Representations of Information Structures

Sanmay Das[1] and Emir Kamenica[2]

***Abstract*— In many problems of design of mechanisms and multi-agent systems, the system designer has control over the information environment. What is the optimal design given the goals of the system designer? We discuss several ways of representing information structures. Each representation simplifies a particular class of optimization problems over information structures; we discuss current and potential applications of these representations.**

I. Introduction

A central question for economics and for multi-agent systems in computer science and artificial intelligence, is how to design the rules of engagement for agent interaction [1]. The intersection between computer science and economics has been particularly fruitful in pursuing questions related to mechanism design. The feedback between disciplines has gone in both directions – the potential availability of highly rational artificially intelligent agents ("*machina economicus"* – a synthetic *homo economicus* – in the language of Parkes and Wellman [1]) that could participate in various types of markets has motivated the study of more complex mechanism design settings [2], while insights from mechanism design have guided AI researchers in multi-agent systems problems, for example team formation [3].

Mechanism design typically assumes the players' information is given, and searches for rules of the game that yield desired outcomes. In many circumstances, however, the information environment itself is under the control of the designer. This is of course particularly important as the information environment can have a substantial impact on the outcome. Hajaj & Sarne examine how e-commerce platforms can benefit from withholding information from customers about all the opportunities available to them [4]. Rochlin & Sarne consider teams of information-gathering agents and show that restricting their ability to share information can lead to improved group outcomes [5]. In the *deliberative auction* setting [6], where agents have the opportunity to acquire information about valuations before entering a bidding process, Brinkman *et al.* show that the dependence structures between agents' signals of the value of the item they are bidding on can lead to qualitatively different equilibrium outcomes of the auction [7]. Chhabra *et al* study the welfare effects of the cost and precision of information on product (or match) quality provided to agents engaging in costly search, and of competition between information providers with different signal qualities [8], [9]. Das & Li examine the relative effects of common and private information about quality in a two-sided matching model with costly interviews [10].

In this line of computer science research, authors usually compare specific information environments rather than consider how to systematically search the space of all information environments so as to identify the best one. The problem of designing the optimal informational environment – introduced by Kamenica and Gentzkow and often termed "Bayesian persuasion" [11] – has recently been studied in a variety of settings in the economics literature. These include Internet advertising [12], communication in organizations [13], bank regulation [14], [15],

*SD acknowledges support from NSF grants 1527037 and 1414452. EK thanks the University of Chicago Booth School of Business for financial support.

[1]Washington University in St. Louis, Dept. of Computer Science and Engineering

[2]University of Chicago Booth School of Business

medical testing [16], medical research [17], government control of the media [18], entertainment [19], and price discrimination [20].

In this paper, we discuss how various ways of representing information can aid the aforementioned problem of information design. We also discuss potential applications to computer science research. One important insight is that viewing information as a discrete number of bits often limits the power of information design.

Formally, given a state space Ω, an informational environment, or an *information structure*, is a map $\pi : \Omega \to \Delta(S)$ given some set of signal realizations S. Identifying the optimal information structure is a difficult computational problem if approached by brute force. Given a state space Ω, the set of all information structures is as large as $(\Delta(\Omega))^{|\Omega|}$.[1] The main goal of this paper is to give an overview of how these types of optimization problems can be simplified by a suitable way of representing all information structures (conflating those that are payoff equivalent). We discuss three such representations, each of which is useful for a particular class of problems.

The first approach represents information structures as distributions of posterior beliefs. Each signal realization induces a posterior belief and hence an information structure gives rise to a distribution of posterior beliefs. Moreover, any distribution of posterior beliefs that on average equals the prior can be induced by some information structure [11]. Hence, whenever the optimization value depends only on the posterior belief, the optimization problem can be recast as a choice of a distribution of posterior beliefs. This formulation of the problem has a nice geometric interpretation that often provides ample intuition about the solution, but the approach does not scale well to large state spaces.

The second approach represents information structures as convex functions. When the state space is large but the value of the optimization problem only depends on the *mean* of the posterior distribution, one may wish to directly optimize over the distribution of the mean. The trouble, however, is that not every distribution of posterior means that on average equals the prior mean can be induced by an information structure. Gentzkow and Kamenica characterize the set of feasible distributions of posterior means by focusing on the integrals of the cumulative distribution functions of the means [21]. This is always a convex function "sandwiched" between the completely uninformative and the fully informative information structure.

The third approach represents information structures as partitions of $\Omega \times [0, 1]$ [22]. This formalization is useful when one considers games where multiple agents provide information to a third party or when one considers the problem of allocating information across various agents.

[1] Kamenica & Gentzkow show that it is without loss of generality to set the cardinality of the signal realization space to be the same as the cardinality of the state space. Then, the set of all information structures has the same cardinality as $(\Delta(\Omega))^{|\Omega|}$ [11].

II. Optimization problems

There is a state space Ω with a typical state denoted ω. The prior on the state is some $\mu_0 \in \Delta(\Omega)$. An information structure is a map $\pi : \Omega \to \Delta(S)$ given some signal realization space S. Without loss of generality we can set $|S| \leq |\Omega|$.

Given some value function $u(\pi)$, we consider a problem of maximizing over all information structures. The next two sections consider two important special cases of such problems.

A. Value over posteriors

Given a prior μ_0 and information structure π, each possible signal realization $s \in S$ leads to a posterior belief μ_s via Bayes' rule. Thus, an information structure π induces a distribution of posterior beliefs denoted $\langle \pi \rangle \in \Delta(\Delta(\Omega))$.

In this subsection we consider optimization problems where the value function can be written as

$$u(\pi) = \mathbb{E}_{\langle \pi \rangle} v(\mu_s)$$

for some function $v(\cdot)$.

Many situations lead to such value functions. For example, it might be the case that the optimization problem is faced by some "Sender" who wishes to influence the action $a \in A$ of a "Receiver". In that case $v(\mu) = \mathbb{E}_\mu [w(a^*(\mu), \omega)]$ where $w(a, \omega)$ is Sender's objective function and $a^*(\mu)$ is Receiver's optimal action when her belief is μ.

Or, consider a situation where a single agent chooses what costly information to obtain prior to taking her action in order to maximize some $w(a, \omega)$. If the cost of acquiring information π can be written as $c(\pi) = \mathbb{E}_{\langle\pi\rangle} k(\mu)$, then $v(\mu)$ is simply $W(\mu) - k(\mu)$ where $W(\mu) \equiv \max_a \mathbb{E}_\mu [w(a, \omega)]$. Gentzkow and Kamenica discuss the cases when $c(\pi)$ can be decomposed this way [23]. These include those where the cost of information is proportional to the reduction in entropy [24].

Whenever the value function can be written as $u(\pi) = \mathbb{E}_{\langle\pi\rangle} v(\mu)$, there is a geometric way of approaching the problem of maximizing over π. Kamenica and Gentzkow show that for any $\tau \in \Delta(\Delta(\Omega))$ such that $\mathbb{E}_\tau [\mu] = \mu_0$ there exists a π s.t. $\tau = \langle\pi\rangle$ [11]. This representation of all information structures as the set of all distributions of posteriors that on average equal the mean implies that the value of the optimization must be $V(\mu_0)$ where V is the concavification of v. A concavification of a function f is the smallest concave function everywhere above f. This result is depicted in Figure 1.

We now turn to some applications of this concavification approach. Lazear examines the question of when police, with limited resources, should announce the roads they are planning to patrol versus keeping them secret in order to best deter speeding [25]. This problem is equivalent to testing in educational settings or when and what to audit to minimize tax fraud. Lazear's analysis, however, only compares revealing the relevant information versus not revealing it. The concavification approach shows, however, that *partial* information revealed through a stochastic signal always dominates full revelation; moreover, it dominates no revelation unless the police resources are sufficiently large. This observation highlights the fact that viewing information in terms of a deterministic number of bits entails a consequential loss of generality.

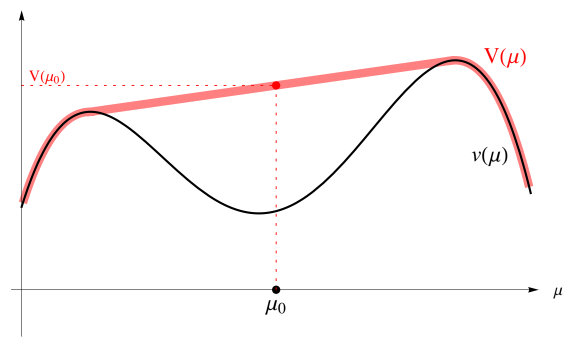

Fig. 1: The concavification of v gives the value of the optimization $V(\mu)$

In computer science, questions related to the optimal deployment of limited resources for problems in security, under the umbrella of "Stackelberg Security Games" have been the subject of recent study. Algorithms that solve for Bayesian Stackelberg equilibria have been a key to deployment in large-scale settings, including officers at airports [26], and air marshals on flights [27]. Rabinovich *et al* have recently put forward proposals based on the Bayesian persuasion approach to Stackelberg Security Games [28]. This is a potentially fertile area for further research on information structures.

Another possible application relates to the problem of *active learning* in machine learning. The idea of active learning is simply that, when building a learning model, instead of receiving the training set complete with labels, with examples in the form $(\mathbf{x}_i, y_i)$, one can query a (potentially noisy) oracle for the label (y_i) corresponding to $\mathbf{x}_i$. Most work on active learning has focused on budgeted models or sample complexity, but there has been some work on cost-sensitive active learning [29] and on so-called "proactive learning" [30] where there are multiple noisy oracles that may have different costs. Reasoning about information structures directly would allow one to approach the problem of *exemplar construction* – if one has the power to construct an $\mathbf{x}_i$ to best achieve some goal, instead of picking from a menu of available examples, what should one construct? Often there are costs associated with constructing training examples, and it costs more to build more accurate training examples – for example, in constructing manufacturing designs for testing, the more realistic the design, the more expensive it is to construct, but the better it models the real-world performance.

The concavification approach described here, however, is of limited use when the state space is large. We now turn to a representation that can be useful when the state space is large, as long as the value function can be written as a function of posterior means.

B. Value over posterior means

Now suppose that ω is a random variable. In this case, it will be more convenient to associate each distribution over Ω with its cumulative distribution function (cdf). Thus we denote the prior by F_0 and (given an information structure) the posterior induced by signal s by F_s. Moreover, we let m_0 denote the prior mean and m_s the mean of F_s. Since each information structure induces a distribution of posteriors, it also induces a distribution over posterior means. Let $G_\pi : \mathbb{R} \to [0,1]$ denote the (cdf of the) distribution of posterior means induced by signal π.

In this subsection we examine cases where the value function can be written as

$$u(\pi) = \mathbb{E}_{G_\pi} v(m_s)$$

for some function $v(\cdot)$.

First note that the concavification of v does *not* yield the solution to this maximization problem because not every distribution of posterior means that on average equals m_0 can be induced by an information structure. To classify the set of feasible distributions, Gentzkow and Kamenica [21], associate with each π the integral of G_π, i.e., $c_\pi(x) = \int_0^x G_\pi(t)\,dt$. If a function c_π is thus obtained from π we say that π *induces* c_π.

We illustrate this definition with some examples. Suppose that F_0 is uniform. Consider a totally uninformative signal $\underline{\pi}$. This signal induces a degenerate distribution of posterior means always equal to $m_0 = \frac{1}{2}$. Hence, $G_{\underline{\pi}}$ is a step function equal to 0 below $\frac{1}{2}$ and equal to 1 above $\frac{1}{2}$. The induced convex function $c_{\underline{\pi}}$ is thus flat on $\left[0, \frac{1}{2}\right]$ and then linearly increasing, with a slope of 1, from $\frac{1}{2}$ to 1. At the other extreme, consider a fully informative signal $\overline{\pi}$ that fully reveals the state. In that case, each posterior has a degenerate distribution with all the mass on the true state and thus $G_{\overline{\pi}} = F_0$. Since F_0 is uniform, $G_{\overline{\pi}}$ is linear, and thus $c_{\overline{\pi}}$ is quadratic: $c_{\overline{\pi}}(x) = \frac{1}{2}x^2$. Finally, consider a "partitional" signal $\mathscr{P}$ that gives a distinct signal realization depending on whether the state is in $\left[0, \frac{1}{2}\right]$, or $\left(\frac{1}{2}, 1\right]$. Then, $G_{\mathscr{P}}$ is a step function and $c_{\mathscr{P}}$ is piecewise-linear. Figure 2 depicts these CDFs and functions.

If we consider an arbitrary signal π, what can we say about c_π? Since G_π is a CDF and thus increasing, c_π as its integral must be convex. Moreover, since any signal π is a garbling of $\overline{\pi}$, we must have that $G_{\overline{\pi}}$ is a mean-preserving spread of G_π (Blackwell 1953); hence, $c_{\overline{\pi}} \geq c_\pi$ by Rothschild and Stiglitz (1970). Similarly, since $\underline{\pi}$ is a garbling of π, G_π is a mean-preserving spread of $G_{\overline{\pi}}$ and thus $c_\pi \geq c_{\underline{\pi}}$. Moreover, these characteristics are not only necessary; they are also sufficient for some c to be induced by an information structure:

Proposition 1: (Gentzkow and Kamenica [21]) Given a function $c : [0,1] \to \mathbb{R}$, there exists an information structure that induces it if and only if it is convex and $c_{\overline{\pi}}(x) \geq c(x) \geq c_{\underline{\pi}}(x)$ $\forall x \in [0,1]$.

Proposition 1 thus provides us with a simple characterization of the distributions of posterior means that can be induced by an information structure. Figure 3 contrasts the space of functions induced by all random variables whose expectation is the prior mean (any convex function

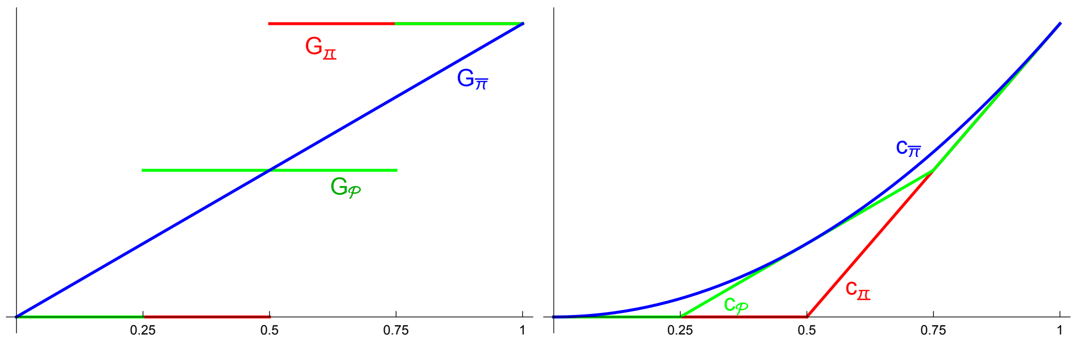

(a) CDFs of posterior means

(b) Functions induced by signals

Fig. 2: Signals as convex functions

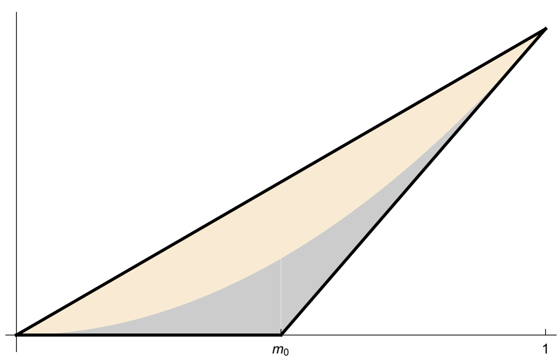

Fig. 3: Feasible distributions of posterior means vs. all random variables with $\mathbb{E}[\tilde{m}] = m_0$

in the lightly shaded area) with the subset of those that represent feasible distributions of the posterior means (any such function in the darker area in the bottom right).

This representation could be used to solve a wide class of optimization problems. For example, an active area of research where we care about a function of the mean of the posterior is costly search. In costly search models with uncertain information available to the searcher, information provision can be thought of as a stage within a larger game. For example, Board and Lu consider a setting where a buyer is searching for an item and each seller can choose its own information disclosure policy [31]. This idea could also inform the work on algorithmic problems faced by matching platforms or information intermediaries in search markets, who must decide what type of information to provide to participants and at what cost. In one-sided models (e.g. Chhabra *et al* [8]), the platform may be able to capture surplus through advertising, for example, and may wish to choose an information structure that maximizes social welfare. In two-sided models (and more general k-sided ones like those studied by Nahum *et al*), the information intermediary is serving as a matchmaker, and may wish to create the best overall set of teams under uncertainty about future arrivals [32]. Nahum *et al* show that more information is not necessarily better in their setting, but do not determine the optimal information structure.

III. Games of information provision

Now, instead of an optimization problem, consider a game where a number of agents choose what information to provide about some common state of the world $\omega \in \Omega$. In this case, modeling the information provided by agent i simply as $\pi_i : \Omega \to \Delta(\Omega)$ is no longer sufficient, since it does not specify the correlation of signal realizations from different agents given the state. For example, if we have $\pi_i = \pi_j$ we do not know whether the information provided by those two agents is redundant or not. Instead, Gentzkow

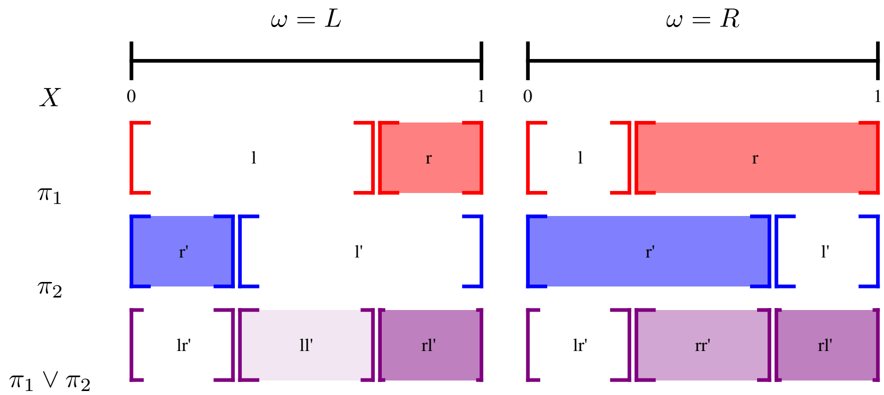

Fig. 4: The join of two signals

and Kamenica define a *signal* as a finite partition of $\Omega \times [0,1]$ s.t. $\pi \subset S$, where S is the set of non-empty Lebesgue measurable subsets of $\Omega \times [0,1]$ [22]. Any element $s \in S$ is a *signal realization.*

With each signal π we associate an S-valued random variable that takes value $s \in \pi$ when $(\omega, x) \in s$. Let $p(s|\omega) = \lambda(\{x| (\omega, x) \in s\})$ and $p(s) = \sum_{\omega \in \Omega} p(s|\omega)\, \mu_0(\omega)$ where $\lambda(\cdot)$ denotes the Lebesgue measure. For any $s \in \pi$, $p(s|\omega)$ is the conditional probability of s given ω and $p(s)$ is the unconditional probability of s.

This representation of information has the benefit of inducing an algebraic structure on the set of signals. The structure allows us to "add" signals together and thus easily examine their joint information content. Let Π be the set of all signals. Let $\trianglerighteq$ denote the refinement order on Π, that is, $\pi_1 \trianglerighteq \pi_2$ if every element of π_1 is a subset of an element of π_2. The pair $(\Pi, \trianglerighteq)$ is a lattice. The join $\pi_1 \vee \pi_2$ of two elements of Π is defined as the supremum of $\{\pi_1, \pi_2\}$.

Note that $\pi_1 \vee \pi_2$ is a signal that consists of signal realizations s such that $s = s_1 \cap s_2$ for some $s_1 \in \pi_1$ and $s_2 \in \pi_2$. Hence, $\pi_1 \vee \pi_2$ is the signal that yields the same information as observing both signal π_1 and signal π_2. In this sense, the binary operation $\vee$ "adds" signals together. The join of two signals is illustrated in Figure 4.

With this representation of information structures, one can not only study games of information provision but also ask questions about how to allocate information across individuals in order to have specific coalitions of those individuals have specific information. This could have immediate application to problems in multi-agent teams, including that of how to optimize multi-agent search teams in the spirit of Rochlin and Sarne [5]. It could also be applied to the design of deliberative auctions [33], [7], in order to find the signal structure for agents that best serves the auction designer's purpose.

IV. Conclusion

The design of the information environment is often as important as mechanism design in determining how the rules of interaction affect the outcomes of multi-agent systems, from economic markets to robot teams. There have been recent developments within economics that enhance our understanding of the various ways to think about the space of information structures. We have summarized three distinct representations of information in this paper, highlighting the types of problems that they make tractable, and speculating about possible applications to problems at the intersection of computer science and economics. In addition to solving specific existing problems, having these tools available may enable the community to engage with new models and computational questions.

References

[1] D. C. Parkes and M. P. Wellman, "Economic reasoning and artificial intelligence," *Science*, vol. 349, no. 6245, pp. 267–272, 2015.

[2] R. P. McAfee, "The design of advertising exchanges," *Review of Industrial Organization*, vol. 39, no. 3, pp. 169–185, 2011.

[3] N. Alon, F. A. Fischer, A. D. Procaccia, and M. Tennenholtz, "Sum of us: Strategyproof selection from the selectors," in *Proc. Conference on Theoretical Aspects of Rationality and Knowledge*, 2011, pp. 101–110.

[4] C. Hajaj and D. Sarne, "Strategic information platforms: Selective disclosure and the price of "free"," in *Proc. ACM Conference on Economics and Computation*, 2014, pp. 839–856.

[5] I. Rochlin and D. Sarne, "Constraining information sharing to improve cooperative information gathering," in *Proc. International Conference on Autonomous Agents and Multiagent Systems*, 2014, pp. 237–244.

[6] K. Larson and T. Sandholm, "Mechanism design and deliberative agents," in *Proc. International Conference on Autonomous Agents and Multiagent Systems*, 2005, pp. 650–656.

[7] E. Brinkman, M. P. Wellman, and S. E. Page, "Signal structure and strategic information acquisition: Deliberative auctions with interdependent values," in *Proc. International Conference on Autonomous Agents and Multiagent Systems*, 2014, pp. 229–236.

[8] M. Chhabra, S. Das, and D. Sarne, "Expert-mediated sequential search," *European Journal of Operational Research*, vol. 234, no. 3, pp. 861–873, 2014.

[9] ——, "Competitive information provision in sequential search markets," in *Proc. International Conference on Autonomous Agents and Multiagent Systems*, 2014, pp. 565–572.

[10] S. Das and Z. Li, "The role of common and private signals in two-sided matching with interviews," in *Web and Internet Economics (WINE)*. Springer, 2014, pp. 492–497.

[11] E. Kamenica and M. Gentzkow, "Bayesian persuasion," *American Economic Review*, vol. 101, no. 6, pp. 2590–2615, 2011.

[12] L. Rayo and I. Segal, "Optimal information disclosure," *Journal of political Economy*, vol. 118, no. 5, pp. 949–987, 2010.

[13] P. Jehiel, "On transparency in organizations," *The Review of Economic Studies*, vol. 82, no. 2, pp. 736–761, 2015.

[14] W. Gick and T. Pausch, "Persuasion by stress testing," 2012, working paper.

[15] I. Goldstein and Y. Leitner, "Stress tests and information disclosure," 2015, working paper.

[16] N. Schweizer and N. Szech, "Optimal revelation of life-changing information," 2013, working paper.

[17] A. Kolotilin, "Experimental design to persuade," *Games and Economic Behavior*, vol. 90, pp. 215–226, 2015.

[18] S. Gehlbach and K. Sonin, "Government control of the media," *Journal of Public Economics*, vol. 118, pp. 163–171, 2014.

[19] J. Ely, A. Frankel, and E. Kamenica, "Suspense and surprise," *Journal of Political Economy*, vol. 123, no. 1, pp. 215–260, 2015.

[20] D. Bergemann, B. A. Brooks, and S. Morris, "The limits of price discrimination," *American Economic Review*, vol. 105, no. 3, pp. 921–957, 2015.

[21] M. Gentzkow and E. Kamenica, "A Rotschild-Stiglitz approach to Bayesian persuasion," *American Economic Review Papers & Proceedings*, 2015, forthcoming.

[22] ——, "Competition in persuasion," 2015, working paper.

[23] ——, "Costly persuasion," *The American Economic Review Papers & Proceedings*, vol. 104, no. 5, pp. 457–462, 2014.

[24] C. A. Sims, "Implications of rational inattention," *Journal of Monetary Economics*, vol. 50, no. 3, pp. 665–690, 2003.

[25] E. P. Lazear, "Speeding, terrorism, and teaching to the test," *The Quarterly Journal of Economics*, vol. 121, no. 3, pp. 1029–1061, 2006.

[26] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus, "Playing games for security: An efficient exact algorithm for solving Bayesian Stackelberg games," in *Proc. International Conference on Autonomous Agents and Multiagent systems*, 2008, pp. 895–902.

[27] D. Korzhyk, Z. Yin, C. Kiekintveld, V. Conitzer, and M. Tambe, "Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness," *Journal of Artificial Intelligence Research*, vol. 41, pp. 297–327, 2011.

[28] Z. Rabinovich, A. X. Jiang, M. Jain, and H. Xu, "Information disclosure as a means to security," in *Proc. International Conference on Autonomous Agents and Multiagent Systems*, 2015, pp. 645–653.

[29] M. Saar-Tsechansky and F. J. Provost, "Active sampling for class probability estimation and ranking," *Machine Learning*, vol. 54, no. 2, pp. 153–178, 2004.

[30] P. Donmez and J. G. Carbonell, "Proactive learning: Cost-sensitive active learning with multiple imperfect oracles," in *Proc. ACM Conference on Information and Knowledge Management*, 2008, pp. 619–628.

[31] S. Board and J. Lu, "Competitive information disclosure in search markets," 2015, working paper.

[32] Y. Nahum, D. Sarne, S. Das, and O. Shehory, "Two-sided search with experts," in *Proc. ACM Conference on Electronic Commerce*, 2012, pp. 754–771.

[33] D. R. M. Thompson and K. Leyton-Brown, "Dominant-strategy auction design for agents with uncertain, private values," in *Proc. AAAI Conference on Artificial Intelligence*, 2011, pp. 745–750.