Table of Contents

Editor's Introduction 1
ARIEL D. PROCACCIA

Notes from the EC'13 Program Chairs 2
R. PRESTON MCAFEE and ÉVA TARDOS

Privacy and Mechanism Design 8
MALLESH M. PAI and AARON ROTH

Setting Equilibrium Prices, Approximately 30
BRENDAN LUCIER

Logit Dynamics: A Model for Bounded Rationality 34
DIODATO FERRAIOLI

Planning and Learning in Security Games 38
FRANCESCO M. DELLE FAVE, YUNDI QIAN, ALBERT X. JIANG, MATTHEW BROWN, and MILIND TAMBE

Generalized Scoring Rules: A Framework That Reconciles Borda and Condorcet 42
LIRONG XIA

Back to Original Frugality 49
RODRIGO A. VELEZ

Constrained Signaling for Welfare and Revenue Maximization 53
SHADDIN DUGHMI, NICOLE IMMORLICA, and AARON ROTH

Editor's Introduction

ARIEL D. PROCACCIA
Carnegie Mellon University

Issue 12.1 of SIGecom Exchanges includes a letter from the program chairs of EC'13, a survey, and six research letters. The purpose of the first letter is to update the community about some of the decisions and discussions surrounding EC'13. The previous Exchanges issue—11.2—included a letter from the SIG chair David Parkes, and going forward the plan is to alternate between these two types of updates, publishing a letter from the SIG chair in each December issue and a letter from the EC chairs in each June issue.

The survey continues a new tradition that started with Eric Budish's survey in issue 11.2. This tradition actually turned out to be very useful—when I want to catch up on a hot area all I need to do is invite a survey, wait a few months, and voilà! This time I invited Mallesh Pai and Aaron Roth to contribute a survey on differential privacy and mechanism design. There has been a lot of action in this space lately, including very recent workshops at Caltech and in NYC, and an upcoming EC workshop organized by Mallesh and Aaron. Their survey is a 22-page masterpiece that is as readable and intuitive as it is comprehensive (the area is still young enough to be covered in full). Tip: don't miss footnote 11!

Puzzle-loving readers may have noticed that issue 11.2 did not include a new puzzle, and indeed the puzzle section has been discontinued. If you sorely miss the puzzles, please drop me an email. I want to thank Daniel Reeves for his contributions as puzzle editor. As usual, Felix Fischer (the SIG's information director) played an important role in putting the issue together; thanks!

Author's address: arielpro@cs.cmu.edu.

Notes from the EC'13 Program Chairs

R. PRESTON MCAFEE
Google
and
ÉVA TARDOS
Cornell University

This note describes our experience as 2013 ACM Electronic Commerce (EC) program chairs and summarizes a survey we ran after the program was complete.

Track System

EC is the leading scientific conference on advances in theory, systems, and applications at the interface of economics and computer science, including applications to electronic commerce. EC'13 continued the 2012 track process involving three non-exclusive focus areas or tracks. The three tracks are:

—Theory and Foundations (TF)

—Artificial Intelligence and Applied Game Theory (AI)

—Experimental, Empirical, and Applications (EA)

We expanded the AI track to include applied game theory, on the principle that applied game theory is usually more like AI than CS theory. The track system represents a compromise. The expansion of the AI track helped more explicitly align the track with an area at the interface of economics and computation, and have the SPC dedicated to this track consist of a mix of computer scientists and economists, analogous to the mix present in the other tracks.

Authors were asked to align their submission with one or two of the tracks. The Program Committees, both senior and regular, were also associated with tracks. The papers submitted to one track were reviewed by PC and SPC members associated with that track, while a paper submitted to two tracks was handled by PC and SPC members in the union of the two tracks, with at least one PC and one SPC from each of the designated tracks.

Allowing authors to align their submission with one or two of the tracks seems to work well for the conference from our own experience, in the opinions expressed at last year's business meeting, and in the opinions of about 90% of those who responded to our survey this year. The main criticism of the complementary 10% was that the tracks overlapped, rendering them confusing. There was certainly some confusion between TF and AI that may be alleviated by SIGecom discussion. In contrast, there was little confusion about experimental and empirical work. The ability to submit to multiple tracks mitigates this confusion. Moreover, the identification of SPC with tracks reduces the ambiguity of submission, based on the SPC's

Authors' addresses: preston@mcafee.cc, eva@cs.cornell.edu

own research. We think the result of the track system is a more diverse program, specifically more papers by AI researchers, than would be created with a single track system. A mix of computer science and economics PC and SPC members in each track results in the desired cross-fertilization across computing and economics that is the hallmark of EC since its inception.

Program Committee and Assignment

Our first task as Program chair was to put together the Senior Program Committee (SPC), and the Program Committee (PC) for each track. Based on last year's submissions, we anticipated that more of the submissions will be for the theory track than either of the two other tracks. In spite of their greater numbers, TF SPCs handled considerably more papers each. In contrast, we aimed to have the PC burden more evenly split. Many of the PC members were suggested by the SPC.

Table 1	SPC	PC	Submissions	Accepted
TF	15	73	167	52
AI	8	33	54	16
EA	8	27	50	16

The call for papers attracted 223 submissions from authors in academia and industry all around the world, 2 of which were eventually retracted. Out of the 221 non-retracted submissions, 50 papers choose dual tracks.

We selected a diverse and capable Program Committee and senior Program Committee who are eminently capable of evaluating any paper in the field. We then optimized the matching of skills and the nature of papers as best we can. Because of the size of the matching problem we used an algorithm described below. Creating the best matching took us several days, but it insured that relevant expertise was available for all 221 papers. We believe that all the papers were carefully evaluated by peers with relevant skills.

Once the papers were submitted, we asked the SPC and PC to indicate the papers they could review, would like to review, or could not review due to a conflict. Both groups also provided keywords about their skills and interests. PC members were limited to blind selection of papers (based on title and abstract only) for confidentiality and asked to mark author conflicts separately.

Based on paper preferences, conflicts and keywords, we used integer programming to generate an initial assignment of papers to PC and SPC members. The main ingredients of this optimization were the constraints of the track system explained above and limits on the number of papers assigned to a PC or SPC, with a heavy weight on paper preferences and a lighter weight on keyword matches.

Each PC member was asked to review of a maximum of 8 papers. Similarly SPC papers also had hard upper bounds on the number of assigned papers depending on the track (higher in TF than in AI and EA).

To make sure each paper gets enough expert reviews, we aimed to have at least one of the assigned PC members and at least one of the assigned SPC very well aligned to each paper. We reviewed the initial assignment, tweaked it, and re-optimized many times to reach the final assignment. We increased the weight on paper preferences so that most reviewers received papers they requested, and most

papers were reviewed by reviewers that requested the paper. The additional tweak involved finding papers without adequate expertise (e.g. because of rare or missing keywords) and forcing an assignment to insure the expertise was available.

Furthermore, we aimed to have each PC and SPC member have at least half of their load come from papers that they specially asked for to make the PC/SPC work both enjoyable and to make the ratings fair, and this was accomplished in most cases.

Review Process

To help the make the ratings more uniform, we announced an interpretation of the ratings associated with journals in economics and computer science. PC members had an extremely short three weeks to review the assigned paper. We owe a special thanks to the diligence of the PC, as almost all reviews arrived by the deadline. The review period was followed by a two day author feedback period, where authors saw the reviews, and were given the opportunity to comment on the reviews and respond to issues raised by them. Authors were limited to 500 words.

Opinion in the community about the author feedback option is decidedly mixed. Many feel that author feedback is useful chance for authors to correct misconceptions by the reviewers, and clarify issues raised. Others feel that author feedback is a waste of time and energy. Author feedback certainly worked well for some papers where the PC and SPC suspected that the paper had a mistake, or where the PC and SPC had questions for the author. In cases where the Program Committee has questions for the author, the PC chairs can also serve as a conduit to obtain answers.

Once we had the reviews and the author feedback, the PC and SPC engaged in two weeks of discussion. The discussion period is the most interesting and valuable part of the EC Program Committee's job. Many of these discussions were quite lively and involved both computer science and economics perspectives, as befits our community. A few key points to comment on:

Most papers had extensive discussions. Usually this discussion is summarized by an additional review (meta-review) provided by one of the assigned SPC members, but unfortunately not all papers received such meta-reviews, which is a social waste because useful information is not communicated the authors.

The review process identified a number of errors in manuscripts, including erroneous proofs. In all cases we insisted that the last submission before the deadline be evaluated or the paper withdrawn. To replace a paper after the deadline is to let the deadline slip; it isn't fair to let some replace and not others. The author feedback system allowed authors to comment on small mistakes, as well as misunderstandings.

It would be desirable to reduce the number of errors in submissions. One way to achieve this is for authors to seek feedback from their colleagues before submission, and avoid the "just in time" production process favored by many computer scientists.

The discussions, involvement of the SPC, and author feedback had a major effect on the final decision. Prior to the start of the discussion period, few SPC members were involved in reviewing. Their role was to lead the discussion and guide the

decision made on the paper. A good number of reviewers adjusted their ratings during the discussion phase, but not all discussions are reflected in such changed ratings. The bigger effect of the discussion was how it shaped the final decision.

The quality of the reviews prompted the most comments in our survey. Overall, most respondents felt that the reviews were reasonable. Only 8% of the respondents were very unsatisfied with the reviews, while 33% were very satisfied. Providing high quality and useful feedback to authors is a hard task both due to the compressed time line of the EC review system, and the broad background of reviewers and authors and the broad range of areas of the conference.

One form of complaint expressed by some in the survey is that reviews, ratings, and decisions are too random. While there is certainly some randomness involved in evaluating papers, the PC members have substantial agreement on overall quality. To assess the agreement empirically, we considered the absolute variation from the mean in the "overall" rating. The three PC members submitted their scores without knowing each other's scores, so these are independently submitted. The mean was 5.47, and average deviation (absolute value) from this mean was 1.54. Thus, the average score was 1.54 away from the global mean of 5.47. In contrast, the average deviation from each paper's mean was only 0.92. So ratings of the same paper were significantly closer to each other than ratings of different papers. After the discussion, and including the SPC rating, these numbers are 1.47 and 0.81, respectively.[1]

Table 2	Initial	Final
Average Absolute Deviation	1.54	1.47
Average Within Paper Deviation	0.92	0.81

An important source of unpredictability in evaluating papers is what the reviewers find interesting. This difference was especially acute when reviewers have very different backgrounds (some are economists, others are computer scientists) because the two disciplines have different perspectives on what is interesting. This mix of economics and computer science in the discussions, as well as the mixed audience of the conference, is designed to result in the desired cross-fertilization across the disciplines.

Program Size and Structure

We decided to keep roughly the same number of parallel and single sessions as last year, accepting 72 papers in the program (same as last year), and the roughly 30% acceptance rate (on each of the three tracks). We also added two keynote speakers. We asked the SPC for suggestions and after some discussion and votes by the SPC, we asked Jon Kleinberg and Al Roth to give keynote talks at EC'13 and both accepted.

The community has strong opinions on what is the right number of papers to accept for EC. Using double sessions for part of the time, and a very compressed

[1]While absolute deviations make the deviation interpretation easier, they make it harder to interpret how much of the variation is explained by paper quality. We find that mean paper quality explains 57% of the total variation in scores, rising to 58% after the discussion. Random chance would produce 33% and 25%, respectively.

program last year's EC'12 increased the accepted papers from 48 at EC'11 to 72 at EC'12. We were strongly encouraged by the community (by discussions at EC'12 business meeting) as well as the SIGecom leadership to keep this higher number of papers at EC'13, which we did. Our survey confirmed the same range of opinions. Out of the 149 responses, the overwhelming majority (104 responses) liked the current mix of single and double session, and the remaining 45 responses were almost evenly split between wanting all single or all parallel sessions. Overall, we think that the 30% acceptance rate and the mix of single and parallel sessions worked well for EC, though many commented that the number of accepted papers should not be increased too fast. The current state appears to be a good compromise between two competing forces:

—Higher acceptance rate helps in building the community, and allows more people to give talks, allows a greater mix of talks, and broader range of topics.

—Lower acceptance rate increases the prestige of the conference and helps increase the quality relative to other venues.

As in the past years, authors of accepted papers can ask that only a one page abstract of the paper appear in the proceedings, along with a URL pointing to the full paper. This option is made available to authors to accommodate the publishing traditions of different fields, where a conference publication precludes an overlapping journal article. Authors were not asked indicate their plan to use this option during the same process.

The ACM EC best paper award rules are set by SIGecom, and are detailed at `http://www.sigecom.org/awardp.html`. One aspect of that was not discussed in the rules is the role of the one page papers. The SIGecom executive committee decided that only full length paper qualify for such awards. The conference can award one or two Best Paper Awards and one or two Best Student Paper Awards to the accepted papers. The Best Paper Award is made irrespective of whether or not a paper is a student paper - a Best Paper that is a student paper is also awarded Best Student Paper. This year, as a student paper won Best Paper, there is only one award.

The award selection process develops in two stages. The program chairs are asked to nominate a small subset of highly rated papers (5 papers this year), some of which are student papers, and are asked to form an Awards Committee from members of the senior Program Committee or the wider research community, who do not have conflicts of interest with the nominated papers. This year's award committee consisted of Susan Athey (Stanford), Vincent Conitzer (Duke), David Easley (Cornell), and Anna Karlin (University of Washington). The winner of the best paper award will be announced at EC in Philadelphia.

Last Word

We believe that the upcoming EC'13 program looks great, and want to thank all the people who contributed to make this happen. We are indebted to Kevin Leyton-Brown and Panos Ipeirotis, the 2012 chairs, for their generous assistance and advice, to Thomas Preuss of Confmaster, and Pooya Jalaly Khalilabadi for their help with running the review system. We want to thank all the authors who

submitted papers, the Program Committee and the Senior Program Committee for their contribution to this process, as well as SIGecom chair David Parkes for his continued helpful advice throughout the process.

Privacy and Mechanism Design

MALLESH M. PAI
Department of Economics, University of Pennsylvania
and
AARON ROTH
Computer and Information Sciences, University of Pennsylvania

This paper is a survey of recent work at the intersection of mechanism design and privacy. The connection is a natural one, but its study has been jump-started in recent years by the advent of *differential privacy*, which provides a rigorous, quantitative way of reasoning about the costs that an agent might experience because of the loss of his privacy. Here, we survey several facets of this study, and differential privacy plays a role in more than one way. Of course, it provides us a basis for *modeling* agent costs for privacy, which is essential if we are to attempt mechanism design in a setting in which agents have preferences for privacy. It also provides a toolkit for controlling those costs. However, perhaps more surprisingly, it provides a powerful toolkit for controlling the stability of mechanisms in general, which yields a set of tools for designing novel mechanisms even in economic settings completely unrelated to privacy.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Algorithms, Economics, Security

Additional Key Words and Phrases: Privacy, Mechanism Design

1. INTRODUCTION

Organizations such as census bureaus and hospitals have long maintained databases of personal information. However, with the advent of the Internet, many entities are now able to aggregate enormous quantities of personal and/or private information about individuals, with the intent to use it for financial gain or even malicious purposes. In reaction, several "privacy advocacy" groups have sprung up, with the intent to move US Congress and other lawmaking bodies to enact laws restricting the ability of private entities to collect and use personal information. Recent decisions by high-profile companies such as Facebook and Google have highlighted issues regarding privacy and brought them into public scrutiny.[1,2]

This interest in privacy is not solely or even largely motivated by the right to privacy as a basic desideratum. Increasingly, private information is explicitly being used for financial gain. In the recent past, companies have experimented with price discriminating against customers based on past purchase history,[3] technology

[1]Facebook has been accused of having a hard to use and frequently changing user interface for users privacy settings.
[2]Several Google projects, most recently their Glass project have drawn controversy, see, e.g. `http://blogs.wsj.com/digits/2013/05/16/congress-asks-google-about-glass-privacy/` .
[3]See, for example, `http://www.cnn.com/2005/LAW/06/24/ramasastry.website.prices/`.

Authors' addresses: `mallesh@econ.upenn.edu`, `aaroth@cis.upenn.edu`. The authors gratefully acknowledge the support of the National Science Foundation award number CCF-1101389.

choices,[4] or social profile.[5] More broadly, there are concerns that the availability of such private information may influence important parts of an individual's life, e.g. access to health insurance or employment opportunities. As a result, issues related to privacy can have a large impact on individual welfare. An understanding of how agents' private data can be used in economic settings is therefore important to guiding policy.

Motivated by these issues, this article is part survey, part position paper and part progress report. To formally study privacy, we have two "toolboxes." The older literature is the large literature on information economics, game theory and mechanism design. The modern literature on "differential privacy," on the other hand, gives a set of tools to reason about and control individual's costs for privacy loss. Combined, we can use these tools both to model settings in which agents have preferences toward privacy, and study mechanisms that trade off individual privacy with social goals. More surprisingly, the latter toolbox allows for the design of novel mechanisms in settings otherwise unrelated to privacy.

To briefly foreshadow the organization of this paper: in the next section, we quickly review the most basic aspects of differential privacy that we will use in this survey. We then study various recent contributions to mechanism design of two sorts. The first kind uses differential privacy as a tool to design novel mechanisms in settings where privacy is not a concern. The second considers the design of mechanisms in settings where agents have privacy concerns, i.e. the level of privacy the mechanism offers enters into agent's utility. Finally, we survey the (limited) literature that provides micro-foundations of preferences for privacy.

2. PRELIMINARIES

This survey is chiefly (but not exclusively) interested in *differential privacy* [Dwork et al. 2006]. Let $\mathcal{T}$ denote some type space, and let $\mathcal{O}$ denote some outcome space. We will write $t \in \mathcal{T}^n$ to denote a vector of n types, using the usual convention of indexing the i'th type by t_i, and the vector of all types *excluding* the i'th type by t_{-i}. We will say that two type vectors $t, t' \in \mathcal{T}^n$ are *neighbors* if there exists some index i such that $t_{-i} = t'_{-i}$: in other words, t and t' only differ in their i'th index. We are now prepared to define differential privacy, which will be a property of *randomized* mappings $M : \mathcal{T}^n \to \mathcal{O}$. We refer to these as *mechanisms*.

Definition 2.1. A mechanism $M : \mathcal{T}^n \to \mathcal{O}$ is ϵ-differentially private if for all pairs of neighboring type vectors $t, t' \in \mathcal{T}^n$, and for all functions $u : \mathcal{O} \to \mathbb{R}^+$:[6]

$$\mathbb{E}_{o \sim M(t)}[u(o)] \leq \exp(\epsilon)\mathbb{E}_{o \sim M(t')}[u(o)].$$

Note that the 'neighbor' relation is symmetric, so by definition, we also have the reverse inequality

$$\mathbb{E}_{o \sim M(t)}[u(o)] \geq \exp(-\epsilon)\mathbb{E}_{o \sim M(t')}[u(o)]$$

[4]See, for example, `http://www.cnn.com/2012/06/26/tech/web/orbitz-mac-users`.

[5]For example, American Airlines offers customers a free day-pass to their premium lounges if they show they are influential on online social media via a "Klout score.". See `https://secure.fly.aa.com/klout/`.

[6]We think of ϵ as being a small constant less than one, and so $\exp(\epsilon) \approx 1 + \epsilon$.

In other words, differential privacy promises that *simultaneously*, for every possible utility function $u : \mathcal{O} \to \mathbb{R}^+$, the unilateral change of a single reported type t_i to a mechanism can have only a small ($\approx 1 + \epsilon$) multiplicative effect on the expected utility of the outcome drawn from the mechanism M. We note that this definition is syntactically different from the standard definition of differential privacy [Dwork et al. 2006], but is easily seen to be equivalent.

We will work with this version of the definition, which is particularly natural in the context of mechanism design. This version of the definition also makes it apparent why differential privacy corresponds to something that one would think of as "privacy." It promises that *regardless of your preferences*, your expected utility is not substantially changed if you decide to participate in the mechanism, compared to not participating (or, say, providing random data). Thus, given the choice to participate in a differentially private computation, you should be willing if given some (small) incentive to do so.[7]

There is a large literature on differential privacy which we will not attempt to survey— we direct the reader to [Dwork and Roth 2013] for an introduction to the area. Here, we mention just one differentially private mechanism: the exponential mechanism of [McSherry and Talwar 2007].

Definition 2.2. The *exponential mechanism* is defined by a range $\mathcal{R}$, a privacy parameter ϵ, and a "quality function" $q : \mathcal{T}^n \times \mathcal{R} \to \mathbb{R}$ which has the property that for all pairs of neighboring type vectors $t, t' \in \mathcal{T}^n$, and for all $r \in \mathcal{R} : |q(t, r) - q(t', r)| \leq \Delta$. We refer to this constant Δ as the *sensitivity* of q. Given an input $t \in \mathcal{T}^n$, the exponential mechanism outputs $r \in \mathcal{R}$ according to the distribution

$$r \propto \exp\left(\frac{\epsilon q(t, r)}{2\Delta}\right).$$

The exponential mechanism is extremely useful due to the following theorem:

Theorem 2.3 [McSherry and Talwar 2007]. *The exponential mechanism is ϵ-differentially private and with probability $1 - \beta$ outputs some $r \in \mathcal{R}$ such that*

$$q(t, r) \geq \max_{r^* \in \mathcal{R}} q(t, r^*) - \frac{2\Delta}{\epsilon}\left(\ln \frac{|\mathcal{R}|}{\beta}\right).$$

In other words, the exponential mechanism is a differentially private mechanism that outputs an element from the range that has quality score that is nearly as high as possible—excepting an additive term which is linear in the sensitivity of the quality score, and only logarithmic in the cardinality of the range of the mechanism.

2.1 Differential Privacy as a Solution Concept

Let us start by recalling a basic notion from mechanism design: dominant strategy truthfulness, also known as strategyproofness. Suppose that agents $i \in \{1, \ldots, n\}$ with types $t_i \in \mathcal{T}^n$ have utility functions $u_i : \mathcal{O} \to [0, 1]$ over outcomes in $\mathcal{O}$ chosen by a mechanism M.

[7] This incentive could take the form of a monetary payment, or could simply be the joy of furthering science, or the love of filling out forms.

Definition 2.4. $M : \mathcal{T}^n \to \mathcal{O}$ is ϵ-approximately dominant strategy truthful if for every player i, for every $t_{-i} \in \mathcal{T}^{n-1}$, and for every $t' \in \mathcal{T}$:

$$\mathbb{E}_{o \sim M(t_i, t_{-i})}[u_i(o)] \geq \mathbb{E}_{o \sim M(t'_i, t_{-i})}[u_i(o)] - \epsilon$$

[McSherry and Talwar 2007] were the first to observe that differential privacy is a stronger guarantee than approximate truthfulness. Note that for $\epsilon \leq 1$, $\exp(\epsilon) \leq 1 + 2\epsilon$ and so the following proposition is immediate.

PROPOSITION 2.5. *If a mechanism M is ϵ-differentially private, then M is also 2ϵ-approximately dominant strategy truthful.*

As a solution concept, this has several robustness properties that strategy proof mechanisms do not. For example, the following is almost immediate from the definition of differential privacy: If M_1 and M_2 are both ϵ-differentially private, and f is any function (including the identity function), then M_3, defined as $M_3(t) = f(M_1(t), M_2(t))$ is 2ϵ-differentially private. This means in particular that the composition of two ϵ-differentially private mechanisms remains 4ϵ-approximately dominant strategy truthful. In contrast, the incentive properties of general strategy proof mechanisms may not be preserved under composition.

Another useful property of differential privacy follows immediately from its definition: suppose that t and $t' \in \mathcal{T}^n$ are not neighbors, but instead differ in k indices. Then we have: $\mathbb{E}_{o \sim M(t)}[u(o)] \leq \exp(k\epsilon)\mathbb{E}_{o \sim M(t')}[u(o)]$. That is, changes in up to k types changes the expected output by at most $\approx (1 + k\epsilon)$, when $k \ll 1/\epsilon$. Therefore, differentially private mechanisms make truthful reporting a $2k\epsilon$-approximate dominant strategy *even for coalitions of k agents* – i.e. differential privacy automatically provides robustness to collusion. Again, this is in contrast to general dominant-strategy truthful mechanisms, which in general offer no guarantees against collusion.

Notably, differential privacy allows for these properties in very general settings *without the use of money!* In contrast, the set of exactly dominant strategy mechanisms when monetary transfers are not allowed is extremely limited.

We conclude with a drawback of using differential privacy as a solution concept as stated, first raised in [Nissim et al. 012b]: not only is truthfully reporting one's type an approximate dominant strategy, *any report* is an approximate dominant strategy! That is, differential privacy makes the outcome approximately independent of any single agent's report. In some settings, this shortcoming can be alleviated. For example, suppose that M is a differentially private mechanism, but that agent utility functions are defined to be functions both of the outcome of the mechanism, *and* of the reported type of the agent: $u_i : \mathcal{O} \times \mathcal{T} \to [0, 1]$. Suppose furthermore that for every outcome o, truthful reporting is a best response. In other words, for all o: $u_i(o, t) \geq \max_{t'_i \in \mathcal{T}} u_i(o, t')$. In this case, it is not hard to verify that the mechanism remains approximately dominant strategy truthful, but it is no longer the case that all reports are approximate dominant strategies.

3. (DIFFERENTIAL) PRIVACY AS A TOOL IN MECHANISM DESIGN

In this section, we show how the machinery of differential privacy can be used as a tool in designing novel mechanisms.

3.1 Warmup: Digital Goods Auctions

To warm up, let us consider a simple special case of the first application of differential privacy in mechanism design, the seminal [McSherry and Talwar 2007]. Consider a *digital goods auction*, i.e. one where the seller has an unlimited supply of a good with zero marginal cost to produce, for example a piece of software or other digital media. There are n unit demand buyers for this good, each with unknown valuation $v_i \in [0,1]$. There is no prior on the bidder valuations, so a natural revenue benchmark is the revenue of the *best fixed price*. At a price $p \in [0,1]$, each bidder i with $v_i \geq p$ will buy. Therefore the total revenue of the auctioneer is

$$\text{Rev}(p, v) = p \cdot |\{i : v_i \geq p\}|.$$

The optimal revenue is the revenue of the best fixed price: $\text{OPT} = \max_p \text{Rev}(p, v)$. This setting is well studied— [Balcan et al. 2005] give a dominant strategy truthful mechanism which achieves revenue at least $\text{OPT} - O(\sqrt{n})$.

We show how a simple application of the exponential mechanism achieves revenue at least $\text{OPT} - O\left(\frac{\log n}{\epsilon}\right)$. That is, the mechanism trades exact for approximate truthfulness, but achieves an exponentially better revenue guarantee. Of course, it also inherits the benefits of differential privacy discussed previously, such as resilience to collusion, and composability.

The idea is to select a price from the exponential mechanism, using as our "quality score" the revenue that this price would obtain. As we have defined it, the exponential mechanism is parameterized by some discrete range.[8] Suppose we choose the range of the exponential mechanism to be $\mathcal{R} = \{\alpha, 2\alpha, \ldots, 1\}$? The size of the range is $|\mathcal{R}| = 1/\alpha$. What have we lost in potential revenue if we restrict ourselves to selecting a price from $\mathcal{R}$? It is not hard to see that

$$\text{OPT}_{\mathcal{R}} \equiv \max_{p \in \mathcal{R}} \text{Rev}(p, v) \geq \text{OPT} - \alpha n.$$

This is because if p^* is the price that achieves the optimal revenue, and we use a price p such that $p^* - \alpha \leq p \leq p^*$, every buyer who bought at the optimal price continues to buy, and provides us with at most α less revenue per buyer. Since there are at most n buyers, the total lost revenue is at most αn.

So how do we parameterize the exponential mechanism? We have a family of discrete ranges $\mathcal{R}$, parameterized by α. For a vector of values v and a price $p \in \mathcal{R}$, we define our quality function to be $q(v, p) = \text{Rev}(v, p)$. Observe that because each value $v_i \in [0,1]$, the *sensitivity* of q is $\Delta = 1$: changing one bidder valuation can only change the revenue at a fixed price by at most $v_i \leq 1$. Therefore, if we require ϵ-differential privacy, from Theorem 2.3, we get that with high probability, the exponential mechanism returns some price p such that

$$\text{Rev}(p, v) \geq (\text{OPT} - \alpha n) - O\left(\frac{1}{\epsilon} \ln\left(\frac{1}{\alpha}\right)\right).$$

Choosing our discretization parameter α to minimize the two sources of error, we

[8] This is not necessary, but simplifies the exposition.

find that this mechanism with high probability finds us a price that achieves revenue

$$\text{Rev}(p, v) \geq \text{OPT} - O\left(\frac{\log n}{\epsilon}\right).$$

Note that if we take (e.g.) $\epsilon = 1/\log(n)$, then we obtain a mechanism that is asymptotically exactly truthful (i.e. as the market grows large, the approximation to truthfulness becomes exact), while still achieving revenue at least $(1-o(1))$OPT, so long as OPT grows more quickly than $\log(n)^2$ with the size of the population n.

Finally, notice that we could make the reported value v_i of each agent i binding. In other words, we could allocate an item to agent i and extract payment of the selected posted price p whenever $v_i \geq p$. If we do this, the mechanism is approximately truthful, because the price is picked using a differentially private mechanism. Additionally, it is not the case that *every* report is an approximate dominant strategy: if an agent over-reports, she may be forced to buy the good at a price higher than her true value.

3.2 Approximately Truthful Equilibrium Selection Mechanisms

We now consider the problem of approximately truthful equilibrium selection, studied in [Kearns et al. 2012]. Roughly speaking, the problem is as follows: suppose we are given a game in which each player knows their own payoffs, but not others' payoffs. The players therefore do not know the equilibrium structure of this game. Even if they did, there might be multiple equilibria, with different agents preferring different equilibria. Can a mechanism offered by an intermediary incentivize agents to truthfully report their utilities and follow the equilibrium it selects?

For example, imagine a city in which (say) Google Navigation is the dominant service. Every morning, each person enters their starting point and destination, receives a set of directions, and chooses his/ her route according to those directions. Is it possible to design a navigation service such that: (1) Each agent is incentivized to report truthfully, and (2) then follow the driving directions provided? Both misreporting start and end points, and truthfully reporting start and end points, but then following a different (shorter) path are to be disincentivized.

Intuitively, our two desiderata are in conflict. In the commuting example above, if we are to guarantee that every player is incentivized to truthfully follow their suggested route, then we must compute an equilibrium of the game in question given players' reports. On the other hand, to do so, our suggested route to some player i must depend on the reported location/ destination pairs of other players. This tension will pose a problem in terms of incentives: if we compute an equilibrium of the game given the reports of the players, an agent can potentially benefit by misreport, causing us to compute an equilibrium of the wrong game.

This problem would be largely alleviated, however, if the report of agent i only has a tiny affect on the actions of agents $j \neq i$. In this case, agent i could hardly gain an advantage through his effect on other players. Then, assuming that everyone truthfully reported their type, the mechanism would compute an equilibrium of the correct game, and by definition, each agent i could do no better than follow the suggested equilibrium action. In other words, if we could compute an approximate equilibrium of the game under the constraint of *differential privacy*, then truthful reporting, followed by taking the suggested action of the coordination device would

be a Nash equilibrium. A moment's reflection reveals that the goal of privately computing an equilibrium is not possible in small games, in which an agent's utility is a highly sensitive function of the actions (and hence utility functions) of the other agents.[9] But what about in large games?

Formally, suppose we have an n player game with action set $\mathcal{A}$, and each agent with type t_i has a utility function $u_i : \mathcal{A}^n \to [0, 1]$. We say that this game is Δ-large if for all players $i \neq j$, vectors of actions $a \in \mathcal{A}^n$, and pairs of actions $a_j, a'_j \in \mathcal{A}$:

$$\left|u_i(a_j, a_{-j}) - u_i(a'_j, a_{-j})\right| \leq \Delta.$$

In other words, if some agent j unilaterally changes his action, then his affect on the payoff of any other agent $i \neq j$ is at most Δ. Note that if agent j changes his own action, then his payoff can change arbitrarily. Many games are "large" in this sense. In the commuting example above, if Alice changes her route to work she may substantially increase or decrease her commute time, but will only have a minimal impact on the commute time of any other agent Bob. The results in this section are strongest for $\Delta = O(1/n)$, but hold more generally.

First we might ask whether we need privacy at all— could it be the case that in a large game, any algorithm which computes an equilibrium of a game defined by reported types has the stability property that we want? The answer is no. As a simple example, consider n people who must each choose whether to go to the beach (B) or the mountains (M). People privately know their types— each person's utility depends on his own type, his action, and the fraction of other people p who go to the beach. A Beach type gets a payoff of $10p$ if he visits the beach, and $5(1-p)$ if he visits the mountain. A mountain type gets a payoff $5p$ from visiting the beach, and $10(1-p)$ from visiting the mountain. Note that this is a large game— each player's payoffs are insensitive in the actions of others. Further, note that "everyone visits beach" and "everyone visits mountain" are both equilibria of the game, regardless of the realization of types. Consider the mechanism that attempts to implement the following social choice rule— "if the number of beach types is less than half the population, send everyone to the beach, and vice versa." It should be clear that if mountain types are just in the majority, then each mountain type has an incentive to misreport as a beach type; and vice versa. As a result, even though the game is "large" and agents' actions do not affect others' payoffs significantly, simply computing equilibria from reported type profiles does not in general lead to even approximately truthful mechanisms.

Nevertheless, [Kearns et al. 2012] are able to give a mechanism with the following property: it elicits the type t_i of each agent, and then computes an α-approximate correlated equilibrium of the game defined by the reported types.[10] It draws an action profile $a \in \mathcal{A}^n$ from the correlated equilibrium, and reports action a_i to each agent i. The algorithm has the guarantee that simultaneously for all players i, the

[9]Positive results are not beyond hope in small games for slightly different settings. See, e.g. [Dziuda and Gradwohl 2012].

[10]A correlated equilibrium is defined by a joint distribution on profiles of actions, $\mathcal{A}^n$. For an action profile a drawn from the distribution, if agent i is told only a_i, then playing action a_i is a best response given the induced conditional distribution over a_{-i}. An α-approximate correlated equilibrium is one where deviating improves an agent utility by at most α.

joint distribution a_{-i} on reports to all players *other than* i is differentially private in the reported type of agent i. This guarantee is sufficient for approximate truthfulness, because it means that agent i cannot substantially change the distribution on actions induced on *the other players* by misreporting his own type.

More specifically, when the mechanism of [Kearns et al. 2012] computes an α-approximate correlated equilibrium while satisfying ϵ-differential privacy, every agent following the honest behavior (i.e. first reporting their true type, then following their suggested action) forms an $(2\epsilon + \alpha)$-approximate Nash equilibrium. This is because, by privacy, reporting your true type is a 2ϵ-approximate dominant strategy, and given that everybody reports their true type, the mechanism computes an α-approximate correlated equilibrium of the true game, and hence by definition, following the suggested action is an α-approximate best response. [Kearns et al. 2012] give mechanisms for computing α-approximate equilibrium in large games with $\alpha = O\left(\frac{1}{\sqrt{n\epsilon}}\right)$. Therefore, by setting $\epsilon = O\left(\frac{1}{n^{1/4}}\right)$, this gives an η-approximately truthful equilibrium selection mechanism for

$$\eta = 2\epsilon + \alpha = O\left(\frac{1}{n^{1/4}}\right).$$

In other words, it gives a mechanism for coordinating equilibrium behavior in large games that is asymptotically truthful in the size of the game, all without the need for monetary transfers.

3.3 Obtaining Exact Truthfulness

So far we have discussed mechanisms that are *asymptotically truthful* in large population games. However, what if we want to insist on mechanisms that are *exactly* dominant strategy truthful, while maintaining some of the nice properties enjoyed by our mechanisms so far: for example, that the mechanisms do not need to be able to extract monetary payments? Can differential privacy help here? It can—in this section, we discuss a special case of a framework laid out by [Nissim et al. 012b] which uses differentially private mechanisms as a building block towards designing exactly truthful mechanisms without money.

The basic idea is simple and elegant. As we have seen, the exponential mechanism can often give excellent utility guarantees while preserving differential privacy. This doesn't yield an exactly truthful mechanism, but it gives every agent very little incentive to deviate from truthful behavior. What if we could pair this with a second mechanism which need not have good utility guarantees, but gives each agent a strict positive incentive to report truthfully, i.e. a mechanism that essentially only punishes non-truthful behavior? Then, we could randomize between running the two mechanisms. If we put enough weight on the punishing mechanism, then we inherit its strict-truthfulness properties. The remaining weight that is put on the exponential mechanism contributes to the utility properties of the final mechanism. The hope is that since the exponential mechanism is approximately strategy proof to begin with, the randomized mechanism can put small weight on the strictly truthful punishing mechanism, and therefore will have good utility properties.

To design punishing mechanisms, [Nissim et al. 012b] work in a slightly non-standard environment. Rather than simply picking an outcome, they model a

mechanism as picking an outcome, and then an agent as choosing a *reaction* to that outcome, which together define his utility. They then give the mechanism the power to *restrict the reactions allowed by the agent based on his reported type.* Formally, they work in the following framework:

Definition 3.1 *The Environment [Nissim et al. 012b].* An environment is a set N of n players, a set of types $t_i \in \mathcal{T}$, a finite set $\mathcal{O}$ of outcomes, a set of reactions R and a utility function $u_i : T \times \mathcal{O} \times R \to [0,1]$ for each agent i.

We write $r_i(t, s, \hat{R}_i) \in \arg\max_{r \in \hat{R}_i} u_i(t, s, r)$ to denote i's optimal reaction among choices $\hat{R}_i \subseteq R$ to alternative s if he is of type t.

A direct revelation mechanism $\mathcal{M}$ defines a game which is played as follows:

(1) Each player i reports a type $t'_i \in \mathcal{T}$.
(2) The mechanism chooses an alternative $s \in \mathcal{O}$ and a subset of reactions for each player $\hat{R}_i \subseteq R$.
(3) Each player chooses a reaction $r_i \in \hat{R}_i$ and experiences utility $u_i(t_i, s, r_i)$.

Agents play so as to maximize their own utility. Note that since there is no further interaction after the 3rd step, rational agents will pick $r_i = r_i(t_i, s, \hat{R}_i)$, and so we can ignore this as a strategic step. Let $\mathcal{R} = 2^R$. Then a mechanism is a randomized mapping $\mathcal{M} : \mathcal{T} \to \mathcal{O} \times \mathcal{R}^n$.

Let us consider the utilitarian welfare criterion: $F(t, s, r) = \frac{1}{n} \sum_{i=1}^{n} u_i(t_i, s, r_i)$, Note that this has sensitivity $\Delta = 1/n$, since each agent's utility lies in the range $[0, 1]$. Hence, if we simply choose an outcome s and allow each agent to play their best response reaction, the exponential mechanism is an ϵ-differentially private mechanism, which by Theorem 2.3, achieves social welfare at least $\mathrm{OPT} - O\left(\frac{\log|\mathcal{O}|}{\epsilon n}\right)$ with high probability. Let us denote this instantiation of the exponential mechanism, with quality score F, range $\mathcal{O}$ and privacy parameter ϵ, as $\mathcal{M}_\epsilon$.

The idea is to randomize between the exponential mechanism (with good social welfare properties) and a strictly truthful mechanism which punishes false reporting (but with poor social welfare properties). If we mix appropriately, then we will get an exactly truthful mechanism with reasonable social welfare guarantees.

Here is one such punishing mechanism which is simple, but not necessarily the best for a given problem:

Definition 3.2 *[Nissim et al. 012b].* The commitment mechanism $M^P(t')$ selects $s \in \mathcal{O}$ uniformly at random and sets $\hat{R}_i = \{r_i(t'_i, s, R_i)\}$, i.e. it picks a random outcome and forces everyone to react as if their reported type was their true type.

Define the *gap* of an environment as

$$\gamma = \min_{i, t_i \neq t'_i, t_{-i}} \max_{s \in \mathcal{O}} \left(u_i(t_i, s, r_i(t_i, s, R_i)) - u_i(t_i, s, r_i(t'_i, s, R_i)) \right),$$

i.e. γ is a lower bound over players and types of the worst-case cost (over s) of mis-reporting. Note that for each player, this worst-case is realized with probability at least $1/|\mathcal{O}|$. Therefore we have the following simple observation:

Lemma 3.3. *For all i, t_i, t'_i, t_{-i}:*

$$u_i(t_i, \mathcal{M}^P(t_i, t_{-i})) \geq u_i(t_i, \mathcal{M}^P(t'_i, t_{-i})) + \frac{\gamma}{|\mathcal{O}|}$$

Note that the commitment mechanism is strictly truthful: every individual has at least a $\frac{\gamma}{|\mathcal{O}|}$ incentive not to lie.

This suggests an exactly truthful mechanism with good social welfare guarantees:

Definition 3.4. The punishing exponential mechanism $\mathcal{M}_\epsilon^P(t)$ defined with parameter $0 \leq q \leq 1$ is, selects the exponential mechanism $\mathcal{M}_\epsilon(t)$ with probability $1-q$ and the punishing mechanism $\mathcal{M}^P(t)$ with complementary probability q.

The following two theorems from [Nissim et al. 012b] show incentive and social welfare properties of this mechanism.

THEOREM 3.5. *If* $2\epsilon \leq \frac{q\gamma}{|\mathcal{O}|}$ *then* $\mathcal{M}_\epsilon^P$ *is strictly truthful.*

THEOREM 3.6. *For sufficiently large* n, M_ϵ^P *achieves social welfare at least*

$$OPT - O\left(\sqrt{\frac{|\mathcal{O}|\log|\mathcal{O}|}{\gamma n}}\right)$$

Note that this mechanism is truthful without the need for payments!

Let us now consider an application of this framework: the facility location game. Suppose that a city wants to build k hospitals to minimize the average distance between each citizen and their closest hospital. To simplify matters, we make the mild assumption that the city is built on a discretization of the unit line.[11] Formally, for all i let $L(m) = \{0, \frac{1}{m}, \frac{2}{m}, \ldots, 1\}$ denote the discrete unit line with step-size $1/m$. $|L(m)| = m+1$. Let $\mathcal{T} = R_i = L(m)$ for all i and let $|\mathcal{O}| = L(m)^k$. Define the utility of agent i to be:

$$u_i(t_i, s, r_i) = \begin{cases} -|t_i - r_i|, & \text{If } r_i \in s; \\ -1, & \text{otherwise.} \end{cases}$$

In other words, agents are associated with points on the line, and an outcome is an assignment of a location on the line to each of the k facilities. Agents can react to a set of facilities by deciding which one to go to, and their cost for such a decision is the distance between their own location (i.e. their type) and the facility that they have chosen. Note that $r_i(t_i, s)$ is here the closest facility $r_i \in s$.

We can instantiate Theorem 3.6. In this case, we have: $|\mathcal{O}| = (m+1)^k$ and $\gamma = 1/m$, because any two positions $t_i \neq t_i'$ differ by at least $1/m$. Hence, we have:

THEOREM 3.7 [NISSIM ET AL. 012B]. M_ϵ^P *instantiated for the facility location game is strictly truthful and achieves social welfare at least:*

$$OPT - O\left(\sqrt{\frac{km(m+1)^k \log m}{n}}\right)$$

This is already very good for small numbers of facilities k, since we expect that $\text{OPT} = \Omega(1)$. We note that for the facility location problem, [Nissim et al. 012b] derive a superior bound using a more refined argument.

[11] If this is not the case, we can easily raze and then re-build the city.

4. THE VALUE OF PRIVACY

In the previous section, we saw that differential privacy can be useful as a tool to design mechanisms, *for agents who care only about the outcome chosen by the mechanism*. We here primarily viewed privacy as a tool to accomplish goals in traditional mechanism design. As a side affect, these mechanisms also preserved the privacy of the reported player types. Is this itself a worthy goal? *Why* might we want our mechanisms to preserve the privacy of agent types?

A bit of reflection reveals that agents might care about privacy. Indeed, basic introspection suggests that in the real world, agents value the ability to keep certain "sensitive" information private, for example, health information or sexual preferences. In this section, we consider the question of how to model this value for privacy, and various approaches taken in the literature.

A first option is to just model value for privacy as a part of the agent's preferences. At one level, this is a satisfactory approach. Agents do seem to value privacy, and this is in the spirit of economic modeling "*De Gustibus non disputandum est.*" However, such a "reduced form" approach may not be helpful in policy analysis.

Recall our original motivation for differential privacy— it quantifies the worst case harm that can befall an agent from revealing his private data. A structural model of how the agent evaluates this harm may, therefore, be helpful in understanding both the individual value of privacy and the social value of privacy policies. For example, consider how an agent values the privacy of his health information. Consider two scenarios, one where health insurers or potential employers can discriminate based on an agent's health history and another where they cannot do so.[12] *Ceteris paribus*, it seems reasonable that the dis-utility he suffers from his health information being made public is different in these two scenarios. A more structural model of preferences for privacy may therefore be more appropriate for understanding, e.g., the social value of privacy policies.

There have been a few notable papers that study privacy policy in dynamic models.[13] Agents' preferences for privacy in a period derive from how other players can use the information revealed against the agent in future periods.

Most of the papers we survey are motivated by repeat purchasers in electronic commerce settings. Information about purchases made by an agent in a setting with limited privacy can be used to learn about his 'payoff type,' and therefore better price discriminate subsequently. An agent understands this and may distort his early purchases, depending on the privacy policy. Similarly, pricing by a profit maximizing seller also depends on the privacy policy. The trade-offs between various privacy policies can be studied by computing the equilibrium welfare and revenue under these policies.

An early paper in this area is [Taylor 2004]. He studies a setting where consumers purchase from firm 1 in period 1 and then firm 2 in period 2. Consumers have

[12] As an aside, we should point out that simply banning discrimination on a certain attribute may not be sufficient to prevent discrimination using other information correlated with that attribute. See for example [Chan and Eyster 2003], and a related definition of fairness [Dwork et al. 2012]

[13] The broader field of information economics studies the value of information a variety of settings too vast to survey here. We restrict attention to papers the explicitly study the value of privacy policies.

additive preferences for the two goods and may have either a high or low value for each good. These values are privately known to them, and for any given customer the two are unconditionally correlated. The paper studies three scenarios. In the first, firm 1 must keep purchases by the consumers private. In the second, firm 1 can sell this information to firm 2, but consumers are unsophisticated. In other words consumers' purchase decisions in period 1 are myopic, not taking into account how this will influence the prices they are offered in period 2. The author shows that firms fare well in the latter relative to the former. Finally, the author considers the equilibrium of a model where consumers are strategic rather than unsophisticated. In this setting, consumers may strategically reduce demand in the first period. This undermines the market for consumer information, and the firms may prefer to commit to a policy of no market for consumer information.

Another early seminal paper is that of [Calzolari and Pavan 2006]. They study a more general setting where an agent with private information sequentially contracts with two principals. The upstream principal may sell information to the downstream principal after contracting with the agent. This 'privacy policy' is a part of the contract offered by the upstream principal, and he can commit to this. The agent is sophisticated and takes this into account. They provide a general characterization of settings in which the upstream principal offers full privacy. This hinges on three conditions being satisfied. Firstly, they require that the agent's trade with the downstream principal is not directly payoff relevant to the upstream principal. Secondly, they require that the agent's valuations across the two principals be positively correlated. Finally, they require that the preferences in the downstream relationship are separable so that the do not depend on the upstream level of trade. If any of these conditions are violated, full privacy need not remain optimal. Surprisingly, the paper shows that the agent may strictly prefer disclosure. In other words, the equilibrium value of privacy may be negative, which runs counter to our intuitions about the value of privacy.

[Conitzer et al. 2012] study a setting where a monopolist seller in a two-period model cannot commit to future prices.[14] Each agent has unit demand in each period, and a private value. In the second period, therefore, the monopolist conditions the price he offers on whether or not the agent bought in the first period. Intuitively, an agent who bought in the first period will face a higher price. The agent realizes this and therefore may not buy in the first period even if it is myopically optimal. In the model, agents may be able to 'buy' privacy, i.e. avoid being identified as past customers, but possibly at a cost. On a similar vein to the previous two papers, they note that in this setting, if this privacy is available for free, all consumers will choose to purchase it, but in equilibrium this will be worse for the agent and better for the monopolist. Increasing the cost of anonymity can actually benefit consumers.

Finally, in a recent paper, [Bergemann et al. 2013] study price discrimination by a monopolist who has some exogenously specified information additional to the prior distribution of the buyers' type. The change in consumer and producer surplus from this additional information can thus be thought of as the value of privacy of this information.

[14]See also [Acquisti and Varian 2005] for a related study.

This suggests that there is much to be done in terms of modeling and understanding preferences for privacy. Our basic intuition, i.e., that private information can be used to price discriminate against an agent, and therefore privacy is "good" for an agent, is not reflected in models where information is revealed by strategic purchases. Indeed, these papers unambiguously suggest that with strategic agents, the value of privacy is negative!

5. MECHANISM DESIGN FOR PRIVACY AWARE AGENTS

Having established that agents might have preferences for privacy, it is worth considering the design of mechanisms that preserve privacy *as an additional goal*, even for tasks such as, e.g. welfare maximization that we can already solve non-privately. As we will see, it is indeed possible to generalize the VCG mechanism to *privately* approximately optimize social welfare in *any* social choice problem, with a smooth trade-off between the privacy parameter and the approximation parameter, all while guaranteeing exact dominant strategy truthfulness.

However, we might wish to go further. In the presence of agents with preferences for privacy, if we wish to design truthful mechanisms, we must somehow model their preferences for privacy in their utility function, and then design mechanisms which are truthful with respect to these new "privacy aware" utility functions. As we have seen with differential privacy, it is most natural to model privacy as a property of the mechanism itself. Thus, our utility functions are not merely functions of the outcome, but functions of the outcome and of the mechanism itself. In almost all models, agent utilities for outcomes are treated as linearly separable, that is, we will have for each agent i,

$$u_i(o, \mathcal{M}, t) \equiv \mu_i(o) - c_i(o, \mathcal{M}, t).$$

Here $\mu_i(o)$ represents agent i's utility for outcome o and $c_i(o, \mathcal{M}, t)$ the (privacy) cost that agent i experiences when outcome o is chosen with mechanism $\mathcal{M}$.

We will first consider perhaps the simplest (and most naïve) model for the privacy cost function c_i, following [Ghosh and Roth 2011]. Recall that for $\epsilon \ll 1$, differential privacy promises that for each agent i, and for every possible utility function f_i, type vector $t \in \mathcal{T}^n$, and deviation $t' \in \mathcal{T}$:

$$|\mathbb{E}_{o \sim M(t_i, t_{-i})}[f_i(o)] - \mathbb{E}_{o \sim M(t'_i, t_{-i})}[f_i(o)]| \leq\approx \epsilon \mathbb{E}_{o \sim M(t)}[f_i(o)].$$

If we view f_i as representing the "expected future utility" for agent i, it is therefore natural to model agent i's cost for having his data used in an ϵ-differentially private computation as being linear in ϵ. That is, we think of agent i as being parameterized by some value $v_i \in \mathbb{R}$, and take:

$$c_i(o, \mathcal{M}, t) = \epsilon v_i$$

where ϵ is the smallest value such that $\mathcal{M}$ is ϵ-differentially private. Here we imagine v_i to represent a quantity like $\mathbb{E}_{o \sim M(t)}[f_i(o)]$. In this setting, c_i does not depend on the outcome o or the type profile t.

Using this naïve privacy measure, we discuss a basic problem in private data analysis: how to collect the data, when the owners of the data value their privacy and insist on being compensated for it. In this setting, there is no "outcome" that agents value, other than payments, there is only dis-utility for privacy loss. We will

then discuss shortcomings of this (and other) measures of the dis-utility for privacy loss, as well as privacy in more general mechanism design settings when agents *do* have utility for the outcome of the mechanism.

Our discussion here is in the context of a specific setting, i.e., the sensitive surveyor's problem. Gradwohl [2012] considers the abstract problem of what social choice functions can be implemented when agents have preferences for privacy. He shows that extensive game forms are useful when agents have privacy concerns. We do not discuss this paper here, and refer interested readers to the original manuscript for details.

5.1 A Private Generalization of the VCG Mechanism.

Suppose we have a general social choice problem, defined by an outcome space $\mathcal{O}$, and a set of agents N with arbitrary preferences over the outcomes given by $u_i : \mathcal{O} \to [0,1]$. We might want to choose an outcome $o \in \mathcal{O}$ to maximize the *social welfare* $F(o) = \frac{1}{n}\sum_{i=1}^{n} u_i(o)$. It is well known that in any such setting, the *VCG* mechanism can implement the outcome o^* which exactly maximizes the social welfare, while charging payments that make truth-telling a dominant strategy. What if we want to achieve the same result, while also preserving privacy? How must the privacy parameter ϵ trade off with our approximation to the optimal social welfare?

Recall that we could use the exponential mechanism to choose an outcome $o \in \mathcal{O}$, with quality score F. For privacy parameter ϵ, this would give a distribution $\mathcal{M}_\epsilon$ defined to be $\Pr[\mathcal{M}_\epsilon = o] \propto \exp\left(\frac{\epsilon F(o)}{2n}\right)$. Moreover, this mechanism has good social welfare properties: with probability $1-\beta$, it selects some o such that: $F(o) \geq F(o^*) - \frac{2}{\epsilon n}\left(\ln \frac{|\mathcal{O}|}{\beta}\right)$. But as we saw, differential privacy only gives ϵ-approximate truthfulness.

However, [Huang and Kannan 2012] show that $\mathcal{M}_\epsilon$ is the solution to the following exact optimization problem:

$$\mathcal{M}_\epsilon = \arg\max_{\mathcal{D}\in\Delta\mathcal{O}} \mathbb{E}_{o\sim\mathcal{D}}[F(o)] + \frac{2}{\epsilon n} H(\mathcal{D})$$

where H represents the *Shannon Entropy* of the distribution $\mathcal{D}$. In other words, the exponential mechanism is the distribution which exactly maximizes the expected social welfare, *plus* the entropy of the distribution weighted by $2/(\epsilon n)$. This result implies that the exponential mechanism is *maximal in distributional range*, and hence can be paired with payments to make it exactly truthful.[15] Moreover, they show how to charge payments in such a way as to preserve privacy. The upshot is that for any social choice problem, the social welfare can be approximated in a manner that both preserves differential privacy, and is exactly truthful.

[Chen et al. 2013] also give an (almost equivalent) private generalization of the VCG mechanism and show conditions under which it is truthful *even taking into account agent preferences for privacy.* We discuss this in Section 5.3.

[15]One way to see this is to view the exponential mechanism as exactly maximizing the social welfare in an augmented setting in with an additional player who cares only about entropy. The payments which make this mechanism truthful are the VCG payments for the augmented game.

5.2 The Sensitive Surveyor's Problem

In this section, we consider the problem of a data analyst who wishes to conduct a study using the private data of a collection of individuals. However, he must *convince* these individuals to hand over their data! Individuals experience costs for privacy loss. The data analyst can mitigate these costs by guaranteeing differential privacy and compensating them for their loss, while trying to get a representative sample of data. We here closely follow a survey of [Roth 2012].

Consider the following stylized problem of the sensitive surveyor Alice. She is tasked with conducting a survey of a set of n individuals N, to determine what proportion of the individuals $i \in N$ satisfy some property $P(i)$. Her ultimate goal is to discover the true value of this statistic, $s = \frac{1}{n}|\{i \in N : P(i)\}|$, but if that is not possible, she will be satisfied with some estimate $\hat{s}$ such that the error, $|\hat{s} - s|$, is minimized. We will adopt a notion of accuracy based on large deviation bounds, and say that a surveying mechanism is α-accurate if $\Pr[|\hat{s} - s| \geq \alpha] \leq \frac{1}{3}$. The inevitable catch is that individuals value their privacy and will not participate in the survey for free. Individuals experience some *cost* as a function of their loss in privacy when they interact with Alice, and must be compensated for this loss. To make matters worse, these individuals are rational (i.e. selfish) agents, and are apt to misreport their costs to Alice if doing so will result in a financial gain. This places Alice's problem squarely in the domain of mechanism design, and requires Alice to develop a scheme for trading off statistical accuracy with cost, all while managing the incentives of the individuals.

As an aside, this stylized problem broadly relevant to any organization that makes use of collections of potentially sensitive data. This includes, for example, the use of search logs to provide search query completion and the use of browsing history to improve search engine ranking, the use of social network data to select display ads and to recommend new links, and the myriad other data-driven services now available on the web. In all of these cases, value is being derived from the statistical properties of a collection of sensitive data in exchange for some payment.[16].

Collecting data in exchange for some fixed price could lead to a biased estimate of population statistics, because such a scheme will result in collecting data only from those individuals who value their privacy less than the price being offered. To obtain an accurate estimate of the statistic, it is therefore natural to consider buying private data using an auction, which was recently considered in [Ghosh and Roth 2011]. There are two obvious obstacles which one must confront when conducting an auction for private data, and an additional obstacle which is less obvious but more insidious. The first obstacle is that one must have a quantitative formalization of "privacy" which can be used to measure agents' costs under various operations on their data. Here, differential privacy provides an obvious tool. For small values of ϵ, because $\exp(\epsilon) \approx (1 + \epsilon)$, it is natural to model agents as having some *linear* cost for participating in a private study. We here imagine that each agent i has an unknown value for privacy v_i, and experiences a cost $c_i(\epsilon) = \epsilon v_i$ when his private data is used in an ϵ-differentially private manner.[17] The second

[16]The payment need not be explicit and/ or dollar denominated— e.g. it may be the use of a "free" service.

[17]As we will discuss later, this assumption can be problematic.

obstacle is that our objective is to trade off with *statistical accuracy*, and the latter is not well-studied objective in mechanism design.

The final, more insidious obstacle, is that an individual's cost for privacy loss may be highly correlated with his private data itself! Suppose we only know Bob has a high value for privacy of his AIDS status, and do not explicitly know, this is disclosive because Bob's AIDS status is likely correlated with his value for privacy. More to the point, suppose that in the first step of a survey of AIDS prevalence, we ask each individual to report their value for privacy, with the intention of then running an auction to choose which individuals to buy data from. If agents report truthfully, we may find that the reported values naturally form two clusters: low value agents, and high value agents. In this case, we may have learned something about the population statistic even before collecting any data or making any payments— and therefore, the agents will have already experienced a cost. As a result, the agents may misreport their value, which could introduce a bias in the survey results. This phenomenon makes direct revelation mechanisms problematic, and distinguishes this problem from classical mechanism design.

5.2.1 *Direct Revelation Mechanisms.* Armed with a means of quantifying an agent i's loss for allowing his data to be used by an ϵ-differentially-private algorithm ($c_i(\epsilon) = \epsilon v_i$), we are almost ready to describe results for the sensitive surveyor's problem. Recall that a differentially private mechanism is some mapping $M : \mathcal{T}^n \to \mathcal{O}$, for a general type space $\mathcal{T}$. It remains to define what exactly the type space $\mathcal{T}$ is. We will consider two models. In both models, we will associate with each individual a bit $b_i \in \{0, 1\}$ which represents whether they satisfy the sensitive predicate $P(i)$, as well as a value for privacy $v_i \in \mathbb{R}^+$.

(1) In the *insensitive value model*, we calculate the ϵ parameter of the private mechanism by letting the type space be $\mathcal{T} = \{0, 1\}$: i.e. we measure privacy cost only with respect to how the mechanism treats the sensitive bit b_i, and ignore how it treats the reported values for privacy, v_i.[18]
(2) In the *sensitive value model*, we calculate the ϵ parameter of the private mechanism by letting the type space be $\mathcal{T} = (\{0, 1\} \times \mathbb{R}^+)$: i.e. we measure privacy with respect to how it treats the pair (b_i, v_i) for each individual.

Intuitively, the insensitive value model treats individuals as ignoring the potential privacy loss due to correlations between their values for privacy and their private bits, whereas the sensitive value model treats individuals as assuming these correlations are worst-case, i.e., their values v_i are just as disclosive as their private bits b_i. [Ghosh and Roth 2011] show that in the insensitive value model, one can derive approximately optimal direct revelation mechanisms that achieve high accuracy and low cost. By contrast, in the *sensitive value model*, no individually rational direct revelation mechanism can achieve any non-trivial accuracy.

Note that here we are considering a setting in which private data and costs are adversarially chosen. If we are willing to assume a known prior on agent costs (but still assume adversarially chosen private bits b_i), then it is possible to improve on the results of Ghosh and Roth [2011], and derive Bayesian optimal mechanisms for

[18]That is, the part of the mapping dealing with reported values need not be differentially private.

the sensitive survey problem as is done in [Roth and Schoenebeck 2012].

5.2.2 *Take it Or Leave it Mechanisms.* Given the impossibility result of Ghosh and Roth [2011] for the sensitive value model, the immediate question is what is possible in this setting. Two methods have been recently proposed, which we briefly summarize here. Both approaches abandon direct revelation mechanisms in favor of mechanisms which offer individuals take-it-or-leave-it offers, but both also require subtle changes in how individuals' privacy preferences are modeled. Readers are directed to [Fleischer and Lyu 2012; Ligett and Roth 2012] for more details.

Circumventing Impossibility with a Sensitive Surveyor. Suppose an individual is made a take it or leave it offer: "If you let us use your bit b_i in an ϵ-differentially private manner, I will give you \$10." If values are correlated with private data, an agent's response might reveal something about his bit b_i beyond that which is revealed through the differentially private computation. To model such correlations, Fleischer and Lyu [2012] assume that each individual's value v_i is drawn independently from one of two priors: $v_i \sim F_x$ if $b_i = x$ for $x = 0, 1$, known to Alice, the surveyor. Under this assumption, Fleischer and Lyu elegantly construct a take-it-or-leave-it offer which an agent can truthfully decide to accept or reject without revealing *anything* about his private bit! The idea is this: Alice may choose some acceptance probability $q \in [0, 1]$. She picks p_0, p_1 such that $\Pr_{v\sim F_0}[v \leq p_0/\epsilon] = \Pr_{v\sim F_1}[v \leq p_1/\epsilon] = q$. Alice can then offer the following take-it-or-leave-it offer to each agent: "If you accept the offer and your (verifiable) bit is x, I will pay you p_x dollars, for $x = 1, 2$." The beauty of this solution is that no matter what private bit the agent has, he will accept the offer with probability q (where the probability is over the corresponding prior) and reject the offer with probability $1 - q$. Therefore, *nothing* can be learned about his private bit from his participation decision, and so he has no incentive not to respond to the offer truthfully. Using this idea, Fleischer and Lyu [2012] develop approximately optimal mechanisms that can be used if the priors F_0 and F_1 are known.

Circumventing Impossibility with an Insensitive Surveyor. What if agent costs are determined adversarially, and there are no known priors? Ligett and Roth [2012] give an alternative solution for this case. To circumvent the impossibility result of Ghosh and Roth [2011], Alice has one additional power: the ability to accost random members of the population on the street, and present them with a take-it-or-leave-it offer. Once individuals are presented with an offer, they are free to accept or refuse as they see fit. However they do not have the option to *not participate*. If they reject the offer, or even just walk away, this is observed by Alice. This can be seen as a weakening of the standard individual rationality condition. Because costs may be correlated with private data, merely by rejecting an offer or walking away, Alice may learn something about the surveyed individual. If the individual rejects the offer, he receives no payment and yet still experiences some cost! This ends up giving a semi-truthfulness guarantee. If Alice makes an offer of p dollars in exchange for ϵ-differential privacy, a rational agent will accept whenever $p \geq \epsilon v_i$. However, rational agents may accept offers that are below their cost—because they will still experience some cost by walking away. But these deviations away from "truthfulness" are in only one direction, and only help Alice, whose

aim it is to compute an accurate population statistic, and does not necessarily care about protecting privacy for its own sake. Ligett and Roth [2012] thus obtain non-trivial accuracy in the sensitive value model by making Alice insensitive to the privacy concerns of the agents she surveys, by making offers that they can refuse, but can't avoid.

5.3 Better Measures for the Cost of Privacy

In the previous section, we took the naive modeling assumption that the cost experienced by participation in an ϵ-differentially private mechanism M was $c_i(o, \mathcal{M}, t) = \epsilon v_i$ for some numeric value v_i. This measure is problematic for several reasons. First, as pointed out by [Nissim et al. 012a], although differential privacy promises that any agent's loss in utility is *upper bounded* by a quantity that is (approximately) linear in ϵ, there is no reason to believe that agents' costs are *lower bounded* by such a quantity. That is, while taking $c_i(o, \mathcal{M}, t) \leq \epsilon v_i$ is well motivated, there is little support for making the inequality an equality. Second, as discussed in [Ligett and Roth 2012], *any* privacy measure which is a deterministic function only of ϵ (not just a linear function) leads to problematic behavioral predictions.

So how else might we model c_i? One natural measure, proposed by [Xiao 2013], is the *mutual information* between the reported type of agent i, and the outcome of the mechanism.[19] For this to be well defined, we must be in a world where each agent's type t_i is drawn from a known prior, $t_i \sim \mathcal{T}$. Each agent's strategy is a mapping $\sigma_i : \mathcal{T} \to \mathcal{T}$, determining what type he reports, given his true type. We could then define

$$c_i(o, \mathcal{M}, \sigma) = I(\mathcal{T}; \mathcal{M}(t_{-i}, \sigma(\mathcal{T})),$$

where I is the mutual information between the random variable $\mathcal{T}$ representing the prior on agent i's type, and $\mathcal{M}(t_{-i}, \sigma(\mathcal{T}))$, the random variable representing the outcome of the mechanism, given agent i's strategy.

This measure has significant appeal, because it represents how "related" the output of the mechanism is to the true type of agent i. However, in addition to requiring a prior over agent types, [Nissim et al. 012a] observe an interesting paradox that results from this measure of privacy loss. Consider a world in which there are two kinds of sandwich breads: Rye (R), and Wheat (W). Moreover, in this world, sandwich preferences are highly embarrassing and held private. The prior on types $\mathcal{T}$ is uniform over R and W, and the mechanism $\mathcal{M}$ simply gives agent i a sandwich of the type that he purports to prefer. Now consider two possible strategies, σ_{truthful} and σ_{random}. σ_{truthful} corresponds to truthfully reporting sandwich preferences (and subsequently leads to eating the preferred sandwich type), while σ_{random} randomly reports independent of true type (and results in the preferred sandwich only half the time). The cost of using the random strategy is $I(\mathcal{T}; \mathcal{M}(t_{-i}, \sigma_{\text{random}}(\mathcal{T})) = 0$, since the output is independent of agent i's type. On the other hand, the cost of truthfully reporting is $I(\mathcal{T}; \mathcal{M}(t_{-i}, \sigma_{\text{truthful}}(\mathcal{T})) = 1$, since the sandwich outcome is now the identity function on agent i's type. However, from the perspective of any

[19] The seminal work of [Xiao 2013] was the first to explore mechanism design with agents who have costs for privacy loss. It proposes several measures of privacy cost, but mutual information was the first, and helped drive subsequent work.

outside observer, the two strategies are indistinguishable! In both cases, agent i receives a uniformly random sandwich. Why then should anyone choose the random strategy? So long as an adversary *believes* they are choosing randomly, they should choose the honest strategy.

[Chen et al. 2013] propose a different approach, not needing a prior on agent types. They propose the following cost function:[20]

$$|c_i(o, \mathcal{M}, t)| = \ln\left(\max_{t_i, t_i' \in \mathcal{T}} \frac{\Pr[\mathcal{M}(t_i, t_{-i}) = o]}{\Pr[\mathcal{M}(t_i', t_{-i}) = o]}\right).$$

Note that if $\mathcal{M}$ is ϵ-differentially private, then

$$\max_{t \in \mathcal{T}^n} \max_{o \in \mathcal{O}} \max_{t_i, t_i' \in \mathcal{T}} \ln\left(\frac{\Pr[\mathcal{M}(t_i, t_{-i}) = o]}{\Pr[\mathcal{M}(t_i', t_{-i}) = o]}\right) \leq \epsilon.$$

That is, we can view differential privacy as bounding the *worst-case* privacy loss over all possible outcomes, whereas the measure proposed by [Chen et al. 2013] considers only the privacy loss for the outcome o (and type vector t) actually realized. Thus, for any differentially private mechanism $\mathcal{M}$, $|c_i(o, \mathcal{M}, t)| \leq \epsilon$ for all o, t, but it will be important that the cost can vary by outcome.

[Chen et al. 2013] then consider the following allocation rule for maximizing social welfare $F(o) = \sum_{i=1}^{n} u_i(o)$.[21] We discuss the case when $|\mathcal{O}| = 2$ (which does not require payments), but the authors analyze the general case (with payments), which privately implements the VCG mechanism for any social choice problem.

(1) For each outcome $o \in \mathcal{O}$, choose a random number r_o from the distribution $\Pr[r_o = x] \propto \exp(-\epsilon|x|)$.
(2) Output $o^* = \arg\max_{o \in \mathcal{O}}(F(o) + r_o)$

[Chen et al. 2013] show that the above mechanism is ϵ-differentially private, and that it is truthful for privacy aware agents, so long as for each agent i, and for the two outcomes $o, o' \in \mathcal{O}$, $|\mu_i(o) - \mu_i(o')| > 2\epsilon$. Note that this will be true for small enough ϵ so long as agent utilities for outcomes are distinct. The analysis is elegant, and proceeds by considering an arbitrary fixed realization of the random variables r_o, and an arbitrary deviation t_i' from truthful reporting for the i'th agent. There are two cases: In the first case, the deviation does not change the outcome o of the mechanism. In this case, *neither* the agent's utility for the outcome μ_i, nor his cost for privacy loss c_i change at all, and so the agent does not benefit from deviating. In the second case, if the outcome changes from o to o' when agent i deviates, it must be that $\mu_i(o') < \mu_i(o) - 2\epsilon$. By differential privacy, however, $|c_i(o, \mathcal{M}, t) - c_i(o', \mathcal{M}, t)| \leq 2\epsilon$, and so the change in privacy cost cannot be enough to make it beneficial.

Finally, [Nissim et al. 012a] take the most conservative approach to modeling costs for privacy. Given an ϵ-differentially private mechanism $\mathcal{M}$, they assume that

$$c_i(o, \mathcal{M}, t) \leq \epsilon v_i,$$

[20] In fact, the upper bound they propose is more general, replacing the $\ln(\cdot)$ used here with any function F_i satisfying certain conditions.
[21] This allocation rule is extremely similar to, and indeed can be modified to be identical to the exponential mechanism.

for some number v_i. This is similar to the linear cost functions of [Ghosh and Roth 2011] that we considered earlier, but crucially, [Nissim et al. 012a] assume only an upper bound. This assumption is satisfied by all of the other models for privacy cost that we have considered thus far. They show that many mechanisms that combine a differentially private algorithm with a punishing mechanism that has the ability to restrict user choices, like those from [Nissim et al. 012b] that we considered in Section 3.3, maintain their truthfulness properties in the presence of agents with preferences for privacy, so long as the values v_i are bounded. Moreover, they go on to show that for a great many distributions from which the values v_i might be drawn (even with *unbounded support*), it is still possible to make truthful reporting a dominant strategy for *almost* all agents, which is often sufficient to get strong welfare guarantees.

6. CONCLUSIONS AND OPEN QUESTIONS

The science of privacy, and its burgeoning connections to game theory and mechanism design are still in their very early stages. As such, there remain more questions than answers at this intersection, but there is already more than enough evidence that the area is rich, and that the answers will be fascinating. We here suggest just a couple of the high-level questions that deserve to be better understood:

(1) Why do people care about privacy? In Section 4 we began to see some answers to this question: people might care about privacy, because of some *repeated interaction*, in which the revelation of their type at one stage of the interaction might cause them quantifiable harm at some later stage of the game. This literature has already provided us with some insights, but only at a very coarse-grained level. The recent literature on differential privacy lets us discuss privacy as a quantitative, rather than just qualitative property. Can we formulate a rich model in which we can directly derive agents' values for (differential) privacy in various settings? This study will surely help us understand how we should in general model agent costs for privacy when trying to design truthful mechanisms for agents who explicitly experience costs for privacy in their utility functions, as we considered in Section 5.

(2) How should we model privacy costs that are not captured by information theoretic measures like differential privacy, but nevertheless seem to have real economic consequences? In particular, how can we model the fact that *search costs* appear to factor into peoples' perceptions of privacy? For example, there has been a recent uproar about Facebook's new Graph Search tool.[22] From the perspective of differential privacy, the addition of this feature, which merely eases the search through information that was already *in principle* publicly available, should have had no incremental privacy cost. Nevertheless, it plainly does. Is there a clean economic model that captures this?

(3) In settings in which agents care about privacy, to what extent does dominant strategy truthfulness remain the right solution concept? Depending on the privacy model chosen, it is no longer clear that the revelation principle holds

[22]See e.g. `http://www.slate.com/blogs/future_tense/2013/01/23/actual_facebook_graph_searches_tom_scott_s_tumblr_a_privacy_wake_up_call.html`

in such settings, because agents have preferences not only over the outcome chosen by the mechanism, but also over the mechanism which chooses the outcome itself. Perhaps in such settings, *better* outcomes can be implemented in Nash equilibrium than can be implemented in dominant strategies?

(4) Differential privacy is clearly a powerful tool for reasoning about noise and stability in mechanism design. Already in Section 3, we saw several examples of results easily derived via differential privacy, which were are not known how to accomplish in any other way. It seems like a particularly promising tool for reasoning about asymptotic truthfulness, which is a compelling second-best solution concept when exactly dominant strategy truthful mechanisms are not known. Can we use differential privacy to map out the power of asymptotically truthful mechanisms? How much more powerful is this class as compared to the set of exactly truthful mechanisms?

(5) In several settings of applied interest, a regulator may wish to reveal summary information about the industry she regulates to guide policy and alleviate information asymmetries. In doing this, she must balance the privacy of individual participants, which arise from competitive concerns, trade secrets etc. Can differential privacy and its variants assist in the design of such information release? How should policy be designed given that the summary used to guide it is privacy-preserving and therefore necessarily coarse? [Flood et al. 2013] propose and study these concerns in the context of the financial industry, and provide an excellent overview of the trade-offs involved.

REFERENCES

Acquisti, A. and Varian, H. R. 2005. Conditioning prices on purchase history. *Marketing Science 24,* 3, 367–381.

Balcan, M.-F., Blum, A., Hartline, J. D., and Mansour, Y. 2005. Mechanism design via machine learning. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*. IEEE, 605–614.

Bergemann, D., Brooks, B., and Morris, S. 2013. The limits of price discrimination. Tech. rep., Cowles Foundation Working Paper, Yale University.

Calzolari, G. and Pavan, A. 2006. On the optimality of privacy in sequential contracting. *Journal of Economic Theory 130,* 1, 168–204.

Chan, J. and Eyster, E. 2003. Does banning affirmative action lower college student quality? *The American Economic Review 93,* 3, 858–872.

Chen, Y., Chong, S., Kash, I. A., Moran, T., and Vadhan, S. P. 2013. Truthful mechanisms for agents that value privacy. *ACN Conference on Electronic Commerce*.

Conitzer, V., Taylor, C. R., and Wagman, L. 2012. Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science 31,* 2, 277–292.

Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ACM, 214–226.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. 2006. Calibrating noise to sensitivity in private data analysis. In *TCC '06*. 265–284.

Dwork, C. and Roth, A. 2013. *The Algorithmic Foundations of Differential Privacy*.

Dziuda, W. and Gradwohl, R. 2012. Achieving coordination under privacy cncerns. Tech. rep., Working paper, Northwestern University.

Fleischer, L. and Lyu, Y.-H. 2012. Approximately optimal auctions for selling privacy when costs are correlated with data. In *ACM Conference on Electronic Commerce*. 568–585.

Flood, M., Katz, J., Ong, S., and Smith, A. 2013. Crptography and the economics of supervisory information: Balancing transparency and confidentiality. Tech. rep.

Ghosh, A. and Roth, A. 2011. Selling privacy at auction. In *ACM Conference on Electronic Commerce*. 199–208.

Gradwohl, R. 2012. Privacy in implementation. Tech. rep., Northwestern University.

Huang, Z. and Kannan, S. 2012. The exponential mechanism for social welfare: Private, truthful, and nearly optimal. In *IEEE 53rd Annual Symposium on the Foundations of Computer Science (FOCS)*. IEEE, 140–149.

Kearns, M., Pai, M. M., Roth, A., and Ullman, J. 2012. Mechanism design in large games: Incentives and privacy. *arXiv preprint arXiv:1207.4084*.

Ligett, K. and Roth, A. 2012. Take it or leave it: Running a survey when privacy comes at a cost. In *Internet and Network Economics*. Springer, 378–391.

McSherry, F. and Talwar, K. 2007. Mechanism design via differential privacy. In *FOCS*. 94–103.

Nissim, K., Orlandi, C., and Smorodinsky, R. 2012a. Privacy-aware mechanism design. In *ACM Conference on Electronic Commerce*. 774–789.

Nissim, K., Smorodinsky, R., and Tennenholtz, M. 2012b. Approximately optimal mechanism design via differential privacy. In *ITCS*. 203–213.

Roth, A. 2012. Buying private data at auction: the sensitive surveyor's problem. *ACM SIGecom Exchanges 11,* 1, 1–8.

Roth, A. and Schoenebeck, G. 2012. Conducting truthful surveys, cheaply. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 826–843.

Taylor, C. R. 2004. Consumer privacy and the market for customer information. *RAND Journal of Economics*, 631–650.

Xiao, D. 2013. Is privacy compatible with truthfulness? In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*. ACM, 67–86.

Setting Equilibrium Prices, Approximately

BRENDAN LUCIER
Microsoft Research New England

We outline recent results, positive and negative, on pricing indivisible goods to approximately maximize social welfare. We describe a relaxation of standard pricing problems in which a seller can bundle goods together prior to sale.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Combinatorial auctions; Walrasian equilibrium; Envy-free

Consider the plight of a retail clothing store owner who must decide how best to sell his wares. The owner takes pride in his store, and wants above all else to distribute his inventory to maximize the welfare of his customers. In principle, the owner could consult with each customer about his or her fashion needs, then ultimately propose a wardrobe and price to each one. While such a business model might be feasible for a super-elite fashion outlet, a more practical (and common) approach is to post prices on individual items and let the customers choose what to purchase at the specified prices. The seller's problem then becomes one of selecting appropriate prices.

In this letter we discuss this classic equilibrium pricing problem, using the clothing store as a running example. We will review some well-known impossibility results, then turn to approximation methods as a way to circumvent them.

First, we must establish what it means to set "appropriate" prices. The common wisdom from freshman-level economics is that prices should equate supply and demand. However, this task is not straightforward in the clothing store setting, where there are many heterogeneous goods and buyers can have idiosyncratic preferences. More concretely, suppose that each item of clothing is an indivisible object, and that each customer has a certain value for any given set of clothes. We will imagine that the store owner is so experienced that he knows the values of all of his customers, and that the customers are rational enough to make the best clothing selection for themselves when faced with a set of prices. Even under these favorable conditions, it is still not immediately obvious how prices should be chosen.

A formal study of such market equilibria goes as far back as Walras in 1874; [Shapley and Shubik 1971] provides a modern market-pricing interpretation. The Walrasian solution concept requires that prices satisfy the following conditions:

(1) Each buyer can obtain his most-demanded set at the given prices, and
(2) Every item with a positive price is sold.

These two conditions define a *Walrasian equilibrium* (WE). The first is a notion of

Author's address: brlucier@microsoft.com

envy-freeness: no customer will envy another for getting the very last pair of pants in the store. The second condition concerns the seller, who should not be left with an unsold suit jacket that a customer would have purchased at a lower price.

A wonderful feature of a Walrasian equilibrium is that it necessarily results in a socially optimal allocation [Bikhchandani and Mamer 1997], as desired by the store owner. Unfortunately for the owner, WE do not always exist. As a simple example, suppose the store has a pair of shoes for sale and that each shoe is an item. Further imagine that two customers, Alice and Bob, are interested in shoes. Alice, a typical customer, values the pair of shoes at \$10 but has no value for either shoe on its own. Bob, on the other hand, wants only a single shoe and will pay up to \$6 for either the left or the right. In this situation, no vector of prices can satisfy the WE conditions: if Alice buys both shoes, then at least one shoe must have price less than \$6 and is therefore also desired by Bob, violating the first condition of WE. If instead Bob gets a shoe for \$6 or less, then the other shoe must be left unsold. The second WE condition implies this unsold shoe must have price \$0, but then Alice would want to buy the pair.

It is perhaps unsurprising that item prices fail to resolve markets with complementary items (like left and right shoes), but the issue extends even further. [Kelso and Crawford 1982] show that a WE always exists when the buyer preferences satisfy a certain gross substitutes condition, which is strictly stronger than submodularity. However, [Gul and Stacchetti 1999] show that this is as far as the existence guarantee can be extended, in the following natural sense. If there is even one customer with values that do not satisfy the gross substitutes condition, a WE can fail to exist. Thus, even for items that have decreasing marginal value, there may not exist item prices that implement a market equilibrium.

Approximations and Unsold Items

Faced with these classic and discouraging impossibility results, we propose relaxing the market-clearing requirements of WE. Specifically, we will drop the second WE condition, allowing items to remain unsold, and require only that customers can buy their most desired sets. This certainly addresses the existence problem, since the seller could set prices so high that no customer desires anything in the store. Of course, this is a very unsatisfactory solution from the perspective of maximizing social welfare. We therefore turn to the paradigm of approximation and ask whether there exist prices that result in approximately optimal social welfare.

It is not too hard to see that the answer is *no* for settings with complements: consider extending the shoe example so that Alice gets value only if she buys many, many items together. Nevertheless, our expectation was that, following recent results in the theory of simultaneous item auctions [Christodoulou et al. 2008; Feldman et al. 2013a], constant factor approximations would be forthcoming when customer preferences are complement-free (or, at least, fractionally subadditive; see [Feige and Vondrak 2006]). However, much to our surprise, this is not the case. There exist simple examples without complements for which item prices fail to achieve non-trivial approximations to the economically efficient outcome. We now describe such an example, from [Feldman et al. 2013b].

In this example, two customers Alex and Betty wish to purchase shirts. Alex

is not particular about shirts, but could use many of them: he would value any number of shirts at \$1 each, regardless of style. However, Alex also has a strong preference for not leaving the store empty-handed, and therefore has a value of \$2 for getting any one shirt. In other words, Alex values any set of $k > 1$ shirts at $\$k$, and values any single shirt at \$2. Betty is also completely indifferent about style, but is only interested in buying at most one shirt: her value is slightly less than \$1 for any single shirt, where the precise value will be described below. It is easily checked that both Alex's and Betty's value functions are subadditive (in fact, they satisfy the stronger condition of fractional subadditivity).

To understand the store owner's pricing problem, consider what would happen if Alex were the only customer. The store owner certainly wants Alex to buy as many shirts as possible. It is tempting to set a price of \$1 on every shirt. At these prices, Alex could indeed buy many shirts for an overall utility (value minus price) of \$0, but he prefers to buy a single shirt for a gain of \$1. Thus, in order to convince Alex to buy all of the shirts in the store, the retailer must set a price less than \$1 on some of them. For example, if the store has m shirts, a price of $\$\left(1 - \frac{2}{m}\right)$ on each shirt would convince Alex to buy them all.

This solution falls apart, however, when we include Betty as a customer. If Betty's value for a shirt is sufficiently close to \$1 — say, $\$\left(1 - \frac{1}{m}\right)$ — she competes with Alex for the cheapest shirt. Indeed, Alex only ever buys multiple shirts if their average price is less than Betty's value, so he and Betty must have conflicting demands. The only way to resolve this tension is to allocate at most one shirt to Alex — say, by setting a price of \$2 on every shirt — which catastrophically reduces the social welfare from $\$m$ to at most a constant.

The Power of Bundling

The example above illustrates that allowing items to go unsold is not enough to enable reasonable approximations. We must therefore consider a further relaxation. To motivate our approach, let us return to the example of selling shoes to Alice and Bob. It is common for stores to offer shoes for sale, and the sort of problem present in our example is certainly not insurmountable in practice. Allowing shoes to remain unsold is not the natural approach — quite the opposite! The more obvious solution is to commit to selling shoes only in pairs. This effectively defines a pair of shoes to be an indivisible object. This is a natural power to afford the seller; after all, as the owner of the goods, he is free to package them however he wishes. Furthermore, this type of bundling does not increase the complexity of the marketplace, since the seller is still posting prices on individual, indivisible items. In fact, the market is arguably simpler, since the owner can only reduce the number of items for sale by pairing shoes.

More generally, we could allow the owner of our clothing store to partition his goods into packages however he sees fit, then treat those packages as the items for sale and set their prices. The goal, then, is to find a partition *and* a set of prices so that the buyer-side envy-freeness condition is met. This solution concept was proposed in [Feldman et al. 2013b], where it was termed Combinatorial Walrasian equilibrium. As it so happens, this additional bundling operation is just what we are looking for: in [Feldman et al. 2013b] we show that, for *arbitrary* customer

preferences, there always exists a partition of the objects and a pricing such that the resulting outcome achieves at least half of the optimal social welfare.

To illustrate this approach, let us return to the example of Alex and Betty buying shirts. In this example, the correct choice for the store owner is to commit to selling all of his stock as a single bundle, for a price of $\$m$. Betty is unwilling to buy at such a high price, and Alex is not tempted to buy only a single shirt (since individual shirts are not for sale). The optimal allocation, where Alex buys everything in the store, becomes the only possible outcome. Effectively, the presence of the insatiable Alex compels the store owner to switch to a wholesale business model.

REFERENCES

Bikhchandani, S. and Mamer, J. W. 1997. Competitive equilibrium in an exchange economy with indivisibilities. *J. of Economic Theory 74,* 2, 385–413.

Christodoulou, G., Kovács, A., and Schapira, M. 2008. Bayesian combinatorial auctions. In *ICALP*. 820–832.

Feige, U. and Vondrak, J. 2006. Approximation algorithms for allocation problems: Improving the factor of 1 - 1/e. In *FOCS*. 667–676.

Feldman, M., Fu, H., Gravin, N., and Lucier, B. 2013a. Simultaneous auctions are (almost) efficient. In *STOC*. To appear.

Feldman, M., Gravin, N., and Lucier, B. 2013b. Combinatorial walrasian equilibrium. In *STOC*. To appear.

Gul, F. and Stacchetti, E. 1999. Walrasian equilibrium with gross substitutes. *J. of Economic Theory 87,* 1, 95–124.

Kelso, Alexander S, J. and Crawford, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica 50,* 6, 1483–1504.

Shapley, L. S. and Shubik, M. 1971. The assignment game i: The core. *International J. of Game Theory 1*, 111–130.

Logit Dynamics: A Model for Bounded Rationality

DIODATO FERRAIOLI
Université Paris Dauphine

We describe logit dynamics, which are used to model bounded rationality in games, and their related equilibrium concept, the logit equilibrium. We also present some results about the convergence time of these dynamics and introduce a suitable approximation of the logit equilibrium. We conclude by describing some interesting future extensions to logit dynamics.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Theory, Economics, Performance

Additional Key Words and Phrases: Bounded Rationality, Equilibrium Concept, Game Dynamics

1. INTRODUCTION

Classical Game Theory assumes that agents have complete knowledge of the game (they know all the players, their set of strategies and their utility functions and they know that other players know, and they know that others know that they know, etc.). It also assumes the players have unlimited computational power to select a strategy that maximizes their utility given the strategies played by other players. However, in many cases players' decisions can be influenced by limited knowledge and limited computational capabilities. Thus, to have a more precise description of phenomena emerging in these settings we need a model that can capture the behavior of agents with *bounded rationality* and that may sometimes make wrong decisions.

An example of such a model is the *logit update rule* [McFadden 1974]. According to this rule, players update their strategies with respect to a parameter β (that represents the level of rationality or knowledge) and the state of the system (i.e., the strategies currently played by the players). Roughly speaking, the logit update rule can be seen as a noisy version of the classical best response update rule, where β is the bias towards choices that are good. A small β represents the situation where players are subject to strong noise or they have very limited knowledge of the game and, therefore, choose their strategies "nearly at random"; a large β, instead, represents the situation where players "almost surely" play the best response.

Blume [1993] introduced the logit update rule in game dynamics through logit dynamics. We will describe these dynamics and discuss some of their properties.

2. LOGIT DYNAMICS AND LOGIT EQUILIBRIA

Consider a strategic game $\mathcal{G} = ([n], S_1, \ldots, S_n, u_1, \ldots, u_n)$, where $[n] = \{1, \ldots, n\}$ is a finite set of players, S_i is the finite set of strategies for player $i \in [n]$, $S = S_1 \times \cdots \times S_n$ is the set of strategy profiles and $u_i \colon S \to \mathbb{R}$ is the utility function of player $i \in [n]$. The *logit dynamics* for a game $\mathcal{G}$ proceed as follows: at each time

Authors' address: diodato.ferraioli@dauphine.fr

step (i) Select one player $i \in [n]$ uniformly at random; (ii) Update the strategy of player i according to the logit update rule with parameter $\beta > 0$ over the set S_i of her strategies. That is, the dynamics select a strategy $s \in S_i$ with probability $\sigma_i(s \mid \mathbf{x}) = e^{\beta u_i(s,\mathbf{x}_{-i})}/Z_i(\mathbf{x})$, where $\mathbf{x} = (x_1, \ldots, x_n) \in S$ is the current strategy profile and $Z_i(\mathbf{x}) = \sum_{z \in S_i} e^{\beta u_i(z,\mathbf{x}_{-i})}$ is the normalizing factor.

These dynamics have been extensively adopted in Economics and, recently, in Computer Science to model the spread of innovation in social networks [Ellison 1993; Young 2000; Montanari and Saberi 2009]. These works focused on the time the dynamics take for hitting a specific pure Nash equilibrium of the game. However, Nash equilibria may not be an adequate solution concept for these dynamics. Indeed, there is always a chance, which is inversely proportional to the rationality level β, that players deviate from these strategy profiles.

To address this issue, Auletta et al. [2010] introduce a new equilibrium concept for logit dynamics, named *logit equilibrium*, describing the long-run behavior of the system (which states appear more frequently in the long run) and defined by a probability distribution over the pure strategy profiles of the game. In fact, it is easy to see that the logit dynamics for a game $\mathcal{G}$ define a Markov chain with the set S of strategy profiles as state space. Then, the logit equilibrium is defined as the stationary distribution of this Markov chain, i.e. the distribution π such that $\pi P = \pi$, where P is the transition matrix of the Markov chain. It is not hard to see that the chain defined by the logit dynamics is ergodic and, hence, any strategic game possesses a logit equilibrium and this is unique. We note that the absence of either of these guarantees is often considered a weakness of pure Nash equilibria.

If the time the dynamics take to reach an equilibrium is long, then the system spends most of its life outside of the equilibrium and thus the relevance of this concept is almost completely lost. For this reason, it is important to bound the time that the dynamics takes to reach either the equilibrium or some suitable approximations of it. The next section provides some results in this direction.

3. CONVERGENCE TIME AND METASTABLE DISTRIBUTIONS

Auletta et al. [2012a] give general bounds on the convergence time of logit dynamics for wide classes of games. These results show that there are games for which the convergence time is bounded from above by a polynomial in the number of players, and by an exponential function in the rationality level and in some structural properties of the game. Thus, as β increases (i.e., players become more rational), the dynamics take longer to reach the equilibrium. Indeed, for high β players tend to play their best response and the system is likely to remain in a pure Nash equilibrium (if any) for a long time. This behavior slows down the convergence to the logit equilibrium, whenever the stationary distribution assigns high probability to other profiles (e.g., this is the case if there are other similar Nash equilibria). Auletta et al. [2012a] also show games for which the convergence time is bounded by a function independent of β. Unfortunately, this function can be exponential in the number of players.

Since the convergence time of the dynamics can be large, it becomes interesting to understand what happens during the transient phase of the dynamics. Is this phase completely chaotic, or can we still spot some regularities? Are we able to describe

the behavior of the system even before stationarity has been reached? Is it possible that on a timescale shorter than the convergence time the chain is "metastable", i.e., it stays close to some subset of the state space, while in a timescale comparable to the convergence time it jumps from one metastable configuration to another?

In order to answer these questions, Auletta et al. [2012] introduce the definition of *metastable distribution*. Roughly speaking, a distribution μ is (ε, T)-metastable for the dynamics if, selecting the starting profile according to μ, the dynamics stays at distance at most ε from μ for at least T steps. Metastable distributions can be seen as approximations of the stationary distribution. Indeed, they remain stable for a time which is long enough for an observer (in computer science terms, we assume this time is super-polynomial), while the stationary distribution remains stable forever. The hope is that whenever there is a starting profile from which the convergence to the logit equilibrium takes too long, then, from that profile, the dynamics should converge quickly to a metastable distribution. Ferraioli and Ventre [2012] take the first steps in this direction.

4. FUTURE DIRECTIONS

Logit dynamics assume that only one player updates her strategy at any time. It would be interesting then to model also players concurrently updating their strategies. Auletta et al. [2012b] give preliminary results in this direction.

Another implicit assumption of logit dynamics is that all players have the same rationality level. However, it would be interesting to extend the analysis of logit dynamics by taking into account that different players may have different levels of rationality depending on different personal attitudes.

Last but not least, it would be extremely interesting to understand in which way we can influence the evolution of a game whose agents are not fully rational, in order to push the system towards desired directions.

REFERENCES

Auletta, V., Ferraioli, D., Pasquale, F., Penna, P., and Persiano, G. 2012a. Convergence to equilibrium of logit dynamics for strategic games. *CoRR abs/1212.1884*. Preliminary version appeared in SPAA 2011.

Auletta, V., Ferraioli, D., Pasquale, F., Penna, P., and Persiano, G. 2012b. Reversibility and mixing time for logit dynamics with concurrent updates. *CoRR abs/1207.2908*.

Auletta, V., Ferraioli, D., Pasquale, F., and Persiano, G. 2010. Mixing time and stationary expected social welfare of logit dynamics. In *Proc. of the 3rd Int. Symp. on Algorithmic Game Theory (SAGT'10)*. LNCS, vol. 6386. Springer, 54–65.

Auletta, V., Ferraioli, D., Pasquale, F., and Persiano, G. 2012. Metastability of logit dynamics for coordination games. In *Proc. of the ACM-SIAM Symp. on Discrete Algorithms (SODA'12)*. SIAM, 1006–1024.

Blume, L. E. 1993. The statistical mechanics of strategic interaction. *Games and Economic Behavior 5*, 387–424.

Ellison, G. 1993. Learning, local interaction, and coordination. *Econometrica 61,* 5, 1047–1071.

Ferraioli, D. and Ventre, C. 2012. Metastability of potential games. *CoRR abs/1211.2696*.

McFadden, D. L. 1974. Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*. Academic Press: New York, 105–142.

Montanari, A. and Saberi, A. 2009. Convergence to equilibrium in local interaction games. In *Proc. of the 50th Ann. Symp. on Foundations of Computer Science (FOCS'09)*. IEEE.

Young, H. P. 2000. The diffusion of innovations in social networks. Economics Working Paper Archive number 437, Johns Hopkins University, Department of Economics.

Planning and Learning in Security Games

FRANCESCO M. DELLE FAVE, YUNDI QIAN, ALBERT X. JIANG,
MATTHEW BROWN, and MILIND TAMBE
University of Southern California

We present two new critical domains where security games are applied to generate randomized patrol schedules. For each setting, we present the current research that we have produced. We then propose two new challenges to build accurate schedules that can be deployed effectively in the real world. The first is a planning challenge. Current schedules cannot handle interruptions. Thus, more expressive models, that allow for reasoning over stochastic actions, are needed. The second is a learning challenge. In several security domains, data can be used to extract information about both the environment and the attacker. This information can then be used to improve the defender's strategies.

Categories and Subject Descriptors: I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search

General Terms: Algorithms, Experimentation, Security; Theory

Additional Key Words and Phrases: Artificial Intelligence, Game Theory

1. INTRODUCTION

In recent years, research in security games has produced a number of approaches that led to the deployment of real world applications for protecting critical infrastructure such as ports, airports and trains [Tambe 2011; Conitzer and Sandholm 2006]. In essence, three types of algorithms were developed: (i) scalable algorithms [Jain et al. 2011; Jain et al. 2013]; (ii) algorithms for games with boundedly rational attackers [Yang et al. 2011; Nguyen et al. 2013]; and (iii) algorithms robust against execution and observation uncertainty [Yin et al. 2011; Yin and Tambe 2012].

With an eye to the future, the deployment of these algorithms in the real world introduces two new important challenges. The first is a planning challenge. As will be described later, several security domains involve significant uncertainty. Hence, the schedules in these domains are often interrupted due to unexpected events. Furthermore, in the real world, a schedule's spatial and temporal constraints are typically continuous dimensions. Hence, more expressive models need to be derived to represent such schedules. In particular, models that allow for reasoning over stochastic actions and continuous dimensions.

The second challenge is a learning challenge. Thus far, data available in most domains has been principally used to define the game's matrices. However, in the newer domains, the planned schedules are frequently interrupted. Hence, the information about the locations and times of interruptions as well as the information about the actual interactions between attackers and defenders can also be used. The former can be used to represent the uncertainty of the environment, whereas the latter can be used to learn the attacker's behavior.

Authors' addresses: {dellefav,yundi.qian,jiangx,matthew.brown,tambe}@usc.edu

Against this background, we present in this letter our initial attempt to address the challenges of planning and learning in the context of security games. In Section 2, we introduce the key concepts related to solving patrolling problems using security games. In Section 3, we present two critical real world problems in which addressing the challenges of planning and learning is the key priority to generate effective schedules. Finally, in Section 4, we discuss some future work and some possible research directions.

2. STACKELBERG GAMES FOR SECURITY PROBLEMS

Protecting critical infrastructure is a challenging task for police and security agencies around the world. Areas such as ports, airports, historical landmarks or locations of political and economic importance are key targets for illegal activities. Example of such activities include fare evasion, burglary and strategic terrorism. Unfortunately, the number of resources available to patrol these domains is typically limited. In addition, adversaries such as terrorists or criminals, will monitor any type of patrolling activity to find and exploit predictable patterns.

To address these shortcomings, game theory provides a method to allocate limited security resources in a selective and randomized fashion. The idea is to cast each problem as a security game, a specialization of a Bayesian Stackelberg game, where a defender (i.e., a security agency) and an attacker (i.e., a terrorist or a criminal) compete over the protection of a number of targets (e.g. buses, trains, forest or marine reserves). In essence, in a security game, if an attacker attacks a target that was covered (protected) by the defender, then the attacker has a worse payoff than if the attacker had attacked the same target when it was not covered. Given these properties, research in security games has produced a number of approaches which are currently being used in several ports and airports of the United States (see [Tambe 2011] for more details).

3. NEW DOMAINS IN SECURITY GAMES

This section presents two different patrolling domains and discusses the way in which they are modeled as security games.

3.1 Patrolling the Los Angeles Metro System

The Los Angeles metro system is a barrier-free transit system. For this reason, fare evasion is a well acknowledged problem for the Los Angeles Sheriff's Department (LASD), the agency responsible for its security. The TRUSTS system was developed to deter such fare evasion and proceeds in two phases [Yin et al. 2012]. First, the spatial and temporal constraints of the problem (e.g., patrol length and train schedules) are encapsulated within a transition graph. Second, the transition graph is used to define a Bayesian Stackelberg game between one defender (the LASD) and multiple types of attackers (the fare evaders).

Schedules produced by the system were deployed on different real world trials. The results showed that the model lacked in accuracy. Indeed, as discussed in Section 1, the schedules were often interrupted due to the uncertainty related to patrolling a metro system. For instance, officers out on patrol would need to arrest an unruly passenger or help out a lost tourist, throwing them off the carefully

constructed schedule. In light of this, the original framework has been recently extended to incorporate uncertainty [Jiang et al. 2013]. The idea is to generalize the transition graph (the first phase) to a Markov decision process (MDP), thus addressing the planning challenge discussed in Section 1. This MDP is then used to define a Bayesian Stackelberg following the second phase of the TRUSTS system described earlier. A pure strategy is produced by solving the game and sampling the randomized strategy. Within this setting, however, the approach produces a Markov strategy which corresponds to a mapping from states to actions. To visualize a schedule then, a mobile application has been developed and is currently being evaluated by the LASD.

3.2 Marine Resources and Forest Protection

The oceans and forests of the world provide a variety of vital natural resources, such as fish and fuelwood, whose unregulated extraction and consumption have become a key concern for security agencies around the world. As a consequence, the intelligent and unpredictable patrolling of these reserves has become a key research challenge within both these domains and security games have been considered as a promising solution concept. Thus far, research has progressed principally in the forest domain. Specifically, the problem is cast as a security game where the forest area is represented as a circular continuous space. Hence, as discussed in Section 1, solving this game requires reasoning over continuous dimensions. An optimal strategy then corresponds to a band patrol capable of maximizing the fully protected pristine area of the forest [Johnson et al. 2012]. In contrast, work on the fish domain has just started. The most interesting feature of this domain is the availability of significant amounts of data on the interactions with the attacker, which, as discussed in the previous section, can be used to improve the defender's strategies.

4. FUTURE CHALLENGES AND RESEARCH DIRECTIONS

The challenges of planning and learning pertain to all the domains discussed in Section 3. In all domains, schedules might be interrupted due to some unexpected events (e.g., writing a citation in a train line, boarding an illegal fisherman's boat). As a consequence, incorporating planning within security games is necessary to model stochastic decision making. In so doing, several issues need to be investigated: new game models and solutions need to be designed and the computational effort required to solve them needs to be evaluated.

Additionally, the introduction of a learning component in security games is necessary to exploit the available domain information. The idea is to extract information from the available data and use it to improve the quality and the effectiveness of the defender's strategies. As a consequence, several key learning challenges such as the amount of data necessary or the computational effort required to learn this information, become relevant to this research. We will consider these challenges in our future work.

REFERENCES

Conitzer, V. and Sandholm, T. 2006. Computing the optimal strategy to commit to. In *Conference on Electronic Commerce (EC)*.

Jain, M., Conitzer, V., and Tambe, M. 2013. Security scheduling for real-world networks. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Jain, M., Tambe, M., and Kiekintveld, C. 2011. Quality-bounded solutions for finite bayesian stackelberg games: Scaling up. In *Int. Conf. on Autonomous Agents and Multiagent Systems*.

Jiang, A. X., Yin, Z., Zhang, C., Tambe, M., and Kraus, S. 2013. Game-theoretic randomization for security patrolling with dynamic execution uncertainty. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Johnson, M. P., Fang, F., , and Tambe, M. 2012. Patrol strategies to maximize pristine forest area. In *Conference on Artificial Intelligence (AAAI)*.

Nguyen, T. H., Yang, R., Azaria, A., Kraus, S., and Tambe, M. 2013. Analyzing the effectiveness of adversary modeling in security games. In *Conf. on Artificial Intelligence (AAAI)*.

Tambe, M. 2011. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press.

Yang, R., Kiekintveld, C., Ordonez, F., Tambe, M., and John, R. 2011. Improving resource allocation strategy against human adversaries in security games. In *International Joint Conference on Artificial Intelligence (IJCAI)*.

Yin, Z., Jain, M., Tambe, M., and Ordonez, F. 2011. Risk-averse strategies for security games with execution and observational uncertainty. In *Conference on Artificial Intelligence (AAAI)*.

Yin, Z., Jiang, A., Johnson, M., Tambe, M., Kiekintveld, C., Leyton-Brown, K., Sandholm, T., and Sullivan, J. 2012. Trusts: Scheduling randomized patrols for fare inspection in transit systems. In *Conference on Innovative Applications of Artificial Intelligence (IAAI)*.

Yin, Z. and Tambe, M. 2012. A unified method for handling discrete and continuous uncertainty in Bayesian stackelberg games. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Generalized Scoring Rules: A Framework That Reconciles Borda and Condorcet

LIRONG XIA
Harvard University

Generalized scoring rules [Xia and Conitzer 08] are a relatively new class of social choice mechanisms. In this paper, we survey developments in generalized scoring rules, showing that they provide a fruitful framework to obtain general results, and also reconcile the Borda approach and Condorcet approach via a new social choice axiom. We comment on some high-level ideas behind GSRs and their connection to Machine Learning, and point out some ongoing work and future directions.

Categories and Subject Descriptors: J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*; I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Computational social choice, generalized scoring rules

1. INTRODUCTION

Social choice theory focuses on developing principles and methods for representation and aggregation of individual ordinal preferences. Perhaps the most well-known application of social choice theory is political elections. Over centuries, many social choice mechanisms have been proposed and analyzed in the context of elections, where each agent (voter) uses a linear order over the alternatives (candidates) to represent her preferences (her *vote*). For historical reasons, we will use *voting rules* to denote social choice mechanisms, though we need to keep in mind that the application is not limited to political elections.[1] Most existing voting rules fall into one of the following two categories.[2]

Positional scoring rules: Each alternative gets some points from each agent according to its position in the agent's vote. The alternative with the highest total points wins. For example, *Borda* is a positional scoring rule where for each vote, the alternative ranked at the ith position gets $m - i$ points, where m is the number of alternatives.

Condorcet consistent rules: Whenever there exists a *Condorcet winner*, it must be the unique winner of the election. A Condorcet winner is an alternative that beats every other alternative in head-to-head comparisons. For example,

[1]More recently, social choice theory has been adopted in many modern computational systems, including but not limited to recommender systems [Ghosh et al. 1999], meta-search engines [Dwork et al. 2001], belief merging [Everaere et al. 2007], crowdsourcing [Mao et al. 2013].
[2]Some popular voting rules do not fall into the two categories, for example the *Single Transferable Vote (STV)*.

Authors' address: lxia@seas.harvard.edu

Maximin (a.k.a. *Simpson-Kramer*) is a Condorcet consistent rule, which selects the alternative that has the highest worst-case head-to-head wins.

One key question in social choice theory is: *Which voting rule is the best*? This is not an easy question, and there has been a long debate over even the meaning of optimality between the advocates of the above two categories, with no clear victory claimed by either side. This can date back to the battle in the 18th century between Jean-Charles de Borda, the inventor of the Borda rule, and Marquis de Condorcet, the inventor of Condorcet consistency.

The classical way to evaluate voting rules in social choice theory is to study their satisfiability of *axiomatic properties* (*axioms* in short), which are desired properties measuring various aspects of voting rules. Unfortunately, no voting rule can satisfy the combination of even a few natural axioms, due to the celebrated Arrow's impossibility theorem [Arrow 1963].[3] Specifically, no positional scoring rule is Condorcet consistent [Fishburn 1974]. So at least the Borda advocates and Condorcet advocates can proudly announce "We are different from the opponent".

1.1 Our Approach

Instead of continuing the Borda vs. Condorcet debate and contrasting existing voting rules, we instead seek for a unified approach by asking the following question:

Do most existing voting rules share some common properties?

Notice that this is in fact a "reverse engineering" question. Knowing these common characteristics helps us understand desired properties of voting rules, so that in the future if we want to design a new voting rule, we can focus on these natural properties. More precisely, we ask the following question:

Is there a framework that reconciles the two categories of voting rules?

A straightforward (and uninformative) answer is affirmative, for example "the class of all voting rules". However, a good framework should be general, covering most existing voting rules, but more importantly, needs to have a good mathematical structure that distinguishes it from an arbitrary voting rule. This means that a good framework should not be too general.

In the rest of the paper, we will introduce the class of generalized scoring rules, and show evidences suggesting that it is a good framework for this purpose.

2. GENERALIZED SCORING RULES

We start with an example of rethinking Borda to illustrate the idea behind the definition. Let $\mathcal{A} = \{a_1, \ldots, a_m\}$ denote a set of m alternatives, and let $\mathcal{L}(\mathcal{A})$ denote the set of all linear orders over $\mathcal{A}$. Let $P = (V_1, \ldots, V_n)$ denote a *preference profile*, where each $V_j \in \mathcal{L}(\mathcal{A})$ represents the vote of agent j. A voting rule r is a mapping that chooses a single winner for any preference profile.[4]

[3]See [Nurmi 1987] for definitions of some natural axioms and a thorough comparison of voting rules in terms of satisfiability of these axioms.

[4]The definition of GSRs can be much more general, but for better presentation we will focus on the classical election setting in this paper.

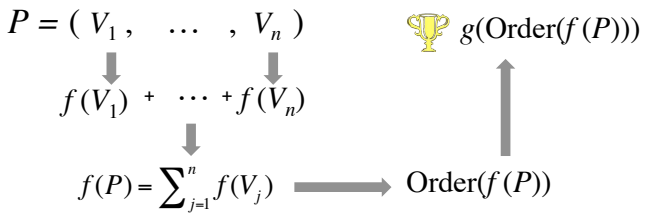

Fig. 1. Illustration of generalized scoring rules.

Example 1 (Rethinking Borda) *For each vote $V \in \mathcal{L}(\mathcal{A})$, we map it to an m-dimensional vector $f(V) \in \mathbb{R}^m$, where the ith component is the number of points a_i obtains in V. Given a preference profile P, we let $f(P) = \sum_{j=1}^{n} f(V_j)$. It is not hard to see that $f(P)$ represents the total points obtained by the alternatives in P. Therefore, the winner is a_i where the i is the largest component of $f(P)$.*

In the above example, we clearly see the following pattern in Borda:

(1) Each vote is mapped to a vector via a function f.
(2) The vectors are summed up to produce a total vector.
(3) The winner is determined by the *order* over the components in the total vector via a function g.

This leads to the definition of generalized scoring rules [Xia and Conitzer 2008], illustrated in Figure 1.[5] Slightly more formally, fix the number of alternatives m, we have a number K that represents the dimensionality of the vectors votes are mapped to by f. Then, a *generalized scoring rule (GSR)*, denoted by $GS(f, g)$, is defined by a pair of function f and g. For any input preference profile P, we perform exactly the above three steps.

(1) Each vote V in P is mapped to a vector $f(V) \in \mathbb{R}^K$.
(2) Let $f(P) = \sum_{V \in P} f(V)$.
(3) The winner is $g(\text{Order}(f(P)))$, where $\text{Order}(f(P))$ is the order over the components in $f(P)$.

***Remark* 2.1** A GSR is defined for a fixed number of alternatives and a variable number of voters.

***Remark* 2.2** By saying that a voting rule r is a GSR, we mean that there exist f and g such that $r = GS(f, g)$. It is possible that different pairs (f_1, g_1) and (f_2, g_2) correspond to the same voting rule.

We have seen in Example 1 that Borda is a GSR, where $K = m$, f is the function described in Example 1, and g simply selects the alternative whose corresponding component is the largest.[6] The next example shows that Maximin, which is Condorcet consistent, is also a GSR.

[5] We use the equivalent definition in [Xia 2012].
[6] Suppose that ties are broken w.r.t. a fixed linear order over $\mathcal{A}$.

Example 2 *To show that Maximin is a GSR, we let*

—$K_M = m(m-1)$*; the components are indexed by pairs* (i,j) *such that* $i,j \leq m$, $i \neq j$.

—$(f_M(V))_{(i,j)} = \begin{cases} 1 & \text{if } a_i \succ_V a_j \\ 0 & \text{otherwise} \end{cases}$

—$g_M(\cdot)$ *simulates Maximin based on the information contained in* $\text{Order}(f_M(P))$.

3. GENERALITY OF GSRS

GSRs are quite general: It was shown by construction that many commonly studied voting rules using fixed-order tie-breaking are GSRs [Xia and Conitzer 2008], including all positional scoring rules, Maximin, Copeland, ranked pairs, Bucklin, and multi-round voting rules including STV, plurality with runoff, Nanson's rule, and Baldwin's rule. Notice that many of these rules are Condorcet consistent.

GSRs are not too general: GSRs are equivalent to the class of voting rules that satisfy the following two axioms [Xia and Conitzer 2009].

—*Anonymity*: r satisfies anonymity if the winner is insensitive to the names of the agents.

—*Finite local consistency (FLC)*: r satisfies FLC if the set of all preference profiles over $\mathcal{A}$ can be partitioned into T parts $\{S_1, \ldots, S_T\}$, such that for any pair of preference profiles (P_1, P_2) that belong to the same partition and $r(P_1) = r(P_2)$, we have $r(P_1) = r(P_1 \cup P_2)$.

***Remark* 3.1** FLC implies *homogeneity*, which says that for any preference profile P and any number $k \in \mathbb{N}$, $r(P) = r(nP)$. Therefore, any voting rule that does not satisfy homogeneity is not a GSR. Among commonly studied voting rules, only Dodgson's rule does not satisfy homogeneity, which means that it is not a GSR.[7]

***Remark* 3.2** Any voting rule that does not satisfy anonymity is not a GSR, including Borda equipped with a non-anonymous tie-breaking mechanism, for example breaking ties using the first voter's vote.

***Remark* 3.3** FLC is an extension of the *consistency* axiom in social choice theory, which is FLC with $T = 1$. Consistency was only previously known to be satisfied by positional scoring rules.[8]

4. WHY ARE GSRS INTERESTING?

Useful in studying the frequency of manipulability. Suppose there are n' manipulators and their favorite alternative a. Let n non-manipulators votes be generated i.i.d. according to some probability distribution. We are interested in the *frequency of manipulability*, which is the probability that the n' manipulators can make a win by voting collaboratively.

[7]However, Dodgson's rule is arguably not a good voting rule since it fails to satisfy many desired axioms, and has a high computational complexity [Brandt 2009].

[8]C.f. Young's insightful axiomatic characterization of positional scoring rules [Young 1975].

For a large class of GSRs, we proved a dichotomy theorem on the frequency of manipulability [Xia and Conitzer 2008]. The theorem states that if the number of manipulators is $o(\sqrt{n})$, then the frequency of manipulability goes to 0 as n goes to infinity; if the number of manipulator is $\omega(n)$, then the frequency of manipulability goes to 1 as n goes to infinity.

The theorem was extended (with slight tweaks) to all GSRs [Mossel et al. 2012], and was also extended to other types of strategic behavior [Xia 2013]. These type of research was generally viewed as negative, because they reconfirm the high-level message that computational complexity is not a strong barrier against manipulation [Faliszewski and Procaccia 2010; Mossel and Rácz 2012]. On the positive side, they suggest that there exists efficient methods for post-election audits by computing the *margin of victory* [Xia 2012].

Reconcile Borda and Condorcet via FLC. Since no Condorcet consistent voting rule satisfies consistency (plus a few other natural axioms), it would be great if a Condorcet consistent voting rule can satisfy a weaker version of consistency. The FLC axiom, which is satisfied by all GSRs, plays such a role, and thus provides a new angle of evaluating Condorcet consistent rules.

At first glance, FLC looks quite abstract, but in fact it has a natural interpretation: each partition S_t can be seen as an abstract "characteristic" of preference profiles. Then, FLC comes down to saying that the voting rule is consistent for preference profiles sharing the same characteristic.

Take Kemeny's rule as an example. It does not satisfy consistency. However, if we define a partition where for every linear order $l \in \mathcal{L}(\mathcal{A})$, S_l is composed of all preference profiles that are closest to l in Kendall tau distance, then Kemeny is consistent within each S_l, since if $P_1, P_2 \in S_l$, then l is the linear order that is closest to $P_1 \cup P_2$ in Kendall tau distance.

Have nice structures and are related to Machine Learning. Mathematically, GSRs are equivalent to *hyperplane rules*, which view all preference profiles in a geometric space and use multiple linear hyperplanes to separate regions for winner determination [Mossel et al. 2012; Xia and Conitzer 2009].

At a high level, GSRs have two interesting connections to Machine Learning. Here a voting rule can be seen as a multi-class classifier, where $\mathcal{A}$ is the set of classes [Procaccia et al. 2009]. A separating hyperplane can be seen as a linear binary classifier.

First, a GSR can be seen as the result of decision making (choosing the winner) based on the position of the input preference profile w.r.t. all hyperplanes. In other words, a GSR classifies a preference profile based on the outputs of all linear binary classifiers. This has been explored in Machine Learning as an effective way to build multi-class classifiers by binary classifiers [Tax and Duin 2002]. Therefore, when designing the g function of a GSR, we may use ideas from the literature on multi-class classifiers.

Another connection is to treat $\mathcal{O}_K$ as the set of features, and f works as the feature abstraction function (though K is not necessarily small). The collective choice is made in an additive manner where "feature values" of the input votes are summed up across the agents. Therefore, when designing the f function of a GSR, we may use techniques developed for feature selection.

How to explore these high-level connections for application is an interesting direction for future research. See [Xia 2013] for some preliminary ideas.

5. CONCLUSION AND FUTURE DIRECTION

In this paper we surveyed some developments in generalized scoring rules, a relatively new class of voting rules for studying social choice. Given the generality and structure of GSRs, there are many directions for future research. In future/ongoing work, we see at least the following directions.

—Develop more general techniques and results for GSRs, for example post-election audits and compilation complexity [Chevaleyre et al. 2009].

—Explore deeper and more practical relationships between GSRs and Machine Learning.

—Study relationship between GSRs and other classes of voting rules, for example distance-based rules [Meskanen and Hannu 2008; Elkind et al. 2011].

6. ACKNOWLEDGMENTS

The author thanks Ariel Procaccia for valuable feedbacks. This work is supported by NSF under Grant #1136996 to the Computing Research Association for the CIFellows Project.

REFERENCES

Arrow, K. 1963. *Social choice and individual values*, 2nd ed. New Haven: Cowles Foundation. 1st edition 1951.

Brandt, F. 2009. Some remarks on Dodgson's voting rule. *Mathematical Logic Quarterly 55*, 460–463.

Chevaleyre, Y., Lang, J., Maudet, N., and Ravilly-Abadie, G. 2009. Compiling the votes of a subelectorate. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*. Pasadena, CA, USA, 97–102.

Dwork, C., Kumar, R., Naor, M., and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proceedings of the 10th World Wide Web Conference*. 613–622.

Elkind, E., Faliszewski, P., and Slinko, A. 2011. Rationalizations of condorcet-consistent rules via distances of hamming type. *Social Choice and Welfare. To appear*.

Ephrati, E. and Rosenschein, J. S. 1991. The Clarke tax as a consensus mechanism among automated agents. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*. Anaheim, CA, USA, 173–178.

Everaere, P., Konieczny, S., and Marquis, P. 2007. The strategy-proofness landscape of merging. *Journal of Artificial Intelligence Research 28*, 49–105.

Faliszewski, P. and Procaccia, A. D. 2010. AI's war on manipulation: Are we winning? *AI Magazine 31,* 4, 53–64.

Fishburn, P. C. 1974. Paradoxes of voting. *The American Political Science Review 68,* 2, 537–546.

Ghosh, S., Mundhe, M., Hernandez, K., and Sen, S. 1999. Voting for movies: the anatomy of a recommender system. In *Proceedings of the third annual conference on Autonomous Agents*. 434–435.

Mao, A., Procaccia, A. D., and Chen, Y. 2013. Better human computation through principled voting. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*. Bellevue, WA, USA.

Meskanen, T. and Hannu, N. 2008. *Power, freedom, and voting*. Springer-Verlag Berlin Heidelberg, Chapter Closeness counts in social choice.

Mossel, E., Procaccia, A. D., and Racz, M. Z. 2012. A smooth transition from powerlessness to absolute power. http://www.cs.cmu.edu/~arielpro/papers/phase.pdf.

Mossel, E. and Rácz, M. Z. 2012. Election Manipulation: The Average Case. *ACM SIGecom Exchanges 11,* 2, 22–24.

Nurmi, H. 1987. *Comparing voting systems*. Springer.

Procaccia, A. D., Zohar, A., Peleg, Y., and Rosenschein, J. S. 2009. The learnability of voting rules. *Artificial Intelligence 173*, 1133—1149.

Tax, D. M. and Duin, R. P. 2002. Using two-class classifiers for multiclass classification. In *Proceedings of the 16th International Conference on Pattern Recognition*. 124–127.

Xia, L. 2012. Computing the margin of victory for various voting rules. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. Valencia, Spain, 982–999.

Xia, L. 2013. How many vote operations are needed to manipulate a voting system. *Arxiv*.

Xia, L. and Conitzer, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*. Chicago, IL, USA, 109–118.

Xia, L. and Conitzer, V. 2009. Finite local consistency characterizes generalized scoring rules. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*. Pasadena, CA, USA, 336–341.

Young, H. P. 1975. Social choice scoring functions. *SIAM Journal on Applied Mathematics 28,* 4, 824–838.

Back to Original Frugality

RODRIGO A. VELEZ
Texas A&M University

We review recent research on frugality of mechanisms for the procurement of a spanning network. Frugality here is defined as the ratio of the maximum price that can be charged to the buyer in some equilibrium to the true minimal cost. Previous negative results are qualified under natural restrictions in costs.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Sciences**]: Economics

General Terms: Economics
Additional Key Words and Phrases: worst case scenario equilibrium analysis; frugality; minimum cost spanning tree problem; price of imperfect competition

1. THE PROBLEM

Consider the problem of a buyer who needs to procure a network spanning a set of three nodes. Let $\{e, f, g\}$ be the complete graph connecting these nodes. Assume that each edge can be exclusively constructed by a different seller, and thus, we can unambiguously refer to sellers e, f, and g. Let c_e, c_f, and c_g be the sellers' costs of constructing the edges, and assume that $c_e \leq c_f < c_g$.

The buyer, with no access to the sellers' technology, and no knowledge of the true costs, will use a "mechanism" to procure the spanning network. Two popular options are the first price auction and the celebrated canonical VCG, or pivotal, mechanism.

In the first price auction the buyer asks the sellers for bids and selects a minimum cost spanning tree for the reported bids. Each selected seller constructs her edge and is paid her bid. In the pivotal mechanism sellers are asked to report their costs. Then, the buyer selects a minimum cost spanning tree for the reported costs. Each selected seller constructs her edge and is paid her reported cost plus the reduction in the cost of a minimum cost spanning tree due to the availability of the seller's edge at her reported cost.

If there is complete information among sellers, limit Nash equilibrium predicts that in the first price auction the efficient tree, i.e., $\{e, f\}$, is built and the buyer pays an amount $2x$ where $x \in [c_f, c_g]$.[1,2] Thus, the maximum aggregate price paid

[1]Pure strategy Nash equilibria may fail to exist in this game depending on the rule that determines the winner of the auction when there are multiple minimal cost spanning trees at the reported bids.

[2]If edge g is selected, seller g must be paid at least c_g and another edge, say f, is not selected. This cannot be true, for seller f would be able to gain at least $c_g - c_f$ by bidding just below c_g. Thus, in equilibrium, seller g bids some amount x and sellers e and f bid just below x. Clearly, $x \geq c_f$, for seller f can guarantee a non-negative payoff by bidding above x. Moreover, $x \leq c_g$, for otherwise seller g would be able to gain by bidding just below x.

Author's address: rvelezca@econmail.tamu.edu

by the buyer in a limit Nash equilibrium of the first price auction is $2c_g$. One can easily see that this is also the payment of the buyer in the truthful equilibrium of the pivotal mechanism.

From the buyer's perspective it is relevant to compare the outcomes of a procurement mechanism with the ideal outcome in which the buyer has access to the sellers' technology. A handy statistic that provides such a comparison is the ratio of the highest predicted payment to the efficient cost [Archer and Tardos 2001; 2002]. In our example, this ratio for both the first price auction and the pivotal mechanism is $2c_g/(c_e + c_f)$. Obviously, this ratio can be arbitrarily high if costs are unrestricted. However, it is at most 2 if costs are metric, i.e., if costs satisfy the familiar triangle inequality, which implies that $c_g \leq c_e + c_f$. This inequality would be satisfied whenever costs are a subadditive function of the "length" of an edge in a metric space.

In [Moulin and Velez 2013], we study the procurement games described above for an arbitrary number of nodes $n > 2$. Surprisingly, the worst case scenario ratio of the maximum predicted payment to the efficient cost in the metric domain, for both first price auction and pivotal mechanism, is essentially that for the triangle, i.e., 2. The ratio increases swiftly when sellers can only bid for a subset of all edges, however.

2. RESULTS

More generally, we consider a buyer who procures a spanning network out of a set of edges that span $n > 2$ nodes. We define the *Price of Imperfect Competition* (PIC) as the ratio of the maximum payment in a limit equilibrium of the first price auction, under complete information, to the efficient cost. (See Section 3 below for the relation of the PIC with the "frugality indices" defined by [Archer and Tardos 2001; 2002] and subsequent literature.) We assume that no seller has monopoly power, i.e., there is no seller who is the exclusive bidder for a cut of the graph of available edges. Our main results are:

(1) Each limit equilibrium of the first price auction is efficient.[3]

(2) Assume that costs are metric and each seller is the exclusive bidder for an edge.
   (a) PIC is at most $n - 1$.
   (b) Assume that the graph of available edges is the complete graph.
      i. PIC is at most and up to 2 if n is odd, and is at most and up to $2\frac{n-1}{n-2}$ if n is even.
      ii. Besides obtaining tightness of our bound, we characterize the trees that can be minimal cost spanning trees for some cost matrix that achieves the upper bound. Let γ be spanning tree.
         A. There is always a metric cost matrix for which γ is a minimal cost spanning tree and for which PIC is 2.
         B. Assume that n is even. There is a metric cost matrix for which γ

[3]This is close to the efficiency of first price auction in matroid markets of [Chen and Karlin 2007], with the difference that we characterize unrefined limit equilibria and allow for each seller to be the exclusive bidder for multiple edges.

is a minimal cost spanning tree and for which PIC is $2\frac{n-1}{n-2}$ if and only if γ contains a perfect matching.

Surprisingly, worst case scenarios do not improve if costs are required to satisfy the stronger substitutability condition requiring that for each three edges forming a triangle $\{e, f, g\}$, $c_e \leq \max\{c_f, c_g\}$. Under this assumption, known as the ultrametric inequality, both 2.b.i and 2.b.iii above hold. However, in contrast to 2.b.ii, given a spanning tree γ, there is an ultrametric cost matrix for which γ is a minimum cost spanning tree and for which PIC is 2, if and only if, γ has at least a leaf edge of which the inner end node is of degree two.
Even though worst case PIC is the same in metric and ultrametric domains, intuitively PIC is lower on average for ultrametric costs. We confirm this in a simple probabilistic model.

(c) Assume that there is only one edge that is not available for purchase by the buyer. If $n = 4$ or $n \geq 6$, then PIC is at most and up to 3. If $n = 5$, then PIC is at most and up to 4.
(d) The worst case PIC increases swiftly when some edges are not available for purchase. We provide a lower bound, which we conjecture is tight, for the worst case PIC as a function of the number of available edges. If at most half of the edges are available for purchase our bound is tight: let $m \leq \frac{n(n-1)}{2}$; the worst case PIC among all metric cost problems in which m edges are available for purchase is the maximum possible, i.e., $n - 1$.
(e) If each seller is the exclusive bidder for exactly one edge, the maximum price charged to the buyer in a limit equilibrium of the first price auction is equal to the payment in the truthful equilibrium of the pivotal mechanism. (Thus, the PIC can be equally interpreted as measuring frugality of the pivotal mechanism.)

(3) Our structural results hold in the more general problem of procuring the basis of a matroid. First, limit Nash equilibria of the first price auction are efficient. Second, define the triangle inequality in a matroid as the requirement that no element of a circuit costs more than half the aggregate cost of the circuit. Then the PIC in the metric domain for a given matroid is bounded above by the dimension of the matroid. Our results illustrate that the whole range of values for the PIC in the metric domain can be achieved by minimum cost spanning tree problems.

3. FRUGALITY

Our work contributes to the literature concerned with the payment of a buyer who needs to procure a team to complete a complex task (see [Karlin et al. 2005] and references within). The main concerns in this literature have been (i) to define frugality indices that measure the extent to which the buyer overpays in a given situation; (ii) to evaluate the frugality of popular mechanisms; and (iii) to design mechanisms that minimize worst case scenario frugality indices [Archer and Tardos 2001; 2002; Talwar 2003; Elkind et al. 2004; Karlin et al. 2005; Chen and Karlin 2007] (see [Nisan et al. 2007, Section 13.5] for a survey). As in our work, the

original frugality ratio was defined as the ratio of the highest predicted price to the efficient cost [Archer and Tardos 2001; 2002]. This ratio is a natural comparison between the strategic procurement situation and the best case situation for the buyer. Unfortunately, the first attempts to study it produced only negative results [Archer and Tardos 2001; 2002; Elkind et al. 2004]. This led to the belief that it was necessary to change the benchmark with respect to which the payment of the buyer is compared [Talwar 2003; Karlin et al. 2005]. The most prominent alternative, proposed by [Karlin et al. 2005], uses as benchmark cost the solution to a linear program that coincides with the pivotal payment in a matroid market. Thus, under this definition, the pivotal mechanism and the first price auction are the ideal of frugality in our environment.

Our work departs from this literature in an important way. We impose natural restrictions in the cost structure, while retaining the original benchmark cost with respect to which the buyer's payment is compared, i.e., the efficient cost. This allows us to obtain the first positive results for the unmodified frugality ratio for the popular first price auction and pivotal mechanisms. With this we open back the question of designing frugal mechanisms in restricted domains.

REFERENCES

Archer, A. and Tardos, E. 2001. Truthful mechanisms for one-parameter agents. In *Proceedings, 42nd IEEE Symposium on the Foundations of Computer Science.* 482–491.

Archer, A. and Tardos, E. 2002. Frugal path mechanisms. In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms.* SODA '02. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 991–999.

Chen, N. and Karlin, A. R. 2007. Cheap labor can be expensive. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms.* SODA '07. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 707–715.

Elkind, E., Sahai, A., and Steiglitz, K. 2004. Frugality in path auctions. In *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms.* SODA '04. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 701–709.

Karlin, A., Kempe, D., and Tamir, T. 2005. Beyond vcg: frugality of truthful mechanisms. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on.* 615–624.

Moulin, H. and Velez, R. A. 2013. The price of imperfect competition for a spanning network. *Games Econ. Behav.*. doi 10.1016/j.geb.2013.03.012.

Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V., Eds. 2007. *Algorithmic Game Theory.* Cambridge University Press, New York.

Talwar, K. 2003. The price of truth: Frugality in truthful mechanisms. In *STACS 2003*, H. Alt and M. Habib, Eds. Lecture Notes in Computer Science, vol. 2607. Springer Berlin Heidelberg, 608–619.

Constrained Signaling for Welfare and Revenue Maximization

SHADDIN DUGHMI
University of Southern California
and
NICOLE IMMORLICA
Microsoft Research
and
AARON ROTH
University of Pennsylvania

We consider auction settings where the seller is constrained in the amount and nature of information he may reveal about the good being sold. This is encountered, for example, in online advertising auctions, where communicating precise details of every viewer to interested advertisers is impractical, costly, and possibly socially undesirable. We initiate the study of *constrained signaling* in such settings, where a seller must choose which information to reveal subject to exogenous constraints on the signaling policy. We consider a seller employing the second-price auction, and present algorithms and hardness results for approximating the welfare and revenue maximizing signaling policies under a variety of constraints.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: General

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Signaling, Auctions, Mechanism Design

1. INTRODUCTION

The study of economic interactions in the presence of information asymmetries has a rich history, beginning with the seminal work of [Akerlof 1970]. Akerlof observed that the information structure of a market — i.e. who has what information regarding the goods and services for sale — can have a profound effect on equilibrium outcomes. Since then, a rich literature has examined the effect of information revelation, also known as *signaling*, in markets (e.g. by [Spence 1973; 2002]) and, on a smaller scale, in auctions.

In this work [Dughmi et al. 2013], we consider signaling in some of the simplest economic settings, namely single item auctions. In addition to their preponderance on the Internet, such as in online advertising, auctions serve as an instructive microcosm in which to study information asymmetry between buyers and sellers in a market, and the effects of signaling. In our model, an auctioneer is looking to sell a single good to one of several buyers, though the various qualities of the good are known to the auctioneer yet ex-ante unknown to the buyers. This arises, for example, in online advertising: a content publisher looking to sell an impression

Authors' addresses: shaddin@usc.edu, nicimm@microsoft.com, aaroth@cis.upenn.edu.

has more information regarding the identity of the viewer behind the impression than do the advertisers bidding for it. We model such information asymmetry by assuming that the good is drawn by nature from a distribution over possible goods, assumed to be common knowledge. The auctioneer then learns the identity of the realized good, and has the opportunity to publicly announce a message, which we refer to as a *signal*, to the buyers before running an auction. We assume that the auctioneer commits to a *signaling scheme* — a policy specifying a distribution over signals for each good — before nature draws the good, and announces said policy to all the buyers. Each signaling scheme and choice of an auction format (first price, second price, etc) then induces a game among the buyers. The auctioneer's problem is that of designing the signaling scheme to optimize his favored objective, commonly his revenue or the welfare of the winning player.

The conventional wisdom on signaling in auctions is that *more information is better*. This principle rings most true in the simplest setting of all: an auctioneer running a second price auction,[1] and interested in maximizing welfare. To see this, observe that when buyers know the identity of the good being sold, a second price auction awards the item to the player who values it most. In contrast, when the employed signaling scheme announces partial information regarding the identity of the good, the second price auction awards the good to the bidder with the greatest *expected* value after a Bayesian update based on the signal, which is not necessarily the optimal point-wise choice. The *linkage principle* of [Milgrom and Weber 1982] reinforces the same popular lore: under some conditions, revealing more information to the players increases the auctioneer's *revenue* as well.[2] It is tempting, therefore, to conclude that the seller's problem of what to reveal is trivial in many settings, as he should simply announce all available information.

Our motivating observation is that complete transparency is often impossible or costly. In online advertising auctions for example, the sheer volume of sales and diversity of viewers make it impractical for a publisher to communicate to advertisers the precise details of every viewer, and the publisher's privacy policy may further constrain the information they reveal. In addition, legal and reputational considerations may restrict a seller from announcing certain signals for certain items: e.g., products in a grocery store must pass a certification process from the corresponding agency to be sold under the "organic" label. Such constraints introduce intricate tradeoffs in choosing *which* information to reveal. Quantifying those tradeoffs inevitably requires examination of these settings with an optimization lens, as we begin to do in this work.

2. SUMMARY OF RESULTS

To begin our study of constrained signaling, we define and analyze two illustrative settings, one with a finite set of items and combinatorial constraints on the signaling

[1]The choice of auction format is not particularly important for this conclusion, though we fix the second price auction for clarity.

[2]However, as pointed out by [Levin and Milgrom 2010], the linkage principle does not apply in many settings, including ones we consider in this work. Recent work in [Emek et al. 2012] and [Bro Miltersen and Sheffet 2012] designs polynomial time algorithms for computing revenue-maximizing signaling schemes in general.

scheme, and another where the set of items is infinite. In both settings, we consider a seller who first invokes a signaling scheme and then runs a second price auction.

(1) In the *combinatorial setting*, an item is drawn by nature from a finite set according to a known prior. We consider two constraints on the signaling policy: a *cardinality constraint* limiting the number of signals used, and a *graph constraint* given as a bipartite graph with items on one side, signals on the other, and edges describing the valid signals for each item.

(2) In the *geometric setting*, items live in a compact subset of Euclidean space. As before, an item is drawn by nature from a common prior. We assume the dimension d of the space is too high to communicate the identity of an item directly, and constrain our signals to strings of length logarithmic in d.

We present algorithms and impossibility results for computing optimal signaling schemes in these settings, after making some assumptions.

(1) In the combinatorial setting, we assume that players' valuations for potential items are known, or drawn from a prior with constant-size support. We present constant-factor approximation algorithms for computing the welfare maximizing signaling scheme under both a cardinality constraint and a graph constraint, and for computing the revenue maximizing signaling scheme under only a cardinality constraint.

(2) In the combinatorial setting, we show that both welfare and revenue are NP-hard to approximate better than a specific constant factor, even assuming given valuations and only a cardinality constraint on the signaling scheme.

(3) In the geometric setting, we assume players' valuations are defined in terms of either the distance or the angle from the realized good to some target good or set of goods, and drawn from a known prior. We present communication-constrained signaling schemes that achieve nearly the optimal welfare, compared even to the optimal unconstrained signaling scheme.

Our results for the combinatorial setting draw heavily on techniques from the literature on submodular function optimization. For the geometric setting, we employ tools from no regret learning and metric embeddings to prove our results.

3. FUTURE DIRECTIONS

The main contribution of this work, we believe, is enabling future study of additional constrained signaling settings. We conclude by mentioning two of them.

(1) In both models we study, the constraint on our signaling scheme is essentially separable over items. What if, instead, products are described by a set of attributes — in online advertising, those may include age, gender, location, and browsing history, and a signaling scheme is constrained to choose a small subset of those attributes, crucially the same for all goods, to advertise prior to sale? Problems of this form are arguably more natural than those we consider in this work, though appear particularly challenging and intimately related to unresolved questions in learning theory.

(2) Signaling can be thought of as a process of classification; specifically of items into subsets, each of which is associated with a signal. When items are associated with entities that have a contractual or service relationship with the seller, as in online advertising, issues of fairness inevitably arise, as described in [Dwork et al. 2012]. A poignant example is the controversy stirred by the recent work of [Sweeney 2013], which illustrates the effect of race on ad delivery in Google's Adsense. To avoid such blowback, a seller may constrain their signaling scheme to treat similar items similarly.

REFERENCES

Akerlof, G. 1970. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*, 488–500.

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

Dughmi, S., Immorlica, N., and Roth, A. 2013. Constrained signaling for welfare and revenue maximization. *CoRR abs/1302.4713*.

Dwork, C., Hardt, M., Pitassi, T., Reingold, O., and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ACM, 214–226.

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*. ACM, 514–531.

Levin, J. and Milgrom, P. 2010. Online advertising: Heterogeneity and conflation in market design. *American Economic Review 100*, 603–607.

Milgrom, P. and Weber, R. 1982. A theory of auctions and competitive bidding. *Econometrica 50*, 1089–1122.

Spence, M. 1973. Job market signaling. *The quarterly journal of Economics 87,* 3, 355–374.

Spence, M. 2002. Signaling in retrospect and the informational structure of markets. *The American Economic Review 92,* 3, 434–459.

Sweeney, L. 2013. Discrimination in online ad delivery. *arXiv preprint arXiv:1301.6822*.