

An Integrated Formal Approach to Usage Control

P. A. Bonatti, L. Sauro, M. Faella, C. Galdi
 Dep. of Electrical Engineering and Information Technology
 Università di Napoli Federico II
 Naples, Italy
 Email: piero.bonatti@unina.it

Abstract—Usage control enforcement is currently voluntary, due to a number of technical difficulties that cannot be addressed by means of purely cryptographic techniques. So, it is commonly argued that purely technical measures should be complemented by surveillance activities and sanctions prescribed by law. The effectiveness of such measures can—and should—be formally analyzed through game theoretic techniques. This paper introduces a framework that integrates both cryptographic techniques and a market model. We provide a first formal analysis of a simplified instance of the framework; it illustrates the integrated methodology and its potential applications, and raises some warnings about the effectiveness of naive approaches.

I. INTRODUCTION

The research works on usage control policies and enforcement have produced an impressive range of expressive policy models and languages, such as [1], [2], [3], [4]. Compliance relies on sophisticated distributed architectures, e.g. [5], [6]. However, compliance is voluntary and enforcement mechanisms can be easily bypassed.

The reason is that cryptographic techniques are well suited to proving that digital resources have been dealt with properly, but they fall short in preventing information from being extracted from its cryptographic envelope and manipulated freely, without any further check.

It is commonly argued that the burden of verifying that similar misuse does not occur should be delegated to authorities and supported by laws, so that abuse can be punished, thereby discouraging the violation of usage control policies. Accordingly, the P3P initiative involved data protection authorities in the standardization and implementation process [7].

However, in the literature we find no formal evidence that a similar approach creates adequate incentives to policy compliance. If good behavior is to be induced through economic sanctions, then a formal analysis can—and should—be based on game-theoretic techniques such as those applied in economics to analyze and regulate markets. This kind of games has already been considered in contexts related to privacy; for example, bargaining techniques have been applied to improve privacy in negotiations involving anonymity and reputation issues [8], and a generalized form of procurement auctions has been proved to moderate—and sometimes minimize—the collection of user profile information [9]. Further game theoretic approaches like [11], [12], [13] focus on the costs

of compliance verification and do not analyze the other issues that may hinder the practical application of measures based on investigations and sanctions.

An integrated approach to usage control, involving both strong cryptographic guarantees and incentive systems, needs to adopt a visionary perspective. Today's authorities and laws are not (yet) equipped for dealing effectively with illegal information markets. In looking for effective and viable solutions, researchers should start by freely imagining possible ecosystems and societies in order to synthesize an effective combination of incentives and technical measures against information misuse, that can be brought to the attention of legislators.

In this paper we adopt a similar perspective. We imagine that information must be kept by law in a suitable format, enriched with sticky policies and provenance information, so that behavior can be reliably proved to be compliant by honest agents. Moreover, as a deterrent to abuse, we assume that some surveillance authority is allowed to inspect the contents of internet connections and digital storage, looking for data items that are exchanged and stored without complying with format requirements or sticky policies.

These assumptions are cast into a mathematical model of an illegal information market, that enables a first estimate of adequate sanctions, and a first tentative assessment of several kinds of surveillance and investigation measures, such as spoofing, sniffing, and storage inspection by means of computer forensics techniques.

More precisely, we show which measures have stronger effects, and estimate target performance values that must be reached in order to destroy the illegal market. Such values can be a useful input to research on computer forensics techniques and monitoring. We will argue that stronger technological support to inspections is needed in order to create effective incentives to privacy-preserving behavior.

The paper is organized as follows: In Sec. II, we introduce the information format required by law in our hypothetical world. The model is abstract and parametric w.r.t. the usage control policy language adopted. In Sec. III, we model the illegal information market and discuss its parameters, including sanctions and the success rate of surveillance measures. In Sec. IV, the model is extended to handle data deletion deadlines. Finally, in Sec. V, we summarize our observations and discuss some interesting directions for further investigation.

This work is partially supported by the National Project *Security Horizons*.

In this first paper, we adopt a simplified market model, that allows us to import some results from game theory. While the accuracy of our estimates is obviously affected by the missing features, the model analysis methodology and the range of its possible outcomes (i.e. the core of the integrated formal approach to usage control) should be expected to scale to more sophisticated frameworks. The limitations of the current model will be discussed in Sec. V.

II. PROVING HONESTY

The agents that comply to a usage policy should be able to prove it reliably by means of cryptographic techniques. Here, as an example, we adopt a simple possible approach to checking release to third parties, based on provenance chains and sticky policies. As in most approaches, we focus on how to prove compliance, rather than supporting investigations on violations; the limitations of this approach will be discussed later on. The approach is parameterized w.r.t. the data to be manipulated and the policy language adopted.

Let \mathcal{D} be a given set of data items, such as user profiles or multimedia contents. Data usage is regulated by policies; let \mathcal{P} denote the set of all possible release control policies. Finally, let \mathcal{A} be the set of agents; as usual, \mathcal{A}^* denotes the set of all finite sequences of agents $\langle a_1, \dots, a_n \rangle$. For all $d \in \mathcal{D}$, for all $\langle a_1, \dots, a_n \rangle \in \mathcal{A}^*$, and all $p \in \mathcal{P}$, we write

$$d; a_1, \dots, a_n \models p$$

iff p permits each a_i to release d to a_{i+1} ($1 \leq i < n$). Accordingly, we assume that, for all $n > 1$,

$$d; a_1, \dots, a_n \models p \text{ implies } d; a_1, \dots, a_{n-1} \models p. \quad (1)$$

The statements $d; a_1 \models p$ are assumed to be true iff a_1 collects d as specified by the policy p , and p is the usage policy published by a_1 .

Now a *checkable (information) item* is an expression

$$c = (d \mid p \mid a_1, \dots, a_n)$$

where $d \in \mathcal{D}$, $p \in \mathcal{P}$, and $\langle a_1, \dots, a_n \rangle \in \mathcal{A}^*$. The intuitive meaning of these components is the following: the data item d is subject to the sticky policy p ; moreover, d has been collected by a_1 and subsequently released by each a_i to a_{i+1} .

We say that c is *legal* iff $d; a_1, \dots, a_n \models p$, and *illegal* otherwise.

A simple implementation of the above abstract framework consists in having each agent a_{n-1} sign each new checkable item $c = (d \mid p \mid a_1, \dots, a_{n-1}, a_n)$ before releasing it to a_n . Formally, each checkable item can be recursively encoded as follows, where $[\sigma]_a$ denotes the string σ signed by agent a , and “|” denotes string concatenation:

$$\text{enc}((d \mid p \mid a_1, a_2)) = [d \mid p \mid a_1 \mid a_2]_{a_1} \quad (2)$$

$$\text{enc}((d \mid p \mid \alpha, a_{n-1}, a_n)) = [\text{enc}((d \mid p \mid \alpha, a_{n-1})) \mid a_n]_{a_{n-1}} \quad (3)$$

This encoding makes it possible to verify cryptographically that:

- 1) p has been associated to d by d 's collector, a_1 ;
- 2) p permits each a_i to release d to a_{i+1} ;
- 3) d and p passed through all a_i ($i < n$).

So, for example, if d is the profile of Ann, and Ann finds out that an unexpected recipient a_n eventually received d , then using $\text{enc}((d \mid p \mid a_1, \dots, a_n))$ Ann can check whether the agent a_1 she interacted with associated the right disclosure policy p to her profile, and whether each a_i was authorized by p to release d to a_{i+1} .

On the other hand, $(d \mid p \mid a_1, \dots, a_n)$ is not signed by a_n and there is no cryptographic evidence that it has been received by a_n . Strictly speaking, there is no evidence that each a_{i-1} directly released d to a_i , either. Finally, nothing prevents a_n to extract d from its cryptographic envelope and distribute it as raw, unprotected data. In the next section we investigate these issues from a market perspective.

The above encoding can be further enhanced so as to guarantee that the owner a_0 of d permits a_1 to collect d and use it according to policy p . For this purpose, (2) shall be replaced with:

$$\text{enc}((d \mid p \mid a_1, a_2)) = [[d \mid p \mid a_1]_{a_0} \mid a_2]_{a_1}. \quad (4)$$

The design of a suitable protocol for the initial agreement of a_0 and a_1 lies beyond the scope of this paper.

III. CRIME AND PUNISHMENT

In order to encourage the intended usage of checkable items, a legislator may want to regulate the collection, disclosure, and storage of data. We start by formalizing what an agent can legally acquire and store.

Each agent a is associated to a set \mathcal{L}_a of checkable items called *legal for a*. \mathcal{L}_a shall satisfy the following axiom:

$$(d \mid p \mid \alpha, a_n) \in \mathcal{L}_a \Leftrightarrow a = a_n \wedge d; \alpha, a_n \models p. \quad (5)$$

Axiom (5) says that a checkable item is legal for a if it is addressed to a , and it satisfies the sticky policy p . By (1), this implies that if the chain of disclosures contains any forbidden step, then a should refrain from accepting the item. Formally, we have: $(d \mid p \mid \alpha, a') \notin \mathcal{L}_{a'} \Rightarrow (d \mid p \mid \alpha, a', a) \notin \mathcal{L}_a$.

Now we can formulate the norms that each agent should obey:

- N1** Agents should not accept, store, or disclose any data item d in unprotected form (i.e. agents should only store and exchange *checkable* items).
- N2** An agent a should accept and store only *legal* checkable items $c \in \mathcal{L}_a$.

The main questions now are:

- How should compliance to the above rules be verified? Which powers should investigators have? How frequently and extensively should agents be inspected?
- In case of violations, which sanctions should be applied?

Here we make a first step towards answering the above questions by formalizing a simple illegal market (where illegal data are exchanged), and analyzing the measures that would make it unprofitable—which are the only serious deterrent to illegal information exchange.

Let the set of agents \mathcal{A} be partitioned into a set C of data collectors and a set R of data recipients. Each collector $a \in C$ has a distinct set of hosts H_a . The information market is represented by a labelled bipartite graph whose nodes are hosts and recipients. Each edge is directed from a host i to a recipient j , and is labelled by a cost c_{ij} representing a quantitative estimate of the risks incurred by i 's owner in releasing illegal information to j from i . Moreover, each recipient $j \in R$ is labelled by a value u_j that formalizes the recipient's expected utility in purchasing illegal information. The utility u_i encompasses also the risks incurred by j .

Each possible transaction between some i and j has a corresponding price π_{ij} . The *profit* for i 's owner is $\pi_{ij} - c_{ij}$, while the profit for j is $u_j - \pi_{ij}$.

To get a finer grained granularity over data items (if needed), one may simply assume that there are multiple instances of the graph, one for each data item (e.g. one graph for each user profile).

This model is isomorphic to the *facility location pricing game* by A. Vetta [10, Ch. 19]. It is a competitive game where each agent acts so as to maximize its own profit. From a strategic perspective, it can be described as follows: (i) each data collector $a \in C$ selects for each recipient $j \in R$ at most one host $i \in H_a$ for disclosing data to j ; (ii) data collectors establish the prices π_{ij} ; (iii) each recipient j choose a provider i and pays π_{ij} . In the following let S be the set of edges (i, j) selected by data collectors in phase (i).

It can be seen that in the stable equilibria of this market, each provider a chooses for each recipient j the host i with minimal cost c_{ij} , and sets the price to $\min_{b \neq a, i' \in H_b} c_{i'j}$ (as any higher price would encourage b to undercut it).

Consequently, the *social welfare* $V(S)$ (which is defined as the sum of each agent's profit) turns out to be:

$$V(S) = \sum_{j \in R} \left(u_j - \min_{(i,j) \in S} c_{ij} \right). \quad (6)$$

Clearly, the illegal information market can be suffocated by pushing $V(S)$ down to 0; by (6), this amounts to ensure that

$$u_j \leq c_{ij} \quad (7)$$

for all hosts $i \in \bigcup_{a \in C} H_a$ and all $j \in R$. The authority can influence this condition by raising risks and sanctions. In order to get into the details, we first introduce a finer grained model of costs and utilities that take the actions of the authority into account.

The expected risks c_{ij} can be defined as follows:

$$c_{ij} = pp_{ij} \cdot s, \quad (8)$$

where pp_{ij} is the probability of being punished and s is the sanction for the distributors of illegal data. In turn, pp_{ij} can be further articulated as:

$$pp_{ij} = ps_{ij} + (pi_j \cdot pe_{ij}), \quad (9)$$

where ps_{ij} is the probability that the messages from i to j are sniffed and decrypted, pi_j is the probability that j is inspected by the investigation authority (and the illegal data item is actually found), and pe_{ij} is the probability that j provides

enough valid evidence to identify the data collector that owns i .

The expected utility u_j of a recipient $j \in R$ can be defined as follows:

$$u_j = \hat{u}_j - (ps_{ij} + pi_j) \cdot s', \quad (10)$$

where \hat{u}_j is the utility of using the purchased data item, s' is the punishment for recipients, and ps_{ij} and pi_j are defined as in (9).

With the above definitions, requirement (7) becomes:

$$\hat{u}_j \leq ps_{ij}(s + s') + pi_j(pe_{ij} \cdot s + s') \stackrel{\text{def}}{=} f. \quad (11)$$

We call f *surveillance factor*.

If $s = s'$, then in order to obtain (7) one should set

$$s \geq \frac{\hat{u}_j}{2ps_{ij} + pi_j(1 + pe_{ij})}. \quad (12)$$

The probabilities ps_{ij} and pi_j are likely to be small, given the large number of connections and hosts that should be monitored and inspected, and the related technical difficulties. This means that s may be several orders of magnitude larger than \hat{u}_j . If \hat{u}_j is not much smaller than the expected damage to the data owner resulting from illegal disclosure, then this may constitute an obstacle to legislators, that might find s out of proportion to the violation: a very high value for s might be deemed inappropriate.

If $s \neq s'$, then by increasing the sanction for recipients one can make the surveillance factor f approach \hat{u}_j (cf. (11)) faster than by increasing the sanction for collectors, because:

$$\frac{\partial f}{\partial s} = ps_{ij} + pi_j \cdot pe_{ij} \leq ps_{ij} + pi_j = \frac{\partial f}{\partial s'}. \quad (13)$$

Similarly, we can compare the relative usefulness of increasing each of the probabilities ps_{ij} , pi_j , and pe_{ij} . We have:

$$\frac{\partial f}{\partial ps_{ij}} = s + s' \quad (14)$$

$$\frac{\partial f}{\partial pi_j} = pe_{ij} \cdot s + s' \quad (15)$$

$$\frac{\partial f}{\partial pe_{ij}} = pi_j \cdot s \quad (16)$$

Clearly, $\frac{\partial f}{\partial ps_{ij}} \geq \frac{\partial f}{\partial pi_j}$ and $\frac{\partial f}{\partial ps_{ij}} \geq \frac{\partial f}{\partial pe_{ij}}$. From a purely mathematical perspective, this means that improving the performance of message sniffing is the fastest way of reducing the profitability of the illegal market.

Moreover, under the reasonable assumption that s and s' have the same order of magnitude and $s' \gg pi_j$, we have that $\frac{\partial f}{\partial pi_j} \geq \frac{\partial f}{\partial pe_{ij}}$. This means that—from a mathematical perspective—the second best way of reducing the welfare of the illegal market is improving the effectiveness of the inspections over recipients.

However, from a technical perspective, increasing pi_j may be significantly easier than increasing ps_{ij} . Just as an example, in inspecting the storage of an organization one may force it to decode encrypted data by law, while the encryption of internet connections typically makes use of session keys that are lost immediately after the connection is closed.

Further difficulties arise when hosts and agents are located in countries that are not subject to the enforcing authority. If recipients cannot be inspected, then $pi_j = 0$, and (11) becomes:

$$\hat{u}_j \leq ps_{ij}(s + s'). \quad (17)$$

In this case, message sniffing and sanctions are the only possible parameters that can be influenced by the enforcement authority.

Finally, let us analyze the benefits of watermarking in this particular market. The effects of watermarking can be modelled by assuming that $pe_{ij} = 1$, that is, once an illegal data item is discovered, the agent that disclosed it can be reliably identified. By (14) and (15), we have that $pe_{ij} = 1$ implies $\frac{\partial f}{\partial ps_{ij}} = \frac{\partial f}{\partial pi_j}$, that is, the improvement of recipient inspection becomes as effective as the improvement of message analysis. Under the same hypothesis, by (13), we have that the sanctions s and s' have the same influence on f , i.e. $\frac{\partial f}{\partial s} = \frac{\partial f}{\partial s'}$.

IV. BOUNDED PERSISTENCY

The framework discussed so far deals only with third party disclosures. In this section, we consider bounded persistency, that is, deletion policies such as “delete d within 30 days”.

First, checkable items should be extended with timestamps t_i that record when each disclosure happened:

$$c = (d \mid p \mid a_1, \dots, a_n \mid t_1, \dots, t_n). \quad (18)$$

Policy semantics should be extended accordingly. We write $d; a_1, \dots, a_n; t_1, \dots, t_n; t \models p$ if all disclosures are legal and the item can still be kept in a_n 's storage at time t . Legal disclosures satisfy all the properties required in the previous section and, moreover, have the property that no a_i releases d to a_{i+1} after the deletion time prescribed by p (in other words $t_{i+1} - t_i$ should not exceed the persistency bound specified by p on a_i).

Now the set of legal items depend also on the current time. Accordingly the set of legal items for a at time t , in symbols \mathcal{L}_{at} , satisfies the axiom:

$$(d \mid p \mid \alpha, a_n \mid \tau) \in \mathcal{L}_{at} \Leftrightarrow a = a_n \wedge d; \alpha, a_n; \tau; t \models p. \quad (19)$$

Norm **N2** is replaced with:

N2' An agent a should accept at time t only checkable items $c \in \mathcal{L}_{at}$ such that $t_n = t$.

N3 An agent a should store at time t only checkable items $c \in \mathcal{L}_{at}$ such that $t_n \leq t$.

If **N2'** is violated, then the message sent from a_{n-1} to a_n does not satisfy the norms. Both the sender and the recipient are guilty and the analysis of (7) is the same as in the previous section.

If **N3** is violated, instead, then only the recipient is guilty. The sender has no responsibility and runs no risks, so $c_{ij} = 0$. The utility of the recipient is

$$u_j = \hat{u}_j - pi_j \cdot s'' \quad (20)$$

where s'' is the sanction for storing data beyond the prescribed limit. It is easy to see that the surveillance factor becomes $f = pi_j \cdot s''$ and that the sanction should satisfy

$$s'' \geq \frac{\hat{u}_j}{pi_j} \quad (21)$$

in order to be effective. Note that, by (11),

$$s + s' \geq \frac{\hat{u}_j}{ps_{ij} + pi_j}. \quad (22)$$

This means that if the sanctions s , s' , and s'' are set to the minimal possible effective values, then $s'' \geq s + s'$, that is, the sanction for persistency violations is higher than the cumulative sanctions for the two agents involved in an illegal disclosure.

V. CONCLUSION

This work is meant to be a first step towards a formal investigation of the interplay between laws, surveillance and cryptography. It raises questions about the hypothesis that checking legal information usage can be simply delegated to some authority without any more specific support from privacy technologies (besides typical proofs of compliance).

We defined a framework where agents are required by law to keep a sticky policy and provenance attached to data items. Such checkable items (whose integrity and reliability are guaranteed by a suitable cryptographic encoding) make it possible for honest agents to prove that they received data items through legal channels only, and that they are satisfying the time bounds on persistency. Dishonest agents are discouraged through surveillance and sanctions. The influence of such enforcement actions are modelled mathematically by a surveillance factor f .

As a first step, we analyzed the effectiveness of enforcement actions in a simplified illegal market where agents are partitioned into data collectors and data recipients. We make the assumption that law permits an investigation authority to inspect the contents of internet connections and private storage, possibly requiring the involved agents to decrypt their data. These two investigation actions (i.e. message sniffing and storage inspection) are modelled in terms of their probability of success. Moreover, we assumed that recipients—when found guilty—may be induced to reveal their illegal information providers. Again, the effectiveness of such measures (that may be affected by lack of valid forensic evidence) are modelled probabilistically. Interestingly, sanctions depend on the kind of violation: those for violating storage limitations are higher than the sanctions for illegal third-party disclosures.

Our model raises the concern that effective sanctions might turn out to be too severe to be acceptable for a standard legal system. In order to verify or disprove this conjecture a quantitative model of negative externalities (i.e. the effects of violations on the principals whom the data refer to) would be needed to check that sanctions are proportioned to violations.

The surveillance factor can be improved also by increasing the performance of monitoring and surveillance. The mathematical analysis of the influence of investigation measures on the surveillance factor, and informal considerations on their technical feasibility suggest that it may be profitable

to focus on the monitoring of data recipients and improve related forensic techniques. However, this would probably induce data recipients to move their hosts to countries where the investigation authority cannot operate.

Improving the frequency and accuracy of surveillance and monitoring may constitute a source of concern, too, due to the invasive nature of these operations that might introduce new kinds of abuse. In order to induce users, providers, and companies to improve their data usage practices, some authority has to be given the power of investigating private information in depth. If such activity could be started only in the presence of violation evidence, then the effectiveness of random inspections would be lost, and the deterrent effect of sanctions weakened.

Further, related difficulties arise from surveillance costs, that are not currently modelled in our framework. However, in the literature several approaches can be found, that aim at good tradeoffs between inspection effectiveness and their costs, e.g. [11], [12], [13].

These observations raise a natural question: can technological solutions such as cryptographic techniques help the inspection game and make it more efficient, so as to mitigate the difficulties outlined above? So far, little has been done in this direction; watermarking is one of such mechanisms.

We found out that in the simplified market, watermarking brings limited benefits, probably because information is not further disclosed by recipients. On the positive side, the derivative of the surveillance factor w.r.t. recipient inspection performance increases.

The analysis carried out in this paper needs to be generalized and extended along several possible directions. First of all, usage policies in general constrain a number of aspects beyond third party disclosures and deletion deadlines. Implicitly, restrictions on what data can be collected are already modelled (through the statements $d; a_1 \models p$). Restrictions on purpose can be regarded as straightforward refinements of the definition of legal recipients. However, the current model is not “dynamic enough” to deal with policies such as counting restrictions (e.g. “copy at most n times”).

Another important extension concerns disclosure chains. Recipients may further distribute the illegal information they purchase. The effects are not yet clear: on the one hand information sales create profit; on the other hand, risks increase. We expect watermarking to have stronger effects in this setting, as indirect distribution appears to raise the risks associated to the first illegal disclosure. However, we expect an ad hoc watermarking method to be necessary.

As an additional difficulty, we remark that illegal markets can hardly be “designed” and regulated. The kind of influence that an authority can exercise on such markets is largely external, which limits the amount of mechanism design techniques that can be applied to induce desired behavior. A deeper understanding of how forensics techniques and cryptoanalysis may support game theoretic approaches is still needed.

The warnings raised in this paper should be regarded as a first result of our formalization effort. As such, they do not make the integrated approach uninteresting; still they highlight

some potential dead ends, and motivate the investigation of other games, and new cryptographic solutions conceived to support those games, rather than being mainly focussed on proving compliance.

REFERENCES

- [1] J. Park and R. S. Sandhu, “The $UCON_{ABC}$ usage control model,” *ACM Trans. Inf. Syst. Secur.*, vol. 7, no. 1, pp. 128–174, 2004.
- [2] L. F. Cranor, *Web privacy with P3P - the platform for privacy preferences*. O’Reilly, 2002.
- [3] M. Hilty, A. Pretschner, D. A. Basin, C. Schaefer, and T. Walter, “A policy language for distributed usage control,” in *ESORICS 2007*, ser. Lecture Notes in Computer Science, J. Biskup and J. Lopez, Eds., vol. 4734. Springer, 2007, pp. 531–546.
- [4] G. Governatori and R. Iannella, “A modelling and reasoning framework for social networks policies,” *Enterprise IS*, vol. 5, no. 1, pp. 145–167, 2011.
- [5] F. Kelbert and A. Pretschner, “Towards a policy enforcement infrastructure for distributed usage control,” in *Proceedings of the 17th ACM symposium on Access Control Models and Technologies*, ser. SACMAT ’12. New York, NY, USA: ACM, 2012, pp. 119–122. [Online]. Available: <http://doi.acm.org/10.1145/2295136.2295159>
- [6] S. Trabelsi, J. Sendor, and S. Reinicke, “PPL: Primelife privacy policy engine,” in *IEEE POLICY 2011*. IEEE Computer Society, 2011, pp. 184–185.
- [7] L. F. Cranor, “The role of privacy advocates and data protection authorities in the design and deployment of the platform for privacy preferences,” Remarks for “The Promise of Privacy Enhancing Technologies” panel at the Twelfth Conference on Computers, Freedom and Privacy, 2002, <http://lorrie.cranor.org/pubs/p3p-cfp2002.html>.
- [8] L. F. Cranor, P. Resnick, and P. Resnick, “Protocols for automated negotiations with buyer anonymity and seller reputations,” *Netnomics*, vol. 2, no. 1, pp. 1–23, 2000.
- [9] P. A. Bonatti, M. Faella, C. Galdi, and L. Sauro, “Towards a mechanism for incentivating privacy,” in *ESORICS*, ser. Lecture Notes in Computer Science, V. Atluri and C. Díaz, Eds., vol. 6879. Springer, 2011, pp. 472–488.
- [10] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007.
- [11] X. Zhao and M. E. Johnson, “Access governance: Flexibility with escalation and audit,” in *HICSS*. IEEE Computer Society, 2010, pp. 1–13.
- [12] J. Blocki, N. Christin, A. Datta, and A. Sinha, “Regret minimizing audits: A learning-theoretic basis for privacy protection,” in *CSF*. IEEE Computer Society, 2011, pp. 312–327.
- [13] —, “Audit mechanisms for provable risk management and accountable data governance,” in *GameSec*, ser. Lecture Notes in Computer Science, J. Grossklags and J. C. Walrand, Eds., vol. 7638. Springer, 2012, pp. 38–59.