

Crowdsourced Bayesian Auctions

Pablo Azar
MIT, CSAIL
Cambridge, MA 02139, USA
azar@csail.mit.edu

Jing Chen
MIT, CSAIL
Cambridge, MA 02139, USA
jingchen@csail.mit.edu

Silvio Micali
MIT, CSAIL
Cambridge, MA 02139, USA
silvio@csail.mit.edu

August 4, 2011

Abstract

We investigate the problem of optimal mechanism design, where an auctioneer wants to sell a set of goods to buyers, in order to maximize revenue. In a Bayesian setting the buyers' valuations for the goods are drawn from a prior distribution \mathcal{D} , which is often assumed to be known by the seller. In this work, we focus on cases where the seller has no knowledge at all, and “the buyers know each other better than the seller knows them”. In our model, \mathcal{D} is not necessarily common knowledge. Instead, each player individually knows a posterior distribution associated with \mathcal{D} . Since the seller relies on the buyers' knowledge to help him set a price, we call these types of auctions *crowdsourced Bayesian auctions*.

For this crowdsourced Bayesian model and many environments of interest, we show that, *for arbitrary valuation distributions \mathcal{D} (in particular, correlated ones)*, it is possible to design mechanisms matching to a significant extent the performance of the optimal dominant-strategy-truthful mechanisms where the seller knows \mathcal{D} .

To obtain our results, we use two techniques: (1) Proper scoring rules to elicit information from the bidders; and (2) a reverse version of the classical Bulow-Klemperer inequality. The first technique lets us build mechanisms *with a unique equilibrium* that guarantee good revenue, even when the players' second and higher-order beliefs about each other are wrong. The second allows us to *upper bound* the revenue of an optimal mechanism with n players by a $(1 - \frac{1}{n})$ fraction of the revenue of the optimal mechanism with $n - 1$ players. We believe that both techniques are new to Bayesian optimal auctions and of independent interest for future work.

1 Introduction

At the core of Bayesian mechanism design is the assumption that the true valuation profile of the players, θ , has been drawn from a distribution \mathcal{D} over the set of all possible valuation profiles. Different assumptions are then made about who knows how much about \mathcal{D} . Most mechanisms work under the traditional “Centralized Bayesian Assumption:” namely, *The Designer knows \mathcal{D}* .

Notably, this is the assumption under which Myerson [27] constructed a truthful mechanism that maximizes revenue. We are instead interested in designing auctions in Bayesian settings in which *the players know each other better than the designer knows them*.

That is, we wish to investigate settings where the auctioneer is not an expert, but might obtain good revenue by leveraging the collective knowledge of the buyers about each other. This requires an assumption about how much players know about \mathcal{D} . For instance, the notion of Bayesian Nash equilibrium presupposes the “Common-Knowledge Bayesian Assumption:” namely, *\mathcal{D} is common knowledge to the players*.

This assumption is of course quite strong, because common knowledge is a non-trivial requirement. Yet, note that Bayesian Nash equilibrium actually requires more. Namely, it requires that

*each player i has **no information about the valuations of his opponents beyond $\mathcal{D}|\theta_i$** , that is, \mathcal{D} ’s posterior corresponding to his own valuation θ_i . Indeed, if i succeeds in gathering additional information, then he might have plenty of incentive to deviate from his Bayesian Nash equilibrium strategy.*¹

This additional requirement is actually strong, since it is very natural for a player in a high-stakes game to gather as much information about his opponents as he can. For example, a bank participating in a treasury auction has an incentive (and the means) to learn much about its rivals before placing its bid.

Removing this additional assumption is our main goal. To do so, we rely on what we term the “Crowd-sourced Bayesian Assumption.” Informally, we assume that the designer has no information about \mathcal{D} , while each player i , in addition to his valuation, *individually* (as opposed to “commonly”) does not necessarily know just $\mathcal{D}|\theta_i$, but may possibly know a posterior distribution possibly more informative than $\mathcal{D}|\theta_i$.

Under the crowdsourced Bayesian assumption, we prove that it is possible to design auctions, all the way from single-good to truly combinatorial, whose revenue almost coincides with that obtainable under the classical centralized-Bayesian assumption. Following the computer science tradition, we are happy to sacrifice a “bit of optimality to increase robustness”.

1.1 Our work in context

Implementation Theory The idea of using players’ knowledge to guarantee the good performance of a mechanism is not new, and indeed is at the center of Implementation Theory. We refer the interested reader to surveys by Maskin and Sjostrom [26] and by Jackson [25]. Our work can be considered part of the non-parametric Bayesian implementation literature for auctions, where the term “non-parametric” implies that the designer does not know the distribution.² For single-good auction environments and independent valuations, Caillaud and Robert [6] show how to obtain optimal revenue where the players have common knowledge of the valuation distribution and the seller knows nothing.³ For public-good projects and independent valuations, Choi and Kim [13] show how to implement the socially efficient outcome when each player’s distribution is known by someone else. Cremer and Riordan [16] show an efficient, budget-balanced and truthful mechanism for the public goods problem, where one player has knowledge about the others, and there are at least two independent players. Their solution concept of Stackelberg Truthfulness is very similar to our own solution concept of two-step dominant strategy truthfulness. To the best of our knowledge, our paper gives the first

¹The mechanism of Cremer and McLean [15] is an example of additional information changing the incentives of the players. This mechanism is further described in section 7.

²We choose to use the term “crowdsourced” because of our use of proper scoring rules, a technique used in other crowdsourcing applications such as prediction markets.

³A more detailed comparison with Caillaud and Robert’s result is given in section 5.

implementation results with *arbitrarily correlated valuations*, and where players may collect extra information on each other beyond what is given by the prior distribution.

Detail-Free mechanisms Mechanisms where the designer does not know the distribution are called *detail-free*, and have been a goal of the mechanism design literature since at least the time of Wilson [33]. Explicit detail-free Bayesian mechanisms for single-unit auctions have been proposed by Segal [32], Baliga and Vohra [2] and Dhangwatnotai, Roughgarden, and Yan [17]. Under the assumption that each player knows his own distribution, the mechanism we present in section 4 obtains more revenue than these existing mechanisms. When valuations are not drawn from distributions, auction mechanisms have been proposed by Goldberg, Hartline, Karlin, Saks and Wright [21], Goldberg and Hartline [20]. Our mechanisms are incomparable to theirs: we obtain more revenue, but only when we make the assumption that a distribution exists.

Robust Mechanism Design This paper focuses on removing the common knowledge assumption, and thus belongs to the Robust Mechanism Design literature initiated by Bergemann and Morris [3]. One of their results shows that many social choice functions which are implementable in a bayesian setting are also dominant strategy implementable (removing the need for common knowledge of the distribution). This result does not hold for environments where monetary payments are restricted. These environments include optimal auctions, where we have a constraint that the auctioneer should not pay out more money than it receives from the players. Thus, we need alternative ways of removing the common knowledge assumption. Chung and Ely [14] show that when the auctioneer knows the distribution but players can have arbitrary beliefs, an auctioneer who wants to maximize his worst-case gain should use a dominant strategy mechanism. This is a justification for using dominant strategy mechanisms to maximize profits, an approach which we follow in this paper.

Our Mechanisms As Black-Box Results And Their Relation To Classical Bayesian Mechanisms Our mechanism from section 3 can be viewed as a “compiler” that takes as input a classical bayesian mechanism and outputs a crowdsourced one. In section 7 we discuss under what assumptions we can leverage the existing classical mechanisms presented by Hartline and Roughgarden [23], Cremer and McLean [15], Papadimitriou and Pierrakos [30], Chawla, Hartline, Malec and Sivan [7] and Dobzinski, Fu and Kleinberg [18], to create new crowdsourced mechanisms.

Other Non-Bayesian Approaches Chen and Micali, in [10] for combinatorial auctions and in [11] for single-good auctions, show how to obtain new and reasonable revenue benchmarks without *any* Bayesian assumption, including ours.⁴ (Of course, although this has generated some confusion in the past, in settings of complete information, one can achieve virtually maximum revenue.⁵)

1.2 Paper outline

Section 2 outlines our model and assumptions. Section 3 contains the proof of our main theorem, which holds for very general environments. Section 4 shows an alternative mechanism for single unit auctions and explain its contributions with respect to existing mechanisms. Section 5 proves that, for 2 players, our mechanisms are virtually optimal with respect to our new assumptions on knowledge. Section 6 provides a reverse version of the Bulow Klemperer inequality that is derived from the proof of our main theorem. Section 7 discusses how our results may be combined with some existing mechanisms and gives final remarks.

⁴In essence, the revenue benchmark of [10] consists of one half of “the maximum revenue that a player is sure he could generate by selling the goods to his opponents via take-it-or-leave-it offers. The revenue benchmark in [11] essentially consists of the second-highest value that a player believes that some player (including himself) is willing to pay for the good. Notice that such benchmark is always greater or equal to the classical second-highest valuation, and sometimes much higher.

⁵For instance, Chen, Hassidim and Micali [12] do precisely that for general settings of quasi-linear utility, which include combinatorial auctions as a special case.




Figure 1: This example shows the knowledge partition of a generic buyer. The buyer’s knowledge partition is $\mathcal{P} = (P_{i,1}, P_{i,2}, P_{i,3})$. He does not know the realized valuation θ exactly, but knows that it lies in the set $S^* = P_{i,1}$ highlighted in red.

2 Model and Assumptions

2.1 Players, Goods and Utilities

There is a finite set of players $N = \{1, \dots, n\}$ and a finite set X of goods. The set of all feasible assignments is denoted by $\mathbb{A} = \{(A_1, \dots, A_n) : A_i \subset X\}$. The set \mathbb{A} satisfies a *downward-closure condition*:

for all $(A_1, \dots, A_n) \in \mathbb{A}$, for all $i \in N$, we have $(A_1, \dots, A_{i-1}, \emptyset, A_{i+1}, \dots, A_n) \in \mathbb{A}$.

Players have quasilinear utility. Each player i has a valuation function $\theta_i : 2^X \rightarrow \mathbb{Z}^+$ that satisfies $\theta_i(\emptyset) = 0$. Player i ’s utility of receiving a set A_i and paying a price p_i is $u_i(A_i, p_i) = \theta_i(A_i) - p_i$. A negative price indicates a payment from the mechanism to the player.

Unless otherwise specified, we denote by $\theta = (\theta_1, \dots, \theta_n)$ a profile of valuations, and by Θ the set of all possible valuation profiles. We use the notation $\Delta(\Theta)$ to refer to the set of all probability distributions over Θ .

These environments are very general. In particular, they include all types of auctions, from single-good to combinatorial. An example environment which is *not* downward-closed is the public goods project, where no agent can be prevented from using the public good.

2.2 Beliefs and the Crowdsourced-Bayesian Assumption

Our knowledge model is similar to that of the Bayesian Implementation literature [24].⁶ There exists a distribution $\mathcal{D} \in \Delta(\Theta)$ from which valuation profiles are drawn. Players’ knowledge is represented by partitions $\mathcal{P} = (S_1, \dots, S_k)$ of the set Θ .

Definition 1 (Knowledge with respect to a partition). *Let θ be the true type profile, \mathcal{P} a partition of Θ and S^* the unique set in \mathcal{P} such that $\theta \in S^*$. A player whose knowledge is represented by a partition \mathcal{P} knows*

⁶There are some differences between our model and that of [24]. In their model, each agent i has an individual prior q_i and the vector of priors (q_1, \dots, q_n) is common knowledge to all agents. In our model, there is only one true prior \mathcal{D} , which may not be known to anyone. Each agent knows their posterior \mathcal{D}_i , but the vector of posteriors is not necessarily common knowledge.

1. The set S^* containing the true type profile.
2. The posterior distribution $\mathcal{D}|S^*$ followed by θ given that it is in the set S^* .

Players know their own valuation. More formally,

Definition 2 (*i*-consistency). A partition \mathcal{P} of Θ is *i*-consistent if, for any set S in \mathcal{P} and any two valuation profiles $\theta, \theta' \in S$, we have $\theta_i = \theta'_i$.

Having established this notation, we state our main assumption on the players' knowledge.

Crowdsourced-Bayesian Assumption: Each player *i* knows a partition \mathcal{P}_i of Θ such that: (1) \mathcal{P}_i is *i*-consistent for each *i*, and (2) when θ is randomly selected according to \mathcal{D} , player *i*'s beliefs are given by $\mathcal{D}_i \triangleq \mathcal{D}|S_i^*$, where S_i^* is the unique set in \mathcal{P}_i containing θ .

Remark The Crowdsourced-Bayesian assumption

- (a) *makes no restriction on \mathcal{D}* : indeed, \mathcal{D} may be any member of $\Delta(\Theta)$, the set of all distributions over Θ (in particular, the players' valuations can be arbitrarily correlated);
- (b) *does not assume any knowledge on the part of the designer*: indeed the designer may know (arbitrarily little or even) nothing about \mathcal{D} and;
- (c) *makes no assumptions about players' higher-order beliefs*: all our results hold even if players' second and higher-order beliefs are wrong.

2.3 Mechanisms and Solution Concepts

A player's type is a pair $(\theta_i, \mathcal{D}_i)$ that specifies his valuation and his beliefs about others' valuations. A direct revelation mechanism is one where players' action space coincides with their typespace.

Definition 3 (Crowdsourced-Bayesian Mechanism). A crowdsourced-Bayesian mechanism is a randomized function $M : \Theta \times \Delta(\Theta)^n \times \{0, 1\}^* \rightarrow \mathbb{A} \times \mathbb{R}^n$ that specifies for any valuation profile θ , belief profile $(\mathcal{D}_1, \dots, \mathcal{D}_n)$ and sequence of coin tosses $r \in \{0, 1\}^*$, an outcome $M(\theta, (\mathcal{D}_1, \dots, \mathcal{D}_n), r) \in \mathbb{A} \times \mathbb{R}^n$.

In this paper, crowdsourced-Bayesian mechanisms are often compared to mechanisms where players only reveal their valuation function θ_i . The corresponding definition is given below.

Definition 4 (Classical Direct Revelation Mechanism). A classical direct revelation mechanism is a randomized function $M : \Theta \times \{0, 1\}^* \rightarrow \mathbb{A} \times \mathbb{R}^n$ that specifies for any valuation profile θ and any sequence of coin tosses $r \in \{0, 1\}^*$ an outcome $M(\theta, r) \in \mathbb{A} \times \mathbb{R}^n$.

Among classical direct revelation mechanisms, we consider those that are dominant strategy truthful.

Definition 5. We say that a Classical Direct Revelation Mechanism M is **dominant strategy truthful** if, whenever θ is the true valuation profile, for all players *i*, for all $\tau_i \neq \theta_i$, for all opponent messages τ_{-i} , for all possible coin tosses r of the mechanism M , we have

$$u_i(M(\theta_i, \tau_{-i}, r)) \geq u_i(M(\tau_i, \tau_{-i}, r)).$$

We say M is strictly dominant strategy truthful if for all τ_i, θ_i there always exists τ_{-i}, r such that the above inequality is strict.

For crowdsourced-Bayesian mechanisms, we introduce a solution concept called *Two-Step Dominant Strategy Truthfulness*.

Definition 6. A Direct Revelation Mechanism is *two-step (strictly) dominant strategy truthful* if, whenever θ is the true valuation profile and $(\mathcal{D}_1, \dots, \mathcal{D}_n)$ is the true vector of beliefs

1. It is a (strictly) dominant strategy for each player to reveal their true valuation.
2. Given that all players are announcing their true valuation, it is a (strictly) dominant strategy for each player to reveal their true beliefs.

An important feature is that mechanisms which are *strictly* two-step dominant strategy truthful have a *unique* equilibrium where players reveal their true valuations and their true beliefs about others. Furthermore, truthfulness is the only strategy that survives two steps of iterated elimination of strictly dominated strategies.

Unless otherwise specified, all mechanisms satisfy the *Individual Rationality Condition*. For all players i , for all true valuations θ_i , for any valuation profile τ_{-i} of the other players' actions, for any coin tosses r of the mechanism \mathcal{M} , $u_i(M(\theta_i, \tau_{-i})) \geq 0$.

2.4 Revenue and Implementation

Let M be a crowdsourced-Bayesian mechanism, let E be an assignment environment and let \mathcal{D} be a distribution. Denote by $\text{REV}(M, E, \mathcal{D})$ the expected revenue generated by M on environment E when valuations are drawn from \mathcal{D} and all players are truthful. We denote by $\text{OPT}(E, \mathcal{D})$ the expected revenue of the optimal classical strictly DST Bayesian mechanism when valuations are drawn from \mathcal{D} and players are truthful. (Each expectation is taken over all possible random choices, that is, over \mathcal{D} and the coin tosses of the corresponding mechanism, if probabilistic.)

A desirable property in a mechanism is that it achieve revenue that is close to optimal. Fix an environment E . For $c < 1$ and $\epsilon > 0$, a mechanism achieves (c, ϵ) – *approximate* revenue if, for all distributions \mathcal{D} it is true that

$$\text{Rev}(M, E, \mathcal{D}) \geq c \cdot \text{OPT}(E, \mathcal{D}) - \epsilon.$$

Definition 7. A crowdsourced-Bayesian mechanism M implements (c, ϵ) – *approximate revenue maximization* if

1. the only strategy that survives iterated elimination of strictly dominated strategies is truthfulness about valuations and beliefs
2. the mechanism (c, ϵ) – *approximates* revenue.

A helpful lemma for implementation tells us that we can convert a weakly dominant strategy truthful auction mechanism into a *strictly* dominant strategy truthful mechanism with an arbitrarily small loss in expected revenue

Lemma 1. Let E be a downward-closed environment and let M be a classical direct revelation dominant strategy truthful mechanism. Then for every $\delta > 0$ there exists a classical direct revelation mechanism M' which is strictly DST and which satisfies

$$\text{Rev}(M', E, \mathcal{D}) \geq (1 - \delta)\text{Rev}(M, E, \mathcal{D})$$

for all distributions \mathcal{D} .

The proof of this lemma is given in an appendix. In the following sections, all DST mechanisms will be assumed to be *strictly* DST unless stated otherwise. Because of the above lemma, the loss of revenue from this assumption can be made arbitrarily small.

2.5 Proper Scoring Rules

Let Ω be a state space and let $\Delta(\Omega)$ be the set of probability distributions over Ω . A scoring rule S is a function, $S : \Omega \times \Delta(\Omega) \rightarrow \mathbb{R}$. A scoring rule S can be used to reward individuals reporting their knowledge (or beliefs) about the world: in particular, giving an individual reporting a probability distribution $\mathcal{D} \in \Delta(\Omega)$ a reward equal to $S(\omega, \mathcal{D})$ whenever the realized state is $\omega \in \Omega$. A scoring rule S is *proper* if an individual maximizes his expected reward by announcing his true knowledge about the world. That is, for any two different probability distributions $\mathcal{D}, \mathcal{P} \in \Delta(\Omega)$,

$$\mathbb{E}_{\omega \leftarrow \mathcal{D}}[S(\omega, \mathcal{D})] \geq \mathbb{E}_{\omega \leftarrow \mathcal{D}}[S(\omega, \mathcal{P})].$$

A recent paper by Gneiting and Raftery is a good survey of proper scoring rules [19]. A scoring rule S is *strictly proper* if the above inequality is strict.

Brier’s scoring rule, denoted by BSR, is defined as follows. For each $s \in \Omega$, denote by $\mathcal{D}(s)$ be the probability assigned to s by \mathcal{D} . Let $\delta_{\omega,s}$ be the indicator function, that is $\delta_{\omega,s} = 1$ if $s = \omega$ and 0 otherwise. The Brier’s Scoring Rule is given by

$$\begin{aligned} \text{BSR}(\omega, \mathcal{D}) &= - \left(\sum_{s \in \Omega} (\delta_{\omega,s} - \mathcal{D}(s))^2 \right) = \\ &= 2\mathcal{D}(\omega) - \|\mathcal{D}\|_2^2 - 1. \end{aligned}$$

Note that Brier’s scoring rule is always *bounded*: indeed, $\text{BSR}(\omega, \mathcal{D}) \in [-2, 0]$ for all ω and \mathcal{D} . (In contrast, Good’s [22] more popular *logarithmic scoring rule* $\text{LSR}(\omega, \mathcal{D}) = \log(\mathcal{D}(\omega))$ is unbounded.)

3 Main Theorem

We show that crowdsourced-Bayesian Mechanisms are in principle as powerful as centralized-Bayesian ones.

Theorem 1. *For any $\epsilon > 0$, for all n -player assignment environments E , there exists a mechanism \mathbb{M} that implements $(1 - \frac{1}{n}, \epsilon)$ -approximate revenue.*

We construct a specific mechanism \mathbb{M} and show it satisfies the requirements of the theorem via a sequence of lemmas. \mathbb{M} first obtains from a player a distribution over the valuations of the other players, and then runs the optimal classical strictly DST mechanism for this $(n - 1)$ -player distribution.

In our description below, numbered steps are taken by the players, and steps marked by letters are steps taken by the mechanism/auctioneer.

Mechanism \mathbb{M}

- a. Choose a player i uniformly at random from $\{1, \dots, n\}$.

Comment. Player i will receive the empty allocation and get a price of zero, but he will be rewarded according to his knowledge. Choosing i deterministically does not affect incentives, but might reduce revenue. We do not know how to make \mathbb{M} two-step DST without “removing one player”.

1. Player i announces a distribution \mathcal{F} over Θ_{-i} .

Comment. Allegedly, \mathcal{F} represents his true knowledge \mathcal{D}_i restricted to the domain Θ_{-i} .

2. Each player $j \neq i$ announces a valuation function $v_j \in \Theta_j$.

Comment. Allegedly, v_j is j ’s true valuation.

- b. Letting \mathcal{M} be the optimal strictly-DST centralized mechanism for the $(n - 1)$ -player distribution \mathcal{F} , run $\mathcal{M}(v_{-i}, \mathcal{F})$ so as to obtain an allocation $A_{-i} = (A_1, A_2, \dots, A_{i-1}, A_{i+1}, \dots, A_n)$, and a price vector $P_{-i} = (P_1, P_2, \dots, P_{i-1}, P_{i+1}, \dots, P_n)$.

Comment. The allocations and prices of the players in $-i$ are determined by \mathcal{M} .

c. Set $A_i = \emptyset$, reward i by setting $P_i = -\frac{\epsilon}{2}(2 + \text{BSR}(v_{-i}, \mathcal{F}))$, and output outcome (A, P) .

Comment. Although player i gets no allocation, he gets a reward according to Brier's scoring rule. The negative price indicates that the mechanism transfers money to player i . Since we insist that our players do not incur negative utilities *ex post*, the boundedness of Brier's scoring rule is crucial here.

Lemma 2. *The mechanism \mathbb{M} is two-step strictly DST.*

Proof. Let i be the player chosen in step a of the mechanism. First we prove that each player $j \neq i$ maximizes his utility by announcing his true valuation in step 2, regardless of the other players' strategies. Indeed, player j gets allocation A_j and price P_j , which are determined by running an optimal strictly DST direct revelation mechanism $\mathcal{M}(v_{-i}, \mathcal{F})$. Since \mathcal{M} is strictly DST, player j maximizes his utility $\theta_j(A_j) - P_j$ by announcing $v_j = \theta_j$, no matter what $v_{-\{i,j\}}$ and \mathcal{F} are. (Notice that it does not even matter whether \mathcal{F} is the true distribution from which θ_{-i} is drawn, because for \mathcal{M} to be classical DST, the players must maximize their utilities by being truthful regardless of the distribution.) Since player j has the same utility in \mathcal{M} and in \mathbb{M} , truthfulness in \mathcal{M} implies truthfulness in \mathbb{M} (about his valuation).

Now we prove that, given that all players $j \neq i$ reveal their true utility types, player i maximizes his expected utility by announcing his true knowledge \mathcal{D}_i . Indeed, player i 's expected utility from announcing distribution \mathcal{F} is $\mathbb{E}_{v_{-i} \leftarrow \mathcal{D}_i}[\frac{\epsilon}{2}(2 + \text{BSR}(v_{-i}, \mathcal{F}))]$. Since Brier's scoring rule is strictly proper, this expectation is maximized if and only if $\mathcal{F} = \mathcal{D}_i$. ■

Remark Note that we need a version of iterated dominance as the solution concept for this mechanism. Let i be the player selected in the first stage. If all other players j are "irrational" and report a fake value (say $v_j = 0$), then it is not dominant for player i to report his true beliefs. Given that all other players report $v_j = 0$, i 's best response is to give a distribution that assigns probability 1 to the $v_{-i} = \vec{0}$.

Lemma 3. *For all assignment environments E and distributions \mathcal{D} ,*

$$\text{REV}(\mathbb{M}, E, \mathcal{D}) \geq (1 - \frac{1}{n})\text{OPT}(E, \mathcal{D}) - \epsilon.$$

Notation Since a direct (possibly randomized) centralized-Bayesian assignment mechanism M may be conceptualized as receiving (the description of) a distribution \mathcal{D} as a separate input, we use a slightly different notation to denote M 's revenue than the one used for decentralized-Bayesian mechanisms. Namely, if M always flips ℓ coins, then letting v be a valuation profile and $r \in \{0, 1\}^\ell$, $M(v, \mathcal{D}, r)$ denotes the unique outcome (A, P) computed by M on inputs v and \mathcal{D} when using coin tosses r ; $\text{REV}(M(v, \mathcal{D}, r)) = \sum_{i=1}^n P_i$; $\text{REV}(M(v, \mathcal{D})) = \mathbb{E}_{r \leftarrow \{0,1\}^\ell} \text{REV}(M(v, \mathcal{D}, r))$; and $\text{REV}(M(\mathcal{D})) = \mathbb{E}_{v \leftarrow \mathcal{D}} \mathbb{E}_{r \leftarrow \{0,1\}^\ell} \text{REV}(M(v, \mathcal{D}, r))$.

Proof.

Consider the following mental experiment. Fix a valuation profile v drawn from the distribution \mathcal{D} , and (if the mechanism is randomized), fix a sequence r of coin flips. Let $\mathcal{M}(v, \mathcal{D}, r)$ be an execution of the optimal mechanism that produces some outcome (A, P) . For any player i , let (A^i, P^i) be an outcome such that $A_i^i = \emptyset$, $P_i^i = 0$, $A_{-i}^i = A_{-i}$, and $P_{-i}^i = P_{-i}$. That is, player i gets the empty allocation and pays zero, and all the other players get the same price and allocation as in (A, P) . Define the mechanism \mathcal{M}_i which, given a valuation profile v , coin flips r , and distribution \mathcal{D} , runs $\mathcal{M}(v, \mathcal{D}, r)$ but gives player i the empty allocation and charges him zero. That is, \mathcal{M}_i produces the outcome $\mathcal{M}_i(v, \mathcal{D}, r) = (A^i, P^i)$. Note that the average revenue of the \mathcal{M}_i mechanisms (taken over the choice of i) is

$$\frac{1}{n} \sum_{i=1}^n \text{REV}(\mathcal{M}_i(v, \mathcal{D}, r)) = \frac{1}{n} \sum_{i=1}^n \sum_{j \neq i} P_j =$$

$$= \frac{n-1}{n} \sum_{j=1}^n P_j = \frac{n-1}{n} \text{REV}(\mathcal{M}(v, \mathcal{D}, r)).$$

Taking expectation over all $v \leftarrow \mathcal{D}$ and all sequences of coin flips r , we get that

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \text{REV}(\mathcal{M}_i(\mathcal{D})) &= \frac{n-1}{n} \text{REV}(\mathcal{M}(\mathcal{D})) = \\ &= \frac{n-1}{n} \text{OPT}(E, \mathcal{D}). \end{aligned}$$

Since the mechanism \mathcal{M} is classical-DST, and each \mathcal{M}_i does not change the outcomes for any player $j \neq i$, we must have that \mathcal{M}_i is dominant-strategy truthful for all players $j \neq i$. Furthermore, it has the property that it produces an allocation A^i such that $A^i_i = \emptyset$, an allocation that does not serve player i .

Now let's restrict ourselves to outcomes that do not serve player i . Since the mechanism \mathcal{M} is optimal, the maximum expected revenue of classical-DST mechanisms not serving player i is obtained by running \mathcal{M} on inputs v_{-i}, \mathcal{D}_{-i} , where v_{-i} is the profile of valuations of all players except i , and \mathcal{D}_{-i} is the distribution of v_{-i} induced by restricting \mathcal{D} to Θ_{-i} . In particular, this generates more expected revenue than running \mathcal{M}_i , which is also a classical-DST mechanism that does not serve player i . We can conclude that

$$\text{REV}(\mathcal{M}(\mathcal{D}_{-i})) \geq \text{REV}(\mathcal{M}_i(\mathcal{D})).$$

To continue, we observe that the expected revenue of optimal mechanisms increases with the precision of their distributions. That is, modeling them as a single universal mechanism \mathcal{M} receiving the relevant distribution as a separate input; letting \mathcal{P} be a partition of the valuation space Θ into events; and assuming that, when the true valuation profile θ is randomly selected from \mathcal{D} , \mathcal{M} 's separate input is $\mathcal{D}|S$ —where S is the unique set in \mathcal{P} containing the realized θ —we have

$$\sum_{S \in \mathcal{P}} \Pr_{\theta \leftarrow \mathcal{D}}(\theta \in S) \cdot \text{OPT}(E, \mathcal{D}|S) \geq \text{OPT}(E, \mathcal{D}).$$

Accordingly, we must have that the revenue of \mathcal{M} does not decrease when given player i 's true knowledge \mathcal{D}_i , which is $\mathcal{D}|S_i$ whenever $\theta \in S_i \in \mathcal{P}_i$. That is,

$$\begin{aligned} \text{REV}(\mathcal{M}, -i) &\triangleq \sum_{S_i \in \mathcal{P}_i} \Pr_{\theta \leftarrow \mathcal{D}}(\theta \in S_i) \cdot \text{REV}(\mathcal{M}((\mathcal{D}|S_i)_{-i})) \\ &\geq \text{REV}(\mathcal{M}(\mathcal{D}_{-i})), \end{aligned}$$

where $(\mathcal{D}|S_i)_{-i}$ is the distribution of θ_{-i} induced by restricting $\mathcal{D}|S_i$ to Θ_{-i} .

The mechanism \mathbb{M} generates revenue by choosing a player i at random and running $\mathcal{M}(\theta_{-i}, \mathcal{D}_i)$ —when the players are truthful, $v_{-i} = \theta_{-i}$ and $\mathcal{F} = \mathcal{D}_i$. Thus the expected revenue (ignoring the reward in step c) of \mathbb{M} is

$$\frac{1}{n} \sum_{i=1}^n \text{REV}(\mathcal{M}, -i),$$

which, based on previous inequalities, is greater than or equal to

$$\frac{1}{n} \sum_{i=1}^n \text{REV}(\mathcal{M}(\mathcal{D}_{-i})) \geq \frac{1}{n} \sum_{i=1}^n \text{REV}(\mathcal{M}_i(\mathcal{D})) = \frac{n-1}{n} \text{OPT}(E, \mathcal{D}).$$

Taking into account the reward given in step c , which is at most ϵ , we can conclude that the expected revenue of \mathbb{M} is $\geq \frac{n-1}{n} \text{OPT}(E, \mathcal{D}) - \epsilon$, as desired. \blacksquare

Theorem 1 follows directly from Lemmas 2 and 3. *Q.E.D.*

Computational Remarks In our proof of Theorem 1, we use the optimality of mechanism \mathcal{M} only to derive two inequalities: namely inequalities $\text{REV}(\mathcal{M}, -i) \geq \text{REV}(\mathcal{M}(\mathcal{D}_{-i}))$ and $\text{REV}(\mathcal{M}(\mathcal{D}_{-i})) \geq \text{REV}(\mathcal{M}_i(\mathcal{D}))$. The first can be interpreted as the following condition: “the more precise the valuation-distribution known to \mathcal{M} , the better \mathcal{M} ’s revenue performance.” The second inequality can be interpreted as the following condition: “ \mathcal{M} generates more revenue when running with all players but i , than when running on all players and then throwing away i ’s payment.” Thus, if \mathcal{M} satisfies the above two conditions and is approximately optimal, with approximation ratio β , then \mathbb{M} will also be approximately optimal, with approximation ratio $\frac{n-1}{n}\beta$.

Our mechanism \mathbb{M} needs to evaluate Brier’s scoring rule. For this, we need to have access to $\mathcal{F}(v_{-i})$ and to compute the norm $\|\mathcal{F}\|_2^2$, which might be computationally expensive when \mathcal{F} is correlated and/or the range of the players’ valuations is big. By contrast, a variant of Good’s logarithmic scoring rule, $\text{LSR}_{a,b}(v_{-i}, \mathcal{F}) = a + b \log(\mathcal{F}(v_{-i}))$ with $a, b > 0$, is still strictly proper, but only requires access to $\mathcal{F}(v_{-i})$. However, although always providing positive utility in expectation for suitable values of a and b , it has the disadvantage that, for some distributions \mathcal{F} and some valuation subprofiles v_{-i} , it can give arbitrarily negative rewards, giving player i negative utility. If positive utility in expectation is good enough, then $\text{LSR}_{a,b}$ is a perfectly suitable scoring rule.

4 The single good auction setting

In this section, we introduce a mechanism \mathbb{L} for single good auctions. The main features of this mechanism are that it aggregates the knowledge of many players and under certain conditions collects more revenue than \mathbb{M} .

We first introduce two tools that are used by \mathbb{L} : a knowledge aggregator **AGG** and the 1 – *lookahead* mechanism \mathcal{L} .

4.1 Knowledge Aggregator

We define the *knowledge aggregator*, **AGG**, to be the function mapping the identity of a player i and a distribution subprofile \mathcal{DK}_{-i} to another distribution as follows.

AGG(i, \mathcal{DK}_{-i})

0. For each $j \neq i$, set S_j to be the support of \mathcal{DK}_j .
1. Set $S = \cap_{j \neq i} S_j$ and $j' = \min\{j : j \neq i\}$.
2. If $S = \emptyset$, then output $\mathcal{DK}_{j'}$.
3. Else, set \mathcal{DK}' to be $\mathcal{DK}_{j'}|S$, and output \mathcal{DK}' .

In essence, **AGG** interprets each \mathcal{DK}_j as the distributional knowledge of player j in a crowdsourced-Bayesian model with distribution \mathcal{D} , and aggregates the individual knowledge of the players in $-i$ so as to reconstruct a refined distribution for i ’s valuation. Notice that in **AGG**, if the knowledge of players in $-i$ disagrees with each other in an apparent way (that is, when $S = \emptyset$), then the function’s output can be arbitrary. Also notice that the choice of the player j' can be arbitrary, because when the input \mathcal{DK} is the “true knowledge” of the players, any choice of j' will lead to the same output \mathcal{DK}' . Here by “true knowledge” we mean that, each \mathcal{DK}_j is obtained from the true distribution \mathcal{D} by first conditioning on some event that is consistent with the true valuation profile, and then conditioning on the knowledge of players in $-i$.

4.2 The 1 – *lookahead* mechanism \mathcal{L}

The 1 – *lookahead* mechanism [31] is a dominant strategy truthful mechanism for single unit auctions that allocates the good to the player with the highest valuation, charging him the highest possible price that

does not depend on his own announcement. The auctioneer gets as input a profile of announced valuations $v = (v_1, \dots, v_n)$ and a distribution of valuations \mathcal{D} that is assumed to be known by the auctioneer.

$\mathcal{L}(v; \mathcal{D})$

1. Each player i announces v_i .
 - a. Let $i^* = \operatorname{argmax}_i v_i$ be a player with maximum announced valuation. If multiple such players exist, choose any one of them.
 - b. Let $p^* = \operatorname{argmax}_p p \cdot \operatorname{Prob}_{\theta \leftarrow \mathcal{D}}(\theta_{i^*} \geq p | \theta_{-i^*} = v_{-i^*} \text{ and } v_{i^*} \geq v_j \forall j)$
 - c. If $v_{i^*} \geq p^*$ player i^* , assign the good to player i by setting $A_{i^*} = 1, P_{i^*} = p^*$. Otherwise, $A_{i^*} = \emptyset, P_{i^*} = 0$.
 - d. For all $j \neq i^*$ set $A_j = \emptyset, P_j = 0$. Output (A, P) .

Ronen [31] showed that the 1-lookahead mechanism is dominant strategy truthful and that $\operatorname{Rev}(\mathcal{L}, \mathcal{D}) \geq \frac{1}{2} \operatorname{OPT}(\mathcal{D})$, even when \mathcal{D} is correlated and the optimal mechanism is not known.

As mentioned in section 2 and proved in appendix A, the mechanism \mathcal{L} can be transformed into a strictly dominant strategy mechanism \mathcal{L}_δ that achieves $(1 - \delta)$ of the revenue obtained by \mathcal{L} , for arbitrarily small δ .

4.3 Our mechanism \mathbb{L}

We give a crowdsourced mechanism \mathbb{L} that is a modification of the 1-lookahead mechanism \mathcal{L} .

Mechanism \mathbb{L}

1. Each player i secretly announces to the auctioneer an alleged valuation v_i .
 - a. Let $i^* = \operatorname{argmax}_i v_i$ be a player with maximum announced valuation. If multiple such players exist, choose any one of them.
2. Each player $j \neq i^*$ announces a distribution $\mathcal{F}_j \in \Delta(\Theta)$.
 - b. Compute $\mathcal{F} = \operatorname{AGG}(i^*, \mathcal{F}_{-i^*})$.
 - c. Set $(A, P) = \mathcal{L}_\delta(v, \mathcal{F})$ to assign the good and determine prices. Break ties in favor of player i^* chosen in step a.
 - e. For each $j \neq i^*$, update $P_j = P_j - \frac{\epsilon}{2n}(2 + \operatorname{BSR}(v_i, \mathcal{F}_j))$. Also update $P_{i^*} = P_{i^*} - \epsilon$.

We can prove a theorem analogous to 1 for mechanism \mathbb{L} .

Theorem 2. *The mechanism \mathbb{L} implements $(\frac{1-\delta}{2}, \epsilon)$ -revenue approximation.*

We prove this theorem via a sequence of lemmas.

Lemma 4. *\mathbb{L} is two-step dominant strategy truthful.*

Proof. This is analogous to the proof of lemma 2. It is strictly truthful for players to announce their valuation truthfully because players are rewarded according to \mathcal{L}_δ , which is a strictly DST mechanism.

Given that player i^* announces his valuation truthfully, all the other players want to announce \mathcal{F}_j truthfully (that is, announcing a distribution \mathcal{F}_j consistent with their true beliefs \mathcal{D}_j). This is because players are rewarded using a scoring rule for announcing their knowledge. It is important to highlight that the players are only rewarded based on v_{i^*} , which is outside their control. *Q.E.D.*

Lemma 5. *For any distribution \mathcal{D} , $\operatorname{Rev}(\mathbb{L}, \mathcal{D}) \geq \frac{1-\delta}{2} \operatorname{OPT}(\mathcal{D}) - \epsilon$.*

Proof. As in lemma 3 , we need to show that revenue increases monotonically with knowledge. Recall that each player j has a partition \mathcal{P}_j over the set Θ of type profiles. Denote by $\mathcal{P} = \bigwedge_{j \in -i^*} \mathcal{P}_j$ the meet⁷ of the partitions of all players except i^* .

We can write the revenue of mechanism \mathbb{L} as

$$Rev(\mathbb{L}, \mathcal{D}) = \sum_{S \in \mathcal{P}} \Pr_{\theta \leftarrow \mathcal{D}}(\theta \in S) \cdot Rev(\mathbb{L}, \mathcal{D}|S).$$

Now note that, for any $S \in \mathcal{P}$, we have $Rev(\mathbb{L}, \mathcal{D}|S) \geq Rev(\mathcal{L}_\delta, \mathcal{D}|S) - \epsilon$. This is because \mathbb{L} charges exactly the same prices as \mathcal{L}_δ , except that it gives n players a reward no larger than $\frac{\epsilon}{n}$. Furthermore, $Rev(\mathcal{L}_\delta, \mathcal{D}|S) > (1 - \delta)Rev(\mathcal{L}, \mathcal{D}|S)$. Thus, we have

$$Rev(\mathbb{L}, \mathcal{D}) \geq \sum_{S \in \mathcal{P}} \Pr_{\theta \leftarrow \mathcal{D}}(\theta \in S) \cdot [(1 - \delta)Rev(\mathcal{L}, \mathcal{D}|S) - \epsilon].$$

The revenue $Rev(\mathcal{L}, \mathcal{D}|S)$ can be written as

$$Rev(\mathcal{L}, \mathcal{D}|S) = \max_p p \cdot Prob_{\theta \leftarrow \mathcal{D}|S}(\theta_{i^*} \geq p | \theta_{-i^*} = v_{-i^*} \text{ and } v_i^* \geq v_j \forall j)$$

If the mechanism \mathcal{L} only knows the distribution \mathcal{D} (instead of the posterior $\mathcal{D}|S$), it charges the price $\hat{p} = \max_p p \cdot Prob_{\theta \leftarrow \mathcal{D}}(\theta_{i^*} \geq p | \theta_{-i^*} = v_{-i^*} \text{ and } v_i^* \geq v_j \forall j)$. The revenue $Rev(\mathcal{L}, \mathcal{D})$ can be written as

$$Rev(\mathcal{L}, \mathcal{D}) = \sum_{S \in \mathcal{P}} Prob_{\theta \leftarrow \mathcal{D}}(\theta \in S) \cdot \hat{p} \cdot Prob_{\theta \leftarrow \mathcal{D}|S}(\theta_{i^*} \geq \hat{p} | \theta_{-i^*} = v_{-i^*} \text{ and } v_i^* \geq v_j \forall j).$$

But \hat{p} is not the optimal price if we have extra information. Indeed, if we condition on the set S then

$$\begin{aligned} & \hat{p} \cdot Prob_{\theta \leftarrow \mathcal{D}|S}(\theta_{i^*} \geq \hat{p} | \theta_{-i^*} = v_{-i^*} \text{ and } v_i^* \geq v_j \forall j) \leq \\ & \max_p p \cdot Prob_{\theta \leftarrow \mathcal{D}|S}(\theta_{i^*} \geq p | \theta_{-i^*} = v_{-i^*} \text{ and } v_i^* \geq v_j \forall j) \end{aligned}$$

This implies that

$$\begin{aligned} & \sum_{S \in \mathcal{P}} \Pr_{\theta \leftarrow \mathcal{D}}(\theta \in S) \cdot [(1 - \delta)Rev(\mathcal{L}, \mathcal{D}|S) - \epsilon] \geq \\ & (1 - \delta) \cdot \sum_{S \in \mathcal{P}} Prob_{\theta \leftarrow \mathcal{D}}(\theta \in S) \cdot \hat{p} \cdot Prob_{\theta \leftarrow \mathcal{D}|S}(\theta_{i^*} \geq \hat{p} | \theta_{-i^*} = v_{-i^*} \text{ and } v_i^* \geq v_j \forall j) - \epsilon \end{aligned}$$

but the right hand side of this inequality is $(1 - \delta)Rev(\mathcal{L}, \mathcal{D}) - \epsilon$ and the left hand side is less than or equal to $Rev(\mathbb{L}, \mathcal{D})$. From Ronen's work [31] we know that $Rev(\mathcal{L}, \mathcal{D}) \geq \frac{1}{2}OPT(\mathcal{D})$. We conclude that

$$Rev(\mathbb{L}, \mathcal{D}) \geq \frac{1 - \delta}{2}OPT(\mathcal{D}) - \epsilon$$

Q.E.D.

4.4 The performance of \mathbb{L} compared to other mechanisms

Comparing \mathbb{M} and \mathbb{L} . The main difference between \mathbb{L} and \mathbb{M} is that \mathbb{L} is *explicit*. That is, we can construct it without appealing to the existence of some possibly unknown optimal classical mechanism \mathcal{M} . Thus, \mathbb{L} is more practical for scenarios such as correlated distributions, where the optimal classical mechanism \mathcal{M} exists but is infeasible to implement.⁸

⁷The meet is defined as follows: given a true type profile θ and a player j , there exists a unique set $S_j(\theta) \in \mathcal{P}_j$ such that $\theta \in S_j(\theta)$. Define $S(\theta) = \bigcap_{j \in -i^*} S_j(\theta)$. The set $\mathcal{S} = \{S(\theta) : \theta \in \Theta\}$ is the *meet* of the partitions. It is well known that \mathcal{S} is itself a partition of Θ and that it is finer than any of the \mathcal{P}_j .

⁸For single-unit auctions and correlated distributions, the mechanism of Cremer and McLean [15] extracts as much revenue as possible. However, it violates the individual rationality condition because truthful players may receive negative utility from this mechanism. Furthermore, it requires that the distribution \mathcal{D} is common knowledge among the players. More details are given in section 7.

IID Valuations and Comparison with Segal [32] and Baliga and Vohra [2] For all distributions \mathcal{D} , the revenue of the mechanism \mathbb{L} is guaranteed to be at least one half of the optimal revenue. However, \mathbb{L} can do better under certain assumptions on \mathcal{D} . For example, if all the valuations are identically and independently distributed, then the optimal auction always sells to the player with the highest valuation. In this case, \mathbb{L} obtains at least $(1 - \delta)OPT(\mathcal{D}) - \epsilon$ revenue for arbitrarily small δ, ϵ and for an arbitrary number of players.

We can compare the performance of \mathbb{L} in this setting with that of existing mechanisms. Segal [32] and Baliga and Vohra [2] show how, in the iid setting, the auctioneer can charge a reserve price to i^* by estimating the distribution from the vector v_{-i^*} . As the number of players n goes to infinity, this estimate becomes more precise and the revenue approaches that of the optimal mechanism. In contrast, our mechanism \mathbb{L} works for any number of players. Furthermore, in the iid setting our crowdsourced-bayesian assumption is implied by the weaker

Self-Knowledge Assumption: *Each player i knows the distribution D from which his valuation is drawn.*

As long as this assumption holds, our mechanism obtains at least as much revenue as existing mechanisms. We acknowledge the possibility that the self-knowledge assumption does not hold. In this case, the existing mechanisms are superior.

5 Optimality of our Mechanism for $n = 2$

In section 3, we showed a crowdsourced-Bayesian mechanism whose revenue is very close to the revenue of the optimal centralized-Bayesian mechanism where players' types correspond to their valuations. Our mechanism achieves this by extending players' types to pairs $(\theta_i, \mathcal{D}_i)$ of valuations and knowledge. An important question is whether our mechanism is optimal among those that allow players to have these extended types.

We show that our mechanism \mathbb{M} is virtually optimal for $n = 2$. For arbitrarily small ϵ , no two-step DST crowdsourced-Bayesian mechanism can achieve more than a $(1 + \epsilon)$ factor of the revenue of \mathbb{M} , even for single-item auctions with independent bidders. The proof consists of showing that no two-step DST crowdsourced-Bayesian mechanism M can achieve revenue higher than $(\frac{1}{2} + \epsilon) \cdot OPT(\mathcal{D})$, where $OPT(\mathcal{D})$ is the revenue of the optimal *centralized*-Bayesian mechanism and ϵ is an arbitrarily small constant. Since \mathbb{M} obtains a profit of at least $\frac{1}{2}OPT(\mathcal{D}) - \epsilon$, these two facts together imply that the difference in revenue between M and \mathbb{M} is at most $\epsilon \cdot (OPT(\mathcal{D}) + 1)$, for arbitrarily small ϵ . Since M was an arbitrary two-step DST crowdsourced-Bayesian mechanism, \mathbb{M} is virtually optimal among this class of mechanisms.

We need to prove now that M cannot obtain revenue higher than $(\frac{1}{2} + \epsilon)OPT(E, \mathcal{D})$. To prove this, we require that players always prefer winning the auction to losing. In a crowdsourced-Bayesian mechanism, losing players can be given cash rewards for their knowledge. These cash rewards must be kept small in order to prevent players from preferring to become informants instead of participating in the auction. Our mechanism \mathbb{M} gives rewards no larger than an arbitrarily small ϵ , so it fits under this class of mechanisms.⁹ When players have a preference for winning, the following theorem holds:

Theorem 3. *Let E be a single-item auction environment with two players. Let M be a two-step DST crowdsourced-Bayesian mechanism where players have a preference for winning. Then for every $\epsilon > 0$ there exists a distribution \mathcal{D} of player valuations such that*

$$Rev(M, E, \mathcal{D}) \leq (\frac{1}{2} + \epsilon)OPT(E, \mathcal{D}),$$

where $OPT(E, \mathcal{D})$ is the revenue of the optimal centralized-Bayesian mechanism.

This can be shown via two lemmas.

⁹Even without this small reward, the informant in \mathbb{M} is chosen at random, so players cannot “choose” to become informants. On the other hand, our mechanism \mathbb{L} in section 4 ensures that players always prefer winning to losing by ensuring that winners always get utility greater than ϵ , and losers (who become informants) never get utility higher than ϵ .

Lemma 6. *Let E be a single-item auction environment with two players. Let M be two-step DST, crowdsourced-Bayesian mechanism in which players have a preference for winning. Then for all distributions \mathcal{D} we have $Rev(M, E, \mathcal{D})$ is bounded above by the revenue of a Vickrey auction with monopoly reserve prices.*

Lemma 7 (Hartline and Roughgarden [23]). *Let E be a single-item auction environment with two players. For every $\epsilon > 0$, there exists a distribution \mathcal{D} of (independent) valuations such that the Vickrey auction with monopoly reserve price does not obtain more than $(\frac{1}{2} + \epsilon)OPT(E, \mathcal{D})$ revenue.*

The two lemmas combined imply the theorem. We prove the first of the lemmas below. The proof of the second can be found in the referenced paper.

Proof. We begin by analyzing a deterministic mechanism. At the end of this proof, we show how to generalize to randomized mechanisms. We denote by θ_i the true valuation of player i , and by \mathcal{D}_i player i 's true belief about the other player's valuation.¹⁰

Let M be a deterministic, two-step DST, crowdsourced-Bayesian mechanism for auctioning a single good between two players. Player 1 announces a valuation $v_1 \in \Theta_1$ and a distribution $\mathcal{F}_1 \in \Delta(\Theta_2)$. Analogously, player 2 announces a valuation $v_2 \in \Theta_2$ and a distribution $\mathcal{F}_2 \in \Delta(\Theta_1)$. Since M is two-step DST, the only strategy of player 1 that survives iterated elimination of dominated strategies is announcing $v_1 = \theta_1$, his true valuation, and $\mathcal{F}_1 = \mathcal{D}_1$, his true belief about θ_2 .

For a pure strategy profile $s = ((v_1, \mathcal{F}_1), (v_2, \mathcal{F}_2))$ let $x(s)$ denote the winner of the auction (with winner 0 if no one gets the good) and let $p(s)$ be the price paid by the winner. We are assuming that the mechanism is ex-post individually rational, so no revenue can be obtained from a player who did not win.

First of all, consider the case where $\mathcal{F}_1 = \mathcal{F}_2$ and both players are truthful. In this case, the players' valuations are iid. As observed in section 4, the mechanism \mathbb{L} obtains $(1 - \delta) \cdot Opt(\mathcal{D}) - \epsilon$ in this environment, for arbitrarily small δ and ϵ . The mechanism \mathbb{L} generates virtually as much revenue as the optimal classical mechanism (which is a Vickrey auction with reserve price in our case of i.i.d bidders). Thus, when distributions are identical, the mechanism should sell the good to the highest value player.¹¹

Now suppose that the profile of strategies is $s = ((v_1, \mathcal{F}_1), (v_2, \mathcal{F}_2))$, with $\mathcal{F}_1 \neq \mathcal{F}_2$. Assume without loss of generality that $v_1 > v_2$. We are going to show that $x(s) \neq 2$. Assume the contrary. Player 1 loses the auction even though he has the highest valuation. However, if he announces \mathcal{F}_1 to be equal to \mathcal{F}_2 , then by the above paragraph he would win the auction. Since player 1 prefers to win, he has an incentive not to announce his true belief about player 2's value, even when both players are truthful about their valuations. This violates the two-step DST property of M , and hence gives us a contradiction.

We conclude that, if a crowdsourced-Bayesian two-step DST mechanism (where players prefer to win than to lose) maximizes revenue, then it can never sell the good to the player with the lowest valuation. Given that it always sells the good to the player with the highest value, it cannot obtain more revenue than the Vickrey auction with reserve price.

The above analysis applied to deterministic mechanisms. If the mechanism instead is randomized (like our mechanism \mathbb{M}), the given revenue bounds still holds. Let M be a randomized and let M_r be the resulting *deterministic* mechanism given that M uses randomness r . The expected revenue of M is $\mathbb{E}_r[Rev(M_r, E, \mathcal{D})]$. None of the mechanisms M_r obtain more revenue than the Vickrey auction with monopoly reserve price. Thus, the revenue of M is also bounded by that of the Vickrey auction with monopoly reserves.

Q.E.D.

Comparison with Caillaud and Robert [6] Remarkably, this optimality implies a separation between what is achievable in our knowledge model and what is achievable when we assume common knowledge of \mathcal{D} . Caillaud and Robert [6] show a mechanism to implement the optimal single-unit auction when players have

¹⁰In general \mathcal{D}_i is player i 's belief about the valuation profile, *including his own valuation*. However, since player i is always certain about his own valuation being θ_i , we interpret \mathcal{D}_i as his belief about the other player's value.

¹¹We are relying on the fact that for iid distributions, no centralized-Bayesian mechanism that achieves the optimal revenue can sell to the player with the lowest valuation.

independent valuations, \mathcal{D} is unknown to the seller, and \mathcal{D} is common knowledge among the players. They show that under these assumptions, they obtain as much revenue as the optimal auction with a knowledgeable seller. In contrast, we have shown above that we cannot obtain more than $(\frac{1}{2} + \epsilon) \cdot OPT(\mathcal{D})$ under the crowdsourced-bayesian assumption when there are only two players. Our mechanism does not guarantee as much revenue as theirs because it uses a weaker assumption on the knowledge of the players.

6 A reverse version of the Bulow-Klemperer inequality

Of independent interest, the proof of theorem 1 implies a reverse version of the Bulow-Klemperer theorem [5]. The Bulow-Klemperer theorem states that for a single good context and symmetric players¹² the Vickrey auction on n players always generates more expected revenue than the optimal auction with $n - 1$ players.

Informally, write this as the inequality $OPT(n - 1) \leq Vickrey(n)$ where $OPT(n - 1), Vickrey(n)$ are the revenues obtained by the optimal and Vickrey auctions on $n - 1$ and n players respectively. We show a reverse direction of this inequality.

Corollary 1 (Reverse Bulow-Klemperer Inequality). *For single-item auctions with i.i.d. valuations drawn from a regular distribution,*

$$OPT(n - 1) \leq Vickrey(n) \leq \frac{n}{n - 1} OPT(n - 1)$$

That is, the optimal auction on $n - 1$ players generates almost as much revenue as the Vickrey auction on n players, losing at most a $\frac{1}{n}$ fraction of the revenue. Indeed, we show that for general downward-closed environments and i.i.d. buyers the optimal auction on $n - 1$ players can generate *more* revenue than the Vickrey auction, disproving the Bulow-Klemperer result for these general environments.

Proof. Note that our proof of theorem 1 relies on the inequality

$$\frac{1}{n} \sum_{i=1}^n Rev(\mathcal{M}(\mathcal{D}_{-i})) \geq \frac{n - 1}{n} OPT(E, \mathcal{D})$$

where $Rev(\mathcal{M}(\mathcal{D}_{-i}))$ is the revenue of the optimal mechanism where player i does not participate and the other players' valuations are drawn from the distribution \mathcal{D}_{-i} . To make this inequality more suggestive, we write E_{-i} to be the environment E with player i removed, and we write $Rev(\mathcal{M}(\mathcal{D}_{-i})) = OPT(E_{-i}, \mathcal{D}_{-i})$. Thus, we can rewrite the above inequality as

$$\mathbb{E}_i OPT(E_{-i}, \mathcal{D}_{-i}) \geq \frac{n - 1}{n} OPT(E, \mathcal{D}).$$

This means that, if we remove a player at random, the revenue of the optimal auction with n players does not drop by more than a $\frac{1}{n}$ fraction. *Q.E.D.*

Our proof applies to general environments and arbitrary symmetric distributions.

Corollary 2. *The expected revenue of a Vickrey auction on an assignment environment with n symmetrically distributed players is always lower than or equal to $\frac{n}{n - 1}$ of the expected revenue of the optimal auction with $n - 1$ players.*

¹²The theorem can be extended to multiple goods and affiliated valuations.

6.1 The Bulow-Klemperer result does not hold for general downward closed environments

For general downward closed environments, the Vickrey auction cannot provide a better than constant approximation to the revenue of the optimal auction. A known result by Hartline and Roughgarden [23] states that, even for i.i.d. buyers, there exists a downward-closed environment E^* where the Vickrey auction cannot achieve more than half the revenue of the optimal auction.¹³

However, for general downward-closed environments with i.i.d. bidders, we have shown that the optimal auction with $n - 1$ players gives a $(1 - \frac{1}{n})$ approximation of the optimal revenue with n players. Thus, the Vickrey auction does not beat the optimal auction on $n - 1$ players for general downward closed environments.

Corollary 3. *Let E be a downward-closed environment, and let \mathcal{D} be a distribution of players' valuations. If $\text{Rev}(\text{Vickrey}(E, \mathcal{D})) < (1 - \frac{1}{n}) \cdot \text{Rev}(\text{OPT}(E, \mathcal{D}))$, then $\text{Rev}(\text{Vickrey}(E, \mathcal{D})) < \mathbb{E}_i \text{OPT}(E_{-i}, \mathcal{D}_{-i})$. In particular, there exists a downward closed environment E^* and an i.i.d. distribution \mathcal{D}^* of regular values such that the optimal auction on $n - 1$ players obtains more revenue than the Vickrey auction on n players. Thus, the Bulow-Klemperer inequality does not hold for general downward closed environments.*

7 Our mechanisms as black-box results

Our mechanism \mathbb{M} as a compiler The main contribution of this work is a very general method for an ignorant seller to extract knowledge from informed buyers. We can think of our method of removing a player at random and incentivizing this player to state his knowledge via a proper scoring rule as a compiler, which can be applied to transform *any classical dominant strategy truthful* mechanism which assumes the seller has some knowledge into a crowdsourced-Bayesian one.

Our method is very robust against errors in players' second and higher-order knowledge. As an example of a procedure that is not robust, consider the case where the distribution \mathcal{D} is common knowledge amongst the players. A seller who would like to know \mathcal{D} could ask all players to report a distribution. If all reports agree, then the seller uses this distribution to determine prices. Otherwise, all players get a large penalty and the mechanism ends. The problem with this approach is equilibrium selection. Any strategy profile where all players report the same distribution \mathcal{F} could be an equilibrium, even if $\mathcal{F} \neq \mathcal{D}$.

In more specific settings, analogous versions of our mechanism \mathbb{L} may be able to incorporate players' knowledge into existing mechanisms. We give a brief discussion in the following paragraphs

Simple Mechanisms The look-ahead auction \mathcal{L} can be considered a “simple” mechanism in that the winner of the auction is determined without the need to know the distribution \mathcal{D} . The only use of the distribution is to set a reserve price. The comparison of simple and optimal mechanisms is a recent and interesting approach in mechanism design. Neeman [28] shows that English Auctions are approximately optimal, performing an analysis similar to Ronen's work on Vickrey auctions. Hartline and Roughgarden [23] study approximately-optimal mechanism design in downward-closed and matroid environments, under the assumptions that the distributions of valuations are independent and satisfy some regularity properties. In particular, they show that Vickrey Auctions with reserve price achieve a constant-factor approximation of optimal revenue for a large class of environments and distributions. They also show matching constant lower bounds in the approximation factor, and sub-constant lower bounds for special kinds of distributions and environments. We can leverage simple mechanisms to create crowdsourced mechanisms for new environments. If a mechanism M uses a distribution \mathcal{D} only to determine prices (and not to determine the allocation to players), then the mechanism can in principle be used to construct a crowdsourced mechanism \mathbb{M} with similar revenue guarantees, analogously to how we constructed the mechanism \mathbb{L} from the classical mechanism \mathcal{L} .

¹³See example 3.3 in the corresponding paper. In their environment E^* there are two groups of buyers, and one cannot simultaneously sell to both.

Extracting the Full Surplus If we accept the weaker notion of ex-interim individual rationality instead of ex-post individual rationality, Cremer and McLean [15] give a classical Bayesian mechanism that extracts the full surplus from the buyers. This mechanism holds under two special conditions: first, the distribution \mathcal{D} must be common knowledge among the players. Second, for each player the $|\Theta_i| \times |\Theta_{-i}|$ matrix whose rows are given by the distributions $\mathcal{D}(\theta_{-i}|\theta_i)$ must be of full rank. Note that this mechanism only works when the designer knows the distribution \mathcal{D} , this distribution is common knowledge among the players, and each player i 's beliefs are given *precisely* by $\mathcal{D}|\theta_i$. Because of this, we cannot apply our techniques to make this mechanism crowdsourced-Bayesian: using scoring rules only incentivizes players to reveal $\mathcal{D}|\theta_i$, but the designer needs to learn the exact distribution \mathcal{D} , which cannot immediately be reconstructed from $(\mathcal{D}|\theta_1, \dots, \mathcal{D}|\theta_n)$ in order to apply Cremer and McLean's mechanism. The impossibility of extracting the full surplus when relaxing the common knowledge assumption has been investigated widely. An example of recent research is [1].

Recent work on correlated and multi-parameter auctions More recently, there have been new developments in algorithmic mechanism design for correlated valuations and for multi-parameter problems. Papadimitriou and Pierrakos [30] have shown that for 2 players there exists a PTAS¹⁴ that approximates the optimal ex-post individually rational auction with correlated valuations, but for 3 or more players achieving more than 99.97 % of the optimal revenue is computationally infeasible.¹⁵ For their 2 player auction, this result can be combined with our technique to remove the assumption that the designer knows the distribution of values.¹⁶

Chawla, Hartline, Malec and Sivan [7] show a sequential-posted price mechanism that achieves a constant-factor approximation of the optimal revenue even in multi-parameter settings. One particular case they consider is where agents can buy only one good, but may have different valuations for different goods. They give a dominant strategy mechanism that achieves a constant approximation factor to the optimal revenue and can be computed in polynomial time. Because the mechanism is DST, it can be combined with our technique to remove the assumption that the seller knows the distribution, in cases where the players know each other better than the seller knows them.

Dobzinski, Fu and Kleinberg [18] show an optimal truthful in expectation mechanism for single-item auctions with correlated distributions, as well as unit-demand multi-item auctions. Their mechanisms can be constructed in time polynomial in the size of the support of \mathcal{D} . Our technique can potentially be used in conjunction with their mechanisms to construct a crowdsourced-Bayesian mechanism for single-item auctions with correlated distributions¹⁷ that is not two-step DST, but two-step *truthful in expectation*. That is, a mechanism where players maximize their expected utility by being truthful, before observing the coin tosses of the mechanism.

Conclusion We have shown how to use scoring rules and a reverse version of the Bulow-Klemperer inequality in order to help an ignorant seller extract knowledge from a group of informed buyers. Beyond auctions, there are many situations in mechanism design where players can have knowledge about each other. We believe the techniques we presented will prove to be useful in these situations as well.

Acknowledgments

We would like to thank Gabriel Carroll for useful comments after oral presentation of this paper, as well as the Program Committee of the Workshop on Bayesian Mechanism Design, whose reviews helped improve this paper.

¹⁴A PTAS is a Polynomial Time Approximation Scheme. For any desired approximation factor ϵ , there exists an efficient algorithm that obtains at least $(1 - \epsilon)$ fraction of the optimal revenue.

¹⁵In particular, it is NP-Hard.

¹⁶Using our technique loses a $\frac{1}{2}$ fraction of the optimal revenue, so this way of removing the auctioneer's knowledge is no better than using our mechanism \mathbb{L} from section 4.

¹⁷Or for unit-demand multi-item auctions.

A Proof of lemma 1

Lemma. *Let E be a downward-closed environment and let M be a classical direct revelation dominant strategy truthful mechanism. Then for every $\delta > 0$ there exists a classical direct revelation mechanism M' which is strictly DST and which satisfies*

$$Rev(M', E, \mathcal{D}) \geq (1 - \delta)Rev(M, E, \mathcal{D})$$

for all distributions \mathcal{D} .

Proof. Construct a lottery mechanism L as follows

$L((\tau_1, \dots, \tau_n))$

1. choose a player i , a set $A_i \subset X$ and a price p_i at random.
2. If $\tau_i(A_i) \geq p_i$, output the outcome $((\emptyset, \emptyset, \dots, A_i, \dots, \emptyset), (0, \dots, 0, p_i - \epsilon, 0, \dots, 0))$ where ϵ is an arbitrary number satisfying $0 < \epsilon < 1$.
3. If $\tau_i(A_i) < p_i$, output the outcome $((\emptyset, \emptyset, \dots, \emptyset), (0, 0, \dots, 0))$.

Now given M, L and δ , construct the mechanism M' as follows

$M'((\tau_1, \dots, \tau_n); \delta)$

1. With probability $1 - \delta$, output $M(\tau)$ and STOP.
2. With probability δ , output $L(\tau)$ and STOP.

In this proof we use the fact that for every set A , $\tau_i(A) \in \mathbb{Z}$. This requirement can be weakened to require $\tau_i(A) \in D$ where $D \subset \mathbb{R}$ is a discrete set. This is true whenever there is a smallest unit of money (such as one cent) on which all bids need to be denominated.

Let θ_i be player i 's true valuation function and let $\tau_i \neq \theta_i$. We show that for all profiles of opponent announcements τ_{-i} and all settings r of the coin tosses of M' we have

$$u_i(M'(\theta_i, \tau_{-i}, r)) \geq u_i(M'(\tau_i, \tau_{-i}, r)),$$

and furthermore this inequality is strict for at least one setting of τ_{-i} and r .

For any sequence of coin tosses $r \in \{0, 1\}^*$ write $A_i(r), p_i(r)$ for the set and price assigned to player i in the mechanism L in step 1. Note that $A_i(r), p_i(r)$ are completely specified by r and do not depend on the actions of the players.

Write $u_i(M'(\theta_i, \tau_{-i}, r)) = (1 - \delta)u_i(M(\theta_i, \tau_{-i}, r)) + \delta L(\theta_i, \tau_{-i}, r)$. Because the mechanism M is weakly dominant strategy truthful, it is true that $(1 - \delta)u_i(M(\theta_i, \tau_{-i}, r)) \geq (1 - \delta)u_i(M(\tau_i, \tau_{-i}, r))$. Thus, it suffices to show that L is strictly dominant strategy truthful.

First, we show that L is *weakly* dominant strategy truthful. That is, for any choice of random coins r the mechanism L weakly incentivizes player i to be truthful. If player i does not get chosen in step 1 of L , then his utility is always zero, regardless of his action. If player i gets chosen in step 1 of L , he is given a choice of purchasing a set A_i at price p_i , which he can accept or reject depending on his announcement. If $p_i < \theta_i(A_i), \tau_i(A_i)$ or $p_i > \theta_i(A_i), \tau_i(A_i)$, the player gets the same utility from announcing θ_i or τ_i . The interesting cases are $\theta_i(A_i) < p_i < \tau_i(A_i)$ and $\tau_i(A_i) < p_i < \theta_i(A_i)$.

Consider the case where $\tau_i(A_i) < p_i < \theta_i(A_i)$. In this case, the utility of player i from choosing θ_i is $\theta_i(A_i) - (p_i - \epsilon) > 0$. If player i chooses τ_i instead, he does not get assigned a good and his utility is zero. For this choice of random coins, player i strictly prefers being truthful.

Now consider the case $\theta_i(A_i) < p_i \leq \tau_i(A_i)$. Since p_i and $\theta_i(A_i)$ are both integers, $p_i \geq \theta_i(A_i) + 1 > \theta_i(A_i) + \epsilon$.

In this case, if player i announces τ_i , his utility will be $\theta_i(A_i) - (p_i - \epsilon) < 0$. If the player announces θ_i , he gets assigned no good and charged a price of zero, so his utility is zero. Thus, player i strictly prefers announcing θ_i for this choice of random coins.

We conclude that L is strictly dominant strategy truthful, and so is M' . Notice $Rev(M', \mathcal{D}, E) = (1 - \delta)Rev(M, \mathcal{D}, E) + (\delta)Rev(L, \mathcal{D}, E)$ and that $Rev(L, \mathcal{D}, E) > 0$. Thus, $Rev(M', E, \mathcal{D}) \geq (1 - \delta)Rev(M, E, \mathcal{D})$. *Q.E.D.*

References

- [1] P. Borelli. On the genericity of full surplus extraction in mechanism design. *Journal of Economic Theory*, 144(3):1320-1332, 2009.
- [2] S. Baliga and R. Vohra. Market research and market design. *Advances in Theoretical Economics*, 3(1):1059-1059, 2003.
- [3] D. Bergemann and S. Morris. Robust mechanism design. *Econometrica*, Vol. 73, No. 6: 1771-1813, 2005.
- [4] G.W. Brier. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78:1-3, 1950.
- [5] J. Bulow and P. Klemperer. Auctions Versus Negotiations. *The American Economic Review*, Vol. 86, No. 1: 180-194, 1996.
- [6] B.J.M. Caillaud and J. Robert. Implementation of the Revenue-Maximizing Auction by an Ignorant Seller./ *Review of Economic Design*, Vol. 9, 127-143, 2005.
- [7] S. Chawla, J.D. Hartline, D.L. Malec and B. Sivan. Multi-parameter Mechanism Design and Sequential Posted Pricing. In *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 311-320. ACM, 2010.
- [8] S. Chawla, J.D. Hartline and B. Sivan. The theory of crowdsourcing.
- [9] J. Chen, A. Hassidim, and S. Micali. Robust Perfect Revenue from Perfectly Informed Players. *Innovations in Computer Science*, pages 94-105, Beijing, 2010.
- [10] J. Chen and S. Micali. A New Approach To Auctions And Resilient Mechanism Design. In *Proceedings of the 41st annual ACM symposium on Theory of Computing*. June 2009, Pages 503-512.
- [11] J. Chen and S. Micali. Mechanism Design with Set-Theoretic Beliefs. To appear in *Foundations of Computer Science (FOCS)*, 2011.
- [12] J. Chen, A. Hassidim, and S. Micali. Robust Perfect Revenue from Perfectly Informed Players. *Innovations in Theoretical Computer Science (ITCS)*, pp. 94-105, 2010.
- [13] J. Choi, T. Kim. A Nonparametric, Efficient Public Good Decision Mechanism: Undominated Bayesian Implementation. *Games and Economic Behavior*, Volume 27, Issue 1, April 1999, Pages 64-85
- [14] K.S. Chung and J.C. Ely. Foundations of Dominant-Strategy Mechanisms. *Review of Economic Studies*, 74(2):447-476, 2007.
- [15] J. Cremer and R. P. McLean. Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions. *Econometrica* Vol. 56, No. 6: 1247-1257, 1988.
- [16] J. Cremer and M. H. Riordan. A Sequential Solution to the Public Goods Problem. *Econometrica* Vol. 53, No. 1 (Jan., 1985), pp. 77-84
- [17] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 129-138. ACM, 2010.

- [18] Dobzinski, S. and Fu, H. and Kleinberg, R. Optimal Auctions with Correlated Bidders are Easy. In Proceedings of the 43rd Symposium of the Theory of Computing, 2011.
- [19] T. Gneiting and A.E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [20] A.V. Goldberg and J.D. Hartline. Competitiveness via consensus. In *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 215–222, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.
- [21] A.V. Goldberg, J.D. Hartline, A.R. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.
- [22] I.J. Good. Rational Decisions. *Breakthroughs in Statistics: Foundations and basic theory*, page 365, 1992.
- [23] J.D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the tenth ACM conference on Electronic commerce*, pages 225–234. ACM, 2009.
- [24] M.O. Jackson. Bayesian Implementation . *Econometrica*. Vol. 59, No. 2, pp. 461-477. 1991.
- [25] M.O. Jackson. A crash course in implementation theory. *Social Choice and Welfare*, 18(4):655–708, 2001.
- [26] E. Maskin and T. Sjostrom. Implementation theory. *Handbook of Social Choice and Welfare*, 1:237–288, 2002.
- [27] R.B. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58, 1981.
- [28] Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.
- [29] Z. Neeman. The relevance of Private Information in Mechanism Design *Journal of Economic Theory*, Volume 117, Issue 1, 55–77. 2004.
- [30] C. Papadimitriou and G. Pierrakos. On Optimal Single-Item Auctions. In Proceedings of the 43rd Symposium of the Theory of Computing, 2011.
- [31] A. Ronen. On approximating optimal auctions. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 11–17. ACM, 2001.
- [32] I. Segal. Optimal pricing mechanisms with unknown demand. *The American economic review*, 93(3):509–529, 2003.
- [33] R. Wilson. Game-theoretic approaches to trading processes. In *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, pages 33–77, 1987.