Approximately Optimal Mechanisms With Correlated Buyer Valuations

by

Michael Albert

Department of Computer Science
Duke University

Date: ______________________
Approved:

____
Vincent Conitzer, Supervisor

____
Kamesh Munagala

____
Pino Lopomo

Thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science in the Department of Computer Science
in the Graduate School of Duke University
2013

Abstract

Approximately Optimal Mechanisms With Correlated Buyer Valuations

by

Michael Albert

Department of Computer Science
Duke University

Date: ______________________

Approved:

____
Vincent Conitzer, Supervisor

____
Kamesh Munagala

____
Pino Lopomo

An abstract of a thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in the Department of Computer Science in the Graduate School of Duke University
2013

Abstract

Cremer and McLean 1985 shows that if buyers' valuations are sufficiently correlated, there is a mechanism that allows the seller to extract the full surplus from the buyers. However, in practice, we do not see the Cremer-McLean mechanism employed. In this thesis, I demonstrate that one reason that the Cremer-McLean mechanism is not implemented in practice is because the mechanism requires very precise assumptions about the underlying distributions of the buyers. I demonstrate that a small mis-estimation of the underlying distribution can have large and significant effects on the outcome of the mechanism. I further prove that the Cremer-McLean mechanism cannot be approximated by a simple second price auction, i.e. there is no approximating factor when using a second price auction with reserve in either outcome or expectation for the Cremer-McLean mechanism. Further, I show that there is no mechanism that approximates the Cremer-McLean mechanism for bidders with regular distributions in a single item auction if the correlation among buyers is not considered. Finally, I introduce a new mechanism that is robust to distribution mis-estimation and show empirically that it outperforms the Cremer-McLean mechanism on average in cases of distribution mis-estimation, and I show that the mechanism can be determined in polynomial time in the number of types of the buyers.

Contents

List of Figures

Acknowledgements

I am greatly indebted to my committee members, professors Dr. Kamesh Munagala and Dr. Pino Lopomo. Further, I am very grateful for the support of my friends and family members throughout my various academic trials and tribulations. Finally, I would like to thank my advisor and chair of my thesis committee, professor Dr. Vince Conitzer, who has been invaluable in the genesis and execution of this thesis.

1

Introduction

Recently, the efficiency of auctions both from a revenue as well as a computational standpoint has become an important question due to the quantity and frequency of auctions performed by large internet advertising companies such as Google and Yahoo. However, there seems to be a potential tradeoff between simple easily computable and implementable auctions and revenue efficient auctions, and this has led to an interest in approximating optimal mechanisms with more easily implementable mechanisms. One area in which this tradeoff seems to be particularly pronounced is the case of "sufficiently" correlated (a term to be made precise later) buyer valuations. Cremer and McLean (1985) shows that when buyer valuations are sufficiently correlated, the seller can extract the full surplus from two or more buyers using a mechanism that I will hereafter refer to as the "Cremer-McLean mechanism". This is a very strong result in mechanism design, and it's more powerful for not relying on the seller to have additional information on the private valuations of the buyers. Instead, it leverages the additional private information that each buyer has about the other to extract full surplus.

The standard assumption in the mechanism design literature is that buyer's valuations are independent. However, there are many situations in which this is likely to be false. One common example of correlated valuations is auctions for oil drilling rights. Before the auction, each bidder conducts a geological survey to estimate the value of the plot for which they are bidding, and since each survey is an estimate of the true potential profit the plot may generate, each bidder knows something about all of the other bidders valuations that the seller doesn't know. Further, the values will also be correlated if there is a potential resale market for a good. If the bidder knows that he can resell the item, the valuations of the other bidders become part of his valuation of the item, creating correlation.

While the Cremer-McLean mechanism gives the maximal revenue, in general we don't see the Cremer-McLean mechanism implemented, and instead we see simple auctions that compensate the buyers for revealing their private information. This might seem to be a case of sellers leaving money on the table. However, there are significant difficulties with implementing the Cremer-McLean mechanism, namely the necessity for a very precise estimation of the probability distribution over types. The mechanism gives all buyers in the market their reservation utility in expectation, and if the distribution over types is mis-estimated, some buyers, ones who know their true types, will exit the market, reducing the expected revenue of the seller. Further, the buyers that remain in the market are those that are receiving greater than or equal to their reservation utility, leading to a further reduction in the expected revenue for the seller.

In practice, the standard mechanism used is typically a first or second price auction with a reserve. In this paper, we will only consider the second price auction with reserve. It is not difficult to see why a second price mechanism is implemented in practice. First, the second price auction mechanism in all situations has a simple dominant strategy, each buyer reports his true valuation. Second, a standard second price auction without reserve is completely prior independent. Third, in the case of independent and identically distributed buyers, the second price auction with reserve can be shown to be revenue optimal. Finally, in many other situations where buyers are not identically distributed it has been shown that the second price auction with reserve approximates the revenue optimal mechanism.

The first question I explore in this thesis is whether a second price auction with reserve can approximate the Cremer-McLean mechanism. This simple auction allows one to sidestep the main difficulty with the Cremer-McLean mechanism, namely the only element that is prior dependent is the reserve price, and it is not overly sensitive to the prior in most instances. The work of Hartline and others have demonstrated that second price auctions perform approximately optimally in a wide variety of settings, and the question that I am ultimately trying to address with this line of inquiry is, "Is the correlation of buyer values fundamental to the design of the mechanism, or is it simply an optimization?" One reason to expect that it may simply be an optimization is that for the case of independent, but non-identically distributed, buyer valuations, a second price auction does perform approximately optimal, i.e. generates revenue within a constant factor of the optimal mechanism. Note that a second price auction can set the reserve dependent on the correlation between buyer types, but the dependence of the mechanism on the correlation is much weaker than the Cremer-McLean mechanism. Therefore, the ability to approximate the Cremer-McLean mechanism by a second price auction would be a statement about the importance of value correlation to mechanism design.

However, I am able to show that there does not exist a performance guarantee for the approximation of the optimal mechanism by the second price auction with reserve, either in outcome or expectation, whether or not we allow the reserve price to be a function of other buyer's valuations. This result is a strong negative result, and it implies that buyer correlation is a fundamental aspect of the mechanism design problem. In spite of the negative results, it is likely that there are situations under which the second price auction with reserve can still be expected to do quite well compared to the optimal mechanism, and my negative results provides guidance for potential additional assumptions that may restrict the prior space to a set under which the second price auction with reserve does approximate the optimal mechanism.

After I demonstrate that a second price auction with reserve cannot approximate the optimal mechanism, I introduce a modification to the Cremer-McLean mechanism which allows for slight mis-estimations of the distribution of types. The principle difficulty with the Cremer-McLean mechanism is that all potential buyers are reduced to their reservation utility in expectation, so if the distribution is incorrectly estimated, it is quite likely that one or more of the buyers will decide not to participate in the mechanism. If that is the case, the expected revenue can change quite severely, especially in the case where there are only two buyers. Once one buyer decides to not participate, the mechanism becomes a take it or leave it mechanism, and the seller can no longer extract full surplus even with a correctly estimated distribution. My robust Cremer-McLean mechanism ensures that no matter the mis-estimation, all buyers will choose to participate in the mechanism.

After developing the robust Cremer-McLean mechanism, I then test it against the standard Cremer-McLean mechanism with varying degrees of distribution mis-estimation. I am able to show that on average, when there are only two buyers, the robust Cremer-McLean mechanism performs significantly better than the standard Cremer-McLean mechanism, especially for small deviations from the true distribution.

The literature that this fits most naturally with is the growing literature on approximating optimal mechanisms with simple auctions pioneered by various people including Jason Hartline and Tim Roughgarden (see Hartline and Roughgarden (1998) and Hartline (2012), among others), in addition to many others. This paper naturally owes much to Myerson (1981), the pioneering paper on optimal auction design. Further, this is also related to Lopomo (1998), which looks at optimal mechanisms where the loser does not pay. This is also related to the field of algorithmic mechanism design, see Nisan and Ronen (2001) among others.

One paper that is particularly close to this thesis is Ronen and Saberi (2002). In Ronen and Saberi (2002), they look at the hardness of optimal auctions, and they are able to demonstrate that no deterministic polynomial time ascending auction can achieve an approximation ratio better then $\frac{3}{4}$. The distribution of types in the example proving the upper bound for the approximation ratio contains dependencies between buyers, i.e. the buyers types are correlated. However, they find that for independent agents, the approximating factor can get quite close to 1. This work differs from my thesis in that I am not as concerned with computational complexity as I am with robustness to mis-estimation of the initial distribution.

Further work has been done on the nature of uncertainty over a buyer's distribution. In particular, the notion of Knightian uncertainty (see Knight (1921)) has been explored in the literature. Lopomo et al. (2009) examines the role of Knightian uncertainty in mechanism design, and they show that full extraction of the surplus is possible in uncertain conditions given restrictions over type space. My thesis looks at a simpler notion of uncertainty in belief. Specifically, I assume that the distribution of types is mis-estimated, but the seller is confident in his mis-estimation.

The thesis proceeds as follows: In chapter 2, I discuss the Cremer-McLean mechanism and the assumption of sufficient correlation in detail. In chapter 3, I briefly discuss the second price auction, and then I show through counterexamples that there is no performance guarantee. In chapter 4, I introduce my robust Cremer-McLean mechanism and I test it's performance empirically. Finally, in chapter 5, I conclude.

2

The Cremer-McLean Mechanism

2.1 Definition and Optimality of the Cremer-McLean Mechanism

The Cremer-McLean mechanism is a solution to a linear program that has been shown to solve the general maximal revenue problem. Cremer and McLean (1985) demonstrates the optimality of the mechanism for the general case where a single monopolistic seller is selling a quantity x of a good to some number of buyers n. In this paper, I will restrict my attention to the case in which the seller has a single indivisible good and there are n sellers.

I will denote the set of agents participating in the mechanism by $N = \{1, .., n\}$ where $n \geqslant 2$, and I will alternatingly call them possible buyers, buyers, and bidders. Each bidder has a risk-neutral valuation for the object from a discrete set of types denoted by the set $M_i = \{v_{(i,1)}, ..., v_{(i,m_i)}\}$ where $v_{(i,j-1)} \leqslant v_{(i,j)}$ for $j = 2..n$. This valuation is private to the agent, and I will denote the buyer's type in general by s_i, i.e. a buyer i has type s_i when his valuation is $v_{(i,s_i)}$.

However, the buyer's valuation for the object are not independent. Let an outcome be denoted by $s = (s_1, ..., s_n)$, the probability of any particular outcome is

$\pi(s)$, and the subjective probability for buyer i is denoted π^i. A crucial assumption of this mechanism is that there is a common prior, i.e. $\pi^i(s_{-i}|s_i) = \frac{\pi(s_i, s_{-i})}{\pi^i(s_i)}$, where $s_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$.

The seller sells the item to bidder i with probability $p_i(\hat{s})$ and charges $x_i(\hat{s})$ for every reported outcome $\hat{s}$. Notice that this mechanism permits there to be a payment from the bidder to the seller in spite of whether or not the bidder wins the item. The seller has a zero valuation for the item. This gives each bidder a utility $u_i(s_i, \hat{s}) = v_i(s_i)p_i(\hat{s}) - x_i(\hat{s})$.

For the case of full information, the seller maximizes revenue by solving the problem:

Definition 1. *The full information mechanism is given by the solution to the following program*

$$
\begin{aligned}
\max_{x_i(s), p_i(s)} & \sum_i x_i(s) \\
& s.t. \\
& v_i p_i(s) - x_i(s) \geqslant 0 \ \forall i \in N \\
& \sum_i p_i = 1 \\
& p_i >= 0 \ \forall i \in N.
\end{aligned}
\tag{2.1}
$$

I will call this problem $R(s)$. Notice that it extracts the full surplus, since it can ensure that for every bidder $u_i(s) = 0$ and the item is always sold in equilibrium. However, this program requires that the seller has full information about every buyer's type. The standard assumption is that the buyers have private information about their own valuations, so the seller cannot solve this program directly.

The program that the seller can solve in order to generate a mechanism for which to sell the item is the program E which yields a solution for the best ex-post Nash Equilibrium from the point of view of the seller:

Definition 2. *The Cremer-McLean mechanism is given by the solution to the following program*

$$\max_{x_i(s),p_i(s)} \sum_s \pi(s)\left[\sum_i x_i(s)\right]$$

$$s.t.$$

$$\sum_{s_{-i}} \pi^i(s_{-i}|s_i)[v_i p_i(s) - x_i(s)] \geqslant 0 \ \forall i \in N, s_i \in M_i$$

$$v_i p_i(s) - x_i(s) \geqslant v_i p_i((t_i, s_{-i})) - x_i((t_i, s_{-i})) \ \forall i \in N, s_i \in M_i, s_{-i} \in M_{-i}, t_i \in M_i$$

$$\sum_i p_i = 1$$

$$p_i >= 0 \ \forall i \in N \tag{2.2}$$

This program maximizes the expected revenue of the seller, not the realized revenue like the program $R(s)$, which maximizes the realized revenue. Further, it does not assume any information about the buyer's private valuations. The solution to program E will always give an optimal ex-post Nash Equilibrium mechanism. However, in general, it will not generate revenue equivalent to the full information program; in general it will be sub-optimal compared to the full information program.

Cremer and McLean 1985 impose the additional assumption:

Assumption 1. *For all $i \in N$, let $\mathbf{\Gamma_i}$ be the following matrix whose rows are indexed by the $k_i = \Pi_{j \in N_{-i}} m_j$ elements of M_{-i}, and whose columns are indexed by the m_i elements of M_i:*

$$\mathbf{\Gamma_i} = \begin{bmatrix} \pi^i(s_{-i} = 1|s_i = 1) & \dots & \pi^i(s_{-i} = 1|s_i = m_i) \\ \vdots & \dots & \vdots \\ \pi^i(s_{-i} = k_i|s_i = 1) & \dots & \pi^i(s_{-i} = k_i|s_i = m_i) \end{bmatrix}$$

For all i*,* $\mathbf{\Gamma_i}$ *has rank* m_i*.*

Note that the above assumption includes the following assumption.

Assumption 2. *For every buyer* i*,* $|M_i| \leqslant |M_{-i}|$*.*

Assumption 1 doesn't require that the correlations have any particular structure or magnitude outside of full rank, and the full rank assumption is due to the necessity for the distribution of probabilities for all of the other bidders to be distinguishable for every draw from bidder i's distribution. If this assumption holds, they are able to show that the programs 2.1 and 2.2 have identical objective values in expectation.

Theorem 1. *Under assumption 1,* $V(E) = \sum_{s \in M} \pi(s) V(R(s))$*.*

Proof. See Cremer and McLean (1985) or Krishna (2009) Chapter 10 for a proof of theorem 1. □

They demonstrate that in expectation, the solution to program E generates the same revenue as the full information program R. This result comes ultimately from the fact that the outcomes do not necessarily satisfy the participation constraint for all of the agents. The participation constraint is only satisfied in expectation, and since all agents are risk neutral this is sufficient. This allows the seller to shift payments in such a way that in expectation his participation constraint is satisfied, but on the valuations where he would be tempted to deviate by claiming to be a lower type, the payoff to the lower type has been shifted downwards by the payments required. This only works in the case of correlated values, because the bidder has additional information about the other bidders types, and so the seller can set the payoffs to end up negative for any situation under which the bidder chooses to lie.

2.2 Form of the Cremer-McLean Mechanism

It's not immediately obvious what form the Cremer-McLean mechanism actually takes from inspection of program 2.2. However, the mechanism can be described as the combination of two standard mechanisms, a standard Vickrey-Clarke-Groves (VCG) mechanism and a lottery or, equivalently since valuations are private, a second price auction and a lottery.

To understand the mechanism, we can consider program 2.2. The VCG mechanism for this situation is always allocate the item to the buyer with the highest valuation, and require the winning buyer to pay the value the second highest bidder would have paid for the item. Ties can be broken arbitrarily. It is easily shown that this mechanism always generates truth telling. Let the VCG mechanism be denoted by $p_i^*(s)$ and $x_i^*(s)$. Then the incentive compatibility constraint is always satisfied for program 2.2 under $p_i^*(s)$ and $x_i^*(s)$.

The expected utility of each buyer under the VCG mechanism is then given by

$$U_i^*(s_i) \equiv \sum_{s_{-i}} \pi^i(s_{-i}|s_i)[v_i p_i^*(s) - x_i^*(s)]. \tag{2.3}$$

Let $\mathbf{u_i^*}$ be the column vector of length m_i giving the values of $U_i^*(s_i)$ for all s_i. Since the matrix $\mathbf{\Gamma_i}$ is of full rank, there exists a column vector $\mathbf{c_i}$ of length k_i such that $\mathbf{\Gamma_i c_i} = \mathbf{u_i^*}$.

Now consider the new mechanism denoted CM given my $p_i^{CM}(s) \equiv p_i^*(s)$ and

$$x_i^{CM}(s) \equiv x_i^*(s) + \mathbf{c_i}(s_{-i}). \tag{2.4}$$

Then the incentive compatibility constraint is still satisfied since the modification to the mechanism does not depend on the buyer's reported type. Further, the item is always allocated in equilibrium and and the expected value of all buyers is zero for all types, i.e. $U_i^*(s_i) = 0$ for all i and s_i. Therefore, this mechanism satisfies all of the constraints of program 2.2 and always allocates the item, so it is optimal. This is the Cremer-McLean mechanism.

Since all buyers have private valuations, though correlated, it can be shown that the VCG mechanism is a second price auction. The additional component of the payment function, $\mathbf{c_i}(s_{-i})$ does not depend on buyer i's reported type. Therefore, buyer i can view this as a lottery over the distribution of others types. So, the Cremer-McLean mechanism is, at its core, a combination of a second price auction with a lottery to enter the auction.

3

The Inapproximability of the Cremer-McLean Mechanism by a Second Price Auction with Reserve

The second price auction with reserve has been shown to be approximately optimal in a strict sense for many classes of buyer beliefs. Further, as shown in chapter 2, the Cremer-McLean mechanism is a combination of a second price auction and a lottery. Given this, one might expect that a second price auction with reserve might be approximately optimal for the Cremer-McLean mechanism as well. However, it can be shown that there is no approximation bound for the second price auction with reserve with respect to the Cremer-McLean mechanism.

3.1 The Second Price Auction with Reserve

A second price auction with reserve is a simple auction with a dominant strategy equilibrium for any distribution of types over the bidders. The mechanism is as follows: Every bidder reports a valuation. If no valuation is above the reserve price, then the item is not sold. If one valuation is above the reserve price, charge that bidder the reserve price and allocate the item to that bidder. If two or more valuations are above the reserve price, charge the bidder with the highest valuation the second highest valuation and allocate the item to that bidder. It is trivial to show that each bidder has a dominant strategy in which they report their true valuations, assuming that there is no collusion among bidders.

For a continuous case with i.i.d. agents, this auction is revenue optimal if the reserve is set to $\phi^{-1}(0)$, where $\phi(\bullet)$ is the virtual value function. If agent valuations are independent, but non-identically distributed, then if the seller sets a vector of reserve prices $\mathbf{r} = (\phi_1^{-1}(0), ..., \phi_n^{-1}(0))$, the second price auction with a reserve is a two approximation to the optimal auction.

The virtual value function is simple to calculate for the case of regular continuous distributions; it is $\phi(v) = v - \frac{1-F(v)}{f(v)}$, where F is the cdf of v. In the case of discrete distributions one can calculate the reserve price directly by finding the monopoly price, which is the value $v_{(i,k)}$ that maximizes $v_{(i,k)}(1 - F(v_{i,k-1}))$, i.e. the expected revenue from selling to a single agent, which can be calculated in time linear in the number of types. Similarly, if values are correlated, one can set the reserve price to the $\text{argmax}_{v_{(i,k)}} v_{(i,k)}(1 - F(v_{i,k-1}|s_{-i}))$, where the seller knows s_{-i} from the reported types of the other agents, i.e. the realized revenue for the second price auction with reserve is at most half of the realized revenue from the optimal auction, see Hartline (2012).

Note that this mechanism is not prior independent; it relies on the distribution for the calculation of the optimal reserve price. So, to some extent, it is susceptible to the same criticism as the Cremer-McLean mechanism. However, the reserve price is much less dependent on the prior, and as you increase the number of bidders, in many cases, the probability of selling at the reserve price decreases further reducing the sensitivity of the outcome to the prior distribution assumption.

3.2 The Inapproximability of the Cremer-McLean Mechanism by a Second Price Auction with Reserve

While the second price auction is approximately optimal for many distributional assumptions over buyers' valuations, I demonstrate through counterexample that there can be no strict approximation bound for the Cremer-McLean mechanism by the second price auction with reserve. However, this negative result gives directions for potential ways over which to restrict the prior distribution space in order to achieve a strict approximation bound.

Theorem 2. *There is no performance guarantee for the approximation of the expected revenue generated by the Cremer-McLean mechanism by the second price auction with reserve.*

Proof. Let there be two bidders with identical valuation sets. Consider the following distribution of values for bidder 1 where the distribution is denoted as 2-tuples with the first entry indicating the marginal probability for him to draw that valuation, $\{(1/2, 2), (1/4, 4), (1/8, 8), ..., (1 - \sum_{i=1}^{k|2^k<m} 1/2^i, m)\}$. Let bidder 2 have the marginal probability distribution given by $\{(\rho_2, 2), (1 - \sum_{i=2}^{k+1|2^k<m} \rho_i, 4), (\rho_3, 8), ..., (\rho_{k+1}, m)\}$, where $\rho_i \geqslant 0$ can be set arbitrarily small, though some must have a nonzero value in order to satisfy assumption 1. Further, let the correlations between the two agents be such that they satisfy Assumption 1. Additionally, assume that

$$\rho_2 + \sum_{i=3}^{k+1|2^k<m} \rho_i 2^i < \sum_{i=2}^{k+1|2^k<m} \rho_i 4, \tag{3.1}$$

The expected type of bidder 2 is given by

$$E(v_2) = \rho_2 + (1 - \sum_{i=2}^{k+1|2^k<m} \rho_i)4 + \sum_{i=3}^{k+1|2^k<m} \rho_i 2^i. \tag{3.2}$$

which by the previous assumption, implies $E(v_2) < 4$.

Note that the Cremer-McLean mechanism has expected revenue bounded below by the expected value of bidder 1, since the expected value of the Cremer-McLean mechanism is just the expected highest value of an agent. Therefore, the Cremer-McLean mechanism has expected revenue $CM \geqslant \sum_{i=1}^{k|2^k \leqslant m} 1/2^i * 2^i + (1 - \sum_{i=1}^{k|2^k<m} 1/2^i)m$. However, as m increases, this diverges towards infinity, implying that the Cremer-McLean mechanism has infinite value.

The second price auction has an upper bound given by the second price auction over the two bidders given by the distribution of values such that bidder $1'$ is identical to bidder 1 and bidder $2'$ is given by $\{(0,2),(1,4),(0,8),...,(0,m)\}$, i.e. the bidder who always values the item at 4. This must give an upper bound for the value of the second price auction over bidders 1 and 2 since the $E(v_{2'}) = 4 \geqslant E(v_2)$. The expected value of the auction over bidders $1'$ and $2'$ can be derived in the standard way given that the bidders are now independent. The optimal reserve price is set to $\text{argmax}_{v_{(i,k)}} v_{(i,k)}(1 - F(v_{i,k-1}))$ for each of the bidders. For bidder $2'$, this value is 4. Note that for any valuation, save for m, the value of $v_{(i,k)}(1-F(v_{i,k-1})) = 2$ for bidder $1'$. Since $m < 2^{k+1}$, where k is defined as above, $m(1 - \sum_{i=1}^{k|2^k<m} 1/2^i) = m/2^k < 2$. So, we can set the reserve at any price save for m for bidder $1'$, to maximize the revenue due a monopolist. I will set the reserve price to 4. This implies that the expected revenue for the second price auction over bidders $1'$ and $2'$ is 4; the item is always sold to one of the bidders at a price of 4. Therefore the value of the second price auction over bidders 1 and 2 is bounded above by 4.

So, the expected value of the Cremer-McLean mechanism diverges to infinity as m increases, and the expected value of the second price auction is always 4. Therefore, there can be no approximation bound for the Cremer-McLean mechanism by the second price auction with reserve. □

The above result leads immediately to the following corollary:

Corollary 1. *There is no performance guarantee for the approximation of the expected revenue generated by the Cremer-McLean mechanism by the second price auction without a reserve.*

The above two results are strongly negative. They indicate that a second price auction with reserve is missing something fundamental about the structure of the problem, and since others have shown that the second price auction approximates these distributions if they are independent, the correlation must be the essential component that is missing.

This result does not imply that second price auctions with reserve always perform poorly in comparison to the optimal auction. It is still likely that a second price auction with reserve is approximately optimal in many situations. However, in the worst case, the second price auction with reserve can do arbitrarily poorly compared to the Cremer-McLean mechanism.

This result does give direction in which to look for distributions that may have provable approximation bounds. The main intuition that this result provides is that the distribution cannot be arbitrarily right skewed. If the family of distributions can be arbitrarily right skewed, than it is likely that a similar argument could be applied. In the future, we must look for distributions that are in some sense "symmetric". This still leaves a large number of potential families of distributions, including correlated normal distributions. Note that the distributions don't have to be symmetric; though it's unproven, it is likely that a lognormal distribution would satisfy the conditions necessary for an approximation bound.

This result can be extended to say something slightly more general about the approximability of the Cremer-McLean mechanism, as given by the following corollary:

Corollary 2. *There is no mechanism that discards the correlation structure of the* *bidders that has an approximation bound for the Cremer-McLean mechanism when bidders have a regular distribution for a single item auction.*

Proof. Since the counter example uses regular distributions, and the second price auction with a reserve has been shown by Hartline and Roughgarden 2009 to be a two approximation for any single item auction with bidders with independent regular distributions, the result follows immediately. □

4

Robust Cremer-McLean

Since it has been demonstrated that a second price auction cannot approximate the Cremer-McLean mechanism, it is useful to consider how one should modify the mechanism in order to increase its robustness and usefulness. However it is worthwhile to examine the specific areas under which the Cremer-McLean mechanism fails to be robust.

4.1 Failure of the Cremer-McLean mechanism

The Cremer-McLean mechanism is a generalized VCG mechanism where the seller is able to drive all the expectation of all buyers that participate in the mechanism to zero through a lottery to enter as discussed in section 2.2. This allows the seller to extract the full surplus. However, it makes the mechanism very sensitive to small mis-estimations of the initial distribution.

To understand why the mechanism is so sensitive to the correct estimation of the initial distribution, it is worthwhile to consider what may happen when the initial distribution is incorrectly estimated. Notice from the program given by equation 2.2 that the distribution enters in in two different places: the objective function and the individual rationality constraints. If the distribution is only slightly incorrectly estimated, a notion that will be defined more fully later, by the seller, then the objective function only changes slightly and in a continuous fashion, leaving little room for a significant change in the outcome due to the incorrectly estimated objective function. However, since the outcome of the Cremer-McLean mechanism is to drive the expected utility for all of the buyers to zero, a slight mis-estimation of the distribution could result in one or more of the buyers no longer considering it optimal to participate in the mechanism, i.e. their expected utility from participating in the mechanism is lower than the outside of option of not participating.

The situation under which a buyer decides not to participate is not well defined, so for the purpose of this thesis, I will make some very specific assumptions about the information structure and the decision process on whether or not to participate. First, I assume that, while the seller incorrectly estimates the distribution, the buyers always know the true distribution and they observe the seller's incorrectly estimated distribution. Since the buyers know the true distribution, and they can observe the seller's estimation, they know what the mechanism will be before participating in the entry lottery. I further assume that the buyers observe their own types (though not other buyers' types) before they make the choice as to whether or not to participate. This will imply that the buyer should only participate when it is individually rational for the buyer to do so.

Since the buyer may choose not to participate based on their own specific draw from the type distribution and the projected mechanism that they will face, one must specify what happens in the case in which the buyer does not participate. For the purpose of this thesis, I assume that the seller still implements the Cremer-McLean mechanism with the remaining participants assuming that there are two or more participants. Note that assumption 1 will hold with probability one for a random distribution assuming that for all buyers still remaining as long as for the set of buyers remaining assumption 2 holds. I assume that this will be the case, and if one requires that all buyers have the same number of states, this is guaranteed to hold. If there is only one participant remaining, then the seller makes a take it or leave it offer that maximizes his expected revenue from a take it or leave it offer.

However, I assume that buyers are somewhat naive in my model. While they know their own distributions and the seller's distribution, they make decisions locally with respect to the mechanism. To be more precise, once the buyer decides not to participate in the mechanism, she cannot decide to re-enter and she does not make her decision based on the probability of others not participating in the mechanism, i.e. she assumes that all other buyers decide to participate.

To reiterate, the buyers see the mis-estimated distribution of the seller and know their own types. Then, they decide whether or not to participate in the mechanism. If some of the buyers decide not to participate, then the seller will implement a mechanism with the remaining buyers. Here, the remaining buyers have an additional choice as to whether or not to participate in the mechanism. This continues until there is a stage in which all buyers remaining decide to participate in the mechanism. At this point, the mechanism is implemented, either Cremer-McLean or a take it or leave it offer with the mis-estimated distribution. I further assume that the seller doesn't update his beliefs about the distribution based on whether or not a buyer chooses to participate.

Here it becomes obvious where the major loss in revenue due to mis-estimating the Cremer-McLean mechanism comes from. The mis-estimation is likely to lead to a decline in the number of participants in the market due to buyers exiting because of an expected value lower than the outside option. There is still a reduction in the revenue by maximizing a mis-specified objective function, but the main reduction in revenue is a consequence of not being able to extract the surplus from that buyer. Further, if there is only one buyer remaining, the mechanism defaults to take it or leave it, which does not extract full surplus, even for a correctly specified distribution.

The assumptions made here are not without loss of generality. They are potentially the most negative assumptions that can be made for buyer exit. If buyers considered the probability of others exiting, this can only increase the probability that they will stay in the mechanism, increasing the expected revenue for the seller, given that they always make a worst case assumption for the exit decision. If sellers update their belief about buyer types given exit, they can only increase their expected revenue. In my simulations in the next section, I generate a lower bound for the expected revenue for the seller, and with different assumptions, the difference between the mis-specified Cremer-McLean mechanism and the robust Cremer-McLean mechanism may not be as pronounced.

4.2 A Specification for a Robust Cremer-McLean Mechanism

Since the main loss in revenue for the Cremer-McLean mechanism is likely the loss due to lack of participation instead of a mis-specified objective function, a robust Cremer-McLean mechanism should ensure that all buyers always participate in the mechanism. This can be accomplished by setting the incentive compatibility constraint such that in the worst case, the expected utility of the buyer is greater than his outside option. However, this requires specific assumptions about the mis-estimation of the distribution.

Assumption 3. *For all states s, the estimated distribution $\pi^*(s) = \pi(s) + \epsilon(s)$ where $\epsilon(s) \in [-\bar{\epsilon}, \bar{\epsilon}]$, i.e. the mis-estimation is strictly bound.*

However, while this bounds the distribution $\pi(s)$, it does not bound the conditional distribution, $\pi(s_{-i}|s_i)$, directly, which is the distribution of interest in the individual rationality constraint. A bound can be calculated as follows.

Theorem 3. *If Assumption 3 holds, then $\pi(s_{-i}|s_i) \in [\underline{\pi}^*(s_{-i}|s_i), \overline{\pi}^*(s_{-i}|s_i)]$ where*

$$\underline{\pi}^*(s_{-i}|s_i) = \max\left(\frac{\pi^*(s_{-i}, s_i) - \bar{\epsilon}}{\sum_{s_{-i'}} \pi^*(s_{-i'}, s_i) + |M_{-i}|\bar{\epsilon}}, 0\right) \tag{4.1}$$

and

$$\overline{\pi}^*(s_{-i}|s_i) = \min\left(\frac{\pi^*(s_{-i}, s_i) + \bar{\epsilon}}{\sum_{s_{-i'}} \pi^*(s_{-i'}, s_i) - |M_{-i}|\bar{\epsilon}}, 1\right). \tag{4.2}$$

Proof. I will show that $\pi(s_{-i}|s_i) \geqslant \underline{\pi}^*(s_{-i}|s_i)$, and $\pi(s_{-i}|s_i) \leqslant \overline{\pi}^*(s_{-i}|s_i)$ can be shown by nearly identical arguments. By the definition of marginal probability,

$$\pi(s_{-i}|s_i) = \frac{\pi(s_{-i}, s_i)}{\sum_{s_{-i'}} \pi(s_{-i'}, s_i)} \tag{4.3}$$

which by the definition of the estimated distribution

$$\pi(s_{-i}|s_i) = \frac{\pi^*(s_{-i}, s_i) - \epsilon(s_{-i}, s_i)}{\sum_{s_{-i'}} (\pi^*(s_{-i'}, s_i) - \epsilon(s_{-i'}, s_i))}. \tag{4.4}$$

However, by assumption 3, $\epsilon(s) \in [-\bar{\epsilon}, \bar{\epsilon}]$, so

$$\pi(s_{-i}|s_i) \geqslant \frac{\pi^*(s_{-i}, s_i) - \bar{\epsilon}}{\sum_{s_{-i'}} (\pi^*(s_{-i'}, s_i) + \bar{\epsilon})} = \frac{\pi^*(s_{-i}, s_i) - \bar{\epsilon}}{\sum_{s_{-i'}} \pi^*(s_{-i'}, s_i) + |M_{-i}|\bar{\epsilon}} \tag{4.5}$$

However, since $\pi(s_{-i}|s_i)$ is a valid probability, $\pi(s_{-i}|s_i) \geqslant 0$, therefore, $\pi(s_{-i}|s_i) \geqslant \underline{\pi}^*(s_{-i}|s_i)$.

Similarly, one can show that $\pi(s_{-i}|s_i) \leqslant \overline{\pi}^*(s_{-i}|s_i)$. □

From theorem 3, a mechanism can be developed that will assure that all types will still participate in the mechanism. This avoids what is likely to be the largest loss of revenue from a mis-estimated distribution. The following mechanism is the robust Cremer-McLean mechanism.

Definition 3. *The robust Cremer-McLean mechanism is given by the solution to the following program*

$$\max_{x_i(s),p_i(s)} \sum_s \pi^*(s) \left[\sum_i x_i(s)\right]$$

$$s.t.$$

$$\sum_{s_{-i}} (P(s_{-i},p)\underline{\pi}^*(s_{-i}|s_i) + (1 - P(s_{-i},p))\overline{\pi}^*(s_{-i}|s_i))[v_i p_i(s_i,s_{-i}) - x_i(s_i,s_{-i})] \geqslant 0$$

$$\forall i \in N, s_i \in M_i, p \in K_i$$

$$v_i p_i(s_i,s_{-i}) - x_i(s_i,s_{-i}) \geqslant v_i p_i((t_i,s_{-i})) - x_i((t_i,s_{-i}))$$

$$\forall i \in N, s_i \in M_i, s_{-i} \in M_{-i}, t_i \in M_i$$

$$\sum_i p_i = 1$$

$$p_i >= 0 \ \forall i \in N \tag{4.6}$$

where P is a matrix of all possible binary strings of length $|M_{-i}|$ where the strings are indexed by $p \in K_i$, and if T is the number of possible binary strings of length $|M_{-i}|$, K_i is the set $\{1, 2, ..., T\}$.

Notice that the mechanism is only slightly, though significantly, modified from the original Cremer-McLean program. Since the incentive compatibility constraint does not include any reference to the prior distribution, it is not affected by a mis-estimation of the prior, and each buyer always finds it optimal to report his true type. Further, the objective function is identical, save that the objective is misspecified due to the mis-estimated distribution. However, the individual rationality constraint now is guaranteed to include the worst case situation for each buyer, due to including, for all types, the constraint that the worst possible outcome for that buyer still guarantees a utility equal to the reservation utility.

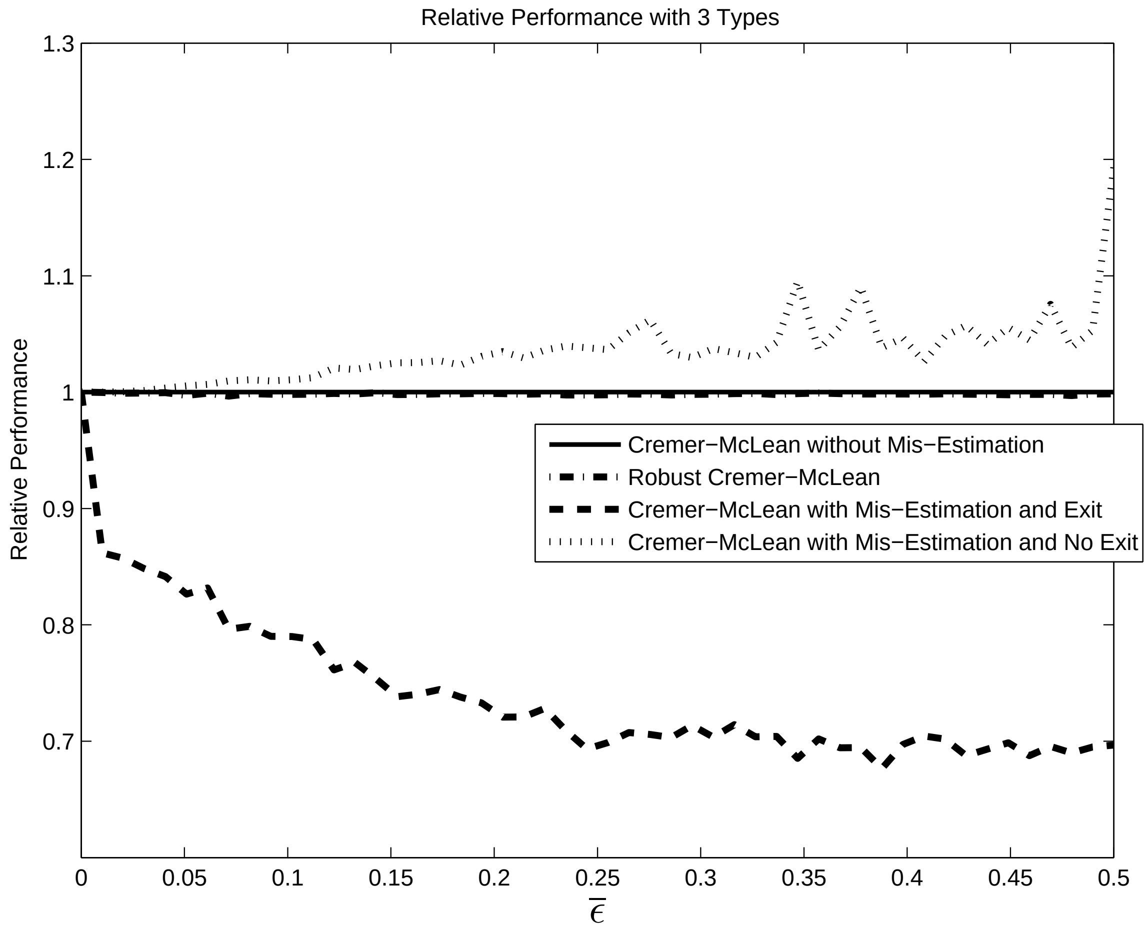

FIGURE 4.1: The relative performance of the robust Cremer-McLean mechanism compared to the optimal mechanism and the two mis-specification assumptions for the optimal mechanism when buyers have three types.

4.3 Empirical Examination of the Robust Cremet-McLean Mechanism

The robust Cremer-McLean mechanism appears impractical to analyze theoretically, so the approach taken in this thesis is to examine the results empirically. The empirical approach consists of generating random distributions over the set of all possible discrete distributions for two buyers of varying numbers of types. Both buyers have identical type spaces, where their valuations range from 0 to a value equal to the number of types minus one, with their type space consisting of the

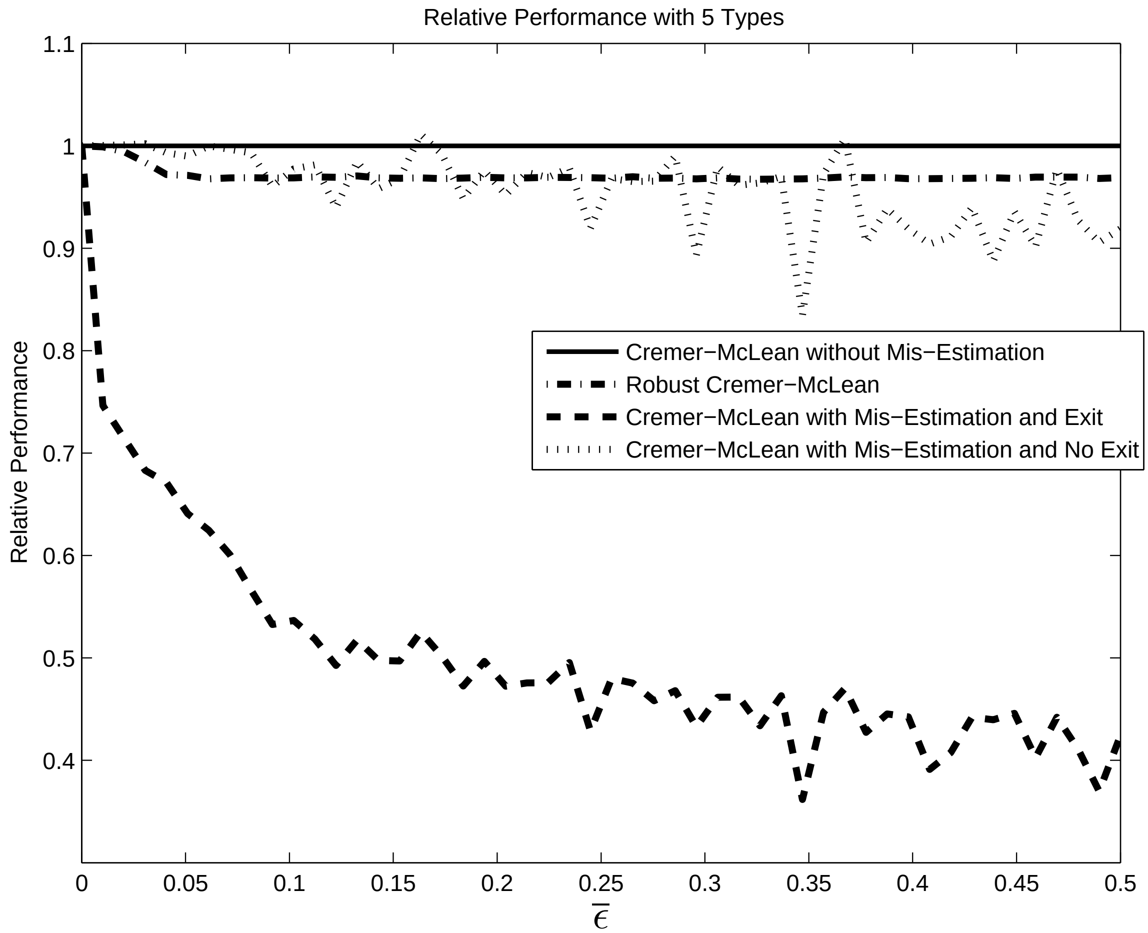

FIGURE 4.2: The relative performance of the robust Cremer-McLean mechanism compared to the optimal mechanism and the two mis-specification assumptions for the optimal mechanism when buyers have five types.

integer values between and including the endpoints. Using asymmetric type spaces does not materially change the results for a variety of examined type spaces. The randomly drawn distribution is then perturbed as in assumption 3, for 50 equally spaced values of $\overline{\epsilon}$ in the interval $[0, .5]$. This is done 2000 times for each value of $\overline{\epsilon}$, and the runs are then averaged. The value is then scaled by the value obtainable by the standard Cremer-McLean mechanism assuming that there is no mis-estimation in the distribution.

The case where there are only two buyers is the only case considered. While

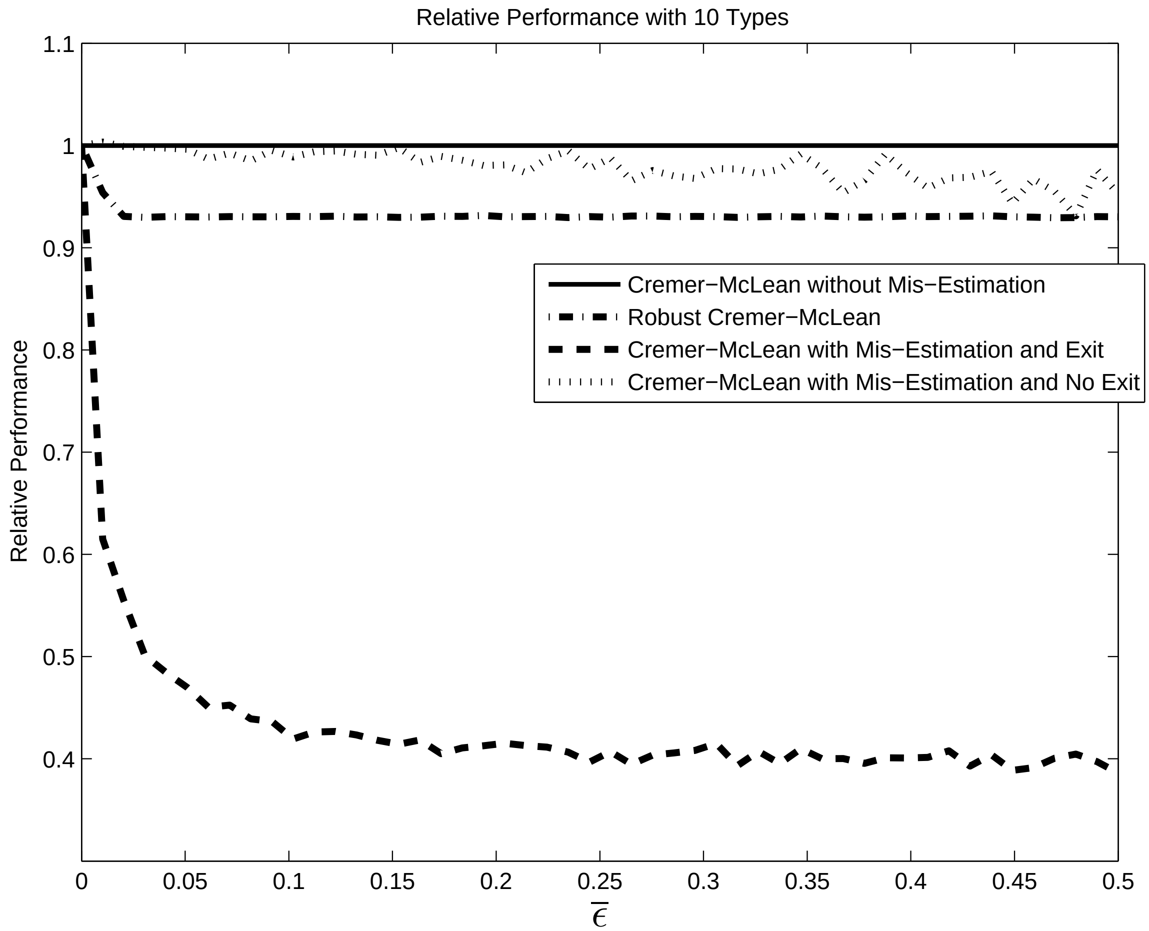

FIGURE 4.3: The relative performance of the robust Cremer-McLean mechanism compared to the optimal mechanism and the two mis-specification assumptions for the optimal mechanism when buyers have ten types.

the robust Cremer-McLean mechanism is specified and sufficiently general for an arbitrarily large number of agents, it becomes computationally infeasible to solve with larger numbers of agents in the current specification due to the exponential number of constraints. A more efficient implementation will be discussed in the following section. However, for all of the empirical results, the less efficient representation was implemented. The implications of this limited number of buyers is important to consider when interpreting the results.

The two buyer case is likely to generate the most stark contrast between the robust

Cremer-McLean mechanism and the mis-estimated Cremer-McLean mechanism with exit for two reasons. First, with only two buyers, if a buyer decides not to participate, the mechanism immediately becomes a take it or leave it offer which fails to extract full surplus even with a correctly estimated distribution, and it becomes more likely that there will be no buyers who choose to participate in the mechanism, leaving the seller with zero revenue. Second, since the bounds of the estimated conditional distribution for the individual rationality constraint are dependent on the number of types for all other buyers (see equation 4.5) the bounds will be tighter for fewer agents, requiring that less unnecessary surplus is likely to be given to the buyers. However, considering only two buyers will provide a good analysis of an upper bound on the relative performance.

Three different cases are compared against the Cremer-McLean mechanism with no mis-estimation. Of course, the robust Cremer-McLean mechanism with a mis-estimated distribution is compared against the Cremer-McLean mechanism. However, each figure also compares the performance of using the standard Cremer-McLean mechanism with the mis-estimated distribution both assuming buyers can exit optimally and assuming that they are forced to participate in the mechanism even if they believe that there expected value is below their reservation utility. This provides insights into both the relative performance of the robust Cremer-McLean mechanism as compared to the situation under which buyers choose to exit optimally as well as the relative importance of buyer exit versus objective mis-estimation.

Figure 4.1 displays results for when buyers have three types, 0, 1, and 2. The first thing to notice is that for all mis-estimation tolerances tested, the robust Cremer-McLean mechanism performs almost identically to the Cremer-McLean mechanism without mis-estimation, maintaining a relative performance of approximately 99% for all tolerances. This is in stark contrast to the Cremer-McLean mechanism with mis-estimation and exit, which falls in relative performance as the mis-estimation bound increases until it seems to stabilize around 70% of the value of the mechanism without mis-estimation. The mis-estimated Cremer-McLean mechanism without exit exhibits an unusual property; on average, it outperforms the Cremer-McLean mechanism without mis-estimation. This may seem unusual on first consideration, given that the Cremer-McLean mechanism provably extracts full-surplus, however, with the mis-estimated mechanism, there is no longer a strict requirement for buyers to have an expected utility of zero. Therefore, with some probability, the buyers have a negative expected utility. In the three type case, it appears that the probability of forcing a buyer below her reservation utility outweighs the cost of a misspecified objective function. However, it does serve to highlight the relative importance of buyer exit and the misspecified objective function, in that buyer exit seems to be the major loss of revenue in the mis-estimated Cremer-McLean mechanism.

Similar patterns are observed in figures 4.2 and 4.3 for buyers with five and ten types respectively. However, two things happen as the number of types for each buyer is increased. First, the relative performance for the robust Cremer-McLean mechanism as well as the Cremer-McLean mechanism with mis-estimation and both exit and no exit all do relatively worse as compared to the optimal mechanism. Second, the robust Cremer-McLean mechanism appears to converge quite quickly to a long run relative performance. The mechanism converges to its long run performance because the binding constraints quickly become constraints where for all positive outcomes to the buyer, the estimated conditional probability is zero, and for all negative outcomes for the buyer, the estimated conditional probability is one. Once that level is reached, the binding constraints no longer change, and the mis-specification of the objective function has a relatively minor effect on the outcome.

These terminal binding constraints reveal an important aspect of the mechanism. If the set of binding constraints is such that for all negative outcomes the estimated conditional probability is one and for all positive outcomes the estimated conditional probability is zero, then this is equivalent to an ex-post individual rationality constraint. The mechanism that is ex-post individually rational has been well studied, however, it does not extract full surplus in expectation.

4.4 Complexity Analysis for the Robust Cremer-McLean Mechanism

One significant hindrance to the feasibility of implementing the robust Cremer-McLean mechanism is the exponential number of constraints. The Cremer-McLean mechanism has a polynomial number of individual rationality constraints, specifically it includes $\sum_{i=1}^{N} |M_i|$ individual rationality constraints. However, the robust Cremer-McLean mechanism contains $\sum_{i=1}^{N} |M_i||K_i|$ individual rationality constraints, where $|K_i|$ is the number of possible binary strings of length $|M_{-i}|$, a non-polynomial function of the input size. Therefore, the robust Cremer-McLean mechanism contains a non-polynomial number of constraints as a function of the input size.

However, the Cremer-McLean mechanism contains a polynomial number of constraints in the input size. This leads to the following theorem.

Theorem 4. *The Cremer-McLean mechanism can be determined in a time polynomial in the number of types of the buyers.*

Proof. Since the Cremer-McLean mechanism's linear program has a constraint set that is polynomial in the number of types of the buyers, it is always a polynomial number of operations to determine if there is a violated constraint, i.e. all constraints can be checked in polynomial time. That is, the separation oracle for the Cremer-McLean mechanism's linear program is to check all constraints for violation. Since a polynomial time separation oracle exists, the ellipsoid algorithm can solve the linear program that determines the Cremer-McLean mechanism with a complexity polynomial in the size of the input. □

The above proof is not valid for the robust Cremer-McLean mechanism since there exist a non-polynomial number of constraints. However, there is a polynomial time algorithm for determining if there exists a violated constraint.

Theorem 5. *For a set of $x_i(s), p_i(s)$ for all $i \in N$ and $s \in M$, there exists a polynomial time algorithm for determining whether there exists a violated constraint for the linear program that determines the robust Cremer-McLean mechanism.*

Proof. Given $x_i(s), p_i(s)$ for all $i \in N$ and $s \in M$, one can check all of the non-individual rationality constraints individually, since there are a polynomial number of non-individual rationality constraints. Assuming that no non-individual rationality constraints are violated, calculate the following values

$$vc(i, s_i) = \sum_{s_{-i}} (I_{(v_i p_i(s) - x_i(s) \geqslant 0)} \underline{\pi}^*(s_{-i}|s_i) + I_{(v_i p_i(s) - x_i(s) < 0)} \overline{\pi}^*(s_{-i}|s_i))[v_i p_i(s) - x_i(s)]$$

$$\forall i \in N, s_i \in M_i \quad (4.7)$$

where I is an indicator variable. Note that this is a polynomial number of evaluations of this equation in the input as well as operations within each evaluation, specifically there are $\sum_{i=1}^{N} |M_i|$ values that must be checked. Note that for $\{i, s_i\}$, $vc(i, s_i)$ is strictly less than or equal to

$$\sum_{s_{-i}} (P(s_{-i}, p) \underline{\pi}^*(s_{-i}|s_i) + (1 - P(s_{-i}, p)) \overline{\pi}^*(s_{-i}|s_i))[v_i p_i(s_i, s_{-i}) - x_i(s_i, s_{-i})].$$

for all $p \in k_i$, where P is defined as in definition 3. Therefore, if $vc(i, s_i) \geqslant 0$ for all $\{i, s_i\}$, there are no violated constraints. However, if there exists some $\{i, s_i\}$ such that $vc(i, s_i) < 0$, then there exists a constraint that is violated, namely the constraint for the set $\{i, s_i, p'\}$ where p' is the index such that $P(s_{-i}, p') = I_{(v_i p_i(s) - x_i(s) \geqslant 0)}$. □

By theorem 5, there is a separation oracle that can determine in polynomial time in the size of the input whether there exists a violated constraint for the linear program that determines the robust Cremer-McLean mechanism. This leads immediately to the following corollary by the same argument as theorem 4.

Corollary 3. *The robust Cremer-McLean mechanism can be determined in a time polynomial in the number of types of the buyers.*

5

Conclusion

While the Cremer-McLean mechanism as outlined in Cremer and McLean (1985) is striking in that it can extract the full surplus from buyers with private information, this thesis demonstrates that it fails in very specific circumstances that are likely to be encountered in any attempt to implement the mechanism in practice. Namely that mis-estimation of the distribution of buyer types will likely lead to buyers not choosing to participate in the mechanism and, therefore, to significant drops in expected revenue. Then several methods are explored to overcome the sensitivity to the estimated distribution.

First, the ability to estimate the optimal mechanism by a second price auction is explored. The second price auction has been shown to be approximately optimal in a variety of mechanism design settings. However, this thesis demonstrates that the second price auction cannot approximate the Cremer-McLean mechanism with any constant approximation factor. Further, there is no mechanism for regular distributions that discards the correlation structure of buyer types that approximates the Cremer-McLean mechanism with a constant approximating factor.

Then, a mechanism is described that approximates the Cremer-McLean mechanism well in expectation and is robust to mis-estimation of the type distribution. Namely, the mechanism is specified such that buyers never choose to not participate in the mechanism. In order to ensure buyer participation, some surplus is not extracted by the seller. This leftover surplus acts as a cushion for the buyer in case of mis-estimation in a manner that decreases his expected utility from participating in the mechanism. This mechanism is called the robust Cremer-McLean mechanism.

Overall, the robust Cremer-McLean mechanism performs significantly better than the mis-estimated Cremer-McLean mechanism with exit, and in expectation, it captures more than 90% of the full surplus for the average distribution in all cases examined. Further, the mis-estimated Cremer-McLean mechanism's performance drop compared to the optimal mechanism is dominated by the exit of buyers, not by the misspecified objective function. In situations where the robust Cremer-McLean mechanism is feasible to compute and there is likely to be mis-estimation in the distribution of types, empirically, the robust Clemer-McLean mechanism seems like a useful alternative to the standard Cremer-McLean mechanism. Though the robust Cremer-McLean mechanism has a non-polynomial number of constraints in the size of the input, this thesis shows that the mechanism can be computed in polynomial time. Therefore, the robust Cremer-McLean mechanism seems to both perform well and be reasonably efficient to compute.

Going forward, it will be worthwhile to determine a theoretical limit on the performance of the robust Cremer-McLean mechanism as well as explore the empirical performance of the robust Cremer-McLean mechanism when there are more than two buyers. Further, there are likely a large class of distributions where a second price auction performs reasonably well compared to the Cremer-McLean mechanism, and determining those distributions would potentially a provide additional theoretical justifications for why the Cremer-McLean mechanism is not used in practice. Finally, if distribution mis-estimation is a large concern, the desired mechanism would be independent of the prior beliefs. It would be valuable to have a clear understanding of what class of distributions a prior independent mechanism could be expected to perform approximately optimally as compared to the Cremer-McLean mechanism.

Bibliography

Aghassi, M. and Bertsimas, D. (2006), "Robust Game Theory," *Mathematical Programming*, 107, 231–273.

Archer, A., Papadimitriou, C., Talwar, K., and Tardos, E. (2003), "An Appoximate Truthful Mechanism for Combinatorial Auction With Single Parameter Agents," in *Proc. 14th ACM Symp. on Discrete Algorithms*, ACM/SIAM.

Bergemann, D. and Morris, S. (2005), "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

Bulow, J. and Roberts, J. (1989), "The Simple Economics of Optimal Auctions," *The Journal of Political Economy*, 97, 1060–1090.

Cremer, J. and McLean, R. (1985), "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent," *Econometrica*, 53, 345–361.

Cremer, J. and McLean, R. (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Economometrica*, 56, 1247–1257.

Groves, T. (1973), "Incentives in Teams," *Econom*, 41, 617–631.

Hartline, J. (2012), "Approximation in Economic Design," `http://users.eecs.northwestern.edu/~hartline/amd.pdf`.

Hartline, J. and Roughgarden, T. (1998), "Simple Versus Optimal Mechanisms," in *Proc. 11th ACM Conf on Electronic Commerce*.

Jehiel, P. and Moldovanu, B. (2003), "Efficient Design with Interdependent Valuations," *Econometrica*, 69, 1237–1259.

Knight, F. (1921), *Risk, Uncertainty, and Profit*, Boston: Houghton Mifflin.

Krishna, V. (2009), *Auction Theory*, Academic Press, 2 edn.

Lopomo, G. (1998), "The English Auction is Optimal Among Simple Sequential Auctions," *Journal of Economic Theory*, 82, 144–166.

Lopomo, G., Rigotti, L., and Shannon, C. (2009), "Uncertainty in Mechanism Design," Revise and Resubmit at Review of Economic Studies.

Myerson, R. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Myerson, R. (1991), *Game Theory: Analysis of Conflict*, Harvard University Press.

Nisan, N. and Ronen, A. (2001), "Algorithmic Mechanism Design," *Games and Economic Behavior*, 35, 166–196.

Ronen, A. (2002), "On Approximating Optimal Auctions," Extended Abstract.

Ronen, A. and Saberi, A. (2002), "On the Hardness of Optimal Auctions," in *Proc. of the 43rd Symposium on Foundations of Computer Science*, pp. 396–405.

Wilson, R. (1997), *Nonlinear Pricing*, Oxford University Press.