

Automated Design of Robust Mechanisms

Michael Albert and Vincent Conitzer

Department of Computer Science
Duke University
Durham, NC 27708, USA
{malbert,conitzer}@cs.duke.edu

Peter Stone

Department of Computer Science
University of Texas at Austin
Austin, TX 78712, USA
pstone@cs.utexas.edu

Abstract

We introduce a new class of mechanisms, *robust mechanisms*, that is an intermediary between ex-post mechanisms and Bayesian mechanisms. This new class of mechanisms allows the mechanism designer to incorporate imprecise estimates of the distribution over bidder valuations in a way that provides strong guarantees that the mechanism will perform at least as well as ex-post mechanisms, while in many cases performing better. We further extend this class to mechanisms that are with high probability incentive compatible and individually rational, ϵ -*robust mechanisms*. Using techniques from automated mechanism design and robust optimization, we provide an algorithm polynomial in the number of bidder types to design robust and ϵ -robust mechanisms. We show experimentally that this new class of mechanisms can significantly outperform traditional mechanism design techniques when the mechanism designer has an estimate of the distribution and the bidder's valuation is correlated with an externally verifiable signal.

Introduction

Auctions are one of the fundamental tools of the modern economy for allocating resources. They are used to allocate online ad space, offshore oil drilling rights, famous artwork, small and medium lift capacity to planetary orbit, government supply contracts, FCC spectrum licenses, and almost limitless numbers of other things, large and small. Further, the sizes of these markets are economically enormous. In 2014, \$10 billion dollars of ad revenue was generated through automated auctions (Interactive Advertising Bureau (IAB) 2015). In 2012, just four government agencies—the Army, the Department of Homeland Security, the Department of the Interior, and the Department of Veteran Affairs—purchased \$800+ million of commercial items through auctions (Government Accountability Office 2013). In 2014, NASA awarded contracts to Boeing and Space-X worth \$4.2 billion and \$2.6 billion respectively through an implicit auction process (NASA 2014). In an ongoing auction, the FCC is expected to allocate between \$60 and \$80 billion worth of broadcast spectrum. Given the economic magnitudes involved, it is crucial that these auctions are implemented optimally, for even small deviations from opti-

mality can lead to millions of dollars worth of lost revenue, inefficiencies in resource allocation, and overspending.

It has long been understood that revenue optimal auction mechanisms are *prior-dependent* (also known as *Bayesian*) mechanisms (Myerson 1981; Cremer and McLean 1985; 1988; Lopomo 2001), i.e., mechanisms that assume some knowledge of the kinds of bidders that are likely to participate. For a seller trying to maximize revenue by selling a single item to independent bidders, she¹ needs to set a reserve price below which she will not sell the item (Myerson 1981), and this reserve price is dependent on her belief about the bidders she is likely to face. However, much of the focus of the algorithmic mechanism design community has been on the approximate optimality of *simple* mechanisms (Bulow and Klemperer 1996; Hartline and Roughgarden 2009; Roughgarden and Talgam-Cohen 2013; Morgenstern and Roughgarden 2015), that is, mechanisms that are either prior-independent or weakly prior-dependent (this distinction will be made clear later on). This focus is due to two factors. First, prior-dependent mechanisms can be very brittle to mis-specified priors (Lopomo 2001; Albert, Conitzer, and Lopomo 2015). That is to say, if a prior-dependent mechanism is constructed using an incorrect prior it can perform much worse than simple mechanisms (Hartline and Roughgarden 2009). Second, competition can be an effective substitute for knowledge of the distribution (Bulow and Klemperer 1996), so instead of implementing prior-dependent mechanisms, practitioners generally implement prior-independent mechanisms under the assumption that there are many bidders, or that it will somehow be possible to *acquire* new bidders

Unfortunately, in many auctions there is no feasible way to acquire more bidders. When NASA is awarding contracts to private space companies, they cannot generate more companies with the expertise to provide lift capacity, nor can Google find more bidders who are interested in advertising on search results for obscure brand-related keywords. Thin auctions are an actively recognized concern for many organizations that use auction mechanisms. A Government Office of Accountability report from 2013 (Government Accountability Office 2013) examining the use of reverse auctions by

¹In this paper, we will use “he” to denote bidders and “she” to denote mechanism designers/sellers.

four governmental organizations, found that of the 19,688 reverse auctions the organizations conducted in 2012 with a total worth of \$800+ million, over one-third had only a *single bidder*. Further, the organizations discussed in the report each voiced concern over the lack of competition as a significant hindrance to the effective application of auctions.

In this paper, we develop a new class of mechanisms, *robust mechanisms*, that allow for some degree of uncertainty in the distribution over the bidders types while still performing better than weakly prior-dependent mechanisms such as ex-post mechanisms. We provide an algorithm for computing an optimal robust mechanism in polynomial time in the number of bidder types by combining techniques from the automated mechanism design literature (Conitzer and Sandholm 2002; 2004; Guo and Conitzer 2010; Sandholm and Likhodedov 2015) and the literature on robust optimization (Bertsimas and Sim 2004; Bergemann and Morris 2005; Aghassi and Bertsimas 2006). This mechanism design algorithm is primarily targeted at applications with few bidders, or thin markets, due to exponential scaling in the number of bidders. However, thin markets are likely to benefit the most from these techniques (Bulow and Klemperer 1996; Hartline and Roughgarden 2009), so we do not consider this a significant negative. We then introduce the notion of ϵ -robust mechanisms, or mechanisms that guarantee with high probability that the standard constraints of incentive compatibility and individual rationality hold but allow for a non-zero chance of violation. Finally, we show experimentally that ϵ -robust mechanisms can significantly outperform other mechanism design procedures when the mechanism designer estimates the distribution over a bidder's type and a correlated external signal.

Preliminaries

We consider a single monopolistic seller auctioning one object, which the seller values at zero, to a single bidder whose valuation is correlated with an external signal. The special case of a single bidder and an externally verifiable signal captures many of the important aspects of this problem while increasing ease of exposition relative to the case of many bidders, and this setting has been used in the literature on correlated mechanism design (McAfee and Reny 1992; Albert, Conitzer, and Lopomo 2015; 2016) for this purpose. The external signal can, but does not necessarily, represent other bidders' bids.

The bidder has a type θ drawn from a finite set of discrete types $\Theta = \{1, \dots, |\Theta|\}$. Further, the bidder has a valuation function $v : \Theta \rightarrow \mathbb{R}^+$ that maps types to valuations for the object. Assume, without loss of generality, that for all $\theta, \theta' \in \Theta$, if $\theta > \theta'$ then $v(\theta) \geq v(\theta')$. The discrete external signal is denoted by $\omega \in \Omega = \{1, 2, \dots, |\Omega|\}$. Throughout the paper, we will denote vectors as bold symbols.

There is a probability distribution, π , over the types of the bidder and external signal where the probability of type and signal (θ, ω) is $\pi(\theta, \omega)$. The probability distribution can be represented in many possible ways, but we will represent it as a matrix. Specifically, the distribution is a matrix of dimension $|\Theta| \times |\Omega|$ whose elements are all positive and sum to one. Note that in contrast to much of the literature on

mechanism design, we do not require that the bidder type be distributed independently of the external signal.

The distribution over the external signal ω given θ will be denoted by the $|\Omega|$ dimensional vector $\pi(\cdot|\theta)$. We will, in many cases, be primarily interested in the conditional distribution over the external signal given the bidder's type, $\pi(\cdot|\theta)$, so we will represent the full distribution as a marginal distribution over Θ , π_θ , and a set of conditional distributions over Ω , $\pi(\cdot|\cdot) = \{\pi(\cdot|1), \pi(\cdot|2), \dots, \pi(\cdot| |\Theta|)\}$. Therefore, if the true distribution is π , we will alternatively represent it as $\{\pi_\theta, \pi(\cdot|\cdot)\}$.

A (direct) revelation mechanism is defined by, given the bidder type and external signal (θ, ω) , 1) a probability that the seller allocates the item to the bidder and 2) a monetary transfer from the bidder to the seller. We will denote the probability of allocating the item to the bidder as $p(\theta, \omega)$, which is a value between zero and one, and the transfer from the bidder to the seller as $x(\theta, \omega)$, where a positive value denotes a payment to the seller and a negative value a payment from the seller to the bidder. We will denote a mechanism as (\mathbf{p}, \mathbf{x}) .

Definition 1 (Bidder's Utility). *Given a realization of the external signal ω , reported type $\theta' \in \Theta$, and true type $\theta \in \Theta$, the bidder's utility under mechanism (\mathbf{p}, \mathbf{x}) is:*

$$U(\theta, \theta', \omega) = v(\theta)p(\theta', \omega) - x(\theta', \omega)$$

Due to the well-known revelation principle (e.g., Gibbons (1992)), the seller can restrict her attention to incentive compatible mechanisms, i.e., mechanisms where it is always optimal for the bidder to truthfully report his valuation. However, incentive compatibility can be specified in multiple ways. For the sake of presentation, we will restrict our focus to two of the most common, *ex-post incentive compatibility* and *Bayesian incentive compatibility*. Ex-post incentive compatible mechanisms guarantee that for any realization of the external signal, the bidder always finds it optimal to report his value truthfully. In contrast, Bayesian incentive compatible mechanisms only guarantee that, given the beliefs of the bidder over the external signal, the bidder will have the highest *expected* utility if he reports truthfully: after seeing the realization of the external signal, he may regret his report.

Definition 2 (Ex-Post Incentive Compatibility). *A mechanism (\mathbf{p}, \mathbf{x}) is ex-post incentive compatible (IC) if:*

$$\forall \theta, \theta' \in \Theta, \omega \in \Omega : U(\theta, \theta, \omega) \geq U(\theta, \theta', \omega)$$

Definition 3 (Bayesian Incentive Compatibility). *A mechanism (\mathbf{p}, \mathbf{x}) is Bayesian incentive compatible (IC) if:*

$$\forall \theta, \theta' \in \Theta, \sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \theta, \omega) \geq \sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \theta', \omega)$$

Bayesian incentive compatibility is a statement about the *beliefs* of the bidder over the external signal, $\pi(\omega|\theta)$. Specifically, it allows the seller to determine payments by *lottery*. The lottery that bidder i faces can be dependent on his valuation, but the lottery itself is over the external signal (see Albert, Conitzer, and Lopomo (2016) for a full discussion).

Bayesian incentive compatibility is a strict relaxation of ex-post in the sense that any mechanism that is ex-post incentive compatible is also Bayesian incentive compatible.

In addition to incentive compatibility, we are interested in mechanisms that are *individually rational*, i.e., it is rational for a bidder to participate in the mechanism. We will define *ex-post individual rationality* (a bidder is *never* worse off by participating in the mechanism) and *ex-interim individual rationality* (the bidder has non-negative *expected* utility for participating in the mechanism). Again, ex-interim individual rationality is a strict relaxation of ex-post.

Definition 4 (Ex-Post Individual Rationality). *A mechanism (p, x) is ex-post individually rational (IR) if:*

$$\forall \theta \in \Theta, \omega \in \Omega : U(\theta, \theta, \omega) \geq 0$$

Definition 5 (Ex-Interim Individual Rationality). *A mechanism (p, x) is ex-interim individually rational (IR) if:*

$$\forall \theta \in \Theta : \sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \theta, \omega) \geq 0$$

We will refer to mechanisms that satisfy ex-post individual rationality and incentive compatibility as *ex-post mechanisms* and mechanisms that satisfy Bayesian incentive compatibility and ex-interim individual rationality as *Bayesian mechanisms*. Bayesian mechanisms are what we have been referring to as prior-dependent mechanisms, while ex-post is weakly prior-dependent, i.e., only the objective function depends on the distribution, not the constraints over incentive compatibility and individual rationality.

To illustrate the importance of prior-dependent mechanisms, it is necessary to review a few important results in the literature on revenue maximization with correlated valuation distributions when the distribution is perfectly known.

Definition 6 (Cremer-McLean Condition). *The distribution over types π , is said to satisfy the Cremer-McLean condition if the set of beliefs associated with the bidder, $\{\pi(\cdot|\theta) : \theta \in \Theta\}$, are linearly independent.*

Theorem 1 (Cremer and McLean (1985)). *If the Cremer-McLean condition is satisfied by the distribution π , then there exists an ex-interim IR and ex-post IC mechanism that extracts the full social surplus as revenue.*

This result due to Cremer and McLean (1985) states that under the apparently reasonable Cremer-McLean condition, i.e., a condition that holds with probability one for a random distribution, the mechanism designer can generate as much revenue in expectation as if she knew the bidder's valuation. This is a remarkable result and it can be relaxed further, for both ex-post and Bayesian IC, by the results in Albert, Conitzer, and Lopomo (2016).

Theorem 2 (Albert, Conitzer, and Lopomo (2016)). *A Bayesian IC and ex-interim IR mechanism can extract full social surplus as revenue if and only if there exists a convex function $G : \mathbb{R}^{|\Omega|} \rightarrow \mathbb{R}$ such that for all $\theta \in \Theta$, $G(\pi(\cdot|\theta)) = -v(\theta)$.*

Consistent Sets of Distributions

While Theorems 1 and 2 make relatively weak assumptions about the distributions in order to guarantee full revenue extraction, they do require that the mechanism designer knows the distribution exactly. If instead of precise knowledge of the distribution of bidder types and external signals the mechanism designer has an imprecise estimate of the distribution, the prior-dependent mechanism can fail to be both incentive compatible and individually rational. This failure can be a significant problem for two reasons. First, if the mechanism is not individually rational bidders will not participate in the mechanism. If the market is thin, the loss of even a single bidder can lead to significant decreases in expected revenue, even relative to simple mechanisms (Bulow and Klemperer 1996). Second, if the mechanism is not incentive compatible, the bidder may optimally choose to mis-report his true valuation, leading both to biases in future estimates of the distribution and difficulty in reasoning about the performance of the mechanism, since it is unclear a-priori how the bidder will report.

It is in this sense that Bayesian incentive compatible and ex-interim individually rational mechanisms are, in general, strongly prior-dependent. The mechanism depends not only on the seller's estimate of the distribution, but also the bidder's belief over the distribution. The consequences of these being mis-aligned is not just slightly lower expected revenue, as would be the case for weakly prior-dependent mechanisms such as a second price auction with reserve; it is a failure of the mechanism to maintain fundamental characteristics (Hartline 2014; Albert, Conitzer, and Lopomo 2015). Therefore, unless the seller has perfect knowledge of the bidder's beliefs, standard mechanism design techniques will leave only the option of using sub-optimal, weakly prior-dependent mechanisms.

A more realistic assumption is that the distribution is not perfectly known, but instead estimated. The seller estimates the distribution π as $\hat{\pi}$. Assume that this estimation is imperfect, and that there exists a set of distributions that are consistent with the estimated distribution.

Definition 7 (Set of Consistent Distributions). *Let $P(A)$ be the set of probability distributions over a set A . Then the space of all probability distributions over $\Theta \times \Omega$ can be represented as the Cartesian product $P(\Theta) \times \prod_{\theta \in \Theta} P(\Omega)$. A subset $\mathcal{P}(\hat{\pi}) = \mathcal{P}(\{\hat{\pi}_\theta, \hat{\pi}(\cdot|\cdot)\}) \subseteq P(\Theta) \times \prod_{\theta \in \Theta} P(\Omega)$ is a consistent set of distributions for the estimated distribution $\{\hat{\pi}_\theta, \hat{\pi}(\cdot|\cdot)\}$ if the true distribution, $\{\pi_\theta, \pi(\cdot|\cdot)\}$, is guaranteed to be in $\mathcal{P}(\hat{\pi})$.*

With a consistent set of distributions, we can relax the notion of ex-interim IR and Bayesian IC by requiring that the mechanism be IR and IC for all distributions in the consistent set. However, since the distribution π is also private information, by the revelation principle, the mechanism designer can also elicit the true distribution from the bidder and condition the mechanism on the reported distribution. Therefore, we modify the definitions of the mechanism, (p, x) , such that they depend not only on the reported type and external signal, but also the reported distribution π' . We similarly modify the definition of bidder utility.

Definition 8 (Robust Individual Rationality). *A mechanism is robust individually rational for estimated bidder distribution $\hat{\pi}$ and consistent set of distributions $\mathcal{P}(\hat{\pi})$ if for all $\theta \in \Theta$ and $\pi \in \mathcal{P}(\hat{\pi})$,*

$$\sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \pi, \theta, \pi, \omega) \geq 0$$

Definition 9 (Robust Incentive Compatibility). *A mechanism is robust incentive compatible for estimated bidder distribution $\hat{\pi}$ and consistent set of distributions $\mathcal{P}(\hat{\pi})$ if for all $\theta, \theta' \in \Theta$ and $\pi, \pi' \in \mathcal{P}(\hat{\pi})$,*

$$\sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \pi, \theta, \pi, \omega) \geq \sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \pi, \theta', \pi', \omega)$$

Note that we can restrict our attention to settings where the bidder only reports $\pi' \in \mathcal{P}(\hat{\pi})$ by setting the allocation probability p to zero if the bidder reports $\pi' \notin \mathcal{P}(\hat{\pi})$.

Robust Mechanisms

Existing mechanism design techniques are inadequate to optimize revenue over these situations. Specifically, Bayesian mechanisms can have very unintuitive formulations consisting of multiple lotteries over the value of the external signal making it unlikely that a standard class of simple intuitive mechanisms will be able to implement revenue efficient robust mechanisms. Therefore, we will combine techniques from automated mechanism design and robust convex optimization to automate the design of robust mechanisms.

Further, while it is theoretically possible to allow bidders to report both their valuations and their beliefs, and design optimal mechanisms given this joint report, standard automated mechanism design techniques require finitely specified input, and we are explicitly interested in infinite sets of distributions. We will simplify the mechanism design process by only considering mechanisms for which the payments, x , and probabilities of allocations, p , depend only on the reported bidder types and the realization of the external signal. While this is not without loss of generality, it will be sufficient to significantly outperform ex-post mechanisms in our experiments.

Definition 10 (Optimal Robust Mechanism). *An optimal robust mechanism given an estimated distribution $\hat{\pi}$ and a consistent set of distributions $\mathcal{P}(\hat{\pi})$ is a mechanism that is an optimal solution to the following program:*

$$\begin{aligned} & \max_{x, p} \sum_{\theta, \omega} \hat{\pi}(\theta, \omega)x(\theta, \omega) \\ & \text{subject to} \\ & \sum_{\omega \in \Omega} \pi'(\omega|\theta)U(\theta, \theta, \omega) \geq 0 \quad \forall \theta \in \Theta, \pi' \in \mathcal{P}(\hat{\pi}) \\ & \sum_{\omega \in \Omega} \pi'(\omega|\theta)U(\theta, \theta, \omega) \geq \sum_{\omega \in \Omega} \pi'(\omega|\theta)U(\theta, \theta', \omega) \\ & \quad \forall \theta, \theta' \in \Theta, \pi' \in \mathcal{P}(\hat{\pi}) \\ & 0 \leq p(\theta, \omega) \leq 1 \quad \forall \theta \in \Theta, \omega \in \Omega \end{aligned}$$

Note that the linear program in Definition (10) still contains an infinite number of constraints over a, potentially, non-convex set, and therefore is in general computationally intractable. However, the following assumption allows computational tractability.

Assumption 1. *The set $\mathcal{P}(\hat{\pi})$ is such that for all $\theta \in \Theta$, $\mathcal{P}(\hat{\pi}(\cdot|\theta))$ is a convex n -polyhedron where n is polynomial in the number of bidder types.*

Assumption 1 includes very reasonable cases such as the case where for all $\theta \in \Theta$ and $\omega \in \Omega$, $\pi'(\omega|\theta) \in [\underline{\pi}(\omega|\theta), \bar{\pi}(\omega|\theta)]$. Further, any set that does not satisfy Assumption 1 can be contained in a set that does. Therefore, we can always make the assumption hold by using a larger consistent set.

Theorem 3. *For a given (p, x) and $\mathcal{P}(\hat{\pi})$ that satisfy Assumption 1, there exists a polynomial time algorithm that determines whether there exists a $\pi'(\cdot|\theta) \in \mathcal{P}(\hat{\pi}(\cdot|\theta))$ such that robust individual rationality or robust incentive compatibility is violated.*

Proof. For each $\theta \in \Theta$, solve the following linear program

$$\begin{aligned} & \min_{\pi'(\cdot|\theta)} \sum_{\omega} \pi'(\omega|\theta)(v(\theta)p(\omega, \theta) - x(\omega, \theta)) \\ & \text{subject to} \\ & \pi'(\cdot|\theta) \in \mathcal{P}(\hat{\pi}(\cdot|\theta)) \end{aligned} \tag{1}$$

Note that in the program (1), (p, x) are no longer variables but coefficients. If (1) has an objective value of less than 0, then the robust IR constraint with distribution π' is violated. If the objective value is at least 0, there is no robust IR constraint violated for θ .

There are $|\Theta|$ linear programs that must be solved, each with a polynomial number of variables and constraints, due to Assumption 1. Therefore, violated robust IR constraints can be generated in polynomial time.

Similarly for robust incentive compatibility, the following program, for all $\theta, \theta' \in \Theta$, finds violated constraints:

$$\begin{aligned} & \min_{\pi'(\cdot|\theta)} \sum_{\omega} \pi'(\omega|\theta)(v(\theta)p(\omega, \theta) - x(\omega, \theta) \\ & \quad - (v(\theta)p(\omega, \theta') - x(\omega, \theta'))) \\ & \text{subject to} \\ & \pi'(\cdot|\theta) \in \mathcal{P}(\hat{\pi}(\cdot|\theta)) \end{aligned}$$

□

Corollary 4. *If $\mathcal{P}(\hat{\pi})$ satisfies Assumption 1, the optimal robust mechanism can be computed in time polynomial in the number of types of the bidder and external signal.*

Proof. By Theorem 3, we can determine whether or not a robust IR or robust IC constraint is violated in polynomial time, and add the constraint to the linear program. There are $2|\Theta||\Omega|$ variables in the linear program in Definition 10, and there are $2|\Theta||\Omega|$ non-IC and IR constraints.

Therefore, by the ellipsoid method, the optimal robust mechanism can be computed in polynomial time (Kozlov, Tarasov, and Khachiyan 1980). □

Note that Corollary 4 states that the optimal robust mechanism is polynomial in the *number of bidder types*, not the number of bidders. The current formulation is exponential in the number of bidders. However, as stated in the introduction, the advantage of prior-dependent mechanisms is primarily for thin auctions, so we do not view this as a significant weakness of this approach.

Since for ex-post mechanisms, incentive compatibility and individual rationality are independent of the distribution, it would be expected that when we have no useful information about the distribution, the optimal robust mechanism should be equivalent to the optimal ex-post mechanism. The following proposition shows that this is indeed the case.

Proposition 5. *If for all $\theta \in \Theta$, and $\omega \in \Omega$, $\mathcal{P}(\hat{\pi})$ is such that the distribution $\pi'(\omega'|\theta) = 1$ if $\omega' = \omega$ and 0 otherwise is in $\mathcal{P}(\hat{\pi})$, then an optimal robust mechanism is an optimal ex-post mechanism for the distribution $\hat{\pi}$.*

Proof. If for all $\theta \in \Theta$ and $\omega \in \Omega$, the distribution such that $\pi'(\omega'|\theta) = 1$ if $\omega' = \omega$ and 0 otherwise is in $\mathcal{P}(\hat{\pi})$, the robust IR constraints contain the following set of constraints

$$v(\theta)p(\theta, \omega) - x(\theta, \omega) \geq 0 \quad \forall \omega \in \Omega, \theta \in \Theta$$

which implies ex-post IR. Conversely, ex-post IR implies robust IR.

By an identical argument, the robust IC constraints imply the ex-post IC constraints, and vice-versa. \square

ϵ -Robust Mechanisms

While, so far we have been assuming that there is a well defined set, $\mathcal{P}(\hat{\pi})$, such that the mechanism designer can guarantee that the true distribution, π , is in the set, this is unlikely to be a realistic assumption in practice. It is far more reasonable that the mechanism designer would have a set such that the true distribution is in the set with some high probability. If this is the case, we can still design mechanisms that are likely to outperform weakly prior dependent mechanisms, such as ex-post mechanisms, by relaxing the requirement that the mechanism be always incentive compatible and always individually rational. We will define the set of ϵ -consistent distributions as follows.

Definition 11 (Set of ϵ -Consistent Distributions). *A subset $\mathcal{P}_\epsilon(\hat{\pi}) = \mathcal{P}_\epsilon(\{\hat{\pi}_\theta, \hat{\pi}(\cdot|\cdot)\}) \subseteq P(\Theta) \times \prod_{\theta \in \Theta} P(\Omega)$ is an ϵ -consistent set of distributions for the estimated distribution $\{\hat{\pi}_\theta, \hat{\pi}(\cdot|\cdot)\}$ if the true distribution, $\{\pi_\theta, \pi(\cdot|\cdot)\}$, is in $\mathcal{P}_\epsilon(\hat{\pi})$ with probability $1 - \epsilon$.*

Now we can define the notion of ϵ -robust individual rationality and incentive compatibility. These definitions are analogous to Definitions 8 and 9.

Definition 12 (ϵ -Robust Individual Rationality). *A mechanism is ϵ -robust individually rational for estimated bidder distribution $\hat{\pi}$ and ϵ -consistent set of distributions $\mathcal{P}_\epsilon(\hat{\pi})$ if for all $\theta \in \Theta$ and $\pi \in \mathcal{P}_\epsilon(\hat{\pi})$,*

$$\sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \pi, \theta, \pi, \omega) \geq 0$$




Figure 1: The performance of the ex-post, ϵ -robust, and Bayesian mechanisms using the estimated distribution. All revenue is scaled by the full social surplus, denoted as 1. Note that the Number of Samples is in log scale. The parameters used were as follows: Correlation = .5, $\epsilon = .05$. Each experiment was repeated 200 times, and the 95% confidence interval is included for the ϵ -robust and ex-post mechanisms. The Bayesian mechanism confidence interval is off the plot.

Definition 13 (ϵ -Robust Incentive Compatibility). *A mechanism is ϵ -robust incentive compatible for estimated bidder distribution $\hat{\pi}$ and ϵ -consistent set of distributions $\mathcal{P}_\epsilon(\hat{\pi})$ if for all $\theta, \theta' \in \Theta$ and $\pi, \pi' \in \mathcal{P}_\epsilon(\hat{\pi})$,*

$$\sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \pi, \theta, \pi, \omega) \geq \sum_{\omega \in \Omega} \pi(\omega|\theta)U(\theta, \pi, \theta', \pi', \omega)$$

Similarly to Definition 10, we can define the optimal ϵ -robust mechanism. Again, for tractability of mechanism design, and to obtain more practical mechanisms, we will restrict attention to mechanisms that only depend on the reported bidder type and the external signal.

Definition 14 (Optimal ϵ -Robust Mechanism). *An optimal ϵ -robust mechanism given an estimated distribution $\hat{\pi}$ and an ϵ -consistent set of distributions $\mathcal{P}_\epsilon(\hat{\pi})$ is a mechanism that is an optimal solution to the following program:*

$$\max_{x(\theta, \omega), p(\theta, \omega)} \sum_{\theta, \omega} \hat{\pi}(\theta, \omega)x(\theta, \omega)$$

subject to

$$\sum_{\omega \in \Omega} \pi'(\omega|\theta)U(\theta, \theta, \omega) \geq 0 \quad \forall \theta \in \Theta, \pi' \in \mathcal{P}_\epsilon(\hat{\pi})$$

$$\sum_{\omega \in \Omega} \pi'(\omega|\theta)U(\theta, \theta, \omega) \geq \sum_{\omega \in \Omega} \pi'(\omega|\theta)U(\theta, \theta', \omega)$$

$$\forall \theta, \theta' \in \Theta, \pi' \in \mathcal{P}_\epsilon(\hat{\pi})$$

$$0 \leq p(\theta, \omega) \leq 1 \quad \forall \theta \in \Theta, \omega \in \Omega$$

Proposition 6. *If $\mathcal{P}_\epsilon(\hat{\pi})$ satisfies Assumption 1, the optimal ϵ -robust mechanism can be computed in time polynomial in the number of bidder types and external signals.*




Figure 2: The performance of the robust and ex-post mechanisms using the estimated distribution. All revenue is scaled by the full social surplus, which is denoted as 1. Shown are the 95% confidence intervals for the robust mechanism. For any variable not explicitly shown the following values were used: Correlation = .5, $\epsilon = .05$

Experimental Results

Throughout the experiments, we have a single bidder with type $\theta \in \{1, 2, \dots, 10\}$ and valuation $v(\theta) = \theta$. The external signal is $\omega \in \{1, 2, \dots, 10\}$. We model the true distribution as a categorical distribution with 10×10 elements, with each element corresponding to a tuple (θ, ω) .

There are not, to our knowledge, standard distributions to test correlated mechanism design procedures available, so we use a discretized bi-variate normal distribution. Specifically, we discretize the area under the bi-variate standard normal distribution between $[-1.96, 1.96]$ in both dimensions as a 10×10 grid and normalize. We chose the bi-variate normal distribution for its broad relevance to many empirically observed distributions and the ability to easily vary the correlation. Note that the bi-variate normal distribution always satisfies the Cremer-McLean condition if the correlation is positive.

To estimate the distribution, we sample from the true distribution and use Bayesian updating with a maximally uninformative Dirichlet prior ($\alpha = [1, \dots, 1]$) to arrive at a Dirichlet posterior over the distribution of bidder types and external signals. We then calculate empirical confidence intervals by sampling from the Dirichlet posterior and observing the $\epsilon/(2 * 10 * 10)$ and $(1 - \epsilon/(2 * 10 * 10))$ quantiles for each element of the conditional distributions $\pi(\omega|\theta)$ and use the quantiles as the ϵ -consistent set. Note that we do not simply use the $\epsilon/2$ and $(1 - \epsilon/2)$ quantiles due to jointly estimating confidence intervals for 100 variables; the expression for the quantiles above is based on applying a union bound.

For our experiments, we solve for the optimal ex-post, ϵ -robust, and Bayesian mechanisms given our estimated distribution $\hat{\pi}$ and our ϵ -consistent set. Given that both the optimal ϵ -robust and Bayesian mechanisms can fail to be incentive compatible and/or individually rational due to the difference between the estimated and true distribution, we compute the optimal action for the bidder: either report truthfully, strategically misreport, or do not participate. We then calculate the revenue accordingly.

In Figure 1, we show the performance of the optimal ex-post, robust, and Bayesian mechanisms using our estimated distribution as we increase the number of samples. We report confidence intervals for both the ex-post mechanisms and the robust mechanisms; however for the Bayesian mechanisms, the confidence intervals were off the chart. Figure 1 demonstrates how badly the Bayesian mechanism performs when the distribution is not exactly known. Even after 10,000 samples from the true distribution, the Bayesian mechanism fails to outperform the ex-post mechanism. By contrast, the optimal ϵ -robust mechanism generates revenue indistinguishable from the ex-post mechanism for low numbers of samples, while significantly outperforming the ex-post mechanism starting at about 10,000 samples.

In Figures 2a and 2b, we vary correlation and ϵ with increasing numbers of samples. As the bidder type and external signal are more highly correlated, the ϵ -robust mechanism requires fewer samples to perform well, Figure 2a. Also, we see that the ϵ -robust mechanism is not very sensitive to the choice of ϵ , Figure 2b, a fact that we attribute to being overly cautious in requiring all elements of the distribution to be in the bounded intervals.

In Figure 2c, we bin some of the external signals together in order to explore the trade-off between estimating a lower dimensional distribution and constructing a mechanism over the full information. Specifically, for the Signals = 2 case we put all of $\omega = \{1, \dots, 5\}$ into one bin and $\omega = \{6, \dots, 10\}$ to a second bin. Note the true signal still has 10 values, we are just binning the observed signal. We find that for a low number of samples, we do much better by binning the external signal, but, while difficult to see on the plot, at higher numbers of samples, it is better to use the full distribution.

Note that we consider the results here to be lower bounds on the performance of optimal ϵ -robust mechanisms. We assume a completely uninformative prior, increasing the required sample size. Further, we have used a naïve distribution estimation procedure, so there is likely significant room to improve upon the estimation.

Conclusion

We have presented a new paradigm in mechanism design that formally addresses the problem of uncertainty about the bidder distribution. We do this by introducing a new notion of individual rationality and incentive compatibility that spans a spectrum between ex-post and Bayesian notions and demonstrate that this class of mechanisms can be optimized over efficiently. We further relax these notions to allow for some probability of mechanisms failing to be IR or IC. While much work needs to be done in exploring the performance of these mechanisms, experimental results suggest that they may generate significantly higher revenue than either ex-post or Bayesian mechanisms when the underlying correlated distribution is estimated.

Acknowledgements

Albert and Conitzer are thankful for support from NSF under awards IIS-1527434 and CCF-1337215, ARO under grants W911NF-12-1-0550 and W911NF-11-1-0332, and a Guggenheim Fellowship.

This work has partially taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (CNS-1330072, CNS-1305287, IIS-1637736, IIS-1651089), ONR (21C184-01), and AFOSR (FA9550-14-1-0087). Peter Stone serves on the Board of Directors of, Cogitai, Inc. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research.

References

- Aghassi, M., and Bertsimas, D. 2006. Robust game theory. *Mathematical Programming* 107(1):231–273.
- Albert, M.; Conitzer, V.; and Lopomo, G. 2015. Assessing the robustness of Cremer-McLean with automated mechanism design. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI '15, 763–769. Austin, TX, USA: AAAI Press.
- Albert, M.; Conitzer, V.; and Lopomo, G. 2016. Maximizing revenue with limited correlation: The cost of ex-post incentive compatibility. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI-16)*.
- Bergemann, D., and Morris, S. 2005. Robust mechanism design. *Econometrica* 73(6):1771–1813.
- Bertsimas, D., and Sim, M. 2004. The price of robustness. *Operations Research* 52(1):35–53.
- Bulow, J., and Klemperer, P. 1996. Auctions versus negotiations. *The American Economic Review* 86(1):pp. 180–194.
- Conitzer, V., and Sandholm, T. 2002. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence*, UAI '02, 103–110. Alberta, Canada: Morgan Kaufmann.
- Conitzer, V., and Sandholm, T. 2004. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, EC '04, 132–141. New York, NY, USA: ACM.
- Cremer, J., and McLean, R. P. 1985. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* 53(2):pp. 345–361.
- Cremer, J., and McLean, R. P. 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56(6):pp. 1247–1257.
- Gibbons, R. 1992. *Game Theory for Applied Economists*. Princeton University Press.
- Government Accountability Office. 2013. Reverse auctions: Guidance is needed to maximize competition and achieve cost savings. Technical report, Government Accountability Office.
- Guo, M., and Conitzer, V. 2010. Computationally feasible automated mechanism design: General approach and case studies. In *Proceedings of the Twentieth AAAI Conference on Artificial Intelligence*, AAAI '10, 1676–1679. Atlanta, GA, USA: AAAI Press. NECTAR track.
- Hartline, J. D., and Roughgarden, T. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, EC '09, 225–234. New York, NY, USA: ACM.
- Hartline, J. D. 2014. Mechanism design and approximation. <http://jasonhartline.com/MDnA/>.
- Interactive Advertising Bureau (IAB). 2015. IAB programmatic revenue report 2014 results. http://www.iab.net/media/file/PwC_IAB_Programmatic_Study.pdf.
- Kozlov, M.; Tarasov, S.; and Khachiyan, L. 1980. The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics* 20(5):223 – 228.
- Lopomo, G. 2001. Optimality and robustness of the English auction. *Games and Economic Behavior* 36(2):219–240.
- McAfee, R. P., and Reny, P. J. 1992. Correlated information and mechanism design. *Econometrica* 60(2):pp. 395–421.
- Morgenstern, J., and Roughgarden, T. 2015. The pseudo-dimension of nearly-optimal auctions. In *NIPS*, Forthcoming.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1):58–73.
- NASA. 2014. NASA chooses american companies to transport U.S. astronauts to international space station. <https://www.nasa.gov/>.
- Roughgarden, T., and Talgam-Cohen, I. 2013. Optimal and near-optimal mechanism design with interdependent values. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, 767–784. New York, NY, USA: ACM.
- Sandholm, T., and Likhodedov, A. 2015. Automated design of revenue-maximizing combinatorial auctions. *Operations Research* 63(5):1000–1025.