Table of Contents

Editor's Introduction

HU FU
University of British Columbia

This July I took over from Shaddin Dughmi as the editor of SIGecom Exchanges. My predecessors have laid a solid base for this newsletter as a playground for vivid communications on research activities in Economics and Computation. I will do my best to continue this great tradition, and welcome ideas on both the contents and innovations in the form or scope of the newsletter, in the hope for even more active engagement with the community.

In this issue, apart from four letters on recent research progress, we have a very informative survey by Shaddin Dughmi on Algorithmic Information Structure Design, an area which has seen exciting progress in recent years. We also have a letter by Nikhil Devanur, reporting on the First Workshop on the Economics of Cloud Computing. As the program chair of the workshop, Nikhil summarized each keynote and contributed talk at the workshop.

I would like to thank Shaddin Dughmi for his help with the transition, and Felix Fischer for putting the issue together.

Author's address: hufu@cs.ubc.ca

Algorithmic Information Structure Design: A Survey

SHADDIN DUGHMI
University of Southern California

Information structure design, also sometimes known as *signaling* or *persuasion*, is concerned with understanding the effects of information on the outcomes of strategic interactions (the descriptive question), and in characterizing and computing the information sharing strategies which optimize some design objective (the prescriptive question). Both questions are illuminated through the lens of algorithms and complexity, as evidenced by recent work on the topic in the algorithmic game theory community. This monograph is a biased survey of this work, and paints a picture of the current state of progress and challenges ahead.

We divide information structure design into *single agent* and *multiple agent* models, and further subdivide the multiple agent case into the *public channel* and *private channel* modes of information revelation. In each of these three cases, we describe the most prominent models and applications, survey the associated algorithms and complexity results and their structural implications, and outline directions for future work.

Categories and Subject Descriptors: F.0 [**Theory of Computation**]: General

General Terms: Algorithmic, Economics, Theory

Additional Key Words and Phrases: Signaling, Persuasion, Information Structures

1. INTRODUCTION

There are two primary ways of influencing the behavior of self-interested agents: by providing *incentives*, or by influencing *beliefs*. The former is the domain of traditional mechanism design, and involves the promise of tangible rewards such as goods and/or money. The latter, the focus of this survey and the subject of a recent flurry of interest in both the computer science and economics communities, involves the selective provision of payoff-relevant information to agents through strategic communication. Such "sweet talk" was estimated by [McCloskey and Klamer 1995] to account for a quarter of all economic activity in the United States in 1995, and the estimate has since been revised to 30% [Antioch 2013]. This is emblematic of the emergence of large-scale social and economic networks, with countless transactions and interactions among asymmetrically-informed parties occurring daily.

The primary object of interest in this topic is the *information structure* of a game of incomplete information. Informally, the information structure determines "who knows what" about the payoff structure of the game, and in doing so determines the set of equilibria. More formally, an information structure maps the *state of nature* θ — a parameter or set of parameters which determines the payoff function of the game — to a *signal* for each agent in the game. The map is typically randomized, and therefore reveals partial and noisy information regarding the payoffs

Author's address: `shaddin@usc.edu`. Supported by NSF CAREER award CCF-1350900 and NSF Grant CCF-1423618.

of various strategies. Like traditional incentive mechanisms, information structures can be studied either descriptively or prescriptively. The latter concerns the task faced by a *principal* who can control agents' access to information and wishes to optimize some objective function at the resulting equilibrium. Also like in mechanism design, the prescriptive question is naturally algorithmic, and studying it with the computational lens both provides structural insights and paves the way towards application. The task of optimizing the information structure in order to further an objective is often referred to as *information structure design*, *signaling*, or *persuasion*. In this survey, we refer to an information structure as a *signaling scheme* when we wish to think of it as a randomized algorithm — implemented by the principal — which takes as input a state of nature and outputs a signal for each agent.

1.1 The Basic Model and Assumptions

We use Θ to denote the family of states of nature, and assume $\theta \in \Theta$ is drawn from a common knowledge prior distribution μ. In all models we consider in this survey, the order of events for a game with n agents is as follows:

(1) The principal *commits* to a signaling scheme φ.
(2) Nature draws $\theta \sim \mu$.
(3) Signals $(\sigma_1, \ldots, \sigma_n) \sim \varphi(\theta)$ are drawn, and agent i learns σ_i.
(4) Agents select their strategies, and receive payoffs as determined by the game and the state of nature.

This survey adopts the perspective of the principal who has an objective in mind, and we focus on the optimization task faced by the principal in Step (1). A few notes on this general setup are in order. First, we note that signals are best thought of not as meaningful strings, but rather as abstract objects. Indeed, a rational agent interprets a signal by virtue of how it is used by the scheme φ, and therefore it has no intrinsic meaning beyond that. Second, it might seem unrealistic that the principal has the power to implement an arbitrarily informative signaling scheme. However, as pointed out in [Kamenica and Gentzkow 2011], this is without loss of generality: we can simply interpret the most informative signal the principal can access as the state of nature. Third, the reader might have noticed that we made no mention of information received by the agents which is out of the control of — and perhaps even unknown to — the principal. This is mostly for simplicity, and in fact some applications of the models we describe do involve agents who receive an exogenous signal, often referred to as the agent's *type*. Fourth — and this is related to the previous point — we restricted our attention to a one-step protocol of information revelation. In games where agents are privately informed, it might in fact be in the principal's interest to engage in a multi-round protocol where the principal and the agents exchange information piecemeal. The study of such protocols and settings is interesting in its own right, yet beyond the scope of this survey.

Last but not least, we justify what might at first appear as the most controversial assumption employed in most of the recent literature on information structure design. This is the *commitment assumption*: we assume that the principal has the power to credibly commit to the signaling scheme before realization of the state of nature. Without the power of commitment, the model becomes one of *cheap talk* (see [Crawford and Sobel 1982; Sobel 2010]). As is common in the recent literature on information structure design, we argue that the commitment assumption is not as unrealistic as it might first seem. One argument, mentioned in [Rayo and Segal 2010], is that commitment arises organically at equilibrium if the game is played repeatedly with a long horizon. This is because in such settings, the principal maximizes his long-term utility by establishing a reputation for credibility. A somewhat different argument, particularly suited for the algorithmic view of information structure design, is that any entity deploying an automated signaling scheme is likely to have a contractual service agreement with the agents, or otherwise have a vested interest in being perceived as a trusted authority. Trusting such an entity with the provision of information is not too unlike trusting an auctioneer to faithfully implement the rules of an auction, or trusting a certificate authority to properly issue digital certificates. In the case of signaling, commitment can involve publishing the source code of the signaling scheme used by the entity; agents can then verify the commitment over time through the use of statistical tests or audits. Additional justifications of the commitment assumption can be found in [Kamenica and Gentzkow 2011].

1.2 Structure of This Survey

We focus on three models, which to our knowledge capture or come close to capturing most recent work on information structure design. In Section 2 we consider the single-agent information structure design problem, also known as *Bayesian persuasion*. This is a special case of the next two models. In Section 3 we consider multiple agents, but a principal constrained to a public communication channel. The third model, considered in Section 4, affords the most power to the principal by permitting private communication between the principal and the individual agents. For all three of these models, we describe the mathematical setup, present structural characterizations (typically of a geometric nature) of the optimal information structure, discuss the state-of-the-art in algorithmic and complexity-theoretic work in that setting, and present open questions. We close this survey by briefly describing variations and extensions of these models in Section 5, and present some concluding thoughts in Section 6.

2. PERSUADING A SINGLE AGENT

Specializing information structure design to the case of a single agent yields the *Bayesian Persuasion* model proposed by [Kamenica and Gentzkow 2011], generalizing an earlier model by [Brocas and Carrillo 2007]. This is arguably the simplest model of information structure design, and the most applied. Indeed, the Bayesian persuasion model has been applied to a number of domains such as bilateral trade [Bergemann et al. 2015], advertising [Chakraborty and Harbaugh 2014], security games [Xu et al. 2015; Rabinovich et al. 2015], medical research [Kolotilin 2015], and financial regulation [Gick and Pausch 2012; Goldstein and Leitner 2013], just to mention a few. In addition to being interesting in its own right, the Bayesian persuasion model serves as a building block for more complex models of information structure design, and illustrates many of the basic principles which we will refer to in future sections of this survey.

2.1 The Model and Examples

In Bayesian persuasion, we adopt the perspective of a *sender* (the principal) looking to persuade a *receiver* (the single agent) to take an action which is desirable to the sender. There is a set A of actions, and the payoff of each action $a \in A$ to both the sender and the receiver is determined by the state of nature $\theta \in \Theta$ — we use $s(\theta, a)$ and $r(a, \theta)$ to denote those payoffs, respectively. We assume θ is drawn from a *common prior distribution* μ, and the sender must commit to a signaling scheme $\varphi : \Theta \to \Delta(\Sigma)$, where Σ denotes some set of signals and $\Delta(\Sigma)$ denotes the family of distributions over Σ. To illustrate this model, we look at a pair of examples.

EXAMPLE 2.1 (ADAPTED FROM [KAMENICA AND GENTZKOW 2011]). *Consider an academic adviser (the sender) who is writing a recommendation letter (the signal) for his graduating student to send to a company (the receiver), which in turn must decide whether or not to hire the student. The adviser gets utility 1 if his student is hired, and 0 otherwise. The state of nature determines the quality of the student, and hence the company's utility for hiring the student. Suppose that the student is excellent with probability* $\frac{1}{3}$*, and weak with probability* $\frac{2}{3}$*. Moreover, suppose that the company gets utility* 1 *for hiring an excellent student, utility* -1 *for hiring a weak student, and utility* 0 *for not hiring. Consider the following signaling schemes:*

—*No information: Given no additional information, the company maximizes its utility by not hiring. The adviser's expected utility is 0.*

—*Full information: Knowing the quality of the student, the company hires if and only if the student is excellent. The adviser's expected utility is* $\frac{1}{3}$*.*

—*The optimal (partially informative) scheme: The adviser recommends hiring when the student is excellent, and with probability just under* 0.5 *when the student is weak. Otherwise, the adviser recommends not hiring. The company maximizes its expected utility by following the recommendation, and the adviser's expected utility is just under* $\frac{2}{3}$*.*

EXAMPLE 2.2 (ADAPTED FROM [DUGHMI AND XU 2016]). *Example 2.1 can be generalized so that the receiver has many possible actions. The adviser has a number of graduating students, and the company must choose to hire one of them. The qualities of the different students are i.i.d.; specifically, each student is equally likely to be weak (W), a short-term achiever (S), and a long-term achiever (L). The company derives utility* 0 *from hiring a* W *student, utility* $1+\epsilon$ *for hiring an* S *student, and utility* 2 *from hiring an* L *student. The adviser, on the other hand, is up for tenure soon and derives utility* 1 *if the company hires an* S *student and* 0 *otherwise. Consider the following signaling schemes:*

—*No information: All students appear identical to the company, which chooses arbitrarily. The hired student is of type* S *with probability* $\frac{1}{3}$*, and therefore the adviser's expected utility is* $\frac{1}{3}$*.*

—*Full information: Knowing the quality of all students, the company hires a student of type* L *whenever one is available. As the number of students grows large, the adviser's utility tends to* 0*.*

$$\begin{array}{lll} \text{maximize} & \sum_{\theta \in \Theta} \mu(\theta) \sum_{a \in A} \varphi(\theta, a) s(\theta, a) & \\ \text{subject to} & & \\ & \sum_{\theta} \mu(\theta)\varphi(\theta, a)(r(\theta, a) - r(\theta, a')) \geq 0, & \text{for } a, a' \in A. \\ & \sum_{a \in A} \varphi(\theta, a) = 1, & \text{for } \theta \in \Theta. \\ & \varphi(\theta, a) \geq 0, & \text{for } \theta \in \Theta, a \in A. \end{array}$$

Fig. 1. Linear Program for Optimal Bayesian Persuasion

—*The optimal (partially informative) scheme: When at least one student has type S, the adviser recommends one of the S students uniformly at random. Otherwise, he recommends a student uniformly at random. Using the fact that the company prefers a student of type S to an equal mixture of types W and L, a simple calculation using Bayes' rule reveals that the company maximizes its expected utility by following the adviser's recommendation. As the number of students grows large, the adviser's utility tends to* 1.

2.2 Characterization of the Optimal Scheme

The reader might notice that, in both Examples 2.1 and 2.2, the signals of the optimal scheme correspond to the different actions, and can be thought of as *recommending* an action to the receiver. Moreover, when the receiver is recommended an action a, this recommendation is *persuasive* in the sense that a maximizes the receiver's expected payoff with respect to the posterior distribution of the state of nature θ induced by the signal a. This is not a coincidence: as observed by [Kamenica and Gentzkow 2011], an argument similar to the *revelation principle* shows that every signaling scheme is equivalent to one which recommends an action subject to such persuasiveness — i.e., we may assume without loss that $\Sigma = A$.[1] This dimensionality-reduction holds not only in this single-agent setting, but also in multi-agent settings when the agents are sent private signals (more on this in Section 4).

Given this characterization, it is not hard to see that the sender's optimal signaling scheme — i.e., the scheme maximizing his expected utility — is the solution to a simple linear program (Figure 1). This LP has a variable $\varphi(\theta, a)$ for each state of nature θ and action a, corresponding to the conditional probability of recommending action a given state θ. Solving this LP is impractical unless the prior distribution μ is of small support and given explicitly, but it nevertheless serves as a useful structural characterization. The LP maximizes the sender's utility, in expectation over the joint distribution of θ and a, subject to persuasiveness. Another way to visualize the feasible region of this LP is instructive: the probability of signal a is $\Pr[a] = \sum_{\theta} \mu(\theta)\varphi(\theta, a)$, and the posterior distribution μ_a on states of nature induced by signal a is given by $\mu_a(\theta) = \frac{\mu(\theta)\varphi(\theta,a)}{\Pr[a]}$. Therefore, a feasible solution of the LP can be thought of as a distribution over posteriors — one per signal — whose expectation equals the prior μ. In other words, if the prior μ is represented by a point in the probability simplex $\Delta = \Delta(\Theta)$, then the signaling scheme corresponds to a way of writing μ as a convex combination of posterior

[1]Persuasiveness has been referred to as *obedience* or *incentive compatibility* in the prior literature on persuasion.

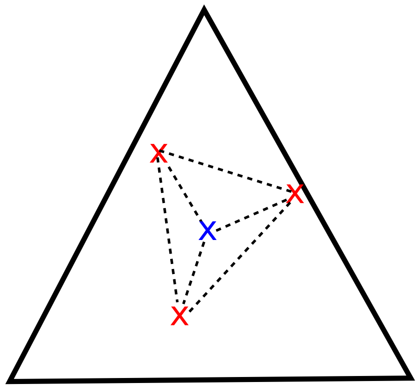

Fig. 2. A Signaling Scheme with Three Signals

distributions in Δ, as illustrated in Figure 2. Persuasiveness is equivalently stated as the constraint that action a is favored by a rational receiver facing the posterior distribution on θ induced by the signal a.

This geometric / LP interpretation reveals some structural properties of optimal signaling schemes. Each posterior distribution $\mu' \in \Delta$ is associated with a preferred action $a = a(\mu')$ for the receiver — this is the action maximizing the receiver's expected utility $\mathbf{E}_{\theta \sim \mu'} r(\theta, a)$.[2] This allows us to plot the sender's utility as a function $f : \Delta \to \mathbb{R}$ of the posterior: $f(\mu') = \mathbf{E}_{\theta \sim \mu'} s(\theta, a(\mu'))$. Let $\widehat{f}$ be the *concave closure*[3] of f; our geometric interpretation implies that the optimal signaling scheme achieves sender utility equal to $\widehat{f}(\mu)$ by optimally decomposing μ into posterior distributions. In other words, optimal signaling can be thought of as computing the concave closure of f (This was observed by [Kamenica and Gentzkow 2011]). This task is nontrivial only when f is neither convex nor concave. If f is concave then the optimal scheme reveals no information (i.e., recommends the receiver's ex-ante preferred action regardless of the state of nature). Whereas if f is convex the optimal scheme reveals all information (i.e., recommends the receiver's ex-post preferred action for each state of nature). For example, f is convex when sender and receiver utilities are always equal, and concave when they sum to zero. In general, this geometric interpretation yields a different bound on the number of signals via Caratheodory's theorem: the optimal scheme needs no more signals than the number of states of nature.

2.3 Independent Distributions and the Connection to Auctions

In mechanism design, the salient properties of optimal policies are often revealed when the underlying uncertainty admits a simple structure. The prime example of this is Myerson's characterization [Myerson 1981] of the optimal single-item auction: when bidder values are independent, the optimal auction maximizes what is known

[2] In case of ties, we break them in favor of the sender.
[3] The concave closure of a function f is the pointwise smallest concave function upperbounding f, obtained by taking the convex hull of the points beneath the plot of f.

as the *virtual welfare*, and when they are i.i.d. this optimal auction is furthermore a second-price auction with a reserve price. This characterization leads to an efficient algorithm implementing the optimal auction when bidders' distributions are independent and given explicitly. The analogous questions in the Bayesian Persuasion model are the following: *Is there an efficient algorithm for optimal persuasion when actions are independent or i.i.d.? Can this algorithm be captured by a simple "rule of thumb" à la virtual welfare maximization?*

The situation for persuasion turns out to be more nuanced than in the case of auctions. An efficient optimal algorithm exists in the case of i.i.d. actions (as in Example 2.2), but there is evidence that it might not exist in the case of independent non-identical actions. To be precise, when we say actions are independent we mean that $s(a) = s(\theta, a)$ and $s(a') = s(\theta, a')$ are independent random variables for different actions $a \neq a'$, and the same for $r(a) = r(\theta, a)$ and $r(a') = r(\theta, a')$. In this case, the distribution μ is fully specified by the marginal distribution of the pair $(s(a), r(a))$ for each action a. We assume that each action's marginal distribution has finite support, and refer to each element of this support as a *type*. We prove the following in [Dughmi and Xu 2016]; n denotes the number of actions and m denotes the number of types for each action.

Theorem 2.3 [Dughmi and Xu 2016]. *There is a polynomial-time (in n and m) algorithm implementing the optimal scheme for Bayesian persuasion with i.i.d. actions.*

Theorem 2.4 [Dughmi and Xu 2016]. *Unless $P = \#P$, there is no polynomial-time (in n and m) algorithm for evaluating the sender's optimal utility in Bayesian persuasion with independent actions.*

Theorem 2.3 relies on a connection to auction theory, and in particular to Border's characterization of reduced-form allocation rules for a single-item auction. The analogy to a single-item auction is as follows: each action corresponds to a bidder, the types of an action correspond to the types of bidders, and recommending an action corresponds to selecting the winning bidder. In this analogy, the joint distribution of recommended action and its type is simply what is known as the *reduced-form allocation rule* of a single-item auction. The space of reduced forms has been characterized by [Border 1991; 2007], and has since been observed to be a polytope admitting an efficient separation oracle [Cai et al. 2012; Alaei et al. 2012]. Therefore, computing the optimal scheme reduces to linear optimization over this polytope, augmented with persuasiveness constraints instead of auction incentive-compatibility constraints. We show that in the i.i.d. setting, persuasiveness constraints are linear in the reduced form, and the resulting linear program can be solved efficiently. This leads to the reduced form allocation rule associated with the optimal signaling scheme, and the optimal scheme can be implemented from that using Bayes' rule. Notably, whereas this algorithm is efficient, it arguably lacks the simplicity of Myerson's optimal auction — the difference is due to replacing the auction incentive-compatibility constraints with persuasiveness constraints.

Theorem 2.4 shows that, unlike in the case of auctions, the tractability of persuasion in the i.i.d. setting does not appear to directly extend to the independent setting. The culprit yet again are the persuasiveness constraints: unlike in the i.i.d.

setting, these can no longer be expressed as linear constraints in the reduced form. Moreover, our result shows that no generalization of the reduced form works either, ruling out a *generalized Border's theorem* for persuasion, in the sense of [Gopalan et al. 2015]. That being said, we note that this result rules out exactly computing the principal's optimal utility from signaling in polynomial time, yet does not rule out efficiently sampling the output of the optimal signaling scheme "on the fly", input by input.

2.4 General Black-Box Distributions

The previous subsection illustrates how simplifying the input model and using computation as a lens can lead to structural insights into optimal schemes. Another approach, common when computer scientists seek unifying and simplifying algorithmic and structural results, is the following: fix an input model which is as general as possible, and design an algorithm which succeeds regardless of the details of the setting. In Bayesian persuasion, the most general model permits an arbitrary joint distribution of sender and receiver payoffs from the various actions, allowing arbitrary correlations between the payoff-relevant variables. To be precise, there are $2n$ payoff-relevant random variables in total, where n is the number of actions: each action is associated with a payoff to both the sender and the receiver. In [Dughmi and Xu 2016], we assume that this joint distribution can be sampled from, and moreover that all the variables lie in a bounded interval (without loss of generality $[0, 1]$), and prove the following.

THEOREM 2.5 [DUGHMI AND XU 2016]. *For general payoff distributions presented as a sampling black box, an ϵ-optimal and ϵ-persuasive scheme (in the additive sense) for Bayesian persuasion can be implemented in time polynomial in n and $\frac{1}{\epsilon}$. Moreover, this bi-criteria loss is inevitable in the worst case: an optimal scheme must be $\Omega(1)$ far from being persuasive, and a persuasive scheme must be $\Omega(1)$ far from optimality.*

The positive statement in Theorem 2.5 concerns a simple scheme based on Monte-Carlo sampling: When our signaling scheme is queried with a state of nature $\theta \in \Theta$, it additionally samples polynomially many times from the prior μ to get a set $S \subseteq \Theta$, solves a relaxed variant of the LP in Figure 1 on the empirical distribution $S \cup \{\theta\}$, and produces a recommendation as suggested by the LP for θ.

Reflecting on Theorem 2.5 and the associated scheme, we can conclude that Bayesian persuasion admits an efficient, approximately-optimal, and approximately-persuasive scheme in arbitrarily general settings. Moreover, this scheme is simple: additional samples S are procured in order to place the query θ "in context" of the prior distribution μ, and the algorithm "pretends" that the prior distribution is the uniform distribution on $S \cup \{\theta\}$. Naturally, this succeeds due to convergence of the LP solution to the optimal solution. The LP in Figure 1 is modified by relaxing the persuasiveness constraints, as it so happens that this prevents "overfitting" to the sample. The negative statement in Theorem 2.5 shows that this relaxation is inevitable for information-theoretic reasons: one can construct examples where the sender — having only imprecise sampling knowledge of the prior distribution — can not be certain whether his preferred recommendation is persuasive to a receiver who has full knowledge of the prior.

2.5 Future Work and Open Questions

We now mention three open questions of an algorithmic flavor pertaining to the Bayesian persuasion model.

Open Question 2.6. *Consider Bayesian persuasion with n independent actions, each having a marginal distribution supported on m types. Is there a polynomial-time (in n and m) implementation of the optimal signaling scheme? If so, what does this algorithm reveal about the structure of the optimal policy?*

Recall that Theorem 2.4 rules out efficiently computing the optimal expected sender utility, yet does not preclude sampling $\varphi^*(\theta)$ for each input θ, where φ^* is an optimal scheme. This is a subtle point, but is not unprecedented: [Gopalan et al. 2015] exhibit simple auction settings in which the optimal revenue of the auctioneer is $\#P$ hard to compute, and yet Myerson's optimal auction can be efficiently implemented input-by-input. Theorem 2.4 implies that an optimal signaling scheme cannot be efficiently computed using linear-programming approaches à la Border's theorem. Therefore, if a similar phenomenon occurs here, the algorithm implementing the optimal signaling scheme would likely reveal some important structure of optimal persuasion, à la Myerson's famous characterization of optimal auctions as virtual-welfare maximizers [Myerson 1981]. This leads right into our next open question.

Open Question 2.7. *Can optimal Bayesian persuasion be described by a simple "rule of thumb" policy, à la virtual welfare maximization from auction theory?*

The results of [Dughmi and Xu 2016] do not answer this question, even in the simplest case of i.i.d. actions. Indeed, the result in Theorem 2.3 invokes Border's theorem to compute the entire reduced form for the optimal scheme, rather than provide a simple input-by-input rule such as virtual welfare maximization.

Our final question concerns moving beyond an explicitly-given list of actions. In a number of natural applications of persuasion, the receiver's action is naturally multidimensional — say a path in a network, a point in space, or an allocation of resources among different projects. In such settings, the actions lie in a vector space, and the receiver faces an optimization problem — say, encoded as a linear or convex program — when choosing their optimal action. When the state of nature determines the objective function of both the sender and the receiver, can we solve for an approximately-optimal signaling scheme in time polynomial in the natural parameters of the problem? Say, in the number of variables and constraints of the linear program rather than the number of its vertices (the possible actions)?

Open Question 2.8. *Consider Bayesian persuasion with multidimensional actions. In what settings can an optimal or near-optimal signaling scheme be computed in time polynomial in the dimensionality of the receiver's optimization problem?*

3. MULTIPLE AGENTS: PUBLIC SIGNAL

Information structure design can get much more intricate when multiple players are involved, particularly if they have heterogeneous beliefs, or if a scheme induces heterogeneous beliefs through revealing different information to different players. In this section, we examine a model which simplifies away such considerations:

	Coop	Defect
Coop	$-1+\theta$, $-1+\theta$	$-5+\theta$, 0
Defect	0, $-5+\theta$	-4, -4

Fig. 3. An incomplete information variant of the prisoners' dilemma.

all players (including our principal) share the same common prior on the state of nature, and we constrain our principal to a public communication channel — i.e., all players in the game receive the same information. This *public signal* model underlies much of the work on multi-agent information structure design, such as [Emek et al. 2012; Bro Miltersen and Sheffet 2012; Guo and Deligkas 2013; Dughmi et al. 2014] in the context of auctions, and [Alonso and Cmara 2016a; 2016b] in the context of voting.

3.1 The Model and Examples

An n-player game of incomplete information specifies a set A_i of actions for each player i, a set Θ of states of nature, and a payoff function $\mathcal{G} : \Theta \times A_1 \times \ldots \times A_n \to \mathbb{R}^n$, where $\mathcal{G}_i(\theta, a_1, \ldots, a_n)$ is player i's payoff when the state of nature is θ and each player j plays action a_j. The game may be represented explicitly via its normal form, or via some implicit representation permitting evaluation of the function $\mathcal{G}$.

We assume that the state of nature θ is distributed according to a *common prior distribution* $\mu \in \Delta(\Theta)$. A principal must commit to a signaling scheme $\varphi : \Theta \to \Delta(\Sigma)$, where Σ is some set of signals. We interpret the (random) output $\sigma \sim \varphi(\theta)$ as a public signal which is received by all players in the game, say through a public communication channel. The payoff function $\mathcal{G}$, the prior μ, and the signaling scheme φ then define a *Bayesian Game* in the classical game-theoretic sense. We naturally assume that players react by playing a Bayesian Nash Equilibrium in this game, possibly according to some domain-specific equilibrium selection rule in the case of multiple equilibria.

We adopt the perspective of the principal, looking to maximize some function in expectation over the realized state of nature and action profiles. At its most general, the principal's utility function is of the form $\mathcal{F} : \Theta \times A_1 \times \ldots A_n \to \mathbb{R}$. We present some examples below to make this model concrete.

EXAMPLE 3.1 [DUGHMI 2014]. *An incomplete information variant of the classical prisoners' dilemma is shown in Figure 3. The game's payoffs are parametrized by a state of nature θ. When $\theta = 0$, this is the traditional prisoners' dilemma in which cooperation is socially optimal, yet the unique Nash equilibrium is the one where both players defect, making both worse off. Assume, however, that θ is uniformly distributed in $[-3, 3]$, and assume the principal wishes to maximize the social welfare. Consider the following public signaling schemes:*

—*No information: The players, being risk neutral, play as if θ equals its expectation of 0. Defection dominates and the social welfare is -8.*

—*Full information: Defection dominates when $\theta < 1$ yielding welfare -8, and*

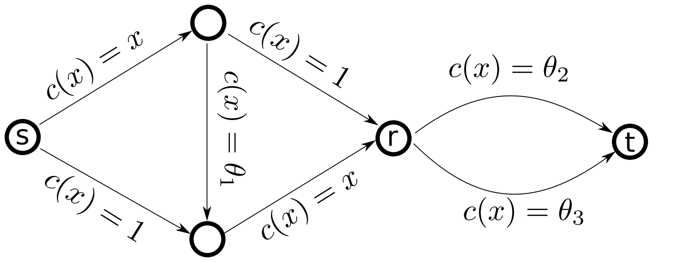

Fig. 4. An incomplete-information Routing Game

cooperation dominates when $\theta \geq 1$ *yielding welfare* $2\theta - 2$*. The expected social welfare is* $-\frac{14}{3}$*.*

—*The optimal (partially informative) scheme: When* $\theta > -1$ *signal* **High***, and otherwise signal* Low*. Given signal* **High***, which is output with probability* $\frac{2}{3}$*, both players' posterior expectation of* θ *is* 1*, and therefore cooperation dominates. Defection dominates when the signal is* Low*. The expected social welfare is* $-\frac{8}{3}$*.*

Example 3.2. *In a* non-atomic selfish routing game *(see [Nisan et al. 2007]), there is continuum of selfish agents looking to travel from a source s to a sink t in a directed network, and each edge of the network is labeled with a congestion function measuring the cost incurred by each agent as a function of the total fraction of flow using that edge. Consider the incomplete-information routing game depicted in Figure 4, in which the congestion functions are determined by a state of nature* $\theta = (\theta_1, \theta_2, \theta_3)$*. This network consists of a variant of the* Braess paradox *network (see [Nisan et al. 2007]), followed by a pair of parallel edges feeding into the sink.*

Suppose that the three components of θ *are i.i.d. random variables, each uniformly distributed on* $[0, 1]$*. Suppose also that the principal wishes to minimize the social cost, i.e. the average congestion experienced by the agents, at equilibrium, assuming agents are risk neutral. Consider the first portion of the journey, from s to r; As in Braess's paradox, the equilibrium routing from s to r is suboptimal when agents believe that the expectation of* θ_1 *is less than* 0.5*, and optimal otherwise. Therefore, it behooves the principal to reveal no information about* θ_1 *— this can be thought of as an* informational Braess's paradox*. In contrast, for the final hop of the journey from r to t the principal optimizes the social cost by revealing* θ_2 *and* θ_3*, as this assists all agents in choosing the lower congestion edge.*

To summarize, the optimal scheme reveals some components of the state of nature and withholds others.

Example 3.3. *In a* probabilistic single-item auction*, an item with a priori unknown attributes is being auctioned to a set of bidders. This arises in online advertising, where advertisers must bid on an impression, and this impression is associated with a web user drawn from a population of web users. [Emek et al. 2012] and [Bro Miltersen and Sheffet 2012] study optimal public signaling policies by a revenue-maximizing auctioneer in this setting, assuming the second-price auction format is fixed. In some settings — such as when bidders have similar preferences and the market is highly competitive — the optimal policy reveals all information* *about the item for sale. In other settings — such as when bidders have idiosyncratic preferences and markets are "thin" — withholding much of the information about the item can increase competition and drive up prices. In general, optimal policies reveal partial information. We refer the reader to [Emek et al. 2012; Bro Miltersen and Sheffet 2012] for a more detailed discussion.*

3.2 Characterization of the Optimal Scheme

As in the case of a single agent, we can identify a signaling scheme $\varphi : \Theta \to \Delta(\Sigma)$ with a way of writing the prior distribution μ as a convex combination of posterior distributions $\{\mu_\sigma : \sigma \in \Sigma\}$ (See Figure 2), where μ_σ is the posterior distribution on the state of nature given the signal σ. Unlike in the case of a single agent, however, we can no longer identify signals with actions. Indeed, each signal $\sigma \in \Sigma$ induces a subgame in which the common prior on the state of nature is μ_σ, and players might play a mixed Nash equilibrium in this subgame.

Fixing an equilibrium selection rule and denoting $\Delta = \Delta(\Theta)$, like in the single-agent case we get an objective function $f : \Delta \to \mathbb{R}$ mapping posterior distributions to the principal's expected utility for the resulting equilibrium. Like in the single-agent case, we can therefore also interpret the optimal scheme as evaluating the concave closure $\widehat{f}$ of f by optimally decomposing the prior μ into posterior distributions. Whereas we can no longer bound the number of signals by the number of actions, the bound from Caratheodory's theorem still holds: the optimal scheme needs no more signals than the number of states of nature.

3.3 Negative Results

Recall that in the single-agent setting of Section 2, a simple linear program expresses the optimal signaling task. Moreover, this LP is of size linear in the normal form of that game — i.e., linear in the number of (state of nature, action) pairs. This simplicity is largely a consequence of the dimensionality-reduction property in the single-agent case, and underlies the positive algorithmic and structural results outlined in Section 2.

It is natural to ask how quickly this structure deteriorates as we move beyond a single agent. The answer for the public signal model: very quickly. A series of works [Dughmi 2014; Bhaskar et al. 2016; Rubinstein 2015] examines signaling in 2-player zero-sum games, and culminates in the following.

THEOREM 3.4 [BHASKAR ET AL. 2016; RUBINSTEIN 2015]. *Consider a 2- player Bayesian zero-sum game with payoffs in* $[0, 1]$*, presented via its normal form, and a principal interested in maximizing the utility of one of the players. The principal's signaling problem is NP-hard [Bhaskar et al. 2016], and moreover does not admit an additive PTAS*[4] *assuming either the* planted clique conjecture *[Bhaskar et al. 2016] or the* exponential time hypothesis *[Rubinstein 2015].*

These impossibility results have significant bite: they hold in a setting where equilibrium selection and computation are a non-issue — all equilibria of a 2-player zero-sum game are equivalent to the efficiently-computable minimax equilibrium.

[4] A Polynomial-Time Approximation Scheme (PTAS) is an algorithm which, given any constant $\epsilon > 0$, computes an ϵ-approximately optimal solution in time polynomial in the size of the instance.

Moreover, maximizing one player's utility can be shown to be no harder than other natural choices of objective function, such as the social welfare (weighted or unweighted).[5] An arguably reasonable reading of these results is that *a simple and general characterization of optimal public signaling is unlikely to exist even in the simplest of multiagent settings.*

We mention two other impossibility results in specific game domains which reinforce this message. [Emek et al. 2012] show that revenue-maximizing public signaling in the second price auction (à la example 3.3) is NP-hard, and [Bhaskar et al. 2016] show that welfare-maximizing public signaling in selfish routing with linear congestion functions (à la example 3.2) is NP-hard to approximate to within a multiplicative factor better than $\frac{4}{3}$ — this is the price of anarchy in this setting, and is trivially achievable by the fully-informative scheme.

3.4 Positive Results: Exploiting "Smoothness"

The impossibility results stated in Theorem 3.4 involve constructing games where a near-optimal signaling scheme must reveal much — but not all — of the information contained in the state of nature. For example, the reductions in Theorem 3.4 feature a state of nature which is a random node in an n-node graph, and a near-optimal scheme must essentially partition the nodes into equivalence classes of size roughly $\sqrt{n}$, in effect revealing roughly half of the bits of information. Informally, this induces a combinatorial search problem which searches for the "right" half of the bits to reveal. It is not entirely surprising, therefore, that this problem can be intractable.

One might wonder if there are natural classes of games in which we can get away with revealing much less information. This would simplify the search problem, making it more computationally tractable. This idea is explored by [Cheng et al. 2015], who identify two "smoothness" properties which seem to govern the complexity of near-optimal signaling schemes. Suppose that each state of nature $\theta \in \Theta$ is a vector of N "relevant parameters" in a bounded interval — for example, in a probabilistic single-item auction N may be the number of different bidder types, and $\theta_t \in [0, 1]$ may be the item's value for bidders of type t. Moreover, suppose that the equilibrium of the game, and therefore the principal's utility, depend only on the posterior expectation of each of the relevant parameters. This is the case in the auction setting (Example 3.3), where risk-neutral bidders bid their expected value for the (random) item being sold. Finally, suppose that the game is "smooth" in two respects:

—α-Lipschitz continuity in L^∞: If each relevant parameter changes by at most ϵ, then the principal's utility does not decrease by more than $\alpha\epsilon$.

—β-Noise stability: Suppose that an adversary corrupts (i.e. changes arbitrarily) a random subset R of the relevant parameters, where no individual parameter is in R with probability more than ϵ. The principal's utility does not decrease by more than $\beta\epsilon$.

[5] It is clear how this is a special case of weighted welfare. Moreover, multiplying player 2's utility by a small constant approximately reduces the problem of maximizing player 1's utility to maximizing the unweighted social welfare.

Theorem 3.5 [Cheng et al. 2015]. *Suppose that a game is α-Lipschitz and β-noise stable in its relevant parameters, and let $\epsilon > 0$ be an approximation parameter. There is an ϵ-approximately optimal signaling scheme (in the additive sense) where each posterior belief is a uniform distribution over $(\frac{\alpha}{\epsilon})^2 \log(\frac{\beta}{\epsilon})$ states of nature. When α and β are constants, such a scheme can be computed in time polynomial in the number of states of nature and the number of relevant parameters, yielding an additive PTAS.*

This theorem implies a PTAS for revenue-maximizing signaling in the probabilistic second-price auction, since the second-price auction is 1-Lipschitz and 2-stable. For the former, observe that changing bids by no more than ϵ can only change the second price by at most ϵ. For the latter, if each bid is corrupted to an arbitrary value with probability at most ϵ, then with probability at least $1 - 2\epsilon$ the top two bids are untouched and the revenue does not decrease.

These two conditions, Lipschitz continuity and noise stability, imply that "small" posterior beliefs suffice for a near-optimal public signaling scheme. The proof of this portion of Theorem 3.5 proceeds by decomposing each posterior belief μ_σ of the optimal scheme into "small" posterior beliefs by sampling; i.e., μ_σ is written as the average of empirical distributions sampled from it. Sampling from a distribution over states of nature leads to (a) high-probability small errors in the relevant parameters, the effect of which is bounded using Lipschitz continuity of the objective; and (b) low-probability large errors, the effect of which is bounded using noise stability of the objective.

Once we can restrict attention to small posteriors, the computational task becomes tractable. Specifically, computing a near-optimal scheme reduces to a brute force search over small posterior distributions, using linear programming in order to assemble the convex decomposition of the prior distribution.

In addition to probabilistic second-price auctions, these ideas have been applied in [Cheng et al. 2015] to signaling in the voting setting of [Alonso and Cmara 2016a] yielding an approximation scheme, and to derive a quasipolynomial-time approximation scheme for signaling in normal-form games.

3.5 Future Work and Open Questions

We mention one direction for future work, concerning the extent to which the restriction to public signals is binding.

Open Question 3.6. *In what classes of games and objectives is a public signaling policy optimal or near optimal, when evaluated against the optimal policy which can send private signals?*

In other words, when does moving to private signals — which we discuss in Section 4 — not buy the principal additional power? An answer to this question in the special class of games with no inter-agent externalities and binary actions is provided by [Arieli and Babichenko 2016] (we discuss their model in detail in Section 4.3).

4. MULTIPLE AGENTS: PRIVATE SIGNALS

We now examine a model which affords the principal the most power, allowing him to tailor his signal to individual players. As in Sections 2 and 3, we still assume that the principal and all the agents share the same common prior on the state of nature. However, since the principal reveals different information to different agents, the agents' posterior beliefs will differ.[6] Though information structure design with private signals has not been very thoroughly explored, particularly algorithmically, most recent applications of private signaling fall under the model we outline in this section. Specifically, we mention the work of [Taneva 2015] who characterizes optimal information structures in two-agent two-action games, [Bardhi and Guo 2016] who study persuading voters in a unanimity election, and [Arieli and Babichenko 2016; Babichenko and Barman 2016] who study persuading multiple agents facing a binary action with no inter-agent externalities.

4.1 The Model and Examples

As in Section 3, we have an n-player game of incomplete information $\mathcal{G} : \Theta \times A_1 \times \ldots \times A_n \to \mathbb{R}^n$, a common prior distribution $\mu \in \Delta(\Theta)$ over states of nature, and an objective function $\mathcal{F} : \Theta \times A_1 \times \ldots A_n \to \mathbb{R}$. We again adopt the perspective of a principal who designs and commits to a signaling scheme with the goal of maximizing $\mathcal{F}$ in expectation, but unlike Section 3 this signaling scheme is of the form $\varphi : \Theta \to \Delta(\Sigma_1 \times \ldots \times \Sigma_n)$, where Σ_i is a set of signals intended for player i. The output of φ is a random signal profile $(\sigma_1, \ldots, \sigma_n)$, where σ_i is sent to payer i via a private channel. Together, $\mathcal{G}$, μ and φ define a Bayesian game.

We present two examples to illustrate the model, and to show how private signaling affords more power to the principal than does public signaling.

Example 4.1 (Adapted from [Arieli and Babichenko 2016]). *As in Example 2.1, consider an academic adviser who is writing a recommendation letter for his student. However, now the student has applied to two fellowship programs, each of which must decide whether or not to award the student a fellowship funding part of his graduate education. Suppose that the student can accept one or both fellowship awards. The adviser, who has enough grant funding for most (but not all) of his student's education, gets utility* 1 *if his student is awarded at least one fellowship, and* 0 *otherwise. As in Example 2.1, the student is excellent with probability* $1/3$ *and weak with probability* $\frac{2}{3}$*, and a fellowship program gets utility* 1 *from awarding an excellent student,* -1 *from awarding a weak student, and* 0 *from not awarding the student. Naturally, a fellowship program makes an award if and only it believes its expected utility for doing so is nonnegative.*

Consider the following signaling schemes:

—*No Information: Neither program makes the award, and the adviser's utility is* 0.

—*Full information: Both programs make the award if the student excellent, and neither makes the award if the student is weak. The adviser's expected utility is* $1/3$.

—*Optimal public scheme: If the student is excellent, the adviser publicly signals "award". If the student is weak, the adviser publicly signals "award" or "don't award" with equal probability. This scheme is the same as the optimal scheme for the single-receiver version of this example (Example 2.1), extended to both*

[6] Equally importantly, their higher order posterior beliefs regarding each other are different and nontrivial.

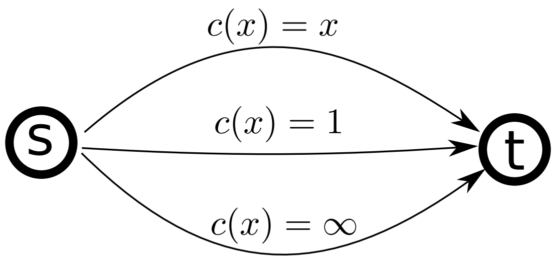

Fig. 5. Adapted Pigou Example

receivers via publicizing the recommendation letter. Therefore, both programs are simultaneously persuaded to award the student the fellowship with probability $2/3$*, and neither makes the award with probability* $1/3$*. The adviser's expected utility is* $2/3$*.*

—*Optimal private scheme: If the student is excellent, the adviser recommends "award" to both fellowship programs. If the student is weak, the adviser recommends "award" to one fellowship program chosen uniformly at random, and recommends "don't award" to the other. Notice that, from an individual program's perspective this is the same as the previous scheme, the difference being that the recommendations are anticorrelated when the student is weak. The result is that both fellowship programs make the award when the student is excellent, and exactly one of the programs makes the award when the student is weak. This yields utility* 1 *for the adviser.*

EXAMPLE 4.2 [CHENG AND XU 2016]. *This example concerns non-atomic selfish routing, as in Example 3.2. Consider the adaptation of the Pigou routing network (see [Nisan et al. 2007]) depicted in Figure 5. The optimal routing evenly splits the agents between the* x *and* 1 *edges leading to a social cost of* $\frac{3}{4}$*, whereas the full-information equilibrium routing sends all agents along the* x *edge leading to a social cost of* 1*. Suppose that nature applies a cyclic permutation to the tuple of congestion functions* $(x, 1, \infty)$ *— i.e. there are three states of nature. The principal can influence the routing by revealing information about the state of nature. Consider the following signaling schemes.*

—*No information: The average congestion is* ∞*, since one third of the agents will end up using the* ∞ *edge.*

—*Full information: The agents avoid the* ∞ *edge, and arrive at an equilibrium routing for the remaining Pigou network. At equilibrium, all travelers use the* x *edge, leading to a social cost of* 1*.*

—*Optimal public scheme: Though it takes proof, it can be shown that no public signaling scheme can do better than revealing full information. Intuitively, the principal's signal must identify at least one of the non-*∞ *edges with certainty, lest any of the agents use the* ∞ *edge.*[7] *If both non-*∞ *edges are identified, then this fixes the cyclic permutation and is equivalent to revealing full information.*

[7]By this we mean that the principal's signal must allow agents to conclude with certainty that

If, however, exactly one of the non-∞ edges is identified, all agents will use that edge in order to avoid the ∞ edge, leading to a social cost of 1 *as well.*

—*Optimal private scheme: The principal identifies one of the non-∞ edges to half the agents, and the other non-∞ edge to the other half. In other words, the principal privately recommends the x edge to a random half of the agents and privately recommends the* 1 *edge to the other half, without revealing whether the recommended edge has congestion function x or* 1*. Each agent is persuaded to follow the recommendation, since deviating from the recommendation lands them on the ∞ edge with probability $\frac{1}{2}$. This is the optimal routing for this network, and yields a social cost of $\frac{3}{4}$.*

4.2 Characterization of the Optimal Scheme

In both examples 4.1 and 4.2, the optimal private scheme recommends an action to each agent, and correlates these recommendations in a manner that optimizes the principal's objective. As in the single agent case (Section 2), this is not a coincidence: a revelation-principle-style argument reveals any signaling scheme is equivalent[8] to one which makes *persuasive* recommendations. In this multi-agent context, we say a recommendation scheme is persuasive if the agent maximizes his expected utility by always choosing the recommended action, assuming all other agents follow the recommendation. Equivalently, a scheme φ is persuasive if the strategy profile where each agent follows the recommendation forms a Bayes-Nash equilibrium of the Bayesian game $(\mathcal{G}, \mu, \varphi)$.

The above discussion might remind the astute reader of the *correlated equilibrium.* In fact, the joint distribution of action profile and state of nature induced by a signaling scheme at equilibrium forms what [Bergemann and Morris 2016] call a *Bayes Correlated Equilibrium (BCE).* Conversely, every BCE is induced by some persuasive scheme. Given a game of incomplete information $\mathcal{G}$ and a prior distribution μ on states of nature, one way to think of a BCE is as follows: It is a correlated equilibrium of $\mathcal{G}$ when we interpret nature as a player in the game,[9] endow the nature player with a trivial (i.e., constant everywhere) payoff function, and constrain the nature player's marginal strategy to be equal to μ. Therefore, like in the case of the correlated equilibrium, the space of BCEs can be expressed by a set of linear inequalities, and optimization over BCEs — equivalently, over signaling schemes — can be written as a linear program.

The LP for optimizing over BCEs is shown in Figure 6. Here, $\boldsymbol{a}$ ranges over action profiles, $\boldsymbol{a_{-i}}$ ranges over action profiles of players other than i, and θ ranges over states of nature. This LP generalizes the single-agent persuasion LP in Figure 1, modulo a simple change of variables. More interestingly, this LP is the same as the LP for optimizing over correlated equilibria (see e.g. [Papadimitriou and Roughgarden 2008]) with the exception of two differences: (a) the nature player has no incentive constraint, (b) the nature player's marginal distribution is constrained to

a particular edge does not have congestion function ∞; the signal may, however, provide no information on whether that same edge has congestion function x or 1.

[8]We say two signaling schemes are equivalent if they induce the same joint distribution of action profile and state of nature at equilibrium.

[9]Under this interpretation, $\mathcal{G}$ becomes a game of *complete* information

$$
\begin{array}{lll}
\max & \sum_{\theta,\boldsymbol{a}} x(\theta,\boldsymbol{a})\mathcal{F}(\theta,\boldsymbol{a}) & \\
\text{s.t.} & & \\
& \sum_{\theta,\boldsymbol{a_{-i}}} x(\theta,a_i,\boldsymbol{a}_{-i})[\mathcal{G}_i(\theta,a_i,\boldsymbol{a_{-i}}) - \mathcal{G}_i(\theta,a_i',\boldsymbol{a_{-i}})] \geq 0, & \text{for } i \in [n], a_i \in A_i, a_i' \in A_i. \\
& \sum_{\boldsymbol{a}} x(\theta,\boldsymbol{a}) = \mu(\theta), & \text{for } \theta \in \Theta. \\
& x(\theta,\boldsymbol{a}) \geq 0, & \text{for } \theta \in \Theta, \boldsymbol{a} \in A_1 \times \ldots \times A_n.
\end{array}
$$

Fig. 6. Linear Program for Finding the Optimal Bayes Correlated Equilibrium

equal the prior distribution. A solution x to this LP (a BCE) corresponds to the persuasive private signaling scheme which, given a state of nature θ, recommends action profile $\boldsymbol{a}$ with probability $\frac{x(\theta,\boldsymbol{a})}{\mu(\theta)}$.

This characterization in terms of the Bayes correlated equilibrium exposes the dual role of an information structure: (1) informing players, which in effect allows them to correlate their actions with the state of nature, and (2) coordinating players by serving as a correlating device (as in the correlated equilibrium).

4.3 The Case of No Externalities and Binary Actions

To our knowledge, there has not been much algorithmic work on the use of private signals to influence agent behavior in multiagent settings. The recent exception is the *private Bayesian persuasion* model introduced by [Arieli and Babichenko 2016] and explored via the computational lens by [Babichenko and Barman 2016]. This model restricts information structure design to games with two simplifying features: (1) One agent's action does not impose an externality (positive or negative) on the other agents, and (2) each agent has a binary choice of action. The no-externality assumption implies that each agent's utility can be written as a function of just the state of nature θ and that particular agent's action, without any dependence on the actions of others. The principal's objective, on the other hand, may depend arbitrarily on the joint action profile of all the agents (as well as the state of nature). Since each agent's action is binary, without loss $A_i = \{0, 1\}$, the principal's objective can be equivalently described by a set function f, where $f(S)$ is the principal's utility if S is the set of agents taking action 1. In most natural examples, action 1 corresponds to adoption of an product or opinion, and 0 corresponds to non-adoption.

This model is illustrated by Example 4.1. The no-externality assumption implies that the principal faces n different Bayesian persuasion problems, one per agent, each with a binary action space. However, solving these problems separately can be suboptimal due to the non-modular dependence of the principal's objective on the agents' actions. In fact, the signaling problem can often be thought of as how to optimally correlate the solutions to the n (single-agent) Bayesian persuasion problems in order the maximize the principal's expected utility. This is indeed the case for Example 4.1: as a result of the adviser's submodular objective function, the optimal scheme anti-correlates the fellowship programs' actions as much as possible, subject to maximizing the marginal probability of a fellowship award in both cases.

More generally, the results of [Arieli and Babichenko 2016] demonstrate that a submodular objective function encourages anticorrelating the agents' recommendations (as in Example 4.1), and a supermodular objective function encourages correlation.[10] As an example of the latter, consider changing the adviser's utility function in Example 4.1 so that the adviser gets utility 1 if both fellowships are awarded and 0 otherwise; some thought reveals that the public scheme which persuades both fellowship programs to simultaneously award with probability $\frac{2}{3}$ becomes optimal. If the principal's objective function is modular (a.k.a. linear) in the set of persuaded agents, such as if the adviser's utility equals the number of fellowships awarded, then it is optimal to solve the n Bayesian persuasion problems separately — i.e., correlation in agents' actions does not affect the principal's utility.

[Babichenko and Barman 2016] go on to examine computing or approximating the principal's optimal signaling policy. In the case of a binary state of nature, they show that the problem of optimally correlating agents' actions is equivalent to that of computing the *concave closure* of the principal's objective, viewed as a set function. This makes a lot of sense, since the concave closure of a set function f maps a profile of marginal probabilities — in our case a probability of persuading each agent to take action 1 — to the maximum expected value of f for a random set S respecting those marginals — in our case, S is the random set of "persuaded" agents, and the optimal choice correlates or anticorrelates agents' actions in order to maximize the expectation of $f(S)$. It is known that the concave closure of a supermodular function can be computed efficiently (see e.g. [Dughmi 2009]), and it is shown in [Babichenko and Barman 2016] that the concave closure of a submodular function can be efficiently approximated to within a factor of $\frac{e}{e-1}$, and this is the best possible assuming $P \neq NP$. This connection leads to the following Theorem.

Theorem 4.3 [Babichenko and Barman 2016]. *Consider the private Bayesian persuasion model with a binary action space and binary state of nature. If the principal's objective function is supermodular (or modular), then there is a polynomial-time optimal signaling scheme. If the principal's objective function is submodular, then there is a polynomial-time $\frac{e}{e-1}$-approximately optimal signaling scheme, and this is the best possible assuming $P \neq NP$.*

4.4 Future Work and Open Questions

With the exception of the work described in Section 4.3, the computational aspects of information structure design with private signals remain largely unexplored territory. For one, [Babichenko and Barman 2016] leave open the natural generalization of their algorithmic questions to a non-binary state of nature. More generally, it remains to explore the algorithmics of private signaling in games with inter-agent externalities and non-binary action spaces.

Open Question 4.4. *When does private signaling admit simple and computationally efficient schemes which are optimal or near optimal?*

Since the optimal private signaling scheme is the solution to a large linear program (Figure 6), this question is most interesting when the game is described via some succinct representation, or given implicitly as an oracle for evaluating the payoff

[10]Though these characterization results are only formally stated for a binary state of nature, the conceptual message seems to hold more generally.

function. One class of games which seems to capture the challenges involved is the class of non-atomic routing games of incomplete information, illustrated in Example 4.2.

5. ADDITIONAL MODELS AND EXTENSIONS

This survey attempted to summarize the dominant models in information structure design, particularly as it relates to recent work which explores computational aspects of the question. However, we inevitably cannot capture the full breadth of work in this area. This section briefly describes a selection of models beyond the three we focus on in this survey.

Motivated by specific applications, a number of works in the computer science community have explored variants and extensions of the basic models from a computational perspective. A pair of papers consider optimal signaling subject to a constraint on the amount of communication (i.e., the number of signals): [Dughmi et al. 2014] consider public signaling subject to a communication constraint in an auction context, and [Dughmi et al. 2016] study Bayesian persuasion subject to a communication constraint both in general and in the context of bilateral trade. Motivated by recommender systems on the Internet, [Kremer et al. 2014; Mansour et al. 2015] consider a multi-armed bandit setting and a principal seeking to persuade a sequence of agents to balance exploration with exploitation over time. This can be viewed as a repeated game which combines information revelation with information acquisition, and the principal's interaction with each individual agent can be viewed as a Bayesian persuasion problem. Motivated by applications to security games, [Conitzer and Korzhyk 2011; Xu et al. 2016] consider a Stackelberg setting where the leader commits to both a mixed strategy and a signaling scheme. The role of the signaling scheme is to reveal different information to different followers about the realization of the leader's strategy, and in doing so to improve the leader's utility.

We restricted attention to a state of nature drawn from a common prior, and a principal who simply discloses information to agents without soliciting information. Two recent works have explored relaxing these assumptions for the Bayesian persuasion model described in Section 2. [Alonso and Camara 2014] characterize the optimal signaling scheme when the sender and receiver have different prior distributions. [Kolotilin et al. 2015] consider a privately-informed receiver, and a sender who first solicits the agent's information and then selectively signals his own information. Under some assumptions, they characterize the optimal combined policy. To our knowledge, neither of these two generalizations of the Bayesian persuasion model has been explored algorithmically.

Finally, we mention that a number of related, but importantly different, models have a long history in the economics community, yet to our knowledge have not been explored from a computational perspective. Most notably, the literature on *cheap talk* does away with the power of commitment, and much of that work analyzes the equilibria of cheap talk games. Also, the literature on *verifiable information* restricts the principal to signals which are meaningful and "honest" in a precise sense. We refer the reader to the survey by [Sobel 2010] which compares and contrasts a variety of these models.

6. CONCLUSIONS

The past few years have seen an explosion of interest in understanding the effects of information on strategic interactions. Since information structure design, like traditional mechanism design, is fundamentally algorithmic, it was inevitable that computer science would have much to say on the topic. This survey illustrates how asking computational questions has led to new structural insights in this area. Moreover, we believe that information structure design has grown into a deep theory waiting for an application, and simple and efficient algorithms will pave the way for such impact.

In addition to the open problems mentioned in Sections 2 through 4, there is much work to be done beyond the three basic models. Specifically, Section 5 suggests a number of related or more expressive models which have not been subjected to rigorous examination through a computational lens. We also believe there is room for experimental work to test the predictive power of this theory. Specifically, to what extent is signaling effective when deployed against human agents? [Azaria et al. 2015] examine this question in a single-agent setting (a slight variant of the Bayesian persuasion model from Section 2), and their findings are encouraging in the pair of domains they consider. More experimental validation of the predictions of information structure design, both in single and multiple agent settings, remains to be done.

REFERENCES

Alaei, S., Fu, H., Haghpanah, N., Hartline, J. D., and Malekian, A. 2012. Bayesian optimal auctions via multi- to single-agent reduction. In *ACM Conference on Electronic Commerce*, B. Faltings, K. Leyton-Brown, and P. Ipeirotis, Eds. ACM, 17.

Alonso, R. and Camara, O. 2014. Bayesian persuasion with heterogeneous priors. *Available at SSRN 2306820*.

Alonso, R. and Cmara, O. 2016a. Persuading voters. *American Economic Review 106,* 11 (November), 3590–3605.

Alonso, R. and Cmara, O. 2016b. Political disagreement and information in elections. *Games and Economic Behavior 100*, 390 – 412.

Antioch, G. 2013. Persuasion is now 30 per cent of us gdp. *Economic Roundup* 1, 1–10.

Arieli, I. and Babichenko, Y. 2016. Private bayesian persuasion. *Available at SSRN 2721307*.

Azaria, A., Rabinovich, Z., Goldman, C. V., and Kraus, S. 2015. Strategic information disclosure to people with multiple alternatives. *ACM Transactions on Intelligent Systems and Technology (TIST) 5,* 4, 64.

Babichenko, Y. and Barman, S. 2016. Computational aspects of private bayesian persuasion. *arXiv preprint arXiv:1603.01444*.

Bardhi, A. and Guo, Y. 2016. Modes of persuasion toward unanimous consent. *Manuscript*.

Bergemann, D., Brooks, B., and Morris, S. 2015. The limits of price discrimination. *The American Economic Review 105,* 3, 921–957.

Bergemann, D. and Morris, S. 2016. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics 11,* 2, 487–522.

Bhaskar, U., Cheng, Y., Ko, Y. K., and Swamy, C. 2016. Hardness results for signaling in bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation.* ACM, 479–496.

Border, K. 2007. Reduced Form Auctions Revisited. *Economic Theory 31,* 1 (April), 167–181.

Border, K. C. 1991. Implementation of Reduced Form Auctions: A Geometric Approach. *Econometrica 59,* 4.

Bro Miltersen, P. and Sheffet, O. 2012. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 234–247.

Brocas, I. and Carrillo, J. D. 2007. Influence through ignorance. *The RAND Journal of Economics 38,* 4, 931–947.

Cai, Y., Daskalakis, C., and Weinberg, S. M. 2012. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*. 459–478.

Chakraborty, A. and Harbaugh, R. 2014. Persuasive puffery. *Marketing Science 33,* 3, 382–400.

Cheng, Y., Cheung, H. Y., Dughmi, S., Emamjomeh-Zadeh, E., Han, L., and Teng, S.-H. 2015. Mixture selection, mechanism design, and signaling. In *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science (FOCS).*

Cheng, Y. and Xu, H. 2016. personal communication.

Conitzer, V. and Korzhyk, D. 2011. Commitment to correlated strategies. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*. AAAI Press, 632–637.

Crawford, V. P. and Sobel, J. 1982. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, 1431–1451.

Dughmi, S. 2009. Submodular functions: Extensions, distributions, and algorithms. a survey. PhD Qualifying Exam Technical Report, Stanford University.

Dughmi, S. 2014. On the hardness of signaling. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science (FOCS).*

Dughmi, S., Immorlica, N., and Roth, A. 2014. Constrained signaling in auction design. In *Proceedings of the 25th ACM Symposium on Discrete Algorithms (SODA).*

Dughmi, S., Kempe, D., and Qiang, R. 2016. Persuasion with limited communication. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 663–680.

Dughmi, S. and Xu, H. 2016. Algorithmic bayesian persuasion. In *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC).*

Emek, Y., Feldman, M., Gamzu, I., Paes Leme, R., and Tennenholtz, M. 2012. Signaling schemes for revenue maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*. 514–531.

Gick, W. and Pausch, T. 2012. *Persuasion by stress testing: Optimal disclosure of supervisory information in the banking sector.* Number 32/2012. Discussion Paper, Deutsche Bundesbank.

Goldstein, I. and Leitner, Y. 2013. *Stress tests and information disclosure.* FRB of Philadelphia Working Paper.

Gopalan, P., Nisan, N., and Roughgarden, T. 2015. Public projects, boolean functions, and the borders of border's theorem. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. EC '15. ACM, New York, NY, USA, 395–395.

Guo, M. and Deligkas, A. 2013. Revenue maximization via hiding item attributes. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. AAAI Press, 157–163.

Kamenica, E. and Gentzkow, M. 2011. Bayesian persuasion. *American Economic Review 101,* 6, 2590–2615.

Kolotilin, A. 2015. Experimental design to persuade. *Games and Economic Behavior 90*, 215–226.

Kolotilin, A., Li, M., Mylovanov, T., and Zapechelnyuk, A. 2015. Persuasion of a privately informed receiver. Tech. rep., Working paper.

Kremer, I., Mansour, Y., and Perry, M. 2014. Implementing the "wisdom of the crowd". *Journal of Political Economy 122,* 5, 988–1012.

Mansour, Y., Slivkins, A., and Syrgkanis, V. 2015. Bayesian incentive-compatible bandit exploration. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. ACM, 565–582.

McCloskey, D. and Klamer, A. 1995. One quarter of gdp is persuasion. *The American Economic Review 85,* 2, pp. 191–195.

MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research 6,* 1, 58–73.

NISAN, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. V. 2007. *Algorithmic Game Theory.* Cambridge University Press, New York, NY, USA.

PAPADIMITRIOU, C. H. AND ROUGHGARDEN, T. 2008. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM) 55,* 3, 14.

RABINOVICH, Z., JIANG, A. X., JAIN, M., AND XU, H. 2015. Information disclosure as a means to security. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS),*.

RAYO, L. AND SEGAL, I. 2010. Optimal information disclosure. *Journal of Political Economy 118,* 5, pp. 949–987.

RUBINSTEIN, A. 2015. Eth-hardness for signaling in symmetric zero-sum games. *arXiv preprint arXiv:1510.04991*.

SOBEL, J. 2010. Giving and receiving advice. In *Econometric Society 10th World Congress.*

TANEVA, I. 2015. Information design. *Manuscript*.

XU, H., FREEMAN, R., CONITZER, V., DUGHMI, S., AND TAMBE, M. 2016. Signaling in bayesian stackelberg games. In *AAAI Conference on Artificial Intelligence (AAAI).*

XU, H., RABINOVICH, Z., DUGHMI, S., AND TAMBE, M. 2015. Exploring information asymmetry in two-stage security games. In *AAAI Conference on Artificial Intelligence (AAAI).*

A Report on the Workshop on the Economics of Cloud Computing

NIKHIL R. DEVANUR
Microsoft Research

This is a report on the first Workshop on the Economics of Cloud Computing, which was held in conjunction with the ACM conference on Economics and Computation (EC).

1. INTRODUCTION

The digitization of the world's businesses, and the movement of this digitization into the cloud is akin to the industrial revolution. It is speculated that cloud computing will be to businesses what mobile computing has been to consumers. This raises a whole slew of questions in economics, most of which are deeply entangled with computer science topics. A half-day workshop on the economic aspets of cloud computing was held in conjunction with the ACM conference on Economics and Computation (EC) 2016 in Maastricht. The goal of the workshop was to be the premier platform to raise the most important research questions, to announce the latest results, to exchange ideas, to learn and to get feedback on the state-of-the-art research in this area. The topics of interest for this workshop were broadly set out to be as follows.

Moving to the Cloud. How are current businesses impacted by moving to a cloud enabled world?

New Markets. What new markets emerge as a result of a cloud enabled world? What new economic models come into play?

Cloud Pricing. What are the different pricing or auction mechanisms to sell cloud computing resources, and the pros and cons of each?

Cloud provisioning. What are best practices in the process of provisioning all the requirements for building a datacenter? What economies of scale can be exploited in running large data centers?

Fair allocation. How should one allocate cloud resources in a fair manner in a shared multi-tenant system?

There were 2 keynote speakers, Noam Nisan and Simon Wilkie, and 6 contributed talks. A call for papers for the contributed talks was circulated, and the following program committee decided which papers are to be presented.

—Nikhil R. Devanur (Program Chair)
—Eric Friedman
—Preston McAfee
—Noam Nisan

Author's address : `nikdev@microsoft.com`

—Eva Tardos

—Adam Wierman

All the details about the workshop can be found at `http://wecc.azurewebsites.net/`. We now present short descriptions of the keynote and the contributed talks.

2. TALKS

2.1 Keynotes

Noam Nisan:ERA: A Framework for Economic Resource Allocation for the Cloud. Noam Nisan from the Hebrew University at Jerusalem opened the talk with his definition of the cloud as a shared computational resource, which is typically a virtual machine at a remote data center. A well designed cloud system should make the most efficient use of the shared resources. For instance, flexible jobs should run during low congestion times, and the most "valuable" jobs must be run during periods of over demand. Simple schemes such as pay-as-you-go and dedicated hosts have obvious inefficiencies. Overcoming these shortcomings requires skills from various disciplines such as computer systems, algorithms, and economics.

The Economic Resource Allocation (ERA) project is a prototype system that is meant to expose cloud design issues at the intersection of all these three disciplines. It is a system for scheduling, reserving and pricing cloud resources. It provides friendly APIs for two interfaces: one user-facing that accepts requests for reservations, which are either accepted at a given price, or rejected; the other interfaces with an existing cloud system and provides it with jobs to run at any point of time. In between these two interfaces sits the ERA algorithm. The system allows plugging and playing with different algorithms, making it easy to compare and contrast them.

The prototype was used to provide a proof of concept for the benefits of combining insights from systems, algorithms, and economics. It showed how a simple economically aware algorithm can significantly improve the efficiency of the system. It provides a unified simulator and a platform over various cloud systems on which to test algorithms; this is a useful tool for future research. The main algorithmic challenges are in predicting future job requests from data, and in making optimal scheduling decisions from these predictions.

Simon Wilkie: The Price of Privacy in the Cloud, or The Economic Consequences of Mr. Snowden. Simon Wilkie from Microsoft Research spoke about estimating the effect of the Snowden revelations on cloud adoption for US based cloud providers. The adoption of the cloud among businesses has been on an upward trajectory ever since the inception of large scale cloud providers in the late 2000s. And then, Snowden's revelations about NSA's spying program in 2013 made consumers who cared about the privacy of their data wary of adopting US cloud providers. What was the effect of this? How much revenue was lost?

Simon and his co-author (Hyojin Song, Microsoft Research) find answers to these intriguing questions by using a global panel dataset of cloud revenues. They build a behavioral model for cloud adoption from the data, and use the non-US based providers as the "control". This then lets them estimate the magnitude of the negative demand shock on US providers due to Snowden. They estimated that the growth rate of US providers decreased by about 11%; this equalled about 18 billion USD in lost revenues. The US providers reacted to this decreased demand by reducing prices, which led to a "price war". An interesting side effect of this price war was that the market share of US providers eventually went up.

2.2 Contributed Talks

Cloud Pricing: The Spot Market Strikes Back. The decision on which model to use for selling cloud resources is a very real and important one for the providers. [Dierks and Seuken 2016] consider whether offering both a fixed price and a dynamic (spot) price can increase profits over offering either of these alone. Previous work by [Abhishek et al. 2012] showed that the answer is that it doesn't, but based on an assumption that the provider has access to an infinite pool of resources. This paper considers the cost of procuring resources, and shows that a hybrid model can indeed be better for profit. The demand is modeled as a stochastic process similar to the queueing theory models. The system is assumed to be at an equilibrium where the supply (the number of servers provisioned) is equal to the demand. (The expected waiting time of a job is below some threshold.) In the current model, the idle instances of the fixed price market cannot be sold on the spot market. Utilizing such idle instances is one of the main attractions of the spot market, and incorporating this into the model seems an important step.

On-Demand or Spot? Selling the Cloud to Risk-Averse Customers. [Hoy et al. 2016] consider essentially the same quesion as the previous talk, but focus on a risk averse model of a consumer. A concave curve determines the utility of a consumer as a function of her surplus. A dual market with both fixed and spot prices works as follows: bidders first decide whether to reserve an instance using the fixed price market. The available supply is then sampled from a given distribution. Any excess supply that remains after allocating all the reserved instances is then sold through an auction resulting in a spot price. They show how this model explains the existence of such a dual market by showing increased revenue/welfare/efficiency compared to markets with a single option.

An alternate direction to tackle this issue of two markets versus one is to consider time sensitivity of consumers. Consumers whose jobs are time sensitive tend to opt for a reservation market while others would prefer the lower prices in a spot market. Another interesting question is how the conclusion is affected by the presence of competitors who sell imperfect substitutes.

Approximately Efficient Cost Sharing via Double Auctions. [Fischer et al. 2016] propose jointly solving the problems of pricing and procurement for cloud resources. This makes sense since any reasonable objective (welfare/revenue) depends on both the demand as well as the cost. Solving each of them separately (assuming the other as fixed), as is done currently, can be sub optimal. They model this as a cost sharing problem and consider a twist to the standard guarantees by allowing an additive as well as a multiplicative term for approximating efficiency. They present a mechanism that is inspired by the double auction of [McAfee 1992]. In a large market setting, this mechanism attains strategyproofness, budget balance and approximate (additive + multiplicative) efficiency, thus bypassing previous impossibility reults. The result is interesting more generally for other cost sharing problems as well.

Pretium: Dynamic Pricing and Traffic Engineering for Timely Inter-Datacenter Transfers. [Jalaparti et al. 2016] consider pricing schemes for data transfers between data centers. The current standard practice for such transfers is to have a fixed price per unit of data, but this is inefficient due to the large temporal variation in the requests. Moreover, it is shown via a survey that customers are quite receptive to the idea of time-of-use pricing and trading off the timing of their transfers for the cost of transfers. The paper shows that a dynamic pricing combined with traffic engineering can significantly increase the efficiency of these systems. The methodology is empirical: they uses traces of data transfer from a large data center and replay them under the different schemes. The prices are calculated using the requests from a reference time window from the past; the optimal dual variables of a welfare maximizing linear program give market clearing prices. The process is also shown to limit the users from gaming the system by showing that certain types of gaming don't help (both theoretically and empirically). This paper is an excellent demonstration that simple economic insights can have significant impact on real systems.

Congestion Games with Mixed Objectives. In allocating a shared network bandwidth among many users, two different objectives have beens studied: latency, and bandwidth. In any large data center there are heterogeneous users among whom some care more about latency (video gaming) while others more about bandwidth (media streaming). (While latency is typically additive, bandwidth is a min or a max.) [Feldotto et al. 2016] consider congestion games where the users have different utility functions of these two cost measures. They show that when the agent preferences satisfy a certain monotonicity assumption, a pure Nash equilibrium always exists. Moreover, a lazy best response dynamics converges to it. In the absence of this assumption, pure Nash may not exist and finding one is NP-Hard. Best response dynamics may cycle. An interesting direction for future research would be to incorporate uncertainties, in preferences as well as realized costs.

Dynamic Games for Market Dominance in the Cloud. Traditionally, providing an online service involved a big fixed cost for setting up the IT infrastructure and almost zero marginal cost thereafter. Cloud computing flipped this cost structure so that it is almost all marginal cost and no fixed cost. Standard competitive models in economics suggest that this lowers the cost of entry and hence lead to the market being shared by many competitors. Often, one sees the market in the technology sector being dominated by one or two companies, contrary to this prediction. [Conley 2016] seeks to explain this discrepancy by considering the presence of venture capital funding. The strategies for a venture capital funded firm and a publicly traded firm differ in that the venture capital funded firm can be way more aggressive in spending money on advertisements and promotions than a publicly traded firm. This is because once a venture capitalist commits to a funding, it becomes rational for the company to spend as much of it as possible to gain market share, whereas a similar strategy would be irrational for a publicly traded firm.

The paper introduces an interesting model of competition among firms, and there is an opportunity to extend it further to capture other aspects of the real markets. For instance, while there are a large number of firms entering the market, there are fewer firms competing for consumers ex post because of either network effects resulting in a winner-take-all situation, or because their exit strategy is to get acquired by a bigger firm.

3. CONCLUSION

With cloud computing fast becoming the de facto way for businesses to handle their IT infrastructure, we expect research into the economic aspects of cloud computing to grow. Given sufficient interest from the community, the workshop could become an annual or a biennial fixture at EC, alongside similar workshops such as the one on ad auctions.

REFERENCES

Abhishek, V., Kash, I. A., and Key, P. 2012. Fixed and market pricing for cloud services. In *Computer Communications Workshops (INFOCOM WKSHPS), 2012 IEEE Conference on*. IEEE, 157–162.

Conley, J. 2016. Dynamic games for market dominance in the cloud.

Dierks, L. and Seuken, S. 2016. Cloud pricing: The spot market strikes back.

Feldotto, M., Leder, L., and Skopalik, A. 2016. Congestion games with mixed objectives. In *International Conference on Combinatorial Optimization and Applications*. Springer, 655–669.

Fischer, F., Kash, I. A., Key, P., and Wang, J. 2016. Approximately efficient cost sharing via double auctions.

Hoy, D., Immorlica, N., and Lucier, B. 2016. On-demand or spot? selling the cloud to risk-averse customers. In *International Conference on Web and Internet Economics*. Springer, 73–86.

Jalaparti, V., Bliznets, I., Kandula, S., Lucier, B., and Menache, I. 2016. Dynamic pricing and traffic engineering for timely inter-datacenter transfers. In *Proceedings of the 2016 Conference on ACM SIGCOMM 2016 Conference*. SIGCOMM '16. ACM, New York, NY, USA, 73–86.

McAfee, R. P. 1992. A dominant strategy double auction. *Journal of economic Theory 56,* 2, 434–450.

Bounded and Envy-free Cake Cutting

HARIS AZIZ, DATA61 and UNSW
and
SIMON MACKENZIE, Carnegie Mellon University

Although a finite envy-free cake-cutting protocol has been known for more than twenty years, it had been open whether a protocol exists in which the number of steps taken by the protocol is bounded by a function of the number of agents. In this letter, we report on our recent results on discrete, bounded, and envy-free cake-cutting protocols.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Theory, Algorithms, Economics

Additional Key Words and Phrases: Fair allocation, Solutions Concepts, Multiagent resource allocation, Computational Complexity

1. INTRODUCTION

The cake cutting problem is a fascinating and fundamental mathematical problem in which the cake is a heterogeneous divisible resource represented by the unit interval [Brams and Taylor, 1996, Robertson and Webb, 1998]. Each of the n agents have additive and non-negative valuations over segments of the interval. The challenge is to query agents about their valuations in an efficient way to find a fair allocation. Originally formalized by Polish mathematician Hugo Steinhaus in the 1940's, the problem has attracted considerable attention in mathematics, computer science and economics. One of the most important criteria for fairness is *envy-freeness*. An allocation is envy-free if no agent would prefer replacing her allocation with another agent's. A cake cutting protocol is termed envy-free if each agent is guaranteed to be non-envious if she reports her real valuations. If a protocol is envy-free, then an honest agent will not be envious even if other agents misreport their valuations.

The most famous envy-free cake cutting protocol is *Divide and Choose* for two agents: one agent is asked to divide the cake into two equally preferred pieces. The other agent is then asked to pick her most preferred piece whereas the cutter gets the remaining piece. In the 1960's, John Selfridge and John Conway independently proposed an envy-free protocol for three agents that requires at most five cuts on the cake. In the early 1990's, Steven Brams and Alan Taylor invented a general finite envy-free cake cutting protocol [Brams and Taylor, 1995]. The protocol can require arbitrary number of steps and cuts on the cake even for four agents. It has been an open problem whether a four-agent bounded envy-free protocol exists or not [Brams and Taylor, 1995, Procaccia, 2013, 2016].

This year we presented a four agent envy-free protocol that requires 203 cuts on

Authors' addresses: haris.aziz@data61.csiro.au, simonm@andrew.cmu.edu

the cake and a total of 584 queries [Aziz and Mackenzie, 2016a].[1] We have now generalized the protocol to any number of agents [Aziz and Mackenzie, 2016b,c]. In this letter, we give an overview of some of the ideas and building blocks of the general protocol.

2. A BIRD'S EYE VIEW OF THE PROTOCOL

In a nutshell, the protocol allocates a large enough portion of the cake in an envy-free manner. After that, it tries to add some small portions of the unallocated cake to the allocated part in a structured and envy-free manner with the goal to reduce the problem to envy-free allocation for a smaller number of agents.

A crucial building block of our protocol is the Core Protocol. The Core Protocol asks one of the n agents—the cutter—to divide the cake into n equally preferred pieces. It then uses the recursive SubCore Protocol to obtain a *neat* allocation for the other agents. In a neat allocation, each agent gets a part of exactly one of the pieces, one agent gets a full piece, and the agents are not envious of each other or of the unallocated pieces. The cutter then gets one of the unallocated and untrimmed pieces.

The SubCore is a general protocol that takes as input agents and pieces of cake where the number of agents is at most the number of pieces. We order the agents and find a neat envy-free allocation for an expanding set of agents. Say that we have found a neat allocation for $m-1$ agents. In that case, $m-2$ pieces could be partially allocated but the other pieces are untrimmed. If the m-th agent thinks that one of the unallocated full pieces is her most preferred, then she is simply given such a piece. Otherwise, we have a situation where m agents are interested in $m-1$ '*contested*' pieces. In such a situation one of the m agents has to be given a most preferred uncontested piece. In order to find such an agent as well as reallocate the $m-1$ contested pieces, we have to do more work and may have to call SubCore recursively for a smaller number of agents.

In the Core Protocol, the cutter agent gets a full piece. Another agent also gets a full piece. So from the cutter's perspective at least $2/n$ of the cake is allocated by one call of the Core Protocol. If we call the Core Protocol with a different cutter each time to further allocate the unallocated cake, we just need n calls of the Core Protocol to obtain an envy-free allocation in which each agent thinks she gets $1/n$ value of the whole cake. This answers an open problem posed by Segal-Halevi et al. [2015] which asks whether there exists a bounded algorithm that returns envy-free partial allocation that is proportional (gives each agent value at least $1/n$ of the whole cake).

Continuing to call the Core Protocol on the updated remaining cake gives no guarantee that the cake will be allocated fully even in finite time. Hence, we need to use other protocols. Throughout the overall protocol, we maintain an envy-free allocation as well as keep track of the updated unallocated cake.

Since the Core Protocol by itself is not powerful enough to allocate all the cake in bounded time, we rely on the idea of *domination.* An agent i *dominates* another agent j if she is not envious of j even if the unallocated cake is given to j. The

[1] Walter Stromquist observed that both the number of cuts and queries can be halved by a simple adjustment.

other protocols are used with the following objective in mind: find a set of agents $N \setminus A \subset N$ such that each agent in the set dominates each agent in $A \subset N$. In order to ensure that each agent in some set $N \setminus A$ dominates each agent in A requires changing the current allocations of the agents as well as the unallocated cake. While we make changes to the allocation, we ensure that the current partial allocation remains envy-free. By identifying such a set $N \setminus A$, we reduce the problem to envy-free allocation for a smaller number of agents. The agents in $N \setminus A$ are not envious whatever the unallocated cake is allocated among agents in N.

3. CONCLUSIONS AND OPEN PROBLEMS

In this letter, we provided a very high-level overview of our bounded envy-free protocol. The protocol has an upper bound that is a power tower of six n's. In the other direction, any envy-free protocol requires at least $\Omega(n^2)$ queries [Procaccia, 2016].[2] There is a lot of work to be done to close the gap between the current upper and lower bound.

We additionally show that even if we do not run our protocol to completion, it can find in at most n calls of the Core Protocol a partial allocation of the cake that achieves proportionality (each agent gets at least $1/n$ of the value of the whole cake) and envy-freeness. It also adds further evidence to the idea popularized by Segal-Halevi et al. [2015] that wasting some resource can lead to much faster fair division algorithms. If we allow for partial allocations, an interesting open problem is the following one: can envy-freeness and proportionality can be achieved in polynomial number of steps? Finally we mention that it is still open whether a bounded and envy-free cake cutting protocol exists for the case of where agents have negative valuations.

Acknowledgments

We thank Hu Fu, editor SIGEcom Exchanges, for inviting us to write this letter.

REFERENCES

H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 454–464. ACM Press, 2016a.

H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. Technical Report arXiv:1604.03655, arXiv.org, 2016b.

H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS)*, 2016c.

S. J. Brams and A. D. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995.

S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution.* Cambridge University Press, 1996.

A. D. Procaccia. Cake cutting: Not just child's play. *Communications of the ACM*, 56 (7):78–87, 2013.

[2] In the second author's PhD thesis, this lower bound has been improved to $\Omega(n^3)$.

A. D. Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2016.

J. M. Robertson and W. A. Webb. *Cake Cutting Algorithms: Be Fair If You Can.* A. K. Peters, 1998.

E. Segal-Halevi, A. Hassidim, and Y. Aumann. Waste makes haste: Bounded time protocols for envy-free cake cutting with free disposal. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 901–908. IFAAMAS, 2015.

Observing Algorithmic Marketplaces In-the-Wild

LE CHEN
and
CHRISTO WILSON
Northeastern University

In this letter, we briefly summarize two recent works from our group that use observational data to study the mechanisms used by two large markets. First, we examine Uber's surge price algorithm, and observe that its incentive model may not be effective at changing driver behavior. Second, we study the adoption of dynamic pricing strategies by sellers on Amazon Marketplace, and investigate how these strategies interact with Amazon's "Buy Box" matching algorithm. We make our data available to the research community.

Categories and Subject Descriptors: K.4.4 [**Computing Milieux**]: Computers and Society—*Electronic Commerce*; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Observational Study, Market Design, Pricing

Additional Key Words and Phrases: Empirical, Amazon, Uber, Ridesharing, Dynamic Pricing

1. INTRODUCTION

Much of the classic literature in economics deals with mechanism design, i.e., the construction of markets that maximize some useful quantity like revenue or welfare. As commerce has moved online, it has become easier to directly apply these ideas from economic theory in practice. One obvious example of this are online advertising auctions, but more broadly, many companies are now experimenting with differential [Mikians et al. 2012; 2013; Hannak et al. 2014] and dynamic pricing [Chen 2016] strategies in contexts ranging from retail to ridesharing.

Although academics are beginning to propose models for modern e-commerce platforms [Banerjee et al. 2015; Fang et al. 2016], we lack a comprehensive empirical understanding of the actual mechanisms adopted by companies in their marketplaces. The opacity surrounding widely used platforms raises fundamental questions for researchers and consumers: what objectives are these systems optimized for, and are they achieving these objectives? What features do they consider? Are the markets fair, and if so, for what definition of fairness?

In our recent work, we attempt to answer these questions through empirical measurements of major online markets. Using observed data, we quantify the basic properties of markets over time, such as number of participants and prices. In some cases, we must develop novel data gathering methodologies to acquire this information. We then leverage this raw data to infer implementation details of the markets themselves, e.g., the weights of key features, or how algorithms discretizes prices across time and physical space. Finally, we examine the implications of deployed mechanisms on market participants.

Authors' addresses: `leonchen@ccs.neu.edu`, `cbw@ccs.neu.edu`

In this letter, we briefly overview two of our recent observational studies of major online marketplaces:

(1) In [Chen et al. 2015], we examined Uber's dynamic pricing algorithm, which is known as "surge pricing". Using 43 emulated copies of the Uber smartphone app, we blanketed midtown Manhattan and downtown San Francisco for about one month in order to collect data about available supply of rides, fulfilled demand, prices, and estimated wait times. This data enabled us to quantify the sensitivity of surge prices to fluctuations in supply and demand, as well as analyze the impact of surges on rider and driver behavior.

(2) In [Chen et al. 2016], we study two types of algorithms on Amazon Marketplace. First, we investigate Amazon's "Buy Box" algorithm, which determines the default seller that will fulfill orders for each product in the market. It is estimated that ∼80% of purchases on Amazon go through the Buy Box [Taft 2014], so understanding this algorithm is key to being competitive on Amazon Marketplace. Second, we examine the dynamic pricing strategies adopted by individual sellers. Although we find that only a small fraction of sellers have adopted dynamic pricing (and that their strategies are relatively unsophisticated), we also observe that *algo sellers* have a significant competitive advantage versus non-algo sellers, especially with respect to winning the Buy Box.

Overall, we view our empirical work as being complementary to theory. Our data can be used to refine existing models, bound their parameters, or evaluate their behavior under realistic conditions. More broadly, our observations about the strategies adopted by businesses in practice can potentially motivate the design of new models. We make the data and code from many of our studies publicly available (additional data is available by request) at: `http://personalization.ccs.neu.edu`.

2. PEEKING BENEATH THE HOOD OF UBER

In this work [Chen et al. 2015], we examined Uber's surge pricing system, which aims to balance the demand for rides with the available supply by varying price dynamically. In the literature these ridesharing systems are conceptualized as traditional two-sided platforms [Banerjee et al. 2015; Fang et al. 2016] serving passengers and drivers. However, in terms of *implementation*, these systems are untraditional: rather than having an open marketplace and allowing the two parties to converge towards equilibrium, ridesharing companies have adopted centralized algorithms that attempt to balance supply and demand. The closed nature of these ridesharing platforms raises questions about whether the dynamic pricing mechanisms are efficient and fair.

2.1 Methods and Data Collection

To collect data for our study, we emulated 43 copies of the Uber smartphone application. By default, Uber's app requests fresh data from Uber's servers every 5 seconds, including: 1) the eight closest available cars to the user (based on GPS coordinates), 2) the current surge multipliers, and 3) the Estimated Wait Times (EWTs) for cars. By carefully spoofing the GPS coordinates for our 43 emulated

Fig. 1: Uber measurement points in downtown San Francisco.

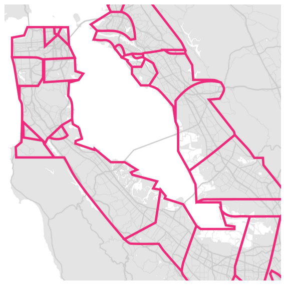

Fig. 2: Uber surge area map for the Bay Area.

users, we were able to place them in a grid throughout a target city, thus enabling us to passively collect data that covered the entire area. Figure 1 shows the measurement grid we used to collect data from downtown San Francisco. Furthermore, we were able to observe when each car became unavailable, which implies that either the driver had logged-off from Uber, or that they accepted a ride request. This enabled us to place an upper-bound on *fulfilled* demand.

Using our methodology, we collected data from midtown Manhattan between April 3–17, 2015, and from downtown San Francisco between April 18–May 2, 2015. Additionally, before collecting data at scale, we performed trials to make sure that our emulated users did not *induce* surge pricing. In these tests, we placed all 43 emulated users at a single, remote GPS coordinate late at night, and did not observe any surge pricing for one hour. We repeated this trial many times at many locations, and never observed surge prices.

2.2 Surge Pricing Algorithm

Based on our observed data, we are able to draw several conclusions about Uber's surge pricing mechanism. The system divides each city into areas (that we suspect are statically defined by human operators), and updates the surge multiplier for each area at five minute intervals. Figure 2 shows the surge areas for a subset of the San Francisco Bay Area.

If we treat surge prices as a time series and calculate the cross-correlation versus other variables, we observe statistically significant correlations between surge prices, available supply, and fulfilled demand at a time delta of -5 minutes. This suggests that Uber's algorithm calculates the surge price $s_t(a)$ at time t in area a using the supply and demand from area a measured over the previous five minute interval. While this demonstrates that Uber's dynamic pricing algorithm is highly responsive, it also means that surge prices are quite noisy (60% of surges last $\leq$10 minutes in our data).

2.3 Incentives

One of our most interesting findings concerns the incentives of the surge pricing system. Recall that Uber's goal is to equalize supply and demand by increasing the former (by incentivizing drivers with high prices) and decreasing the latter. To understand if drivers and customers are responding to these incentives in the expected way, we model the behavior of each Uber driver as a discreet-time Markov

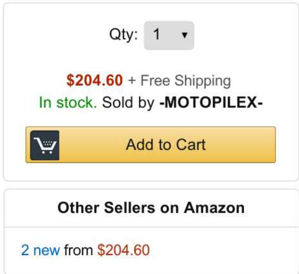

Fig. 3: Example Buy Box from Amazon Marketplace.

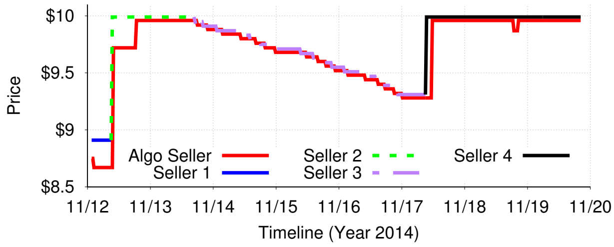

Fig. 4: Observed algorithmic seller that consistently underprices other sellers by several cents.

chain, where the state of a driver at time $t + 5$ minutes is determined by its state at t. In this model, "state" encompasses the physical location of the car and the surge prices in all areas containing and surrounding the driver. Possible state transitions include remaining stationary, accepting a ride request, or driving into an adjacent area. Using this model, we calculate the probability of state transitions at times when all areas have equal surge prices and when exactly one area of the city has surge multiplier $\geq$0.2 higher than all other areas (i.e., drivers have a strong monetary incentive to travel to the surging area).

Our results paint a complicated picture of Uber's dynamic pricing system. As expected, we find that high surge multipliers reduce demand. Bookings decrease by 7% on average in the surging area (compared to when it is not surging), while drivers that do not get booked increases by 14%. However, we also find that drivers are 13% less likely to drive into the area that is surging (compared to times when all surge multipliers are equal); in fact, the number of drivers who *leave* the surging area increases by 14%! These results suggest that the surge mechanism is not effective at incentivizing drivers. Our findings echo qualitative findings by Lee et al. who interviewed Uber drivers, and found that veteran drivers find it futile to "chase the surge" [Lee et al. 2015]. Our results also stand in contrast to Uber's own with respect to the benefits of the surge system [Hall and Krueger 2015].

3. ALGORITHMIC PRICING ON AMAZON MARKETPLACE

In this study [Chen et al. 2016], we examine two separate types of algorithms on Amazon Marketplace. *First*, we examine Amazon's Buy Box algorithm. This algorithm determines, for each product in the Marketplace, which seller's offer price will be shown to customers (and consequently, which seller makes the sale when the product is purchased). Figure 3 shows an example of the Buy Box for a product. In essence, the Buy Box algorithm functions as a matching mechanism between buyers and sellers in the Marketplace, and therefor it must balance the interests of customers (low prices, good service), third party sellers (revenue), and Amazon itself (revenue, overall health of the platform). However, despite the critical importance of the Buy Box algorithm, little is known about it beyond online folk wisdom.

Second, we investigate dynamic pricing strategies adopted by sellers on Amazon Marketplace. Amazon offers APIs that allow sellers to track competitors' prices in real-time and respond with their own price changes. Theoretically, sellers may also adopt strategies that increase their chances of winning the Buy Box. Although subscription-based tools like Feedvisor and Sellery have made dynamic pricing tools widely available to third party sellers on Amazon, it is unclear how widely dynamic pricing has been adopted, or what strategies are used by sellers.

3.1 Data Collection

To bootstrap our study, we crawled roughly four months of data from Amazon Marketplace. We chose 1000 best-selling products that were each offered by >1 seller, and crawled them every 25 minutes. Each time the crawler visited a product, it recorded the seller and price in the Buy Box, as well as up-to two additional pages of sellers offering that product (each page contains up-to 10 sellers, sorted roughly from low-to-high price). We chose to wait 25 minutes between crawls and limit the number of seller pages visited as a tradeoff between recency and completeness: Amazon implements strict rate limits, so more frequent visits (or more pages per visit) would have forced us to crawl fewer products overall. Furthermore, we could not use Amazon's APIs to collect data since the only way to get price updates for products is to list them for sale. We crawled data in two phases, between September 15–December 8, 2014, and between August 11–September 21, 2015.

3.2 The Buy Box

To investigate the features behind the Buy Box algorithm, we trained a Random Forest (RF) classifier to predict Buy Box winners (given a list of offers for a specific product). The intuition behind this process is that if we can train an accurate predictor, then it is likely that the feature weights in our model correspond closely to those used by the actual Buy Box algorithm. We input seven features into RF classifier, including each sellers': offer price relative to the lowest offer for the given product, average customer rating, positive feedback percentage, and enrollment in the Fulfilled By Amazon (FBA) program.

After performing 10-fold cross validation, we found that our RF classifier was able to predict Buy Box winners with 75–85% accuracy (depending on the total number of offers for a given product). In contrast, a simple predictor that always chooses the seller with the lowest offer only achieves 50-60% accuracy. This demonstrates that our RF classifier does a reasonable job of approximating the Buy Box algorithm, and that price alone is not the sole feature used by the true algorithm. Indeed, by examining the Gini coefficients associated with features in our RF model, we find that while the price feature has the highest weight, customer feedback also has significant weight.

3.3 Dynamic Pricing

To identify sellers on Amazon that use dynamic pricing, we look for sellers whose offer prices have high correlation over time with a specific *target* price time series. Example targets include the lowest overall price for a given product, or Amazon's offer price for that product. Intuitively, this methodology attempts to identify sellers whose offer price is pegged to an observable benchmark over time.

Using this methodology, we identify XXX sellers who we are confident have adopted dynamic pricing. Although this only represents 2.4% of the sellers in our dataset, their listings cover 51% of the products we crawled. 70% these *algo* *sellers* choose to peg their offer price within $1 above the lowest available price for each product. However, despite setting higher prices than competitors, we observe that algo sellers much more likely to win the Buy Box than non-algo sellers. This clearly demonstrates that sellers who adopt automation are at a competitive advantage versus sellers who do not. Finally, we note that algo sellers are responsible for the vast majority of price and Buy Box changes on Amazon Marketplace; we even observe a small number of products with thousands of price changes over the course of a month.

4. CONCLUSION

As commerce moves online, the opportunities to construct fluid, dynamic marketplaces increase. In some cases, like online display advertising, the structure and mechanisms in these new markets are relatively well understood. However, in other cases, like ridesharing and e-commerce, the algorithms being adopted by industry are opaque.

The overarching goal of our work is to increase the transparency of algorithms in online markets. This can help consumers and producers make more informed choices about how to best optimize their behavior on these platforms. We also hope that our work is beneficial to the theory community, as a starting point for evaluating existing models, or even motivating new designs.

REFERENCES

Banerjee, S., Johari, R., and Riquelme, C. 2015. Pricing in Ride-Sharing Platforms: A Queueing-Theoretic Approach. In *Proc. of EC*.

Chen, L., Mislove, A., and Wilson, C. 2015. Peeking Beneath the Hood of Uber. In *Proc. of IMC*.

Chen, L., Mislove, A., and Wilson, C. 2016. An Empirical Analysis of Algorithmic Pricing on Amazon Marketplace. In *Proc. of WWW*.

Chen, M. K. 2016. Dynamic pricing in a labor market: Surge pricing and flexible work on the uber platform. In *Proc. of EC*.

Fang, Z., Huang, L., and Wierman, A. 2016. Boosting sharing economy: Social welfare or revenue driven? *CoRR abs/1604.01627*.

Hall, J. V. and Krueger, A. B. 2015. An analysis of the labor market for uber's driver-partners in the united states.

Hannak, A., Soeller, G., Lazer, D., Mislove, A., and Wilson, C. 2014. Measuring Price Discrimination and Steering on E-commerce Web Sites. In *Proc. of IMC*.

Lee, M. K., Kusbit, D., Metsky, E., and Dabbish, L. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers.

Mikians, J., Gyarmati, L., Erramilli, V., and Laoutaris, N. 2012. Detecting price and search discrimination on the internet. In *Proc. of HotNets*.

Mikians, J., Gyarmati, L., Erramilli, V., and Laoutaris, N. 2013. Crowd-assisted search for price discrimination in e-commerce: First results. In *Proc. of ACM CoNEXT*.

Taft, D. K. 2014. Amazon buy box: The internet's $80 billion sales button. eWeek. `http://www.eweek.com/enterprise-apps/slideshows/amazon-buy-box-the-internets-80-billion-sales-button.html`.

Ellipsoids for Contextual Dynamic Pricing

MAXIME C. COHEN
Stern School of Business, New York University
and
ILAN LOBEL
Stern School of Business, New York University
and
RENATO PAES LEME
Google Research NYC

We study the dynamic pricing problem faced by a firm selling differentiated products. At each period, the firm receives a new product, which is described by a vector of features. The firm needs to choose prices, but it does not know a priori the market value of the different features. We first consider an algorithm that we call PolytopePricing, but prove that it incurs worst-case regret that scales exponentially in the dimensionality of the feature space. We then consider a closely related algorithm, EllipsoidPricing, and show it incurs low regret with regards to both the time horizon and the dimensionality of the feature space. For more details, we refer the reader to our full paper.

1. INTRODUCTION

In this letter, we briefly survey the results of our recent paper, [Cohen et al. 2016]. We also discuss recent developments in the literature surrounding our paper.

We consider an online market where a firm sells highly differentiated products to its buyers. In each period, a new product arrives and the selling firm must set a price for it. Each product is characterized by a vector of features (or contexts) that determine its market value. The firm knows the features of the products, but it does not know a priori the value of the different features. Our paper aims to understand what is a good pricing policy for balancing learning and earning in such a contextual setting.

Our problem is motivated by online markets such as Airbnb and the market for online ads. In these markets, every product is unique in its attributes. In the case of Airbnb, a product is a stay in a particular listing on a specific date. The features therefore represent both characteristics of the listing, such as location and amenities, as well as the check-in/check-out dates. In the market for online ads, a product is an impression which is sold to potential advertisers. The features that determine the market value of an impression include the IP address and the relevant cookie data, which might contain information such as gender, age and browsing history.

Authors' addresses: maxime.cohen@stern.nyu.edu, ilobel@stern.nyu.edu, renatoppl@google.com

2. THE MODEL

Consider a seller that receives products in an online fashion. At every period $t = 1, ..., T$, a new product arrives with a set of features $x_t \in \mathcal{X} \subset \mathbb{R}^d$ such that $||x_t||_2 \leq 1$. The market value of each feature is given by $v_t = \theta' x_t$, where θ is a d-dimensional vector unknown to the seller. The seller chooses a price p_t as a function of x_t. If the seller chooses a price below or equal to v_t, a sale occurs and she earns p_t. If the seller chooses a price p_t above v_t, no transaction occurs.

The seller knows initially only that θ belongs to a bounded convex set K_1, where $||\theta||_2 \leq 1$. If a sale occurs at time t, the seller updates her uncertainty set according to $K_{t+1} = K_t \cap \{\theta \in \mathbb{R}^d : \ \theta' x_t \geq p_t\}$. Similarly, if a sale does not occur at time t, the seller updates her uncertainty set following $K_{t+1} = K_t \cap \{\theta \in \mathbb{R}^d : \ \theta' x_t \leq p_t\}$. The seller's problem is to choose a pricing policy that minimizes her worst-case regret, which is given by:

$$\text{REGRET} = \max_{\theta \in K_1, \ \{x_t\}} \ \sum_{t=1}^{T} \Big[\theta' x_t - p_t I\{\theta' x_t \geq p_t\}\Big]. \qquad (1)$$

Our model assumes that both θ and the feature vectors $\{x_t\}$ are chosen adversarially by nature. Good policies under this metric would therefore be robust to changes or seasonal fluctuation in the pattern of arrival of the feature vectors over time.

The model as described above assumes a linear deterministic relationship between the features and the market values. We also briefly discuss below how to address more general market value functions, including both noisy valuation models and some commonly used nonlinear models.

3. CONTEXTUAL PRICING ALGORITHMS

If the seller knew the value of θ, she could maximize her revenue by simply choosing $p_t = \theta' x_t$ at each period. However, since the seller does not know the value of θ, she must balance exploration and exploitation. Since the optimal decision at each period depends on a context vector x_t, our problem is a special case of a contextual bandit problem ([Auer 2003]). We could therefore use an off-the-shelf algorithm for contextual bandits, such as [Agarwal et al. 2014]. Such an algorithm would have suboptimal performance (polynomial rather than logarithmic regret in T) since it does not take advantage of the underlying linear structure of our pricing problem. Consequently, our aim is to construct an algorithm for this problem with good performance with respect to both the time horizon T and the dimensionality d. Our first attempt is a multidimensional version of binary search that we call PolytopePricing.

The PolytopePricing Algorithm. At each period t, we have access to the vector of features x_t and we know that $\theta \in K_t$. A natural first question is to ask whether we can predict the value $v_t = \theta' x_t$ with reasonable accuracy. The lowest and highest possible values of v_t are given by:

$$\underline{b}_t = \min_{\hat{\theta} \in K_t} \hat{\theta}' x_t \qquad \text{and} \qquad \bar{b}_t = \max_{\hat{\theta} \in K_t} \hat{\theta}' x_t.$$

For a given accuracy parameter $\epsilon > 0$, we can say that we know the value of v_t with ϵ-accuracy if and only if $\bar{b}_t - \underline{b}_t \leq \epsilon$.

This gives rise to a parameterized algorithm PolytopePricing(ϵ). If $\bar{b}_t - \underline{b}_t \leq \epsilon$, then we should choose an exploit price since we know the market value of product t with ϵ-accuracy. The natural choice for an explore is $p_t = \underline{b}_t$ since no price above this level guarantees a sale. If $\bar{b}_t - \underline{b}_t > \epsilon$, then we should use an explore price. A natural explore price is the one inspired by binary search: $p_t = (\bar{b}_t + \underline{b}_t)/2$.

Unfortunately, as shown in Theorem 1 below, this algorithm does not scale well with the dimensionality of the feature space. This algorithm becomes problematic in high dimensions because the explore price might remove too small a fraction of the uncertainty set.

Theorem 1. *For any parameter $\epsilon > 0$, the algorithm* PolytopePricing *suffers worst-case regret $\Omega(1.2^d)$.*

The EllispoidPricing Algorithm. A natural follow-up question is whether we can "fix" PolytopePricing by ensuring that we remove a sufficiently large volume of the uncertainty set per explore period. A way to accomplish this objective is to borrow ideas from the ellipsoid method from optimization theory ([Khachiyan 1979]).

The EllispoidPricing algorithm works exactly as the PolytopePricing, except it has an additional step. At the end of each period t, we replace our convex set K_t by its Löwner-John ellipsoid E_t. The Löwner-John ellipsoid of a convex set is the smallest ellipsoid that contains that set. It turns out that this simple modification to the algorithm is sufficient to ensure a good regret performance.

Theorem 2. *The worst-case regret of the* EllipsoidPricing *algorithm with parameter $\epsilon = d^2/T$ is $O(d^2 \ln(T/d))$.*

To prove Theorem 2, we combine two ideas. From the work of Khachiyan, we know that an algorithm that iteratively cuts an ellipsoid and then replaces the remaining half-ellipsoid by its own Löwner-John ellipsoid yields an exponentially fast volume reduction. This idea alone is not sufficient to prove our result since our theorem requires us to control the radii of the ellipsoid too, not merely the volume. Note that bounding the radii of the period t ellipsoid allows us to bound the difference $\bar{b}_t - \underline{b}_t$, and thus determines the accuracy of our knowledge of the value of v_t. This now brings us to the second part of the proof. Building on the linear algebra machinery of [Wilkinson 1965], we prove that the radii of our ellipsoids never become too small. This follows from the fact that our algorithm never uses an explore price in a direction that is already small; it uses an exploit price instead.

The EllipsoidPricing algorithm not only has a good worst-case regret guarantee, it is also computationally efficient. Both key operations — optimizing a linear function over an ellipsoid and finding the Löwner-John ellipsoid of a half-ellipsoid — require only matrix-vector multiplications.

Nonlinear and noisy market values. The model above assumes the market value is a deterministic linear function of the feature vector. Our algorithm can be extended to more general models. If the market value function can be expressed as $v_t = f(\theta' \phi(x_t))$, where f is a non-decreasing Lipschitz continuous function and $||\phi(\cdot)||_2 \leq 1$, then Theorem 2 still applies up to the Lipschitz constant.

If the market value is a noisy function of the feature vector, i.e., $v_t = \theta' x_t + \delta_t$ for a noise term δ_t, then we need to slightly modify the EllipsoidPricing algorithm before applying it. The key idea here is to add a safety margin when updating the uncertainty set. Instead of cutting the ellipsoid through its center, we can cut the ellipsoid in a way that leaves more than half of the original volume. This is called a shallow cut. In [Cohen et al. 2016], we show the details of how shallow cuts can be used to handle both bounded adversarial noise and i.i.d. Gaussian noise.

4. RECENT DEVELOPMENTS IN THE LITERATURE

There is a large body of literature on dynamically adjusting prices both with and without contextual information. We refer to [Cohen et al. 2016] for an extensive discussion. Here, we focus on papers that have appeared in the last few months and modify the model of [Cohen et al. 2016] in interesting new directions.

Lower bound. One natural question is whether our EllipsoidPricing algorithm has optimal regret. The answer is no. [Kleinberg and Leighton 2003] showed that in a one-dimensional (non-contextual) version of our problem, the optimal regret is $\Theta(\ln \ln T)$. The best known lower bound for our problem is thus $\Omega(d \ln \ln T)$.

In recent work, [Lobel et al. 2016] make progress towards closing this gap by showing that there exists an algorithm that incurs $O(d \ln(dT))$ worst-case regret. This algorithm also runs in polynomial time, but is far more complex and harder to implement than EllipsoidPricing. It requires computing approximate centroids of high-dimensional sets as well as projecting and cylindrifying polytopes.

Stochastic versions. Another direction is to consider a stochastic version of our model, as opposed to our adversarial framework. This approach harks back to [Amin et al. 2014], who proposed using stochastic gradient descent. Recently, different authors have taken this problem in very different directions. [Javanmard and Nazerzadeh 2016] show that a regularized maximum likelihood algorithm can be used to solve this problem and obtain strong performance bounds as a function of the sparsity of the feature space. In a different direction, [Qiang and Bayati 2016] show that a method as simple as greedy least squares regression performs well since the contexts (or covariates in their paper) ensure the seller performs sufficient exploration.

REFERENCES

Agarwal, A., Hsu, D., Kale, S., Langford, J., Li, L., and Schapire, R. E. 2014. Taming the monster: A fast and simple algorithm for contextual bandits. In *Proceedings of ICML*.

Amin, K., Rostamizadeh, A., and Syed, U. 2014. Repeated contextual auctions with strategic buyers. In *Proceedings of NIPS*. 622–630.

Auer, P. 2003. Using confidence bounds for exploitation-exploration trade-offs. *The Journal of Machine Learning Research 3*, 397–422.

Cohen, M., Lobel, I., and Leme, R. P. 2016. Feature-based dynamic pricing. Working paper, New York University.

Javanmard, A. and Nazerzadeh, H. 2016. Dynamic pricing in high-dimensions. Working paper, University of Southern California.

Khachiyan, L. G. 1979. A polynomial algorithm in linear programming. *Doklady Akademii Nauk SSSR 244*, 1093–1096.

Kleinberg, R. and Leighton, T. 2003. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Proceedings of FOCS*. 594–605.

LOBEL, I., LEME, R. P., AND VLADU, A. 2016. Multidimensional binary search for contextual decision-making. Working paper, New York University.

QIANG, S. AND BAYATI, M. 2016. Dynamic pricing with demand covariates. Working Paper, Stanford University.

WILKINSON, J. H. 1965. *The Algebraic Eigenvalue Problem*. Vol. 87. Clarendon Press Oxford.

Settling the Complexity of Computing Approximate Two-Player Nash Equilibria

AVIAD RUBINSTEIN
UC Berkeley

In our recent paper [Rubinstein 2016] we rule out a PTAS for the 2-Player Nash Equilibrium Problem. More precisely, we prove that there exists a constant $\epsilon > 0$ such that, assuming the Exponential Time Hypothesis for PPAD, computing an ϵ-approximate Nash equilibrium in a two-player $n \times n$ game requires time $n^{\log^{1-o(1)} n}$. This matches (up to the $o(1)$ term) the algorithm of Lipton, Markakis, and Mehta [Lipton et al. 2003].

Categories and Subject Descriptors: F.2 [**ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY**]:

General Terms: Algorithms; Economics, Theory

Additional Key Words and Phrases: Nash equilibrium; Computational complexity; PPAD

Introduction

Equilibrium is a ubiquitous assumption for modeling non-cooperative game, as well as markets, traffic, biospheres, and many other systems. Once players are at equilibrium they have no incentive to deviate. But how do they arrive at an equilibrium in the first place? This question has been studied by economists for over half a century (e.g. [Brown 1951; Robinson 1951; Lemke and Howson 1964; Scarf 1967]), but a general recipe is yet to be found. In recent decades, it was considered under the *Lens of Computation* by looking at the surrogate: "is there an efficient algorithm for computing an equilibrium?" Breakthrough results of [Daskalakis et al. 2009] and [Chen et al. 2009] proved that such an algorithm does not exist (assuming PPAD $\neq$ P). Furthermore, if a central, specially designed algorithm fails to find an equilibrium, it is even less likely that distributed, selfish agents will naturally converge to one. This puts the entire solution concept in doubt.

For the past decade, the main open question in this field was whether the computational intractability results extend to approximate equilibria. We had good reasons to hope that they don't, i.e. that two-player Nash admits a PTAS: there was a series of improved approximation factors in polynomial time [Kontogiannis et al. 2009; Daskalakis et al. 2009; 2007; Bosse et al. 2010; Tsaknakis and Spirakis 2008] and several approximation schemes for special cases [Kannan and Theobald 2007; Daskalakis and Papadimitriou 2009; Alon et al. 2013; Barman 2015]. Yet most interesting are two inefficient algorithms for two-player Nash:

—the classic Lemke-Howson algorithm [Lemke and Howson 1964] finds an exact Nash equilibrium in exponential time; and

Authors' addresses: aviad@eecs.berkeley.edu

—a simple algorithm due to [Lipton et al. 2003] finds an ϵ-Approximate Nash Equilibrium in time $n^{O(\log n)}$.

Although the Lemke-Howson algorithm takes exponential time, it has a special structure which places the problem inside the complexity class PPAD [Papadimitriou 1994]; i.e. it has a polynomial time reduction to the canonical problem EndOfALine:

Definition 1 (EndOfALine *[Daskalakis et al. 2009]).* Given two circuits S and P, with m input bits and m output bits each, such that $P(0^m) = 0^m \neq S(0^m)$, find an input $x \in \{0,1\}^m$ such that $P(S(x)) \neq x$ or $S(P(x)) \neq x \neq 0^m$.

Proving hardness for problems in PPAD is notoriously challenging because they are *total*, i.e. they always have a solution, so the standard techniques from NP-hardness do not apply. By now, however, we know that exponential and polynomial approximations for two-player Nash are PPAD-complete [Daskalakis et al. 2009; Chen et al. 2009], and so is ϵ-approximation for games with n players [Rubinstein 2015b].

ϵ-approximation for two-player Nash is unlikely to have the same fate: otherwise, the quasi-polynomial algorithm of [Lipton et al. 2003] would refute the Exponential Time Hypothesis for PPAD:

Hypothesis 2 (ETH for PPAD [Babichenko et al. 2016]). *Solving* EndOfALine *requires time* $2^{\tilde{\Omega}(n)}$.[1]

Thus the strongest hardness result we can hope to prove (given our current understanding of complexity) is a quasi-polynomial hardness that sits inside PPAD, and this is precisely the main result of [Rubinstein 2016]:

Theorem 3 (2 players [Rubinstein 2016]). *There exists a constant* $\epsilon > 0$ *such that, assuming ETH for* PPAD*, finding an* ϵ*-Approximate Nash Equilibrium in a two-player* $n \times n$ *game requires time* $T(n) = n^{\log^{1-o(1)} n}$.

Quasi-fine-grained Complexity

Two-player Nash equilibrium belongs to a growing class of fundamental problems that admit a quasi-polynomial time approximation algorithms, and also have matching conditional lower bounds on the running time. Those include problems of relevance to the SIGecom community, such as ϵ-best ϵ-Nash equilibrium [Braverman et al. 2015; Deligkas et al. 2016], Densest k-subgraph [Braverman et al. 2017; Manurangsi 2016], signaling in a zero-sum game [Rubinstein 2015a; Bhaskar et al. 2016], and community detection [Rubinstein 2017].

For all of those problems, the birthday repetition framework [Aaronson et al. 2014] gives a reduction size of $N \approx 2^{\sqrt{n}}$. Assuming the exponential time hypothesis (ETH) [Impagliazzo et al. 2001], approximating 3-SAT requires time $T(n) \approx 2^n \approx N^{\log N}$; hence the quasi-polynomial lower bound. The same blowup in instance size occurs in the proof of Theorem 3.

Unfortunately, it is not clear how to apply the birthday repetition framework to Nash equilibrium because we don't have an equivalent of the PCP Theorem for

[1] As usual, n is the size of the description of the instance, i.e. the size of the circuits S and P.

PPAD. (But recently [Babichenko et al. 2016] conjectured what a "PCP for PPAD" could look like - and proving it remains an important open problem.) The actual proof of Theorem 3 in [Rubinstein 2016] circumvents this obstacle and does not explicitly use birthday repetition. As a result, it is quite involved and requires tools from the studies of PCP, locally decodable codes, pseudorandomness, etc. In particular, this is the first time that such ideas are used for hardness of approximation inside PPAD.

Can almost everyone be almost happy?

An ϵ-Approximate Nash Equilibrium (ϵ-ANE) is a (mixed) strategy profile for which each player plays an ϵ-best response; i.e. she can gain at most (an additive) ϵ by deviating. This is the standard and most-studied notion of approximate Nash equilibrium, and it is indeed very natural for two-player games. However, for some settings with *many players*, requiring that the ϵ-best response condition holds for *every* player may be too restrictive. Recently, [Babichenko et al. 2016] introduced a more relaxed notion of (ϵ, δ)-WeakNash Equilibrium, where we only require that a $(1 - \delta)$-fraction of the players play ϵ-best response, while the remaining δ-fraction may play arbitrarily.

En route to proving Theorem 3 we obtain impossibility results for the latter WeakNash relaxation:

THEOREM 4 (n PLAYERS [RUBINSTEIN 2016]). *There exist constants $\epsilon, \delta > 0$ such that finding an (ϵ, δ)-WeakNash Equilibrium...*

Query Complexity. *requires $2^{\Omega(n)}$ oracle calls to the payoff tensor; and*

Computational Complexity. *is PPAD-hard given a succinct description of the payoff tensor.*

We note that the former result on query complexity resolves an open question posed by Hart and Nisan [Hart and Nisan 2013], Babichenko [Babichenko 2016], and Chen et al. [Chen et al. 2017]. Furthermore, in subsequent joint work with Babichenko [Babichenko and Rubinstein 2016], we extend this result to a lower bound on **communication complexity** (where each player knows her own utilities).

REFERENCES

AARONSON, S., IMPAGLIAZZO, R., AND MOSHKOVITZ, D. 2014. Am with multiple merlins. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on.* IEEE, 44–55.

ALON, N., LEE, T., SHRAIBMAN, A., AND VEMPALA, S. 2013. The approximate rank of a matrix and its algorithmic applications: approximate rank. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013.* 675–684.

BABICHENKO, Y. 2016. Query Complexity of Approximate Nash Equilibria. *J. ACM 63,* 4, 36:1–36:24.

BABICHENKO, Y., PAPADIMITRIOU, C. H., AND RUBINSTEIN, A. 2016. Can almost everybody be almost happy? In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016.* 1–9.

BABICHENKO, Y. AND RUBINSTEIN, A. 2016. Communication complexity of approximate nash equilibria. *CoRR abs/1608.06580.*

BARMAN, S. 2015. Approximating nash equilibria and dense bipartite subgraphs via an approximate version of caratheodory's theorem. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015.* 361–369.

BHASKAR, U., CHENG, Y., KO, Y. K., AND SWAMY, C. 2016. Hardness results for signaling in bayesian zero-sum and network routing games. 479–496.

BOSSE, H., BYRKA, J., AND MARKAKIS, E. 2010. New algorithms for approximate nash equilibria in bimatrix games. *Theor. Comput. Sci. 411,* 1, 164–173.

BRAVERMAN, M., KUN-KO, Y., RUBINSTEIN, A., AND WEINSTEIN, O. 2017. ETH hardness for densest-k-subgraph with perfect completeness. In *SODA*. To appear.

BRAVERMAN, M., KUN-KO, Y., AND WEINSTEIN, O. 2015. Approximating the best nash equilibrium in $n^{o(\log n)}$-time breaks the exponential time hypothesis. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*. 970–982.

BROWN, G. W. 1951. Iterative solutions of games by fictitious play. *Activity Analysis of Production and Allocation*.

CHEN, X., CHENG, Y., AND TANG, B. 2017. Well-supported versus approximate nash equilibria: Query complexity of large games. In *To appear in ITCS*.

CHEN, X., DENG, X., AND TENG, S.-H. 2009. Settling the complexity of computing two-player Nash equilibria. *J. ACM 56,* 3.

DASKALAKIS, C., GOLDBERG, P. W., AND PAPADIMITRIOU, C. H. 2009. The complexity of computing a Nash equilibrium. *Commun. ACM 52,* 2, 89–97.

DASKALAKIS, C., MEHTA, A., AND PAPADIMITRIOU, C. H. 2007. Progress in approximate nash equilibria. In *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), San Diego, California, USA, June 11-15, 2007*. 355–358.

DASKALAKIS, C., MEHTA, A., AND PAPADIMITRIOU, C. H. 2009. A note on approximate nash equilibria. *Theor. Comput. Sci. 410,* 17, 1581–1588.

DASKALAKIS, C. AND PAPADIMITRIOU, C. H. 2009. On oblivious ptas's for nash equilibrium. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*. 75–84. Full version available at http://arxiv.org/abs/1102.2280.

DELIGKAS, A., FEARNLEY, J., AND SAVANI, R. 2016. Inapproximability results for approximate nash equilibria. *CoRR abs/1608.03574*.

HART, S. AND NISAN, N. 2013. The query complexity of correlated equilibria. *CoRR abs/1305.4874*.

IMPAGLIAZZO, R., PATURI, R., AND ZANE, F. 2001. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci. 63,* 4, 512–530.

KANNAN, R. AND THEOBALD, T. 2007. Games of fixed rank: a hierarchy of bimatrix games. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*. 1124–1132.

KONTOGIANNIS, S. C., PANAGOPOULOU, P. N., AND SPIRAKIS, P. G. 2009. Polynomial algorithms for approximating nash equilibria of bimatrix games. *Theor. Comput. Sci. 410,* 17, 1599–1606.

LEMKE, C. E. AND HOWSON, J. T. 1964. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics 12,* 2, 413–423.

LIPTON, R. J., MARKAKIS, E., AND MEHTA, A. 2003. Playing large games using simple strategies. In *EC*. 36–41.

MANURANGSI, P. 2016. Almost-Polynomial Ratio ETH-Hardness of Approximating DENSEST k-SUBGRAPH with Perfect Completeness. *CoRR abs/1611.05991.*

PAPADIMITRIOU, C. H. 1994. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci. 48,* 3, 498–532.

ROBINSON, J. 1951. An iterative method of solving a game. *Annals of Mathematics 54*, 296–301.

RUBINSTEIN, A. 2015a. Eth-hardness for signaling in symmetric zero-sum games. *CoRR abs/1510.04991*.

RUBINSTEIN, A. 2015b. Inapproximability of nash equilibrium. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*. 409–418.

RUBINSTEIN, A. 2016. Settling the Complexity of Computing Approximate Two-Player Nash Equilibria. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*. 258–265.

RUBINSTEIN, A. 2017. Detecting communities is hard and counting them is even harder. In *To appear in ITCS*.

SCARF, H. 1967. The approximation of fixed points of a continuous mapping. *SIAM Journal of Applied Mathematics 15*, 1328–1343.

TSAKNAKIS, H. AND SPIRAKIS, P. G. 2008. An optimization approach for approximate nash equilibria. *Internet Mathematics 5,* 4, 365–382.