Summaries: Assessing the Robustness of Cremer-McLean Mechanism Design and Maximizing Revenue with Limited Correlation: The Cost of Ex-Post Incentive Compatibility

Prabhat Nagarajan

February 2016

Abstract and Introduction

A result by Cremer and McLean in 1985 shows that "if buyers' valuations are sufficiently correlated, there exists a mechanism that allows the seller to extract the full surplus from efficient allocation as revenue". This paper uses automated mechanism design to examine the sensitivity of the Cremer-McLean result when its main technical assumption is relaxed. The Cremer-McLean result assumes that each buyer valuation corresponds to a unique conditional distribution over external signal(s). It also assumes that these conditional distributions are linearly independent. This paper explores the relaxation of these assumptions by allowing multiple buyer types to have the same distribution over the signal(s)(but different valuations). In this paper, the authors partition the set of buyer types into subsets, where types within the same subset have identical conditional distributions. Ultimately, this paper finds that as the size of these subsets increase, the optimal revenue converges to that of a second-price auction with reserve.

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

In this paper, the setting is as follows: A single buyer, a single external signal, and a single indivisible good that the seller values at 0. This simple case can be extended to the case of more than one buyer. The risk-neutral buyer has a valuation for the object, drawn from the set $V = \{v_1, v_2, ..., v_{|V|}\} \subset [0, \overline{v}]$, where $v_i < v_{i+1}$ for all i. The buyer's type(valuation along with possible additional information) is denoted $\theta \in \Theta = \{1, 2, ..., |\Theta|\}$, and $v(\theta)$ denotes the valuation that is associated with type θ. In addition to the buyer, there is an external signal $\omega \in \Omega = \{1, 2, ..., |\Omega|\}$, where $|\Omega| \geq |V|$. The joint distribution over θ and ω is given by $\pi(\theta, \omega)$. We have the probability $q(\hat{\theta}, \omega)$ that the buyer is allocated the item, and the payment of the buyer $m(\hat{\theta}, \omega)$, where $\hat{\theta}$ is the reported outcome, and ω is the observed signal.

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

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

We here partially present a definition from the paper.

Definition 2. (The Optimal Mechanism Under Incomplete Information) Optimal mechanisms under incomplete information correspond to optimal solutions to the linear program that maximizes the expected revenue of the seller (sum of $m(\theta, \omega)\pi(\theta, \omega)$), with the constraints that the probability of the buyer being allocated the item must be in the interval $[0, 1]$, along with incentive compatibility and individual rationality constraints ($U(\theta, \theta) \geq 0$ and $U(\theta, \theta) \geq U(\theta, \hat{\theta})$).

The Cremer-McLean result has the following assumption: The matrix τ of conditional probabilities $\pi(\omega|\theta)$ where each column is the conditional distribution for a single signal ω, has full rank, that is, the columns are linearly independent.

Theorem. ***Cremer and McLean (1985).*** *Under the above assumption, there exists a solution to the linear program in Definition 2 with an objective value of:* $\sum_{\theta,\omega} \pi(\theta,\omega)m(\theta,\omega) = \sum_{\theta,\omega} \pi(\theta,\omega)v(\theta)$.

Essentially, such a mechanism extracts the buyer's expected valuation as expected payment.

Relaxing Assumption 1

The authors relax Assumption 1 by allowing a range of possible conditional distributions for each unique buyer valuation. Additionally, these distributions can overlap, which also breaks assumption 1. Formalizing these relaxations, the type space is partitioned into sets $X_1, ..., X_k$, such that any two types θ and θ' have the same conditional distribution if and only if they are contained in the same X_i. Note that, if assumption 1 was not relaxed, each type θ would be its own set $X_i = \{\theta\}$. It is careful to note, however, that the conditional distributions of the X_i are linearly independent. This is formalized in the following assumption:

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

A More Efficient Algorithm

Solving the linear program from Definition 2 will yield this optimal mechanism, however, it is computationally difficult. The authors present a novel algorithm for computing the optimal revenue in this setting. In this summary we go over the general notions of the algorithm. However, this requires some definitions and theorems.

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

$$\sum_{\omega \in \Omega} p_\omega H(\boldsymbol{p}, \omega) \geq \sum_{\omega \in \Omega} p_w H(\boldsymbol{p'}, \omega).$$

A strict proper scoring rule is one such that the equation above is only satisfied with equality when $\boldsymbol{p'} = \boldsymbol{p}$.

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

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

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

The intuitive approach towards defining our mechanism is to obtain the optimal reserve price mechanism for each of our X_i, and then add a strictly proper scoring rule (with expected value of 0 for truthful reporting), removing the incentive to misreport as being in a different X_i.

Definition 5 (**(Summary)**). *Let* q_i *and* m_i *be defined for any* i, $1 \leq i \leq k$ *such that for a signal* ω *and type* $\theta \in X_i$, $q_i(\theta, \omega)$ *and* $m_i(\theta, \omega)$ *are equivalent to the optimal reserve price mechanism for the types in* X_i. *We then have a strictly convex function:*

$$G^*(\boldsymbol{p}) = \sum_{\omega \in \Omega} (p_\omega^2 - a_\omega p_\omega).$$

Essentially G^* maps a lottery $\boldsymbol{p}$ over the signals, to the sum for all signals ω, of p_ω^2 (probability of that signal in the lottery $\boldsymbol{p}$) less $a_\omega p_\omega$. What is a_ω? The a_i are given by any solution of $M\boldsymbol{a} = \boldsymbol{b}$, where M is the $k \times |\Omega|$ matrix where row i contains the distribution of types in X_i over the signals. The vector $\boldsymbol{b} \in \mathbb{R}^k$ is defined with the jth element being:

$$\sum_\omega (\pi_j(\omega))^2.$$

*We then define a strictly proper scoring rule H^* using G^* as in the Theorem by **Gneiting and Raftery**. We then compute κ, the minimum loss from misreporting the distribution. We then compute a function $G(\boldsymbol{p})$, and obtain H using G as in the **Gneiting and Raftery** Theorem. We then let $m^*(\theta, \omega) = -H(\pi_i, \omega)$ for $\theta \in X_i$. The authors finally define the mechanism as $q(\theta, \omega) = q_i(\theta, \omega)$ and $m(\theta, \omega) = m^*(\theta, \omega) + m_i(\theta, \omega)$. While we omit many details, the paper (**Assessing the Robustness of Cremer-McLean with Automated Mechanism Design (Albert, Conitzer, Lopomo)**) contains full details.*

Theorem. *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 omit the proof. However, we outline the steps required to prove this theorem. First, the authors prove that the mechanism is well defined and that the scoring rules are proper. They then proceed by showing that the mechanism is incentive compatible and individually rational. After showing that the mechanism is a feasible solution to the linear program outlined in Definition 2, the authors proceed to proving the optimality of this mechanism. □

Theorem (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 in the algorithm is solving $M\mathbf{a} = \mathbf{b}$. Since the columns in M are linearly independent, we can factor it using QR decomposition and can solve for $\mathbf{a}$ in at most $O(|\Omega|^2 k)$ steps. □

Experimental Results

As the number of buyer types increases, the authors compare the runtimes of the linear program from Definition 2, the algorithm from Definition 5, and the algorithm for computing only revenue. The resulting graphs indicate that The linear program runs much slower than the Definition 5 algorithm, which runs significantly slower than the revenue-only algorithm. Additionally, the authors measure the sensitivity of relaxing the the assumptions of the Cremer-McLean mechanism, and its results on seller revenue. As the degree of overlap in the conditional distributions increases, the revenue decreases exponentially from full surplus (Cremer-McLean) and converges to the revenue obtained from a pure reserve price mechanism.

Maximizing Revenue with Limited Correlation: The Cost of Ex-Post Incentive Compatibility

We now summarize a related paper, published by the same authors as the one above, continuing much of the research from the previous paper. This paper has many of the same definitions as the previous paper, and we may refer back to them. Additionally, note that CM refers to Cremer-McLean.

Abstract and Introduction

In this paper, the authors "explore the implications of Bayesian versus ex-post IC in a correlated valuation setting". They "generalize the full extraction results to settings that do not satisfy the assumptions of CM". The authors "demonstrate that the expected revenue from the optimal ex-post IC mechanism guarantees at most a $(|\Theta| + 1)/4$ approximation to that of a Bayesian IC mechanism, where $|\Theta|$ is the number of bidder types". The authors also show that "the average expected revenue achieved by Bayesian IC mechanisms is significantly larger than that for ex-post IC mechanisms". This paper assumes ex-interim individual rationality and examines "the case of a single bidder whose valuation is correlated with an external signal that both the seller and bidder observe ex-post". However, the results obtained for a single bidder can be easily extended to multiple bidders". This paper characterizes the "necessary and sufficient conditions for full surplus extraction in both Bayesian and ex-post IC settings that are a significant relaxation from those of Cremer and McLean".

Notation and Review of Cremer and McLean

Definition 2 ((Ex-Post Incentive Compatibility))**.** A *mechanism* (q, m) *is* ex-post incentive compatible (IC) *if:*

$$\forall \theta, \hat{\theta} \in \Theta, \omega \in \Omega : U(\theta, \theta, \omega) \geq U(\theta, \hat{\theta}, \omega)$$

Definition 3 ((Bayesian Incentive Compatibility).)**.** A *mechanism* (q, m) *is* Bayesian incentive compatible(IC) *if:*

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

Observation: A mechanism is ex-post IC $\implies$ this mechanism is Bayesian IC.

Definition 4 ((Ex-Interim Individual Rationality))**.** A *mechanism* (q, m) *is* ex-interim individually rational (IR) *if:*

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

Note that ex-interim IR is assumed for the CM full surplus extraction result. It ensures that the bidder has a non-negative *expected* utility, but does not disallow negative utility. In fact, a result by Lopomo in 2001 shows that if ex-interim IR is replaced by a condition guaranteeing non-negative utility (ex-post IR), then full surplus extraction is impossible in the general case.

Definition 5 ((Optimal Mechanisms).)**.** *A mechanism* (q, m) *is an* optimal ex-post mechanism *if under the constraint of ex-interim individual rationality and ex-post incentive compatibility it maximizes the following:*

$$\sum_{\theta, \omega} m(\theta, \omega) \pi(\theta, \omega)$$

A mechanism that maximizes the above under the constraint of ex-interim individual rationality and Bayesian incentive compatibility is an optimal Bayesian mechanism.

Definition 6 ((Full Social Surplus Extraction as Revenue).)**.** *We say that a mechanism* extracts the full social surplus as revenue in expectation *if there exists a (Bayesian or ex-post) mechanism such that:*

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

Necessary and Sufficient Conditions for Full Surplus Extraction as Revenue

In this section, the authors characterize the conditions required for full social surplus extraction as revenue, regardless of the correlation structures. The results even guarantee full surplus extraction when a subset of the conditional beliefs are not linearly independent.

Lemma 1. *A mechanism* (q, m) *extracts full surplus if and only if* $q(\theta, \omega) = 1$ *and* $U(\theta, \theta) = 0$ *for all* $\theta \in \Theta, \omega \in \Omega$.

Proof. The item is always allocated. $U(\theta, \theta) = 0 \implies$ the bidder paid his value for the item, which by our definition, means that our mechanism extracts full surplus. □

Theorem 2. *For a given (π, V, Ω), full surplus extraction is possible for an ex-post incentive compatible mechanism if and only if there exists a linear function $G : \mathbb{R}^{|\Omega|} \to \mathbb{R}$ such that $G(\pi(\bullet, \theta)) = -v(\theta)$.*

Proof. Proof omitted. □

Theorem 3 ((Full Surplus Extraction with Bayesian IC).)**.** *For a given (π, V, Ω), full surplus extraction is possible for a Bayesian incentive compatible mechanism if and only if there exists a convex function $G : \mathbb{R}^{|\Omega|} \to \mathbb{R}$ such that $G(\pi(\bullet|\theta)) = -v(\theta)$.*

Proof. Proof omitted. □

The key components of Theorems 2 and 3 are their characterizations of the conditions required for full extraction. Theorem 2 tells the exact setting of an ex-post incentive compatible mechanism in which full surplus extraction is possible. Similarly, Theorem 3 tells us the exact setting that a Bayesian IC mechanism can obtain full surplus extraction.

Theorem 4. *The expected revenue generated by an optimal ex-post mechanism guarantees at most a $(|\Theta| + 1)/4$ approximation to the expected revenue generated by an optimal Bayesian mechanism.*

Proof. Proof omitted. □

Simulation Results

Note: With probability 1, if $|\Omega| = |\Theta|$, then the CM assumption is satisfied, and full social surplus can be extracted. The experimental results indicate that optimal Bayesian mechanisms converge to full extraction with fewer external signals than optimal ex-post mechanisms. That is, the optimal ex-post mechanism generates less revenue until $|\Omega| = |\Theta|$. Furthermore, the experimental results indicate that as the correlation between the bidder's type and the external signal approaches 1, both the optimal Bayesian and ex-post mechanisms approach full surplus extraction. However, for all correlation values, the optimal Bayesian mechanism performs better.

Conclusion

The empirical results seem to indicate that Bayesian incentive compatibility is very important when there are few signals and when the correlation is low. Given the empirical results of the authors, they believe that when implementing optimal mechanisms through learning beliefs of bidders over the external signals, focusing on Bayesian IC mechanisms would benefit.