

Assessing the Robustness of Cremer-McLean with Automated Mechanism Design*

Michael Albert

The Ohio State University
Fisher School of Business
2100 Neil Ave., Fisher Hall 844
Columbus, OH 43210, USA
Michael.Albert@fisher.osu.edu

Vincent Conitzer

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

Giuseppe Lopomo

Duke University
The Fuqua School of Business
100 Fuqua Drive
Durham, NC 27708, USA
glopomo@duke.edu

Abstract

In a classic result in the mechanism design literature, Cremer and McLean (1985) show that if buyers' valuations are sufficiently correlated, a mechanism exists that allows the seller to extract the full surplus from efficient allocation as revenue. This result is commonly seen as "too good to be true" (in practice), casting doubt on its modeling assumptions. In this paper, we use an automated mechanism design approach to assess how sensitive the Cremer-McLean result is to relaxing its main technical assumption. That assumption implies that each valuation that a bidder can have results in a unique conditional distribution over the external signal(s). We relax this, allowing multiple valuations to be consistent with the same distribution over the external signal(s). Using similar insights to Cremer-McLean, we provide a highly efficient algorithm for computing the optimal revenue in this more general case. Using this algorithm, we observe that indeed, as the number of valuations consistent with a distribution grows, the optimal revenue quickly drops to that of a reserve-price mechanism. Thus, automated mechanism design allows us to gain insight into the precise sense in which Cremer-McLean is "too good to be true."

Introduction

One of the chief problems in mechanism design, specifically in the case of a monopolistic seller with one or more buyers, is to provide incentives for buyers to share their private information. Typically this requires the seller to share some of the expected surplus from allocating the item with the buyer(s) in exchange for his private information. However, Cremer and McLean (1985) show that if buyers' valuations satisfy a correlation condition, then this is not necessary: the seller is able to extract the full surplus in expectation. This result is highly regarded, but it has perhaps not had the impact one might have expected. But why not?

A standard assumption in the mechanism design literature is that buyer valuations are independently distributed, and therefore, knowing the valuation of one agent gives no information about other agents' valuations. This independence is clearly violated in auctions in which there is a common value

component to the item that is estimated by each of the buyers individually, such as those for oil drilling rights. More generally, for any item for which there is a resale market, an agent's valuation will generally include a component that is dependent on the valuation of the rest of the community, and as a consequence, will be correlated to the extent that potential resale is a motivation for purchase. Hence, assuming some degree of correlation seems quite reasonable.

The correlation assumption in the Cremer-McLean result requires that every buyer valuation corresponds to a unique conditional distribution over the external signal(s) (e.g., the other buyer's valuations), and further, that these conditional distributions are linearly independent of each other. In this paper, we explore a relaxation of this correlation assumption. Namely, we let a buyer type consist of a valuation and a distribution over the external signal(s), and allow multiple types to have the same distribution (but different valuations). More specifically, buyer types are partitioned into subsets, where within each subset all types have the same conditional distribution over the external signal(s), but within the subset, the conditional distributions are identical, implying that within such a subset, the buyer's valuation is independent of the external signal(s). This setup allows us to gradually move away from the Cremer-McLean assumption and analyze how quickly the revenue drops as we do so, thereby assessing the brittleness of the Cremer-McLean result.

To achieve this, we first obtain some structural insights into optimal mechanisms for our relaxed setting, generalizing the corresponding insights from Cremer-McLean. These insights lead us to an algorithm for computing the optimal revenue that is much more efficient than solving the full linear program. This allows us to run much larger simulations of distributions with zero partial correlation over subsets of various sizes.

We find that as the size of the subsets of buyer types corresponding to the same conditional distribution increases the optimal revenue quickly converges to that achievable under a significantly simpler mechanism, namely a second price auction with a reserve. Thus, in addition to increasing our understanding of how to solve for optimal mechanisms under a broader class of distributions, our result informs the practicality of full surplus extraction along the lines of Cremer and McLean (1985). Our zero partial correlation assumption can be interpreted as there being

*Conitzer thanks NSF and ARO for support under grants CCF-1101659, IIS-0953756, CCF-1337215, W911NF-12-1-0550, and W911NF-11-1-0332.

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

some degree of uncertainty about the buyer's beliefs conditional on his valuation, as seems likely to be the case in practice. Therefore, our result that the revenue from the optimal mechanism quickly converges to that of simple mechanisms helps explain why the Cremer-McLean mechanism is rarely, if ever, implemented. As such, our result contributes to the literature on the optimality of simple, robust mechanisms (Hartline and Roughgarden 1998; Lopomo 1998; Ronen and Saberi 2002). Methodologically, our work also fits within the framework of *automated mechanism design* (Conitzer and Sandholm 2002; 2004), and more specifically the use thereof to build intuition for classical results in mechanism design (Guo and Conitzer 2010; Likhodedov and Sandholm 2005).

Full Surplus Extraction with Correlation (Review of Cremer and McLean (1985))

In Cremer and McLean (1985), the optimal mechanism for the general case where a single monopolistic seller is selling a quantity x of a good to some number of buyers n is shown to extract the full surplus as revenue, in expectation. For the sake of clarity and exposition, in this paper, we restrict our attention to a simplified setting where the seller has a single indivisible good that the seller values at zero, there is a single buyer, and there is an external signal (which, in this case, does not correspond to another buyer's valuation—but this will work just as well for Cremer-McLean) that is correlated with the buyer's valuation. Our analysis can be naturally extended to the case of more than one buyer.

The (risk-neutral) buyer has a valuation for the object drawn from a discrete set of types denoted by the set $V = \{v_1, v_2, \dots, v_{|V|}\} \subset [0, \bar{v}]$ where $v_i < v_{i+1}$ for all i . The buyer's type, which may include information in addition to his valuation, is denoted by $\theta \in \Theta = \{1, 2, \dots, |\Theta|\}$, where $v(\theta)$ denotes the valuation associated with this type.

In addition to the buyer, there is a discrete external signal $\omega \in \Omega = \{1, 2, \dots, |\Omega|\}$ where $|\Omega| \geq |V|$. The buyer's type may be correlated with the external signal, with the joint distribution over θ and ω given by $\pi(\theta, \omega)$. Hence, for our purposes, a type θ for a buyer consists of a valuation $v(\theta)$ as well as a conditional distribution over the external signal, $\pi(\omega|\theta)$. It is assumed that the buyer and the seller possess a common prior.

A (direct-revelation) mechanism is defined by the probability that the seller will allocate the item to the buyer, $q(\hat{\theta}, \omega)$, and a monetary transfer from the buyer to the seller, $m(\hat{\theta}, \omega)$, for every buyer-reported outcome $\hat{\theta} \in \Theta$ and observed $\omega \in \Omega$. The mechanism is permitted to require a payment from the buyer to the seller whether or not the bidder is allocated the item, though we do require individual rationality, i.e., the buyer should in expectation receive non-negative utility for any type.

Definition 1 (Buyer's Utility). *Given true type $\theta \in \Theta$ and reported type $\hat{\theta} \in \Theta$, the buyer's expected utility under mechanism (q, m) is:*

$$U(\theta, \hat{\theta}) = \sum_{\omega} (v(\theta)q(\hat{\theta}, \omega) - m(\hat{\theta}, \omega))\pi(\omega|\theta) \quad (1)$$

In the case of full information, the seller observes the buyer's type and chooses a mechanism such that $q(\theta, \omega) = 1$ and $m(\theta, \omega) = v(\theta)$ for all $\theta \in \Theta$ and $\omega \in \Omega$. This ensures that the seller's revenue is optimal; he is able to extract all surplus from the buyer. Hence, the seller's expected revenue under full information is $\sum_{\theta, \omega} v(\theta)\pi(\theta, \omega)$.

In the case of incomplete information, the seller is only able to directly observe the realization of the outside signal ω , so he must induce the agent to report his private information. The seller can use the following linear program to find an optimal (revenue-maximizing) mechanism:

Definition 2 (The Optimal Mechanism Under Incomplete Information). *Optimal mechanisms under incomplete information correspond to optimal solutions to the following linear program:*

$$\max_{q(\theta, \omega), m(\theta, \omega)} \sum_{\theta, \omega} m(\theta, \omega) \pi(\theta, \omega) \quad (2)$$

subject to

$$U(\theta, \theta) \geq 0 \quad \forall \theta \in \Theta \quad (3)$$

$$U(\theta, \theta) \geq U(\theta, \hat{\theta}) \quad \forall \theta \in \Theta, \hat{\theta} \in \Theta \quad (4)$$

$$0 \leq q(\theta, \omega) \leq 1 \quad \forall \theta \in \Theta, \omega \in \Omega. \quad (5)$$

An optimal solution to the program in Definition 2 is guaranteed to maximize the expected revenue of the seller. However, in general, it will not generate revenue equivalent to the case of full information; the seller will have to share some of the surplus with the buyer. Cremer and McLean (1985) impose the following additional assumption:

Assumption 1. *For all $\theta \in \Theta$, let Γ be the following matrix whose rows are indexed by the $|\Omega|$ elements of Ω , and whose columns are indexed by the $|\Theta|$ elements of Θ :*

$$\Gamma = \begin{bmatrix} \pi(1|1) & \cdots & \pi(|\Omega|1) \\ \vdots & \ddots & \vdots \\ \pi(1||\Theta|) & \cdots & \pi(|\Omega||\Theta|) \end{bmatrix} \quad (6)$$

Γ has rank $|\Theta|$.

Assumption 1 does not require that the correlations have any structure or magnitude other than full rank. Under this assumption, they show that the expected revenue from the program in Definition 2 and the case of full information are identical, i.e., the seller can extract the full surplus.

Theorem 1 (Cremer and McLean (1985)). *Under Assumption 1, there exists a solution to the program in Definition 2 with an objective value of:*

$$\sum_{\theta, \omega} \pi(\theta, \omega)m(\theta, \omega) = \sum_{\theta, \omega} \pi(\theta, \omega)v(\theta) \quad (7)$$

A proof of Theorem 1 can be found in Cremer and McLean (1985) or Krishna (2009). Intuitively, what allows the seller to extract so much surplus is her ability to severely punish the buyer for not reporting his value truthfully, by providing a payoff that is dependent on the realization of the external signal. Thanks to Assumption 1, the seller is able to construct such payoffs in a way that incentivizes the buyer both to report his true type and to pay his reported valuation for the item in expectation. The techniques we use below will make it clearer why Assumption 1 allows this.




Figure 1: Each dot represents one discrete buyer type. On the horizontal axis is the space of conditional distributions over the external signal (collapsed to one dimension). Assumption 1 implies that no two buyer types are aligned vertically.




Figure 2: Each conditional distribution over the external signal is associated with more than one buyer valuation. Assumption 1 is not satisfied because vertically aligned types have linearly dependent (in fact, identical) conditional distributions.

Relaxing Assumption 1

Assumption 1, while being satisfied by a “random” distribution with probability one, in practice adds significant fragility to the Cremer-McLean mechanism. The mechanism relies on a very precise specification of the joint distribution over buyer types and external states, being able to associate with each buyer type a unique conditional distribution over the external signal. In practice, it will be hard to estimate this distribution so precisely. In particular, the buyer may well perceive the conditional distribution over the external signal to be slightly different.

Figure 1 illustrates the crucial feature of Assumption 1 that there are no two types with the same conditional distribution over the external signal but with different valuations. In contrast, we allow that there exists a range of possible conditional distributions for each unique buyer valuation. We also allow for the possibility that these distributions overlap, i.e., a conditional distribution may not uniquely identify the buyer valuation, breaking Assumption 1. This is illustrated in Figure 2, where valuations are consistent with multiple conditional distributions, and vice versa. Formally,

we let the type space be partitioned as $\Theta = \bigcup_{i=1}^k X_i$ (where $X_i \cap X_j = \emptyset$ when $i \neq j$), such that θ and θ' lead to the same conditional distribution over Ω if and only if they are in the same X_i . We do hold on to part of Assumption 1, in that we assume that the distinct distributions corresponding to different X_i are linearly independent.

Assumption 2. Define $\pi_i(\omega) = \pi(\omega|\theta)$ for any $\theta \in X_i$. $\{\pi_1, \dots, \pi_k\}$ are linearly independent.

A More Efficient Algorithm

Again, whether or not Assumptions 1 and/or 2 are satisfied, solving the linear program in Definition 2 will give an optimal mechanism for the seller. However, solving the LP in a straightforward way does not scale to large instances of the problem. E.g., the algorithm by Karmarkar (1984) gives a runtime bound of $O((|\Theta||\Omega|)^{3.5}L)$, where L is the length of the encoding. This still significantly limits what we can do, particularly when the external signal is multidimensional (for example, predicting multiple other valuations). In this section, we develop a significantly more efficient algorithm for finding the optimal revenue in our setting. To do so, we must first define some useful concepts. First, we define a reserve price mechanism on a restricted subset of the types.

Definition 3 (Reserve Price Mechanism). A mechanism (q, m) is a reserve price mechanism when restricted to subset of types $\Theta' \subseteq \Theta$ if the following holds. There exists some $\theta^* \in \Theta'$ such that for all $\omega \in \Omega$, (1) for all $\theta \in \Theta'$ with $v(\theta) \geq v(\theta^*)$, $q(\theta, \omega) = 1$ and $m(\theta, \omega) = v(\theta^*)$, and (2) for all $\theta \in \Theta'$ with $v(\theta) < v(\theta^*)$, $q(\theta, \omega) = 0$ and $m(\theta, \omega) = 0$. Note that this mechanism is optimal among such mechanisms if and only if

$$\theta^* \in \arg \max_{\theta^* \in \Theta'} \sum_{\theta, \omega} \pi(\theta, \omega) v(\theta^*) \mathbb{1}_{\{v(\theta) \geq v(\theta^*)\}}. \quad (8)$$

Next, we recall the definition of a proper scoring rule, of which we will make use in what follows. Informally, a proper scoring rule rewards a forecaster in a way that makes it optimal for him to reveal his true subjective distribution.

Definition 4 (Proper Scoring Rule). Let $\Omega = \{1, \dots, |\Omega|\}$ be a sample space consisting of a finite number of mutually exclusive events and let $P = \{\mathbf{p} = (p_1, \dots, p_{|\Omega|}) : p_1, \dots, p_{|\Omega|} \geq 0, p_1 + \dots + p_{|\Omega|} = 1\}$ be the set of probability distributions over Ω . Then a scoring rule is a set of $|\Omega|$ functions $H(\cdot, \omega) : P \rightarrow \mathbb{R}$ for $\omega \in \Omega$. A proper scoring rule is one such that for all $\mathbf{p} = (p_1, \dots, p_{|\Omega|})$ and $\mathbf{p}' = (p'_1, \dots, p'_{|\Omega|})$:

$$\sum_{\omega \in \Omega} p_{\omega} H(\mathbf{p}, \omega) \geq \sum_{\omega \in \Omega} p_{\omega} H(\mathbf{p}', \omega) \quad (9)$$

A strictly proper scoring rule is one such that Equation 9 is only satisfied with equality when $\mathbf{p}' = \mathbf{p}$.

Proper scoring rules have been elegantly characterized:

Theorem 2 (Gneiting and Raftery (2007)). A scoring rule is proper if and only if there exists a convex function $G : P \rightarrow \mathbb{R}$ such that

$$H(\mathbf{p}, \omega) = G(\mathbf{p}) - \sum_{\omega' \in \Omega} p_{\omega'} G'_{\omega'}(\mathbf{p}) + G'_{\omega}(\mathbf{p}) \quad (10)$$

where $G'_\omega(\mathbf{p})$ is the ω th component of the subgradient of G . H is strictly proper if and only if G is convex. Note that $\sum_{\omega \in \Omega} p_\omega H(\mathbf{p}, \omega) = G(\mathbf{p})$ —that is, G gives the expected reward for truthful reporting.

We are now in a position to define our mechanism. Intuitively, we first obtain the optimal reserve price mechanism for each X_i , and subsequently add a strictly proper scoring rule—one that gives an expected value of 0 for truthfully reporting any of the π_i —to remove the incentive to misreport into a different X_i . The formal description follows. That this mechanism is well defined (under Assumption 2) will be shown in the subsequent theorem.

Definition 5. Let q_i and m_i be defined for any i ($1 \leq i \leq k$), such that for $\omega \in \Omega$ and $\theta \in X_i$, $q_i(\theta, \omega)$ and $m_i(\theta, \omega)$ are equivalent to the optimal reserve price mechanism for the subsets of type in X_i . (Note that (q_i, m_i) as a whole may fail to be incentive compatible, because an agent may misreport to a different X_i .) Then, define the following strictly convex function:

$$G^*(\mathbf{p}) = \sum_{\omega \in \Omega} (p_\omega^2 - a_\omega p_\omega) \quad (11)$$

where the a_i are given by any solution to the following:

$$\begin{bmatrix} \pi_1(1) & \dots & \pi_1(|\Omega|) \\ \vdots & \ddots & \vdots \\ \pi_k(1) & \dots & \pi_k(|\Omega|) \end{bmatrix} \begin{bmatrix} a_1 \\ \vdots \\ a_{|\Omega|} \end{bmatrix} = \begin{bmatrix} \sum_{\omega} (\pi_1(\omega))^2 \\ \vdots \\ \sum_{\omega} (\pi_k(\omega))^2 \end{bmatrix} \quad (12)$$

Then, define a strictly proper scoring rule H^* using G^* as in Equation 10, and let

$$\kappa = \min_{j \neq l, j, l \in \{1, \dots, k\}} \left| \sum_{\omega} \pi_l(\omega) H^*(\pi_j, \omega) \right|$$

be the minimum loss from misreporting the distribution. Let $K = \bar{v}/\kappa$, let $G(\mathbf{p}) = KG^*(\mathbf{p})$, and define H as in Equation 10 using G . Let $m^*(\theta, \omega) = -H(\pi_i, \omega)$ for $\theta \in X_i$.

Finally, for $\theta \in X_i$ and $\omega \in \Omega$, our mechanism is defined as $q(\theta, \omega) = q_i(\theta, \omega)$ and $m(\theta, \omega) = m^*(\theta, \omega) + m_i(\theta, \omega)$.

Theorem 3. If Assumption 2 holds, then the mechanism in Definition 5 is well defined and constitutes an optimal solution to the linear program in Definition 2.

Proof. We first show that the mechanism is well defined and that the scoring rules are proper. First, note that G^* as given by Equation 11 is strictly convex. Since π_1, \dots, π_k are linearly independent by Assumption 2, Equation 12 has at least one solution. Also by Equation 12, for all $i \in \{1, \dots, k\}$, we have

$$G^*(\pi_i) = \sum_{\omega \in \Omega} (\pi_i(\omega)^2 - a_\omega \pi_i(\omega)) = 0$$

Because G^* is strictly convex, by Theorem 2, H^* is strictly proper, so $\sum_{\omega} \pi_l(\omega) H(\pi_j, \omega) < 0$ for all $l \neq j$, implying $\kappa > 0$. Therefore, K is well defined and positive. Since $K > 0$, G is strictly convex, implying that H is also a strictly proper scoring rule with expected value 0 for truthfully reporting one of the π_i .

We next show that the mechanism satisfies individual rationality and incentive compatibility. Let the buyer's true type be $\theta \in X_i$ and his reported type by $\hat{\theta} \in X_j$. By Equation 1,

$$U(\theta, \hat{\theta}) = \sum_{\omega} (v(\theta)q_j(\hat{\theta}, \omega) - m_j(\hat{\theta}, \omega) - m^*(\hat{\theta}, \omega))\pi(\omega|\theta),$$

and if we define

$$U'(\theta, \hat{\theta}) = \sum_{\omega} (v(\theta)q_i(\hat{\theta}, \omega) - m_i(\hat{\theta}, \omega))\pi(\omega|\theta),$$

then $U(\theta, \hat{\theta}) = U'(\theta, \hat{\theta}) + \sum_{\omega} \pi(\omega|\theta)H(\pi_j, \omega)$. If the buyer reports truthfully ($\hat{\theta} = \theta$), then $U(\theta, \theta) = U'(\theta, \theta)$ because H gives expected value 0 for truthfully reporting one of the π_i . Since $U'(\theta, \theta)$ is just the utility of a buyer facing a reserve price, Equation 3, the individual rationality constraint in the program in Definition 2, is satisfied.

We now move on to incentive compatibility. Suppose $\theta \neq \hat{\theta}$; then, there are two possibilities. First, assume that $\hat{\theta} \in X_i$ as well. Then, $U(\theta, \hat{\theta}) = U'(\theta, \hat{\theta})$. $U'(\theta, \hat{\theta})$ is the utility for a buyer under a reserve price mechanism, so Equation 4, the incentive compatibility constraint in the program in Definition 2, is satisfied, since reserve price mechanisms satisfy incentive compatibility. Second, assume that $\hat{\theta} \in X_j$ with $j \neq i$. Because U' consists of reserve price mechanisms, the maximum that a buyer can hope to profit under U' from misreporting his true type is \bar{v} . So,

$$\begin{aligned} U(\theta, \hat{\theta}) &= U'(\theta, \hat{\theta}) + \sum_{\omega} \pi(\omega|\theta)H(\pi_j, \omega) \\ &\leq U'(\theta, \theta) + \bar{v} + \sum_{\omega} \pi(\omega|\theta)H(\pi_j, \omega) \\ &\leq U'(\theta, \theta) + \bar{v} + (\bar{v}/\kappa) \sum_{\omega} \pi(\omega|\theta)H^*(\pi_j, \omega) \\ &\leq U'(\theta, \theta) = U(\theta, \theta) \end{aligned}$$

by the definition of κ . Therefore, the incentive constraint defined by Equation 4 is satisfied.

We have thus shown that our mechanism constitutes a feasible solution to the program in Definition 2. All that remains to show is that it is optimal. The objective value of our mechanism (Equation 2) is:

$$\begin{aligned} \sum_{\theta, \omega} m(\theta, \omega)\pi(\theta, \omega) &= \sum_{\theta, \omega} (m'(\theta, \omega) + m^*(\theta, \omega))\pi(\theta, \omega) \\ &= \sum_{\theta, \omega} m'(\theta, \omega)\pi(\theta, \omega) + \sum_{\theta} \pi(\theta) \sum_{\omega} m^*(\theta, \omega)\pi(\omega|\theta) \\ &= \sum_{\theta, \omega} m'(\theta, \omega)\pi(\theta, \omega) \end{aligned} \quad (13)$$

where the last equality follows from the fact that H gives expected value 0 for truthfully reporting one of the π_i . Therefore, the expected revenue of our mechanism is a weighted combination of the expected revenues of a set of optimal reserve price mechanisms, one for each X_i . More precisely, conditional on the buyer's true type being in X_i , the expected revenue of our mechanism is that of the optimal reserve price mechanism for the conditional distribution on the

valuations in X_i . However, this is the maximum revenue the seller could hope for, even if she directly observed X_i . This is because an optimal reserve price mechanism is revenue optimal for a single buyer whose valuation is independent of all observable signals (Myerson 1991). It follows that our mechanism maximizes expected revenue. \square

It is worth pointing out that Cremer-McLean follows as the special case where $|X_i| = 1$ for all i . In this case, of course, the reserve price for X_i will be set to the unique valuation in X_i .

Definition 5 corresponds straightforwardly to an algorithm for computing its mechanism.

Theorem 4 (Runtime). *Under Assumption 2, the algorithm corresponding to Definition 5 can be used to find an optimal mechanism in $O(|\Omega|^2 k)$ steps.*

Proof. The most computationally intensive step of the algorithm is computing the coefficients of G^* using Equation 12, which can be done in at most $O(|\Omega|^2 k)$ steps using QR decomposition. \square

However, the precise solution to Equation 12 does not affect the optimal revenue. Hence, if we are only interested in computing the optimal revenue (without explicitly finding a mechanism that attains this revenue), then we only need to compute the optimal reserve price mechanisms. This is indeed the situation in which we find ourselves, with our goal of assessing how quickly the optimal revenue drops off as we relax Assumption 1 from Cremer-McLean.

Theorem 5 (Runtime, optimal revenue only). *Under Assumption 2, the optimal objective value of the program in Definition 2 can be calculated in $O(|\Omega||\Theta|)$ steps by computing each m_i as in Definition 5 for $i \leq i \leq k$.*

Proof. From Equation 13, it follows that the optimal objective value of the program in Definition 2 can be calculated using just the values of m_i , where each m_i is calculated separately for each X_i ($i \in \{1, \dots, k\}$). The most computationally intensive step is calculating the conditional distributions, π_i , to determine the subsets X_i , and the marginal distribution over types, $\pi(\theta)$; these can be computed in $O(|\Omega||\Theta|)$ steps.

Then we can calculate the revenue and reserve price for the optimal reserve price mechanism in $O(|X_i|)$ steps for each X_i . To see this, start with $\theta_{\max} \in X_i$ such that $v(\theta_{\max}) \geq v(\theta)$ for all $\theta \in X_i$. Then, if we set $v(\theta_{\max})$ as the reserve price and $\pi_{\max} = \pi(\theta)$, then the revenue is $\pi_{\max} v(\theta_{\max})$. Record this value, type, and π_{\max} , and then repeat the same procedure with the next highest type, i.e., $\theta'_{\max} \in X_i \setminus \{\theta_{\max}\}$, and set $\pi'_{\max} = \pi_{\max} + \pi(\theta'_{\max})$. Repeat until $X'_i = \emptyset$. The maximum value realized across all types is the revenue for an optimal reserve price mechanism under X_i , requiring $O(|X_i|)$ steps to calculate. The revenue from the mechanism as a whole is just the sum of the revenues of the individual optimal reserve price mechanisms. If we do this for every X_i , it takes $O(\sum_i |X_i|) = O(|\Theta|)$ steps. Therefore, the revenue from the optimal mechanism can be calculated in $O(|\Omega||\Theta|)$ steps. \square




Figure 3: Runtimes of the various algorithms. This demonstrates the scaling in buyer types of our algorithms compared to the naive method, so there is no distributional overlap (i.e., $|\Theta| = |\Omega| = k$, or Assumption 1 holds). Distributions and valuations are randomly generated 20 times for each point and runtimes are averaged. Note the logarithmic scale for runtime.¹

Experimental Results

To explore the effects of relaxing Assumption 1, we simulate increasing amounts of overlap in conditional distributions, along the lines of Figure 2. We perform simulations using problem instances generated under two conditions, one where the instances are deterministically specified and one where they are randomly generated.

For the deterministically specified problem instance, we create a base set of conditional distributions, π_i , indexed by $i \in \{1, \dots, |V|\}$, implying $k = |V|$. First, we define π'_i , a linear function of $\omega \in \Omega = \{1, \dots, |\Omega|\}$ such that $\pi'_i(1) = 1\frac{1}{2} - \epsilon_i$ and $\pi'_i(|\Omega|) = 1\frac{1}{2} + \epsilon_i$, where each ϵ_i is the i th point on a straight line of evenly spaced points from $-\epsilon$ to ϵ . These k functions are simply linear functions with a range bounded by $[1, 2]$ where as i increases, the slope of the function increases. We then set $\pi(\omega) = \log(\pi'(\omega))$, to ensure that linear independence is satisfied (the choice of $\log(\bullet)$ is arbitrary; any non-linear transformation will do), and we normalize to create a valid probability distribution. The range of the linear functions, $[1, 2]$ is chosen to ensure that all values are positive after this non-linear transformation. We end up with conditional distributions such that for all $i \in \{1, \dots, |V| - 1\}$, π_{i+1} first order stochastically dominates π_i . We calculate the marginal probability of any particular buyer type, $\pi(\theta)$, by assigning a uniform probability equal to $1/|\Theta|$. Valuations are given by the set $\{1, 2, \dots, |V|\}$.

¹We use a core i7 3770 CPU with 8 GB of memory. The linear program in Definition 2 is solved using CPLEX with a .lp file specifying the full linear program. The algorithms corresponding to Definition 5 and calculating optimal revenue only (as in Theorem 5) are computed using Matlab R2013b. In Figure 4, CPLEX required all of system memory to specify the linear program after an overlap of 10 causing significant thrashing, so we discarded any further runtimes for the linear program.




Figure 4: Runtimes of the various algorithms when we hold fixed the number of buyer valuations ($|V| = 20$) and vary the degree of overlap from 0 to complete. Probabilities and valuations are randomly generated 20 times for each point and runtimes are averaged. Note that $|k| = |V| = 20$, and $|\Theta| = |\Omega|$ varies from k (no distributional overlap) to $k|V|$ (full distributional overlap, or $|V|$ types in each of k subsets). Note the logarithmic scale for runtime.¹

This specification of the conditional distributions gives us a natural interpretation of distributional “closeness”: specifically, two distributions are close if their value of ϵ_i is close. We increase overlap in a way analogous to Figure 2, where the degree of overlap would be two in the figure (two additional marginal distributions on each side of the original distribution).

To demonstrate that our results are not driven by the details of our deterministic specification, we also generate random instances. In these, both the marginal and conditional distributions’ probabilities are drawn randomly from the uniform distribution over $[0, 1]$ and then normalized. The elements of V are drawn uniformly, without replacement, from the set $\{1, 2, \dots, 500\}$.

Figures 3 and 4 demonstrate the improvements in runtime for both calculating the full mechanism using the algorithm corresponding to Definition 5 and just calculating the revenue only as in Theorem 5 relative to solving the linear program from Definition 2. Our improvement in runtime is especially dramatic for increasing distributional overlap: the runtimes for both revenue only and the full mechanism increase very slowly, while the linear program fails due to memory constraints for a relatively small problem size.

In Figure 5, we demonstrate the value of the optimal mechanism as we go from no distributional overlap (Assumption 1 holds), with a relative performance of 1, to full overlap (i.e., each X_i contains a type corresponding to every valuation). We also report the scaled value of an optimal pure reserve price mechanism on the full set of types for comparison (that is, not a separate reserve price for each X_i , but just a single reserve price overall). As can be seen in the figure, once a small degree of overlap is introduced, the revenue from the optimal auction rather quickly converges to that of a simple optimal reserve price mechanism.




Figure 5: Revenue comparisons. We have $|V| = 50$, $|\Omega| = 2500$, and $\epsilon = .4$ (for deterministically specified). Probabilities and valuations are generated as described in the text for both the deterministically specified and randomly generated instances (relative performance is effectively identical). We generate random data 20 times for each point and average the relative performance. On the vertical axis, a value of one denotes full surplus extraction, as in Cremer and McLean (1985), and .55 is the value of a pure reserve price mechanism (the average scaled value in the simulations).

Conclusion

In practice, there are many situations where a buyer’s valuation is correlated with an external signal. However, the accuracy required of the seller and buyer’s common prior for the Cremer and McLean (1985) mechanism is likely impractical. In addition to increasing our understanding of mechanism design with non-independent buyer valuations in general, our relaxation of the Cremer-McLean assumption allows us to investigate what happens in an environment with more limited accuracy. Figure 5 suggests that as the accuracy goes down, the gains relative to a simple reserve price auction, which does not use the correlation information at all, rapidly disappear. We view our result as further justifying the use in practice of such simple, robust mechanisms.

To be able to run these experiments, we introduced a new algorithmic approach to calculate the optimal mechanism for a single buyer with a valuation that is correlated with an external signal. This algorithm is correct if our Assumption 2 (which is a relaxation of the original Cremer-McLean assumption, Assumption 1) holds. This algorithm is significantly faster than the standard approach of solving the linear program directly, both under the Cremer-McLean assumption (as evidenced by Figure 3) and under Assumption 2 (Figure 4), and the decreased runtime allowed us to explore much larger problem instances than would otherwise have been possible. As such, our work illustrates how automated mechanism design, with the help of algorithmic tools developed for the specific mechanism design problem at hand, allows us to develop new high-level insight into the problem.

References

- Conitzer, V., and Sandholm, T. 2002. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 103–110.
- Conitzer, V., and Sandholm, T. 2004. Self-interested automated mechanism design and implications for optimal combinatorial auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 132–141.
- Cremer, J., and McLean, R. 1985. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* 53:345–361.
- Gneiting, T., and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477):359–378.
- Guo, M., and Conitzer, V. 2010. Computationally feasible automated mechanism design: General approach and case studies. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 1676–1679. NECTAR track.
- Hartline, J., and Roughgarden, T. 1998. Simple versus optimal mechanisms. In *Proc. 11th ACM Conf on Electronic Commerce*.
- Karmarkar, N. 1984. A new polynomial-time algorithm for linear programming. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing, STOC '84*, 302–311. New York, NY, USA: ACM.
- Krishna, V. 2009. *Auction Theory*. Academic Press, 2 edition.
- Likhodedov, A., and Sandholm, T. 2005. Approximating revenue-maximizing combinatorial auctions. In *AAAI*, volume 5, 267–274.
- Lopomo, G. 1998. The english auction is optimal among simple sequential auctions. *Journal of Economic Theory* 82:144–166.
- Myerson, R. 1991. *Game Theory: Analysis of Conflict*. Harvard University Press.
- Ronen, A., and Saberi, A. 2002. On the hardness of optimal auctions. In *Proc. of the 43rd Symposium on Foundations of Computer Science*, 396–405.