On the Value of Private Information

Jon Kleinberg
Cornell University
kleinber@cs.cornell.edu

and

Christos H. Papadimitriou
U.C. Berkeley
christos@cs.berkeley.edu

and

Prabhakar Raghavan
Verity
pragh@verity.com

Abstract

As individuals increasingly take advantage of on-line services, the value of the private information they possess emerges as a problem of fundamental concern. We believe that the principle underlying privacy should be simple: Individuals are entitled to control the dissemination of private information, disclosing it as part of a transaction only when they are fairly compensated. We show how this principle can be made precise in several diverse settings — in the use of marketing surveys by a vendor designing a product; and in the design of collaborative filtering and recommendation systems, where we seek to quantify the value of each user's participation. Our approach draws on the analysis of coalitional games, making use of the core and Shapley value of such games as fair allocation principles.

1 Introduction

Even though privacy is arguably one of the most important problems — indeed, missions — facing Computer Science today, a number of its most crucial facets have remained largely unexplored. Privacy issues form a wide spectrum that ranges from the safeguarding of highly secret data to the concerns of Internet users about giving away their personal e-mail addresses to on-line merchants. While the former end of the spectrum — providing unambiguous protection for secrets — forms the domain of cryptography and information security, the latter end raises a qualitatively different set of issues. Indeed, it forces us to quantify the circumstances under which it is worthwhile for an individual to voluntarily relinquish private information — and the extent to which he or she should be compensated for doing so — and it is this question which has received very little attention in theoretical (or practical) research.

One of the reasons for this lack of research is that privacy is notoriously hard to define and capture rigorously. As Hal Varian points out [10], when asked how important privacy is to them, people are likely to answer "*extremely* important," but then be at a loss when asked to define privacy. At best, they come up with answers boiling down to different things: To some privacy means "don't interrupt my dinner," to some "don't embarrass me," while to others "don't overcharge me."[1] In today's Internet "privacy" usually means a rough commitment by companies to divulge to others only "anonymized aggregates" (terms never defined precisely) of one's information [11]; we argue below that this is at best misguided and limited. There has recently been interest by some economists (see, for example, [10, 4, 1]) and legal scholars [7, 8] in exploring theoretically the nature of privacy and developing a theory of the subject; the present paper is a contribution to this fledgling but important area.

We believe that the principle underlying privacy should be simple: Individuals are entitled to control the dissemination of private information, where we can restrict the latter term to mean all records of the individual's participation in transactions. (Thus, for example, the present paper began as a conference submission, disseminated to co-participants in the submission process, specifically, TARK's program committee chair.) Privacy is a problem today because such private information is not controlled by individuals who produced it by participating in transactions, or by the agents and transaction co-participants to whom such control has been yielded explicitly, but by entities that happen to be in opportunistic possession of such information, such possession being usually an unintended byproduct of the transaction. Decisions about further dissemination of such information are made by these entities without consultation and appropriate compensation (or charging, see below) of the rightful owner of the information. We believe that it is this *"externality"* that lies at the root of the privacy problem. Private information in today's world (and Internet) is intellectual property (see [7]) that typically bears negative royalty (such as dinner interruptions and spam).

We do not believe that such dissemination is always detrimental to the individual's interests and should be prohibited. If you just bought a PC from a merchant, it may be to your interest that printer manufacturers know that you now need a printer, since such disclosure would reduce "search costs" [10]; it is in fact to your advantage for such information to be transmitted unaggregated, and you could even agree to pay more for your PC if such

[1] It is in fact debatable whether privacy is a rational concern that can be studied mathematically, or a primordial, deeply seated human value akin to religion and freedom that is beyond such study. Obviously, by writing this paper we side with the former point of view.

disclosure service were included (see [1] for an extensive theoretical treatment that supports this possibility). But you may not be eager for the merchant to also announce your credit information or your history of previous purchases (even if it is aggregated with sales information of other clients), or for your university's sponsored research office to divulge to printer manufacturers that your grant's budget for a printer is $500, for example. Such disclosures should be made only after an explicit agreement and possibly appropriate compensation. In a similar vein, services such as search engines, marketing surveys and recommendation systems promise to use private information obtained only for a user's "own good." Such claims should actually be proved, and in fact users who benefit more should compensate users who contributed useful information but benefited less.

In this paper we propose that in certain situations it is possible to precisely quantify these notions, by modeling the value of private information. We give several simplified and stylized examples in which the value of information supplied by individuals to a transaction can be evaluated quite accurately for fair compensation, and in fact in ways that act as incentives for individuals to truthfully participate in the transaction. In our examples we treat each transaction as a coalitional game [5, 6] and consider its core and Shapley value as examples of fair allocation principles.

We consider cases in which information is collected by a single monolithic enterprise, as well as cases in which information is pooled in a distributed fashion. Our first example is an instance of the former scenario — a *marketing survey* in which individuals announce their preferences among several possible products (or catalogs, services, etc.), and exactly one such product will be adopted. The vendor's payoff increases with the size of the majority supporting the outcome, while each individual wants to have his/her own preference adopted. In this example we show that imputations in the core favor participants in the majority, and, rather surprisingly, so in general does the Shapley value (except in cases of a "very close vote" which, as is well known, never happen in real life). Indeed, when there is a clear majority in the preferences of the individuals, the Shapley value confers a vanishingly small quantity on individuals outside this majority.

A quite different setting is that of a *recommendation system*. Here, there is no single vendor; rather, individuals pool their knowledge of interesting items with the goal of learning about new items from this collective knowledge base. Predictably, the Shapley value in this game rewards "novel" contributions (the core is empty except in the uninteresting case of disjoint contributions). We also consider an extension of this to a collaborative filtering situation in which agents take recommendations about items they do not know.

2 Three Examples

Marketing Survey

Consider a vendor who is conducting a survey of n consumers regarding their preferences among k possible versions of a planned product. Assume that each consumer has a single clear preference among the k alternatives, and that it is to her own advantage if her preferred alternative is adopted. (The case in which multiple votes are allowed, or the top s winners are adopted, can be handled by an analysis completely parallel to what follows.) The vendor decides to adopt the alternative that attracts the most votes (for simplicity, we say the lower-indexed item wins in case of ties). The payoff to the vendor is equal to the size of the majority (presumably because it represents the size of the market for the item), while each consumer in the winning majority gets a payoff of one; all other consumers get a payoff of zero.[2]

To formulate this as a coalitional game [5, 6], we consider $n+1$ agents (the n consumers $\{1,\ldots,n\}$ and the vendor, named 0), and the payoff of any coalition S (subset of the set of $n+1$ agents) is determined by the situation restricted to that subset: If S contains the vendor then the payoff $v(S)$ is twice the size of the majority of the preferences of the consumers in S, and if it does not contain the vendor then $v(S)$ is zero. Our goal is to find a "fair" way of subdividing $v([n+1])$, the total payoff of the game. That is, we are going to determine when a vector $x \in \Re_+^{n+1}$ with total sum of its components equal to $x([n+1])$ (such vectors are called *imputations*) will be considered "fair." (We denote the n-dimensional simplex $\{x \in \Re_+^{n+1} : \sum_i x_i = v([n=1])\}$ by I for "imputations"; for a vector $x \in I$ and $S \subseteq [n+1]$ we denote $\sum_{i\in S} x_i$ by $x(S)$.)

Over the past six decades, a great number of such notions of "fairness," usually called *solution concepts,* have been proposed, championed, and criticized in the Economics literature (see [6] for an extensive discussion). In this paper we focus on the two that are simpler and best known: the *core* and the Shapley value. A vector $x \in I$ is in the core if for every coalition S, we have $x(S) \geq v(S)$. That is, if an imputation is in the core, no coalition has an incentive to secede in order to assure a better payoff. The core is a rather conservative notion of fairness —as a result, in many natural games it is empty. For the marketing survey game, it is not hard to see that the following holds:

Proposition 2.1 *An imputation is in the core iff all consumers not in the winning majority have* 0 *components, the vendor has payoff* $x_0 \geq 2r$*, where* r *is the size of the largest minority, and the payoffs* $y_1 \leq y_2 \leq \cdots \leq y_m$ *of the majority consumers obey* $\sum_{i=1}^{j} y_i \geq 2 \cdot j - x_0$ *for all* $j = 1,\ldots,m$.

[2] Alternately, one could think of the situation as one in which the n consumers are willing to contract the supply of the product with the vendor.

That is, the consumers in the majority can split among them an amount up to twice the margin of their "victory" —in the case of a tie they also get zero. An equal split among winners is always possible. Also, the imputation in which the vendor gobbles down the whole payoff is always in the core.

Another important solution concept is the *Shapley value*. According to the Shapley value, each agent will be awarded an amount equal to the *average contribution of this agent to the payoff of the group at the time of his or her arrival*, where the average is taken over all arrival orders of the agents. Mathematically, it is the imputation x where

$$x_i = \frac{1}{n!} \sum_{\pi \in S_n} v(S(\pi, i)) - v(S(\pi, i) - \{i\}),$$

with $S(\pi, i)$ denoting the set of all agents who, under arrival order π, arrived not later than i. The Shapley value can be shown to be the only notion of fairness compatible with three (natural yet controversial) axioms, see [5, 6].

In many settings, the Shapley value — despite its complicated appearance — can be represented as a simple closed-form expression. This does not appear to be the case here; but using large-deviation bounds [2] we can provide strong asymptotic bounds for the Shapley value. In what follows, we assume that n (the number of consumers) grows while k (the number of alternatives) is fixed (or grows much more slowly than n). We divide the n customers into k *categories* $C_0, C_1, \ldots, C_{k-1}$, each corresponding to the subset voting for a single alternative, and we suppose these are indexed so that $|C_0| \geq |C_1| \geq \cdots \geq |C_{k-1}|$. We let $m = |C_0|$ denote the size of the largest category (the majority), and $r_i = |C_i|$ for $i > 0$ denote the sizes of the remaining categories. For a constant $\epsilon > 0$, we say that there is a *clear majority with margin* ϵ if $m - r_1 \geq \epsilon n$ — that is, the majority and the largest minority categories differ by $\Omega(n)$ votes.

We now claim the following

Proposition 2.2 *The Shapley value of the vendor always satisfies $x_0 \geq m$. If for some constant $\epsilon > 0$, there is a clear majority with margin ϵ, then the consumers in the majority category C_0 have payoff $x_i = 1 - o(1)$, the consumers outside the majority category have payoff $x_i = o(1)$, and the vendor has payoff $x_0 = m(1 + o(1))$.*

The $o(1)$ terms in the above proposition are parametrized by the constant ϵ — the extent to which the majority is "clear."

We now sketch a proof of this proposition. To see that $x_0 \geq m$, observe that the expected number of members of C_0 arriving before the vendor is precisely $m/2$, and so at the point of the vendor's arrival, the payoff increases from 0 to an expected value of at least m.

Now suppose there is a clear majority with margin ϵ, let $p = m/n$, and let $q = r_1/n$; thus we have $p - q \geq \epsilon$, and in particular this implies $p \geq \epsilon$. For a permutation π of the agents, let $\overline{\pi}(a, b)$ denote the set of agents arriving between positions a and b of π. We say that a permutation π is *t-good* if for all $t' \geq t$, the largest category in $\overline{\pi}(1, t')$ is a subset of C_0. By Hoeffding's bounds [2], there is a constant c (depending on ϵ) so that if π is a random permutation, and $t' \geq c \log n$, then $|\overline{\pi}(1, t') \cap C_0| \geq (p - \epsilon/3)t'$ holds with high probability. Similarly, for $t' \geq c \log n$ and $i > 0$, we have $|\overline{\pi}(1, t') \cap C_i| \leq (q + \epsilon/3)t'$ with high probability. Taking the union bound over all C_i and all $t' \geq c \log n$ (and using the fact that $p - q \geq \epsilon$), it follows that a random permutation π is $(c \log n)$-good with high probability.

Consider any customer i, and a random permutation π. Consider the event $\mathcal{E}$ that π is $(c \log n)$-good and $i \in \overline{\pi}((c \log n), n)$; by the arguments above, $\mathcal{E}$ holds with high probability. Now, let $\mathcal{F}$ denote the event that in π, the vendor comes before customer i; the probability of $\mathcal{F}$ is $1/2$, and hence the probabilities of $\mathcal{E} \cap \mathcal{F}$ and $\mathcal{E} \cap \overline{\mathcal{F}}$ are each at least $1/2 - o(1)$. Conditioned on the event $\mathcal{E} \cap \mathcal{F}$, the contribution of i to the payoff is 2 if $i \in C_0$, and 0 otherwise; conditioned on $\mathcal{E} \cap \overline{\mathcal{F}}$, the contribution of i is 0 in all cases. Combining this with the contribution of i on permutations for which $\mathcal{E}$ does not hold, we have $x_i = 1 - o(1)$ for $i \in C_0$ and $x_i = o(1)$ for $i \notin C_0$.

It is interesting to notice that, in the absence of certainty that his or her preference is not in the majority, a consumer would have an incentive to participate in the survey under the Shapley value or any of the symmetric payoffs in the core (with the exception of the one in which all proceeds go to the vendor).

Recommendation Systems

Consider a recommendation system in which individuals pool information about interesting items (such as books; see [3] for theoretical work on recommendation systems). Suppose first that each individual likes all items, and benefits from finding out about each one of them. A simple model would have an agent i identified with a set B_i of the items he or she knows about. The payoff for a coalition $S \subseteq \{1, \ldots, n\}$ is $v(S) = |\cup_{i \in S} B_i|$: the total benefit to the coalition is the total number of items collectively known by its members.

In this case it is easy to see the following:

Proposition 2.3 *(a) The core is always empty, unless all B_i's are disjoint, in which case it is $\{(|B_1|, \ldots, |B_n|)\}$.*

(b) The Shapley value of agent i is

$$\sum_{j \in B_i} \frac{1}{|\{k : j \in B_k\}|}.$$

In other words, each agent's payoff in the Shapley value is, naturally enough, the total "novelty" of the items the agent contributes, where the novelty of an item is inversely proportional to the number of agents who know about it.

Collaborative Filtering

To introduce a somewhat more sophisticated model of preferences and the ways in which participants benefit from the announcement of preferences, suppose now that preferences are given in the form of an $n \times m$ matrix B, where each entry B_{ij} can be 1 (agent i likes item j) or -1 (agent i does not like item j) or 0 (agent i does not know item j). This setting is motivated by formalizing classic ideas from collaborative filtering (see for instance [9]).

In such work (see also [3]), there is an underlying set of preferences that users have for items; some of these preferences are patently known, while others are latent. The process of collaborative filtering aims at unearthing the latent preferences. To model this, we assume that the matrix B is the result of replacing certain entries of another $-1, 1$ matrix B' by zeros: B' captures the intrinsic preferences of the individuals, be they known to them or not. Let b_i denote the ith row of B. If $B_{ij} = 0$, agent i can get advice about item j from the agent $a(i) = \arg\max_{i' \neq i} b_i \cdot b_{i'}$, the agent i' who maximizes the "agreement" captured by the inner product $b_i \cdot b_{i'}$ (in case of ties, assume that $a(i)$ is the smaller-indexed agent). The payoff from this advice is 1 if i ends up agreeing with $a(i)$'s opinion on item j (that is, if $B'_{ij} = B_{a(i)j}$, and -1 if they disagree; if $B_{a(i)j} = 0$, the payoff is zero. Thus, the total payoff of the advice $a(i)$ can give i is $(b'_i - b_i) \cdot b_{a(i)}$.

To avoid unrealistic and degenerate worst-case situations with highly negative payoffs, with participants "luring" one another into undesirable items, we assume the following prior for B: Agents belong to a small set $\{t_1, \ldots, t_k\}$ of *well-separated types*. A type is a vector $t_i \in \{1, -1\}^m$, and any two types are separated by a large Hamming distance $h \geq c \cdot m$: $t_i \cdot t_j \leq m - 2h$ for $i \neq j$. An agent of type t (that is, a row b_i of the matrix B of type t) is generated by flipping, independently, each coordinate of t from 1 to -1 or vice-versa with a small probability p (thus determining whether agent i likes item j, whether he or she knows it or not); we denote the resulting vector by b'_i. Then each entry of b'_i is masked by a 0 with probability q, and the result is the row b_i of B.

In this framework we can prove the following:

Lemma 2.4 *Assume that all pairs of types are well separated in the Hamming metric, as above. Then there is a constant $a > 0$ such that with probability at least $1 - n^2 e^{-am}$, every agent consults another agent of his own type.*

According to the lemma, with high probability, we have the superposition of k independent games, one for each type. We are therefore justified in considering the case of only one type ($k = 1$). For any set S of participants, all of the same type, we denote by $v(S)$ the total payoff of this set, that is, the sum $\sum_{i \in S}(b'_i - b_i) \cdot b_{a(i)}$

We can prove the following:

Proposition 2.5 *Fix an integer d between 0 and m, and consider an agent whose vector $b'_i \in \{-1, 1\}^m$ is a distance d away from the type t ($|\{j : b'_{ij} \cdot t_j = -1\}| = d$. The expectation of b_i's Shapley value, given d, is a non-increasing function of d.*

In other words, in this formulation "typical" agents benefit their environment the most. This is in some contrast with our conclusions in the simpler recommendation system.

3 Discussion

We saw that in several cases the value of information offered by agents in a realistic if stylized scenario can be estimated in a clean and arguably objective fashion. We see this as modest, preliminary, exploratory work, whose main function is to point out, by dint of simplified and attractive examples, that issues of the fair value of information can be usefully studied within a theoretical framework. We hope that more research will further enrich and advance this effort, and eventually help address in a rational and practical manner the problem of privacy.

Acknowledgements. We acknowledge support from the David and Lucile Packard Foundation; the NSF under grants CCR-9701399 and IIS-0081334; and an ONR Young Investigator Award.

References

[1] Giacomo Calzolari and Alessandro Pavan, "Optimal design of privacy policies," manuscript, January 2001.

[2] Wassily Hoeffding, "Probability inequalities of Sums of Bounded Random Variables," *J. Amer. Stat. Assoc., 58* (1963), 13.

[3] R. Kumar, P. Raghavan, S. Rajagopalan, A. Tomkins, "Recommender systems: a probabilistic analysis," *Proc. 39th IEEE Symp. on Foundations of Computer Science*, 1998.

[4] Kenneth C. Laudon, "Markets and privacy," *Comm. ACM, 39* (1996), 104.

[5] Andreu Mas-Colell, Michael D. Winston, and Jerry R. Green, *Microeconomic Theory*, Oxford, 1995.

[6] Martin Osborne and Ariel Rubinstein, *A Course in Game Theory*, MIT Press, 1994.

[7] Pamela Samuelson, "Privacy as Intellectual Property?" in [8].

[8] *Stanford Legal Review, 52* (July 2000), special issue on privacy.

[9] Upendra Shardanand, Pattie Maes, Social Information Filtering: Algorithms for Automating "Word of Mouth," *Proc. ACM SIGCHI Conf. on Human Factors in Computing Systems*, 1995, 210-217.

[10] Hal Varian, "Economic Aspects of Personal Privacy," in *Privacy and Self-Regulation in the Information Age*, NTIA report. Available at
http://www.sims.berkeley.edu/ hal/people/hal/papers.html

[11] See for example:
Amazon.com, "Purchase circles," at
http://www.amazon.com/exec/obidos/tg/browse/-/468604/105-3945687-9154318.
America On-Line, "AOL Privacy Policy," at
http://legal.web.aol.com/policy/aolpol/privpol.html.
Microsoft, "Information about privacy on microsoft.com,"
http://www.microsoft.com/info/privacy.htm.