine the weights w^* that minimize the least squares norm:

$$w^* = \arg \min_w \|(Aw - b)_+\|^2,$$

where $(Aw - b)_+$ is the vector whose i th component equals

$$\max\{(Aw - b)_i, 0\}.$$

6. EVALUATION

To test our agent in a real privacy sensitive situation, we conducted a lab study in which we asked participants to select permissions that give access to personal data on their smartphones. In an effort to make the experiments as realistic as possible, we analyze users’ responses to exposing their personal privacy-sensitive information available through their own phone, combined with real monetary incentives to share their private data online.

Our main goal is to explore whether the negotiation agent we described above can effectively and understandably perform permission negotiations on behalf of a user. In particular, the agent’s aim is not to persuade or influence the user’s decision making, but to select and follow a quoting strategy that is accurate; i.e., close to the user’s wishes.

6.1 Experimental Design

The agent negotiates in the background and then presents the agreement to the user, enabling the user to override the decision if necessary (i.e., called flexible autonomy, which other studies have shown to be important [4]). It is important to note that the user is not aware of (and is not told about) the agent and simply perceives the outcome of the negotiation as a default setting. Such default settings as a design tool have been well studied [55]. This provides the ultimate litmus test for the effectiveness of the agent. To evaluate this, we introduced an additional treatment where the default settings are chosen randomly, and compared this to the treatment with the agents. In what follows we refer to these two designs as the *Random* and *Agent* designs respectively.

We employed a between-participants design in which the participants were divided between the *Random* and the *Agent* treatment. We maximized direct comparability between the groups through a matched subjects design with exactly the same number of privacy types in both treatments (as explained in more detail in Section 6.5). In addition, we tested 8 different negotiation scenarios (detailed below in Section 6.4) in which we varied the number of points that the user would receive. All users participated in each scenario, enabling the agent to learn from previous interactions.

6.2 Participants

We recruited 66 participants (20 female and 46 male) from the university, targeting student groups through invitation flyers and social media outreach. Participants were undergraduate, masters or PhD students from a variety of disciplines (e.g., Engineering, Languages, Business and Management, Law, Health and Social Sciences, and Geography). Their age ranged from 17 to 35 (Mean: 21.3, SD: 3.4).

6.3 Permissions

We focused on the five most often used permissions (as ranked by [34]) that can be acquired and mined from users’ smartphones.



Figure 3: A participant in the lab study, negotiating data access on her own smartphone.

These are: access to all *Contacts*, *Messages*, *Apps*, *Photos*, and *Browsing history* stored on the smartphone (as detailed in Figure 2a). Whenever access is granted, three randomly-selected, unique data points from each data type are collected after each negotiation scenario. The five permissions result in 32 (2^5) possible partial offers in $\Omega|_S$. We set the quoting cost C to 10 points, while no sharing results in no reward, and so that the reserve value is $r = 0$.

6.4 Reward Scenarios and Opponent Strategy

To explore varying reward levels, each participant engaged in 8 negotiations with different scenarios with a low, medium and high maximum reward of 25, 50, and 100 points respectively (every 100 points corresponds to £1). For any partial offer o , the stochastic completion X_o is defined by the possible points offered by the opponent, thereby inducing Y_o and the corresponding c.d.f. H (see Section 4.2). Depending on the number of enabled permissions N and the maximum reward $M \in \{25, 50, 100\}$, the opponent completed the quote by offering a uniformly random number of points between $\max(0, M(N-1)/5)$ and $M \cdot N/5$. This approach ensures that share more permissions will result in higher reward. We used a balanced Latin square design to determine the order of the scenarios to cancel out interaction effects. The full details of the scenarios are available separately (see Acknowledgement).

6.5 User Privacy Types

We derived each user's initial *privacy type* from a 3-question survey designed by Westin [57] that measures attitudes and concerns about privacy. The privacy type classifies the user into three categories, which, from most to least protective of their privacy are: *Fundamentalists*, *Pragmatists*, and *Unconcerned*.

However, it is well known that people's declared attitudes towards privacy have a limited effect on their actual behavior (this phenomenon is called the *privacy paradox*). Indeed, consistent with other works, we found little predictive value from Westin's categorization into the three canonical privacy type [16]. Therefore, we only use the reported type to mitigate the cold start problem of the first scenario. After that, we use a variant of Westin's three classes based on behavior, in which we classify the participants by their sharing actions. Specifically, after every scenario, we count the number of items that were shared without regret. Users are classified as *Behavioral Fundamentalists* if this amounts to less than 33% of the total, *Behavioral Pragmatists* between 33%-66%, and *Behavioral Unconcerned* otherwise. For example: suppose a user has just entered scenario number 4 and has so far shared 2

text messages and 3 photos out of 15 possible items. Assuming the user regretted sharing both text messages, their score would be $\frac{5-2}{15} = 20\%$ and hence they would be assigned the privacy profile associated with *Behavioral Fundamentalists*.

6.6 The Agent

The agent uses the optimal quoting strategy described in Section 4.3 to select an optimal default setting for the user while minimizing cost. Note that the agent is not aware of any of the quotes before they are requested; it makes its decision based solely on the user's behavioral privacy type and its stochastic utility model.

To establish the priors of its utility model, the agent classified all users in the *Random* treatment (sample size: 33) according to their privacy type and learned the weights of each permission for each type of user using the techniques described in Section 5.2. In this way, the permission weights of the user utility model are learned offline, but the user classification into a privacy type is performed online, using the technique described in the previous section.

6.7 Procedure

Each experiment started with a set of instructions explaining the features of the app and the experiment procedure. Specifically, participants were informed that they will receive a cash payment at the end, based on their total points, and that the more data they share, the more points they gain. It was emphasized that any data they shared would be made publicly accessible on a website.

Participants installed the app from Google's Play Store and entered basic demographics data (e.g., age, gender and department), together with answers to the three-question privacy survey from Westin to derive the participant's reported privacy type. Once the participants were assigned to an experimental condition, the actual study started, in which participants were required to make permission decisions and review the shared data alternately for eight rounds (Figure 3).

We collected both quantitative and qualitative data through a number of techniques. User interactions with the app were automatically recorded by measuring the following dependent variables throughout the experiment: *the number of quotes* and *changes to the default settings* that participants made in each scenario, the *average time taken* to complete the scenarios, and the *final reward* received by participants. Afterwards, participants were asked to complete a post experimental questionnaire for measuring users' experience with the app and their sensitivity about sharing each type of data. We also administered a NASA Task Load Index (NASA-TLX) assessment [19] for evaluating the workload perceived by users in the course of the experiment.

We conducted exit interviews with 11 participants, who were selected based on their privacy attitudes and patterns in their interactions with the app, such as those who often (or seldom) made changes to the default settings. The interviews were audio-recorded, and later partially transcribed based on key questions that aimed to clarify their actions during the sessions.

At the end of the study, participants received a reward in cash, varying between £5 and £10, based on the number of points that they earned in the eight scenarios.

7. RESULTS

During our experiments, participants shared 3090 units of data (e.g. contacts or text messages) sampled from their smartphones, out of a total of 343709 available items. Our participant pool consisted of 15% Fundamentalists, 79% Pragmatists, and 6% Unconcerned (reported type), which is broadly consistent with the overall American public [31]. We present our analysis below.

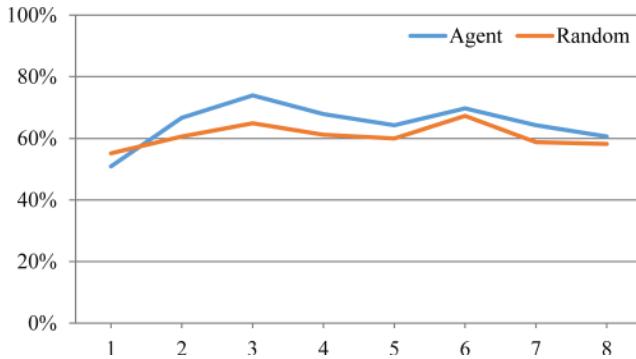


Figure 4: Accuracy of the default setting in random and agent treatments in each scenario.

	Contacts	Messages	Apps	Photos	Browser	
Random	r	-0.79	-0.78	-0.43	-0.72	-0.26
	p	<0.01	<0.01	0.012	<0.01	0.142
Agent	r	-0.59	-0.76	-0.45	-0.73	-0.39
	p	<0.01	<0.01	<0.01	<0.01	0.024

Table 2: Results of correlation analysis between reported sensitivity and actual sharing behavior

7.1 User Understanding

It was important that participants acted based on genuine privacy concerns, and so we informally assessed participant understanding of the study methodology and the app itself by asking them to explain the purpose of the points, and what would happen to the data that we collected. We checked whether participants believed our claim that the data they shared would be published, and if they understood the link between the number of points that they were awarded and the amount of money that they would receive.

All of the interviewed participants consistently explained, in their own words, the link between points and final payment. Participants' reluctance to share data despite the tangible reward for doing so, coupled with their responses to the questions in the interview phase suggest that the scenario was successful in creating a genuine belief that their decision to share would have real privacy implications.

7.2 Agent Accuracy

We calculated the accuracy of the agent's choices by comparing the number of changes that the users made to the default settings for sharing each data type in each scenario. Figure 4 shows that in all scenarios except the first one (when the agent is still relying on Westin's *reported* privacy type), the users made less changes to the default settings in the Agent treatment. In the following rounds, the agent's choices for the default settings were significantly more accurate at accommodating the users' privacy preferences than the randomly assigned default settings (*one-tailed t-test*, $p=0.013$).

Our experimental setup, in which participants were not aware of the agent, allows us to be confident that this accuracy is the result of correctly predicting preference, rather than a tendency of participants to "go along with" suggestions that they know are made by an agent. Although we detect some bias resulting from the defaults, this is apparent in both conditions. This is because the *Random* treatment serves as a fair base case, as any default bias influ-

ences both treatments equally: since defaults were set randomly, we would expect them to be aligned with user preferences 50% of the time; in fact, the slightly higher percentage of around 60% show that the defaults exert some influence and act as a means to promote exploration of the different options.

This result is further supported by the analysis of the subjective answers given by the users for one of our questionnaire statements, reading: "Overall, the default settings on what data to share were appropriate for me". The users of agent treatment were significantly more content with the default settings (Mdn:5, *one-tailed t-test*, $p=0.01$), than the users who were exposed to the random default settings (Mdn:2).

The interviews provide qualitative accounts from some participants about how they felt about the defaults. In particular, participants gave explanations as to how the defaults had influenced their choices. Some suggested that the defaults had led to voluntary exploratory behavior, for instance **P133** who said:

"On one, photos was preticked and the points were a lot higher, so I assumed that meant that photos were worth more."

Others reported that the defaults led to them making mistakes, suggesting (for instance) that in some rounds they had forgotten to alter some of the settings before accepting the quote.

We wanted to see if participants had noticed – or expected – the presence of an agent in the experiment. We asked them whether the default settings in each scenario were similar to their preferences, and whether they felt the defaults had improved during the experiment. None said they noticed increasing accuracy during the session, and they typically felt that the chosen defaults were of mixed accuracy, for instance **P291** said that:

"They [the defaults] were not too different to my preferences ... but they were very varied and changed every time."

It was apparent during the interviews that participants did not infer the existence of a learning process within the app, but that they were uncertain about the provenance of the defaults. Although some, such as P133, did express a belief that the defaults might somehow be linked to the value of the data. This was not restricted to those participants in the agent condition (where this was relationship was true) - P133 was themselves a participant in the manual negotiation condition, where the defaults were not linked to payoff.

7.3 User Behavior and Experience

We found that the agent does not introduce any new biases to user behavior or changes in user experience, which suggests that the agent was successful in inconspicuously assisting the user in picking the right options. Our overall workload evaluation through NASA-TLX Raw did not reveal any significantly different results for the random and agent-suggested default settings on key workload indicators such as *mental demand*, *temporal demand*, *performance*, *frustration level*, *effort*, and *overall time taken*. Furthermore, the users of agent-suggested settings earned an amount of points (and hence, received an amount of money) that was not significantly different from the random setting users.

We were at first surprised to see no less overall workload for the agent treatment, but some reflection on the design principles of the agent helps to understand this. Due to our setup that provides overriding control to the user, the negotiation is far from fully automated: users are unaware of the agent and still tend to check and amend every proposed setting. We anticipate very different results in a fully automated setting (see also our discussion in Section 8).

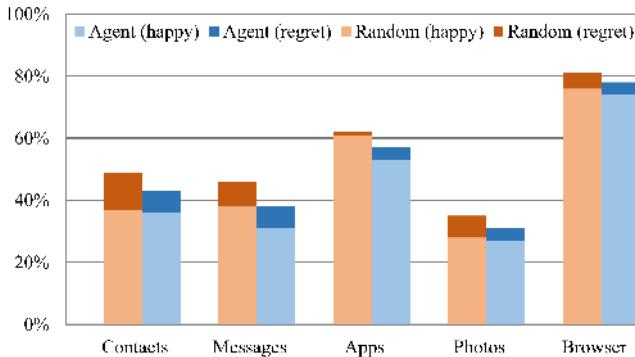


Figure 5: Sharing percentages for every permission in both treatments, together with the relative percentage of the sharing decisions for which the user felt, upon reflection, happiness or regret.

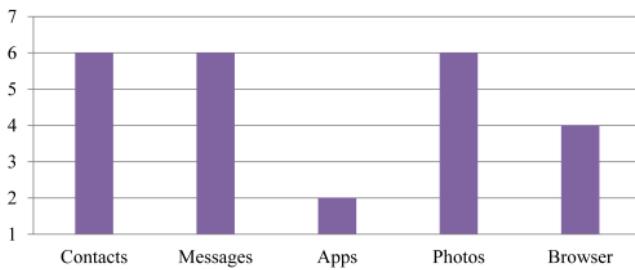


Figure 6: Medians of the self-reported sensitivity scores for sharing different data types, ranging from 1 (very low), to 7 (very high).

Similarly, the sharing percentages for every permission and retrospective feelings (happy/regret) were comparable in both treatments, and aligned with their preferences. Figure 5 shows the overall sharing choices that the users made during the experiment. In both cases, users were least concerned about sharing their browsing history, while photos were shared the least. Figure 6 displays the medians of the all users' reported sensitivity for each data type. When we compare the reported sensitivity of each data type with users' sharing decisions depicted in Figure 5, we can clearly see that the behavior of the users in both groups are aligned with their privacy preferences. Using the correlation between participants' choices and preferences through the Pearson's correlation coefficient (r), we obtain a negative correlation for all data types from the users of both treatments (see Table 2). The correlations are mostly very strong except for browser and apps, which indicates that users were generally able to act out their wishes in both treatments.

The only notable difference in user behavior induced by the agent is in the number of quotes requested. In the agent treatment, users made 223 quotes throughout the experiment ($M:6.75$, $SD:2.16$), while the random treatment users quoted 194 times in total ($M:5.87$, $SD:3.13$). The qualitative interview results with respect to the defaults suggest that defaults provide a means of prompting *exploration* of offers that the user had not previously considered. Importantly, participants did not necessarily regret the decisions that had been influenced by the defaults, and this seems to be an important finding in itself. There is perhaps a risk that a highly "accurate" agent, given the user's own uncertainty about their preferences, might in effect stifle exploration and reflection. Conceptually,

this is similar to filter bubble effect [38] observed on social media, whereby users are disproportionately exposed to views that they already agree with. Except that, in this case, users are only presented with defaults that match their existing behavior. This may be coined the 'padded room effect', in which mechanisms intended to decrease discomfort or improve safety, actually prevent exploration and inhibit potentially beneficial preference change.

8. CONCLUSIONS AND FUTURE WORK

In an ever more connected world, there is a pressing need for automating privacy negotiation that can make meaningful decisions on the user's behalf while minimizing their burden. Our work is a first step towards agent-supported permission negotiation, with results derived from a user study using actual, private data.

We present a novel agent-based approach that is able to automatically negotiate between users and services, while optimally balancing between agreement utility and the sum of incurred quoting costs. We show through a proof-of-concept interaction design that an agent can accurately automate privacy decisions on human user behalf in line with normal user behavior, while learning from feedback during a review phase. Moreover, our interaction design leaves control fully into the hands of the user by allowing them to override the agent decision – an important feature in a privacy setting with invariably sensitive personal data.

In particular, our results show that the deals negotiated by the agents are more accurate than our baseline in that the resulting agreements are better aligned with the user's actual preferences. Although the agent's deployment does not result in less effort from the user, we hypothesize this is partly a consequence of its unobtrusive design (i.e. users were unaware of the agent). Our hypothesis is that, if the user is made aware of an agent working on their behalf and builds trust over time, the user will be inclined to concede more autonomy. Also, it might be important for the agent to articulate why it has reached a particular decision – for instance that the payoff for sharing a particular data type was very low.

These results provide useful insights and lessons for the design of effective agents for automated consenting decisions and point to several avenues of future work. In particular, while the proposed agent is quite general, our experimental setting is limited to reasoning about data types. Other issues that need to be considered during negotiation are, for example, the recipient, retention period, purpose, quality, and privacy risks. In addition, we noticed from the interviews that when expressing regret, users often did so for specific data points (e.g. specific contacts or photos). Therefore, it is clear that a model based on permissions alone is too coarse to accurately capture the privacy preferences. Combining a semi-autonomous agent with a more meaningful classification of data (perhaps using signals such as location, time of day, and relation to other people) is another avenue that warrants further exploration. Finally, the agent's user model was based on three privacy types. A more personalized model, derived from e.g. apps installed on the phone and other indicators, could increase the accuracy of the deal negotiated by the agent even further.

Acknowledgement

This work is supported by the EPSRC Meaningful Consent in the Digital Economy project (EP/K039989/1) and is funded through the ERA-Net Smart Grids Plus project Grid-Friends, with support from the European Union's Horizon 2020 research and innovation programme. Data and scenarios are available at: <http://doi.org/10.5258/SOTON/405394>. Study approved by University of Southampton FPSE Ethics Committee (ref: 18082).

REFERENCES

[1] Mark S. Ackerman, Lorrie Faith Cranor, and Joseph Reagle. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, EC '99, pages 1–8, New York, NY, USA, 1999. ACM.

[2] Alessandro Acquisti and Jens Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 2(2005):24–30, 2005.

[3] Alessandro Acquisti and Jens Grossklags. Uncertainty, ambiguity and privacy. In *Fourth workshop on the economics of information security (WEIS 2005)*, pages 2–3, 2005.

[4] Alper Alan, Enrico Costanza, Joel Fischer, Sarvapali D. Ramchurn, Tom Rodden, and Nicholas R. Jennings. A field study of human-agent interaction for electricity tariff switching. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 965–972, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

[5] Tim Baarslag, Alper T. Alan, Richard C. Gomer, Ilaria Liccardi, Helia Marreiros, Enrico H. Gerding, and M.C. Schraefel. Negotiation as an interaction mechanism for deciding app permissions. In *Proceedings of the 2016 CHI Conference: Extended Abstracts on Human Factors in Computing Systems*, CHI EA '16, pages 2012–2019, New York, NY, USA, 2016. ACM.

[6] Tim Baarslag and Enrico H. Gerding. Optimal incremental preference elicitation during negotiation. In *Proceedings of the Twenty-fourth International Joint Conference on Artificial Intelligence*, IJCAI'15, pages 3–9. AAAI Press, 2015.

[7] Tim Baarslag, Enrico H. Gerding, Reyhan Aydoğan, and M.C. Schraefel. Optimal negotiation decision functions in time-sensitive domains. In *2015 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT)*, volume 2, pages 190–197, Dec 2015.

[8] Tim Baarslag, Mark J.C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. Measuring the performance of online opponent models in automated bilateral negotiation. In Michael Thielscher and Dongmo Zhang, editors, *AI 2012: Advances in Artificial Intelligence*, volume 7691 of *Lecture Notes in Computer Science*, pages 1–14. Springer Berlin Heidelberg, 2012.

[9] Tim Baarslag, Mark J.C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, 30(5):849–898, 2016.

[10] M. Bennicke and P. Langendorfer. Towards automatic negotiation of privacy contracts for internet services. In *The 11th IEEE International Conference on Networks, 2003. ICON2003.*, pages 319–324. IEEE, 2003.

[11] M. Bennicke and P. Langendorfer. Towards automatic negotiation of privacy contracts for internet services. In *The 11th IEEE International Conference on Networks, 2003. ICON2003.*, pages 319–324. IEEE, 2003.

[12] Elisa Burato and Matteo Cristani. The process of reaching agreement in meaning negotiation. In Ngoc Thanh Nguyen, editor, *Transactions on Computational Collective Intelligence VII*, pages 1–42. Springer-Verlag, Berlin, Heidelberg, 2012.

[13] Siqi Chen, Haitham Bou Ammar, Karl Tuyls, and Gerhard Weiss. Optimizing complex automated negotiation using sparse pseudo-input gaussian processes. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 707–714, Richland, SC, 2013. International Foundation for Autonomous Agents and Multiagent Systems.

[14] Vivying Cheng, Patrick Hung, and Dickson Chiu. Enabling Web Services Policy Negotiation with Privacy preserved using XACML. In *2007 40th Annual Hawaii International Conference on System Sciences (HICSS'07)*, pages 33–33. IEEE, 2007.

[15] EC European Commission et al. Proposal for a regulation of the european parliament and of the council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (general data protection regulation). *COM (2012) 11 final, 2012/0011 (COD)*, Brussels, 25 January 2012, 2012.

[16] Janna Lynn Dupree, Richard Devries, Daniel M. Berry, and Edward Lank. Privacy personas: Clustering users via attitudes and behaviors toward security practices. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 5228–5239, New York, NY, USA, 2016. ACM.

[17] Shaheen S. Fatima, Michael J. Wooldridge, and Nicholas R. Jennings. Multi-issue negotiation under time constraints. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 143–150, New York, NY, USA, 2002. ACM.

[18] Jianye Hao and Ho-fung Leung. CUHK agent: an adaptive negotiation strategy for bilateral negotiations over multiple items. In Ivan Marsa-Maestre, Miguel A. Lopez-Carmona, Takayuki Ito, Minjie Zhang, Quan Bai, and Katsuhide Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 171–179. Springer, Japan, 2014.

[19] Sandra G Hart. NASA-task load index (NASA-TLX); 20 years later. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 50, pages 904–908. Sage Publications, 2006.

[20] Luke Hutton and Tristan Henderson. “I didn’t sign up for this!”: Informed consent in social network research. In *International AAAI Conference on Web and Social Media*, 2015.

[21] L. Ilany and Y. (K.) Gal. The simple-meta agent. In I. Marsa-Maestre, M.A. Lopez-Carmona, T. Ito, M. Zhang, Q. Bai, and K. Fujita, editors, *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535 of *Studies in Computational Intelligence*, pages 197–200. Springer, Japan, 2014.

[22] Johnson Iyilade and Julita Vassileva. P2U: A Privacy Policy Specification Language for Secondary Data Sharing and Usage. In *2014 IEEE Security and Privacy Workshops*, pages 18–22. IEEE, may 2014.

[23] I. Jang and H. S. Yoo. Personal information classification for privacy negotiation. In *Computer Sciences and Convergence Information Technology, 2009. ICCIT '09. Fourth International Conference on*, pages 1117–1122, Nov 2009.

[24] I. J. Jang, W. Shi, and H. S. Yoo. Policy negotiation system architecture for privacy protection. In *Networked Computing and Advanced Information Management, 2008. NCM '08*.

Fourth International Conference on, volume 2, pages 592–597, Sept 2008.

[25] Y. Kalyani and C. Adams. Privacy negotiation using a mobile agent. In *2006 Canadian Conference on Electrical and Computer Engineering*, pages 628–633, May 2006.

[26] Shogo Kawaguchi, Katsuhide Fujita, and Takayuki Ito. Compromising strategy based on estimated maximum utility for automated negotiating agents. In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations, Series of Studies in Computational Intelligence*, pages 137–144, Berlin, Heidelberg, 2012. Springer-Verlag.

[27] Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives*. Cambridge University Press, 1976.

[28] Nadin Kokciyan. Privacy management in agent-based social networks. In *AAAI Conference on Artificial Intelligence*, 2016.

[29] Sarit Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, Oct 2001.

[30] Andreas Krause and Eric Horvitz. A utility-theoretic approach to privacy in online services. *Journal of Artificial Intelligence Research*, 39:633–662, 2010.

[31] Ponnurangam Kumaraguru and Lorrie Faith Cranor. Privacy indexes: a survey of westin’s studies. *ISRI Technical Report*, 2005.

[32] Haifei Li, David Ahn, and Patrick C. K. Hung. Algorithms for automated negotiations and their applications in information privacy. In *Proceedings of the IEEE International Conference on E-Commerce Technology*, CEC ’04, pages 255–262, Washington, DC, USA, 2004. IEEE Computer Society.

[33] Yanhuang Li, Nora Cuppens-Boulahia, Jean-Michel Crom, Frederic Cuppens, and Vincent Frey. Reaching Agreement in Security Policy Negotiation. In *2014 IEEE 13th International Conference on Trust, Security and Privacy in Computing and Communications*, pages 98–105. IEEE, sep 2014.

[34] Ilaria Liccardi, Joseph Pato, and Daniel J Weitzner. Improving user choice through better mobile apps transparency and permissions analysis. *Journal of Privacy and Confidentiality*, 5(2):1, 2014.

[35] M. Maaser and P. Langendoerfer. Automated negotiation of privacy contracts. In *29th Annual International Computer Software and Applications Conference (COMPSAC’05)*, volume 1, pages 505–510 Vol. 2, July 2005.

[36] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. The MIT Press, 1st edition, 1994.

[37] Bill Parducci and Hal Lockhart. eXtensible Access Control Markup Language (XACML) Version 3.0. *OASIS Standard*, (January):1–154, 2013.

[38] Eli Pariser. *The filter bubble: What the Internet is hiding from you*. Penguin UK, 2011.

[39] Charith Perera, Susan Y. L. Wakenshaw, Tim Baarslag, Hamed Haddadi, Arosha K. Bandara, Richard Mortier, Andy Crabtree, Irene C. L. Ng, Derek McAuley, and Jon Crowcroft. Valorising the IoT databox: creating value for everyone. *Transactions on Emerging Telecommunications Technologies*, 2016.

[40] Elena Popescu. On the approximation of inconsistent inequality systems. *Analele Științifice ale Universității Ovidius*, 11(2):109–118, 2003.

[41] S. Preibusch. Implementing privacy negotiation techniques in e-commerce. In *Seventh IEEE International Conference on E-Commerce Technology (CEC’05)*, pages 387–390, July 2005.

[42] SÃ‰ren Preibusch. Privacy negotiations with p3p. In *W3C Workshop on Languages for Privacy Policy Negotiation and Semantics-Driven Enforcement*, 2006, 2006.

[43] Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus. Negochat: A chat-based negotiation agent. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS ’14, pages 525–532, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

[44] Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

[45] S. Sadki and H. El Bakkali. An approach for privacy policies negotiation in mobile health-cloud environments. In *Cloud Technologies and Applications (CloudTech), 2015 International Conference on*, pages 1–6, June 2015.

[46] Sarah Spiekermann. Online information search with electronic agents: drivers, impediments, and privacy issues, 2001. Zsfassung in dt. Sprache.

[47] A. C. Squicciarini, E. Bertino, E. Ferrari, and I. Ray. Achieving privacy in trust negotiations with an ontology-based approach. *IEEE Transactions on Dependable and Secure Computing*, 3(1):13–30, Jan 2006.

[48] C. Stach and B. Mitschang. Privacy management for mobile platforms – a review of concepts and approaches. In *2013 IEEE 14th International Conference on Mobile Data Management*, volume 1, pages 305–313, 2013.

[49] Jose M Such, Agustín Espinosa, and Ana García-Fornes. A survey of privacy in multi-agent systems. *The Knowledge Engineering Review*, 29(03):314–344, 2014.

[50] Jose M. Such and Michael Rovatsos. Privacy policy negotiation in social media. *ACM Trans. Auton. Adapt. Syst.*, 11(1):4:1–4:29, February 2016.

[51] Salah-Eddine Tbahriti, Brahim Medjahed, Zaki Malik, Chirine Ghedira, and Michael Mrissa. Meerkat - A Dynamic Privacy Framework for Web Services. In *2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, volume 1, pages 418–421. IEEE, aug 2011.

[52] Janice Y Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22(2):254–268, 2011.

[53] Hal R. Varian. *Economic Aspects of Personal Privacy*, pages 101–109. Springer US, Boston, MA, 2009.

[54] D. D. Walker, E. G. Mercer, and K. E. Seamons. Or best offer: A privacy policy negotiation protocol. In *Policies for Distributed Systems and Networks, 2008. POLICY 2008. IEEE Workshop on*, pages 173–180, June 2008.

[55] Jason Watson, Heather Richter Lipford, and Andrew Besmer. Mapping user preference to privacy default settings. *ACM Trans. Comput.-Hum. Interact.*, 22(6):32:1–32:20, November 2015.

[56] Martin L Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

[57] A Westin. Privacy on & off the internet: What consumers want. Technical report, Tech. Report for Privacy & American

Business. Hackensack, NJ: Privacy & American Business, 2001.

[58] Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nicholas R. Jennings. Iamhaggler: A negotiation agent for complex environments. In Takayuki Ito, Minjie Zhang, Valentin Robu, Shaheen Fatima, and Tokuro Matsuo, editors, *New Trends in Agent-based Complex Automated Negotiations*, Studies in Computational Intelligence, pages 151–158, Berlin, Heidelberg, 2012. Springer-Verlag.

[59] A. Yassine and S. Shirmohammadi. An intelligent agent-based framework for privacy payoff negotiation in virtual environments. In *Computational Intelligence in Virtual Environments, 2009. CIVE '09. IEEE Workshop on*, pages 20–25, March 2009.

[60] A. Yassine and S. Shirmohammadi. Measuring users' privacy payoff using intelligent agents. In *2009 IEEE International Conference on Computational Intelligence for Measurement Systems and Applications*, pages 169–174, May 2009.

[61] Abdulsalam Yassine and Shervin Shirmohammadi. Privacy and the market for private data: A negotiation model to capitalize on private data. In *2008 IEEE/ACS International Conference on Computer Systems and Applications*, pages 669–678. IEEE, mar 2008.