

Scoring Auctions with Coarse Beliefs

Joseph Feffer*

April 1, 2025

Abstract

This paper studies a simplicity notion in a mechanism design setting in which agents do not necessarily share a common prior. I develop a model in which agents learn from a coarse signal structure. After receiving information, the agents have uncertainty about the environment in which they play and about their opponents' higher-order beliefs. A mechanism admits a coarse beliefs equilibrium if agents can play best responses even with this uncertainty. I study the existence of coarse beliefs equilibria in multidimensional scoring auctions. I fully characterize a property that determines whether coarse beliefs equilibria exist under canonical signal structures. I also show that considering these signal structures is without loss as, in general, auctions that satisfy this property admit equilibria under more coarse signal structures than those that do not. I then find a simple, primitive condition on the auction's rules to identify auctions that satisfy this property and use the condition to classify real-world scoring auctions by their strategic simplicity.

*Stanford Graduate School of Business. Email: feffer@stanford.edu. I am grateful to Mohammad Akbarpour, Benjamin Brooks, Songzi Du, Benjamin Golub, Klajdi Hoxha, Mike Ostrovsky, Andrzej Skrzypacz, Eric Tang, and Filip Tokarski and seminar participants at Stanford University and Microsoft Research New England for helpful comments and suggestions.

1 Introduction

It has been notoriously difficult for the standard toolbox of economic theory to guide design in multidimensional auction settings. Optimality criteria often yield mechanisms that are either too intractable to characterize or too complex to implement in practice Asker and Cantillon [2010]. However, in some settings, it seems plausible that agents do not have enough information about their environment to be able to play in many of the mechanisms considered in this optimization problem. Agents may know some information about their industry or about their opponents but they are unlikely to form the precise common priors that they are assumed to possess in typical models. With such coarse beliefs, even the most sophisticated agents may fail to identify an action as optimal in certain games. When this happens, considering the strategic limitations caused by the imprecision of agents' information could provide a meaningful constraint on the set of multidimensional auction formats a designer should consider. After all, revenue projections mean little if agents cannot play the equilibria upon which they are based. Thus, in this paper, I study sensible relaxations on agents' information and characterize the set of auctions in which they can play.

Public procurement auctions represent an ideal setting in which to study strategizing with coarse beliefs. First, these auctions are almost always conducted in pay-as-bid formats that do not admit dominant-strategy or ex post equilibria. Thus, participants' strategies must depend nontrivially on their beliefs. Second, competing firms are generally sophisticated; while their beliefs may not be fully precise, they often hire market researchers and auction consultants to help them acquire relevant information and bid intelligently. Third, auctioneers in public market design settings often care about strategic simplicity for participants.¹ Lastly, governments often care about non-price features of proposed projects, turning a simple auction setting into a complex multidimensional one. Thus, the simplicity notion studied in this paper may be a useful criterion to narrow the auctioneer's design problem.

Specifically, I model a scenario in which a buyer (female) seeks to obtain a good from one of two potential sellers (male)². The good can be provided at heterogeneous quality levels, so the buyer has preferences over both the quality of the good she obtains as well as the price at which she procures it. She wants to pick a mechanism from a class of auctions called first-score auctions. These are the generalizations of first-price auctions to

¹A 2021 EU directive listed transparency and fairness as core principles for auction procedure. Strategically simple mechanisms are more transparent because the mapping from a firm's fundamentals to their bidding strategy is more clear. They are fairer because they require less sophistication and less market research to form best-response strategies. See Verens [2021] for more information.

²The restriction to 2 participants is entirely for ease of exposition.

this procurement setting and are commonly used in practice. Before the auction, each seller has private information about his own cost structure. He does not know the state of the industry, which is the distribution from which his opponent’s type is drawn, but receives coarse information about the state from an exogenous signal structure. Because signal realizations are coarse, each seller has uncertainty about the environment in which he plays and about his opponents’ beliefs. I characterize the first-score auction formats in which both of these forms of uncertainty are strategically irrelevant. In the equilibria of such mechanisms called coarse beliefs equilibria (CBE), there is a bid in the auction that is a best-response at all possible resolutions of this uncertainty. Coarse beliefs equilibrium is an intermediate solution concept that attempts to incorporate both the robustness of ex post equilibrium and the general existence of Bayes-Nash equilibrium. In a coarse beliefs equilibrium, a player will not know what his opponent has learned; he will simply know that his opponent is collecting and using information intelligently.

Results. There are two types of signal structures to which I devote special attention. The first is a moment-based signal structure, meant to approximate what agents can learn through an idealized sampling and estimation process. The second is a history-dependent signal structure meant to resemble what agents can learn through publicly available data on prior auctions. In my first main result, I show that when agents learn from a moment-based signal structure, a CBE exists if and only if the mechanism satisfies a property called *the fixed-order property*. Similarly, in my second main result, I show that when agents learn from histories, CBE existence is guaranteed only when the fixed-order property is satisfied. This property describes settings when, in equilibrium, the auction’s ex post allocation function is *fixed*, regardless of the type distribution. That is, if type A is given the contract over type B at one distribution, then A will be given the contract over B at every possible type distribution. *The fixed-order property* is a natural property that is satisfied in common auction formats—like first-price, second-price, and all-pay auctions—when agents’ private information is one-dimensional. Lastly, I show that auctions that satisfy the fixed-order property admit the greatest number of CBE with respect to general convex signal structures. Informally, this implies that auctions satisfying this property admit equilibria under more coarse signal structures than those that do not.

Having established the importance of the fixed-order property in CBE existence, I next study the auctions that satisfy it. I start by showing that, while multidimensional settings generally do not admit some of the nice properties of one-dimensional ones, these properties are restored exactly when the fixed-order property is satisfied. Theorem 4 states that this property holds if and only if first-score and “second-score” auctions satisfy a Payoff Equivalence result. In general, the multidimensional forms of first and second-price auctions do

not implement the same interim allocation rules. This makes first-score auctions harder to study than their one-dimensional counterparts because we cannot understand them through a similar dominant strategy mechanism. However, when the fixed-order property is satisfied, payoff equivalence is restored and strategies and expected outcomes of an auction are easier to analyze.

For studying CBE, the connection between first and “second-score” auctions has important implications. First, it gives a nice interpretation of the information that agents need to strategize in a CBE. Second, it allows us to derive a simple condition on the rules of a first-score auction that can distinguish mechanisms that admit CBE and those that do not. This condition is an easily testable function of the auction’s rules and, in particular, does not require explicit computation of an equilibrium. Applying this condition, I show that generic scoring rules do not satisfy this criterion and only a strict subset of scoring rules used in practice do. Thus, considering the simplicity of strategies in this setting could help designers by restricting the set of mechanisms they must consider implementing. Lastly in the appendix, I endogenize the signal structures from which agents learn by studying a prior-free information acquisition model. I then provide microfoundations for the concept of a coarse beliefs equilibrium.

Related literature. This paper primarily contributes to the literature on simplicity in mechanism design. While most works in this literature (like Li [2017]) study failures of contingent reasoning in extensive form games, my work instantiates a more general definition of simplicity advanced by Li and Dworzak [2024]. These authors study settings in which the relaxation of a standard assumption of agent behavior, in general, leads to strategic confusion. They define a simple mechanism as one in which this confusion either is never generated or is easily resolvable. In this paper, the strategic confusion stems from the coarseness of agents’ learning process. Of papers that study specific forms of Li and Dworzak [2024]’s general simplicity, Börgers and Li [2019] and Pernoud and Gleyze [2023] similarly consider belief-related notions of simplicity. The first paper considers a setting in which agents have precise first-order beliefs but cannot form accurate higher-order beliefs. They thus call a mechanism simple if agents’ strategies do not depend on these higher-order beliefs. The second studies a form of simplicity when agents do not fully know their values and have no incentives to learn about their opponents’ values. They find this restriction is severely limiting as in almost all practically relevant settings, agents are incentivized to learn about opponents’ values to understand what they should learn about their own. The conclusions of both papers are quite negative, essentially stating that agents can only play in dictatorial games to satisfy the relevant belief-related notion in simplicity. In contrast, the simplicity criterion advanced here allows for nontrivial mechanisms to admit equilibria with respect to many realistic

signal structures.

Methodologically, this paper relates to the robust implementation literature in its relaxation of the common prior assumption and modeling of belief structures. The microfoundation of my solution concept mirrors that of robust equilibrium discussed in Bergemann and Morris [2005]. Additionally, while not as stark, my results resemble those of Jehiel et al. [2006] who show that ex post equilibria generically do not exist in multidimensional settings. ? overturns this result in important contexts, showing that these equilibria exist commonly in multidimensional auction settings. My result, then, can be viewed as again constraining the set of multidimensional mechanisms that admit robust equilibria, though I view this result positively as it narrows the designer’s search problem. Lastly, Yamashita and Zhu [2022] show that ex post implementable mechanisms are optimal when the ordering of types’ allocations is independent of the realization of a state. Though I do not study optimality, the fixed-order property in my work closely resembles this condition.

Viewing ex post implementation as overly demanding, this paper adds to the growing literature on belief restrictions in mechanism design. My approach relates to three natural ways in which this has been done. The first, developed by Lopomo et al. [2021] considers local, rather than global, forms of robustness to distributional uncertainty. My approach in modeling convex information sets proceeds in a similar fashion. However, I obtain markedly different results because I make less demanding assumptions on the richness of this local robustness. The second assumes agents’ beliefs are defined by a finite number of moments. Brooks and Du [2023] and Gretschko and Mass [2024] study robust versions of the optimal auction design problem where the auctioneer’s information about agents’ values is captured by moments. In modeling moment-based signal structures, I take a similar approach but apply these considerations to the agents’ problem. In this way my approach more closely resembles Ollár and Penta [2017], who consider moment-based belief restrictions in studying full implementation of allocation rules³ Lastly, belief restrictions can be modeled through the assumption that agents learn through observing data from a public dataset. Camara [2022] and Liang [2021] both study mechanism design under this restriction though with different aims. In modeling history-dependent signal structures, my approach resembles these authors’ ideas though importantly, I abstract from noise and other practical considerations stemming from finite data.

Finally, in terms of application, this paper contributes to the literature on the use of a type of public procurement auction called a first-score auction. The most well-known scoring auction papers such as Che [1993] and Asker and Cantillon [2008] study simple,

³The concept studied by Ollár and Penta [2017] relates closely to the idea of Δ -rationalizability defined in Battigalli and Siniscalchi [2003].

quasilinear scoring auctions and model sellers with simple cost functions. Later work, such as Wang and Liu [2014], Dastidar [2014], and Hanazono et al. [2016] characterize equilibria in general scoring auctions with more general cost structures and multidimensional types. My model considers a middle-ground setting with general auctions rules but simple, additively separable cost structures. This allows for a nontrivial characterization of strategies that generalizes the familiar strategies of one-dimensional auctions. Lastly, I contribute to the empirical literature on estimating type distributions from auction data, extending ideas on inverting score distributions in Guerre et al. [2000], Athey and Haile [2002], and Hanazono et al. [2016]. I show that the equilibria of certain auctions are more stable than others which may imply better identification of underlying type distributions in the former settings.

2 Model

2.1 The Designer

A buyer runs a procurement auction in which she cares about both the quality of a good she receives, $q \in [0, 1]$, and the price she has to pay for it, p . To obtain this good, she runs a type of auction called a first-score auction.⁴ To define this, I first define a scoring rule:

Definition 1. *A scoring rule Φ is a smooth function from $\mathcal{C} := \mathbb{R}_+ \times [0, 1] \rightarrow \mathbb{R}$ that maps $(p, q) \mapsto \Phi(p, q)$ and is strictly increasing in q and strictly decreasing in p .*

Definition 2. *A **first-score auction** is defined as follows:*

1. *The buyer announces a **scoring rule** Φ .*
2. *The participants submit bids which are proposed contracts of asking price and good quality, (p, q) .*
3. *The winner is the participant who submitted the highest-scoring bid, evaluated by Φ . He is asked to supply the contract that he bid.*

First-score auctions are the generalization of first-price auctions to multidimensional settings; a scoring rule is used to map multidimensional contracts into one-dimensional values called scores and then an analogue of a first-price auction is run on the scores. These auctions are the most common allocation mechanism used in this type of procurement. Examples of scoring rules used in practice are the following:⁵

⁴These auctions are generally conducted in a sealed-bid static format. In the United States, this format is required by the Federal Code of Regulations. See Federal Acquisition Regulatory Council [1983].

⁵See Bergman and Lundberg [2013] and Hanazono et al. [2016] for more details on where these rules are used and of other scoring rules used in practice.

1. **quasilinear** scoring rules — $\Phi(p, q) = \phi(q) - p$ for some concave function $\phi : [0, 1] \rightarrow \mathbb{R}$.
2. **price-quality ratio (PQR)** scoring rule — $\Phi(p, q) = -\frac{p}{q}$,
3. **quality discount (QD)** scoring rules — $\Phi(p, q) = -p(\bar{Q} - q)$, for some $\bar{Q} > 1$.

In the first two examples, the buyer’s choice of scoring rule reflects a cost-benefit and a cost efficiency objective, respectively. Though the third case has no natural interpretation as a governmental priority, it is still easily comprehensible to participants: sellers will be ranked on their asking price with greater discounts given to them for greater provisions of quality.

2.2 Participants

There are 2 sellers in the auction, each with private information about his cost structure. I denote sellers’ types by $(m, f) \in \Theta := [\underline{m}, \bar{m}] \times [\underline{f}, \bar{f}] \subset \mathbb{R}_{++}^2$. The cost of producing the good with quality level $q \in [0, 1]$ takes an additively separable form, equal to

$$m \cdot q^\eta + f$$

for some convexity parameter $\eta \in [1, \infty)$. In words, when $\eta = 1$, the components of a participant’s type represents his marginal cost for providing quality and fixed cost for providing the lowest-quality version of the good respectively. I denote a general element of Θ by θ .

Both agents’ types are drawn i.i.d. according to a type distribution $g_0 \in \mathcal{G}$ where \mathcal{G} is the space of strictly positive, square-integrable density functions over Θ . Formally,

$$\mathcal{G} = \{g \in \mathcal{L}^2(\Theta, \mu) : g > 0, \int g(\tau) d\mu = 1\}$$

where μ is the Lebesgue measure over Θ . The agents do not know g_0 nor do they have a prior over \mathcal{G} . After learning their types and before participating in the auction, agents can acquire information about the type distribution. More specifically, there is a common knowledge signal structure $I : \Theta \rightrightarrows \mathcal{G}$ that maps each type to a partition of the state space. Each agent of type θ privately learns the set $E \in I(\theta)$ with $g_0 \in E$. The two types of signal structures I study for the majority of this paper are the following:

Example 1. *Agents acquire information coarsely by learning the values of a finite number of (generalized) moments of the type distribution. This approach is frequently taken in the robust mechanism design literature, as seen in Brooks and Du [2023] and Gretschko and Mass [2024] for example.*

Definition 3. Let $M : \mathcal{G} \rightarrow \mathbb{R}$ be defined by

$$M(g) = \mathbb{E}_{\theta_{-i} \sim g}[\zeta(\theta_{-i})]$$

for some bounded measurable function $\zeta : \Theta \rightarrow \mathbb{R}$. This function M is called a **moment**. Adopting standard notation, I denote the space of these moments by \mathcal{G}^* .

Under a **finite-moment signal structure**, each agent of type θ learns the realizations of the moments $M_1^\theta, \dots, M_n^\theta \in \mathcal{G}^*$ at the true type distribution, and restricts his beliefs to states that could provide such realizations. In other words, for any $E \in I(\theta)$,

$$M_k^\theta(g_1) = M_k^\theta(g_2)$$

for all $k \in \{1, \dots, n\}$ and $g_1, g_2 \in E$.

The set of such signal structures is denoted as \mathcal{I}^{FM} .

Example 2. Agents acquire information by reviewing historical data about previous iterations of the auction. Let there be a countable number of periods in which an auction with the same scoring rule has occurred. In period t , there are n_t short-lived bidders whose types are drawn i.i.d from the distribution g_0 . We assume the number of bidders is bounded ($1 < n_t < N$ for some $N \in \mathbb{R}_+$). In each period, some publicly available data is revealed:

$$d_t \subset \left\{ (n_t, \phi_{(1)}^t), (n_t, p_{(1)}^t, q_{(1)}^t) \right\}$$

where n_t is the realization of the number of bidders and $\phi_{(1)}^t$ and $(p_{(1)}^t, q_{(1)}^t)$ are random variables denoting the score bid and contract bid respectively of the winning bidder. Each agent learns the full history of publicly revealed information:

$$H = (d_t)_{t \in \mathbb{N}}$$

The data available in each history and a sequence of $(n_t)_{t \in \mathbb{N}}$ define a signal structure. Realizations of the results of the auction, that is realizations $\phi_{(1)}^t$ or $(p_{(1)}^t, q_{(1)}^t)$ for each $t \in \mathbb{N}$ correspond to a signal. Under a **history-dependent signal structure**, each bidder restricts his beliefs to states that could provide such realizations under the assumption that agents in the history played Bayes-Nash equilibrium strategies with the true distribution as a common prior. The set of such signal structures is denoted by \mathcal{I}^H .

A similar approach to agent learning under this type of signal structure is taken in estimating bidders' values in the industrial organization literature. (See Guerre et al. [2000], Athey and Haile [2002], and Hanazono et al. [2016]) for example.

I study the signal structures in \mathcal{I}^{FM} and \mathcal{I}^H for most of this paper but discuss more general ones after characterizing equilibrium existence for finite-moment and history-dependent signal structures.

2.3 Strategies in Coarse Beliefs

I start by discussing the familiar notion of Bayes-Nash equilibrium in this setting before defining *coarse beliefs equilibrium*, the primary solution concept of this paper. Under the following mild assumptions adapted from Hanazono et al. [2016], there is a unique, symmetric Bayes-Nash equilibrium in the game in which agents share a correctly-specified common prior, g :

Assumption 1 (Regularity). *Fix a scoring rule Φ . Let $u(s, \theta)$ be the indirect (ex post) utility type θ receives from bidding score s and winning. Let $P(s, q)$ be the price needed to obtain a score s when bidding a quality q under this scoring rule.*

- (a) *For all s , the function $\frac{\partial}{\partial m} \left(\frac{u(s, \theta)}{u_1(s, \theta)} \right)$ is continuous in θ and is strictly negative.*
- (b) *For all $s \in \text{Im } \Phi$, we have that $P(s, q) < \infty$ for all $q \in [0, 1]$. Further, we have that $P(s, q)$ is strictly convex in q , with smooth, strictly increasing derivative. Additionally, $P_q(s, 0) < -\bar{m}$ and $P_q(s, 1) > -\underline{m}$.*

Assumption 1(a) guarantees a form of single-crossing holds in this multidimensional setting and that pure-strategy equilibria will always exist as a result. Further, it tells us the unique distribution of scores bid in equilibrium has a density function, denoted $h_g(s)$. Assumption 1(b) guarantees that any type's optimal contract bid is unique.

I denote equilibrium strategies by $(p^*, q^*) : \Theta \times \mathcal{G} \rightarrow \mathcal{C}$, Mathematically,

$$(p^*, q^*)(\theta, g) \in \arg \sup_{p, q \in \mathcal{C}} \mathbb{P}_{\theta_{-i} \sim g}[\Phi(p, q) \geq \Phi(p^*, q^*)(\theta_{-i}, g)] \cdot (p - m \cdot q^n - f)$$

Note the dependence of strategies on the type distribution g in this notation. In standard mechanism design models, this distribution is fixed and so dependence of strategies on it is suppressed. We would like to consider settings in which agents can play symmetric equilibrium strategies without fully knowing the type distribution at all $g \in \mathcal{G}$.

Fixing an agent, under any signal structure that is not fully revealing, there are some states at which he will not fully learn the type distribution; after learning, he will retain some uncertainty about the true state. The solution concept introduced in this paper corresponds to cases in which this uncertainty is strategically irrelevant.

Definition 4. An auction has a **coarse beliefs equilibrium (CBE)** with respect to the signal structure I if there exist strategies $(p^I, q^I) : \Theta \times I(\theta) \rightarrow \mathcal{C}$ such that for all $\theta \in \Theta$ and $E \in I(\theta)$,

$$(p^I, q^I)(\theta, E) \in \arg \sup_{p, q \in \mathcal{C}} \mathbb{P}_{\theta_{-i} \sim g} [\Phi(p, q) \geq \Phi(p^I, q^I)(\theta_{-i}, I(\theta_{-i}, g))] \cdot (p - m \cdot q^n - f) \quad \forall g \in E$$

where $I(\theta_{-i}, g)$ is the element of $I(\theta_{-i})$ that contains g .

For each type of agent, at every signal realization, there is an action that is optimal no matter what the true state actually is. Learning the signal gives agents enough information about the state to strategize optimally. In Appendix D, this solution concept is formally microfounded as a robust equilibrium of a game in which each agent has no prior over the state and uses his signal realization to discipline his set of possible beliefs.

Coarse beliefs implementation can be thought of as an intermediate between Bayes-Nash equilibrium and ex post equilibrium in both a literal and philosophical sense. If I is fully-revealing, then this solution concept corresponds to Bayes-Nash equilibrium; If I reveals no information, then it corresponds to ex post equilibrium. In terms of interpretation, coarse beliefs implementation mirrors the concerns relating to robustness of equilibrium predictions that motivates the study of ex post equilibrium. However, ex post equilibrium does not exist in many settings. By allowing for agents to have some information, coarse beliefs implementation allows for greater existence of robust equilibria. The following proposition characterizes when coarse beliefs equilibria can exist:

Proposition 1. A coarse beliefs equilibrium exists with respect to the signal structure I if and only if

$$(p^I, q^I)(\theta, E) = (p^*, q^*)(\theta, g_1) = (p^*, q^*)(\theta, g_2)$$

for all $E \in I(\theta)$ and $g_1, g_2 \in E$.

It is immediate from Proposition 1 that familiar auctions admit coarse beliefs equilibria with respect to moment-based information technologies:

Example 3. Consider a first-price auction between 2 participants. They have private values $v \in [0, 1]$ drawn i.i.d from a distribution with density $g \in \mathcal{G}$, have quasilinear utility functions, and are risk-neutral expected profit maximizers. Let each agent learn from a finite-moment signal structure in which the type v learns

$$M_1(g) = \int x \cdot \mathbb{I}_{x \leq v} g(x) dx \quad \text{and} \quad M_2(g) = \int x \cdot \mathbb{I}_{x \leq v} g(x) dx$$

Bayes-Nash equilibrium strategies for both players are to submit bids $b^* : [0, 1] \times \mathcal{G}$ defined by

$$b^*(v, g) = \frac{\int x \cdot \mathbb{I}_{x \leq v} g(x) dx}{\int \mathbb{I}_{x \leq v} g(x) dx} = \frac{M_1(g)}{M_2(g)}. \quad (1)$$

By definition, $M_i(g') = M_i(g)$ for all g' in the same information set as g . Thus, this game admits a coarse beliefs equilibrium with respect to this signal structure.

Similarly, second-price and all-pay auctions in which agents possess one-dimension of private information admit coarse beliefs equilibria with respect to 1-moment information technologies. This suggests that agents' information about their environment does not need to be too precise to form adequate best-responses. These auctions additionally always admit coarse beliefs equilibria with respect to history-dependent technologies. Neither of these results will not hold generically in the multidimensional scoring auction setting.

In the following sections, I study the existence of coarse beliefs equilibria with respect to the signal structures defined above. For the sake of tractability, I make the following assumptions about strategies:

Assumption 2 (Differentiability). *Let $s^* = \Phi \circ (p^*, q^*) : \Theta \times \mathcal{G} \rightarrow \mathbb{R}$ be the score bidding strategy in equilibrium. Then for all types $\theta \in \Theta$ and $g \in \mathcal{G}$, we have that $s^*(\theta, g)$ is differentiable with respect to fixed costs.*

3 Characterization

In this section, I provide an interpretable condition for when a first-score auction admits a coarse beliefs equilibrium with respect to a signal structure. I then define a property called the *fixed-order property* and show that it is necessary for auctions to admit CBE's with respect to canonical signal structures:

Theorem 1. *A first-score auction can be implemented in coarse beliefs with respect to some n -moment signal structure if and only if it satisfies the fixed-order property. In such cases, it is sufficient for agents to learn just 2 moments of the type distribution. When this property is not satisfied, no finite number of moments is sufficient for agents to form equilibrium strategies.*

Theorem 2. *A first-score auction can be implemented in coarse beliefs with respect to all history-dependent signal structures if and only if it satisfies the fixed-order property.*

In Section 3.1, I characterize equilibrium strategies in a first-score auction for general scoring rules. In Section 3.2, I prove the results above and additionally show that auctions

that satisfy the fixed-order property admit CBE's with respect to the largest number of general information structures.

3.1 First-Score Strategizing

I start by describing Bayes-Nash equilibrium strategies in general first-score auctions. To discuss these strategies, I first define second-score auctions, the generalization of second-price auctions to this multidimensional setting. After this, I define an equivalence relation on the set of types, calling two types equivalent if they bid the same scores in equilibrium. With these two components, I then characterize Bayes-Nash strategies.

In one dimension, studying second-price auctions helps us understand strategies in a first-price auction. An analogous result holds here.

Definition 5. A *second-score auction* is defined as follows:

1. The buyer announces a **scoring rule** Φ .
2. The participants submit bids which are proposed contracts of asking price and quality, (p, q) .
3. The winner is the participant who submitted the highest-scoring bid, evaluated by Φ . He is asked to supply a contract whose score equals that of the second-highest bid.

I see second-score auctions far less frequently than first-score auctions in practice. Nevertheless, these auctions remain useful as analytical tools. It is important to note that, in a second-score auction, the winner does not need to provide the second-highest scoring contract that was submitted; they can provide a contract that is optimal for them given that it attains the requisite score. This means that if two types have very different cost structures—e.g. one (the other) has a high (low) marginal cost and has a low (high) fixed cost—the distance between their types and the difference between their resulting bids does not penalize either of them. Because of this feature, the incentives in a second-score auction resemble those of a second-price auction. As a result, their equilibria have similar properties:

Proposition 2. (Hanazono et al. [2016]) A second-score auction with scoring rule Φ that satisfies Assumption 1 has a unique equilibrium in weakly dominant strategies. These strategies, $(p_{BE}, q_{BE}) : \Theta \rightarrow \mathcal{C}$ represent the highest-scoring contracts a type can provide while still obtaining a weakly positive profit.

Because scoring rules are strictly increasing in price, a type obtains exactly zero profit from the contract he bids in a second-score auction. For this reason, I refer to these bids as

break-even (BE) contracts. As discussed earlier, second-score auctions and break-even contracts are useful for describing strategies in first-score auctions, as in the one-dimensional case. However, unlike in the one-dimensional case, this result is not due to revenue equivalence. In general, revenue equivalence does not hold in this setting because the equilibria in the two auctions do not always have the same allocation rules. To discuss the allocation rule of a first-score auction, I define a notion of an equilibrium order over the set of types.

Definition 6. *At the type distribution g , an equilibrium of a scoring auction **implements the order** \succeq_g over Θ where*

$$s^*(\theta_1, g) \geq s^*(\theta_2, g) \iff \theta_1 \succeq_g \theta_2.$$

In words, one type is greater than another at a distribution g if, in the equilibrium at that distribution, the former type wins the contract over the latter. The relation \succeq_g is a total order as the transitivity and completeness of the order used to compare scores extend to the order \succeq_g . I sometimes use the phrases equilibrium structure or equilibrium order to refer to the type ordering defined by \succeq_g . In a second-score auction, this equilibrium order is the same for all $g \in \mathcal{G}$. This is generically not the case in a first-score auction.

This order induces an equivalence relation over the set of types that bid the same score. I denote this relation by \sim_g . The equivalence classes of this relation are analogous to one-dimensional “pseudotypes” discussed in the scoring auction literature.⁶ As in those contexts, a distribution over 1-dimensional pseudotypes rather than that over 2-dimensional types is all that is necessary for describing strategies.

Proposition 3. *Fix $g \in \mathcal{G}$. For all $(m, f) \in \Theta$, this type’s equilibrium strategy can be represented as*

$$(p^*, q^*)(m, f, g) = (p_{BE}, q_{BE})(m, f_{(2)})$$

where

$$f_{(2)} = \frac{\int z \cdot \mathbb{1}_{z \geq f} g_m([(m, z)]_g) dz}{\int \mathbb{1}_{z \geq f} g_m([(m, z)]_g) dz}. \quad (2)$$

Here, $g_m([(m, z)]_g)$ refers to the pushforward density of g onto the set of equivalence classes when parametrized by the types with marginal cost m .

Proposition 3 reduces strategizing in a first-score auction into choosing a type to imitate and then simply bidding that type’s break-even contract. The type to mimic can be interpreted as the expectation of the second-highest bidder’s type under the pushforward

⁶Asker and Cantillon [2008] discuss pseudotypes in quasilinear scoring auctions while Hanazono et al. [2016] define pseudotypes in general first-score settings.

distribution. This decomposition of first-score auction generalizes strategies used in one-dimensional, first-price settings. In those auctions, the Bayes-Nash strategy is to bid the break-even bid of a different type, where this break-even bid is simply that type’s value. The type to imitate can be expressed as an expected conditional order statistic in Equation (1) which resembles that in Equation (2).

There are two key observations at work in proving Proposition 3. The first results from the additive separability in firms’ cost structures: conditional on bidding a score s , agents with the same marginal cost pick the same contract because their fixed cost does not effect their maximization problem. Mathematically, if $h_g(s)$ is the density of the score distribution in equilibrium,

$$\arg \max_{(p,q) \in C(s)} h_g(s) \cdot (p - m \cdot q^n - f_1) = \arg \max_{(p,q) \in C(s)} h_g(s) \cdot (p - m \cdot q^n - f_2).$$

where

$$C(s) = \{(p, q) \in \mathcal{C} : \Phi(p, q) = s\}.$$

This immediately allows the description of a general bid as some other type’s break-even bid. Deriving an expression for this other type requires an additional observation, illustrated in Figure 1. Intuitively, if types θ_1 and θ_2 are equivalent, then from the perspective of one agent, it does not matter whether their opponent has type θ_1 or θ_2 . More formally, equilibrium strategies best-respond to an equilibrium distribution of scores, $H_g(s)$. Moving mass between types in the same equivalence class, then, does not change any type’s equilibrium strategies because this distribution remains unchanged. Thus, to study strategies, it is without loss to move all mass in an equivalence class to a type that has some fixed marginal cost and then consider the resulting auction in which an agent’s only private information is his fixed cost.

3.2 CBE Existence

Returning to the existence of coarse beliefs equilibrium, Proposition 1 tells us that CBE exist if and only if the strategy in Proposition 3 can be expressed as a function of the signal realization. Equivalently, agents can strategize with coarse information if the “imitation type” in Equation 2 is constant for all type distributions in an information set. I begin this section by adapting notation from Section 3.1 to discuss this.

First, fixing a marginal cost $m \in [\underline{m}, \bar{m}]$, I define the projection function induced by the equivalence relation \sim_g :

$$\rho_m(\tau, g) : \Theta \rightarrow \mathbb{R}$$

$$\theta \mapsto f \text{ such that } (m, f) \in [\theta]_g.$$




Figure 1: On the left are three equivalence classes of types in an equilibrium. To any type of agent, the distribution of density within these equivalence classes does not matter. He only cares about the aggregate density of types in that class who bid the same score. Thus, it is without loss to consider a setting in which all of this mass is pushed to one point, as shown on the right. This then reduces the problem to a one-dimensional first-price auction.

While the notation in Equation 2 is useful for showing the relationship between equilibrium strategies of first-price and first-score auctions, this current notation is more useful for analyzing these auctions as the type distribution varies. With it, I can write that

$$\int \mathbb{1}_{z \geq f} g_m([m, z]_g) dz = \int \mathbb{1}_{\tau <_g \theta} g(\tau) d\tau \text{ and } \int z \cdot \mathbb{1}_{z \geq f} g_m([m, z]_g) dz = \int \rho_m(\tau, g) \mathbb{1}_{\tau <_g \theta} g(\tau) d\tau.$$


Using this notation, I define

$$f_{(2)}(m, f, g) := \frac{\int \rho_m(\tau, g) \mathbb{1}_{\tau <_g \theta} g(\tau) d\tau}{\int \mathbb{1}_{\tau <_g \theta} g(\tau) d\tau}.$$

The requirement discussed above, that the imitation type $f_{(2)}(\theta, g)$ must be the same for all $g \in E$, means that for a type to bid the same at two different distributions, the average “competitiveness” of the types he faces should be the same. Because each signal realization E is a convex set when $I \in \mathcal{I}^{FM} \cup \mathcal{I}^H$ and strategies are differentiable, this intuition must hold locally. This yields the following result:

Lemma 1 (Equal Competitiveness). *A first-score auction admits a coarse beliefs equilibrium with respect to $I \in \mathcal{I}^{FM} \cup \mathcal{I}^H$ if and only if for all $\theta \in \Theta$ and $E \in I(\theta)$,*

$$g_1, g_2 \in E \implies \int (\rho(\tau, g_1) - \psi(\theta, g_1)) \cdot \mathbb{1}_{\tau <_{g_1} \theta} \cdot (g_2(\tau) - g_1(\tau)) d\tau = 0 \quad (3)$$



(a) Let the equilibrium at the distribution g induce the solid black iso-score lines depicted. Consider a perturbation of g that shifts mass from the red to the blue types.

(b) This perturbation shifts mass from lower-scoring g -equivalence classes to higher-scoring ones. This increases the “average competitiveness” of the auction.

Figure 2: These panels depict a violation of the Equal Competitiveness Lemma. No distribution in the direction of this shift from g can be in the same information set as g .

where

$$\rho(\tau, g) := \frac{\int_{\underline{m}}^{\overline{m}} \rho_m(\tau, g) dm}{\overline{m} - \underline{m}} \text{ and } \psi(\theta, g) := \frac{\int \rho(\tau, g) \mathbb{1}_{\tau <_g \theta} g(\tau) d\tau}{\int \mathbb{1}_{\tau <_g \theta} g(\tau) d\tau}.$$

In general, the difficulty in studying the imitation type $f_{(2)}(\theta, g)$ at different states $g \in \mathcal{G}$ is that the equilibrium order, and thus the function $\rho_m(\cdot, g)$ can change arbitrarily with g . The Equal Competitiveness Lemma essentially states that locally, and on average, these changes are not first-order. As a result, a set of complicated equalities to check whether a CBE exists can be reduced to a set of simpler moment conditions that the functions $g_2 - g_1$ must satisfy for all $g_1, g_2 \in \mathcal{G}$. I call these functions $g_2 - g_1$ directions and let the set of directions \mathcal{G}_0 be defined as

$$\mathcal{G}_0 = \{v \in \mathcal{L}^2(\Theta, \mu) : \int v(\tau) d\tau = 0\}.$$

Figure 2 graphically shows the implication of violations of Lemma 1.

The Equal Competitiveness Lemma tells us that there are strong requirements for the existence of CBE. If an information set E is n -dimensional, this lemma gives a $2n$ -dimensional set of equations that must be satisfied in a CBE. This immediately gives three general insights into when CBE are more likely to exist:

1. **Each information set is small.** Larger information sets lead to both more con-

straints that the directions in E must satisfy and a larger set of directions that must satisfy these constraints.

2. **Each information set is specifically tailored to the auction.** Each distribution $g \in E$ defines both new directions and new moments (indexed by $\rho(\cdot, g)$) in the set of constraints that must be satisfied. Thus, the set E must include distributions that are similar not only from an ex ante perspective but also in the sense that their equilibria in the given first-score auction are related.
3. **The auction’s equilibrium structure should change in nice ways.** If the set of functions $\{\rho(\cdot, g) : g \in E\}$ is higher-dimensional, then fewer sets of directions can satisfy the constraints posed by Lemma 1 for all such functions.

The cases of finite-moment and history-dependent signal structures will show respectively that when information sets are large or not highly relevant, the designer is unable to implement any CBE unless the equilibrium structure does not change at all. Specifically, in such cases, these equilibria only exist when the set $\{\rho(\cdot, g) : g \in \mathcal{G}\}$ is a singleton. Because of its importance, I define this property formally below:

Definition 7. *A first-score auction satisfies **the fixed-order property** if for all $\theta_1, \theta_2 \in \Theta$, if θ_1 is awarded the contract over θ_2 in equilibrium at some distribution $g \in \mathcal{G}$ then θ_1 is awarded the contract over θ_2 at all distributions in \mathcal{G} . Mathematically, a first-score auction satisfies the fixed-order property if for all $g_1, g_2 \in \mathcal{G}$,*

$$\theta_1 \succeq_{g_1} \theta_2 \iff \theta_1 \succeq_{g_2} \theta_2.$$

Aside from the current discussion, the fixed-order property is naturally appealing; the fixed-order property indicates that there is a concrete notion of one type being a “stronger bidder than” another. Viewing an allocation function as a choice rule, this property relates to an Independence of Irrelevant Alternatives condition in a social choice setting or a substitutability condition in a matching one because the designer’s preference of one type over another does not depend on any other factors. Additionally, many common auction formats satisfy this property. When there is only one dimension of private information, there is a canonical total order over the set of types. Popular auction formats—like first-price, second-price, and all-pay auctions—implement allocation rules that are monotonic with respect to this order. This definition of monotonicity is independent of the type distribution so the fixed-order property is satisfied. With multidimensional private information, there is no canonical total order over type spaces so this property is not readily implied by monotonicity

requirements. Thus, multidimensional auctions satisfying the fixed-order property generalize the nice structure of equilibria in one-dimensional settings.

3.2.1 Finite-Moment Signal Structures

Theorem 1. *A first-score auction can be implemented in coarse beliefs with respect to some n -moment signal structure for $n < \infty$ if and only if it satisfies the fixed-order property. In the affirmative case, this auction has a coarse beliefs equilibrium with respect to a 2-moment signal structure.*

With this type of signal structure, agents generally cannot learn enough about the equilibrium type ordering to be able to strategize properly. Thus, the only auctions in which they are able to play coarse strategies are ones in which they do not need to learn this equilibrium structure as it is always fixed. When this holds, each of the expressions in Equation (2) is a moment; thus a 2-moment signal structure suffices.

To see this mathematically, consider an n -moment signal structure in which type θ 's information partition is defined by

$$M_k(g) = \int \zeta_k(\tau) g(\tau) d\tau$$


for $k \in \{1, \dots, n\}$. By the definition of this signal structure, if $v \in G_0$ satisfies

$$M_k(v) = 0 \quad \forall k \in \{1, \dots, n\}$$

then the direction v preserves information sets. That is, for any $g \in \mathcal{G}$, if $g \in E$, then $g + v \in E$. Applying the Equal Competitiveness Lemma to this information set and these two distributions, we have that

$$M_k(v) = 0 \quad \forall k \in \{1, \dots, n\} \implies \int (\rho(\tau, g) - \psi(\theta, g)) \cdot \mathbb{1}_{\tau \prec_g \theta} \cdot (v(\tau)) d\tau = 0 \quad \forall g \in \mathcal{G}.$$

The set of directions $v \in \mathcal{G}_0$ that satisfy the left-hand side of this implication is large, as it has finite codimension n . Thus, the set of directions satisfying Equation (3) must be similarly large. This implies that the set of functions $\{\mathbb{1}_{\tau \prec_g \theta} : g \in \mathcal{G}\}$ must be finite-dimensional, which happens if and only if the fixed-order property holds.



(a) This panel depicts the equilibrium iso-score lines when there are 3 participants. The red types bid more weakly than the corresponding blue types do when there is less competition.

(b) In weaker competition, the red types bid more weakly than the corresponding blue ones. Iso-score lines if the red (blue) types are playing are dotted in red (blue).

Figure 3: Coarse Beliefs Equilibria do not always exist when agents learn from history-dependent signal structures.

3.2.2 History-Dependent Signal Structures

Theorem 2. *A first-score auction can be implemented in coarse beliefs with respect to all history-dependent signal structures if and only if it satisfies the fixed-order property.*

With a history-dependent signal structure, agents learn significantly more information than they do with moment-dependent ones. Instead of solely knowing a finite number of moments, agents generally can learn the full equilibrium distribution of scores for some number of participants. However, if this number of participants is not equal to 2, then generically, this information that agents learn does not relate closely enough to the current setting for them to be able to strategize correctly.

As an example, consider a history at which $d_t = (3, \phi_{(1)}^t)$ for all $t \in \mathbb{N}$. Bidders can only strategize in the current auction if the equilibrium score distribution when there are 3 bidders reveals enough information about that in the 2-bidder setting. This fails unless an auction satisfies the fixed-order property. Figure 3 illustrates the intuition behind this result. Let the red circles and corresponding blue circles in Figure 3(a) denote types that bid the same score in the 3-bidder setting. Further, let the red types “bid less competitively” at

weaker distributions than the blue types. That is, for low scores, the ratio

$$\frac{u_1(s, \theta)}{u(s, \theta)}$$

is lower for red types than their corresponding blue types. Thus, the former have a greater incentive to increase their bids at these scores than blue types in the weaker, 2-bidder setting. Figure 3(b) shows the different equilibria structures in the 2-bidder setting that can result depending on whether the other bidder is a red type or a blue type. Thus, information about whether an opponent is a red or blue type meaningfully changes an agent’s strategy; however, this information is not contained in the data.

More formally, recalling the proof of Proposition 4, moving mass along the iso-score lines in 3(a) does not change the equilibrium score distribution in the 3-bidder setting. Thus, any distributions that differ by such changes are contained within the same information set. For a coarse beliefs equilibrium to exist, these mass movements must also not change bidding incentives in the 2-bidder setting. Thus, for any type distribution, the equilibrium iso-score lines in the 2-bidder setting must be the same as those in the 3-bidder setting. This only holds when the fixed-order property is satisfied.

3.2.3 General CBE Existence

As alluded to above, the importance of the fixed-order property in allowing for more CBE pertains to general signal structures, not solely those discussed above. To see this, we start by defining a space of admissible signal structures:

Definition 8. (a) A signal structure I is **coarse** if for all $\theta \in \Theta$, there exists $E \in I(\theta)$ such that

$$|E| > 1$$

(b) A signal structure I is **convex** if for all $\theta \in \Theta$ and $E \in I(\theta)$,

$$g_1, g_2 \in E \implies \lambda g_1 + (1 - \lambda)g_2 \in E$$

for any $\lambda \in [0, 1]$.

Let the set of coarse, convex signal structures as \mathbb{I} .

Coarseness is a weak requirement; if it is not satisfied, then the signal structure is fully revealing. In such cases, this environment reduces to the one in which g_0 is a common prior. Convexity is also a natural assumption as if an agent cannot distinguish g_1 from g_2 , then he should also not be able to distinguish g_1 from a state that “resembles” g_1 more than g_2 does.

All finite-moment signal structures are coarse and convex. Though all history-dependent signal structures are convex, not all of them are coarse as a history may provide sufficient information for agents to learn the type distribution entirely.

For a scoring rule Φ , let

$$\mathcal{I}(\Phi) = \left\{ I \in \mathbb{I} : \text{the auction with rule } \Phi \text{ admits a CBE w.r.t. } I \right\}.$$

I would like to compare these sets across different scoring rules and show that if Φ satisfies the fixed-order property, the size of this set is maximal. One approach would be to try to show that for any other scoring rule Φ' ,

$$\mathcal{I}(\Phi') \subset \mathcal{I}(\Phi).$$

However, this statement is in general not true. As discussed above, for a signal structure $I' \in \mathcal{I}(\Phi')$, the information contained in I' must be tailored to the scoring rule Φ' . Thus, unless the scoring rules Φ and Φ' are similar in some meaningful sense, $I' \in \mathcal{I}(\Phi)$. However, for any $I' \in \mathcal{I}(\Phi')$, there is an $I \in \mathcal{I}(\Phi)$ such that I is “as coarse as” I' . I formalize this notion through the following definitions:

Definition 9. A bijection $\alpha : \mathcal{G} \rightarrow \mathcal{G}$ is **an isomorphism between two partitions** Π_1, Π_2 of \mathcal{G} if for all sets $E \subset \Pi_1$, there exists some set denoted $\alpha(E) \in \Pi_2$ such that

$$g \in E \iff \alpha(g) \in \alpha(E).$$

Definition 10. An isomorphism $\alpha : \mathcal{G} \rightarrow \mathcal{G}$ between two partitions Π_1, Π_2 of \mathcal{G} is **convexity-preserving** if whenever $\lambda g_1 + (1 - \lambda)g_2 \in E$ for all $\lambda \in [0, 1]$, we have that

$$\lambda \alpha(g_1) + (1 - \lambda)\alpha(g_2) \in \alpha(E) \quad \text{for all } \lambda \in [0, 1].$$

Definition 11. Two signal structures $I_1, I_2 \in \mathbb{I}$ are **isomorphic** if for all types $\theta \in \Theta$, there exists a convexity-preserving isomorphism between the partitions $I_1(\theta)$ and $I_2(\theta)$.

The first part of the definition of isomorphisms specifies that one signal structure is as coarse as another. For every type and every information set of that type in one signal structure, there is an information set of the same cardinality in the other. The second part of the definition requires that some structure of the state space is not perturbed through this bijection. Additionally, in our case, the Equal Competitiveness Lemma implies that existence of coarse beliefs equilibria amounts to the satisfaction of linear moment conditions.

Thus, requiring that isomorphisms preserve the linearity within information sets is essential for CBE existence under one signal structure to relate to CBE existence under another.

With this definition, I can formalize the statement above. While in general $\mathcal{I}(\Phi')$ is not a subset of $\mathcal{I}(\Phi)$, it is isomorphic to a subset of $\mathcal{I}(\Phi)$ and this subset is strict whenever Φ' does not satisfy the fixed-order property:

Theorem 3. *Let the first-score auction with respect to Φ_1 satisfy the fixed-order property. Then, for any Φ_2 ,*

$$I \in \mathcal{I}(\Phi_2) \implies \exists I' \in \mathcal{I}(\Phi_1) \text{ with } I' \text{ isomorphic to } I.$$

Additionally, if Φ_2 does not satisfy the fixed-order property,

$$\exists I \in \mathcal{I}(\Phi_1) \text{ such that } I' \text{ isomorphic to } I \implies I' \notin \mathcal{I}(\Phi_2).$$

4 The Fixed-order Property

In the previous section, I formed an equivalence between mechanisms admitting CBE's with respect to common signal structures and those satisfying the fixed-order property. One consequence of this result is that it translates the definition of CBE's, a statement defined on action spaces, to one defined on type spaces. While thinking about strategies in terms of action spaces leads to more interpretability, doing so leads to more analytical difficulties than considering type spaces. With the fixed-order property, I can take a direct revelation approach and characterize CBE's using standard mechanism design tools. First, I define effort to be the convex transformation of quality, q^n , that appears in participants' cost structures. With this reparametrization, agents have linear utility structures. I then define a symmetric direct revelation mechanism $(x, y, t) : \Theta^2 \times \mathcal{G}$ where

- $x(\theta_1, \theta_2, g)$ is the probability θ_1 is awarded the contract over θ_2 at the distribution g ,
- $y(\theta_1, \theta_2, g)$ is the effort θ_1 provides conditional on beating θ_2 at the distribution g ,⁷
- $t(\theta_1, \theta_2, g)$ is the transfer θ_1 receives conditional on beating θ_2 at the distribution g ,

I use capital letters to denote interim allocations, writing

$$Z(\theta_1, g) = \mathbb{E}_{\theta_2 \sim g}[z(\theta_1, \theta_2, g)]$$

⁷In a first-score auction, this does not depend on θ_2 because of the chosen auction format. I use this notation to discuss second-score auctions as well so dependence on θ_2 is included here.

for $z \in \{x, y, t\}$.

In equilibrium at a specific distribution g , for all types $\theta = (m, f)$, I must have that the usual individual rationality and incentive compatibility constraints hold:

$$X(\theta, g)(T(\theta, g) - mY(\theta, g) - f) \geq X(\hat{\theta}, g)(T(\hat{\theta}, g) - mY(\hat{\theta}, g) - f) \quad \forall \hat{\theta} \in \Theta, \quad (\text{IC})$$

$$X(\theta, g)(T(\theta, g) - mY(\theta, g) - f) \geq 0. \quad (\text{IR})$$

The fixed-order property holds if and only if the ex post allocation rule is independent of the type distribution, or $x(\theta_1, \theta_2, g) =: x(\theta_1, \theta_2)$ for all $g \in \mathcal{G}$. This restricts the set of ex post allocation rules that can be implemented; some allocation rules can be implemented at some distribution in \mathcal{G} but do not yield an IC mechanism at other distributions. Considering this stronger form of incentive compatibility yields Theorem 4.

Theorem 4. *A first-score auction with scoring rule Φ satisfies the fixed-order property if and only if it implements the same interim outcomes as the second-score auction with the same scoring rule at all distributions $g \in \mathcal{G}$.*

With one dimension of private information, the Revenue Equivalence Theorem gives a quick derivation of the first-price auction’s 2-moment strategies. In this setting, when a first-score auction has 2-moment strategies, the appropriate generalization of the Revenue Equivalence Theorem provides the same guarantee. Additionally, Theorem 4 gives more interpretability of the expressions in Equations 2. Because these first-score auctions implement the same order as their correspond second-score auctions, the equilibrium structure is defined by the “break-even contract order” with respect to Φ .

Theorem 4 is not very difficult to prove. Informally, as the distribution of the opponent’s type gets more competitive—as it puts more density on types with lower costs—an agent’s probability of winning and his gains from shading his bid drop. Thus, his optimal bid approaches his break-even contract regardless of his type. As a result, the implemented equilibrium order approaches the break-even contract order. Then, because the equilibrium structure must be the same at all distributions, the break-even contract order must hold at all distributions.

Applying this theorem gives the following testable condition:

Proposition 4. *A first-score auction with scoring rule Φ admits a coarse beliefs equilibrium if and only if its break-even effort function, $e_{BE} := q_{BE}^n : \Theta \rightarrow [0, 1]$ is linear in fixed cost.*

Revenue Equivalence and Proposition 3 imply that for every type, taking break-even efforts and applying expectations must commute. A converse to Jensen's inequality tells us that this implies linearity. Importantly, this test does not require computing an equilibrium. Break-even strategies are easily verifiable given any scoring rule, as shown below:

Corollary 1. *First-score auctions with quasilinear scoring rules or price-per-quality ratio scoring rules admit CBE's. First-score auctions with a quality discount scoring rule do not.*

Proof. (a) In a first-score auction with a quasilinear scoring rule, e_{BE} minimizes

$$\phi(e^{1/\eta}) - p \text{ s.t. } p - e \cdot m - f \geq 0.$$

Taking a first-order condition gives

$$\phi'(e_{BE}^{1/\eta}) \cdot \left(\frac{1}{\eta} e_{BE}^{1/\eta-1} \right) - m = 0.$$

e_{BE} does not depend on fixed costs so it is linear in fixed cost (with coefficient 0).

(b) In a first-score auction with a PQR scoring rule and $\eta > 1$, e_{BE} minimizes

$$\frac{p}{e^{1/\eta}} \text{ s.t. } p - em - f \geq 0.$$

The minimand of this expression is

$$e_{BE}(m, f) = \frac{f}{(\eta - 1)m}$$

which is linear in f .

(c) In a first-score auction with a quality discount scoring rule, e_{BE} minimizes

$$p(1 - e^{\frac{1}{\eta}}) \text{ s.t. } p - em - f \geq 0.$$

The minimand of this expression satisfies

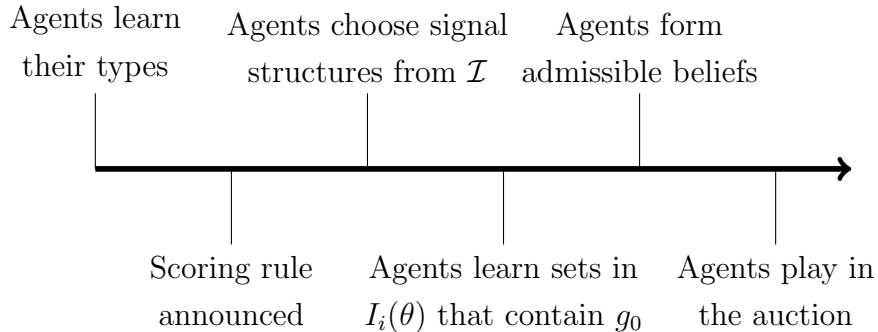
$$m = e_{BE}(m, f)^{\frac{1}{\eta}-1} \cdot \left[\left(1 + \frac{1}{\eta}\right) \cdot e_{BE}(m, f) \cdot m + \frac{f}{\eta} \right]$$

and is nonlinear in f .

□

5 Discussion

Endogenous Information Acquisition. The definition of coarse beliefs equilibria can be extended to one in which agents endogenously acquire information. In this model, prior-free agents participate in a game with two stages. In the first, they acquire information about the type distribution from a common knowledge information technology \mathcal{I} which is a set of signal structures. Familiar examples of information technologies include the finite-moment information technology, \mathcal{I}^{FM} , or the singleton, history-dependent signal structure $\{I^H\}$ for some history H . In the second stage of the game, the agents participate in a scoring auction. This model is detailed fully in Appendix D but is summarized in the figure below:




In general, the beliefs an agent forms before the auction will meaningfully affect their strategies; under one admissible belief, he should bid a certain contract and under another, he should bid differently. This is worrying because agents have no rational method of choosing between such beliefs. In other words, the coarseness of the learning process and the imprecision of beliefs generates a source of strategic confusion that the agents in this model are unable to resolve. As before, we look for equilibria in which this confusion is not generated.

Definition 12 (Informal). *A mechanism has a **Coarse Beliefs Equilibrium (CBE)** with respect to an information technology \mathcal{I} if for all types, at every realization of the moments they learn, there is an action that is optimal at every admissible belief.*

This definition of a coarse beliefs equilibrium coincides with that in which signal structures are exogenous. Thus, for example, auction formats that always admit CBE with respect to finite-moment or history-dependent information technologies are those that satisfy the fixed-order property. The only part of this model that is common knowledge is the information technology, or the type of learning process in which each agent can partake. The main result of Appendix D is that when they exist, CBE implement the same equilibria as when agents possess common prior. As famously argued by Hayek [1945], markets solve the

problem of limited and dispersed knowledge by providing incentives to individuals to report this information truthfully. Mechanisms that admit CBE perform a similar role in settings in which beliefs matter. They solve the problem of limited beliefs by providing incentives for agents to learn about their environment and then make a report dependent on both their payoff types and what they have learned. In a CBE, the auction mechanism aggregates this information and yields an equilibrium outcome.

Fully Informative Histories. The discussion of histories in Section 3.2.2 focused on those that generated coarse information structures. However, there are many histories, like those that publicly announce winning contracts in every period, that fully reveal the true type distribution. Even in such cases, it may be preferable for a designer to select an auction format that satisfies the fixed-order property. The models of signal structures discussed earlier are idealized; in the case of histories, the availability of an infinite amount of data allows for clean descriptions of agents' information. In practice, bidders can only obtain a finite amount of data. As the amount of data grows large, agents' strategies will converge more quickly when the fixed-order property is satisfied than when it is not. To see this, consider the steps taken by such a firm in strategizing:



When the fixed-order property is not satisfied, firms must estimate the underlying state. As discussed in Section 3.2.2, an agent's optimal bid can differ at two states that have the same probabilities of generating the equilibrium score distribution seen in the data. Thus, agents must use the additional information they have about winning contracts to estimate the underlying state before forming a bid. However, when the fixed-order property holds, this second estimation process is unnecessary; learning the equilibrium score distribution from the data is sufficient for strategizing in the current auction. In such cases, the noise that would otherwise be introduced in estimating the state is not generated. This leads to lower errors and more precise bidding for the same amount of information when the fixed-order property holds.

Functional Form Robustness. There may be settings in which the assumption of additive separability of the cost structure is unrealistic. In such cases, Proposition 4, the characterization of the fixed-order property through scoring rules, will no longer hold. However, results that concern auction structure and equilibrium existence, Theorems 3 and 4, do not not rely on the modeling of the cost structure here. Thus, the ideas of this paper can still be applied to guide design choices in such settings.

References

- John Asker and Estelle Cantillon. Properties of scoring auctions. *Rand Journal of Economics*, 39(1):69–85, 2008.
- John Asker and Estelle Cantillon. Procurement when price and quality matter. *Rand Journal of Economics*, 41(1):1–34, 2010.
- Susan Athey and Philip Haile. Identification of standard auction models. *Econometrica*, 70(6):2107—2140, 2002.
- Pierpaolo Battigalli and Marciano Siniscalchi. Rationalization and incomplete information. *Advances in Theoretical Economics*, 3(1):1073–1116, 2003.
- Dirk Bergemann and Steven Morris. Robust mechanism design. *Econometrica*, 73:1521–34, 2005.
- Mats A. Bergman and Sofia Lundberg. Tender evaluation and supplier selection methods in public procurement. *Journal of Purchasing and Supply Management*, 19(2):73–83, 2013.
- Tilman Börgers and Jiangtao Li. Strategically simple mechanisms. *Econometrica*, 87(6):2003–2035, 2019.
- Benjamin Brooks and Songzi Du. Maxmin auction design with known expected values. Working paper, 2023.
- Modibo Camara. Mechanism design with a common dataset. In *EC '22: Proceedings of the 23rd ACM Conference on Economics and Computation*, page 558. Association for Computing Machinery, 2022.
- Yeon-Koo Che. Design competition through multidimensional auctions. *Rand Journal of Economics*, 24(4):668 – 680, 1993.
- Krishnendu G. Dastidar. Scoring auctions with non-quasilinear scoring rules. Working paper, 2014.
- Federal Acquisition Regulatory Council. Sealed bidding. Code of Federal Regulations, 1983.
- Vitali Gretschko and Helene Mass. Worst-case equilibria in first-price auctions. *Theoretical Economics*, 19(1):61–93, 2024.

- Emmanuel Guerre, Isabelle Perrigne, and Quang Vuong. Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574, 2000.
- Makoto Hanazono, Yohsuke Hirose, Jun Nakabayashi, and Masanori Tsuruoka. Theory, identification, and estimation for scoring auctions. Working paper, 2016.
- John Harsanyi. Games with incomplete information played by bayesian play ers, i-iii. *Management Science*, 14:159–182, 320–334, 486–502, 1968.
- Friedrich. A. Hayek. The use of knowledge in society. *The American Economic Review*, 35(4):519–530, 1945.
- Phillippe Jehiel, Moritz Meyer-ter Vehn, Benny Moldovanu, and William Zame. The limits of ex post equilibrium. *Econometrica*, 74(3):585–610, 2006.
- Jiangtao Li and Piotr Dworzak. Are simple mechanisms optimal when agents are unsophisticated? Working paper, 2024.
- Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, 2017.
- Annie Liang. Games of incomplete information played by statisticians. Working paper, 2021.
- Giuseppe Lopomo, Luca Rigotti, and Chris Shannon. Uncertainty in mechanism design. Working paper, 2021.
- J-F. Mertens and S. Zamir. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29, 1985.
- Mariann Ollár and Antonio Penta. Full implementation and belief restrictions. *American Economic Review*, 107(8):2243–2277, 2017.
- Agathe Pernoud and Simon Gleyze. Informationally simple incentives. *Journal of Political Economy*, 131(3):802–837, 2023.
- Rudolf Verens. The state of implementing procurement procedures in eu agencies: enhancing transparency and assessing flexibility. Directorate-General for Internal Policies, 2021.
- Mingxi Wang and Shulin Liu. Equilibrium bids in practical multi-attribute auctions. *Economics Letters*, 123(3):352–355, 2014.
- Takuro Yamashita and Shugang Zhu. On the foundations of ex post incentive-compatible mechanisms. *American Economic Journal: Microeconomics*, 14(4):494–514, 2022.

A Proofs from Section 3.1

Throughout this appendix, it will be useful to think of the extended type space $\bar{\Theta} := [\underline{m}, \bar{m}] \times \mathbb{R}$ and extend all relevant definitions to this space. We do this because given some type $(m_1, f_1) \in \Theta$, it may be that the type in the same equilibrium equivalence class with marginal cost $m_2 \in [\underline{m}, \bar{m}]$ does not have fixed cost in $[\underline{f}, \bar{f}]$. We'll consider strategies of types in $\bar{\Theta}$ even though the only types with positive density are those in Θ .

In this section I prove Proposition 3 and its corollary. I start with a lemma that will be useful in the proof of Proposition 3.

Lemma 2. *Fix a type $\theta = (m, f) \in \Theta$. There is some type $(m, f_{(2)}) \in \bar{\Theta}$ such that $s^*(m, f) = s_{BE}(m, f_{(2)})$ where $s_{BE} = \Phi \circ (p_{BE}, q_{BE})$.*

Proof. Firstly, we have that the score bid by this type in equilibrium satisfies

$$s^*(m, f) \in [s_{BE}(m, \bar{f} + \bar{m}), s_{BE}(m, f)].$$

To see this, we have that $s^*(m, f) \leq s_{BE}(m, f)$ because bidding above this break-even score guarantees a strictly negative payoff conditional on winning the auction and a strictly positive probability of winning. A profitable deviation is to simply bid the break-even score. In equilibrium, the lowest-type agent in Θ will bid his break-even score and win with probability 0. Any other type will bid above this score to obtain a strictly positive expected payoff. Thus, it suffices to show that

$$s_{BE}(m, \bar{f} + \bar{m}) \leq s_{BE}(\bar{m}, \bar{f}).$$

If \bar{p}, \bar{q} is the break-even contract of the type $(m, \bar{f} + \bar{m})$, then because $\bar{q} \leq 1$,

$$0 \leq \bar{p} - m \cdot \bar{q}^n - (\bar{f} + \bar{m}) < \bar{p} - \bar{m} \cdot \bar{q}^n - \bar{f}$$

The type (\bar{m}, \bar{f}) can obtain the score $s_{BE}(m, \bar{f} + \bar{m})$ with a contract that provides positive payoff so $s_{BE}(m, \bar{f} + \bar{m}) \leq s_{BE}(\bar{m}, \bar{f})$.

Now, we have that the break-even score for a type (m, f) is the solution to the maximization problem⁸

$$\max_{p, q \in \mathcal{C}} \Phi(p, q) \text{ s.t. } p = m \cdot q^n + f.$$

⁸It is without loss to restrict the constraint to the equality case because the scoring rule is strictly decreasing in price.

The constraint correspondence is compact-valued and continuous in (m, f) so by Berge's Theorem of the Maximum, the solutions to this maximization problem are continuous in (m, f) . Then, if we fix a marginal cost m and apply the Intermediate Value Theorem to the continuous function $s_{BE}(m, \cdot) : [\underline{f}, \bar{f} + \bar{m}] \rightarrow \mathbb{R}$, there must be some $f_{(2)}$ satisfying the desired condition. \square

Now, we turn attention to Proposition 3.

Proposition 3. *Fix $g \in \mathcal{G}$. For all $\theta \in \Theta$, the equilibrium strategy can be represented as*

$$(p^*, q^*)(m, f, g) = (p_{BE}, q_{BE})(m, f_{(2)})$$

where

$$f_{(2)} = \frac{\int z \cdot \mathbb{1}_{z \geq f} g_m([(m, z)]_g) dz}{\int \mathbb{1}_{z \geq f} g_m([(m, z)]_g) dz}.$$

Proof. As discussed in the main text, we would like to project our auction with 2 dimensions of private information into a more familiar one-dimensional analogue where all types have marginal cost m . From Proposition 2 of Hanazono et al. [2016], we have that at any $g \in \mathcal{G}$, there exists a continuous density function of scores bid in equilibrium, $h(s)$ (because we are dealing with a fixed g currently, we suppress dependence of h on g in the notation). Now, fixing an m , we define a new distribution over types $g_m : \mathbb{R} \rightarrow \mathbb{R}$ as follows:

$$G_m(f) := H(s^*(m, f, g))$$

$$g_m(f) := h(s^*(m, f, g)) \cdot \left. \frac{\partial s^*(m, f', g)}{\partial f'} \right|_{f'=f}.$$

We now consider an auction in which types' fixed costs are their only private information and are drawn i.i.d according to G_m . In this auction, the agent with cost f who receives their bid b upon winning attains the expected utility

$$\mathbb{P}(\text{win}|b) \cdot (b - f).$$

From standard results on symmetric first-price auctions, in the unique symmetric Bayes-Nash equilibrium of this auction, bids satisfy

$$G_m(f) \cdot b^{*'}(f) + g_m(f) \cdot (b^*(f) - f) = 0. \tag{4}$$

Now consider the original scoring auctions. If the type (m, f) bids the score s , then best-

responding implies that

$$H(s) \cdot u_1(s, m, f) + h(s) \cdot u(s, m, f) = 0 \quad (5)$$

where $u(s, \theta)$ is the indirect utility to type θ from bidding s . If we let $b(f) = p^*(m, f, g) - m \cdot q^{*\eta}(m, f, g)$, then

$$u(s, m, f) = b(f) - f \text{ and } u_1(s, m, f) \cdot s_2^*(m, f, g) = b'(f).$$

Simplifying gives us that Equations 4 and 5 are equivalent. Thus,

$$b(f) = b^*(f) := \frac{\int z \cdot \mathbb{1}_{z \geq f} g_m(f) dz}{\int \mathbb{1}_{z \geq f} g_m(f) dz}.$$

From Lemma 2, the contract $(p^*(m, f, g), q^*(m, f, g))$ is the break-even strategy of some type $(m, f_{(2)})$. We finish by noting that by the definition of a break-even strategy,

$$b^*(f) = p^*(m, f, g) - m \cdot q^{*\eta}(m, f, g) = p_{BE}(m, f_{(2)}) - m \cdot q_{BE}^\eta(m, f_{(2)}) = f_{(2)}.$$

□

B Proofs From Section 3.2

Lemma 3. *Let $\theta \in \Theta$ and $g_1, g_2 \in \mathcal{G}$. If $s^*(\theta, g_1) = s^*(\theta, g_2)$, then*

$$s^*(\theta', g_1) = s^*(\theta', g_2) = s^*(\theta, g_1)$$

for all $\theta' \in [\theta]_{g_1}$. Equivalently,

$$[\theta]_{g_1} = [\theta]_{g_2}.$$

Proof. Let h_1, h_2 be the equilibrium score distributions at g_1 and g_2 respectively. Optimizing expected utility implies that

$$H_1(s) \cdot u(s, \theta') + h_1(s) \cdot u_1(s, \theta') = 0$$

for all $\theta' \in [\theta]_{g_1}$ where $s = s^*(\theta, g_1)$. From the given,

$$H_1(s) \cdot u(s, \theta) + h_1(s) \cdot u_1(s, \theta) = H_2(s) \cdot u(s, \theta') + h_1(s) \cdot u_1(s, \theta') = 0.$$

Thus, $\frac{h_1(s)}{H_1(s)} = \frac{h_2(s)}{H_2(s)}$ and

$$H_2(s) \cdot u(s, \theta') + h_2(s) \cdot u_1(s, \theta') = 0$$

for all $\theta' \in [\theta]_{g_1}$. There is a unique type with a given marginal cost for which this FOC holds. Thus, it is sufficient for concluding that $\theta' \in [\theta]_{g_2}$. \square

This will be applied in the proof of the following lemma from the main text:

Proof. If type θ 's score bid is unaffected by perturbation in the direction $v = g_2 - g_1$ then so are the bids of all types in the equivalence class $[\theta]_{g_1}$ by Lemma 3. If all of these strategies are unchanging under perturbation by v , then by Proposition 3

$$\frac{\partial \left(\frac{\int \rho_m(\tau, g_1 + \epsilon v) \mathbb{1}_{\tau \prec_{g_1 + \epsilon v} \theta}(\tau) d\tau}{\int \mathbb{1}_{\tau \prec_{g_1 + \epsilon v} \theta}(\tau) d\tau} \right)}{\partial \epsilon} \Big|_{\epsilon=0} = 0$$

for all m . Simplifying and noting that $\frac{\partial \int (\mathbb{1}_{\tau \prec_{g_1 + \epsilon v} \theta}(\tau) - \mathbb{1}_{\tau \prec_{g_1} \theta}(\tau)) g_1(\tau) d\tau}{\partial \epsilon} \Big|_{\epsilon=0} = 0$ yields

$$\underbrace{\int \frac{\partial \rho_m(\tau, g_1 + \epsilon v)}{\partial \epsilon} g(\tau) \mathbb{1}_{\tau \prec_{g_1} \theta} d\tau}_{(\dagger)} + \int \rho_m(\tau) \mathbb{1}_{\tau \prec_{g_1} \theta} v(\tau) d\tau - f_{(2)}(m, \theta, g_1) \int \mathbb{1}_{\tau \prec_{g_1} \theta} v(\tau) d\tau = 0.$$

Integrating uniformly over m and simplifying yields that

$$\int (\dagger) dm + \int \rho(\tau) \mathbb{1}_{\tau \prec_{g_1} \theta} g_2(\tau) d\tau - f_{(2)}(\theta, g_1) \int \mathbb{1}_{\tau \prec_{g_1} \theta} g_2(\tau) d\tau = 0$$

We will show that the first integral is equal to 0. To gain intuition for what this means, let $(m, f) \sim_g (m', f')$ but let $(m, f) \prec_{g+\epsilon v} (m', f')$. Then

$$\rho_{m'}(m, f) < f'.$$

But necessarily, we also have that

$$\rho_m(m', f') > f.$$

This fact that one type becoming less competitive inherently means another type becomes more competitive leads to this term canceling out on average. Examining this more formally, we have that by the definition of ρ_m ,

$$s^*(m, \rho_m(\tau, g + \epsilon v), g_1 + \epsilon v) = s^*(\tau, g_1 + \epsilon v).$$

If at $\epsilon = 0$, we have that $\rho_m(\tau, g_1) = f_1$, then implicitly differentiating yields

$$\frac{\partial s^*(m, f_1, g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} + \left(\frac{\partial s^*(m, f, g_1 + \epsilon v)}{\partial f} \Big|_{f=f_1} \right) \left(\frac{\partial \rho_m(\tau, g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} \right) = \frac{\partial s^*(\tau, g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0}.$$

Rearranging gives us that

$$\frac{\partial \rho_m(\tau, g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} = \frac{\frac{\partial s^*(\tau, g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} - \frac{\partial s^*(m, f_1, g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0}}{\frac{\partial s^*(m, f, g_1)}{\partial f} \Big|_{f=f_1}}.$$

The terms on the RHS exist because of the differentiability assumption of equilibrium strategies. Returning to the expression (†), recall that, fixing a marginal cost m , we can consider the space Θ as the set of equivalence classes $[(m, f)]_g$. We can thus rewrite this integral as

$$\int \frac{\partial \rho_m(\tau, g_1 + \epsilon v)}{\partial \epsilon} g(\tau) \mathbb{1}_{\tau \prec_g \theta} d\tau = \int \int \left(\frac{\partial \rho_m(m', \rho_{m'}(m, f), g_1 + \epsilon v)}{\partial \epsilon} dm' \right) \mathbb{1}_{(m, f) \prec_g \theta} g_m(f) df$$

where the density $g_m(f)$ satisfies

$$g_m(f) = h(s^*(m, f, g_1)) \cdot s_2^*(m, f, g_1)$$

where h is the equilibrium score density. Abusing notation, let $\rho_m : \mathbb{R} \rightarrow \bar{\Theta}$ be the type with marginal cost m that bids s at the distribution g . Then we can rewrite (†) as

$$\int \int \left(\frac{\partial s^*(\rho_{m'}(s), g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} - \frac{\partial s^*(\rho_m(s), g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} dm' \right) \mathbb{1}_{s \leq s^*(\theta, g_1)} h(s) ds.$$

From above, this expression must be negative for all m . Integrating (†) over all m however gives us

$$\int \int \int \left(\frac{\partial s^*(\rho_{m'}(s), g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} - \frac{\partial s^*(\rho_m(s), g_1 + \epsilon v)}{\partial \epsilon} \Big|_{\epsilon=0} dm' dm \right) \mathbb{1}_{s \leq s^*(\theta, g_1)} h(s) ds = 0.$$

This gives the desired result. □

B.1 Proofs From Section 3.2.1

The goal of this section is to prove Theorem 1. There are 2 lemmas proven beforehand to ease in exposition. This is divided into 2 separate results; one for each direction.

Proposition 5. *If an auction satisfies the fixed-order property then it admits a CBE with respect to a 2-moment signal structure.*

Proof. Fixing a marginal cost $m \in [\underline{m}, \bar{m}]$, we can define a projection function that maps the equivalence classes of types into the type with marginal cost m :

$$\rho_m : \Theta \rightarrow [\underline{f}, \bar{f} + \bar{m}]$$

$$\theta \mapsto f \text{ such that } (m, f) \in [\theta]_g$$

This is a well-defined function because no two types with the same marginal cost will be in the same equivalence class. We can write the expression in Equation 2 as

$$f_{(2)} = \frac{\int \rho_m(\tau) \cdot \mathbb{1}_{(m,f) \succeq \tau} g(\tau) d\tau}{\int \mathbb{1}_{(m,f) \succeq \tau} g(\tau) d\tau}$$

which is a smooth function of 2 moments⁹. The differentiability of the functions p_{BE} and q_{BE} results from the smoothness of the scoring rule.

To see that no such auction can have precision less than 2, no first-score auction can have precision 0 by Proposition 3 because strategies depend on the distribution g . To show an auction satisfying the fixed-order property cannot have an informational precision of 1, fix a type θ and assume it has a strategy with a precision of 1, described by the function $\zeta_\theta : \Theta \rightarrow \mathbb{R}$. Fix a distribution $g_1 \in \mathcal{G}$ and let $v \in \mathcal{G}_0$ satisfy

$$\int \zeta_\theta(\tau) v(\tau) d\tau = 0.$$

In words, the strategy of type θ should not change when any distribution is changed in the direction v . Then it must be that

$$\int (\rho_m(\tau) - f_{(2)}) \mathbb{1}_{(m,f) \succeq \tau} v(\tau) d\tau = 0$$

where $f_{(2)} = \frac{\int \rho_m(\tau) \cdot \mathbb{1}_{(m,f) \succeq \tau} g_1(\tau) d\tau}{\int \mathbb{1}_{(m,f) \succeq \tau} g_1(\tau) d\tau}$. Otherwise, from Proposition 3, this type's strategy would change when the distribution g_1 changes in the v direction. Now pick some distribution g_2

⁹Note that we suppress dependence of \succeq on g when the fixed-order property is satisfied.

so that

$$f_{(2)} \neq \frac{\int \rho_m(\tau) \cdot \mathbb{1}_{(m,f) \succeq \tau} g_2(\tau) d\tau}{\int \mathbb{1}_{(m,f) \succeq \tau} g_2(\tau) d\tau}.$$

Then changing g_2 in the direction v should change type θ 's strategy by Proposition 3 but the strategy with precision 1 does not change, giving a contradiction. \square

Proposition 6. *If a first-score auction admits a CBE with respect to an n -moment signal structure for $n < \infty$, then it satisfies the fixed-order property.*

Proof. Assume there exists a type $\theta \in \text{Int}(\Theta)$ and two distributions g_0 and g_1 such that $[\theta]_{g_0} \neq [\theta]_{g_1}$. Then without loss of generality, let there be two types $(m, f_1) \in [\theta]_{g_0}$ and $(m, f_{n+2}) \in [\theta]_{g_{n+2}}$ with $f_{n+2} < f_1$. Let $g_\lambda := \lambda g_0 + (1 - \lambda) g_1$. By the differentiability of scores with respect to the distribution, there exist $0 = \lambda_1 < \lambda_2 < \dots < \lambda_{n+2} = 1$ such that

$$(m, f_i) \in [\theta]_{g_i}$$

with $f_j < f_i$ if $j > i$. By the continuity of the score bidding function, there exists a small open neighborhood U_i around (m, f_i) such that the set

$$S_i := \left(U_i \cap \{ \tau \in \Theta : \theta \succ_{g_{\lambda_i}} \tau \} \right) \subset \{ \tau \in \Theta : \theta \succ_{g_{\lambda_i}} \tau \text{ but } \theta \prec_{g_{\lambda_j}} \tau \text{ for all } j < i \}$$

and S_i has positive measure for each $i \in \{1, \dots, n+2\}$. In words, the set S_i contains the (competitive) types against which θ wins the auction at the distribution g_{λ_i} but not at g_{λ_j} for lower j . Then, there must be constants $b_{i,j} \in \mathbb{R}$ with $b_{i,i} > 0$ such that

$$\int_{S_j} \left((\rho(\tau, g_{\lambda_i}) - \psi(\theta, g_{\lambda_i})) \mathbb{1}_{\tau \prec_{g_{\lambda_i}} \theta} \right) d\tau = b_{i,j}$$

for all $j \leq i$ and

$$\int_{S_i} \left((\rho(\tau, g_{\lambda_i}) - \psi(\theta, g_{\lambda_i})) \mathbb{1}_{\tau \prec_{g_{\lambda_j}} \theta} \right) d\tau = 0$$

for all $j > i$. We will show that this leads to a violation of the Equal Competitiveness Lemma. Specifically, the above equations define $n+2$ linearly independent functionals which cannot be spanned by the n functionals M_1, \dots, M_n . Thus, there exists a vector in the kernel of M_k for each k that does not satisfy the constraint of the Equal Competitiveness Lemma. To see this, we can define $\mu_1, \dots, \mu_{n+2} \in \mathbb{R}$ satisfying

$$\sum_{i=1}^{n+2} \mu_i = 0 \text{ and } \sum_{i=1}^{n+2} M_k(S_i) \mu_i = 0 \text{ for all } k$$

where $M_k(S_i) := \int_{S_i} \zeta_k(\tau) d\tau$.

There is at least a one-dimensional space of solutions to this system of equations. Picking some nontrivial vector in this space, let \underline{i} be the lowest index at which μ_i is non-negative

and consider $v \in \mathcal{G}_0$ with $v(\tau) = \begin{cases} \mu_i & \tau \in S_i \\ 0 & \text{else.} \end{cases}$ Then we have that

$$\int \left((\rho(\tau, g_{\lambda_{\underline{i}}}) - \psi(\theta, g_{\lambda_{\underline{i}}})) \mathbb{1}_{\tau < g_{\lambda_{\underline{i}}}} \right) v(\tau) d\tau = \sum b_{\underline{i}, j} \mu_j = b_{\underline{i}, \underline{i}} \mu_{\underline{i}} \neq 0.$$

□

B.2 Proofs from Section 3.2.2

Proof. From the logic in the text, the iso-score lines at the distribution g in the 2 bidder setting must be the same as those in the 3 bidder setting. We can construct a sequence of distributions g_n for which

$$H_{g_i}^3 = H_{g_{i+1}}^2$$

H_g^n denotes the equilibrium score distribution of the auction where there are n bidders, each of whose type is drawn from a distribution with density g_i . From the observation above, the iso-score lines must be the same across all of these distributions, or

$$\sim_{g_1} = \sim_{g_n}$$

for all n . As $n \rightarrow \infty$, the probability that any type in $\text{Int}(\Theta)$ wins approaches 0. From the logic in Section 4, this implies that \sim_{g_1} is the break-even contract order and this auction must satisfy the fixed-order property. □

B.3 Proofs from Section 3.2.3

For a scoring rule Φ , let the recommendation signal structure $\mathbb{I}^*(\Phi)$ be defined by $I(\theta) = \cup E_{s, \theta}$ where

$$g \in E_{s, \theta} \iff s = s^*(\theta, g).$$

Each information set in $I^*(\Phi)$ corresponds to a score and types should simply bid the score corresponding to their signal realization. I will prove the following result:

Proposition 7. *For all types $\theta \in \Theta$ and any two scoring rules Φ and Φ' , the signal structures*

$I^*(\Phi)(\theta)$ and $I^*(\Phi')(\theta)$ are isomorphic. However, $I^*(\Phi)$ is convex if and only if Φ satisfies the fixed-order property.

Proof. If $I^*(\Phi)$ is convex, then the Equal Competitiveness Lemma applies. If Φ satisfies the fixed-order property, then this linear constraint defines a hyperplane in \mathcal{G} which is convex. Assuming this property is not satisfied, then there exist a type $\theta = (m, f) \in \Theta$, an information set $E \in I(\theta)$, and two distributions $g_1, g_2 \in E_{s,\theta}$ such that $\rho_m(\tau, g_1)$ and $\rho_m(\tau, g_2)$ differ on a set of positive measure. Let v be a direction that moves mass along the iso-score lines defined by \sim_{g_2} such that

$$\text{sign}(v(\tau)) = \text{sign}(\rho_m(\tau, g_2) - \rho_m(\tau, g_1)).$$

Perturbing g_1 in the direction of v makes the auction more competitive and thus raises the bid of type θ . However, by the equal competitiveness lemma, this perturbation should leave θ 's bid the same. Thus, this set cannot be convex. \square

Theorem 3 will result from this proposition. All signal structures that admit CBE are refinements of $\mathcal{I}^*(\Phi)$. Thus, the only possible signal structure that can be isomorphic to $\mathcal{I}^*(\Phi)$ and admit a CBE is $\mathcal{I}^*(\Phi')$. When this set is not convex, there is then no isomorphic element to $\mathcal{I}^*(\Phi)$.

For the reverse direction, let $E'_{s,\theta} \in \mathcal{I}^*(\Phi')$ and let $E_{s,\theta} \in \mathcal{I}^*(\Phi)$ with Φ satisfying the fixed-order property. $E_{s,\theta}$ is a (shifted) linear subspace of $\mathcal{L}^2(\Theta, \mu)$ with codimension 2. Thus, a linear projection from $L^2(\Theta, \mu) \rightarrow U$ for some linear subspace U of the same codimension followed by a linear isomorphism from $U \rightarrow E_{s,\theta}$ preserves the linearity and convexity of any subset of \mathcal{G} and thus any subset of $E'_{s,\theta}$. It suffices to pick U to ensure the $E'_{s,\theta}$ is isomorphic with its image under the projection. This happens when no two points in $E'_{s,\theta}$ are separated by one of the vectors normal to U . We define $U \subset \mathcal{G}_0$ and so that its other normal vector shifts mass from each type to one with the same marginal cost and a higher fixed cost. This unambiguously lowers scores. Thus, no two elements of $E'_{s,\theta}$ can be separated by this vector. This completes the construction of a bijection from $E'_{s,\theta} \rightarrow E_{s,\theta}$. Similar constructions can produce an isomorphism from an element of $\mathcal{I}(\Phi')$ to one of $\mathcal{I}(\Phi)$.

C Proofs from Section 4

In this section, I provide a proof of Theorem 4 and Proposition 4.

Theorem 4. *A first-score auction with scoring rule Φ is informationally coarse if and only if it implements the same interim outcomes as the second-score auction with the same scoring rule at all distributions $g \in \mathcal{G}$.*

Proof. Fix a scoring rule Φ such that the first-score auction defined with this rule satisfies the fixed-order property. Consider two types θ_1, θ_2 with the same break-even scores and assume for the sake of contradiction that $\theta_1 \succ \theta_2$ in the first-score auction's equilibrium structure. Consider the sequence of distributions $(g_n)_n \in \mathcal{G}$ defined by

$$g_n = \frac{1}{n}g_0 + \left(1 - \frac{1}{n}\right) \cdot u$$

where u is the density of a uniform distribution over the set $\{\theta' \in \Theta : \theta_1 \succ \theta' \succ \theta_2\}$.

Let s_1 be the score bid by θ_1 at the distribution u and let \bar{U} be the payoff of type θ_2 from bidding this score, conditional on winning. Note that $s_1 < s_{BE}(\theta_1) = s_{BE}(\theta_2)$ so \bar{U} is strictly positive. Note that for any type's equilibrium bid (p, q) , we must have that

$$0 \geq p - (\bar{m}, \bar{f}) \cdot (q^n, 1) \implies \bar{m} + \bar{f} \geq p - \theta \cdot (q^n, 1)$$

for any $\theta \in \Theta$. Thus, conditional on winning, type θ_2 's payoff is bounded above by $\bar{m} + \bar{f}$.

As n grows, $X(\theta_2, g_n) \rightarrow 0$ but $X(\theta_1, g_n) \rightarrow 1$. Thus, let $N \in \mathbb{R}$ large enough such that

$$\frac{X(\theta_1, g_N)}{X(\theta_2, g_N)} > \frac{\bar{m} + \bar{f}}{\bar{U}}.$$

This implies that expected utility to type θ_2 will be strictly greater from bidding s_1 than their bid in the mechanism. As a result, this mechanism is not incentive-compatible, giving a contradiction.

This means that the first-score auction and the second-score auction have the same interim contract allocation rule X . In both auctions, all types in the equivalence class $[(\bar{m}, \bar{f})]$ receive zero utility. By the envelope theorem, the slope of the indirect utility functions in both auction formats are the same upon changing fixed costs. Thus, fixing marginal costs and applying the Fundamental Theorem of Calculus on each of these lines yields that the interim utility functions from the two auction formats are the same. Then, also by the envelope theorem, the interim effort allocations are equal to the partial derivative of this utility function with respect to marginal cost, which are the same across these two mechanisms. \square

Proposition 4. *A first-score auction with scoring rule Φ is informationally coarse if and only if the break-even effort function, $e_{BE} := q_{BE}^\eta : \Theta \rightarrow [0, 1]$ is linear in fixed cost.*

Proof. Fix such a first-score auction and a type $(m, f) \in \Theta$. Recall that given a score to bid, all types with the same marginal cost will optimize by picking the same contract.

As a result, if the other participant's type is in $[(m, z)]$ for $(m, z) \in \bar{\Theta}$, an agent with marginal cost m will provide the contract $(p_{BE}(m, z), q_{BE}(m, z))$ in a second-score auction. Thus, the expected effort of this type in this auction, at any distribution $g \in \mathcal{G}$, is

$$\int e_{BE}(m, z) \cdot \mathbb{1}_{z \geq f} g([(m, z)]) dz.$$

From Proposition 1, the expected effort of this type in a first-score auction is

$$e_{BE} \left(m, \frac{\int z \cdot \mathbb{1}_{z \geq f} g([(m, z)]) dz}{\int \mathbb{1}_{z \geq f} g([(m, z)]) dz} \right) \cdot \int \mathbb{1}_{z \geq f} g([(m, z)]) dz.$$

These two expressions must be equal for a scoring rule that satisfies the fixed-order property, by Theorem 2. In a converse to Jensen's inequality, we'll show that these two expressions being equal for all x implies that e_{BE} must be linear. Firstly, we have that $e_{BE}(m, \cdot)$ is smooth because Φ is. This means its second-derivative must be continuous. Now assuming, e_{BE} is not linear, there must be some subinterval of $[\underline{f}, \bar{f}]$ on which this function is strictly convex or concave. Let $[f_1, f_2]$ be the highest such interval and WLOG let the function be strictly convex here. Then we have that e_{BE} is convex on $[f_1, \bar{f}]$ so by the strict form of Jensen's inequality,

$$e_{BE} \left(m, \frac{\int z \cdot \mathbb{1}_{z \geq f'} g([(m, z)]) dz}{\int \mathbb{1}_{z \geq f'} g([(m, z)]) dz} \right) < \frac{\int e_{BE}(m, z) \cdot \mathbb{1}_{z \geq f'} g([(m, z)]) dz}{\int \mathbb{1}_{z \geq f'} g([(m, z)]) dz}.$$

for all $f' \in [f_1, f_2)$ giving us a contradiction. Thus, e_{BE} must be linear in fixed cost. \square

D Endogenous Information Acquisition

In this section, I provide a formal model of information acquisition and belief formation in a coarse beliefs equilibrium. The model in this section is more general than that required to discuss concepts in Sections 2 and 3. First, I relax assumptions on the number of and symmetry between agents. Second, I model agents playing a general game with potentially more dependence of one agent's utility on another's type and action than in a first-score auction. Third, I allow for agents to learn from a general information technology. In the main result of this section, I relate coarse prior equilibria to equilibria of common prior games and show that for moment-based information technologies, a coarse prior equilibrium is equivalent to an equilibrium in n -moment strategies.

D.1 Model

Overview. There is a finite number of agents n who play in a predetermined mechanism. At the start of the game, agent i 's private information is captured by the payoff type $\theta_i \in \Theta$ for $i \in \{1, \dots, n\}$. There is a state of the world $g_0 \in \Delta(\Theta)^n$ which captures the true distribution of payoff types in a large population from which these agents are drawn. Each agent's payoff type is, then, drawn independently according to the density $g_{0,i} \in \Delta(\Theta)^{10}$. Agents know that $\mathbf{g}_0 \in \mathcal{G} \subset \Delta(\Theta)^n$ but they have no priors over \mathcal{G} . They are, however, able to learn coarse information about \mathbf{g}_0 . The structure of their learning process will be described in more detail below. Once the information acquisition stage is completed, the agents play in the mechanism, which is defined by the set of actions A , set of outcomes Y , and a function from $A^n \rightarrow Y$. Each agent's ex post utility from the mechanism is captured by the indirect utility function

$$U_i(\boldsymbol{\theta}, \mathbf{a}).$$

In summary, including the information acquisition stage, they participate in a two-stage game whose timing is defined as follows:

1. The rules of the mechanism are announced, agents learn their payoff types, and then agents form arbitrary beliefs.
2. Each agent privately selects a feasible signal structure which is a partition of the state space \mathcal{G} .
3. Each agent receives a private signal which is the element of his selected partition that contains the true state g_0 .
4. Each agent updates his beliefs given his signal realization.
5. The agents participate in the announced mechanism.

Given that the choice of initial beliefs is uninformed, my notion of simplicity will correspond to an equilibrium in which this choice is strategically irrelevant for every possible realization of signals. In other words, for every belief that is reasonable for an agent to hold, his best-response should be the same.

In formalizing these concepts, I will first discuss the structure of agents' learning process, then I will turn to agents' prior and posterior beliefs, and finally, I will present my solution concept.

¹⁰Note that this formulation implies conditional independence. The realization of θ_i yields no information about other agents' components of the state.

Signal Structures. There is a set of (non-fully revealing) partitions of the state space \mathcal{I} . This set, called the information technology, is common knowledge. After his payoff type is realized, each agent privately selects a partition from \mathcal{I} as his signal structure. Because there is uncertainty about other agents' payoff types, one agents' off-path behavior—that is, his choice of a signal structure at an unrealized payoff type—will affect another's agents on-path actions. Thus, for each possible payoff type, an agent chooses a signal structure. Agent i 's information acquisition strategy is captured by the function $I_i : \Theta \rightarrow \mathcal{I}$. For each payoff type θ_i , the image of this map, $I_i(\theta_i)$ forms a partition of the state space. Each agent receives a signal and learns the set in $I_i(\theta_i)$ containing the true state g_0 . Fixing a payoff type θ_i , agent i will have different beliefs depending on the realization of the signal. As a result, the payoff type space Θ is insufficient for fully describing the beliefs agents hold before playing in the mechanism. Thus, for each agent we define a new type space, $\mathcal{T}_i \subset \Theta \times 2^{\mathcal{G}}$ whose elements are the terms

$$(\theta_i, E_i)$$

for all $\theta_i \in \Theta$ and $E_i \in I_i(\theta_i)$. I refer to $t_i \in \mathcal{T}_i$ as agent i 's belief type. I write $\theta(t_i)$ as t_i 's corresponding payoff type and $E(t_i)$ as the set of distributions admissible to t_i . Note that this definition of the belief type space depends on the signal structure \mathbf{I} . When necessary, I write the type space as $\mathcal{T}_i(I_i)$ but suppress this dependence when clear.

Beliefs. Each agent enters the game with arbitrary beliefs about the state and their opponents' types. Rather than modelling these explicitly, I follow Harsanyi [1968] and model this implicitly through the construction of a belief structure (see Mertens and Zamir [1985] for more detail on the equivalence of these two approaches). Let $\pi_i \in \Delta(\mathcal{G})$ be agent i 's (arbitrary) first-order belief¹¹. After choosing the signal structure I_i in the first stage and learning the realization E_i , each agent updates his first-order beliefs using Bayes rule. In other words, this signal realization induces the posterior first-order beliefs $\tilde{\pi}_i : \mathcal{T}_i \rightarrow \Delta(\mathcal{G})$ defined by

$$\mathbb{P}_{\tilde{\pi}_i(\theta_i, E_i)}(A) = \frac{\mathbb{P}_{\pi_i}(A \cap E_i)}{\mathbb{P}_{\pi_i}(E_i)}$$

for all Lebesgue measurable sets $A \subset \mathcal{G}$. In general, defining first-order beliefs is insufficient for fully specifying belief hierarchies. However, because the state in this model is a distribution over opponents' types, each first-order belief structure uniquely extends to a full hierarchy of beliefs as expressed below:

¹¹Important measurability considerations are obscured here as in general the set \mathcal{G} can be infinite-dimensional. To deal with this, we first restrict to distributions whose supports admit some finite parametrization. For simplicity, we then restrict the set of possible priors to measures that are absolutely continuous with respect to the Lebesgue measure over this parameter space.

Lemma 4. *The first-order belief structure $\tilde{\pi}$ uniquely defines beliefs of the form $\hat{\pi}_i : \mathcal{T}_i(I_i) \rightarrow \Delta(G \times \mathcal{T}_{-i}(I_{-i}))$ for all i .*

Proof. By the definition of the state, each $g = (g_1, \dots, g_n)$ defines a unique distribution over opponents' payoff types, which is an element of $\Delta(\Theta_{-i})$. By the definition of an information technology, for each agent j and payoff type θ_j , there is a unique element of $I_j(\theta_j)$ that contains g . In other words, there is a unique $t_j \in I_j(\theta_j)$ such that $g \in E(t_j)$. Thus, fixing a state g , which is a distribution over payoff types, we can extend this to a distribution over belief types as follows:

$$\mathbb{P}_g(A) := \mathbb{P}_g\left(\{\theta(t_{-i}) : t_{-i} \in A, g \in E(t_j) \forall j \neq i\}\right)$$

for $A \subset \mathcal{T}_{-i}(I_{-i})$. We have uniquely mapped each element $g \in \mathcal{G}$ to a distribution over belief types. We can extend a distribution over states to a joint distribution over states and payoff types by interpreting the definition of each state in the support as a distribution over payoff types. We can then extend this joint distribution over states and payoff types to one over states and belief types with the definition above. □

Solution Concept. Once agents have formed posterior beliefs, they are expected utility maximizers. Thus, fixing some posterior belief structure $\tilde{\pi}$ and opponents' action profile $a_{-i} : \mathcal{T}_{-i} \rightarrow A^{n-1}$, an agent with type t_i prefers action a_1 over a_2 in the second stage of the game if

$$\mathbb{E}_{\tilde{\pi}}[U(\theta(t_i), \theta(t_{-i}), a_1, a_{-i}(t_{-i}))] \geq \mathbb{E}_{\tilde{\pi}}[U(\theta(t_i), \theta(t_{-i}), a_2, a_{-i}(t_{-i}))].$$

To define preferences over the whole game, intuitively an agent should prefer one signal structure to another if it allows him to choose an action that gives him a higher expected utility:

Definition 13 (Preferences). *Fix prior beliefs π and fix the strategy of other agents to be $\mathbf{I}_{-i} : \Theta^{n-1} \rightarrow \mathcal{I}^{n-1}$ and $a_{-i} : \mathcal{T}_{-i}(\mathbf{I}_{-i}) \rightarrow A^{n-1}$. Then agent i prefers the strategy (I_i^1, a^1) to (I_i^2, a^2) if for all payoff types $\theta_i \in \Theta$,*

$$\mathbb{E}_{\pi}[U(\theta_i, \theta(t_{-i}(g)), a^1(t_i^1(g)), a_{-i}(t_{-i}(g)))] \geq \mathbb{E}_{\pi}[U(\theta_i, \theta(t_{-i}(g)), a^2(t_i^2(g)), a_{-i}(t_{-i}(g)))]$$

where $t_{-i}(g)$ represents opponents' belief types at g and $t_i^j(g)$ represents agent i 's belief type at g under the signal structure $I_i^j(\theta_i)$. If this holds, we write that

$$(I_i^1, a^1) \succeq_{i, (\mathbf{I}_{-i}, a_{-i}), \pi} (I_i^2, a^2)$$

Because each agent arbitrarily chooses beliefs at the beginning of the game, his preferences should not be dependent on this random selection. Thus, we can refine this definition of preferences to one that holds at all beliefs:

Definition 14 (Strong Preferences). *Agent i strongly prefers the strategy (I_i^1, a^1) to (I_i^2, a^2) if for all possible belief structures $\tilde{\pi}$, we have that*

$$(I_i^1, a^1) \succeq_{i,(\mathbf{I}_{-i}, a_{-i}), \tilde{\pi}} (I_i^2, a^2).$$

In this case, we write

$$(I_i^1, a^1) \succeq_{i,(\mathbf{I}_{-i}, a_{-i})} (I_i^2, a^2).$$

I use the term preference but this relation is not a preference relation in the standard sense. In particular, it is not rational because it is not complete. One strategy may be better than another at some but not all prior belief structures. As a result, there may not be a unique maximal element according to this ordering. In general, the existence of best-responses depends both on the definition of the information technology and the mechanism. Nevertheless, when best responses do exist, a corresponding equilibrium notion can be defined.

Definition 15. *A mechanism admits a **coarse beliefs equilibrium** with respect to the information technology \mathcal{I} if for each agent, there is an information acquisition strategy $I_i^* : \Theta \rightarrow \mathcal{I}$ and a second-stage strategy $a_i^* : \mathcal{T}(I_i^*) \rightarrow A$ such that*

$$(I_i^*, a_i^*) \succeq_{i,(I_{-i}^*, a_{-i}^*)} (I, a)$$

for all other strategies $I : \Theta \rightarrow \mathcal{I}$ and $a : \mathcal{T}(I) \rightarrow A$.

The simplest example of a coarse beliefs equilibrium is when the information technology is fully uninformative. In this case, an agent's equilibrium strategy must be a best response at any belief given equilibrium strategies of his opponents. In this case, this concept is the same as a robust equilibrium. In general, the concept of a coarse beliefs equilibrium can be viewed as a weakening of the concept of a robust equilibrium. The space of feasible beliefs is restricted through the receipt of signals but after that, it requires an action be optimal at all reasonable belief structures.

Remark. The common knowledge assumptions in this model are substantial weakenings of those in the Bayes-Nash model. In particular, this model does not require agents share a common prior or even common knowledge of some information structure. All that is needed is common knowledge about the information technology \mathcal{I} . I view this as less restrictive than specific knowledge about belief hierarchies.

D.2 Results

As mentioned above, the concept of a CBE relates closely to robust equilibrium. As shown in Bergemann and Morris [2005], in certain classes of games, a mechanism satisfies robust implementability if and only if it is interim-implementable on all common prior type spaces. An analogous result will hold here.

To begin, I define the common prior game at $g \in \mathcal{G}$ to be similar to the one defined above with the sole change that it is common knowledge that all agents believe the true distribution is g . All agents are assumed to be expected utility maximizers so they play pure Bayes-Nash strategies written as $a_i^{CP}(\theta_i, g)$ ¹². The main result of this section can be stated as follows:

Theorem 5. *A mechanism admits a coarse beliefs equilibrium defined by I_i^* and $a_i^* : \mathcal{T}(I_i^*) \rightarrow A$ if and only if*

$$a_i^*(t_i) = a_i^{CP}(\theta_i, g_1) = a_i^{CP}(\theta_i, g_2)$$

for all $t_i := (\theta_i, E_i) \in \mathcal{T}(I_i^*)$ and $g_1, g_2 \in E_i$.

The proof of Theorem 5 is quite straightforward. For the forward direction, if a strategy is optimal at all beliefs, then it must be optimal at all plausible common prior beliefs. For the reverse direction, if an action is optimal point-wise at every type distribution g , then it must be optimal at any belief that can be expressed as a mixture over these distributions.

Proof. (\Rightarrow) Starting with a coarse beliefs equilibrium, fix a type $t_i = (\theta_i, E_i) \in \mathcal{T}_i(I_i^*)$ and a distribution $g \in E_i$. For $i \in \{1, \dots, n\}$, consider a belief structure satisfying

$$\tilde{\pi}_i(t_j) = \delta_g$$

for all $t_j = (\theta_j, E_j) \in \mathcal{T}_j(I_j^*)$ such that $g \in E_j$. Under this belief structure, to type t_i , it will appear that the type distribution being g is common knowledge. By the definition of a common prior equilibrium and robust preferences, this immediately implies that at this specific belief structure, each agent's action $a_i^*(t_i)$ is a best-response to other agents' equilibrium actions. Thus, this forms an equilibrium under the belief that g is a common prior.

(\Leftarrow) Fix agent i and assume that for all $g \in \mathcal{G}$, opponent j with payoff type θ_j plays the strategy $a_j^{CP}(\theta_j, g)$. I will show agent i can do no better than adopting the common prior strategy. Let $t_i^*(g)$ be type θ_i 's belief type under the information acquisition strategy I_i^* and

¹²I assume a unique, pure strategy Bayes-Nash equilibrium in this section solely for ease of exposition. These results hold more generally.

let $a^*(t_i^*(g)) = a_i^{CP}(\theta_i, g)$. This is well-defined because, by the given,

$$t_i^*(g) = t_i^*(g') \implies a_i^{CP}(\theta_i, g) = a_i^{CP}(\theta_i, g').$$

We have that by definition,

$$a_{CP}(\theta_i, g) \in \arg \max_{a \in A} \mathbb{E}_g[U(\theta(t_i), \theta_{-i}, a, a_{CP}(\theta_{-i}, g))]$$

Thus, for any prior belief π and alternative strategy that yields the action $a(g)$ at the type distribution g , we have that

$$\mathbb{E}_{g \sim \pi} \left[\mathbb{E}_g[U(\theta(t_i), \theta_{-i}, a(g), a_{CP}(\theta_{-i}, g))] \right] \leq \mathbb{E}_{g \sim \pi} \left[\mathbb{E}_g[U(\theta(t_i), \theta_{-i}, a^*(t_i^*(g)), a_{CP}(\theta_{-i}, g))] \right].$$

The RHS is the expected utility of playing the strategy (I^*, a^*) . Because this holds for all π , by definition this strategy is robustly optimal and thus this strategy profile forms a CBE. \square

Instead of considering this two-stage game and many different belief structures, it suffices to characterize equilibrium actions in common prior games to understand which equilibria can be implemented in coarse beliefs. Implementability of an equilibrium depends on the information technology to which agents have access. I will consider from now on the symmetric case in which all agents' types are drawn from the same distribution. Slightly abusing notation, I will call this distribution g_0 and let $\mathcal{G} = \Delta(\Theta)$. Under this simplification the results from Sections 3 can be adapted to this setting through the following observation:

Proposition 8. *A mechanism admits a coarse beliefs equilibrium with respect to the information technology \mathcal{I} if and only if there exists $I^* \in \mathcal{I}$ such that the mechanism admits a coarse beliefs equilibrium with respect to I^* .*