

Efficient Competition through Cheap Talk: The Case of Competing Auctions*

Kyungmin Kim[†] and Philipp Kircher[‡]

December 2014

Abstract

We consider a large market where auctioneers with private reservation values compete for bidders by announcing cheap-talk messages. If auctioneers run efficient first-price auctions, then there always exists an equilibrium in which each auctioneer truthfully reveals her type. The equilibrium is constrained efficient, assigning more bidders to auctioneers with larger gains from trade. The choice of the trading mechanism is crucial for the result. Most notably, the use of second-price auctions (equivalently, ex post bidding) leads to the non-existence of any informative equilibrium. We examine the robustness of our finding in various dimensions, including finite markets and equilibrium selection.

JEL Codes: C72, D82, D83.

Keywords: Competing auctions; cheap talk; directed search

*We are grateful to Yeon-Koo Che, Andrew Clausen, Srihari Govindan, Johannes Hörner, Navin Kartik, John Kennes, George Mailath, Guido Menzio, Dale Mortensen, Roger Myerson, Andrew Postlewaite, József Sákovics, Robert Shimer, Art Shneyerov, and seminar participants at various places for their helpful suggestions and comments. Kircher acknowledges support from ERC grant 284119. This paper supersedes an earlier draft that circulated under the title “Efficient cheap talk in directed search: on the non-essentiality of commitment in market games.”

[†]University of Iowa, kyungmin-kim@uiowa.edu.

[‡]University of Edinburgh, philipp.kircher@ed.ac.uk.

1 Introduction

We investigate the potential of cheap-talk communication in the canonical competing auctions framework. In general, auctioneers with larger gains from trade are willing to accept higher monetary demands in exchange for higher trading probabilities. How such incentives operate when auctioneers publicly commit to reserve prices, and what their aggregate consequences are, have been thoroughly investigated in the literature. In this paper, we demonstrate that even if auctioneers' announcements are cheap talk, competition can drive auctioneers to truthfully reveal their private reservation values, which leads to efficient market outcomes. Unlike in reserve price posting models, the trading mechanism is crucial for the result: full revelation arises under first-price auctions, but cheap talk can never be informative, e.g., under second-price auctions.

The literature on competing auctions has provided two fundamental insights. First, if auctioneers compete by posting reserve prices in a large market, then each auctioneer's optimal reserve price is driven down to her reservation value. In other words, in a competitive environment, an efficient auction is the revenue-maximizing auction mechanism for each auctioneer.¹ Second, if each auctioneer runs an efficient auction and all auctioneers' reservation values are publicly observable, then the resulting market outcome satisfies desirable efficiency properties. In particular, auctioneers who are more eager to trade attract more bidders, and the gross surplus in the market is maximized subject only to inherent market (search) frictions.² Together, they imply that reserve price competition yields socially efficient outcomes in large markets, even if auctioneers have private information about gains from trade.³

Cheap talk offers a plausible alternative to reserve price posting. Both enable pre-match communication, which is necessary for an efficient assignment of bidders to auctioneers. They differ in that the latter provides explicit incentives for the other side, while the former relies only on implicit incentives: a more attractive reserve price directly improves bidders' payoffs, whereas cheap-talk messages, by definition, cannot do that. The effects of cheap-talk communication have been broadly investigated in information economics, but its potential to induce efficient competition in market environments has not garnered much attention yet. In our model, each auctioneer possesses private information about gains from trade. This generates a potential

¹See, among others, McAfee (1993), Peters (1997b), and Peters and Severinov (1997). The former two papers go one step further and prove that an efficient auction is an optimal selling mechanism among all incentive-compatible mechanisms that ask bidders to report only their values.

²See, for example, Julien, Kennes and King (2005, 2006). This efficiency result is also highlighted in the studies that establish the connection between competing auctions and competitive search. See, e.g., Kultti (1999), Shimer (1999), and Julien, Kennes and King (2000).

³In the online appendix, we show that the efficiency result of reserve price posting does not hold generally in finite markets. This is a new finding in the literature. Julien, Kennes and King (2002) establish efficiency of reserve price posting for the case of two auctioneers and two bidders. We show that their result does not extend beyond their special case in the presence of auctioneer heterogeneity. This result is similar to the inefficiency result in finite directed search with heterogeneity (see Galenianos, Kircher and Virág (2011)).

role for cheap-talk communication as well as raises the question of how much information can be transmitted through cheap-talk messages. In this regard, our model extends earlier work that establishes efficiency in decentralized markets without incomplete information (e.g., Julien, Kennes and King (2005, 2006)).

Introducing cheap talk into an otherwise canonical competing auctions setting is also a logical response to the common criticism that binding reserve price announcements are often absent in real markets. The criticism was recognized from the beginning of the competing auctions literature, as Peters (1997b) states: “The prediction that sellers will offer auctions is not so implausible. [...] However, [...] sellers do not publish commitments to reserve prices.” This concern has been confirmed in applied work. For example, Piga and Zanza (2004) study 18 organizations from the EU Procurement learning lab representing different EU countries and conclude that “many organizations consider the reserve price as an estimate that is not necessary to be disclosed to bidders.”⁴ More broadly, it is well-documented in the auction literature that reserve prices are not publicly revealed in various auction markets (see, e.g., Ashenfelter, 1989; Li and Perrigne, 2003; Bajari and Hortacısu, 2004).

We first show that a fully revealing equilibrium exists under the following first-price auction format:⁵ after observing auctioneers’ announcements, each bidder quotes a price to an auctioneer, and then each auctioneer decides which quote, if any, to accept. Given full revelation, the resulting market outcome coincides with the efficient outcome in existing competing auctions models. This demonstrates that the main insights from the competing auctions literature extend to a setting where auctioneers compete only with implicit incentives. This might be particularly useful in environments where reserve price posting is demanding.⁶

The details of the auction mechanism are crucial for the result, unlike in reserve price posting models where any standard auction format can be adopted. In particular, the result about full revelation does not extend to second-price auctions in which the resulting price reflects the second-best bidder’s value.⁷ In fact, with second-price auctions, there does not exist any

⁴According to Piga and Zanza (2004), only a minority of organizations allow their constituents to post binding reserve prices. In addition, those organizations acknowledge encouraging bidder participation, which is at the center of our theoretical investigation, as the rationale behind their rule.

⁵We note that, unlike competing mechanisms models, the auction rule is exogenously given in our model. In other words, the trading mechanism is not designed by each auctioneer. In Section 5, we explain why this is unavoidable. Although restrictive in theory, it is consistent with the use of the default trading rule in various markets. For example, in the EU procurement market, the default trading rule is the open sealed-bid first-price auction in which procurers “must advertise the contract in the Official Journal, allow any interested firm to submit a bid, and evaluate the bids as received without entering into negotiations” (PwC, London Economics, and Ecorys, 2011). A similar rule applies to most procurements in the US.

⁶For example, if bidders differ in quality and this is observable but not verifiable (see Section 5), each auctioneer needs to specify different reserve prices for different bidder types (i.e., a reserve price schedule rather than a single reserve price), whereas cheap-talk communication requires only that each auctioneer transmits her own type.

⁷In our model, each auctioneer’s reservation value is private information. Since there is a positive probability that there is only one bidder, the standard second-price auction format cannot be directly implemented. Still, the second-price auction outcome can be replicated through ex-post bidding (i.e., English auction).

informative equilibrium, not to mention a fully revealing equilibrium. Given full revelation, due to revenue equivalence, all standard auctions achieve the same efficient market outcome. However, different auctions provide different incentives for auctioneers at the communication stage and, therefore, have different implications for the informativeness of cheap talk.

To understand these results, first consider the first-price auction. Truth-telling can be sustained for the following reason: in equilibrium bidders must be indifferent among auctions. This implies that any bid below two auctioneers' reservation values must have the same probability of winning at both auctions. This arises only when the distributions of expected bids at both auctions are identical, except that the one with the higher reservation value gets additional bids above the other's reservation value but below her own. It follows that understating gains from trade only eliminates some bids that are acceptable to the auctioneer and, therefore, is strictly unprofitable, while exaggerating gains from trade invites more bids only beyond the auctioneer's reservation value, which is weakly unprofitable. For the failure of the second-price auction, recall the well-known result from auction theory that its price distribution has the same expected value as, but is more dispersed than, the one induced by the first-price action. Consider an auctioneer who exaggerates her value. She can reject any price above her reservation value. This means that she benefits from the upside of more dispersion, while being shielded from the downside. Since truth-telling incentives are tight with the first-price auction, the deviation is strictly profitable with the second-price auction, which unravels full revelation.

Our positive result with the first-price auction is robust in several dimensions, although it also has some limitations. Among others, our result holds even in finite markets. Precisely, we show that an efficient fully revealing equilibrium exists in finite markets, provided that the market size is sufficiently large and the set of procurers' types is finite. This is in contrast with the generic inefficiency result with reserve price posting (see footnote 3). In finite markets, each procurer has non-negligible market power, leading to certain distortions (e.g., too low reserve prices in procurement auctions). In our cheap-talk model, the inability to commit restricts the scope of a procurer's exerting his market power, which helps improve market efficiency and might contribute to the prevalence of cheap talk in market design.⁸

Although we focus on competing auctions environments, the insights from our analysis have implications beyond the narrow confines of our model. To begin with, the equivalence between

⁸This result offers a unique perspective to a practical issue in real-world market design. In several markets, whether initial calls (advertisements) should be cheap talk (hidden reserve prices) or contain binding reserve prices remains an active market design issue. For example, the institutions within the EU Public Procurement Learning Lab debate whether to adopt hidden or public reserve prices in their respective markets (see Piga and Zanza, 2004). Similarly, while most private freelance sites (e.g., MyHammer.co.uk) allow only signals of the "desired budget", some online applications such as Qluso experiment with binding reserve prices. Our large market result that cheap talk can be as efficient as reserve price posting might explain their co-existence. In addition, our finite market result that cheap talk can outperform reserve price posting can be used to explain why cheap talk might retain a real edge in some markets.

competitive search (price posting) and competing auctions has been established by several studies (see footnote 2). Therefore, our result immediately implies that cheap talk, combined with a certain trading mechanism such as the first-price auction, can also serve as an alternative to price commitments in competitive search models. More generally, pre-match communication via public messages plays a crucial role for the functioning of many decentralized markets. It may involve commitment to the terms of trade, as in competing auctions and competitive search, or not, as in this paper. Although both are plausible, the latter has received disproportionately less attention. Our result is only one further step toward understanding the role of cheap-talk competition in decentralized markets, and we expect to see more progresses along this line.

Our work highlights how cheap talk can facilitate trade in decentralized market environments and interacts with different trading mechanisms. In this regard, among many studies following the seminal contribution by Crawford and Sobel (1982), it is particularly close to Menzio (2007), who considers the same cheap talk environment as ours, but adopts bilateral bargaining as the trading mechanism.⁹ In his model, some information can be transmitted, but there never exists a fully revealing equilibrium, and constrained efficiency cannot be achieved. Our result implies that imperfect communication and the failure of constrained efficiency in his model are driven by the specific trading protocol adopted, not by the non-binding nature of cheap talk per se. More broadly, our paper is also related to a subset of the cheap-talk literature that study how cheap talk influences outcomes in strategic trading environments.¹⁰ Our analysis is qualitatively different from those in existing work, mainly because of the nature of the subsequent competition and the focus on large decentralized markets. Nevertheless, as explained in Section 5, equilibrium selection criteria developed in the literature naturally apply to our environment and provide a formal justification for our focus on the fully revealing equilibrium.

The remainder of the paper is organized as follows. Section 2 introduces the basic model. Section 3 presents and proves our main result. In Section 3, we also discuss the effects of adopting an alternative trading protocol in our framework. We consider a finite market version of the model in Section 4 and conclude by discussing the robustness and the limitations of our results in various dimensions in Section 5.

⁹Two more related contributions are Kim (2012) and Albrecht, Gautier and Vroman (2014). The former considers a common value counterpart to this paper and shows that, although there does not exist a fully revealing equilibrium, there may exist a partially revealing equilibrium. The partially revealing equilibrium is more efficient than the babbling equilibrium, but is not constrained efficient. The latter considers a model of directed search with limited commitment. In their setup, a seller is committed to her asking price, because she must accept any price above, but not fully committed, because she is free to accept any price below.

¹⁰Farrell and Gibbons (1989) and Matthews and Postlewaite (1989) are seminal contributions. See Ye (2007), Chakraborty and Harbaugh (2010), Gardete (2013), Quint and Hendricks (2014), and Koesler and Skreta (2014) for some recent contributions.

2 The Model

We lay out our model in the context of procurement markets. The model can be easily translated into other contexts, such as product or labor markets.

Agents and preferences. There is a continuum of procurers in the market, whose measure is normalized to 1. Each procurer has one project to complete. Once completed, the project yields the procurer utility v . The value of the project v is private information to each procurer, but it is commonly known that the values are distributed over a set $\mathcal{V} \subset \mathbb{R}_+$ according to the distribution function F : $F(v)$ denotes the measure of procurers whose types are below v . The set \mathcal{V} has a finite minimal element \underline{v} and a finite maximal element \bar{v} , where $\underline{v} < \bar{v}$.

Each procurer must hire a contractor to run the project. There is a continuum of contractors available in the market, whose measure is given by β . Each contractor has the capacity to run only one project. For analytical tractability as well as the transparency of the driving forces for our results, we focus on the case where contractors are homogeneous and have an identical opportunity cost of working $c(\leq \underline{v})$. We explain how to accommodate contractor heterogeneity in the online appendix. The notation becomes significantly more complicated, but the main logic goes through unchanged.

All agents maximize their expected utility and are risk neutral. If a procurer with project value v hires a contractor at price p , then the procurer receives utility $v - p$, while the contractor obtains utility $p - c$. All agents who fail to match receive utility 0.

Market interaction. Agents match according to the following sequence. *Communication*: each procurer publicly announces a cheap-talk message drawn from the set \mathcal{M} . For our main purpose, it suffices to assume that \mathcal{M} coincides with \mathcal{V} . *Search*: after observing all announced messages, each contractor selects a procurer. *Trading*: each contractor submits a quote specifying his asking price to the selected procurer. We assume that contractors do not observe whether there are other contractors who selected the same procurer. Subsequently, each procurer decides which contractor, if any, to hire at the quoted price.

Remark 1 *The overall market interaction closely resembles that of the canonical competing auctions model. The main difference is that announced reserve prices are binding obligations that directly restrict the subsequent trading game, while cheap-talk messages in our model can affect only contractors' beliefs about procurers' types and, therefore, only indirectly their search and bidding actions. Note that messages can take various forms: a message may represent the asking price commonly observed in housing markets and on-line marketplaces, an intended budget sometimes announced in private procurements, a wage range often announced in labor markets, or a verbal description of the value of the job.*

Strategies and equilibrium. We take the large market game approach and assume that each agent's payoff depends only on his own action and the distributions of other agents' actions. The interaction between an individual agent's action and aggregate distributions is specified in a familiar way in the literature on competitive search and competing auctions.¹¹

Procurers' communication strategies are described by a joint distribution function $Q : \mathcal{M} \times \mathcal{V} \rightarrow [0, 1]$, where $Q(m, v)$ denotes the measure of procurers whose types are below v and announce a message below m .¹² Contractors observe procurers' announcements, but not their types. To accommodate this, denote by $Q_{\mathcal{M}}$ the marginal distribution of Q over \mathcal{M} , so that $Q_{\mathcal{M}}(m)$ represents the measure of procurers who announce a message below m , and by $\mathcal{Q}_{\mathcal{M}}$ the set of all feasible marginal distributions. Contractors' (search and bidding) strategies are described by a joint distribution function $P : \mathcal{M} \times \mathbb{R}_+ \times \mathcal{Q}_{\mathcal{M}} \rightarrow [0, 1]$, where $P(m, b; Q_{\mathcal{M}})$ denotes the proportion of contractors who choose a procurer below m and quote a price below b given procurers' announcements $Q_{\mathcal{M}}$. Contractors' beliefs about procurers' types are represented by a function $\mu : \mathbb{R}_+ \times \mathcal{M} \times \mathcal{Q}_{\mathcal{M}} \rightarrow [0, 1]$, where $\mu(b; m, Q_{\mathcal{M}})$ denotes the proportion of procurers whose types are *strictly* below b among those who announced message m given $Q_{\mathcal{M}}$. Procurers' optimal hiring strategies are straightforward, so we omit their formal description: each procurer hires the contractor who offers the lowest price, provided that the price is below her own value. For completeness, we assume that if the lowest price is submitted by multiple contractors, then the procurer selects each of them with equal probability.

Fix a pair of distributions $(Q_{\mathcal{M}}, P)$ and a belief system μ . To formally specify agents' payoffs, it is convenient to define the queue length for each pair (m, b) . For each b , define $\lambda(\cdot, b; Q_{\mathcal{M}}, P)$ to be the Radon-Nikodym derivative of $\beta P(\cdot, b; Q_{\mathcal{M}})$ with respect to $Q_{\mathcal{M}}$.¹³ In other words, with a slight abuse of notation,

$$\lambda(m, b; Q_{\mathcal{M}}, P) \equiv \frac{\beta dP(m, b; Q_{\mathcal{M}})}{dQ_{\mathcal{M}}(m)}. \quad (1)$$

Economically, the queue length $\lambda(m, b; Q_{\mathcal{M}}, P)$ represents the expected number of quotes (equivalently, contractors) weakly below b for a procurer with message m . We denote by $\lambda_-(m, b; Q_{\mathcal{M}}, P)$ the queue length of quotes strictly below b (i.e., $\lambda_-(m, b; Q_{\mathcal{M}}, P) \equiv \lim_{b' \nearrow b} \lambda(m, b'; Q_{\mathcal{M}}, P)$), by $d\lambda(m, b; Q_{\mathcal{M}}, P)$ that of quotes exactly equal to b (i.e., $d\lambda(m, b; Q_{\mathcal{M}}, P) \equiv \lambda(m, b; Q_{\mathcal{M}}, P) - \lambda_-(m, b; Q_{\mathcal{M}}, P)$), and by $\lambda_{\mathcal{M}}(m; Q_{\mathcal{M}}, P)$ the queue length of all quotes associated with message

¹¹Our probability assignment and subsequent payoff formulas particularly resemble those in Shi (2002), Shimer (2005), and Peters (2010), each of whom studies a directed search model with ex ante worker (contractor) heterogeneity. In the current context, contractors' ex ante heterogeneity is replaced by ex post heterogeneity in their quotes.

¹²Throughout the paper, we restrict attention to pure communication strategies. This implies that there exists a function $M : \mathcal{V} \rightarrow \mathcal{M}$ such that (m, v) belongs to the support of Q if and only if $M(v) = m$.

¹³As usual, the Radon-Nikodym derivative is well-defined only up to a zero measure set. We assume a selection rule that uniquely selects it and obeys the right-continuity of λ everywhere.

m (i.e., $\lambda_{\mathcal{M}}(m; Q_{\mathcal{M}}, P) \equiv \lim_{b \rightarrow \infty} \lambda(m, b; Q_{\mathcal{M}}, P)$).

Following the standard specification in the literature, we assume that the number of quotes associated with each (m, b) follows a Poisson distribution with parameter $\lambda(m, b; Q_{\mathcal{M}}, P)$.¹⁴ In particular, the probability that a procurer with message m does not receive any quote weakly below b is given by $e^{-\lambda(m, b; Q_{\mathcal{M}}, P)}$. Combining this with the fact that the procurer does not accept any price above v gives the procurer's expected payoff of the following form:

$$V(m, v; Q_{\mathcal{M}}, P) = \int_{b \leq v} (v - b) d(1 - e^{-\lambda(m, b; Q_{\mathcal{M}}, P)}). \quad (2)$$

Now consider a contractor who selects a procurer with message m and offers b to the procurer. He has a chance of winning only when no other contractor quotes strictly below b , which occurs with probability $e^{-\lambda(m, b; Q_{\mathcal{M}}, P)}$. In addition, he needs to be selected among those who offer the same price, whose probability is given by $(1 - e^{-d\lambda(m, b; Q_{\mathcal{M}}, P)})/d\lambda(m, b; Q_{\mathcal{M}}, P)$.¹⁵ We follow the convention that this probability is equal to 1 if there is no atom at b . Finally, even if b is the lowest price, it is accepted only when it does not exceed the procurer's value and, therefore, with probability $1 - \mu(b; m, Q_{\mathcal{M}})$. Combining all of these gives the contractor's expected payoff of the following form:

$$U(m, b; Q_{\mathcal{M}}, P, \mu) = e^{-\lambda(m, b; Q_{\mathcal{M}}, P)} \frac{1 - e^{-d\lambda(m, b; Q_{\mathcal{M}}, P)}}{d\lambda(m, b; Q_{\mathcal{M}}, P)} (1 - \mu(b; m, Q_{\mathcal{M}}))(b - c). \quad (3)$$

If there is no atom on b , then $U(m, b; Q_{\mathcal{M}}, P, \mu)$ reduces to $e^{-\lambda(m, b; Q_{\mathcal{M}}, P)}(1 - \mu(b; m, Q_{\mathcal{M}}))(b - c)$.

The above payoffs are well-defined only on the support of G . In standard competitive search and competing auctions, it is crucial to specify agents' out-of-equilibrium beliefs and resulting payoffs. In the current model with cheap talk, it is not necessary at this point. As well-known in the literature, out-of-equilibrium beliefs can always be assigned in a way not to unravel a given equilibrium. For example, all out-of-equilibrium messages can be interpreted as identical

¹⁴The micro-foundation behind this standard specification is as follows. Suppose a finite number of procurers posted an identical message m , a finite number of quotes are sent to those procurers, and the ratio of those bids to the procurers is given by λ . Further assume that contractors cannot coordinate one another and, therefore, their quotes are randomly distributed across the procurers. Then, the number of quotes each procurer receives follows a binomial distribution. As the numbers of procurers and contractors tend to infinity, while their ratio remains at λ , the binomial distribution converges to a Poisson distribution with parameter λ . See, e.g., Peters and Severinov (1997), Julien, Kennes and King (2000), and Burdett, Shi and Wright (2001) for formal arguments and excellent discussions. In Section 4 and the online appendix, we formally consider a finite version of the model and prove the convergence of finite-market payoffs to those assumed here.

¹⁵This probability incorporates the fact that his probability of winning is $1/n$ when there are $n - 1$ other contractors who offer the same price. Formally, the expression comes from the following fact:

$$\sum_{n=1}^{\infty} \frac{e^{-d\lambda}(d\lambda)^{n-1}}{(n-1)!} \frac{1}{n} = \frac{1}{d\lambda} \sum_{n=1}^{\infty} \frac{e^{-d\lambda}(d\lambda)^n}{n!} = \frac{1}{d\lambda} \left(\sum_{n=0}^{\infty} \frac{e^{-\lambda}(d\lambda)^n}{n!} - e^{-d\lambda} \right) = \frac{1 - e^{-d\lambda}}{d\lambda}.$$

to one equilibrium message, or one can simply change the strategy profile (e.g., through uniform randomization) so that there is no unused message. The meaning and interpretation of out-of-equilibrium messages play a crucial role for equilibrium selection, which we discuss in Section 5 and formally study in the online appendix.

Without further refinements, equilibrium simply requires that each action in the support of the equilibrium distributions should be a best response, and beliefs should be consistent with the given strategy profile whenever possible.

Definition 1 *An equilibrium of the market game is a pair of distribution functions Q and P and a belief system μ such that*

- *Procurer Optimality: $(m, v) \in \text{supp}Q$ implies $V(m, v; Q_{\mathcal{M}}, P) \geq V(m', v; Q_{\mathcal{M}}, P), \forall m' \in \mathcal{M}$.*
- *Contractor Optimality: $(m, b) \in \text{supp}P(\cdot, \cdot; Q_{\mathcal{M}})$ implies $U(m, b; Q_{\mathcal{M}}, P, \mu) \geq U(m', b'; Q_{\mathcal{M}}, P, \mu), \forall (m', b') \in \mathcal{M} \times \mathbb{R}_+$.*
- *Belief Consistency: for each $m \in \mathcal{M}$, $\mu(\cdot; m, Q_{\mathcal{M}})$ is a conditional probability distribution over \mathcal{V} , obtained from Q .¹⁶*

For most parts of the paper, we focus on fully revealing equilibria in which all procurers' types are fully revealed through their announcements. Formally,

Definition 2 *A fully revealing equilibrium is an equilibrium (Q, P, μ) in which if (m, v) is in the support of Q , then (m, v') is not in the support of Q for any $v' \neq v$.*

In what follows, without loss of generality, we restrict attention to the fully revealing equilibrium in which each procurer truthfully announces her own type. In the fully revealing equilibrium, the distribution of announced messages $Q_{\mathcal{M}}$ coincides with the distribution of procurer types F . In addition, contractors' beliefs about procurers' types μ are trivial. From now on, for notational simplicity, we drop $Q_{\mathcal{M}}$, P , and μ from the arguments of all equilibrium objects and use, for example, $\lambda(m, b)$, $\lambda_{\mathcal{M}}(m)$ and $V(m, v)$ to denote $\lambda(m, b; Q_{\mathcal{M}}, P)$, $\lambda_{\mathcal{M}}(m; Q_{\mathcal{M}}, P)$ and $V(m, v; Q_{\mathcal{M}}, P)$, respectively.

3 The Main Result

We begin by establishing the constrained efficient benchmark. We then characterize the fully revealing equilibrium of the model. We show that the equilibrium always exists and achieves constrained efficiency. Finally, we discuss the effects of adopting an alternative trading protocol in our framework.

¹⁶Formally, $\mu(b; m, Q_{\mathcal{M}})$ is the left-continuous Radon-Nikodym derivative of $Q(\cdot, b)$ with respect to $Q_{\mathcal{M}}$.

3.1 Constrained Efficiency

We consider the (utilitarian) social planner who wishes to maximize gross surplus, knowing procurers' types (or dictating truthful revelation) but being constrained to search frictions as in the decentralized market. Specifically, the social planner is restricted to match procurers and contractors only probabilistically. In tight analogy with the equilibrium variables, we denote by $P^*(v)$ the proportion of contractors who are sent to procurers with types below v and by $\lambda_{\mathcal{M}}^*(v)$ the queue length of all contractors for a type v procurer (formally defined as the Radon-Nikodym derivative of βP^* with respect to F). The planner is constrained by the same Poisson trading probabilities as in the decentralized market. In particular, the probability that no contractor approaches a type v procurer is given by $e^{-\lambda_{\mathcal{M}}^*(v)}$. Under this constraint, the planner wishes to maximize gross output with the optimal choice of P^* from the set of probability distributions over \mathcal{V} :

$$\max_{P^*} \int_{\mathcal{V}} (1 - e^{-\lambda_{\mathcal{M}}^*(v)})(v - c), \text{ subject to } P^*(\bar{v}) \leq 1.$$

Since there is a one-to-one relationship between P^* and $\lambda_{\mathcal{M}}^*$, the maximization problem can be equivalently written as

$$\max_{\lambda_{\mathcal{M}}^*} \int_{\mathcal{V}} (1 - e^{-\lambda_{\mathcal{M}}^*(v)})(v - c), \text{ subject to } \int_{\mathcal{V}} \lambda_{\mathcal{M}}^*(v) dF(v) \leq \beta.$$

The objective function is concave, while the constraint is linear in $\lambda_{\mathcal{M}}^*$. Therefore, the first-order condition of the corresponding Lagrangian with multiplier μ and the binding constraint fully characterize the planner's optimal solution.

Proposition 1 *There exists a unique solution to the social planner's problem: the optimal queue length $\lambda_{\mathcal{M}}^*(v)$ for type v procurers is given by*

$$e^{-\lambda_{\mathcal{M}}^*(v)}(v - c) \leq \mu, = \mu \text{ if } \lambda_{\mathcal{M}}^*(v) > 0,$$

where μ is the value that satisfies $\int_{c+\mu}^{\bar{v}} \ln\left(\frac{v-c}{\mu}\right) dF(v) = \beta$.

Proof. See the appendix. ■

The shadow value μ measures the marginal cost of assigning one more contractor to each procurer. The corresponding marginal benefit stems from the fact that each procurer may not be selected by any contractor. Since a contractor contributes to social surplus if and only if he selects a procurer who would otherwise not have any contractor, the marginal benefit is equal to $e^{-\lambda_{\mathcal{M}}^*(v)}(v - c)$ for a type v procurer. Constrained efficiency requires that the marginal benefit and the marginal cost coincide for each procurer type such that $v - c > \mu$. Naturally, the optimal $\lambda_{\mathcal{M}}^*$ is increasing, meaning that procurers with higher values attract more contractors on average and have a higher probability of completing their projects.

3.2 Fully Revealing Equilibrium

We say that an equilibrium (Q, P, μ) is constrained efficient if the total surplus in the economy equals the surplus in the planner's solution. This arises in the decentralized market if and only if the equilibrium total queue lengths coincide with the planner's solution (i.e., $\lambda_{\mathcal{M}}(v) = \lambda_{\mathcal{M}}^*(v)$ for all $v \in \mathcal{V}$) and only acceptable price quotes are made (i.e., no price quote exceeds v). The main result of the paper is as follows:

Theorem 1 *The fully revealing equilibrium always exists and is constrained efficient.*

The theorem consists of two parts, the existence of the fully revealing equilibrium and constrained efficiency of the equilibrium. The efficiency result, as explained shortly, can be obtained by combining existing insights in search theory and auction theory, while the existence result is, to our knowledge, based on new insights. In order to highlight our unique contribution, we first prove the existence result and then illustrate the efficiency result.

3.2.1 Contractors' Bidding Strategies in the Fully Revealing Equilibrium

The following lemma provides a characterization of the equilibrium queue lengths for each procurer type, $\lambda(v, \cdot)$, using necessary conditions for contractors' equilibrium bidding strategies. It identifies the minimal and the maximal prices, and shows that the bid distribution for each procurer type has a full convex support and no mass point (i.e., $\lambda(v, \cdot)$ is a continuous and strictly increasing function). No mass point exists, because a contractor would always find it profitable to undercut any mass point. Given this, the convexity of the support follows from contractors' continuous trade-off between price and trading probability.

Lemma 1 *Denote by $\underline{b}(v)$ and $\bar{b}(v)$ the minimal element and the maximal element of the support of $\lambda(v, \cdot)$, respectively. In the fully revealing equilibrium, for each $v \in \mathcal{V}$ such that $\lambda_{\mathcal{M}}(v) > 0$, $\underline{b}(v) = c + e^{-\lambda_{\mathcal{M}}(v)}(v - c)$, $\bar{b}(v) = v$, and $\lambda(v, \cdot)$ is a continuous and strictly increasing function such that for any $b \in [\underline{b}(v), \bar{b}(v)]$,*

$$U(v, b) = e^{-\lambda(v, b)}(b - c) = U(v, \bar{b}(v)) = e^{-\lambda_{\mathcal{M}}(v)}(v - c).$$

Proof. See the appendix. ■

Applying Lemma 1 to (2), the expected payoff of a type v procurer in the fully revealing equilibrium is given by

$$V(v, v) = (1 - e^{-\lambda_{\mathcal{M}}(v)} - \lambda_{\mathcal{M}}(v)e^{-\lambda_{\mathcal{M}}(v)})(v - c).$$

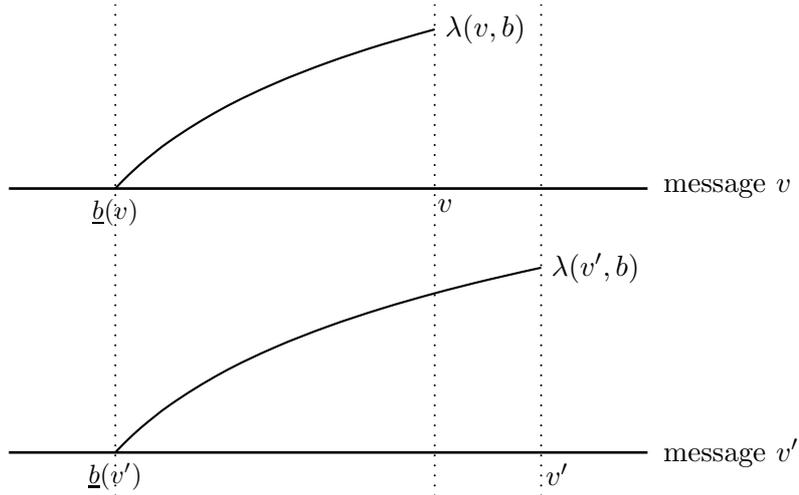


Figure 1: The figure depicts the expected queue lengths $\lambda(v, \cdot)$ and $\lambda(v', \cdot)$, where $v < v'$.

One can intuitively understand the payoffs through the revenue equivalence between first- and second-price auctions.¹⁷ Suppose a type v procurer runs the second-price procurement auction with reserve price equal to v . Then, a contractor obtains the entire surplus $v - c$ if and only if there is no other competitor, while the procurer extracts the entire surplus if and only if there are at least two contractors. Therefore, a contractor's expected payoff is equal to $e^{-\lambda_{\mathcal{M}}(v)}(v - c)$, while the procurer's expected payoff is equal to $(1 - e^{-\lambda_{\mathcal{M}}(v)} - \lambda_{\mathcal{M}}(v)e^{-\lambda_{\mathcal{M}}(v)})(v - c)$.

3.2.2 Existence of the Fully Revealing Equilibrium

We now establish the existence of the fully revealing equilibrium. Specifically, we show that in the fully revealing equilibrium, for any $v, v' \in \mathcal{V}$ such that $v < v'$, a type v procurer has no incentive to deviate and announce v' , and vice versa for a type v' procurer. We focus on the only interesting case where both $\lambda_{\mathcal{M}}(v')$ and $\lambda_{\mathcal{M}}(v)$ are strictly positive.¹⁸

First, notice that the minimal price quotes for v and v' have to be identical, that is, $\underline{b}(v) = \underline{b}(v')$ (see Figure 1). This follows from the equilibrium contractor optimality: a contractor must be indifferent between selecting a type v procurer and a type v' procurer. Therefore,

$$U(v, \underline{b}(v)) = \underline{b}(v) - c = \underline{b}(v') - c = U(v', \underline{b}(v')).$$

From now on, we denote by \underline{b} the common minimal price (i.e., $\underline{b} \equiv \underline{b}(v) = \underline{b}(v')$).

Second, we argue that the price distribution facing a procurer (i.e., $1 - e^{-\lambda(m, b)}$) is independent

¹⁷Revenue equivalence continues to hold with a stochastic number of bidders, provided that bidders are risk neutral. See, for example, McAfee and McMillan (1987).

¹⁸ $\lambda_{\mathcal{M}}(v') = 0$ implies $\lambda_{\mathcal{M}}(v) = 0$, that is, neither type receives any price quote. If $\lambda_{\mathcal{M}}(v) = 0$, then, for the reason given shortly, all price quotes exceed v (i.e., $\underline{b}(v') > v$). Therefore, a type v procurer has no incentive to deviate to v' .

of her announcement over the interval $[\underline{b}, v]$. In other words, in Figure 1, $\lambda(v, b) = \lambda(v', b)$ for any $b \in [\underline{b}, v]$. To see this, notice that a contractor's expected payoff depends only on his own price and the probability of winning, not directly on the procurer's type, as long as he quotes a price below the procurer's value. Therefore, for any price b below $v (< v')$, a contractor can be indifferent between v and v' only when his probability of winning is equalized between them. Since the probability of winning ($e^{-\lambda(m,b)}$) is isomorphic to the expected queue length $\lambda(m, b)$, this also means that the expected queue lengths for a procurer are independent of her announcement. Formally, given $b \in [\underline{b}, v]$, contractor optimality requires that a contractor must be indifferent between v and v' :

$$U(v, b) = e^{-\lambda(v,b)}(b - c) = e^{-\lambda(v',b)}(b - c) = U(v', b),$$

which implies $\lambda(v, b) = \lambda(v', b)$. Note that the total queue length does depend on a procurer's announcement (that is, $\lambda_{\mathcal{M}}(v) \neq \lambda_{\mathcal{M}}(v')$): more quotes are sent to procurers with message v' . However, contractors quote relatively higher prices to those procurers. For b below $v (< v')$, these two effects exactly cancel each other out. In other words, additional quotes following v' , relative to v , are only above v . Therefore, for $b \in [\underline{b}, v]$, the probability that a procurer hires a contractor at a price below b is independent of her announcement.

Finally, we consider the effects of a procurer's deviation on the support of contractors' bidding strategies. If a type v procurer deviates to v' , then the maximal price quote changes from $\bar{b}(v) = v$ to $\bar{b}(v') = v'$, and vice versa. We argue that this effect does not make any deviation strictly profitable. First, consider a type v' procurer's deviation to v (downward deviation). In this case, the procurer no longer receives price quotes between v and v' . Consequently, the deviation is strictly unprofitable. Now, consider a type v procurer's deviation to v' (upward deviation). In this case, the procurer may receive price quotes between v and v' . However, those prices exceed her value and, therefore, are irrelevant to her expected payoff: she will simply reject all those prices. Consequently, the (upward) deviation does not strictly increase her expected payoff either. Overall, each procurer strictly prefers truth-telling to downward deviations and weakly to upward deviations.¹⁹

¹⁹Formally, combining (2) with the fact that $\lambda(m, b)$ is independent of m ,

$$V(v', v) = \int_{\underline{b}}^{\min\{v', v\}} (v - b)d(1 - e^{-\lambda(v', b)}) = \int_{\underline{b}}^{\min\{v', v\}} (v - b)d(1 - e^{-\lambda(v, b)}).$$

Therefore, the difference between the equilibrium payoff and the deviation payoff is equal to

$$V(v, v) - V(v', v) = \int_{\min\{v', v\}}^v (v - b)d(1 - e^{-\lambda(v, b)}).$$

Obviously, this expression is strictly positive if $v > v'$ (downward) and equal to 0 if $v < v'$ (upward).

3.2.3 Constrained Efficiency of the Fully Revealing Equilibrium

We now show that the fully revealing equilibrium is constrained efficient. Specifically, we prove that the equilibrium total queue lengths $\lambda_{\mathcal{M}}(v)$ coincide with the social planner's solution $\lambda_{\mathcal{M}}^*(v)$.

The total queue length for a type v procurer can be positive only when a contractor cannot do strictly better than selecting the procurer. In order to derive a formal condition, define the equilibrium utility $u \equiv \max_{v,b} U(v,b)$. Following the literature, we refer to u as contractors' *market utility*. Contractor optimality requires that

$$\lambda_{\mathcal{M}}(v) > 0 \text{ only when } U(v,v) = e^{-\lambda_{\mathcal{M}}(v)}(v-c) = u.$$

Rewriting this with the feasibility condition (which corresponds to $P(\bar{v},\bar{v}) = 1$) yields

$$e^{-\lambda_{\mathcal{M}}(v)}(v-c) \leq u, = u \text{ if } \lambda(v) > 0, \text{ and } \int_{\mathcal{V}} \lambda_{\mathcal{M}}(v) dF(v) = \beta.$$

The conditions are identical to those for constrained efficiency in Proposition 1, once u is replaced with μ . It is an immediate corollary that the fully revealing equilibrium outcome is unique and achieves constrained efficiency.

This efficiency result can be understood by combining existing insights in the literature, in particular, constrained efficiency of the second-price auction and the revenue equivalence between first- and second-price auctions. It is well-known that an agent makes a socially efficient choice if he receives as much as he contributes to social surplus. In a two-sided market environment as ours, but without private information, Julien, Kennes and King (2005, 2006) show that if each procurer runs the second-price auction with reserve price equal to her value, then constrained efficiency is achieved. Intuitively, a contractor generates surplus $v - c$ if there is no other contractor, and nothing if there is at least one other contractor. He fully internalizes this if he extracts the entire surplus when he is the only contractor in a meeting and receives no surplus when there are other contractors.²⁰ This ex post outcome can be exactly replicated if each procurer runs the efficient second-price auction. Agents' ex-post payoffs in our game are not identical to those from the second-price auction. However, by revenue equivalence, their interim payoffs are identical, which is sufficient for efficient search by contractors.

²⁰This is partly due to the urn-ball matching assumption. Under urn-ball matching, a contractor's meeting a procurer does not affect other contractors' chances of meeting the same procurer. In other words, the probability that a procurer is selected by n additional contractors is independent of whether the procurer is selected by a particular contractor or not. If this property is violated (i.e., a contractor's meeting a procurer increases or decreases the procurer's chance of meeting other contractors), then his marginal social contribution needs to be adjusted in order to reflect relevant externalities. In other words, a contractor may receive more or less than his social contribution to the local interaction. See Eeckhout and Kircher (2010b) and Lester, Visschers and Wolthoff (2014) for more detailed discussions.

3.3 Alternative Trading Protocols

In principle, any trading protocol can be considered at the trading stage of our model. Although it does not seem feasible to analyze all possible trading protocols, the preceding analysis allows us to shed some light on the effects of adopting an alternative trading protocol in our framework. In particular, the following result is straightforward from the discussions above.²¹

Proposition 2 *In the market communication game with an alternative trading protocol, an equilibrium is constrained efficient if and only if (i) it is fully revealing for all procurer types such that $\lambda_{\mathcal{M}}^*(v) > 0$ and (ii) for some non-negative constant a , the expected payoff of a contractor who selects a type v procurer is equal to $U(v) = ae^{-\lambda_{\mathcal{M}}^*(v)}(v - c)$ for all $v \in \mathcal{V}$.*

Proof. See the appendix. ■

We use this result to show that two common trading protocols, one corresponding to the bilateral bargaining considered in Menzio (2007) and the other to the second-price auction (equivalently, ex post bidding), fail to deliver constrained efficiency in our framework. Assuming full revelation, in the former, a contractor obtains a fixed fraction of the surplus, provided that he is selected by the procurer. In the latter, he asks for the whole surplus if there is no other contractor and none if there are other contractors.

Corollary 1 (Bilateral bargaining) *Suppose that once contractors select procurers, each procurer randomly selects a contractor and then plays a bargaining game with the contractor in which the contractor's bargaining power is given by $\alpha \in [0, 1]$. The trading protocol fails to achieve constrained efficiency.*

Proof. Suppose the efficient fully revealing equilibrium exists. Consider a contractor who selects a type v procurer. He is selected as the bargaining partner with probability $(1 - e^{-\lambda_{\mathcal{M}}^*(v)})/\lambda_{\mathcal{M}}^*(v)$ (see footnote 15), and thus his expected payoff is given by $\alpha(v - c)(1 - e^{-\lambda_{\mathcal{M}}^*(v)})/\lambda_{\mathcal{M}}^*(v)$. In this case, the second condition in Proposition 2 can be satisfied only when $\alpha = a = 0$. This implies that the expected payoff of a type v procurer who announces v' is equal to $V(v', v) = (1 - e^{-\lambda_{\mathcal{M}}^*(v')})(v - c)$. Since this is strictly increasing in v' , the fully revealing equilibrium cannot exist. ■

Corollary 2 (Ex post bidding) *Consider the same competitive bidding environment as in the baseline model, but now suppose that contractors observe the number of competitors before they quote prices. The fully revealing equilibrium cannot exist in such a setting.*

²¹Notice that it suffices that contractors' expected payoffs are *proportionally* revenue-equivalent to those from our baseline model. It is easy to show that the exact revenue equivalence (i.e., $a = 1$) is required if one would allow for endogenous entry on either side: procurers might enter at some fixed cost, or contractors could have a search cost.

Proof. Suppose the efficient fully revealing equilibrium exists. In the equilibrium, each contractor quotes as much as the announced value to each procurer if alone, while the price is driven down to c if there are at least two contractors, following the usual Bertrand-competition logic. Since a procurer can reject a price above his value, the expected payoff of a type v procurer who announces $v'(> v)$ is given by $V(v', v) = (1 - e^{-\lambda_{\mathcal{M}}^*(v')} - \lambda_{\mathcal{M}}^*(v')e^{-\lambda_{\mathcal{M}}^*(v')})(v - c)$. The term inside the parenthesis is the probability of receiving at least two contractors, which is strictly increasing in v' . This implies that the fully revealing equilibrium cannot exist. ■

The analysis of the main model also suggests a condition for the existence of the fully revealing equilibrium. To this end, restrict attention to trading protocols that are revenue-equivalent to the first-price auction conditional on full revelation, so that the second condition in Proposition 2 is satisfied and we only need to check the first condition. In addition, suppose that a procurer who deviates from truthful announcement faces the same price distribution as the type she pretends to be, but can reject prices above her value.²² The possibility of rejection breaks the revenue equivalence off the equilibrium paths, which results in different incentive properties for different trading mechanisms.

Once attention is restricted to this class of trading protocols, the existence of the fully revealing equilibrium boils down to whether the price dispersion for each procurer type is more or less dispersed than that induced by the first-price auction. To understand this claim, recall that a procurer faces the following trade-off at the communication stage. On the one hand, if she overstates her value, then she can attract relatively more contractors. On the other hand, each contractor quotes a relatively higher price. If a procurer must hire a contractor after the deviation, only the expected price matters and truth-telling becomes optimal, as in canonical competing auctions models. If communication is cheap talk, then a procurer can refuse to hire a contractor. This shields her from high prices and, therefore, limits the negative effect of the deviation. The more dispersed prices are, the more a procurer can gain from low prices, while being shielded from high ones. The first-price auction in our baseline model serves as a tight benchmark, because the two effects exactly cancel each other out and procurers are indifferent over upward deviations. Consequently, if a trading protocol generates more price dispersion (e.g., the second-price auction), then the shielding effect dominates and truth-telling cannot be sustained. To the contrary, if a trading protocol yields less price dispersion, then truth-telling becomes strictly optimal, and thus the fully revealing equilibrium exists.²³

²²For example, our first-price auction model generates a continuous distribution $1 - e^{-\lambda(v, \cdot)}$ for a type v procurer, while the ex post bidding (second-price auction) above leads to a discrete distribution for each procurer type: c with probability $1 - e^{-\lambda_{\mathcal{M}}(v)} - \lambda_{\mathcal{M}}(v)e^{-\lambda_{\mathcal{M}}(v)}$ and v with probability $e^{-\lambda_{\mathcal{M}}(v)}$.

²³See the previous version of this paper for a formal proof. In the previous version, we also provide an example of trading protocol that induces a single price (i.e., a degenerate price distribution) for each procurer type. A strict, fully revealing equilibrium exists with the trading protocol.

4 Finite Markets

In this section, we present a finite market version of the model. As in other work, this provides a micro-foundation for the payoff specification used in the large market model. In addition, it allows us to examine the robustness of our large market results. This latter aspect is more important than in other models, because procurers' truth-telling incentives are weak in our large market model and, therefore, can be potentially unravelled in finite markets. We focus on illustrating key ideas, relegating all formalities, including formal descriptions of agents' strategies and market equilibrium, a characterization of the constrained efficiency benchmark, and the equilibrium convergence to the large market equilibrium, to the online appendix.

Setup. There are N procurers and M contractors, where M is the smallest integer such that $M \geq \beta N$. At the beginning of the game, procurers' values are independently and identically drawn from \mathcal{V} . We restrict attention to the case where the set of procurer types \mathcal{V} is finite: $\mathcal{V} \equiv \{v^1, \dots, v^I\}$, where $v^i < v^{i+1}$ for each $i \leq I - 1$.²⁴ We denote by v_n procurer n 's type, by \vec{v} a profile of all procurers' types, and by \vec{v}_{-n} a profile of all procurers' types save procurer n 's. All other specifications of the model are identical to those in Section 2.

We focus on perfect Bayesian equilibria in which each procurer truthfully announces her type. As in other finite models, we further restrict attention to the symmetric equilibrium in which all contractors play an identical search and bidding strategy. We denote by $p_n(b; \vec{v})$ the probability that each contractor selects procurer n and quotes a lower price than b and by $p_n(\vec{v})$ the probability that each contractor selects procurer n (i.e., $p_n(\vec{v}) \equiv \lim_{b \rightarrow \infty} p_n(b; \vec{v})$) given the announced value profile \vec{v} .²⁵ Finally, we denote by $u(\vec{v})$ contractors' market utility.

The result. The following result is a finite market analogue to Theorem 1:

Proposition 3 *In finite markets, the fully revealing equilibrium is always constrained efficient. The equilibrium exists, with strict incentives, whenever the market size (N) is sufficiently large.*

²⁴Note that we do not impose any restriction on \mathcal{V} in the large market model. The restriction on the type space in finite markets allows us to obtain clear-cut results on the fully revealing equilibrium. For the reason clarified shortly, if the set \mathcal{V} is continuous, then we would have to look for interval partitioned equilibria as in, for example, Crawford and Sobel (1982). It would still be possible to construct a sequence of equilibria in finite markets that converges to the fully revealing equilibrium in the corresponding large market, in the sense that the size of each partition element shrinks to zero, but an exact characterization is quite involved. Alternatively, using our results, one can approximate the fully revealing equilibrium with a continuous type space, with a double sequence of finite type sets that approach the continuous type space and finite markets whose size grows to infinity.

²⁵For notational parsimony, we use \vec{v} to represent the dependence of contractors' strategies on procurers' announcements. In general, this is not equivalent to conditioning only on the distribution of announced messages (which corresponds to $Q_{\mathcal{M}}$ in the large market model). In the online appendix, we introduce precise restrictions on agents' strategies and beliefs, so that contractors' strategies indeed depend only on the distribution of announced messages, as in the large market model. For now, it suffices to interpret \vec{v} as a reduced-form expression of the distribution of announced messages.

The result has two important implications. First, it demonstrates that our large market results have a desirable robustness property. The large market assumption is made only for analytical tractability, but is ultimately ad hoc. Although the outcome continuity in the market size typically guarantees the robustness of large market results, it cannot be taken for granted in our model, because agents' incentives over upward deviations are weak in large markets. The result proves that in large but finite markets, agents' incentives are well-aligned in a way not to disrupt our large market results. In fact, agents' incentives over upward deviations are strict in finite markets, which strengthens the case for the fully revealing equilibrium.

Second, the efficiency result contrasts with the inefficiency result in finite competing auctions (see footnote 3 and the online appendix). Our cheap-talk model achieves efficiency because the nature of cheap talk limits the scope for procurers' deviations, thereby weakening the market-power forces that drive inefficiencies with competition in finite markets.²⁶

Constrained efficiency. We first illustrate that the logic for constrained efficiency goes through unchanged even in finite markets. By the same reasoning as in Lemma 1, the expected payoff of a contractor who selects procurer n is equal to the probability that no other contractor selects procurer n times his monopolistic profit: $(1 - p_n(\vec{v}))^{M-1}(v_n - c)$. This implies that the following condition is necessary for contractor optimality:

$$(1 - p_n(\vec{v}))^{M-1}(v_n - c) \leq u(\vec{v}), = u(\vec{v}) \text{ if } p_n(\vec{v}) > 0. \quad (4)$$

Notice that, as in the large market setting, the left-hand side can be interpreted as the marginal social contribution generated by a contractor who selects procurer n : again, a contractor contributes to social surplus only when he selects a procurer who would otherwise not be able to produce. Therefore, his marginal social contribution is the probability that no other contractor selects the procurer $(1 - p_n(\vec{v}))^{M-1}$ times gains from trade $(v_n - c)$. The equilibrium condition suggests that a contractor selects a procurer only when he can generate at least as much social surplus as selecting any other procurer, which coincides with the condition for social optimum.

Existence of the fully revealing equilibrium. Suppose procurer n has project value v , but deviated and announced v' instead. The deviation affects the procurer's expected payoff in two ways. First, as in large markets, contractors adjust their search and bidding strategies

²⁶In this regard, our finite market result is closely related to Hernando-Veciana (2005). He considers a finite competing auctions model in which the type space is finite and, importantly, each procurer (seller) is restricted to commit to only a reserve price that belongs to the type space. He proves an analogous result to ours: provided that the market size is sufficiently large, there is an equilibrium in which each procurer sets a reserve price equal to her own value. The resulting market outcome coincides with our fully revealing outcome and, therefore, is constrained efficient. In both his model and our model, the efficiency result is driven by the limited scope for procurers' deviations imposed by the finite type assumption. The difference is that the limit to agents' deviations is exogenously given in his model, while it endogenously arises in our model, due to the nature of cheap talk.

($p_n(\cdot; v', \vec{v}_{-n})$ instead of $p_n(\cdot; v, \vec{v}_{-n})$). Second, unlike in large markets, the deviation changes contractors' market utility from $u(v, \vec{v}_{-n})$ to $u(v', \vec{v}_{-n})$, because each procurer is not negligible in a finite market. We show that the former has the same effects on the deviating procurer's expected payoff as in large markets, while the latter makes downward deviations more profitable and upward deviations less profitable.

Fix $\vec{v}_{-n} \in \mathcal{V}^{N-1}$ and consider the probability that the lowest price quote to procurer n is weakly below b , where b is in the support of $p_n(\cdot; \vec{v})$. If procurer n 's announcement is v_n and $p_n(\vec{v}) > 0$, then the probability is given by

$$1 - (1 - p_n(b; \vec{v}))^M = 1 - \left(\frac{u(\vec{v})}{b - c} \right)^{\frac{M}{M-1}}, \quad (5)$$

where the equality is due to (4). Notice that the probability would be independent of v_n if contractors' market utility $u(v_n, \vec{v}_{-n})$ were independent of v_n . This is the first effect mentioned above. When a procurer deviates upward, relatively more contractors would select the procurer, but each of them would quote a relatively higher price. As in large markets, these two effects cancel each other out. Combined with the fact that a procurer's announcement determines the upper bound of price quotes, this makes a procurer strictly prefer truth-telling to downward deviations and indifferent over upward deviations.

In finite markets, contractors' market utility $u(\vec{v})$ does depend on each procurer's announcement. It is easy to show that $u(\vec{v})$ is increasing in each of its arguments: higher values imply more surplus available in the market.²⁷ Combining this with (5), it follows that procurer n 's probability of receiving a price below b is decreasing in her own announcement v_n . Intuitively, a contractor would quote a higher price when he has a higher reservation value. Therefore, each procurer has an incentive to lower contractors' market utility by understating her own value.

Combining the two effects, it is clear that a procurer strictly prefers truth-telling to an upward deviation: the first effect makes a procurer indifferent over upward deviations, while the second one makes him strictly prefer truth-telling. The (un)profitability of downward deviations is ambiguous in general, because the two effects work in the opposite direction. However, when the market size N is sufficiently large, the former direct effect strictly dominates the latter market utility effect: the former effect is essentially independent of the market size, while each procurer's market power reduces as the market size increases and vanishes in the limit. We conclude that a procurer also strictly prefers truth-telling to downward deviations, provided that the market

²⁷From (4), one can derive the following equation for $u(v)$:

$$N - 1 = \sum_n \min \left\{ \left(\frac{u(\vec{v})}{v_n - c} \right)^{\frac{1}{M-1}}, 1 \right\}.$$

size is sufficiently large.

5 Robustness and Limitations

We conclude by discussing the robustness as well as the limitations of our main result in several other dimensions.

Equilibrium selection. As in other cheap-talk games, our market game also suffers from equilibrium multiplicity. Indeed, in the online appendix, we fully characterize the set of all interval partitional equilibria, which includes both the babbling equilibrium and the fully revealing equilibrium as special cases.²⁸ This raises the question of whether the fully revealing equilibrium is any more salient than other equilibria. In the online appendix, we show that, provided that the type space \mathcal{V} is finite, the fully revealing equilibrium is the only interval partitional equilibrium that satisfies *neologism proofness* in the sense of Farrell (1993). We also consider a modified model in which procurers face the cost of lying and prove that under a certain monotonicity assumption, the fully revealing communication strategy uniquely survives iterated deletion of weakly dominated strategies. Insofar as these are reasonable selection criteria in cheap-talk games, they point to the prominence of the fully revealing equilibrium in our model.

Risk aversion. Our main model generates price dispersion for each procurer type. This suggests that the fully revealing equilibrium would fail to achieve contained efficiency if agents are risk averse. Indeed, if all agents are strictly risk averse, then competitive search models with price (wage) posting Pareto dominates our baseline model, as the former yields a single price for each procurer type conditional on hiring a contractor.²⁹ It also raises the question of whether the fully revealing equilibrium continues to exist with risk-averse agents, because its existence relies on a particular form of price dispersion induced by the first-price auction. In the online appendix, we show that the existence of the fully revealing equilibrium is independent of agents' risk aversion: it always exists whether contractors or procurers (or both) are strictly risk averse. This suggests some weak notion of robustness: if agents are nearly risk neutral, then our market game can induce a nearly constrained efficient outcome.

²⁸An equilibrium is interval-partitional if each message is selected by an interval of types and none of these types selects any other message. Formally, an equilibrium is interval-partitional if there exists a disjoint partition $\{\Gamma(m)\}_{m \in \mathcal{M}}$ of the type space such that each $\Gamma(m)$ is an interval and for any $v \in \Gamma(m)$, $(m', v) \in \text{supp}Q$ only when $m' = m$. Each procurer's indifference over all upward deviations puts a particular structure on interval partitional equilibria. See Proposition A2 in the online appendix.

²⁹Even though there is no price dispersion for each procurer (firm), competitive search models also do not deliver constrained efficiency under risk aversion. Intuitively, there is too much price dispersion between employed and unemployed contractors (workers). This fact has been exploited to provide an efficiency rationale for a positive level of unemployment insurance benefits. See, for example, Acemoglu and Shimer (1999a) and Boostani, Gervais and Siu (2014).

Interim information. The information available to contractors after search but before bidding is crucial for our results. In our model, contractors offer a price without receiving any further information. As explained in Section 3.3, if each contractor learns the presence of competitors for the same procurer, then our fully revealing result unravels.³⁰ It also unravels if procurers' types are observable after search. In this case, each procurer has only the incentive to maximize the expected queue length, and thus their announcements cannot be informative.

Trading mechanism design by auctioneers. The literature on competing mechanism design endogenizes the trading mechanism used in markets by allowing each auctioneer to optimally design her own trading mechanism and publicly announce it (see, among others, McAfee (1993) and Peters (1997b)). In our cheap talk environment, trading mechanism design by procurers leads to a trivial result. To be precise, suppose that each procurer designs her own trading mechanism. Since their public announcements are cheap talk, they can credibly reveal their trading mechanisms only after contractors select them. But then, the optimal trading mechanism and the consequent market equilibrium are trivial: each procurer will extract the entire surplus, for example, by making a take-it-or-leave-it offer at price c . It then follows that all procurers wish to maximize the expected queue length at the communication stage, and thus there cannot exist any informative equilibrium.

Bidder heterogeneity. The assumption of homogeneous contractors permits a particularly simple characterization of the model and allows us to highlight the main driving forces for our results. However, bidder heterogeneity is arguably one of the reasons for the common use of auctions. In the online appendix, we show that all of our results carry over to the case where contractors differ either in their opportunity cost of working c or in their observable skills in the absence of complementarity between procurers' and contractors' types.³¹ Although the notation becomes significantly more complicated, the main logic for the homogeneous case goes through unchanged. The specification with heterogeneous skills is particularly intriguing for two reasons. First, it captures quality concerns prevalent in real procurement markets. Presumably, it is the primary reason why both EU and US allow procurers to make the final decision based on quality as well as price. Second, it epitomizes one advantage of cheap talk over reserve price posting. The optimal reserve price depends on each contractor's type. Therefore, with reserve

³⁰Even if contractors do not directly observe the presence of competitors, if procurers can credibly disclose the information, then the same unraveling occurs, following the familiar full disclosure logic in the literature on persuasion games. Of course, if the information about competitors is not verifiable (i.e., it is also a cheap talk), then the unraveling does not occur as procurers would always pretend to have many procurers and, therefore, could not credibly reveal information about the true number of procurers.

³¹For the more general case where contractors' types can interact non-additively with procurers' types, as in Shi (2002), Shimer (2005), and Peters (2010), we have less clear intuition. While our exact technique no longer applies, some of our basic insights are likely to carry over.

price posting, it is necessary for each procurer to announce the reserve price schedule for all contractor types. With cheap talk, however, it suffices that each procurer announces her true value, which is likely to be a lot simpler than posting the full reserve price schedule.

Appendix

Proof of Proposition 1. The objective function is strictly concave, while the constraint is linear in $\lambda_{\mathcal{M}}^*$. Therefore, $(\lambda_{\mathcal{M}}^*, \mu)$ is the solution if and only if

$$e^{-\lambda_{\mathcal{M}}^*(v)}(v - c) \leq \mu, = \mu \text{ if } \lambda_{\mathcal{M}}^*(v) > 0, \text{ and } \int_{\mathcal{V}} \lambda_{\mathcal{M}}^*(v) dF(v) = \beta.$$

Combining the two conditions,

$$\int_{\mathcal{V}} \max \left\{ \ln \left(\frac{v - c}{\mu} \right), 0 \right\} dF(v) = \int_{c+\mu}^{\bar{v}} \ln \left(\frac{v - c}{\mu} \right) dF(v) = \beta.$$

The left-hand side is continuous and strictly decreasing in μ . In addition, it approaches ∞ and 0 as μ tends to 0 and $\bar{v} - c$, respectively. Therefore, there exists a unique value of μ that satisfies this equation. The uniqueness of $\lambda_{\mathcal{M}}^*(v)$ follows from its explicit solution. ■

Proof of Lemma 1. If $\lambda_{\mathcal{M}}(v) = \infty$, then the result is straightforward. Suppose $\lambda_{\mathcal{M}}(v) < \infty$. A key observation is that there is no atom in the support of $\lambda(v, \cdot)$. Suppose $\lambda(v, \cdot)$ has a jump at b . First observe that $c < \underline{b}(v) \leq b$. This is due to the fact that a contractor has no competitor with probability $e^{-\lambda_{\mathcal{M}}(v)}$, and thus his expected payoff is bounded below by $e^{-\lambda_{\mathcal{M}}(v)}(v - c)$. Given this, a contractor can obtain a strictly higher expected payoff by quoting a price slightly below b , as his price would decrease slightly, while his winning probability would jump. Formally, for ε sufficiently small but positive,

$$U(v, b - \varepsilon) = e^{-\lambda(v, b - \varepsilon)}(b - \varepsilon - c) > U(v, b) = e^{-\lambda(v, b)} \frac{1 - e^{-d\lambda(v, b)}}{d\lambda(v, b)}(b - c),$$

whenever $d\lambda(m, b) > 0$. This contradicts the fact that b is an optimal price for the contractor.

Given this observation, all the results follow from the following two facts. First, the contractor who quotes the highest price wins the project if and only if there is no competitor (which occurs with probability $e^{-\lambda_{\mathcal{M}}(v)}$), and thus his price must be optimal conditional on him being the only quote provider to the procurer (and, therefore, be equal to v). Second, in equilibrium a contractor must be indifferent over all prices in the support of his bidding strategy, and thus $U(v, b) = U(v, v) = e^{-\lambda_{\mathcal{M}}(v)}(v - c)$ whenever b is in the support of $\lambda(v, \cdot)$. ■

Proof of Proposition 2. Constrained efficiency requires that for any $v, v' \in \mathcal{V}$, $\lambda_{\mathcal{M}}(v) = \lambda_{\mathcal{M}}^*(v) \neq \lambda_{\mathcal{M}}(v') = \lambda_{\mathcal{M}}^*(v')$, unless $\lambda_{\mathcal{M}}^*(v) = \lambda_{\mathcal{M}}^*(v') = 0$. This property can be satisfied only when all procurer types with $\lambda_{\mathcal{M}}^*(v) > 0$ are fully revealed.

The following conditions are necessary for contractor optimality:

$$U(v) \leq u, = u \text{ if } \lambda_{\mathcal{M}}(v) > 0, \text{ and } \int_{\mathcal{V}} \lambda_{\mathcal{M}}(v) dF(v) = \beta,$$

where $u \equiv \max_v U(v)$. If $U(v) = ae^{-\lambda_{\mathcal{M}}^*(v)}(v - c)$ for all $v \in \mathcal{V}$, then the equilibrium conditions clearly coincide with those for constrained efficiency in Proposition 1: it suffices to replace u with $a\mu$. Conversely, suppose there does not exist such a . This means that for some $v, v' \in \mathcal{V}$, $e^{-\lambda_{\mathcal{M}}(v)}(v - c) \neq e^{-\lambda_{\mathcal{M}}(v')}(v' - c)$. It follows that $\lambda_{\mathcal{M}}(v) \neq \lambda_{\mathcal{M}}^*(v)$ or $\lambda_{\mathcal{M}}(v') \neq \lambda_{\mathcal{M}}^*(v')$. ■

References

- Acemoglu, Daron and Robert Shimer**, “Efficient unemployment insurance,” *Journal of Political Economy*, 1999a, 107 (5), 893–928.
- Albrecht, James, Pieter A. Gautier, and Susan Vroman**, “Directed search in the housing market,” *mimeo*, 2014.
- Ángel Hernando-Veciana**, “Competition among auctioneers in large markets,” *Journal of Economic Theory*, 2005, 121 (1), 107–127.
- Ashenfelter, Orley**, “How auctions work for wine and art,” *Journal of Economic Perspectives*, 1989, 3 (3), 23–36.
- Bajari, Patrick and Ali Hortaçsu**, “Economic insights from internet auctions,” *Journal of Economic Literature*, 2004, 42 (2), 457–486.
- Boostani, Reza, Martin Gervais, and Henry E. Siu**, “Optimal unemployment insurance in a directed search model,” *mimeo*, 2014.
- Burdett, Kenneth, Shouyong Shi, and Randall Wright**, “Pricing and matching with frictions,” *Journal of Political Economy*, 2001, 109 (5), 1060–1085.
- Chakraborty, Archishman and Rick Harbaugh**, “Persuasion by cheap talk,” *American Economic Review*, 2010, 100 (5), 2361–2382.
- Crawford, Vincent P. and Joel Sobel**, “Strategic information transmission,” *Econometrica*, 1982, 50 (6), 1431–1451.
- Eeckhout, Jan and Philipp Kircher**, “Sorting versus screening: search frictions and competing mechanisms,” *Journal of Economic Theory*, 2010b, 145 (4), 1354–1385.
- Farrell, Joseph**, “Meaning and credibility in cheap-talk games,” *Games and Economic Behavior*, 1993, 5 (4), 514–531.
- and **Robert Gibbons**, “Cheap talk can matter in bargaining,” *Journal of Economic Theory*, 1989, 48 (1), 221–237.

- Galenianos, Manolis, Philipp Kircher, and Gábor Virág**, “Market power and efficiency in a search model,” *International Economic Review*, 2011, 52 (1), 85–104.
- Gardete, Pedro M.**, “Cheap-talk advertising and misrepresentation in vertically differentiated markets,” *Marketing Science*, 2013, 32 (4), 609–621.
- Julien, Benoit, John Kennes, and Ian King**, “Bidding for labor,” *Review of Economic Dynamics*, 2000, 3 (4), 619–649.
- , – , and – , “Auctions beat posted prices in a small market,” *Journal of Institutional and Theoretical Economics*, 2002, 158 (4), 548–562.
- , – , and – , “Ex post bidding and efficient coordination unemployment,” *Canadian Journal of Economics*, 2005, 38 (1), 174–193.
- , – , and – , “The Mortensen rule and efficient coordination unemployment,” *Economics Letters*, 2006, 90 (2), 149–155.
- Kim, Kyungmin**, “Endogenous market segmentation for lemons,” *RAND Journal of Economics*, 2012, 43 (3), 562–576.
- Koesller, Frédéric and Vasiliki Skreta**, “Sales Talk,” *mimeo*, 2014.
- Kultti, Klaus**, “Equivalence of auctions and posted prices,” *Games and Economic Behavior*, 1999, 27 (1), 106–113.
- Lester, Benjamin, Ludo Visschers, and Ronald Wolthoff**, “Meeting technologies and optimal trading mechanisms in competitive search markets,” *mimeo*, 2014.
- Li, Tong and Isabelle Perrigne**, “Timber sale auctions with random reserve prices,” *The Review of Economics and Statistics*, 2003, 85 (1), 189–200.
- Matthews, Steven A. and Andrew Postlewaite**, “Pre-play communication in two-person sealed-bid double auctions,” *Journal of Economic Theory*, 1989, 48 (1), 238–263.
- McAfee, R. Preston**, “Mechanism design by competing sellers,” *Econometrica*, 1993, 61 (6), 1281–1312.
- and **John McMillan**, “Auctions with a stochastic number of bidders,” *Journal of Economic Theory*, 1987, 43 (1), 1–19.
- Menzio, Guido**, “A theory of partially directed search,” *Journal of Political Economy*, 2007, 115 (5), 748–769.
- Peters, Michael**, “A competitive distribution of auctions,” *Review of Economic Studies*, 1997b, 64 (1), 97–123.
- , “Noncontractible heterogeneity in directed search,” *Econometrica*, 2010, 78 (4), 1173–1200.
- and **Sergei Severinov**, “Competition among sellers who offer auctions instead of prices,” *Journal of Economic Theory*, 1997, 75 (1), 141–179.

- Piga, Gustavo and Matteo Zanza**, “An exploratory analysis of public procurement practices in Europe,” *Challenges in Public Procurement: An International Perspective*, 2004.
- PwC, London Economics, and Ecorys**, *Public Procurement in Europe: Cost and Effectiveness*, report for the European Commission, 2011.
- Quint, Daniel and Kenneth Hendricks**, “Indicative bidding in auctions with costly entry,” *mimeo*, 2014.
- Shi, Shouyong**, “A directed search model of inequality with heterogeneous skills and skill-biased technology,” *Review of Economic Studies*, 2002, 69 (2), 467–491.
- Shimer, Robert**, “Job auctions,” *Princeton University and NBER*, 1999.
- , “The assignment of workers to jobs in an economy with coordination frictions,” *Journal of Political Economy*, 2005, 113 (5), 996–1025.
- Ye, Lixin**, “Indicative bidding and a theory of two-stage auctions,” *Games and Economic Behavior*, 2007, 58 (1), 181–207.