

Implementing Privacy Negotiations in E-Commerce

Sören Preibusch

German Institute for Economic Research,
Königin-Luise-Str. 5, 14191 Berlin, Germany
spreibusch@diw.de

Abstract. This paper examines how service providers may resolve the trade-off between their personalization efforts and users' individual privacy concerns. We analyze how negotiation techniques can lead to efficient contracts and how they can be integrated into existing technologies to overcome the shortcomings of static privacy policies. The analysis includes the identification of relevant and negotiable privacy dimensions for different usage domains. Negotiations in multi-channel retailing are examined as a detailed example. Based on a formalization of the user's privacy revelation problem, we model the negotiation process as a Bayesian game where the service provider faces different types of users. Finally an extension to P3P is proposed that allows a simple expression and implementation of negotiation processes. Support for this extension has been integrated in the Mozilla browser.

1 Introduction

Online users are facing a large and increasing complexity of the web, due to its size and its diversity. In online retailing, stores are constantly expanding their assortments in width, depth and quality levels, making it impossible for users to examine all possible alternatives. The user may be offered effective guidance through automated recommender systems; her appreciation for personalized websites [11], and their economical benefits for service providers could be verified empirically [2, 10, 13].

Several recommendation strategies have been developed in the past, an overview can be found in [14]. All these techniques have in common a need of personal data, either by explicit collection or by inferring them. Thus a personalized (or user-adaptive) system intrinsically has to deal with privacy issues, especially if personal data is stored and not volatile or only saved for the current session. A common way for websites to communicate their privacy principles is to post "privacy policies".

Our contribution is to depict how negotiation techniques can overcome current drawbacks of static privacy policies, and how these negotiations can be implemented using existing technologies (reviewed in section 2). In section 3 and 4, we examine negotiable privacy dimensions and present the optimization calculi of the user and the service provider respectively, based on a formalization of privacy negotiations. Section 5 explains how privacy negotiations can be implemented using P3P and illustrates a negotiation scenario in multi-channel retailing. The paper concludes with a summary and outlook in section 6.

2 Related Work

The privacy-personalization trade-off as presented above has led to several approaches both in research and in practice, among them are the Platform for Privacy Preferences developed by the World Wide Web Consortium (W3C) [19], and the Enterprise Privacy Authorization Language (EPAL) developed by IBM [7]. P3P is a recommendation since 2002 and aims “to inform Web users about the data-collection practices of Web sites” [20]. P3P has become widely adopted by service providers but it remains restricted to the “take-it-or-leave-it” principle: The service provider offers a privacy policy; the potential user either can accept it or to reject it as a whole. A negotiation process between the involved parties is not intended. Although the first drafts of the P3P specification included negotiation mechanisms, these parts had been removed in favour of easy implementation and early and wide adoption of the protocol. The latest P3P 1.1 specification [20] does not mention negotiations either.

In addition to the P3P specification, the W3C conceived APPEL1.0, A P3P Preference Exchange Language 1.0 [18]. APPEL is a language “for describing collections of preferences regarding P3P policies between P3P agents”. APPEL is primarily intended as a transmission format and a machine-readable expression of a user’s preferences. Given a P3P privacy policy, it may be evaluated against a user-defined ruleset to determine if her preferences are compatible with the service provider’s intentions for data. Though standard behaviours and basic matching operations are supported by APPEL, its applications are still limited and the capability of expressing negotiation strategies is explicitly excluded from the language’s scope. Using APPEL as a negotiation protocol is neither supported by its semantics nor is the language designed for this purpose.

EPAL allows enterprises to express data handling practices in IT systems. The developed policies are intended “for exchanging privacy policy in a structured format between applications or enterprises” [7]. The language focuses on the internal business perspective, and is not intended for customers to express their own privacy preferences. Although EPAL is not suited for the direct dialogue with the end-user – which is needed for negotiation – privacy guarantees towards customers can sometimes be deduced from the stated internal procedures and then be expressed in P3P policies.

In parallel to the development of privacy-related technologies and research both in online and offline IT-based transactions, negotiation has been studied in various disciplines. The bases had been set up in game theory, where negotiation is modelled as a bargaining game [8, 16]. Recent influences have arisen with the increasing importance of autonomous agents and collaborative computing [4]. Frameworks for carrying out negotiations have been developed [12]. The rapid development of the Grid and service-based IT-architectures on the technical side, and the enduring process outsourcing to third parties on the economic side, combined with mobile and ubiquitous computing will make Privacy Negotiation Technologies gain in importance in the near future [9, 21].

3 Privacy Negotiations

Thompson states that negotiations are an “interpersonal decision-making process necessary whenever we cannot achieve our objectives single-handedly” [17]. Especially in the case of integrative negotiations, negotiations can unleash the integrative potential that lies in conflicting interests and preferences and turn it into efficient contracts: two major shortcomings of current online privacy handling mechanisms can be overcome if privacy negotiation processes are implemented during the transaction between the service provider (seller) and the user (buyer):

The first shortcoming is the “take-it-or-leave-it” principle, i.e. the user can only accept or refuse the provider’s proposal as a whole. The provider is always the one who moves first, he makes the initial offer; the user cannot take the initiative.

The second shortcoming is the “one-size-fits-all” principle: once the service provider has designed its privacy policy, it will be proposed to all interested users – no matter what their individual preferences are. There may be users who would have accepted offers with less privacy protection and would have agreed to the provider’s proposal even if more personal data would have been asked. Thus, the provider fails to tap the users’ full potential.

3.1 Individualized Privacy Policies

Adopting a broader view and extending the analysis from a single service provider to the whole market, providers specializing on different privacy levels may be an idea. Since the amount of service providers (as discrete units) is much smaller than the amount of potential privacy preferences, which can be seen as quasi-continuous due to the large number of gradations for all considerable privacy dimensions, a specialization is not trivial.

Consider n service providers and $m \gg n$ users having different privacy levels with a known distribution. Hence, a given service provider will target more than one privacy level. This may be implemented by giving the users the choice between a set of usage scenarios corresponding to different amounts of personal data to be collected. As the differences between these usage scenarios have to be clearly communicated and the maintenance of one scenario induces costs for the service provider, the set of scenarios will be limited in size to a few possibilities.

A current example of this strategy is the search site A9.com, a wholly owned subsidiary of Amazon.com, Inc. A9.com offers a highly personalized version of its services, which is the standard setting. Users more concerned about their privacy can switch to an alternative service where the data collection and use is limited to a minimum and no personalization is implemented.

The notable difference between the offered privacy levels is part of the service provider’s user discrimination strategy and aims at a successful self-selection of the potential users. Thence, even under market-driven specialization and alternative usage scenarios, the user still faces fixed policies and the main problem persists.

3.2 Negotiable Privacy Dimensions

Apparently, as it is not feasible to negotiate the entire privacy policy, one important aspect is to identify relevant and negotiable privacy dimensions. We define a *privacy dimension* as one facet of the multi-dimensional concept ‘user privacy’. For each dimension, different discrete revelation levels exist, monotonously associated with the user’s willingness to reveal the data. Privacy dimensions can be identified at different degrees of granularity.

Based on the semantics of P3P, a priori all non-optional parts of a P3P privacy STATEMENT are possible negotiable privacy dimensions: The RECIPIENT of the data, the PURPOSE for which the data will be used, the RETENTION time and what kind of DATA will be collected. As shown by empirical studies, these four generic dimensions (recipient, purpose, retention, and data) reflect privacy aspects users are concerned about. Moreover, they are in accordance with European privacy legislation [5, 6].

It is obvious, that the importance of each of the four dimensions as perceived by the users as well as their respective willingness to provide information, depends on the thematic domain of the service. Some recent work proposed to negotiate the recipient of the data in different application scenarios, among them are medical help [21], distance education [22], and online retailing [4]. We will focus on negotiating the amount of data to be revealed (see section 3.4).

3.3 Privacy vs. Personalization – User’s Individual Utility Calculus

In order to model the user’s individual trade-off between personalization and privacy, we present it as a utility maximization problem, taking into account different overall sensitivity levels towards privacy and different importance one may assign to a specific privacy dimension. The formalization allows solving the negotiation game presented in section 4, giving the service provider the opportunity to choose its optimal strategy.

We denote the user’s utility by U, using the following notations:

- D^n is a n-dimensional privacy space and $d_i \in D$ are its privacy dimensions
- a_i is the user’s data revelation level on dimension d_i
- a_i^T is a threshold indicating the minimum required data the user must reveal
- α_i is a weighting of dimension d_i
- γ indicates the user’s global privacy sensitivity
- R is the discount provided by the service provider
- P are other non-monetary personalization benefits
- B is the base utility by the execution of the contract

Using this notation, the user’s utility can be expressed by:

$$U(.) = -\gamma \cdot \prod_{i=1}^n a_i^{\alpha_i} + P(a_1, \dots, a_n) + R(a_1, \dots, a_n) + B. \tag{1}$$

In case that the user is not willing to provide sufficient data for the contract to be executed, the base utility B and the discount R will be zero (2). The user gets the personalization benefits P even if the involved parties do not conclude on a contract.

In case P is less than the negative utility the user gets from providing the necessary data, the user will prefer unpersonalized usage of the services (3).

$$R(a_1, \dots, a_n) = 0 \wedge B = 0 \iff \exists i : a_i < a_i^T. \tag{2}$$

$$P(a_1, \dots, a_n) < -\gamma \cdot \prod_{i=1}^n a_i^{\alpha_i} \implies \text{unpersonalized usage preferred.} \tag{3}$$

As the ability to identify a user individually (identity inference, also known as triangulation) does not increase linearly when more data is provided, we use a *Cobb-Douglas utility function* instead of an additive composition for the user’s disutility of data revelation. Two other important characteristics of this utility expression in the context of privacy awareness are discussed at the end of this section.

The *thresholds* a_i^T are set by the service provider and are usually openly communicated. In implementations, hints like ‘required field’, ‘required information’ or form fields marked by an asterisk are common practice. The necessity can be deduced from the nature of the transaction: It is obvious that an online bookstore cannot achieve postal delivery if the user refuses to provide her shipping address. It is to note that in this model, the kind of privacy dimensions is not fixed: The purpose as well as the recipient can be privacy dimensions. In the case of shipping, the threshold for the recipient dimension may be the company itself (no third-party logistics company used) and the minimum purpose the user has to agree upon may be postal delivery.

The *weightings* α_i for each of the privacy dimensions as well as the global privacy sensitivity γ are private information of the user and constitute her type. The same holds for the valuation of the non-monetary *personalization benefits* P and the *base utility* B , but these two components can be neglected in the further analysis: First, users tend to only value additional personalization benefits, known solutions will shortly be seen as a standard service and thus there will be no special appreciation. Nevertheless, some personalization benefits may remain. In case of classical implementations such as active guidance, purchase suggestions based on purchase or service usage history, product highlighting or implicit search criteria, the personalization improves the perceived service quality. Through the active support, the user can save search time and simultaneously the matching quality between her preferences and the store’s offers increases: These savings can be seen as monetary benefits and thus subsumed under the variable R . This is especially appropriate, as increased matching quality only becomes effective in case the product is purchased (and R is zero in case of no contract). The base utility can be neglected as it does not depend on the data revelation levels. Hence, the user’s type is determined by α_i and γ .

As mentioned above, the multiplicative structure of the *Cobb-Douglas utility function* allows a good expression of inference threats. In addition, there are two other interesting characteristics in the context of profile data, related to each other. First, the different privacy dimensions are not perfectly substitutable (e.g. the user’s telephone number and her e-mail address constitute two possible ways to contact the user but they are not completely interchangeable). Second, different to an additive composition, the substitution rate between two privacy dimensions (which yields here to -

α_i/α_j) is not constant or independent from the current level of revealed data: it decreases with the amount of data already provided.

The influences of the different parts on the user's utility function are described by the partial derivatives and their interpretations shown below:

$\partial U / \partial a_i \leq 0$: Any privacy infringement reduces the user's utility except in the case where she does not care.

$\partial U / \partial R > 0$: The user appreciates discounts.

$\partial R / \partial a_i \geq 0$: But the service provider is only willing to grant discounts in case he gets some personal information in return. The case $\partial R / \partial a_i = 0$ is applicable for a privacy dimension irrelevant in the current transaction scenario or (more restricted) for which the service provider does not honour revelation.

$\partial P / \partial a_i \geq 0$: The more data the service provider can access, the better the personalization will be.

3.4 Negotiating the 'data'-Dimension

While the recipient may be the relevant negotiation dimension for distance education or health services, we propose the extent and amount of shared data as negotiation dimension for online retailing. First, the willingness of customers to provide personal information is mainly determined by the service provider's reputation, who is the (nonnegotiable) initial recipient of the data. Second, disclosure practices are often determined business processes (e.g. outsourced billing services or delivery by third-party companies). Third, the relevance of the retention time is rated considerably less important [1]. Finally, all data carries with it a more or less pronounced intrinsic purpose that cannot be subject to a negotiation (e.g. phone numbers are used for personal contact and telemarketing). Hence, negotiating the kind of data seems appropriate in the case of online retailing.

Generally spoken, for a type of data to become part of the negotiation process, it must at least meet the following criteria:

- the user must be able to provide the data
- the data must not be off-topic; the user should see at least a slight reason for the necessity of providing it
- it must not be indispensable for the execution of the contract, either by its nature or by the level of detail (i.e. no negotiations for $a_i < a_i^T$)
- the service provider must gain the user's favour for collecting the data, i.e. if the data can be smoothly collected without the user's consent, there is no need for negotiating (for example the request time can be collected automatically)

The empirical findings of [1] allow establishing a cardinal ordering of types of data according to the willingness of user's to provide the information. Ackermann et al. found significant differences in comfort level across the various types of information, implying weighting factors α_i in the user's utility function constituting one aspect of the user's type. The other aspect, the global privacy sensitivity expressed by γ , will be examined in the following section.

4 The Service Provider's Perspective

4.1 Facing Different Types of Users

The service provider is confronted with different types of customers that have various global privacy sensitivity levels, and may rate the importance of one kind of data differently. Efficient customer value extraction is based on a combination of discrimination and negotiation techniques. Discrimination relies on the identification of different groups of customers having the same (or a comparable) type. [1] identified three types: the 'privacy fundamentalists', the 'pragmatics', and the 'marginally concerned' users. [15] distinguishes the pragmatic majority into 'profiling averse' and 'identity concerned' users, hence establishing four user clusters.

Table 1 summarizes the four types whose distribution is assumed to be common knowledge; the characteristics are deduced from [1] and [15].

Table 1. User typology

User type	Characteristics, Important factors
θ_{PF} (fundamentalists)	γ near 1 extremely concerned about any use of their data
θ_{PA} (profile averse)	γ around 0.5 sensitive about equipment, salary, hobbies, health or age
θ_{IC} (identity concerned)	γ around 0.5 sensitive about addresses, phone or credit card numbers
θ_{MC} (marginally concerned)	γ near 0 generally willing to provide data

4.2 Modelling the Negotiation Process

Various methods for modelling negotiation processes exist, some more influenced by computer science (e.g. using finite state machines), others more influenced by micro-economics. We will adopt a game-theoretic approach, examining two possible negotiation scenarios: a sequential game as framework and a simultaneous game that may be played on every step. [3] has examined negotiation protocols in different contexts: customer anonymity (or not), complete knowledge of the service provider's strategy (or not) and no transaction costs for both parties (or not).

Assuming that the service provider can reliably identify privacy fundamentalists for example by means of web usage mining technologies, he will propose a take-it-or-leave-it offer to fundamentalists as in most cases, the valuation of hiding personal data will be higher than the discounts the service provider can offer; the inequality (3) becomes binding. This results in a subgame that can be solved by standard procedures. In contrast, the three other types are indistinguishable for the service provider. Cf. figure 1 where the nodes for the types θ_{PA} , θ_{IC} and θ_{MC} are in the same information set whereas the node for type θ_{PF} is apart.

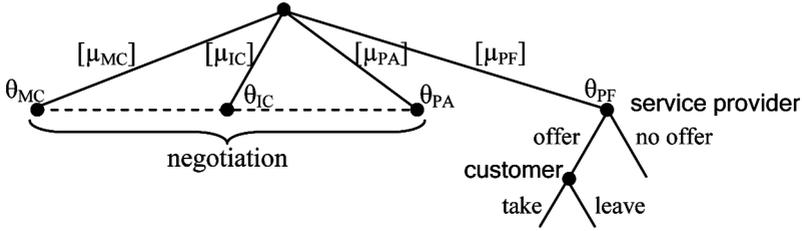


Fig. 1. The service provider negotiates with three types he cannot distinguish

The service provider’s strategy is a function that associates discounts to data revelation level vectors ($D^n \rightarrow \text{ran}(R)$). Determining the service provider’s best strategy results in solving the following optimization problem: For users being drawn from a known distribution (with the probabilities as depicted in figure 1), maximize the total profit. The total profit is the revenue generated by the whole population minus the granted discounts, minus the costs for implementing the personalization, and minus other costs. Latter encompass in particular customers that are lost during the negotiation process by cancelling (e.g. due to psychological reasons or just because they feel overstrained). This maximization is subject to constraint of the users’ participation constraint ($U(\cdot) - B > 0$) and the constraints (2) and (3). We deliberately refrain from a detailed solution, as rigorously integrating the service provider’s cost structure would go beyond the scope of this paper.

The framework for the negotiation process is a dynamic game where the service provider has high bargaining power: He opens the negotiation with a basic offer, consisting of a small discount and a few personal data (the threshold) to be asked. This constitutes the fallback offer in case the user does not want to enter negotiation. In case the user accepts, she will be presented another offer with a higher discount and more data to be asked. On every step, the user may *cancel* (i.e. no contract or the fallback solution are implemented), *continue* (i.e. reveal more data or switch to another privacy dimension) to the next step or *confirm* (i.e. the reached agreement will be implemented).

This wizard-like structure is strategically equivalent to a set of offers as (data, discount)-tuples from which the user can choose one. However, a sequential implementation allows better guidance, better communication of the benefits in providing the data and instantaneous adaptation of the strategy. Note that for a given offer, the requested data are always a subset of the requested data of the previous offer, even if the customer only enters the additional information (monotonously increasing revelation level for a given dimension). The service provider can also implement more alternatives for one step, so that the user can choose which data she will provide (for example the service provider can ask either for the home address or the office address). This is particular useful for addressing different weightings of privacy dimensions that are equivalent for the service provider. Implementations may offer the multiple privacy dimensions sequentially. A switch to another dimension is performed in case the user refuses to provide further data or the service provider is not interested in a higher detail level for the current dimension. A current implementation is described in section 5.2.

In this basic case, the service provider grants a fixed discount on every single step, which is cumulated along the process. A more sophisticated procedure could also include the service provider's concessions into the negotiation process, e.g. by a simultaneous game on every stage: the user indicates the minimum discount she wants to get for revealing the data and the service provider indicates the maximum discount he wants to grant. Problems will arise as the service provider's maximal willingness can be overt due to the unlimited number of times one or several anonymous users can play this simultaneous game.

5 Implementation

5.1 Integrating Privacy Negotiations into P3P

The negotiation process as described in the previous section can be implemented using the extension mechanism of P3P, which can be used both in a policy reference file and in a single privacy policy. The extensions in the privacy policies will not be optional, but in order to ensure backward compatibility, these extended policies will only be referenced in an extension of the policy reference file. Hence, only user agents capable of interpreting the negotiation extension will fetch extended policies.

In a P3P policy, two extensions can be added: a `NEGOTIATION-GROUP-DEF` in the `POLICY` element, and a `NEGOTIATION-GROUP` in the `STATEMENT` element. The mechanism is comparable to the tandem of `STATEMENT-GROUP-DEF` and `STATEMENT-GROUP` in P3P 1.1 [20]. A `NEGOTIATION-GROUP-DEF` element defines an abstract pool of alternative usage scenarios. One or several statements (identified by the attribute `id`) code a possible usage scenario; the pool membership is expressed by the `NEGOTIATION-GROUP` extension in the statement (attribute `groupid`), which describes relevant parameters of the given scenario, such as the benefits for the user. This mechanism allows to establish a n:m-relation between statements and negotiation groups. The fallback contract can be indicated via the `standard`-attribute of the `NEGOTIATION-GROUP-DEF` element. The following example illustrates the usage: users of a bookstore can choose between e-books sent by email and hard copy books, shipped by postal delivery:

Example of an extended P3P policy, including the proposed elements `NEGOTIATION-GROUP-DEF` and `NEGOTIATION-GROUP` (XML namespaces omitted)

```
<POLICY>
...
<EXTENSION optional="no">
  <NEGOTIATION-GROUP-DEF id="delivery" standard="delivery_hardcopy"
    short-description="Choosing delivery medium" /> </EXTENSION>
...
<STATEMENT>
  <EXTENSION optional="no">
    <NEGOTIATION-GROUP groupid="delivery" id="delivery_ebook"
      name="delivery as e-book" benefits="10% discount"/></EXTENSION>
    ...
  <DATA-GROUP><DATA ref="#user.home-info.online.email"/></DATA-GROUP>
</STATEMENT>...
<STATEMENT>
  <EXTENSION optional="no">
```

```

<NEGOTIATION-GROUP groupid="delivery" id="delivery_hardcopy"
  name="delivery as hard copy" benefits="robust hard-cover" />
</EXTENSION>
...
<DATA-GROUP><DATA ref="#user.name" />
<DATA ref="#user.home-info.postal" /></DATA-GROUP>
</STATEMENT>
...
</POLICY>

```

Note that the benefits need to be displayed concisely by the user agent. The human-readable privacy policy and other information resources on the site must work hand in hand with the P3P policy. The exhaustive machine-readable coding of the benefits is a remaining challenge – especially for multi-dimensional phenomena other than just a reduced purchase price.

5.2 Example: Negotiations in Multi-channel Retailing

In addition to the introductory example of the previous section, we want to outline a possible privacy negotiation for a multi-channel retailer. The scenario is as follows: Besides using the service provider’s e-shop, customers can find the nearest store by entering their ZIP code into the store locator. A check of majority is done before these website services can be used.

Two privacy dimensions can be identified: the user’s age (d_1), with the revelation levels: {year (Y), year and month (YM), year and month and day (YMD)}, and the user’s address (d_2), with the revelation levels: {city (C), city and ZIP code (CZ), city and ZIP code and street (CZS)}. The revelation thresholds are $a_1^T = \text{year}$ (for the majority check) and $a_2^T = \text{city and ZIP code}$ (for the store locator). Possible negotiation outcomes are depicted in the left part of figure 2. Using the user’s utility function as defined in equation (1), we can draw iso-utility curves: the user’s disutility increases when moving to the upper right corner, as the revelation levels increase.

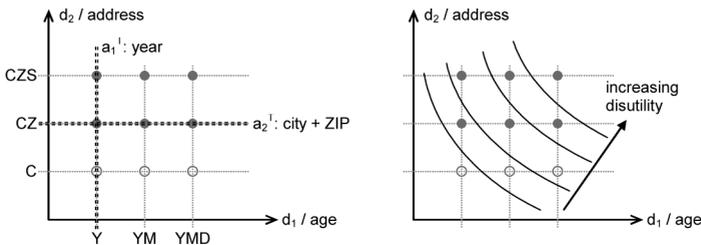


Fig. 2. Possible negotiation outcomes. Contracts whose data revelation levels are below the thresholds cannot be reached (marked by unfilled dots) (left). User’s iso-utility curves corresponding to different revelation levels (right).

Based on these two figures, the service provider develops its strategy: he chooses the discounts he will grant to the customer for each of the six possible contracts, “labelling” them with the $R(\cdot)$ function (that maps from D^n to discounts). Hence, he can code the negotiation space by six statements in an extended P3P policy.

The customer's user agent fetches this policy and serves as a negotiation support system, displaying possible alternatives (a human-readable communication of the data handling practices as coded in the statements along with negotiation benefits) from which the user can choose one. We have integrated this negotiation support into the Mozilla browser, thence extending its P3P support: a site's privacy policy can be accessed via the "Policy", "Summary" and "Options" buttons in the "Page Info" dialog, directly available from the status bar. Extending the chrome components, we have added a "Negotiate" button: a modal dialog is opened, summarizing the negotiable privacy dimensions (d_i) and the possible realizations (a_i) with drop-down menus. The implementation relies on XUL and JavaScript, uses the Mozilla APIs and integrates seamlessly into the user agent. As the proposed extension to P3P is not restricted to a specific privacy dimensions, neither is the implementation. Any privacy dimension can be negotiated as long it can be expressed using the P3P data scheme.

6 Conclusion and Further Work

This paper has presented the necessity of negotiation about privacy principles in a relationship between service provider and customer. Negotiating allows a better matching between the seller's needs and the buyer's disclosure restraint and helps to reduce the trade-off between personalization and privacy. Modelling the user's individual utility maximization can take into account the multi-dimensionality of privacy; the service provider may wish to reduce the negotiation space in a way that suits the given business scenario. The incremental revelation of data by the user can be strategically reduced to a choice from a set of alternatives. Using the extension mechanism of P3P, there is no limitation in coding these alternatives even for complex cases involving diverse privacy dimensions: We proposed two new elements that follow the structure of the current P3P 1.1 grouping mechanisms and allow software-supported negotiations in E-Commerce. Software support of the extension was added to the Mozilla browser, integrating privacy negotiations seamlessly into the user agent.

Future work will focus on the practical implementation of privacy negotiation techniques on large scale public websites. We are currently investigating which user interface design best fulfils the usability requirements and how negotiable privacy dimensions are best visualized. Moreover, a taxonomy should be developed to allow a machine-readable coding of the user's benefits for a negotiation alternative. A remaining question is whether users feel more concerned about their privacy when an explicit negotiation process is started. This increasing sensitivity could make take-it-or-leave-it offers more favourable for the service provider.

References

1. Ackerman, M. S., Cranor, L.F., Reagle, J.: Privacy in E-commerce: Examining User Scenarios and Privacy Preferences, First ACM Conference on Electronic Commerce, Denver, CO (1999) 1-8
2. Cooperstein, D., Delhagen, K., Aber, A., Levin, K.: Making Net Shoppers Loyal, Forrester Research, Cambridge (1999)

3. Cranor, L. F., Resnick, P.: Protocols for Automated Negotiations with Buyer Anonymity and Seller Reputation, *Netnomics*, 2(1), 1-23 (2000)
4. El-Khatib, K.: A Privacy Negotiation Protocol for Web Services. Proceedings of the International Workshop on Collaboration Agents: Autonomous Agents for Collaborative Environments (COLA) (2003)
5. European Parliament, Council of the European Union: Directive 2002/58/EC on privacy and electronic communications. Official Journal of the European Communities, 31.7.2002, L 201, 37–47 (2002)
6. European Parliament, Council of the European Union: Regulation (EC) No 45/2001 of the European Parliament and of the Council of 18 December 2000. Official Journal of the European Communities, 12.1.2002, L 8, 1–22 (2002)
7. International Business Machines Corporation: Enterprise Privacy Authorization Language (EPAL 1.2), W3C Member Submission 10 November 2003 (2003)
8. Karrass, C. L.: Give and Take: The Complete Guide to Negotiating Strategies and Tactics. HarperCollins Publishers, New York, NY (1993)
9. Kurashima, A., Uematsu, A., Ishii, K., Yoshikawa, M., Matsuda, J.: Mobile Location Services Platform with Policy-Based Privacy Control (2003)
10. Peppers, D., Rogers, M., Dorf, B.: The One to One Fieldbook. New York, Currency Doubleday (1999)
11. Personalization Consortium: Personalization & Privacy Survey (2000)
12. Rebstock, M., Thun, P., Tafreschi, O.A.: Supporting Interactive Multi-Attribute Electronic Negotiations with ebXML. Group Decision and Negotiation. 12 (2003) 269–286
13. Schafer, J.B., Konstan, J., Riedl, J.: Recommender Systems in E-Commerce (1999)
14. Schafer, J.B., Konstan, J., Riedl, J.: Electronic Commerce Recommender Applications, *Journal of Data Mining and Knowledge Discovery*. 5, 115–152 (2000)t
15. Spiekermann, S.: Online Information Search with Electronic Agents: Drivers, Impediments, and Privacy Issues (2001)
16. Ståhl, I.: Bargaining Theory. Stockholm: The Economics Research Institute (1972)
17. Thompson, L.L.: The Mind and Heart of the Negotiator. 3rd edn. Pearson Prentice Hall, Upper Saddle River, New Jersey (2005)
18. W3C, A P3P Preference Exchange Language 1.0 (APPEL1.0), W3C Working Draft 15 April 2002, <http://www.w3.org/TR/P3P-preferences> (2002)
19. W3C, The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation 16 April 2002, <http://www.w3.org/TR/P3P/> (2002)
20. W3C, The Platform for Privacy Preferences 1.1 (P3P1.1) Specification”, W3C Working Draft 4 January 2005, <http://www.w3.org/TR/2005/WD-P3P11-20050104/> (2005)
21. Yee, G., Korba, L.: Feature Interactions in Policy-Driven Privacy Management. Proceedings from the Seventh International Workshop on Feature Interactions in Telecommunications and Software Systems (FIW’03) (2003)
22. Yee, G., Korba, L.: The Negotiation of Privacy Policies in Distance Education. Proceedings. 4th International IRMA Conference (2003)