

The Value of Privacy

Optimal Strategies for Privacy Minded Agents

Sieuwert van Otterloo
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF
sieuwert@bluering.nl

ABSTRACT

Agents often want to protect private information, while at the same acting upon the information. These two desires are in conflict, and this conflict can be modeled in strategic games where the utility not only depends on the expected value of the possible outcomes, but also on the information properties of the strategy an agent uses. In this paper we define two such games using the information theory concepts of entropy and relative entropy. For both games we compute optimal response strategies and establish the existence of Nash equilibria.

Categories and Subject Descriptors

I.2.11 [ARTIFICIAL INTELLIGENCE]: Distributed Artificial Intelligence—*Multiagent systems*; K.4.4 [COMPUTERS AND SOCIETY]: Electronic Commerce—*Security*

General Terms

Theory

Keywords

game theory, privacy, utility, information theory, security, entropy

1. INTRODUCTION

Information is valuable, and thus agents do not always want to give it away. Both organisations and individuals often want to keep certain information private. At the same time they might want to act upon it. Does this reveal the information? In this paper we study how agents should act if they want to maximize their utility, while at the same time not giving away too much information. We do this by defining games in which the utility for each agent does not only depend on the payoff of the chosen action, but also on the information properties of the used strategy. These games can be applied to the following situations.

- Supermarkets and e-commerce shops register which customer buys what. Customers know this and even assist in this process by using so-called ‘bonus cards’ (Albert Heijn) or ‘club cards’ (Tesco). Nevertheless many customers are worried about their privacy. They would prefer it if the shop knew less about them. Customers can do something to minimize the knowledge of the shop. First of all they can make their shopping less regular (i.e. randomly buy items so that the shop is not sure which products the customer actually uses). Secondly they can sign up for more than one card(account) or swap cards between each other. On the Internet, deleting cookies with random intervals and using a different IP number can have the same effect.
- In a second prize auction it is optimal to bid exactly as much as you think the item is worth [11]. However, you might have spent a lot of time to estimate the value of the item, so you do not want to reveal your estimate. Since your bid has to be public, it seems that you might do better by bidding slightly random. By modeling this as a minimal information game, one can compute how one should randomise. A similar argument applies when you send out an artificial agent to do your shopping. If the agent is sent over an insecure network, everyone can inspect the source code and thus the bidding strategy of the agent. You might not want to send an agent that is exactly optimal for your preferences, in order to hide your preferences.
- Many public places are now guarded by closed circuit television systems. If you come to one such place regularly, the camera attendants learn a lot about your habits and thus about you. You feel this as a breach of your personal privacy, and decide to hide your habits by changing your behaviour often, for instance by going to different shops in a different order every time. This situation can also be modeled as a minimal information game. Again one can translate this example to the domain of artificial agents and the Internet.
- Consider now the case of a criminal who wants to steal from a shop guarded by a closed circuit television system. He wants to look like a regular shopper, but has different goals. He thus wants to behave so that he can steal the most, while at the same time appear to be a normal shopper. This can be modeled as a most normal game.

In this paper we define the two types of games mentioned, the minimal information game and the most normal game. As the similar setting of the last two examples suggest, these two games are related. From these examples it should also be clear that we assume that the strategies that agents use are publicly known. This assumption makes our results stronger (if you have privacy while your strategy is public, you will have even more privacy when you can keep your strategy secret).

Privacy has received a lot of attention from economists or in a legal setting. Some key sources have been collected on a website [1]. This paper differs from these economic papers for two reasons. First of all we only deal with personal information privacy, whereas the word ‘privacy’ also has other meanings. The second difference is that these papers try to explain the need for personal privacy in terms of economic utility. Odlyzko for instance relates privacy and price discrimination [14]. This paper is written under the assumption that privacy is a fundamental value, that is not instrumental to any gain. Privacy itself is a good cause that can be enjoyed directly.

The games defined in this paper use a soft approach towards information. They deal with probabilities explicitly, and can make subtle distinctions between possible, likely and almost certain events. This soft approach can be contrasted to the hard approach of logic and model checking. When taking a hard approach in protocol analysis, one is only interested in what is possible and what not, with a complete disregard for the relative likelihoods of different outcomes. Both the soft and the hard approach have been used for multi agent systems. The use of epistemic logic to understand the game of Cluedo [17] is an example of the hard approach, as well as other logical approaches to reasoning about knowledge and knowledge change [8, 4, 16, 3, 19]. Recent work on privacy preserving auctions [5] and work on the Dining cryptographer problem [6] or the Russian Cards problem [17, 18] can also be classified as ‘hard’. At the same time there is some work on reasoning about uncertainty [9, 10] that combines logic and a soft approach to information. The soft approach is more precise than the hard approach and in certain circumstances this is an advantage. The hard approach can tell us that agents do best by randomising their strategy, but does not indicate the exact probabilities of an optimal strategy. On the other hand the higher level of abstraction of the hard approach makes it easier to interpret the results.

The layout of this paper is as follows. Section 2 describes a detailed example problem. The next section, section 3, introduces basic information theory notions such as entropy. Then we introduce strategic games in section 4. In section 5 we define minimal information games, and calculate the best strategies in these games. In section 6 we do the same for most normal games. Then we present our conclusions in section 7. Finally section A of the appendix contains a technical result that is not essential to the main argument of this paper.

2. EXAMPLE

The following problem serves as an example. Alice (agent 1) needs to buy one box of breakfast cereals every week. Every week she is faced with the following choice: whether

to buy Allgrain(A), Barley(B) or Cornflakes(C). Alice is not indifferent to which brand she eats. In fact she likes A better than B and B better than C , as is indicated by the following matrix of utilities.

action	A	B	C
utility	3.0	2.0	1.0

If Alice is solely interested in maximising her expected utility, she should buy A every day. However Alice knows that the shop is watching her shopping behaviour closely, and she is concerned about her privacy. She decides that the decision that she makes should be private, and she can achieve this by flipping a coin and letting her decision depend on this coin flip. This way the shop cannot predict her decision.

Alice first attempts to use the following random strategy.

action	A	B	C
probability	0.98	0.01	0.01

If Alice uses this strategy, then the shop does not know anything about her decision: all three actions may occur with positive probability. At the same time her expected payoff is still very high, because the suboptimal actions occur with a very low probability. Problem solved, so it seems. But this is not the whole story. Even though the shop does not gain any knowledge, it does gain information from this strategy. If the shop learns, from repeated observation, that Alice uses this strategy, then it is quite certain that she will buy A . Therefore the shop has gained quite a lot of information. Therefore the indicated strategy is not the right strategy if one analyses the situation using information theory.

One can argue that if Alice is concerned about her privacy, then that fact should be represented in her utility function. This is not possible, because the utility function can only express properties of single actions, whereas privacy is a property of the whole strategy. One could also decide to include an extra player that tries to guess Alice’s actions. It is however not clear how one should estimate all the variables that one needs for this larger game. These considerations have convinced us that it is easier to treat privacy as an independent aspect of an agent’s utility.

3. INFORMATION THEORY

Information theory is the field of science that deals with the measurement of information [7]. It has applications in signal processing, communication networks, cryptography and error correction codes. In this paper we use information theory, and its central notion entropy, to estimate the amount of information in strategies. Strategies will be modeled as stochastic variables ranging over a finite set of actions, so we define entropy over stochastic variables. The entropy of a stochastic variable is the amount of randomness in, the disorder of, or uncertainty about the value that the variable will take. The concept of entropy has been introduced by Shannon [15]. We define the following function $f(x, y)$. Let

lg be the base 2 logarithm.

$$f(x, y) = \begin{cases} 0 & \text{if } x = 0 \text{ and } y = 0 \\ \infty & \text{if } x > 0 \text{ and } y = 0 \\ -x \lg y & \text{if } x \geq 0 \text{ and } y > 0 \end{cases}$$

For a random variable \mathbf{X} we define the entropy $E(X)$, which is measured in bits, in the following way.

$$E(\mathbf{X}) = \sum_k f(p(X = k), p(X = k))$$

A random variable X with values in the domain $\{1, 2, \dots, m\}$ can be specified by giving a vector of length m with the probabilities of each value: $(p(X = 1), p(X = 2), \dots, p(X = m))$. For a mixed strategy, the numbers $\{1, 2, \dots, m\}$ represent the available actions. A requirement for probability measures on stochastic variables is that the probabilities should add up to 1. We can thus only use vectors x that indeed add up to 1. Define the sets \mathbf{P}^m and \mathbf{Q}^m .

$$\mathbf{P}^m = \{x \in [0, 1]^m \mid \sum_i x_i = 1\}$$

$$\mathbf{Q}^m = \{x \in (0, 1]^m \mid \sum_i x_i = 1\}$$

The set \mathbf{P}^m contains all vectors of length m that add up to 1, and \mathbf{Q}^m contains all vectors that add up to 1 and do not take the value 0. The set \mathbf{Q}^m is important in some of the proofs, but often we work with the more general set \mathbf{P}^m . We can apply the notion of entropy to probability vectors $x \in \mathbf{P}^m$.

$$E(x) = \sum_k f(x_k, x_k)$$

In the context of strategies, a strategy with a higher entropy leaves observers with more uncertainty, and thus gives the agent that uses that strategy more privacy. Below we give five examples of entropy. The example strategy vectors can all be seen as strategies over three basic actions. A strategy (a, b, c) contains the probability a of selection the first action, b for the second action and c for the third.

$$E((1/3, 1/3, 1/3)) = 1.585 \text{ bits}$$

$$E((0.5, 0.25, 0.25)) = 1.5 \text{ bits}$$

$$E((0.5, 0.5, 0)) = 1 \text{ bit}$$

$$E((0.98, 0.01, 0.01)) = 0.161 \text{ bits}$$

$$E((1.0, 0, 0)) = 0 \text{ bits}$$

Pure strategies, in which only one action gets a positive probability, have an entropy of zero bits. The entropy function is bounded. It cannot be negative, and a vector x of length m can have at most an entropy of $\lg m$. It has this entropy if all the entries x_i are equal to $1/m$, thus if the vector represents a stochastic variable with a uniform distribution.

The second idea that we use from information theory is *relative entropy* [7]. The function $r(x, y)$ can be used to compare two probability vectors $x, y \in \mathbf{P}^n$. The underlying idea is that $r(x, y)$ measures how much difference one would notice if probability vector x is used instead of y for selecting actions. In order to compute this difference, we add up the differences for each action k . Using Bayes' law one can derive that the relative likelihood of strategy x instead of strategy y when observing that action k is chosen is x_k/y_k . This

observation is the motive behind the following definition.

$$r(x, y) = \sum_k f(x_k, y_k/x_k)$$

The function r almost behaves as a norm or distance function. It is never negative and only zero if $x = y$. It is infinite if for some k it is the case that $x_k > 0$ and $y_k = 0$. The only difference between this function and a distance or norm function is that r is not symmetric. In many cases $r(x, y) \neq r(y, x)$.

$$r((0.5, 0.5), (0.75, 0.25)) = 0.2075 \text{ bits}$$

$$r((0.75, 0.25), (0.5, 0.5)) = 0.1887 \text{ bits}$$

$$r((0.9, 0.1), (0.75, 0.25)) = 0.1045 \text{ bits}$$

$$r((0.75, 0.25), (0.9, 0.1)) = 0.1332 \text{ bits}$$

If x has a higher entropy than x' , then on average for a random vector y it is the case that $r(y, x) < r(y, x')$. It is harder to notice a difference between y and a high entropy vector x than to notice a difference between y and a low entropy vector x' .

4. STRATEGIC GAMES

Games can be presented in different forms. A very natural but detailed form is as an extensive game. In this form there are a number of decision points in each play of the game, and the outcome is determined by all these decisions. This model is too detailed for our purposes. Therefore we study games in strategic or normal form. In this form, each agent has a number of strategies available at the beginning of the game, and each agent independently picks a strategy. We can calculate the utility of each agent in the game directly, without going into details which actions have been played. The general definition for an n -agent normal form game is the following. We let Σ be the set of all agents, and assume that $\Sigma = \{1, 2, \dots, n\}$ for some $n > 0$.

DEFINITION 1. A game G is a tuple $(\Sigma, \{S\}_\Sigma, p)$ where for each $X \in \Sigma$ the set S_X is a set of strategies for agent X , and $p : (S_1 \times \dots \times S_n) \rightarrow \mathbb{R}^\Sigma$ is a utility function.

Each agent tries to maximize its utility. The sets of strategies do not have to be finite. A vector $\vec{s} = (s_1, \dots, s_n)$ is a strategy vector for game G if $G = (\Sigma, \{S\}_\Sigma, p)$ and for all i we have $s_i \in S_i$. If $t \in S_j$ then we define $[s_{-j}, t] = [s_1, \dots, s_{j-1}, t, s_{j+1}, \dots, s_n]$ as the strategy vector where s_j is replaced by t . For example $[(a, b, c)_{-2}, d] = (a, d, c)$.¹

We assume that every agent X always has a finite number of basic actions m_X to choose from, and that the total utility of a strategy somehow depends on the payoff of each action. The payoff of each action is typically given in the form of a matrix A . Since the number of agents may be larger than two, we extend the idea of a matrix to the following

¹It is a game-theoretic convention that s_{-j} denotes the vector s with the j th element removed. Thus $(a, b, c)_{-2} = (a, c)$. The construct $[s, x]$ is used to denote the vector s with x inserted in an appropriate place: $[(a, c), d] = (a, d, c)$. Determining what the appropriate place is can be difficult, therefore I only define the combination of these two constructs.

definition of a multi-matrix. A $m_1 \times m_2 \dots \times m_n$ multi-matrix is a function A such that for each vector $i_1 i_2 \dots i_n$ with $i_j \in \{1, \dots, m_j\}$ and $X \in \{1, \dots, n\}$, the function A returns a real number $A^X(i_1 i_2 \dots i_n) \in \mathbb{R}$. The expression $A(i_1 i_2 \dots i_n)$ denotes a real vector $v \in \mathbb{R}^\Sigma$ such that $v_1 = A^1(i_1 i_2 \dots i_n)$, $v_2 = A^2(i_1 i_2 \dots i_n)$ etcetera.

For a given multi-matrix A one can define different games. The simplest type of game is the pure strategy game. In this game the strategy of each agent X consists of a single action a_X and the payoff is then $A(a_1 \dots a_n)$. This definition does not allow agents to play randomly. For our purposes this definition is thus too restrictive. In a mixed strategy game, the strategy of an agent is a probability distribution over the available actions. The payoff is the expected (weighted average) value of A . This type of game is defined in the next definition.

The shorthand $A_i^X(\vec{s})$ denotes the expected payoff of action i for agent X when the other agents use strategies from \vec{s} . It can be defined in the following way. Define the set $V_i^X = \{\vec{v} | v_Y \in S_Y, v_X = i\}$. Thus this set contains the pure strategy profiles in which agent X selects action i .

$$A_i^X(\vec{s}) = \sum_{\vec{v} \in V_i^X} (s_{v_1} \dots s_{v_{X-1}} s_{v_{X+1}} \dots s_{v_n}) A^X(\vec{v})$$

DEFINITION 2. Let A be a $m_1 \times m_2 \dots \times m_n$ multi-matrix. The mixed strategy game $M_m(A)$ of A is a tuple $(\Sigma, \{S_X\}, \mathfrak{U})$ where $\Sigma = \{1, 2, \dots, n\}$, the strategy sets are $S_X = \mathbf{P}^{m_X}$ and $\mathfrak{U}^X(\vec{s}) = \sum_i s_i^X A_i^X(\vec{s})$

The fact that agents can play mixed strategies is explicitly defined in this definition of a mixed strategy game. We assume that all agents are equipped with random number generators (coins, dice or whatever) so that they can randomize their behavior exactly as specified in their strategy.

The central question in game theory has always been the question about the ‘solution’ of a certain game. Intuitively the solution is the strategy vector containing the best possible strategy for each agent. However not every game has a unique solution in this sense. Therefore game theorists work with different solution concepts. One of the best known is the Nash Equilibrium. Every mixed strategy game has a Nash equilibrium, but very often it is not unique.

For the next definition we need the function argmax that returns all inputs that maximize a given function. $\operatorname{argmax}_x f(x) = \{x | \neg \exists y : f(x) < f(y)\}$ We use the function argmax to define what a ‘good’ strategy is: A good strategy is a strategy that returns a maximal utility. The function b^X returns the best response strategies for agent X for a given game and strategy vector.

DEFINITION 3. Let $(\Sigma, \{S\}_\Sigma, p)$ be a game and $\vec{s} \in (\prod_X S_X)$ a strategy profile. The best response $b(\vec{s}) = b^1(\vec{s}) \times \dots \times b^n(\vec{s})$ is defined by

$$b^X(\vec{s}) = \operatorname{argmax}_t \mathfrak{U}^X([s_{-X}, t])$$

The set $b(\vec{s})$ thus contains the strategy vectors t such that t_X is optimal if all opponents Y use the strategy s_Y . In a decision theory problem we could assume that the strategy of the opponents is fixed. The set $b^X(\vec{s})$ is the set of best decisions for agent X . In game theory things are not that simple, because the other agents might want to change their strategy once they learn that X uses a strategy in $\mathfrak{U}^X(\vec{s})$. However this interaction is nicely captured by the definition of a Nash equilibrium.

DEFINITION 4. Let $(\Sigma, \{S\}_\Sigma, \mathfrak{U})$ be a game and $\vec{s} \in (\prod_X S_X)$ a strategy profile. The vector \vec{s} is a Nash Equilibrium iff $\vec{s} \in b(\vec{s})$

Every mixed strategy game has at least one Nash equilibrium [13]. There has been some discussion in the literature whether the notion of a Nash equilibrium needs to be refined. Several refinements have been proposed, but none of them have the appealing simplicity of the Nash equilibrium.

5. MINIMAL INFORMATION GAMES

The next definition of a minimal information game aims to capture the following situation. Agents choose a mixed strategy with two goals in mind. First of all they want a high payoff. Secondly they want privacy. They feel that they have more privacy if others are more uncertain about the action they will choose, and thus they prefer strategies with a high entropy. These games thus model the situation where agents have a fundamental desire for privacy.

We have to specify how the agent would like to trade privacy against payoff. This is governed by a parameter $\alpha > 0$ that indicates the value of privacy. It expresses how much expected payoff the agent is willing to trade against a bit of privacy. The higher α , the more the agent values privacy.

DEFINITION 5. Let A be a $m_1 \times m_2 \dots \times m_n$ multi-matrix and $\alpha > 0$. The minimal information game $M_i^\alpha(A)$ is a tuple $(\Sigma, \{S_X\}, \mathfrak{U})$ where $\Sigma = \{1, 2, \dots, n\}$, the strategy sets are $S_X = \mathbf{P}^{m_X}$ and $\mathfrak{U}^X(\vec{s}) = \sum_i s_i^X A_i^X(\vec{s}) + \alpha E(s^X)$

The parameter α regulates how much all the agents value the fact that there is uncertainty over their next action. If we would allow $\alpha = 0$, then the game becomes a mixed strategy game: $M_i^0(A) = M_m(A)$. As α approaches infinity, the actual payoff becomes less and less important. It would have been possible to choose α differently for each agent, but this would have made the definition less clear. One can always scale the utilities in such a way that one value for α works for all agents.

As an example, we consider the shopping game from the introduction. This game has only one agent, that has three options A, B, C with respective payoffs 3, 2, 1. The optimal strategies for the minimal information game with different values of α is given in the next table. It also lists the utility of s that the agent would get in the mixed strategy game $M_m(A)$ for the given strategy s and the utility that the agent would get in the minimal information game $M_i^\alpha(A)$.

α	p_1	p_2	p_3	$M_m(A)$	$M_i^\alpha(A)$
0.1	0.999	$4 \cdot 10^{-5}$	$2 \cdot 10^{-9}$	3.0	3.0
0.5	0.876	0.117	0.015	2.852	3.168
1.0	0.665	0.244	0.090	2.575	3.775

The best payoff that the agent can get is 3.0 by only choosing the first action. However this would result in no privacy, because if the agent would use this strategy, then any observer knows beforehand what the agent will do every day. For a low value of α the utility of s in $M_i^\alpha(A)$ is very close to this optimal value of 3. For higher values, the average payoff without entropy becomes lower. We could call this the *cost* of privacy. From the table we can see that if the agent values privacy at one unit per bit (α is expressed in units per bit) then the agent does best by paying 0.425 in order to obtain 0.775 bits of privacy.

The question is of course how we can calculate the strategies that maximize the utility in minimal information games. For the linear functions of the mixed strategy games this is trivial, but for more complicated functions this can be difficult.

THEOREM 1. *Let $M_i^\alpha(A)$ be a minimal information game and \vec{s} a strategy profile. The set $b^X(\vec{s})$ is a singleton $\{b\}$ such that*

$$b_i = \frac{2^{\alpha^{-1}A_i^X(\vec{s})}}{\sum_k 2^{\alpha^{-1}A_k^X(\vec{s})}}$$

PROOF. Let $M_i^\alpha(A) = (\Sigma, \{S\}_\Sigma, \mathfrak{U})$ be a minimal information game. We have to prove that the set $b^X(\vec{s})$ contains one element, and that that element is described by the given formula. We first show that all points in $b(\vec{s})$ are interior points. Then we derive an equation that any best response must satisfy, and show that this equation has a unique solution, namely the one given in the theorem.

Take any vector $\vec{x} \in \mathbf{P}^n \setminus \mathbf{Q}^n$. We are going to show that there is a better vector \vec{y} , and thus \vec{x} is not a best response. There is some i such that $x_i = 0$ and some j such that $x_j \neq 0$. We will show that there is some ϵ such that $\vec{y} = [x_{-i}, \epsilon]_{-j}, x_j - \epsilon]$ is a better vector: $\mathfrak{U}^X(\vec{y}) > \mathfrak{U}^X(\vec{x})$. To show this, note that the utility function \mathfrak{U}^X is continuous and differentiable. Note further that $\frac{\delta}{\delta x_i} \mathfrak{U}^X(\vec{x}) = +\infty$ and $\frac{\delta}{\delta x_j} \mathfrak{U}^X(\vec{x}) < +\infty$. Therefore, for sufficiently small ϵ , the gain from raising x_i outweighs the potential loss from lowering x_j . Therefore for sufficiently small ϵ we have that $\mathfrak{U}^X(\vec{y}) > \mathfrak{U}^X(\vec{x})$ and thus $\vec{x} \notin b(\vec{s})$.

Now suppose that $b \in b^X(\vec{s})$. We know that $b \in \mathbf{Q}^n$. Take $i, j \in \{1, 2, \dots, m\}$ as two different indices. Since b is optimal, it should not be possible to increase \mathfrak{U}^X by increasing b_i while decreasing b_j , and therefore for any optimal point it must be the case that $\frac{\delta}{\delta b_i} \mathfrak{U}^X([\vec{s}_{-X}, b]) = \frac{\delta}{\delta b_j} \mathfrak{U}^X([\vec{s}_{-X}, b])$. We can use this as a starting point for the following link of equations. First we compute the derivative $\frac{\delta}{\delta b_i} \mathfrak{U}^X([\vec{s}_{-X}, b])$.

$$\begin{aligned} \frac{\delta}{\delta b_i} \mathfrak{U}^X([\vec{s}_{-X}, b]) &= \\ \frac{\delta}{\delta b_i} \left(\sum_j b_j A_j^X([\vec{s}_{-X}, b]) + \alpha E(\vec{b}) \right) &= \\ A_i^X(\vec{s}) + \alpha \frac{\delta}{\delta b_i} (E(\vec{b})) &= \\ A_i^X(\vec{s}) + \alpha (-\lg b_i - \lg e) &= \\ A_i^X(\vec{s}) - \alpha \lg b_i - \alpha \lg e & \end{aligned}$$

Using this derivative one can reduce the equation given above in the following way.

$$\begin{aligned} \frac{\delta}{\delta b_i} \mathfrak{U}^X([\vec{s}_{-X}, b]) &= \frac{\delta}{\delta b_j} \mathfrak{U}^X([\vec{s}_{-X}, b]) & \Leftrightarrow \\ A_i^X(\vec{s}) - \alpha \lg b_i &= A_j^X(\vec{s}) - \alpha \lg b_j & \Leftrightarrow \\ A_i^X(\vec{s}) - A_j^X(\vec{s}) &= \alpha \lg b_i - \alpha \lg b_j & \Leftrightarrow \\ \frac{2^{A_i^X(\vec{s})}}{2^{A_j^X(\vec{s})}} &= \frac{b_i^\alpha}{b_j^\alpha} \end{aligned}$$

Since $b \in \mathbf{P}^n$ it must be the case that b sums up to $\sum_i b_i = 1$. For any $b \in b(\vec{s})$ one can find some positive constant c such that $b_i = c \cdot 2^{\alpha^{-1}A_i^X(\vec{s})}$. It now follows from the above equation that for any b_j it is the case that $b_j = c 2^{\alpha^{-1}A_j^X(\vec{s})}$. We can now calculate $\sum_k b_k = 1 = c \sum_k 2^{-\alpha A_k^X(\vec{s})}$ and thus we know that $\frac{1}{c} = \sum_k 2^{\alpha^{-1}A_k^X(\vec{s})}$. Thus we have proven that there is a unique point $b \in b^X(\vec{s})$ which satisfies the equation in theorem 1 \square

THEOREM 2. *Every minimal information game $M_i^\alpha(A)$ has a Nash equilibrium.*

PROOF. Let f be the function from $S_1 \times \dots \times S_n$ to $S_1 \times \dots \times S_n$ that returns the strategy vector with the best responses for each agent. Thus f is the function that for each x returns the unique point $f(x)$ such that $f(x) \in b(x)$. The previous theorem shows that this is a continuous function. The set $S_1 \times \dots \times S_n$ is topological isomorphic to some closed sphere \mathbb{B}^m . We can now use Brouwer's fixed point theorem, which tells us that every continuous function $f : \mathbb{B}^m \rightarrow \mathbb{B}^m$ must have a point x with $f(x) = x$ [2]. We thus obtain a strategy vector x with $f(x) = x$, and thus a point x such that $x \in b(x)$. This point is a Nash equilibrium. \square

6. MOST NORMAL STRATEGIES

So far we have discussed the situation in which the agents try to protect their privacy against an opponent interested in their next action. In this section we look another situation, in which agents try to hide their preferences. The assumption is here that an average strategy for 'normal' users is given. One agent however has different preferences from the normal users, but does not want to be identified as not normal. Therefore the agent is searching for a strategy that appears as normal as possible and maximizes its payoff at the same time.

We approach the problem in exactly the same way as we have approached the first problem. We define most normal

games $G_n^\alpha(A)$ that depend on a parameter α expressing how important normal behaviour for the agent is.

DEFINITION 6. Let A be a $m_1 \times m_2 \dots \times m_n$ multi-matrix, let $\alpha > 0$, and let \vec{t} be a strategy vector for the game $M_m(A)$. The most normal game $M_n^\alpha(A, \vec{t})$ is a tuple $(\Sigma, \{S_X\}, \mathfrak{U})$ where $\Sigma = \{1, 2, \dots, n\}$, the strategy sets are $S_X = \mathbf{P}^{m^X}$ and $\mathfrak{U}^X(\vec{s}) = \sum_i s_i^X A_i^X(\vec{s}) - \alpha r(s^X, t^X)$

The parameter α again determines the trade-off between selecting actions with a high payoff and acting normal.

THEOREM 3. Let $M_n^\alpha(A, \vec{t})$ be a most normal game and \vec{s} a strategy profile for this game. The set $b^X(\vec{s})$ is a singleton $\{b\}$ such that

$$b_i = \frac{t_i^X 2^{\alpha^{-1} A_i^X(\vec{s})}}{\sum_k t_k^X 2^{\alpha^{-1} A_k^X(\vec{s})}}$$

PROOF. Let $M_n^\alpha(A, \vec{t})$ be a most normal game, \vec{s} a strategy profile and $X \in \Sigma$ an agent. Suppose that $b \in b^X(\vec{s})$ is the best response for agent X and let i be one of B 's actions. If $t_i = 0$ and $b_i \neq 0$, then the relative entropy becomes infinite, and the utility thus infinitely low. This cannot be optimal, thus if b maximizes the utility, then $t_i = 0$ implies $b_i = 0$. Thus in this case the optimal point is not an interior point. It follows that if $t_i = 1$, then for any optimal strategy b we must have $b_i = 1$.

Consider now the case where $t_i > 0$. We calculate the derivative of the relative entropy function.

$$\frac{\delta}{\delta b_i} r(b, t^X) = \frac{\delta}{\delta b_i} \sum_i -b_i (\lg t_i^X - \lg b_i) = \lg b_i + \lg e - \lg t_i^X$$

We see that if $b_i > 0$ approaches zero, then this derivative becomes negative infinity. If b_i is sufficiently small, then we would lower the utility $\mathfrak{U}^X([\vec{s}_{-X}, b])$ by decreasing b_i further. Therefore for any optimal value of b , it cannot be the case that $t_i > 0$ and $b_i = 0$.

Since we have shown that $t_i = 0$ implies $b_i = 0$, it remains for us to find the optimal vector in the space $S = \{b \in [0, 1]^m \mid \sum_i b_i = 1 \wedge (t_i = 0 \rightarrow b_i = 0)\}$. The previous argument has shown that b is an interior point of this set S . Such points can only be optimal if $\frac{\delta}{\delta b_i} \mathfrak{U}^X([\vec{s}_{-X}, b]) = \frac{\delta}{\delta b_j} \mathfrak{U}^X([\vec{s}_{-X}, b])$ for any pair i, j with $t_i, t_j > 0$. The next computation will show that there is a unique point satisfying this condition. Since any continuous function on a closed domain must have a maximum, this point b will maximize agent X 's utility in the normal form game.

First we calculate the derivative.

$$\begin{aligned} \frac{\delta}{\delta b_i} \mathfrak{U}^X([\vec{s}_{-X}, b]) &= \\ A_i^X(\vec{s}) - \alpha \frac{\delta}{\delta b_i} r(b, t^X) &= \\ A_i^X(\vec{s}) - \alpha (\lg b_i + \lg e - \lg t_i^X) &= \\ A_i^X(\vec{s}) - \alpha \lg b_i - \alpha \lg e + \alpha \lg t_i^X & \end{aligned}$$

Now find the points b where the derivatives $\frac{\delta}{\delta b_i} \mathfrak{U}^X$ and $\frac{\delta}{\delta b_j} \mathfrak{U}^X$ are equal.

$$\begin{aligned} \frac{\delta}{\delta b_i} \mathfrak{U}^X([\vec{s}_{-X}, b]) &= \frac{\delta}{\delta b_j} \mathfrak{U}^X([\vec{s}_{-X}, b]) && \Leftrightarrow \\ A_i^X(\vec{s}) - \alpha \lg b_i + \alpha \lg t_i^X &= A_j^X(\vec{s}) - \alpha \lg b_j + \alpha \lg t_j^X && \Leftrightarrow \\ \alpha \lg(b_i/b_j) &= A_i^X(\vec{s}) - A_j^X(\vec{s}) + \alpha \lg(t_i^X/t_j^X) && \Leftrightarrow \\ \frac{b_i}{b_j} &= \frac{t_i^X 2^{\alpha^{-1} A_i^X(\vec{s})}}{t_j^X 2^{\alpha^{-1} A_j^X(\vec{s})}} \end{aligned}$$

Again we can choose c such that $b_i = ct_i^X 2^{\alpha^{-1} A_i^X(\vec{s})}$ and show that $1/c = \sum_k t_k^X 2^{\alpha^{-1} A_k^X(\vec{s})}$. This leads to the next formula.

$$b_i = \frac{t_i^X 2^{\alpha^{-1} A_i^X(\vec{s})}}{\sum_k t_k^X 2^{\alpha^{-1} A_k^X(\vec{s})}}$$

This formula gives us $b_i = 1$ if $t_i = 1$, and $b_i = 0$ if $t_i = 0$. Therefore this formula gives us the optimal strategy for any normal form game. \square

Discussion

One consequence of the theorem is the following observation. If a certain action i is not considered by normal agents ($t_i^X = 0$) then the non-normal agent should not consider action i either ($b_i = 0$). If one had used a hard, logical approach one could have reached the same conclusion. In the most extreme case one can consider the case where normal agents use a pure strategy. In that case the non-normal agent has to use the same pure strategy. If the non-normal agent values all actions equally, he also does best by copying the normal strategy. In all other cases the best strategy for the non-normal agent is different. Apparently the agent does best by always taking some risk and getting a higher utility.

7. CONCLUSION

We have defined two new kinds of games. First of all minimal information games, in which agents want to maximize the uncertainty that observers have over their next move. Secondly most normal games, in which agents want to behave as similar as possible to an existing 'normal' agent, while maximizing their payoff. In order to do so we borrowed the concepts entropy and relative entropy from information theory. In two theorems we have shown what the optimal best responses are in these games. These turn out to be unique in each situation, and to depend continuously on the payoff matrix and the opponent strategies. From this continuity one can derive that Nash equilibria exist in these games.

Minimal information games can be used to analyse situations with privacy-minded agents. If agents attach some value of privacy, the best strategy always gives them some privacy.

In most normal games, the situation is slightly more complicated. How well the non-normal agent X can do, depends very much on the strategy that normal agents use. If the normal agents use a pure strategy, then X has no choice but to adopt the same strategy. The situation however becomes a lot better if the normal agents are privacy-minded. In that

case they choose a high-entropy strategy, and this leaves the wanting-to-be-normal agent a lot of room to pursue its own agenda.

One can extend the work in these games in several ways. First of all it would be interesting to look at experimental data, to see whether most-normal or minimal-information strategies are used in the real world. Secondly one could implement these strategies in order to obtain privacy. The question is then whether the soft approach to privacy is what users want. A small simulation is available at: www.bluering.nl/siewert/programs/privacysim/simprivacy.html

On a theoretical side, it seems that these games give approximations to the Nash equilibrium with very nice technical properties. Two of these properties are continuity of the best response function and the fact that best responses are always interior. In the appendix of this paper we already use minimal information games to define a refinement of the Nash equilibrium, as an example how these properties are technically useful.

8. ACKNOWLEDGMENTS

I would like to thank my colleagues in Liverpool and Amsterdam and the anonymous referees for their comments and suggestions. Specifically I would like to mention Steve Phelps, Karl Tuyls and Jelle Zuidema for reading (earlier versions of) this paper.

9. REFERENCES

- [1] A. Acquisti. The economics of privacy, 2004. <http://www.heinz.cmu.edu/acquisti/economics-privacy.htm>.
- [2] M. Aigner and G. Ziegler. *Proofs from the Book*. Springer-Verlag, 1999.
- [3] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science*, pages 100–109, Florida, October 1997.
- [4] A. Baltag, L. Moss, and S. Solecki. The logic of public announcements, common knowledge and private suspicions. Originally presented at TARK 98, accepted for publication in *Annals of Pure and Applied Logic*, 2002.
- [5] F. Brandt and T. Sandholm. (im)possibility of unconditionally privacy-preserving auctions. In *Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 810–817, 2004.
- [6] D. Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of Cryptology*, 1:65–75, 1988.
- [7] T. Cover and J. Thomas. *Elements of Information Theory*. John Wiley & Sons, 1991.
- [8] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about knowledge*. The MIT Press: Cambridge, MA, 1995.

- [9] J. Halpern. *Reasoning about uncertainty*. The MIT Press: Cambridge, MA, 2003.
- [10] B. Kooi. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information*, 12:381–408, 2003.
- [11] V. Krishna. *Auction theory*. Academic Press, 2002.
- [12] R. Myerson. Refinements of the nash equilibrium concept. *International Journal of Game Theory*, 7:73–80, 1978.
- [13] J. Nash. Non-cooperative games. *Annals of mathematics*, 54:286–295, 1951.
- [14] A. Odlyzko. Privacy, economics, and price discrimination on the Internet. In *Fifth International Conference on Electronic Commerce(ICEC)*, pages 73–80, 2003.
- [15] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, July, October 1948. <http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf>.
- [16] W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, pages 1167–1174, Bologna, Italy, 2002.
- [17] H. P. van Ditmarsch. The russian cards problem. *Studia Logica*, 75(4):31–62, 2003.
- [18] S. van Otterloo, W. van der Hoek, and M. Wooldridge. Model checking a knowledge exchange scenario. *Applied Artificial Intelligence*, 18:937–952, 2004.
- [19] S. van Otterloo and M. Wooldridge. On the complexity of knowledge condition games. In *Proceedings of the second European workshop on Multi Agent Systems(EUMAS)*, Barcelona, 2004.
- [20] E. Weisstein. Compact set, 2004. from Mathworld - A Wolfram web resource.

APPENDIX

A. EQUILIBRIUM REFINEMENTS

In this appendix a refinement of the Nash equilibrium is defined in order to make some technical observations.

By introducing minimal information games we have introduced a game with a new kind of utility function. For small values of α the game $M_i^\alpha(A)$ is very similar to the mixed strategy game $M_m(A)$. One can, with some imagination, see a Nash equilibrium x of $M_i^\alpha(A)$ as a solution of $M_m(A)$. In that case, one has a new solution concept for mixed strategy games $M_m(A)$. Such a solution x of some game $M_i^\alpha(A)$ is not a Nash equilibrium of $M_m(A)$, but an approximation of it. How good this approximation is depends on the parameter α . We can define a Nash equilibrium by letting α approach zero. This way, we can define a ‘minimal information’ equilibrium.

DEFINITION 7. *The strategy profile x is a minimal information equilibrium of $M_m(A)$ iff there is a sequence $\alpha_0, \alpha_1, \dots$ of positive numbers such that $\lim_{i \rightarrow \infty} \alpha_i = 0$, a sequence x_0, x_1, \dots such that x_i is a Nash equilibrium of $M_i^{\alpha_i}(A)$ and $\lim_{i \rightarrow \infty} x_i = x$.*

THEOREM 4. *Every mixed strategy game $M_m(A)$ has a minimal information equilibrium.*

PROOF. Define the sequence β_0, β_1, \dots by $\beta_i = 1/i$. This sequence converges to zero. By theorem 2 each game $M_i^{\beta_i}(A)$ has some Nash equilibrium y_i . The strategy space $S_1 \times \dots \times S_n$ is a closed and bounded subset of \mathbb{R}^m for some m . Therefore, since any closed and bounded subset of \mathbb{R}^m is compact [20] we derive that every sequence in $S_1 \times \dots \times S_n$ has some converging subsequence. Let x_0, x_1, \dots be a converging

subsequence of y_0, y_1, \dots and let x be the limit of $\lim_{i \rightarrow \infty} x_i$. Let $\alpha_0, \alpha_1, \dots$ be the corresponding subsequence of β_0, β_1, \dots , so that x_i is a Nash equilibrium of $M_i^{\alpha_i}(A)$. When α converges to infinity, the utility function of $M_i^{\alpha_i}(A)$ converges uniformly to the utility function of $M_m(A)$. Since x_i is always maximizing each agents utility in $M_i^{\alpha_i}(A)$, it must be the case that x maximizes the utility of $M_m(A)$ for each agent. Therefore x is a Nash equilibrium of $M_m(A)$. \square

Every minimal information equilibrium is a proper equilibrium as defined by Myerson, and therefore it is also a trembling hand perfect equilibrium [12]. These refinements can thus be motivated (if one wants to) by an appeal to privacy minded agents. Perhaps there are other applications where one needs a response concept that selects interior solution points, for instance to avoid division by zero. In that case the minimal information best responses seem suitable.