


Mechanisms With Costly Knowledge

by

Atalay M. Ileri

B.S. Computer Engineering
Bilkent University, 2014



Submitted to the Department of Electrical Engineering and Computer Science
in Partial Fulfillment of the Requirements for the Degree of

Master of Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2016

© 2016 Massachusetts Institute of Technology. All rights reserved

Signature redacted

Signature of Author _____

Department of Electrical Engineering and Computer Science

May 5, 2016

Signature redacted

Certified by _____

Silvio Micali

Associate Department Head of EECS

Thesis Supervisor

Signature redacted

Accepted by _____

Leslie A. Kolodziejcki

Professor of Electrical Engineering

Graduate Officer

Mechanisms With Costly Knowledge

by

Atalay M. Ileri

Submitted to the Department of Electrical Engineering and Computer Science
on May 5, 2016 in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Computer Science

ABSTRACT

We propose investigating the design and analysis of game theoretic mechanisms when the players have very unstructured initial knowledge about themselves, but can refine their own knowledge at a cost.

We consider several set-theoretic models of “costly knowledge”. Specifically, we consider auctions of a single good in which a player i 's only knowledge about his own valuation, θ_i , is that it lies in a given interval $[a, b]$. However, the player can pay a cost, depending on a and b (in several ways), and learn a possibly arbitrary but shorter (in several metrics) sub-interval, which is guaranteed to contain θ_i .

In light of the set-theoretic uncertainty they face, it is natural for the players to act so as to minimize their regret. As a first step, we analyze the performance of the second-price mechanism in regret-minimizing strategies, and show that, in all our models, it always returns an outcome of very high social welfare.

Thesis Supervisor: Silvio Micali

Title: Associate Department Head of EECS

1 Introduction

It is a traditional tenet of mechanism design that each player i knows himself *perfectly*. This assumption strikes us as an oversimplification of our daily experience. Except for whatever innate knowledge we may have, most of us incur a considerable personal cost for the knowledge we acquire. Moreover, we seldom reach “perfect knowledge”. Whether on an individual basis or a collective basis (e.g., in Science) we simply *refine* our knowledge. Accordingly, we believe it is important to investigate mechanism design when the players are very self uncertain, but have the ability of reducing their own uncertainty at a *cost*. In this paper, we put forward a first class of “costly knowledge models” and use it to analyze the efficiency guarantees of second-price mechanism.

Our Goals

Distributional approaches naturally come to mind to model costly knowledge. Indeed, the traditional way to model uncertainty is via probability distributions. Although we encourage others to work on distributional models of costly knowledge, and we may choose to do so ourselves in the future, in this paper, to explore a less-travelled road, and start investigating *set-theoretic* models of costly knowledge.

Our first such model is purposely very basic; essentially, it is a costly version of binary search. More complex models are discussed in section 5.

Our first result analyzes the efficiency guarantees of the second price mechanism in all of our models. We believe and hope that the viability of these models can be also established for more complex mechanisms, new or old.

Our First Model

A player i has no information about the true valuations of his opponents. Moreover, at each point in time, the only information he has about his own true valuation, θ_i , consists of a “knowledge” interval $[a, b]$ guaranteed to contain θ_i .

If $[a, b] \neq [0, 0]$, i may “refine” such knowledge interval by paying a cost inversely proportional to $\delta = \frac{b-a}{b}$ and learn whether θ_i lies in the first or the second half of $[a, b]$.

We respectively refer to $b - a$ and δ as the (absolute) *uncertainty* and the (relative) *inaccuracy* of i ’s knowledge.

Remarks

- *Incomplete Preferences*. Having set-theoretic uncertainty about his true valuation θ_i , a player i may not be able to ‘compare’ all possible outcomes. If his current knowledge about θ_i consists of an interval $[a, b]$, then he is only sure that he prefers an outcome ω to another outcome ω' if his utility for ω is greater than or equal to his utility for ω' , no matter which element of $[a, b]$ may be his true valuation.

Of course, one might always and more simply assume that our players have complete preferences (e.g., that they maxmin preferences [22]), or that they behave as if —say— $\theta_i = \frac{a+b}{2}$). However,

our players have incomplete preferences.

- *Enlarged Mechanisms.* Explicitly or not, our model *de facto* turns every given normal-form mechanism into one of extensive form. To be sure, all players continue to submit their reports “simultaneously”. However, prior to submitting his report, based on his current knowledge, each player can choose whether or not to refine it. Formally, therefore, he chooses his own report by executing a single-player extensive-form (sub)mechanism.
- *Defining Accuracy.* Another natural choice for defining the inaccuracy of a knowledge interval $[a, b]$ of a player i is $\delta' = \frac{b-a}{a+b}$. In this case, δ represents the percentage with which i knows his own true valuation. Indeed, letting $x = \frac{a+b}{2}$, $\theta_i \in [a, b]$ is equivalent to say that

$$“\theta_i = x \pm \delta'x”.$$

These two (and other) ways of defining inaccuracy are essentially semantically equivalent. However, choosing $\delta = \frac{b-a}{b}$ enables us to express the efficiency guarantees of the second-price mechanism more easily in *all* our models of costly knowledge.

- *Inaccuracy-Based Cost.* Of course, a player i 's cost for refining a knowledge interval $[a, b]$ could be defined as a function of the interval's uncertainty, rather than its inaccuracy, but we find the latter choice more meaningful. For instance, consider two possible knowledge intervals for i : $[100, 200]$ and $[10^6, 10^6 + 100]$. In both cases, the uncertainty is 100. Intuitively, however, the effort required from i to determine whether θ_i is in the first half of the interval (or not) is much smaller in the first case than in the second.
- *Simplicity vs. Robustness.* In general, given a short interval $[a, b]$, a player i may rarely be sure that his true valuation θ_i belongs to $[a, b]$. The situation, however, is very different if $[a, b]$ is an interval that i obtains after investing a lot of effort in learning θ_i .¹ In any case, in the classical model, i is assumed to know θ_i exactly. Accordingly, even perfect self knowledge could be achievable (although without cost, or via a process preceding the execution of a given mechanism).

Still, when θ_i lies very very close to the middle, m , of i 's current knowledge interval, it may be unlikely that, by means of a single (and possibly low-cost) step, i can disambiguate whether $\theta_i \leq m$ or not. Indeed, our main reason to analyze first the basic knowledge model is conceptual and calculation simplicity, not robustness.

As we shall see in section 5, however, the social-welfare guarantee of the second-price mechanism remains the same when the refinement of a given knowledge interval yields an arbitrary interval, of half the size, that continues to include θ_i . That is, our analysis continues to hold in a model that is more robust (in that no point of a knowledge interval constitutes a “point of discontinuity”), but more complex (indeed, non-deterministic).

¹Indeed, for most goods, i can be confident, without investing any effort, that his valuation lies between 0 and one quadrillion dollars. And his subsequent confidence in smaller and smaller uncertainty about θ_i proceeds from the investments he makes to clarify the actual value of θ_i .

- *Individualized Costs.* As we shall see, two players whose knowledge interval is the same, may have different costs for refining it.

Rationality It is continuously debated whether players should be modeled as utility maximizers or as regret minimizers. When all players have available dominant strategies, the debate is mute, because a dominant strategy is both the only undominated strategy and the only regret-minimizing one. But when this is not the case, or when dominant-strategy mechanisms are essentially useless, it becomes important to take into consideration mechanisms implementing a given social-choice function, both in undominated strategies and in regret-minimizing ones. After all, utility maximizers will only play weakly undominated strategies, and regret-minimizers will only play regret-minimizing strategies, and both sets of strategies are guaranteed to be non-empty, at least in every mechanism in which every player has finitely many pure strategies (hardly a restriction in practice). In light of our model, of course, the classical notions of an undominated strategy and a regret-minimizing strategy are naturally extended so as also to take into account the set-theoretic knowledge the players have about themselves.

A First Result

Our main aim is to put forward a new model, more precisely, a new class of models. But when proposing new models, it is also important to prove that they are amenable to rational analysis.

Our only theorem, for now, proves that the second-price mechanism can be analyzed in all our costly knowledge models, and, in fact, continues to guarantee high social welfare in regret minimizing strategies.

In light of the set-theoretic self uncertainty a player faces in our model, it is natural for him to act so as to minimize his regret. Nonetheless, we find it important to point out that no mechanism can hope to deliver significant social welfare in dominant strategies or in undominated strategies. For notational simplicity, let us point out both facts, assuming that there are just two players.

Denote by I_i the initial knowledge of a player i . Then,

For all two-player dominant-strategy mechanism M (in which each player has a strategy guaranteeing him utility at least 0, all refinement costs are incurred by the players, and all payments are made by the players) and for all possible I_1 and I_2 , there exist true valuations $\theta_1 \in I_1$ and $\theta_2 \in I_2$ for which M misses the maximum social welfare at least by the length of $I_1 \cap I_2$.

The above statement implies that the performance of any dominant-strategy mechanism is quite poor in what we consider the typical case: namely, when the players have essentially the same initial knowledge $[a, b]$, so that the length of the intersection of their initial knowledge is approximately $b - a$. In fact, when this is the case, the above statement implies that no dominant-strategy mechanism can do better than assign the good to a random player. The above statement is not hard to prove: indeed, its proof is sketched in this footnote.²

²First, it is easy to show that the revelation principle holds in our model. Accordingly, we might focus

Unfortunately, the impossibility of guaranteeing significant social welfare also applies in practice to undominated-strategies mechanisms. In fact, as long as they are *finite* (i.e., assigning finitely many pure strategies to each player³), such mechanisms cannot guarantee a social welfare higher than that obtainable by assigning the good to a random player. More precisely,

For all finite mechanism M (probabilistic or not), all $\ell > 0$, all $\varepsilon \in (0, \ell)$, and all initial knowledge such that I_i and I_{-i} whose self uncertainty is ℓ and whose intersection has length ε , there exist a profile of true valuation θ and a profile s of undominated strategies for which M misses the maximum social welfare by $\ell - \varepsilon/2$.

The proof of this second statement easily follows from the Undominated Intersection Lemma (i.e., lemma E.1) of [12].⁴

Since dominant- and undominated-strategy mechanisms are not very useful, we now consider regret and show that it is possible to guarantee excellent social welfare in regret-minimizing strategies. In fact, without having to look very far, we prove that it suffices to consider the second-price mechanism!

Theorem 1 (Informal): *For all initial knowledge intervals $[a_i, b_i]$, if the players are regret minimizers, then the second-price mechanism misses the actual maximum social welfare by at most $O(\sqrt{b_{\max}})$, that is, the square root of the maximum possible initial valuation.*

A precise statement of Theorem 1 and a sketch of its proof, relative to our first costly knowledge model, can be found in section 4. Extensions of Theorem 1 to our other costly knowledge models are discussed in section 5.

Remark

We note that, due to a lemma of [10], our main result also continues to hold when the players are utility maximizers who resort to regret minimization solely to refine further their set of undominated strategies (if it indeed contains multiple strategies).

2 Related Work

Our model of self-uncertainty can be considered a special case of the one put forward a century ago by Knight [27], and later on refined by Bewley [3]. In their model, the only knowledge of a player i consists of a set of distributions, from one of which θ_i has been drawn. (Indeed, since our utility functions are convex, a risk-neutral player i may *de facto*

our attention to direct dominant-strategy truthful mechanisms (i.e., mechanisms in which each player i is allowed to report an interval, and reporting truthfully I_i is i 's best strategy). Second, it is trivial to show that a dominant-strategy truthful mechanism that does not inject money into the system and providing each player with a strategy guaranteeing him a non-negative utility will *never* (i.e., no matter what the players' reports may be) refine the knowledge of any player. Third, and finally, it is trivial to check that our result statement holds when the mechanism does not refine the player's knowledge.

³This technical restriction is very mild in practice, since *de facto* players have finitely many strategies. Certainly this is the case for all mechanisms played via computers.

⁴In our setting, this powerful lemma can be informally simplified as follows. If I_i and I'_i are two intervals, representing the knowledge of a player i about herself, containing at least two valuations in common, then i 's corresponding set of undominated strategies contain at least one strategy in common.

shrink each such distribution to its expected value, so that his knowledge about himself consists of a set, that is, the mentioned generalization of our self-uncertainty model.) Knightian uncertainty, however, does not envisage the possibility of decreasing uncertainty at a cost.

Knightian uncertainty has been extensively studied in decision theory, see for example [1, 18, 35, 34, 16, 39, 22, 31, 5, 4]. There is also a quite rich literature on equilibrium with incomplete preferences (e.g., [32, 21, 41, 20, 37]) and on ambiguous mechanisms (e.g., [17, 6]). All this literature, however, is not directly relevant to our work.

A mechanism for rent extraction with Knightian uncertainty has been studied by Lopomo, Rigotti and Shannon [29], but in a model quite different from ours, and without considering any notion of knowledge refinement. (See also [30].)

Kiekintveld et al. [26] study a two-player Stakelberg game in which a *defender* needs to allocate his resources when he has uncertainty, also modeled as an interval, about the *attacker's* payoffs. However, the authors do not envision any self uncertainty for the defender or the attacker, nor any way to reduce this uncertainty, with or without cost. Our model of self-uncertainty (generalized from intervals to sets) actually coincides with that used by Chiesa, Micali, and Zhou to analyze single-good auctions [9] and multi-unit auctions [12] in undominated strategies, as well as unrestricted combinatorial auctions, both in undominated strategies and in regret-minimizing ones [11]. They too, however, do not consider any type of knowledge refinement.

In a distributional model, uncertainty reduction, with or without cost, has been studied by Thompson and Leyton-Brown for a variety of auctions [43]. The same authors prove, in a separate paper [44], that, when the players can refine their own valuations, dominant strategy single-good auctions coincide with sequential posted-price auctions. Celis et al. [8] study revenue auctions in which the players can (non-adaptively) pay an upfront cost and reduce their own uncertainty to a predetermined desired level. Celis et al. [7] investigate revenue maximization in auctions in which a buyer (1) knows a distribution from which his own true valuation has been drawn, and (2) can, at a cost, privately refine his distribution. However, these works' uncertainty is not Knightian like ours. Indeed, they assume that every player knows the distribution from which his true valuation has been drawn, and can elicit a signal allows him to further constrain this distribution.

The the notion of regret minimization was introduced by Savage [38] (reinterpreting Wald [45]) and refined by Milnor [33] and recently Stoye [42]. Many empirical works provide evidence of regret minimization, for instance [2, 14, 13, 25]. Coricelli et al. [15] actually argue that neurological evidence exists for human propensity for regret minimization. Regret minimization has been used to analyze mechanisms, for instance in [28, 19, 40, 24, 36]. Halpern and Pass [23] put forward the solution concept of *iterated regret minimization*, and argue that it is indeed a compelling one. Our solution concept lies between traditional and iterated regret minimization. In the latter concept, the players assume that their opponents are regret minimizers. In ours, the players assume nothing about their opponents' rationality.

3 Preliminaries

Below we define our refinable knowledge model as a costly version of the basic binary search process. (More general models are considered in Section 5.)

3.1 The Basic Costly Knowledge Model

For each player $i \in N$,

- The *true valuation* of i , θ_i , is a non-negative real.
A *knowledge interval* for i is an interval of non-negative reals containing θ_i .
If $[a, b]$ is such an interval, then its *uncertainty* is $b - a$, and its *inaccuracy* is $\frac{b-a}{b}$.
- R_i is the *refining* function that, given a knowledge interval I , of i , $I = [a, b]$, returns $[a, \frac{a+b}{2}]$ if $\theta_i \leq \frac{a+b}{2}$, and $[\frac{a+b}{2}, b]$ otherwise.⁵
We refer to each evaluation of R_i as a *refinement*, to the interval $R_i(I)$ as a *refinement* of I , and to evaluating R_i on I as *refining* I . When the player i is clear, we may write R instead of R_i .
- Player i starts with an *initial knowledge interval*, denoted by $I_i^0 = [a_i^0, b_i^0]$ or more simply by $I_i = [a_i, b_i]$.
The *knowledge of i after t refinements*, $I_i^t = [a_i^t, b_i^t]$, is the interval obtained from I_i by iterating t times the function R . (I.e., $I_i^1 = R(I_i)$, $I_i^2 = R(R(I_i))$, ...)
- i 's cost for refining an interval I of inaccuracy δ , $C_i(I)$, is $\frac{c_i}{\delta}$, where c_i is a positive constant known to i .
- *Self Uncertainty.* In our model, a player i has no information about the true valuations of his opponents. Moreover, at each point in time, the only information he has about his own true valuation of the good, θ_i , consists of an interval $[a, b]$ guaranteed to contain θ_i . We refer to such an interval $[a, b]$ as i 's (current) "knowledge interval", or more simply as i 's (self) "knowledge", to $b - a$ as i 's current (self) uncertainty, and to the ratio $\delta = \frac{b-a}{b}$, if $[a, b] \neq [0, 0]$, as the inaccuracy of i 's knowledge, or the interval $[a, b]$.
- *Refinements.* A player i has always the option of refining his current knowledge, $[a, b]$, at a cost $C_i(a, b)$. By exercising this option, i 's new knowledge becomes $[a, \frac{a+b}{2}]$ if $\theta_i \leq \frac{a+b}{2}$, and $[\frac{a+b}{2}, b]$ otherwise.
(Of course, $[\frac{a+b}{2}, b]$ implies that $\theta_i \in (\frac{a+b}{2}, b]$, but for uniformity sake all knowledge intervals are closed.)
- *Costs.* A player's refinement cost is proportional to the inaccuracy of his knowledge. If the inaccuracy of $[a, b]$ is δ , then $C_i(a, b) = \frac{c_i}{\delta}$, where c_i is a player-dependent positive constant.

⁵More precisely, but "less uniformly", if $R([a, b]) = [\frac{a+b}{2}, b]$, then $\theta_i \in (\frac{a+b}{2}, b]$.

3.2 Single-Good Auctions in Our Model

As usual, an auction consists of (1) a context, specifying the set of possible outcomes, the players (including their initial knowledge), and their preferences over outcomes, and (2) a mechanism, specifying the players' strategies and how strategies lead to outcomes.

Our Contexts

- An outcome consists of
 - an *allocation*, a profile of bits a , $\sum_i a_i \leq 1$, where for each player i , $a_i = 1$ if and only if the good is allocated to i ;
 - a profile of *prices*, $p \in \mathbb{R}^n$; and
 - a profile of *total refinements*, r , where r_i is the number of times player i has refined his knowledge.
- The utility of a player i , with valuation v_i and initial knowledge I_i^0 , for an outcome $\omega = (a, p, r)$, is

$$u_i(\omega) \triangleq a_i \cdot v_i - p_i - \sum_{t=0}^{r_i-1} C_i(I^t).$$

Our Mechanisms

We consider mechanisms in which (like in the second-price mechanism), each player i reports a bid $\beta_i \geq 0$ simultaneously with his opponents. However, i 's set of pure strategies, S_i , does not coincide with the set of possible bids, $\mathbb{R}_{\geq 0}$, but with the set of strategies in the single-player extensive-form (sub)mechanism pictorially described below




Figure 1: The Bidding Submechanism of Player i

That is, in the bidding submechanism of player i , the decision nodes, D_0, D_1, \dots , are pictorially represented by “empty circles” and correspond to the number of times i has refined his knowledge interval. The terminal nodes are represented by “full circles”, and correspond to i 's actual bids.

A bit more precisely,

- Player i starts executing his bidding submechanism at the decision node D_0 , the “root”, where the information available to him consists of his initial knowledge, I_i^0 .
- At every decision node D_t , the information available to i consists of the knowledge interval I_i^t , and i 's action set is $\{R_i\} \cup \mathbb{R}_{\geq 0}$.
- At a decision node D_t , if he chooses an action $\beta_i \in \mathbb{R}_{\geq 0}$, then i terminates executing his bidding submechanism and reports only the bid β_i to the mechanism. Else (if i chooses the refining action R_i), the decision node D_{t+1} is reached and the bidding submechanism continues.

In every auction mechanism we consider, the underlying bidding mechanism for each player i is as above. Accordingly, the set of all pure strategies of i , S_i , always coincides with the set of all functions from $\{D_0, D_1, \dots\}$ to $\{R_i\} \cup \mathbb{R}_{\geq 0}$.

The outcome of a mechanism M under a strategy profile s (i.e., the outcome produced by M relative to the bid profile β corresponding to s) is $M(s)$.

3.3 Regret

In a mechanism M , relative to i 's initial knowledge I_i ,

- i 's regret for a strategy $s_i \in S_i$ is

$$\begin{aligned} \text{reg}_i(I_i, s_i) &\triangleq \max_{v_i \in I_i} \max_{s_{-i} \in S_{-i}} \max_{s'_i \in S_i} \\ &u_i(v_i, M(s'_i, s_{-i})) - u_i(v_i, M(s_i, s_{-i})). \end{aligned}$$

- i 's set of regret-minimizing strategies is

$$\text{REG}_i(I_i) \triangleq \{s_i : s_i = \underset{s'_i \in S_i}{\text{argmin}} \text{reg}_i(I_i, s'_i)\}.$$

The set of all profiles of regret-minimizing strategies of M , relative to an initial knowledge profile I , is

$$\text{REG}(I) \triangleq \text{REG}_1(I_1) \times \dots \times \text{REG}_n(I_n).$$

3.4 Social Welfare

As usual, social welfare and maximum social welfare are defined relative to the true valuation profile θ (independent of the fact that in our model each player i may have inaccurate knowledge about his own θ_i .) Indeed, The social welfare of an outcome $\omega = (a, p, r)$ is

$$\text{SW}(\omega) \triangleq \theta_i \text{ if } a_i = 1.$$

The maximum social welfare is

$$\text{MSW} \triangleq \max_{\omega} \text{SW}(\omega) = \max_i \theta_i.$$

4 A First Theorem

The (Non-Deterministic) Second Price Mechanism⁶

The second price mechanism is the normal form mechanism non-deterministically allocating the good to a player reporting the highest valuation, and choosing the price of a player i to be the second highest reported valuation, if the good has been allocated to i , and 0 otherwise.

Notation

We denote the second-price mechanism by SP , the set of players by N , the number of players by n , and let $-i \triangleq N \setminus \{i\}$ for every player i .

Relative to an initial knowledge profile $I = ([a_1, b_1], \dots, [a_n, b_n])$, lexicographically breaking ties if needed, let

$$\phi \triangleq \operatorname{argmax}_{i \in N} c_i b_i.$$

(As we shall see, ϕ is the player with the highest final uncertainty. That is, the player which, at the decision node in which he finally decides to bid rather than refining his own knowledge, has the longest knowledge interval. The length of that interval will actually be $2\sqrt{c_\phi b_\phi}$.)

Theorem 1 *In the basic costly knowledge model, for all initial knowledge profiles $I = ([a_1, b_1], \dots, [a_n, b_n])$, all true valuation profiles $\theta \in [a_1, b_1] \times \dots \times [a_n, b_n]$, and all strategy profiles $s \in \text{REG}(I)$,*

$$SW(\text{SP}(s)) \geq MSW - 2\sqrt{c_\phi b_\phi}.$$

Proof Sketch

We begin by proving that the set of regret-minimizing strategies of each player i , $\text{REG}_i(I_i)$, consists of a single strategy, s_i^* , and actually explicitly computing such s_i^* . To this end, assume for a moment that knowledge refinement is no longer an option, and consider a regret minimizing player i , whose (current) knowledge interval is $I'_i = [a'_i, b'_i]$. Then, i must, for each possible bid β'_i , $\text{reg}_i(I'_i, \beta'_i)$, and report a bid

$$\beta_i^* \in \operatorname{argmin}_{\beta_i} \text{reg}_i(I'_i, \beta_i).$$

Lemma 1 $\beta_i^* = \frac{a'_i + b'_i}{2}$.

Proof of Lemma 1.

We prove Lemma 1 by considering the function $\text{reg}_i(I'_i, \cdot)$ and showing the following two properties:

⁶Should the second-price mechanism break ties deterministically, or at random, regret may not be well defined. (In particular, there may not be a strategy subprofile for a player i 's opponents which maximizes i 's regret. Even relying on suprema, rather than maxima, is problematic.)

(a) $reg_i(I'_i, \cdot)$ is strictly decreasing in the domain $\left[0, \frac{a'_i + b'_i}{2}\right]$; and

(b) $reg_i(I'_i, \cdot)$ is strictly increasing in the domain $\left[\frac{a'_i + b'_i}{2}, \infty\right)$.

Proof Sketch of Property (a).

Consider a generic bid α_i , $0 \leq \alpha_i \leq \frac{a'_i + b'_i}{2}$.

Let $\bar{\alpha}_i$, $\bar{\alpha}_{-i}$, and $\bar{\theta}_i$ respectively be a bid of i , a bid subprofile of the opponents of i , and a valuation of i in $[a'_i, b'_i]$ “achieving the regret of α_i ”. That is,

$$(\bar{\alpha}_i, \bar{\alpha}_{-i}, \bar{\theta}_i) \in \underset{(\alpha'_i, \alpha'_{-i}, \theta'_i)}{\operatorname{argmax}} u_i((\alpha'_i, \alpha'_{-i}), \theta'_i) - u_i((\alpha_i, \alpha'_{-i}), \theta'_i).$$

It can also be seen (by a tedious case analysis) that for any such $\bar{\alpha}_i$, $\bar{\alpha}_{-i}$, and $\bar{\theta}_i$,

$$\max_{j \in -i} \bar{\alpha}_j = \alpha_i, \quad \bar{\alpha}_i = \alpha_i + 1, \quad \text{and} \quad \bar{\theta}_i = b'_i.$$

Thus, the second price mechanism

- under the strategy profile $(\alpha_i, \bar{\alpha}_{-i})$, may assign the good to an opponent of i , so that i 's (worst case) utility is 0, and
- under the strategy profile $(\bar{\alpha}_i, \bar{\alpha}_{-i})$, assigns the good to i , so that i 's utility is $b'_i - \alpha_i$.

Accordingly,

$$\begin{aligned} reg_i(I'_i, \alpha_i) &= \\ u_i((\bar{\alpha}_i, \bar{\alpha}_{-i}), \bar{\theta}_i) - u_i((\alpha_i, \bar{\alpha}_{-i}), \bar{\theta}_i) &= \\ b'_i - \alpha_i. \end{aligned}$$

The above equality indeed shows that $reg_i(I'_i, \alpha_i)$ decreases with α_i . □

Since the bid α_i could coincide with $\beta_i^* = \frac{a'_i + b'_i}{2}$, our last equality also shows that the regret of β_i^* is

$$reg_i(I'_i, \beta_i^*) = \frac{b'_i - a'_i}{2}. \tag{1}$$

Proof Sketch of Property (b).

The proof of property (b) is essentially symmetric to that of property (a).⁷ □

Since properties (a) and (b) hold, so does Lemma 1. ■

Consider now the following two options for a regret-minimizing player whose current knowledge interval is $[a'_i, b'_i]$.

⁷In particular, $\max_{j \in -i} \bar{\alpha}_j = \alpha_i$, $\bar{\alpha}_i = \alpha_i - 1$, and $\bar{\theta}_i = a'_i$.

- *Option 1: i bids immediately without any refinement.*

Accordingly, lemma 1 tells us that i 's bid is β_i^* . Thus, equation (1) implies that i 's regret in this option is $\frac{b'_i - a'_i}{2}$, that is, half of his current uncertainty.

- *Option 2: i refines his knowledge once, and then bids.*

By so doing, his total regret would be the sum of (1) the cost of the refinement and (2) the regret of bidding after he learns his new knowledge interval $[a''_i, b''_i]$. By definition, the above cost is $\frac{c_i b'_i}{b'_i - a'_i}$. And again by Lemma 1 and equation 1 the above regret is $\frac{b'_i - a'_i}{4}$ (because the refinement halves i 's original uncertainty).

Accordingly, i will choose option 2 if and only if

$$\frac{c_i b'_i}{b'_i - a'_i} + \frac{b'_i - a'_i}{4} < \frac{b'_i - a'_i}{2}.$$

That is, after some manipulations, if and only if

$$2\sqrt{c_i b'_i} < b'_i - a'_i.$$

Applying this principle from the very beginning of i sub-mechanism, when i 's knowledge is (in our notation) $[a_i, b_i]$, we see that i has a single regret-minimizing strategy, s_i^* , so defined: For all decision node D_t ,

$$s_i^*(D_t) = \begin{cases} R_i & \text{if } b_i^t - a_i^t > 2\sqrt{c_i b_i^t} \\ \frac{a_i^t + b_i^t}{2} & \text{otherwise.} \end{cases}$$

Having understood the unique regret minimizing strategy s_i^* of each player i , we are ready to finish sketching the proof of our theorem. The best way to do so is pictorially.

It is easy to see that the worst scenario for the realized social welfare of the second-price mechanism is that illustrated in the figure below.




Figure 2:

In the above picture, red is the color associated to the player with the maximum true valuation, player m . The red interval represents m 's "final knowledge", $[a_m^*, b_m^*]$, that is the knowledge interval of m when, executing the strategy s_m^* , he decides to place the bid β_m^* . (Indeed notice that β_m^* lies in the middle of m 's final knowledge.) The red dot indicates the actual value of his true valuation within this interval. The corresponding "situation" for the player with the highest final uncertainty, player ϕ , is instead depicted in blue. Our

analysis of the regret-minimizing strategies s_i^* 's guarantees to upper-bound the length of the red interval in terms of m 's initial knowledge interval: that is,

$$b_m^* - a_m^* \leq 2\sqrt{c_m b_m}.$$

Analogously, the length of the blue interval is upper-bounded in terms of ϕ 's initial knowledge interval as follows:

$$b_\phi^* - a_\phi^* \leq 2\sqrt{c_\phi b_\phi}.$$

The fact that the blue dotted line is slightly on the right of the red dotted line wants to indicate that $\beta_\phi^* > \beta_m^*$. Accordingly, the good is allocated to ϕ , so that the social welfare actually realized by second-price mechanism in this case is $\sqrt{c_\phi b_\phi} + \sqrt{c_m b_m}$ smaller than the maximum one. Since, by definition, $c_\phi b_\phi \geq c_m b_m$, we conclude that, as we wanted to show,

$$SW(\mathbb{SP}(s^*)) \geq MSW - 2\sqrt{c_\phi b_\phi}.$$

■

5 Generalizations

Alternative Choices of Inaccuracy

As already mentioned, another natural choice for the inaccuracy of a knowledge interval $[a, b]$ is $\delta' = \frac{b-a}{a+b}$.

A slightly more general choice is

$$\delta'' = \frac{b-a}{g_i([a, b])}.$$

For such δ'' ,⁸ the social-welfare guarantee of the second-price mechanism of course continues to be expressed in terms of the players' longest "final knowledge interval": namely,

$$SW(\mathbb{SP}(s^*)) \geq MSW - \max_{i \in N} 2\sqrt{c_i \cdot g_i([a_i^*, b_i^*])}.$$
⁹

A General Notion of Inaccuracy

At the highest level, we believe that a function f provides a suitable definition of inaccuracy if it satisfies the following three conditions: for all knowledge intervals $[a, b]$, $[a', b]$, and $[a, a']$ such that $0 \leq a < a' < b$,

1. $f([a, b]) \in (0, 1)$;
2. $f([a, b]) > f([a', b])$; and
3. $f([a, a']) < f([a, b])$.

⁸And continuing to assume that the cost of refinement is proportional to the inaccuracy, that is, $C_i([a, b]) = \frac{c_i}{\delta''}$.

⁹That is, in the above inequality, as in our proof sketch of Theorem 1, $[a_i^*, b_i^*]$ continues to denote the knowledge interval of player i when, executing his regret-minimizing strategy s_i^* , he decides to bid rather than refining his knowledge. In other words, $[a_i^*, b_i^*]$ is the longest knowledge interval $[a, b]$ of i such that

- $[a, b]$ is contained in i 's initial knowledge interval, and
- $2\sqrt{c_i \cdot g_i([a, b])} \geq (b - a)$.

Inaccuracy-Lowering Refinements

In one-dimensional problems, the solution space consists of an interval, and binary search proceeds by shortening the solution space (by a fixed amount).

The problem is also one-dimensional in our setting: indeed, every player seeks to figure out his own true valuation in his current knowledge interval. To mimic binary search as close as possible, we have considered refinements that yield knowledge intervals of shorter lengths. Of course decreasing the length of a knowledge interval by a precise amount also implies decreasing its inaccuracy, but not by a predictably precise amount. Thus, still in the spirit of binary search, one may consider refinements explicitly aimed at decreasing the inaccuracy of a knowledge interval by a given amount. After all, inaccuracy may be a more meaningful dimension to consider in general, not just for meaningfully defining refinement cost.

Accordingly, still within a set-theoretic framework, letting $[a, b]$ be a knowledge interval for a player i and δ be its inaccuracy, let us discuss the following two costly knowledge models that focus on decreasing inaccuracy.

1. *The Accuracy-Bisecting Model.* In this model, at a cost of $C_i(a, b) = \frac{c_i}{\delta}$, $R_i([a, b])$ returns an *arbitrary* knowledge interval of i of inaccuracy $\frac{\delta}{2}$.

Fortunately, essentially the same proof of Theorem 1 shows that, once again, the social welfare guarantee of the second-price mechanism is

$$SW(\text{SP}(s^*)) \geq MSW - 2\sqrt{c_\phi b_\phi}.$$

2. *The Chosen-Accuracy Model.* In this model, for every chosen $\delta' < \delta$, i can, in a single step, refine $[a, b]$ to obtain an arbitrary knowledge interval whose inaccuracy is δ' .

A reasonably natural choice for the cost of such a refinement is to be proportional to the difference of the new and the old inaccuracy: that is, to be equal to

$$c_i \cdot \left(\frac{1}{\delta'} - \frac{1}{\delta} \right).$$

In this model and cost function, essentially the same proof of Theorem 1 (with only a different algebraic manipulation) shows that the social welfare guarantee of the second-price mechanism again is

$$SW(\text{SP}(s^*)) \geq MSW - \sqrt{2c_\phi b_\phi}.$$

General Cost Functions

Abstractly, we believe that every cost function in the chosen-accuracy model should satisfy the following conditions: let $1 \geq \delta > \delta' > \delta'' > 0$, then

1. *Cost is always positive.*

$$C(\delta, \delta') > 0.$$

2. *Cost function must satisfy triangle inequality.*

$$C(\delta, \delta') + C(\delta', \delta'') \geq C(\delta, \delta'').$$

3. *Cost increases with initial inaccuracy.*

$$C(\delta, \delta'') > C(\delta', \delta'').^{10}$$

4. *Cost decreases with desired inaccuracy*

$$C(\delta, \delta'') > C(\delta, \delta').^{11}$$

and, preferably only,

5. *Totally erasing uncertainty has infinite cost.*

$$C(\delta, 0) = \infty.$$

General Inaccuracy and General Cost

Finally, let us consider the chosen-accuracy model, with general inaccuracy and general cost function, and denote by $R_i^{\delta'}$ the non-deterministic function that, given a knowledge interval of i of inaccuracy $\delta > \delta'$, returns an arbitrary knowledge (sub)interval of i of inaccuracy δ' .

In this general model, a player i may have multiple regret-minimizing strategies. Let s_i^* be one such strategy, $[a_i^t, b_i^t]$ the knowledge interval of i at a decision node D^t , δ the inaccuracy of $[a_i^t, b_i^t]$. Then an argument similar to that used in our proof sketch of Theorem 1 shows that

- $s_i^*(D^t) = R_i^{\delta'}$

if there exists $\delta' \in (0, \delta)$ such that

$$b_i^t - a_i^t > \max_{[a,b] \in \mathcal{I}} (2C_i(\delta, \delta') + (b - a))$$

where \mathcal{I} is the set of all subintervals of $[a_i^t, b_i^t]$ which contain θ_i and whose inaccuracy is δ' .

¹⁰Equivalently, for differentiable cost functions,

$$\frac{\partial C(\delta, \delta')}{\partial \delta} > 0$$

¹¹Equivalently, for differentiable cost functions,

$$\frac{\partial C(\delta, \delta')}{\partial \delta'} < 0.$$

- $s_i^*(D^t) = \frac{a_i^t + b_i^t}{2}$ otherwise.

It is then clear that the social-welfare guarantee of the second-price mechanism continues to have the following form

$$SW(\text{SP}(s^*)) \geq MSW - \max_i (b_i^* - a_i^*).$$

Above, once again, $[a_i^*, b_i^*]$ is player i 's final knowledge interval, relative to the strategies s_i^* , in the current model.

Refining Our Knowledge About Our Opponents

In general, a player i may also have some knowledge about the valuations of his opponents. Initially, such knowledge of i may be very coarse, but again i may be able to refine it, although at a cost potentially much higher than that he needs to refine his knowledge about his own valuation.

(For instance, i may hire some employees of his opponents, and analyze the data they report. Alternatively, he may collect financial data about his opponents from the public domain, an operation that may be very expensive, since i does not know exactly where to look for the relevant data).

In principle, for the second-price mechanism, by refining a little his knowledge about the valuations of his opponents, i may save the money necessary to refine a lot his knowledge about himself. For instance, by incurring a modest cost, i may learn that the minimum valuation of one of his opponent is larger than any valuation he may have, in which case, i should not invest a penny in improving his knowledge about his own valuation. Yet, even if i has the ability to refine his knowledge about his opponents at a cost, it can be seen that i 's regret-minimizing strategy never takes advantage of this ability.

We believe, however, that a sufficiently general costly knowledge model should include a player's ability to refine his knowledge about his opponents, and that such an ability will in fact be crucial when properly analyzing some complex mechanisms.

6 Conclusions

In general, knowledge has a cost. Accordingly, we cannot always count on the players having already paid the cost necessary to exactly learn their valuations prior to playing the mechanisms we design.

Including the players' cost of knowledge acquisition in the analysis of a mechanism may thus make our predictions about the outcomes the mechanism produces more realistic and accurate.

Distilling costly knowledge models that are more realistic than the ones considered in this paper may be challenging. And so may be the analysis of new and old mechanisms in such models. We welcome both challenges.

References

- [1] AUMANN, R. J. Utility theory without the completeness axiom. *Econometrica* 30, 3 (July 1962), 445–462.
- [2] BECK, M. J., CHORUS, C. G., ROSE, J. M., AND HENSHER, D. A. Vehicle purchasing behaviour of individuals and groups: regret or reward? *Journal of Transport Economics and Policy (JTEP)* 47, 3 (2013), 475–492.
- [3] BEWLEY, T. F. Knightian decision theory. Part I. *Decisions in Economics and Finance* 25, 2 (2002), 79–110. Earlier version appeared as a discussion paper no. 807 of the Cowles Foundation at Yale University, November 1986.
- [4] BODOH-CREED, A. L. Ambiguous beliefs and mechanism design. *Games and Economic Behavior* 75, 2 (2012), 518–537.
- [5] BOSE, S., OZDENOREN, E., AND PAPE, A. Optimal auctions with ambiguity. *Theoretical Economics* 1, 4 (December 2006), 411–438.
- [6] BOSE, S., AND RENOU, L. Mechanism design with ambiguous communication devices. *Econometrica* 82, 5 (2014), 1853–1872.
- [7] CELIS, L. E., GKLEZAKOS, D. C., AND KARLIN, A. R. On revenue maximization for agents with costly information acquisition. In *Automata, Languages, and Programming*. Springer, 2013, pp. 484–495.
- [8] CELIS, L. E., KARLIN, A. R., LEYTON-BROWN, K., NGUYEN, C. T., AND THOMPSON, D. R. M. Approximately revenue-maximizing auctions for deliberative agents. In *AAAI* (2012).
- [9] CHIESA, A., MICALI, S., AND ZHU, Z. A. Mechanism design with approximate valuations. In *Proceedings of the 3rd Innovations in Theoretical Computer Science conference* (2012), ACM, pp. 34–38.
- [10] CHIESA, A., MICALI, S., AND ZHU, Z. A. Bridging Utility Maximization and Regret Minimization. *ArXiv e-prints* (Mar. 2014).
- [11] CHIESA, A., MICALI, S., AND ZHU, Z. A. Knightian self uncertainty in the vcg mechanism for unrestricted combinatorial auctions. In *Proceedings of the fifteenth ACM conference on Economics and Computation* (2014), ACM, pp. 619–620. For a full version of the paper, see <http://people.csail.mit.edu/zeyuan/paper/2014-EC.pdf>.
- [12] CHIESA, A., MICALI, S., AND ZHU, Z. A. Knightian analysis of the vickrey mechanism. *Will appear in Econometrica* (2015).
- [13] CHORUS, C. Random regret minimization: an overview of model properties and empirical evidence. *Transport Reviews* 32, 1 (2012), 75–92.

- [14] CHORUS, C. G., ARENTZE, T. A., AND TIMMERMANS, H. J. Spatial choice: a matter of utility or regret. *Environment and Planning Part B* 36, 3 (2009), 538–551.
- [15] CORICELLI, G., CRITCHLEY, H. D., JOFFILY, M., O'DOHERTY, J. P., SIRIGU, A., AND DOLAN, R. J. Regret and its avoidance: a neuroimaging study of choice behavior. *Nature neuroscience* 8, 9 (2005), 1255–1262.
- [16] DANAN, E. Randomization vs. selection: How to choose in the absence of preference? *Management Science* 56 (March 2010), 503–518.
- [17] DI TILIO, A., KOS, N., AND MESSNER, M. The design of ambiguous mechanisms. Tech. rep., 2012.
- [18] DUBRA, J., MACCHERONI, F., AND OK, E. A. Expected utility theory without the completeness axiom. *Journal of Economic Theory* 115, 1 (March 2004), 118–133.
- [19] ENGELBRECHT-WIGGANS, R. The effect of regret on optimal bidding in auctions. *Management Science* 35, 6 (1989), 685–692.
- [20] FON, V., AND OTANI, Y. Classical welfare theorems with non-transitive and non-complete preferences. *Journal of Economic Theory* 20, 3 (June 1979), 409–418.
- [21] GALE, D., AND MAS-COLELL, A. An equilibrium existence theorem for a general model without ordered preferences. *Journal of Mathematical Economics* 2, 1 (March 1975), 9–15.
- [22] GILBOA, I., AND SCHMEIDLER, D. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 2 (April 1989), 141–153.
- [23] HALPERN, J. Y., AND PASS, R. Iterated regret minimization: A new solution concept. *Games and Economic Behavior* 74, 1 (2012), 184–207.
- [24] HYAFIL, N., AND BOUTILIER, C. Regret Minimizing Equilibria and Mechanisms for Games with Strict Type Uncertainty. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence* (July 2004), pp. 268–277.
- [25] JOSEPHS, R. A., LARRICK, R. P., STEELE, C. M., AND NISBETT, R. E. Protecting the self from the negative consequences of risky decisions. *Journal of personality and social psychology* 62, 1 (1992), 26.
- [26] KIEKINTVELD, C., ISLAM, T., AND KREINOVICH, V. Security games with interval uncertainty. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems* (2013), International Foundation for Autonomous Agents and Multiagent Systems, pp. 231–238.
- [27] KNIGHT, F. H. *Risk, Uncertainty and Profit*. Houghton Mifflin, 1921.
- [28] LINHART, P. B., AND RADNER, R. Minimax-regret strategies for bargaining over several variables. *Journal of Economic Theory* 48, 1 (1989), 152–178.

- [29] LOPOMO, G., RIGOTTI, L., AND SHANNON, C. Uncertainty in mechanism design, 2009. <http://www.pitt.edu/~luca/Papers/mechanismdesign.pdf>.
- [30] LOPOMO, G., RIGOTTI, L., AND SHANNON, C. Knightian uncertainty and moral hazard. *Journal of Economic Theory* 146, 3 (2011), 1148 – 1172. Incompleteness and Uncertainty in Economics.
- [31] MACCHERONI, F., MARINACCI, M., AND RUSTICHINI, A. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74, 6 (2006), 1447–1498.
- [32] MAS-COLELL, A. An equilibrium existence theorem without complete or transitive preferences. *Journal of Mathematical Economics* 1, 3 (December 1974), 237–246.
- [33] MILNOR, J. W. Games against nature. In *Decision processes*, R. M. Thrall, C. H. Coombs, and R. L. Davis, Eds. John Wiley & Sons, Inc., 1954.
- [34] NASCIMENTO, L. Remarks on the consumer problem under incomplete preferences. *Theory and Decision* 70, 1 (January 2011), 95–110.
- [35] OK, E. A. Utility representation of an incomplete preference relation. *Journal of Economic Theory* 104 (2002), 429–449.
- [36] RENOU, L., AND SCHLAG, K. H. Minimax regret and strategic uncertainty. *Journal of Economic Theory* 145, 1 (Jan. 2010), 264–286.
- [37] RIGOTTI, L., AND SHANNON, C. Uncertainty and risk in financial markets. *Econometrica* 73, 1 (01 2005), 203–243.
- [38] SAVAGE, L. J. The theory of statistical decision. *Journal of the American Statistical association* 46, 253 (1951), 55–67.
- [39] SCHMEIDLER, D. Subjective probability and expected utility without additivity. *Econometrica* 57, 3 (May 1989), 571–87.
- [40] SELTEN, R. Blame avoidance as motivating force in the first price sealed bid private value auction. In *Economics Essays in Honor of Werner Hildenbrand*. Springer, 1989, pp. 333–344.
- [41] SHAFER, W., AND SONNENSCHNEIN, H. Equilibrium in abstract economies without ordered preferences. *Journal of Mathematical Economics* 2, 3 (December 1975), 345–348.
- [42] STOYE, J. Axioms for minimax regret choice correspondences. *Journal of Economic Theory* 146, 6 (2011), 2226–2251.
- [43] THOMPSON, D. R. M., AND LEYTON-BROWN, K. Valuation uncertainty and imperfect introspection in second-price auctions. In *PROCEEDINGS OF THE NATIONAL CONFERENCE ON ARTIFICIAL INTELLIGENCE* (2007), vol. 22, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, p. 148.

- [44] THOMPSON, D. R. M., AND LEYTON-BROWN, K. Dominant-strategy auction design for agents with uncertain, private values. In *AAAI* (2011).
- [45] WALD, A. Statistical decision functions. *The Annals of Mathematical Statistics* 20, 2 (1949), pp. 165–205.