

Towards a mechanism for incentivating privacy

Piero A. Bonatti, Marco Faella, Clemente Galdi, and Luigi Sauro

Università di Napoli “Federico II”

Abstract. The economic value of rich user profiles is an incentive for providers to collect more personal (and sensitive) information than the minimum amount needed for deploying services effectively and securely. With a game-theoretic approach, we show that provider competition can reduce such information requests. The key is a suitable mechanism, roughly reminiscent of a Vickrey auction subject to integrity constraints. We show that our mechanism induces rational providers to ask exactly for the user information strictly necessary to deliver their service effectively and securely. In this framework, maximal attribute disclosures become more difficult to achieve.

1 Introduction

Web sites frequently ask their users for personal, sensitive information before granting access to their full functionalities. For example, some of the most popular web sites and services for e-commerce and social networking collect user name, birth date, gender, detailed address, credit card information, and—in some cases—even sex preferences, and political and religious views. Some of these fields can be easily aggregated to form *quasi identifiers* [7, 19, 16], that is, combinations of attributes such that the probability of their matching a single person is high. (i.e., the values of such attributes uniquely identify an individual with high probability). Releasing fake data is generally not an appropriate or sufficient privacy-preserving measure, as the correctness of some fields may be essential for correct functionality (e.g. credit card information for a purchase and home address for parcel delivery). The ongoing deployment of digital IDs and other cryptographically verifiable documents will further exacerbate this problem.

The above information requests are not parsimonious, in general. Rich user profiles have a significant economic value that constitutes an incentive for increasing the amount of user information collected (and for its disclosure to third parties). Therefore, providers are encouraged to ask for more information than the minimum required to deploy a given service effectively and securely.

Competition may contrast this trend. Indeed, a significant number of users have expressed concern over privacy during online interactions (e.g. user complaints have already influenced Facebook’s privacy policy and services), and analysts say that privacy may become a factor of competition [5, 17, 9].

In this paper we take these analyses seriously and move the first steps towards mechanisms that may increase privacy by exploiting competition between providers. More precisely, such mechanisms should encourage providers to be *truthful*, that is, to ask nothing more than the minimum information sets necessary for correct and secure service deployment. Moreover, such approaches should supply users with *all* of the alternative ways of fulfilling the provider's policy, so that users can choose the alternative that they prefer from a privacy viewpoint.

Setting up a widely applicable mechanism of this kind is, of course, a very complex task. Open questions include at least the following: How should users and providers interact? (E.g. direct interaction vs. mediation by trusted third parties; single interactions vs. prolonged negotiations.) What is the interplay of information disclosure and nonfunctional service properties such as cost and quality of service? In particular, can providers compensate for more invasive information requests with such nonfunctional properties? Can the truthful mechanisms developed by economists be adapted to the scenarios described so far?

A complete answer to all of these questions lies beyond the scope of this paper (and any single paper of standard size). As a first step, here we focus on the last of the above questions, which concerns an essential prerequisite for the mechanisms of our interest and, more generally, for all the transactions where "payments" consist in user information disclosures. In these application scenarios, the payment means is naturally discrete and partially ordered, either in quantitative terms (e.g., according to set inclusion), or in qualitative terms (e.g., based on sensitivity). On the contrary, all standard mechanisms developed in microeconomics are based on totally ordered payment domains, e.g., money. Therefore, reconstructing such mechanisms in order to fulfill the requirements of our scenarios is a problem of general interest—and requires nontrivial changes in the underlying mathematics, as discussed below.

The mechanism studied in this paper is analogous to an auction: The user is the auctioneer; the providers "buy" the user's preference and "pay" for it by reducing the amount of (economically valuable) user information requested. The auction mechanism should be *truthful*, that is, information requested by rational selfish providers should match what is *actually* needed to deploy the service correctly and securely.

Perhaps, the most famous truthful auction mechanism is the second-price auction introduced by Vickrey [20]. It is a "one-shot" mechanism (bidders submit their offers in parallel, then the auctioneer evaluates them); the best offer "wins", and the price payed is the second best offer. As we anticipated, some of the main technical differences between Vickrey's mechanism and ours concern the range of bids and utility functions, that here is discrete and partially ordered

rather than continuous and totally ordered. Of course, in a partial ordering it is not immediately clear how to generalize the notion of second price.

Even if a user may use a same service multiple times, a one-shot mechanism without memory of previous decisions may be appropriate in many cases. In particular, consider any transaction with no monetary costs, that is, where the only “cost” for the user consist of the personal information released. Usually, after such information has been disclosed and the user’s profile created, subsequent service usage requires no further disclosures, so (from a privacy perspective) after the first information release the service can be used for free. Consequently, there is no need for further auctions until services change.

The paper is organized as follows: Section 2 introduces the formal framework and describes its formal properties. Related work is discussed in Section 3. Section 4 concludes the paper with a final discussion of the results and some interesting directions for further work. Proofs can be found in the appendix.

2 Formal framework

The formal framework is relative to an arbitrary but fixed (implicit) service of interest to the user. The set of agents A is identified with an initial segment of the natural numbers: $A = \{0, 1, 2, \dots, N\}$. The user is represented by 0 and the providers by $1 \leq i \leq N$. We assume that the services deployed by the providers are all equivalent from the user’s perspective. The set of information items that can be requested and released before service access is $C = \{c_1, \dots, c_z\}$. By analogy with trust negotiation frameworks [2, 21, 22], information items will be sometimes called *credentials*. The powerset of C is denoted by $\mathcal{P}(C)$.

Credential sets may have different sensitivity. The sensitivity order is modelled with a strict partial order $<$. When a credential set r_2 is more sensitive than a credential set r_1 , we write $r_1 < r_2$. Let $r_1 \leq r_2$ iff either $r_1 < r_2$ or $r_1 = r_2$. We assume that $r_1 \subseteq r_2$ implies $r_1 \leq r_2$ (intuitively, by enlarging information sets, their sensitivity can only increase).

Several concepts, including policies, will be based on *thresholds* over $\mathcal{P}(C)$, that is, sets of sets of credentials $\theta \subseteq \mathcal{P}(C)$ such that $\theta \neq \emptyset$ and for all distinct $r, r' \in \theta$, $r \not\subseteq r'$. Θ denotes the set of all thresholds over $\mathcal{P}(C)$.

The user’s policy is a threshold $pol_0 \in \Theta$. Intuitively, pol_0 represents the maximal sets of information items that the user is willing to disclose to access the service. Formally, a request $r \subseteq C$ is *admissible* iff $\exists r' \in pol_0 : r \subseteq r'$; we denote with $adm(pol_0)$ the set of admissible requests.

Example 1. Suppose $pol_0 = \{\{login, passw\}, \{card_num, exp_date\}\}$. This policy means that the user is willing to disclose either her login-password pair or her credit card number and expiration date. A request for the credit card number

alone is admissible, too, as $\{card_num\}$ is a subset of the second element of pol_0 and hence $\{card_num\} \in adm(pol_0)$. On the contrary, $\{login, card_num\}$ is not admissible, because it is not contained in any element of pol_0 . \square

Symmetrically, each provider i has a policy $pol_i \in \Theta$ that encodes the minimal (alternative) sets of information items that suffice to deliver the service securely and effectively. Formally, a credential set $r \subseteq C$ fulfills pol_i (and grants access to the service) iff $\exists r' \in pol_i : r \supseteq r'$; in the following $ful(pol_i)$ is the set of all $r \subseteq C$ that fulfill pol_i . The *policy profile* is the vector $\mathbf{pol} = \langle pol_0, \dots, pol_N \rangle$.

A provider i may decide to ask users for more information than what is prescribed by pol_i ; the actual information request is called a *strategy* and corresponds to what is traditionally called policy in standard access control frameworks (because it determines which conditions must be fulfilled to access the service). Formally, a *strategy profile* for \mathbf{pol} is any vector $\mathbf{req} = \langle req_1, \dots, req_N \rangle$ such that (i) $req_i \in \Theta$, and (ii) $req_i \subseteq ful(pol_i)$.

Each strategy req_i is the information request that provider i submits to the user. Each credential set $r \in req_i$ is an alternative way of fulfilling i 's request, that is, the user must release a set of credentials $r' \in ful(req_i)$ in order to access the service deployed by i . The requests of provider i are required (by the second condition above) to fulfill the minimal requirements imposed by pol_i . Therefore the user's response r' is guaranteed to satisfy pol_i , too. Note that req_i might omit some ways of accessing the service. That is, for some $r \in pol_i$ there may be no $r' \in req_i$ such that $r' \supseteq r$. In this way, a provider may force the user to disclose credentials that are of greater interest for the provider.

Example 2. With reference to Example 1, a provider that can technically support both account-based and pay-per-use access would have the policy: $pol_1 = \{\{login, passw\}, \{card_num, exp_date\}\}$. The two members of pol_1 represent the two minimal information sets that need to be collected in order to grant the service. The request of provider 1 may in general be different. For instance, if $req_1 = \{\{card_num, exp_date, birth_date\}\}$, the provider is trying to (i) force the user to disclose her credit card information rather than her account information, and (ii) obtain the user's birth date, which is not strictly necessary to service delivery. Note that the only element of req_1 is a superset of the second element of pol_1 , therefore the request fulfills pol_1 . However, it is not admissible w.r.t. pol_0 . Now suppose that $req_1 = \{\{login, passw\}, \{card_num, exp_date, birth_date\}\}$. In this case, due to the additional alternative $\{login, passw\}$, req_1 is admissible for pol_0 and fulfills pol_1 . The user can release the set $\{login, passw\}$ and access the service. \square

A strategy req_i is *truthful* iff $req_i = pol_i$, that is, the provider's requests match the actual minimal requirements for secure and effective service delivery.

Finally, each agent i is associated to a preference relation \preceq_i . We assume that \preceq_0 is \leq (that is, the user's goal is minimizing the sensitivity of disclosed information), while for all providers i , \preceq_i can be either \leq or \subseteq . In the former case, i 's goal is maximizing the sensitivity of the information acquired from users,¹ while in the latter case i 's goal is maximizing the amount of such information. Let $\vec{\preceq}$ be the vector $\langle \preceq_0, \dots, \preceq_N \rangle$.

Now a (*full*) *profile* is a triple $\pi = \langle \mathbf{pol}, \mathbf{req}, \vec{\preceq} \rangle$ where \mathbf{pol} is a policy profile and \mathbf{req} a strategy profile for \mathbf{pol} . The set of all full profiles will be denoted by Π .

2.1 Selection and response mechanism

We need preliminary definitions. The set of optimal admissible requests in a profile π is:

$$\mathit{opt}(\pi) = \min_{\preceq} \left(\bigcup_{j=1}^N \mathit{req}_j \cap \mathit{adm}(\mathit{pol}_0) \right),$$

where $\min_{\preceq}(X)$ denotes the set of minimal elements of X according to \preceq , that is, $\min_{\preceq}(X) = \{r \in X \mid \forall r' \in X, r' \not\prec r\}$.

The user prefers those providers that make minimal information requests. Formally, let the provider i be a *candidate winner* in π iff $\mathit{opt}(\pi) \cap \mathit{req}_i \neq \emptyset$. The set of candidate winners is denoted by $\mathit{cw}(\pi)$.

Each candidate winner i is associated to a set of possible *responses*, $\mathit{res}(\pi, i)$. Possible responses are credential sets that must satisfy both the user's policy and the provider's request, that is, for all $r \in \mathit{res}(\pi, i)$, $r \in \mathit{adm}(\mathit{pol}_0) \cap \mathit{ful}(\mathit{req}_i)$. Different specific definitions of $\mathit{res}(\pi, i)$ yield different properties in terms of robustness (i.e., lack of unnecessary transaction failures) and amount of information released. Before discussing the alternatives, let us fix the decision making process (provider selection and response selection), consisting of two steps:

1. choose a provider $i \in \mathit{cw}(\pi)$
2. choose a response $r \in \mathit{res}(\pi, i)$.

If $\mathit{res}(\pi, i) = \emptyset$, then the transaction fails. To simplify the discussion, let us assume that the above choices are made at random with uniform probability (different distributions can be adopted, though).

¹ The rationale is that information value is often correlated with sensitivity. For simplicity, in this first paper we assume a shared (objective) measure of sensitivity – e.g. based on statistics about the identification power of attribute aggregates, cf. quasi-identifiers [7] – so that \preceq may be regarded as common knowledge. Generalizations are discussed in Sec. 4.

So far, the framework is similar to an auction with some extra constraints posed by policies. In Vickrey's auctions, truthfulness is achieved by setting the price to the second best offer. In this framework, a direct counterpart of this idea consists in defining $res(\pi, i)$ as the set of best requests made by all the providers $j \neq i$. In order to formalize this idea we need a few more auxiliary definitions. First, let $opt_{-i}(\pi)$ denote the best admissible requests of the providers $j \neq i$:

$$opt_{-i}(\pi) = \min_{\leq} \left(\bigcup_{\substack{j \neq i \\ 1 \leq j \leq N}} req_j \cap adm(pol_0) \right).$$

Then, let $res_0(\pi, i)$ be the set of all the best admissible requests of the providers $j \neq i$ that satisfy i 's request, that is: $res_0(\pi, i) = opt_{-i}(\pi) \cap ful(req_i)$.

By setting $res(\pi, i) = res_0(\pi, i)$ one obtains a mechanism that easily leads to failures, because the requests of the winners might not fulfill each other.

Example 3. Suppose that the user's policy permits the simultaneous disclosure of her credit card number (*card_num*), its security code (*sec_code*), user name (*name*), and birth date (*birth_date*), i.e., $pol_0 = \{\{card_num, sec_code, name, birth_date\}\}$. Let the requests of providers 1 and 2 be $req_1 = \{\{card_num, name, sec_code\}\}$ and $req_2 = \{\{card_num, name, birth_date\}\}$, respectively. If req_1 and req_2 are not comparable with respect to $<$, then both requests are \leq -minimal and admissible, so the set of candidate winners is $cw(\pi) = \{1, 2\}$. However, the request of provider 1 does not fulfill the request of provider 2 and viceversa. Then $res_0(\pi, 1) = res_0(\pi, 2) = \emptyset$ and the transaction fails no matter which of the two providers is chosen. \square

This problem can be mitigated by adding to the pool of replies the largest admissible requests, that is, the members of pol_0 . Let

$$\begin{aligned} res_1(\pi, i) &= \min_{\leq} (opt_{-i}(\pi) \cup pol_0) \cap ful(req_i), \\ res_2(\pi, i) &= (opt_{-i}(\pi) \cup pol_0) \cap ful(req_i). \end{aligned}$$

The first definition (res_1) still leads to a failures; for instance, in Example 3, it is equivalent to res_0 because $\min_{\leq} (opt_{-i}(\pi) \cup pol_0) = opt_{-i}(\pi)$, for $i = 1, 2$. The second definition (res_2) does not lead to any failure in Example 3; however, the user has to disclose all releasable credentials (*card_num*, *sec_code*, *name*, *birth_date*). In general, $res_2(\pi, i)$ contains (at least) all the elements of pol_0 that cover some request of i , therefore some maximal disclosable set of credentials can always be released with probability greater than 0. Then we move over to a more parsimonious definition (in terms of maximal disclosures). The idea consists in "interpolating" intermediate requests between the optimal requests

of all $j \neq i$. Such interpolation constitutes a “vault” above req_i , from which possible responses can be selected. Formally, we set

$$V(\pi, i) = \{r \subseteq C \mid r \in adm(pol_0) \wedge r \in ful(req_i) \wedge \forall r' \in opt_{-i}(\pi). r' \not\prec r\}$$

$$res(\pi, i) = \max_{\preceq_i} V(\pi, i).$$

Intuitively, $V(\pi, i)$ contains all sets of credentials that are simultaneously (i) releasable by the user, (ii) acceptable by provider i , and (iii) not more sensitive than any other request. In order to achieve truthfulness (i.e., make sure that providers have no incentive to request more credentials than needed), $res(\pi, i)$ selects the maximal elements of $V(\pi, i)$, w.r.t. the preference relation of provider i . Compare this definition with the standard second price approach: There, the winner pays the minimum price that is not worse (i.e., smaller) than any other offer; analogously, in our framework, the winner gets a maximal response that is not worse (i.e., more sensitive) than any other offer, and satisfies both the user’s policy and the winner’s request.

We are going to study in depth the framework based on this definition. Before starting its formal analysis, note that in the case of Example 3, assuming for simplicity that \leq equals \subseteq , we have

$$res(\pi, 1) = V(\pi, 1) = \{\{card_num, name, sec_code\}\} = req_1$$

$$res(\pi, 2) = V(\pi, 2) = \{\{card_num, name, birth_date\}\} = req_2.$$

Hence, the selected provider receives nothing more than what it asked for (as opposed to what happens with res_2). In general, however, the user may have to release more than what the winning provider asks for (this is the price to pay for truthfulness). At the end of this section we will characterize a wide class of scenarios in which no unnecessary information is disclosed.

The first formal property of the mechanism defined by res concerns its robustness: Unlike res_0 and res_1 , it introduces no unnecessary failures. In other words, all candidate winners can be given a response:

Theorem 1. *For all $i \in cw(\pi)$, it holds $res(\pi, i) \neq \emptyset$.*

It is easy to verify that a similar property holds for res_2 . So the second formal property of interest concerns a comparison of the two robust strategies res and res_2 with respect to the amount of credentials potentially released. It can be shown that res is generally more parsimonious than res_2 :

Theorem 2. *For all $r \in res(\pi, i)$ there exists $r' \in res_2(\pi, i)$ such that $r \subseteq r'$.*

Next we investigate the effectiveness of the provider selection mechanism in reducing the amount of information disclosed in the worst case. A first question

related to this issue is: Under which circumstances can a maximal releasable set of credentials (i.e., a member of pol_0) be disclosed?

Theorem 3. *Let $r \in pol_0$, $r \in res(\pi, i)$ only if there exists $x \in (opt(\pi) \cap req_i)$ such that $x \subseteq r$ and for all the other providers $j \neq i$ and for all $r' \in req_j$, it holds $r' \not\subseteq r$.*

Note that if $r' \subseteq r$, then $r' \not\subseteq r$ implies $r' = r$. Then Theorem 3 says that r can be released to i if either there is no competition on the option r (i.e., the other providers' requests are not compatible with r), or the competitors ask exactly for r . Consequently, it appears that competition makes maximal disclosures more difficult to achieve systematically, at least in the absence of detailed information about the user's policy.

Another interesting, related problem is characterizing the circumstances under which the user may have to release all disclosable credentials at once (as it may happen in the formal trust negotiation frameworks studied in the literature).

Corollary 1. *Assume that i makes an admissible request ($req_i \cap adm(pol_0) \neq \emptyset$). Then $\bigcup_{r \in pol_0} r$ can be disclosed to provider i iff the following conditions hold: (i) $pol_0 = \{r\}$, and (ii) $r \in req_j$ for all providers $j \neq i$ such that $req_j \cap adm(pol_0) \neq \emptyset$.*

According to the first condition in the above corollary, if $|pol_0| > 1$, then it is impossible to release $\bigcup_{r \in pol_0} r$. The reason is clear: whenever $|pol_0| > 1$, the user's policy encodes some integrity constraints that forbid the disclosure of arbitrary unions of disclosable credentials. For example, if $pol_0 = \{\{birthday\}, \{address\}\}$ then both birth date and address can be separately disclosed, but their union is considered too sensitive to be released. Now suppose that $|pol_0| = 1$. Condition 2 says that either i has no competitors (no other provider j makes any admissible request), or i 's competitors all ask for r (which is unlikely in practice unless pol_0 is public). This shows how competition helps in reducing complete credential disclosures.

Example 4. Suppose that a user is willing to execute payments either by using her credit card or by bank transfer. In the first case she permits the simultaneous disclosure of her credit card number ($card_num$) and its security code (sec_code). For the latter payment form, she is willing to provide the unique ID (id) associated to the bank transfer and her own bank account information (acc_info). Formally $pol_0 = \{\{card_num, sec_code\}, \{id, acc_info\}\}$. The user can select among three providers for executing a given payment whose requests are, respectively: $req_1 = \{\{card_num\}\}$, $req_2 = \{\{card_num\}, \{id\}\}$, $req_3 = \{\{card_num, sec_code\}\}$. In such context, there is clear competition among all servers since everyone allows credit card payments. Server 1 and 2 only require $card_num$ while server 3 requires both $card_num$ and sec_code .

So, in a parsimonious selection, the user would prefer server 1 or 2. For the second payment method there is no competition at all. Indeed only server 2 allows bank transfers and requires the unique transaction identifier in order to accept the payment. Let π be a profile describing such a scenario. The set of optimal admissible requests is $opt(\pi) = \min_{\subseteq}(\bigcup_{j=1}^N req_j \cap adm(pol_0)) = \{\{card_num\}, \{id\}\}$. Then the set of candidate winners is $cw(\pi) = \{1, 2\}$. Let us focus on server 1, and assume that $\leq = \subseteq$. The set of optimal admissible requests made by all providers except 1 is $opt_{-1}(\pi) = \{\{card_num\}, \{id\}\}$. Then the response to server 1 is $res(\pi, 1) = V(\pi, 1) = \{\{card_num\}\}$. Note that the response contains no elements of pol_0 , because provider 1 competes with provider 2 (whose request $\{card_num\}$ expunges $\{card_num, sec_code\}$ from $V(\pi, 1)$). Similarly, for provider 2 we have $opt_{-2}(\pi) = \{\{card_num\}\}$, $V(\pi, 2) = \{\{card_num\}, \{id\}, \{id, acc_info\}\}$, and $res(\pi, 2) = \{\{card_num\}, \{id, acc_info\}\}$. Thus, when different servers compete within a specific element of pol_0 , (here, $\{card_num, sec_code\}$), such element is not entirely disclosed to any server. On the other hand, if exactly one server makes a request compatible with some element in $r \in pol_0$, i.e., if there is no competition within r ($r = \{id, acc_info\}$ in the example), then r is fully disclosed with nonzero probability. \square

2.2 Rational strategies

We are left to characterize the strategies adopted by ideally rational providers. We consider two kinds of providers: Providers of the first kind are mainly interested in attracting new customers, i.e. their primary goal is maximizing the probability of being selected (or probability of winning) $pw(\pi, i)$, where $pw(\pi, i) = 1/|cw(\pi)|$ if $i \in cw(\pi)$, and $pw(\pi, i) = 0$ otherwise. As a secondary goal, these providers prefer those strategies that better meet their preference \leq_i . Providers of the second kind invert the above priorities. A new player in a given application domain is likely to be a player of the first kind. Similarly, providers whose main income is based on advertisement are likely to be providers of the first kind. On the contrary, when the utility of service usage is dominated by the value of user profiles, providers should be expected to be agents of the second kind.

In order to formalize the perfect strategies for these providers we need some auxiliary notions. First, one needs to compare different responses w.r.t. provider preferences. For this purpose, \leq_i should be extended from credential sets to *sets* of credential sets (i.e., the range of res).

Definition 1. For all $\rho, \rho' \subset \mathcal{P}(C)$, let $\rho \leq_i^* \rho'$ iff

1. for all $r \in \rho$, there exists $r' \in \rho'$ such that $r \leq_i r'$, and
2. for all $r' \in \rho'$ and $r \in \rho$, $r' \not\leq_i r$.

In other words ρ' is preferable to ρ if for all possible responses in ρ there exists an equally preferred or better response in ρ' (according to i 's preferences) and none of the responses in ρ' is less preferable than any response in ρ .

Next, we need a handy way of replacing the strategy of an agent: For all strategy profiles \mathbf{req} and all providers $1 \leq i \leq N$, let

$$\mathbf{req}[i \leftarrow req'] = \langle req_1, \dots, req_{i-1}, req', req_{i+1}, \dots, req_N \rangle,$$

and for all profiles $\pi \in \Pi$, let

$$\pi[i \leftarrow req'] = \langle \mathbf{pol}, \mathbf{req}[i \leftarrow req'], \overset{\rightarrow}{\leq} \rangle.$$

Finally, let $pol \in \Theta$ be an arbitrary but fixed policy. The set of profiles where $pol_i = pol$ is denoted by Π_{pol}^i . A strategy for pol is any $req \in \Theta$ such that for all $r \in req$, $r \in \text{ful}(pol)$. Now the optimal strategies for the two kinds of agents can be formalized.

Definition 2. A strategy req for pol is a dominant attraction strategy for i with respect to pol iff for all $\pi \in \Pi_{pol}^i$

1. $pw(\pi, i) \leq pw(\pi[i \leftarrow req], i)$, and
2. if $pw(\pi, i) = pw(\pi[i \leftarrow req], i)$ then $res(\pi, i) \leq_i^* res(\pi[i \leftarrow req], i)$.

Definition 3. A strategy req for pol is a dominant investigation strategy for i with respect to pol iff for all $\pi \in \Pi_{pol}^i$

1. $res(\pi, i) \leq_i^* res(\pi[i \leftarrow req], i)$, and
2. if $res(\pi[i \leftarrow req], i) = res(\pi, i)$ then $pw(\pi, i) \leq pw(\pi[i \leftarrow req], i)$.

A few explanations are in order here. The universal quantification over Π_{pol}^i , whose only invariant is $pol_i = pol$, means that strategy req is optimal w.r.t. all the other possible strategies req_i that might be adopted by i , and this holds in all possible contexts (i.e., no matter what the policies and strategies of the other agents are). Our mechanism yields the desired result: the truthful strategy $req = pol$ is the best strategy a provider can adopt, under both priorities.

Theorem 4. For all $pol \in \Theta$ and all providers i , the unique dominant attraction strategy for i w.r.t. pol is pol itself.

Theorem 5. For all $pol \in \Theta$ and all providers i , the unique dominant investigation strategy for i w.r.t. pol is pol itself.

In game-theoretic terms, the previous two theorems prove that being truthful is a *dominant strategy equilibrium* (DSE) [13], i.e., no matter what the other agents' policies and strategies are, being truthful is always the best response. Every DSE is in particular a Nash equilibrium.

One may wonder whether gaining information on the behavior of the other agents could allow a provider to increase either its winning probability or the amount of credentials received from the client. The answer is negative, regardless of the extra information available to the provider. Indeed, any gained information corresponds to restricting the set of possible profiles Π_{pol}^i in the definition of dominant attraction (resp., investigation) strategy (Definition 2 and 3, respectively). Clearly, by applying any such restriction, the set of dominant attraction (resp., investigation) strategies may only increase. However, since any pair of dominant strategies dominate each other, it is straightforward to see from Definition 2 and 3 that all dominant strategies give rise to the same probability of winning and the same response from the client.

The presence of rational (and hence truthful) providers may induce minimal disclosures in a framework that, in general, releases to providers more information than what they ask for. For simplicity, we analyze this issue in scenarios where all providers have the same policy. In practice, this assumption is naturally satisfied when provider policies are determined by the same technological constraints—for example, all providers supporting VISA credit card payments must provide VISA's servers with the same information for credit card validation (as in Example 5 below).

Theorem 6. *If all providers have the same policy pol and provider i is truthful, then for all $j \neq i$, $res(\pi, j) \subseteq pol$.*

In informal terms, the above theorem ensures that under the uniform policy hypothesis, one rational (i.e., truthful) provider is sufficient to ensure that all *other* providers do not receive more information than what their common policy requires.

Example 5. Consider an e-commerce application with the same credentials as in Example 3. Assume that all vendors use the same underlying financial institution that requests the credit card number and either the owner's name or the credit card security code, so the providers all share the same policy $pol_i = \{\{card_num, name\}, \{card_num, sec_code\}\}$. Suppose that server i is truthful, while j has strategy $req_j = \{\{card_num, name, birth_date\}, \{card_num, sec_code\}\}$, i.e., in addition to the credit card number and owner's name, j requests also her birth date. The policy of the client is $pol_0 = \{\{card_num, name, sec_code, birth_date\}\}$, i.e., all credentials are simultaneously releasable. Now j

is a candidate winner (as it makes the request $\{card_num, sec_code\}$). However, even if j is selected to deliver the service, in accordance with Theorem 6 it will *not* receive the user's birth date, even if it is releasable by the client. Indeed, we have $opt_j(\pi) = \{\{card_num, name\}, \{card_num, sec_code\}\}$, and both $V(\pi, j)$ and $res(\pi, j)$ contain only the set of credentials $\{card_num, sec_code\}$. \square

As a corollary to Theorem 6, when there are at least *two* rational providers, no provider receives more than what the policy requires.

Corollary 2. *If all providers have the same policy pol and there are two truthful providers, then for all providers i it holds $res(\pi, i) \subseteq pol$.*

Consider the following scenario, for example. It is inspired by real flight reservation portals. Kayak and Momondo ask for no information; tickets are purchased directly from airline companies. On the contrary, eDreams asks for a rich user profile that is then used to make a request to airline companies. Note that eDream user profiles comprise attributes that are not mandatory for airline companies. The following is a formalization of this scenario:

Example 6. Assume that $pol_1 = pol_2 = pol_3 = \{\emptyset\}$, $req_1 = req_2 = \{\emptyset\}$, and $req_3 = \{\{name, address, phone_num, email\}\}$. Note that providers 1 and 2 are truthful and provider 3 is not. Clearly, only providers 1 and 2 are candidate winners. The response is $\{\emptyset\}$, that is, the user releases no information. Note that the same result is obtained when $pol_3 \neq \{\emptyset\}$ (e.g., $pol_3 = req_3$); thus, in future work, it may be interesting to relax the hypothesis of Corollary 2. \square

3 Related work

To the best of our knowledge, the approach introduced in the above sections has no analogue in the literature. Standard access control frameworks place no constraints on policies and set up no mechanisms for reducing the extension and sensitivity of user profiles. In the trust negotiation area, privacy is mainly pursued by having disclosed information match as precisely as possible the provider's request, see for example [18, 14, 10]. Another work aims at preventing providers from inferring which *types* of credentials are owned by the user when a transaction fails [21]. A major difference between our approach and theirs is that those works do not attempt to influence the providers' requests. There, providers' policies do not represent minimal functional and security requirements (unlike our pol_i s); they should rather be considered as part of the specification of req_i . Not only can trust negotiation policies ask for unnecessary information; some interoperable negotiation strategies may further inflate requests by forcing the user agent to disclose all releasable credentials in the

attempt to keep the negotiation alive [22, 2]. Therefore, the issue of minimizing the difference between req_i and pol_i is not tackled (and there is no precise counterpart of pol_i).

In [8], privacy is tackled in the context of auctions (including second-price auctions). Their purpose is the opposite of ours, namely, minimizing the amount of information that *bidders* have to disclose in order to let the auctioneer compute the optimal outcomes. Payments are fully traditional: continuous and totally ordered. Similarly, another economically inspired model [11] proposes both an estimate of the value of private information and a fair compensation for such information release that may induce *users* to release richer and correct information about their personal preferences.

Concerning the many auction models introduced in the past, the main technical difference is that utility functions and bids always range over a totally ordered domain such as the set of real numbers. Then truthfulness can be obtained with a straightforward implementation of the second-price idea, without resorting to more complex notions such as vaults.

4 Discussion and perspectives

4.1 Current achievements

We provided a first formal evidence that the potential competition between equivalent applications may enhance privacy and reduce the amount of personal information requested and collected by application providers. As a starting point, we focussed on scenarios where the competing services are equivalent with respect to functionalities and quality, and the only cost for the user consists of the personal information released. Flight reservation applications provide concrete examples of such scenarios, cf. Example 6 and the preceding paragraph. We argued that one-shot mechanisms (like second-price auctions) are appropriate for such scenarios; indeed, after using the service a first time, subsequent usage is “free” from the point of view of information disclosure; therefore, re-using the same service is always an optimal choice for the user until either demand or offer change.

We proved that a suitable one-shot mechanism regulating provider selection and information disclosure induces truthful behavior in selfish rational providers, resulting in minimal information requests and in a wider range of choices for the user. It is important to note that essentially these results still hold even if agents know each other’s policies and strategies.

When functional and security requirements are the same for all providers (e.g., because such requirements are determined by exogenous technological constraints), two rational (and hence truthful) providers suffice to minimize the

amount of user information disclosed to *any* provider—and interestingly, if only one provider is rational, then it is the only provider that may receive a non-minimal response. In some cases, the hypothesis that all policies are the same can be relaxed, as shown in Example 6. Another set of results shows that extracting all releasable credentials (or any maximal disclosable set of credentials) from a user is a difficult task, whose systematic achievement, in practice, requires some knowledge of the user’s policy.

Competition may be exploited to address a weakness of automated disclosure techniques, such as OpenID profile sharing and automated trust negotiation. These approaches require a policy to decide when a user profile can be automatically transferred to a new web service, or when an information request is reasonable in a given context. Formalizing a policy that decides on behalf of the user whether a provider is collecting a reasonable set of user attributes is a very difficult task. A mechanism inducing a spontaneous moderation of provider requests may turn out to be less expensive and/or more reliable.

Concerning trust negotiation, it is known that interoperable negotiation strategies are vulnerable to attacks that force a peer to release progressively all disclosable information (cf. [22, 2]). It seems that the current one-shot mechanism can address this problem: First, the user agent negotiates with all the providers’ agents using a method that does not actually disclose credentials until the end of the negotiation (as in the protocol introduced in [21]). At the end of the negotiations, before really sending credentials, the user can choose among the successful negotiations those with minimal information disclosure, and compute the response with our mechanism, thereby inducing rational providers to reduce information requests.

4.2 Possible variations to the current framework

Several aspects of the framework can be modified without affecting our results. For example, it is not hard to see that different probability distribution can be used in choosing winner and response. As an example of possible applications, skewed distributions over candidate winners may address additional user preferences over providers. Similarly, the preference relations \leq_i^* employed to compare responses and define dominant strategies can be modified in various ways, in order to model providers with different risk attitudes. For instance, an optimistic (resp. pessimistic) provider may consider only the \leq_i -maximal (resp. \leq_i -minimal) elements of $res(\pi, i)$ (the current definition considers all of $res(\pi, i)$). Our results hold for all such agents.

4.3 Generalizing preferences

The current theoretical framework should be extended and complemented along several directions. In general, providers compete also on properties such as cost and quality of service. They can easily be modelled by extending credential sets to richer sets, including items that represent the additional quantities of interest. This affects some of the assumptions we adopt in this paper.

For example, in the extended framework, preference relations rank aggregates of credentials, money, and quality of service, thereby reflecting a tradeoff between privacy-related costs and other costs and benefits; consequently, user preferences have a more “personal” nature and it would not be reasonable to make the simplifying assumption that \leq is based on an objective, shared sensitivity measure. As a consequence, provider preferences could not be restricted to \subseteq and \leq , and the effects of this generalization should be formally analyzed. In this context, users may publish their preference relation to get personalized offers; it should be verified whether a rational user should be truthful, and whether preference disclosure may constitute a privacy violation in itself.

As preference specifications become more sophisticated, the need may arise for usability-enhancing techniques. For instance, a coarse-grained preference relation in a pure information disclosure scenario (where service re-use has no additional “costs”) may induce the user to always select the same provider (lock-in effect). The articulated approach to monitoring, refining, and learning preferences introduced in [15] for access control policies can perhaps be adapted to our framework.

4.4 Repeated auctions

Another consequence of modelling features such as money and quality is that each call to a same service may have additional costs; then it makes sense to repeat the service selection process and abandon the one-shot approach. We conjecture that, in general, a sequence of independent selections may lead to globally suboptimal disclosures (as any greedy strategy). A formal analysis of iterative selections is an important step in our agenda, that may start from the literature on repeated (or sequential) procurement auctions, e.g. [12, 1]. It is also interesting to evaluate “globalized” selections over bundles of services as an attempt to optimize information disclosure for a set of commonly used services. It is known that this optimization problem is tractable in some cases [4].

4.5 Deployment

Last but not least, the need is felt for an articulated analysis of deployment solutions; for example, providers may want to make sure that the service selec-

tion mechanism is carried out correctly, i.e. users are not cheating and actually disclose an element of the vault. Some application contexts may admit trusted intermediary services. Portals similar to Kayak, Momondo etc. are interesting candidates to fill in this role, that may create new business opportunities and models. In the absence of trusted third parties (i.e., the user is the auctioneer), auctions can be implemented using Secure Multiparty Computation. The instantiation of general MPC constructions can be inefficient; it is, however, possible to design less complex specialized MPC protocols implementing the described variation of second price auctions, cf. [3]. Alternatively, one can resort to ad hoc protocols where credential requests are eventually revealed to all providers (*commitments* and *blind signature* [6] may be employed for this purpose).

Acknowledgments. The authors would like to thank Alessandro Bonatti for helpful preliminary discussions, and Adam J. Lee and the anonymous referees for their detailed and stimulating feedback.

References

1. J. Bae, E. Beigman, R. A. Berry, M. L. Honig, and R. V. Vohra. Sequential bandwidth and power auctions for distributed spectrum sharing. *IEEE Journal on Selected Areas in Communications*, 26(7):1193–1203, 2008.
2. S. Baseline, P. Bonatti, and M. Faella. On interoperable trust negotiation strategies. In *IEEE POLICY'07*, pages 39–50. IEEE Computer Society, 2007.
3. P. Bogetoft, I. Damgård, T. P. Jakobsen, K. Nielsen, J. Pagter, and T. Toft. A practical implementation of secure auctions based on multiparty integer computation. In *Financial Cryptography and Data Security, 10th Int. Conf., revised selected papers*, pages 142–147, 2006.
4. P. A. Bonatti and P. Festa. On optimal service selection. In A. Ellis and T. Hagino, editors, *Proc. of the 14th int.conf. on World Wide Web, WWW 2005*, pages 530–538. ACM, 2005.
5. A. Broache. Competition is good for search privacy, report says. CNET News, http://news.cnet.com/Competition-is-good-for-search-privacy,-report-says/2100-1029_3-6201468.html, August 8 2007.
6. D. Chaum. Blind signatures for untraceable payments. In *Advances in Cryptology - Crypto '82*, pages 199–203. Springer Verlag, 1983.
7. T. Dalenius. Finding a needle in a haystack - or identifying anonymous census records. *Journal of Official Statistics*, 2(3):329–336, 1986.
8. J. Feigenbaum, A. D. Jaggard, and M. Schapira. Approximate privacy: foundations and quantification (extended abstract). In D. C. Parkes, C. Dellarocas, and M. Tennenholtz, editors, *ACM Conference on Electronic Commerce*, pages 167–178. ACM, 2010.
9. E. Gray. FTC to boost competition in privacy protection. *Global Competition Review*, September 23 2010.
10. Y. He, M. Zhu, and C. Zheng. An efficient and minimum sensitivity cost negotiation strategy in automated trust negotiation. *Comp. Sci. and Soft. Eng., Int. Conf. on*, 3:182–185, 2008.
11. J. Kleinberg, C. H. Papadimitriou, and P. Raghavan. On the value of private information. In *TARK '01: Proceedings of the 8th conference on Theoretical aspects of rationality and knowledge*, pages 249–257, San Francisco, CA, USA, 2001. Morgan Kaufmann.

12. R. Lutton and P. R. McAfee. Sequential procurement auctions. *Journal of Public Economics*, 31(2):181–195, 1986.
13. M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
14. F. Paci, D. Bauer, E. Bertino, D. M. Blough, and A. C. Squicciarini. Minimal credential disclosure in trust negotiations. In E. Bertino and K. Takahashi, editors, *Digital Identity Management*, pages 89–96. ACM, 2008.
15. N. M. Sadeh, J. I. Hong, L. F. Cranor, I. Fette, P. G. Kelley, M. K. Prabaker, and J. Rao. Understanding and capturing people’s privacy policies in a mobile social networking application. *Personal and Ubiquitous Computing*, 13(6):401–412, 2009.
16. P. Samarati. Protecting respondents’ identities in microdata release. *IEEE Transactions on Knowledge and Data Engineering*, 13(6):1010–1027, 2001.
17. A. Schwartz and A. Cooper. Search privacy practices: A work in progress. Center for Democracy and Technology report, August 2007.
18. A. C. Squicciarini, E. Bertino, E. Ferrari, F. Paci, and B. M. Thuraisingham. PP-trust-X: A system for privacy preserving trust negotiations. *ACM Trans. Inf. Syst. Secur.*, 10(3), 2007.
19. L. Sweeney. Guaranteeing anonymity when sharing medical data, the Datafly system. *Journal of the American Medical Informatics Association*, 1997.
20. W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
21. W. H. Winsborough and N. Li. Protecting sensitive attributes in automated trust negotiation. In S. Jajodia and P. Samarati, editors, *WPES*, pages 41–51. ACM, 2002.
22. T. Yu, M. Winslett, and K. E. Seamons. Supporting structured credentials and sensitive policies through interoperable strategies for automated trust negotiation. *ACM Trans. Inf. Syst. Secur.*, 6(1):1–42, 2003.

Appendix: Technical Lemmas and Proofs

This Appendix contains the proofs for the theorems presented in the paper, supported by a series of technical lemmas.

Lemma 1. *If $opt(\pi) \cap req_i \neq \emptyset$, then $res(\pi, i) \neq \emptyset$.*

Proof. By definition, $res(\pi, i) = \max_{\leq_i} V(\pi, i)$, where $V(\pi, i) = \{r \subseteq C \mid r \in adm(pol_0) \wedge r \in ful(req_i) \wedge \forall r' \in opt_{-i}(\pi). r' \not\prec r\}$. It is sufficient to show that $V(\pi, i) \neq \emptyset$.

Let $r \in opt(\pi) \cap req_i$. By definition $r \in opt(\pi)$ implies that $r \in adm(pol_0)$. Similarly, $r \in req_i$ implies that $r \in ful(req_i)$. Assume, by contradiction, that exists $r' \in opt_{-i}(\pi)$ such that $r' < r$. If this is the case, then $r \notin opt(\pi)$. \square

Theorem 1. *For all $i \in cw(\pi)$, it holds $res(\pi, i) \neq \emptyset$.*

Proof. Theorem follows by Lemma 1 and by observing that, for all $i \in cw(\pi)$ it holds that $opt(\pi) \cap req_i \neq \emptyset$. \square

Theorem 2. *For all $r \in res(\pi, i)$ there exists $r' \in res_2(\pi, i)$ such that $r \subseteq r'$.*

Proof. Assume $r \in res(\pi, i)$. By definition, $r \in adm(pol_0)$ and $r \in ful(req_i)$. From $r \in adm(pol_0)$, it follows that there exists $r' \in pol_0$ such that $r \subseteq r'$. Then from $r \in ful(req_i)$, derive $r' \in ful(req_i)$, and hence $r' \in pol_0 \cap ful(req_i)$. The theorem follows immediately from the definition of $res_2(\pi, i)$. \square

Theorem 3. *Let $r \in pol_0$, $r \in res(\pi, i)$ only if there exists $x \in req_i$ such that $x \subseteq r$ and for all the other providers $j \neq i$ and for all $r' \in req_j$, it holds $r' \not\prec r$.*

Proof. First of all, observe that, since $r \in pol_0$ then $r \in adm(pol_0)$. Furthermore, it is immediate that, $V(\pi, i) \subseteq ful(req_i)$. By definition, if $r \in res(\pi, i)$, then $r \in V(\pi, i)$ and, thus, there exists a request $x \in req_i$ such that $x \subseteq r$. Assume now that for some provider $j \neq i$ and some request $r' \in req_j$, $r' < r$. On one hand, there exists in $opt_{-i}(\pi)$ a request r'' such that $r'' \leq r'$ and hence $r'' < r$. On the other hand, since $r \in res(\pi, i)$, for all $r' \in opt_{-i}(\pi)$, $r' \not\prec r$ (*absurdum*). \square

Corollary 1. *Assume that i makes an admissible request ($req_i \cap adm(pol_0) \neq \emptyset$). Then $\bigcup_{r \in pol_0} r$ can be disclosed to provider i iff the following conditions hold: (i) $pol_0 = \{r\}$, and (ii) $r \in req_j$ for all providers $j \neq i$ such that $req_j \cap adm(pol_0) \neq \emptyset$.*

Proof. Since the client releases exactly one element in pol_0 , the only way in which the client could release $\bigcup_{r \in pol_0} r$ is that pol_0 contains exactly one set r . Then, the proof follows easily from Theorem 3. \square

In the following results, we need two auxiliary relations: for all $\theta, \theta' \in \Theta$ let $\theta \subseteq \theta'$ iff $\forall r' \in \theta' \exists r \in \theta : r \subseteq r'$. Note that for all providers i , it holds $pol_i \subseteq req_i$. Similarly, let $\theta \leq \theta'$ iff $\forall r' \in \theta' \exists r \in \theta : r \leq r'$.

Lemma 2. *For all $r \in req_i \cap opt(\pi)$ there exists $r' \subseteq r$ such that $r' \in pol_i \cap opt(\pi[i \leftarrow pol_i])$.*

Proof. Let $r \in req_i \cap opt(\pi)$. Since $pol_i \subseteq req_i$, there exists $r' \in pol_i$ s.t. $r' \subseteq r$; moreover $r \in opt(\pi)$ implies $r \in adm(pol_0)$ and then $r' \in adm(pol_0)$. Assume per adsurdum that there exists $r'' \in opt(\pi[i \leftarrow pol_i])$ s.t. $r'' < r'$. Clearly, r'' cannot belong to pol_i because pol_i is by definition a threshold; therefore there exists a provider $j \neq i$ such that $r'' \in req_j \cap adm(pol_0)$, but in this case, as req_j is the same in both π and $\pi[i \leftarrow pol_i]$, we would have $r'' < r$ and hence $r \notin opt(\pi)$ against the hypothesis. Therefore, $r' \in opt(\pi[i \leftarrow pol_i])$. \square

Corollary 3. *For all $\pi \in \Pi$ and providers i , we have $opt(\pi[i \leftarrow pol_i]) \leq opt(\pi)$.*

Corollary 4. *For all $\pi \in \Pi$, if $i \in cw(\pi)$ then $i \in cw(\pi[i \leftarrow pol_i])$.*

Lemma 3. *For all $\pi \in \Pi$, if $i \in cw(\pi)$ then $cw(\pi[i \leftarrow pol_i]) \subseteq cw(\pi)$.*

Proof. It suffices to show that for all servers $j \neq i$, $j \in cw(\pi[i \leftarrow pol_i])$ implies $j \in cw(\pi)$. Let r be in $req_j \cap opt(\pi[i \leftarrow pol_i])$. By definition, $r \in adm(pol_0)$. Furthermore, because of opt minimality, for all $r' \in opt(\pi[i \leftarrow pol_i])$, we have $r' \not\leq r$. Since $opt(\pi[i \leftarrow pol_i]) \leq opt(\pi)$ (Corollary 3), for all $r'' \in opt(\pi)$, $r'' \not\leq r$. Therefore, $r \in req_j \cap opt(\pi)$. \square

Lemma 4. *Let $pol, req \in \Theta$. If $pol \subseteq req$ then $ful(req) \subseteq ful(pol)$.*

Proof. It suffice to show that $\forall r \in ful(req)$ it holds that $r \in ful(pol)$. Let $r \in ful(req)$. By definition, there exists $x \in req$ such that $x \subseteq r$. Since $pol \subseteq req$, it follows that there exist $y \in pol$ such that $y \subseteq x$. By transitivity, $y \subseteq r$ and, thus $r \in ful(pol)$. \square

Lemma 5. *For all $\pi \in \Pi$ and providers i it holds that $V(\pi, i) \subseteq V(\pi[i \leftarrow pol_i], i)$.*

Proof. Let us denote by $X = \{r \subseteq C \mid r \in adm(pol_0) \wedge \forall r' \in opt_{-i}(\pi). r' \not\leq r\}$. We notice that X is independent from pol_i and req_i . Given the above definition we can write $V(\pi, i) = ful(req_i) \cap X$. Similarly, it is possible to write $V(\pi[i \leftarrow pol_i], i) = ful(pol_i) \cap X$. Since, by Lemma 4, $ful(req_i) \subseteq ful(pol_i)$, it follows that $V(\pi, i) \subseteq V(\pi[i \leftarrow pol_i], i)$. \square

Lemma 6. For all $\pi \in \Pi$ and providers i it holds that $res(\pi, i) \leq_i^* res(\pi[i \leftarrow pol_i], i)$.

Proof. We observe that, by definition of $res(\pi, i)$, the thesis of this lemma can be written as $\max_{\leq_i} V(\pi, i) \leq_i^* \max_{\leq_i} V(\pi[i \leftarrow pol_i], i)$. If $res(\pi, i) \neq res(\pi[i \leftarrow pol_i], i)$, we need to show that both conditions in Definition 1 are met, that is:

1. For each $x \in \max_{\leq_i} V(\pi, i)$ there exists $y \in \max_{\leq_i} V(\pi[i \leftarrow pol_i], i)$ such that $x \leq_i y$. Let $x \in \max_{\leq_i} V(\pi, i)$.
The following cases may arise.
 - $x \in \max_{\leq_i} V(\pi[i \leftarrow pol_i], i)$: the property holds by reflexivity of relation $\leq_i \in \{\leq, \subseteq\}$.
 - $x \notin \max_{\leq_i} V(\pi[i \leftarrow pol_i], i)$: Notice that $x \in V(\pi, i)$ and, by Lemma 5, $x \in V(\pi[i \leftarrow pol_i], i)$. Thus, there exists $y \in \max_{\leq_i} V(\pi[i \leftarrow pol_i], i)$, such that $x <_i y$.
2. Assume there exist $x \in \max_{\leq_i} V(\pi, i)$ and $y \in \max_{\leq_i} V(\pi[i \leftarrow pol_i], i)$ such that $y <_i x$. Since $x \in \max_{\leq_i} V(\pi, i)$, then $x \in V(\pi, i)$. By Lemma 5, $x \in V(\pi[i \leftarrow pol_i], i)$ that implies $y \notin V(\pi[i \leftarrow pol_i], i)$.

□

Theorem 4. For all $pol \in \Theta$ and all providers i , the unique dominant attraction strategy for i w.r.t. pol is pol itself.

Proof. First we prove that $pol_i = pol$ is a dominant attraction strategy (membership), then that for all strategies $req_i \neq pol_i$ there exists a full profile $\pi \in \Pi_{pol}^i$ such that $pw(\pi[i \leftarrow req_i], i) < pw(\pi, i)$, therefore req_i is not a dominant attraction strategy (uniqueness).

(Membership) Let π be a full profile. By the (contrapositive of) Corollary 4 and Lemma 3, it is straightforward to see that $pw(\pi, i) \leq pw(\pi[i \leftarrow pol_i], i)$. By Lemma 6, $res(\pi, i) \leq_i^* res(\pi[i \leftarrow pol_i], i)$ always holds, therefore in particular when $pw(\pi, i) = pw(\pi[i \leftarrow pol_i], i)$.

(Uniqueness) Consider $req_i \neq pol_i$, for some $r \in req_i$ and $r' \in pol_i$, $r' \subset r$. Choose $\pi \in \Pi_{pol}^i$ such that (i) $pol_0 = \{r'\}$; (ii) for two providers $i \neq j$, the requests of i and j in π are $\{r'\}$; (iii) for all the other providers $k \neq i, j$, it holds $r' \notin req_k$. Clearly, as $r' \subset r$, $r \notin adm(pol_0)$ and since req_i is a threshold, for the other $r'' (\neq r) \in req_i$, $r'' \notin adm(pol_0)$. This implies that $opt(\pi[i \leftarrow req_i]) \cap req_i = \emptyset$ and hence $pw(\pi[i \leftarrow req_i], i) = 0$. On the contrary, due to (iii), $\{r'\} = opt(\pi)$ and hence $pw(\pi, i) > 0$. □

Theorem 5. For all $pol \in \Theta$ and all providers i , the unique dominant investigation strategy for i w.r.t. pol is pol itself.

Proof. (Membership) Let π be a full profile. By Lemma 6, $res(\pi, i) \leq_i^* res(\pi[i \leftarrow pol_i], i)$ always holds. By the (contrapositive of) Corollary 4 and Lemma 3, it follows that $pw(\pi, i) \leq pw(\pi[i \leftarrow pol_i], i)$ always holds, in particular when $res(\pi, i) = res(\pi[i \leftarrow pol_i], i)$.

(Uniqueness) Assume a request req_i for the server i such that $req_i \neq pol_i$ and choose a full profile π as in Theorem 4. Since $opt(\pi[i \leftarrow req_i]) \cap req_i = \emptyset$, then $res(\pi[i \leftarrow req_i], i) = \emptyset$, whereas from $i \in cw(\pi)$ and Lemma 1, $res(\pi, i) \neq \emptyset$. Therefore, $res(\pi[i \leftarrow req_i], i) <_i^* res(\pi, i)$. \square

Theorem 6. *If all providers have the same policy pol and provider i is truthful, then for all $j \neq i$ it holds $res(\pi, j) \subseteq pol$.*

Proof. Since $res(\pi, j) = \max_{\leq_j} V(\pi, j)$, it is sufficient to show that $V(\pi, j) \subseteq pol$. We preliminarily observe that since for every j , $pol \subseteq req_j$, the truthfulness of player i (i.e., $req_i = pol$) implies $opt_{-j}(\pi) = opt(\pi) \subseteq pol \cap adm(pol_0)$.

By definition, if $r \in V(\pi, j)$ the following hold: (a) $r \in adm(pol_0)$, (b) $r \in ful(req_j)$ and (c), for every $r' \in opt_{-j}(\pi)$, $r' \not< r$.

On the other hand, since $opt_{-j}(\pi) = opt(\pi)$ is a subset of the policy pol , any request that strictly larger than an element of pol cannot belong to $V(\pi, j)$. Formally, assume by contradiction that there exists $r \in V(\pi, j)$ such that $r \in (ful(req_j) \cap adm(pol_0)) \setminus pol$. Since $opt_{-j}(\pi) \subseteq pol$, $r \notin opt_{-j}(\pi)$ and there exists $r' \in opt_{-j}(\pi)$ such that $r' < r$, that implies $r \notin V(\pi, j)$. Hence, we conclude that for all $r \in res(\pi, j) \subseteq V(\pi, j)$ it holds $r \in pol$. \square