Efficient Competition through Cheap Talk:

Competing Auctions and Competitive Search without Ex Ante Price Commitment*

Kyungmin Kim† and Philipp Kircher‡

November 2013

Abstract

We consider a frictional two-sided matching market in which one side uses public cheap-talk announcements so as to attract the other side. We show that if the first-price auction is adopted as the trading protocol, then cheap talk can be perfectly informative, and the resulting market outcome is efficient, constrained only by search frictions. We also show that the performance of an alternative trading protocol in the cheap-talk environment depends on the level of price dispersion generated by the protocol: If a trading protocol compresses (spreads) the distribution of prices relative to the first-price auction, then an efficient fully revealing equilibrium always (never) exists. Our results identify the settings in which cheap talk can serve as an efficient competitive instrument, in the sense that the central insights from the literature on competing auctions and competitive search continue to hold unaltered even without ex ante price commitment.

JEL Codes: C72, D82, D83.

Keywords: Directed search; competitive search; commitment; cheap talk.

*We are grateful to Yeon-Koo Che, Andrew Clausen, Srihari Govindan, Navin Kartik, John Kennes, George Mailath, Guido Menzio, Dale Mortensen, Roger Myerson, Andrew Postlewaite, Jószef 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 has 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.

"The prediction that sellers will offer auctions is not so implausible. [..] However, [..] sellers do not publish commitments to reserve prices." Peters (1997b, pp. 118-119)

"However, a potential disadvantage of these models - indeed, of any model that assumes posting - is that it is a strong assumption to say that agents commit to the posted terms of trade." Rogerson, Shimer and Wright (2005, p. 976)

1 Introduction

A key ingredient in the functioning of decentralized markets is the way trading partners choose each other. Pre-match communication via public messages (advertising) plays a crucial role in this regard. It directly affects the efficiency of the market if it allows those with higher gains from trade to achieve higher trading probabilities. Committing to better offers is one channel to achieve this, since it directly improves the terms of trade for the other side and, therefore, attracts more counter-parties. For example, in auction settings a seller might want to attract more buyers by offering a lower reserve price. This channel is the main object of study in the large literature on competing auctions and competitive search.[1] As the beginning quotes show, however, even the pioneers of the literature criticize the realism of the commitment assumption, thereby challenging the insights in their own work and in the large subsequent literature that builds on it.[2] Following upon this criticism, we decouple price commitment from pre-match communication and consider a setting in which public announcements are purely *cheap talk*. We examine whether cheap talk can serve as a device to induce efficient competition in decentralized markets, and identify the settings in which the insights from the previous literature remain valid.

We consider a frictional two-sided matching market where agents of one side use cheap-talk announcements so as to attract agents of the other side. The market proceeds in three stages: communication stage, search stage, and trading stage. First, each agent of one side (seller, firm,

[1]Foundational work on competing auctions and competitive search goes back to Peters (1984, 1991, 1997a,b, 2000), McAfee (1993), Peters and Severinov (1997), and Burdett, Shi and Wright (2001). Several recent papers have provided further micro-foundations (e.g., Virág, 2010) as well as extended the foundational work in various dimensions, including multi-unit auctions (Geromichalos, 2012), incentive provision upon matching (Guerrieri, 2008), and coexistence of random and directed search (Lester, 2011). The insights and frameworks from the literature have been applied to study, for example, long-run unemployment outcomes (e.g., Moen, 1997; Mortensen and Wright, 2002; Delacroix and Shi, 2006; Shi, 2009; Fernández-Blanco, 2013), labor market dynamics (e.g., Rudanko, 2009, 2011; Menzio and Shi, 2010, 2011), investment and unemployment insurance models (e.g., Acemoğlu and Shimer, 1999a,b), and wage inequality (e.g., Shi, 2002; Shimer, 2005; Basov, King and Uren, 2011). In monetary economics the methodology has been popularized by Rocheteau and Wright (2005), and in finance by Guerrieri, Shimer and Wright (2010) and Guerrieri and Shimer (2013).

[2]Indeed, empirical papers highlight that commitment is often absent: In many auction markets reserve prices are not publicly revealed (see, e.g., Ashenfelter, 1989; Li and Perrigne, 2003; Bajari and Hortaçsu, 2004). In real estate and used goods markets asking prices seem to lack commitment because they typically differ from actual transaction prices (see, e.g., Horowitz, 1992; Farmer and Stango, 2004; Merlo and Ortalo-Magné, 2004). Finally, job announcements often contain important information, yet explicit wage commitments are rare.

procurer) publicly announces a cheap-talk message to attract a counter-party. Then, each agent of the other side (buyer, worker, contractor) independently decides which one of the senders to approach. Finally, trade takes place according to a pre-determined protocol. Announcements are cheap talk as long as they do not directly affect the subsequent trading protocol. In turn, with cheap-talk communication, unlike models in which the trading rule is set by announcements, it is necessary to separately specify the trading protocol. In this context, questions about the interaction between cheap talk and the trading protocol become central: Does there exist a trading protocol that allows agents to differentiate themselves in the preceding communication stage in a way that achieves efficient market outcomes? What are the necessary properties of such a trading protocol? This paper addresses these questions. It helps to understand in which environments one might expect efficient competition through cheap talk and in which ones not, and which trading protocol a market designer might want to adopt if he had the choice.

Our first main result concerns the existence of an optimal trading protocol. We show that such a protocol exists by focussing on the following trading protocol, which seems particularly well suited and realistic in procurement markets and is given in that context.[3] Once procurers make their announcements (communication), each contractor decides which project to compete for (search) and quotes a price to the procurer of that project without observing other competitors (trading protocol). Then, each procurer decides whether to accept the best offer or not. From an individual procurer's viewpoint, this is equivalent to her running a first-price procurement auction with a stochastic arrival of contractors and a private reserve price. From now on, we refer to this trading protocol simply as the first-price auction. We show that if this trading protocol is adopted at the last trading stage, then there exists a fully revealing equilibrium and the equilibrium outcome is efficient - constrained only by search frictions. In a large market it generates the same expected payoffs and market outcomes as models with full commitment.

The intuitions behind the first result are as follows. For efficiency, recall that if each procurer runs the second-price auction, then each contractor receives exactly as much as his ex post marginal social contribution and, therefore, fully internalizes search externalities. The first-price auction does not yield the same ex post outcome as the second-price auction, but, via revenue equivalence, does provide the same interim expected payoffs, which is sufficient for search efficiency. For existence, notice that a deviating procurer faces the following trade-off: If

[3]First-price auctions are indeed prevalent in procurement markets. European Union Directive 2004/18/EC specifies that a contract has to be given based either on the lowest price or on the economically most advantageous tender. In a study of 540,000 procurement cases, with value of 1.4 trillion Euro over the 2006-2010 period, PwC, London Economics, and Ecorys (2011) find that 82% of cases and 75% of transaction value were contracted through sealed-bid procurement auctions. 73% of contracts and 52% of transaction value were awarded through open sealed-bid auctions (open for any qualified bidder - bids occur prior to knowing who else bids). Similarly, the US Federal Acquisition Regulation lists sealed bids as the first of three possible ways for awarding federal contracts, with winning bid either the lowest "fixed-price" or "fixed-price with economic price adjustment." The latter will be important in our model when quality heterogeneity is present.

she overstates her gains from trade, then she would attract relatively more contractors, but each contractor would quote a relatively higher price. We show that, coupled with procurers' option to reject all the prices, this trade-off is exactly balanced in the case of the first-price auction.

Our second main result concerns a necessary property of other optimal trading protocols in the cheap-talk environment. We show that the performance of an alternative trading protocol ultimately depends on the level of price dispersion generated by the protocol, and our baseline protocol, the first-price auction, provides a tight benchmark for the level of price dispersion. Formally, consider an alternative trading protocol that is revenue-equivalent to the first-price auction. If the distribution of prices resulting from the protocol is a strict mean-preserving contraction (spread) of that from the first-price auction, then an efficient fully revealing equilibrium always (never) exists.[4] This allows us to identify a class of optimal trading protocols as well as evaluate the effects of adopting an alternative trading protocol in the cheap-talk environment, without repeating the characterization of the model. For example, the second-price auction spreads the distribution of prices more than the first-price auction. Therefore, an efficient fully revealing equilibrium never exists with the second-price auction. If a trading protocol yields no price dispersion, then a fully revealing equilibrium always exists.[5]

To understand this second result, recall the trade-off facing a deviating procurer presented above. The literature on competing auctions demonstrated that if each procurer must commit to her announcement of posted reserve price, then the positive effect of an upward deviation, attracting relatively more contractors, is dominated by the negative effect, each contractor quoting a relatively higher price (see, in particular, Peters, 1997b). In our cheap-talk environment, each procurer has an option to reject all the prices, which weakens the negative effect and, therefore, makes the (un)profitability of the deviation ambiguous. Now notice that the option is more valuable, the more frequently do contractors quote prices above the procurer's value. This implies that the more spread the distribution of prices is, the more valuable is the option and, therefore, the more profitable is the deviation. The result follows once this insight is combined with the fact that the procurer's trade-off is exactly balanced in the case of the first-price auction.

One caveat regarding our second result is that it concerns only trading protocols which are revenue-equivalent to the first-price auction. We focus on them because the revenue equivalence is a necessary requirement to obtain efficiency after truthful type revelation, and thus they constitute the only class of trading protocols with hope for efficient outcomes. Yet, other trading protocols could also be considered in our environment. Whereas those mechanisms would not be ideal to study the full potential of cheap-talk communication, they might represent other stylized features of trading protocols common to some markets. Menzio (2007) is particularly relevant in

[4] We introduce a notion of *strict* mean-preserving spreads (contractions) in Section 3.2. Roughly, they mean that a distribution is more or less dispersed than the other *globally*.

[5] We provide a bargaining example in Section 3.2.

this regard.[6] He considers a bilateral alternating offer bargaining protocol in essentially the same environment as ours, and shows that cheap talk can transmit some information and improves on random search, but does not achieve efficiency. Our results and his results complement each other. Our results imply that both imperfect communication and constrained inefficiency in his model are driven by the specific trading protocol, not by the non-binding nature of cheap talk per se. Conversely, his results imply that our second result applies only to trading protocols that would induce efficiency if truthful communication were attainable: his (inefficient) bargaining protocol generates no price dispersion, yet a fully revealing equilibrium does not exist.

One important implication of our results is that ex-ante commitment to the terms of trade is not essential to the insights from previous work on competing auctions. In large markets, constrained efficiency can be achieved whether pre-match communication entails reserve price commitment or not. Furthermore, the expected payoffs of all agents in our cheap-talk environment are exactly identical to those when procurers can commit to reserve prices (e.g., McAfee, 1993; Peters, 1997b; Julien, Kennes and King, 2000). Therefore, one can apply all lessons from commitment models *unaltered* even when ex-ante commitment to the terms of trade is not present - for trading protocols that appear natural in the context of competing auctions. To the extent that the first-price auction (or any trading protocol with a mean-preserving contraction of trading prices) also appears plausible in other buyer-seller or labor market settings, the equivalence in expected market outcomes carries over to such settings. Models of directed search usually consider homogeneous buyers (or workers) and competition through price (wage) commitments (e.g., Peters, 1984, 1991, 1997a, 2000; Montgomery, 1991; Burdett, Shi and Wright, 2001). Our setting again delivers identical expected payoffs to all agents in those settings. This holds even for more complicated wage contracts for workers with heterogeneous productivity (e.g., subclasses of Shi, 2002; Shimer, 2005).

Although we consider an abstract two-sided matching environment, our results have implications for 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, 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), with eight siding with the former relative to six for the latter. Similarly, while most private freelance sites (e.g., MyHammer.co.uk) allow only signals of the "desired budget", which are not binding, some online applications such as Qluso experiment with binding reserve prices. Our result that competition via cheap talk can be as efficient as competition with (re-

[6] Menzio (2007) relates mainly to competitive search rather than to competing auctions, partly because bargaining with two-sided incomplete information raises various conceptual problems, and focuses exclusively on the case of homogeneous workers (contractors). We first consider the same case, for expositional clarity, but later expand towards additional heterogeneity. Another related contribution is a model of directed search with limited commitment by Albrecht, Gautier and Vroman (2013). 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.

serve) prices in large markets might explain their co-existence. In the later part of the paper, we also analyze finite markets and show settings in which cheap talk can lead to a strict equilibrium with full type revelation and efficient market outcomes, with convergence to our limit economy. Interestingly, in finite markets (reserve) price commitment is strictly less efficient due to market power.[7] These explain why cheap talk might retain a real edge in some markets.

In relation to the extensive cheap-talk literature pioneered by Crawford and Sobel (1982), our work focusses particularly on the extent to which cheap-talk communication can facilitate trade and competition in large market settings. In this regard, our work is related to Farrell and Gibbons (1989) and Matthews and Postlewaite (1989), both of which study whether cheap-talk communication can influence trading outcomes in double auctions environments. Our analysis is qualitative different from the ones in those two as well as other existing cheap-talk models, mainly because of the nature of the subsequent competition. Nevertheless, as shown in Section 4.1, 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 model. Section 3 presents and proves our two main results. Section 4 provides further characterizations for the baseline model with the first-price auction: We characterize a class of equilibria in large markets, study the corresponding finite markets where incentives against deviations turn out to be strict, and explain how to accommodate contractor heterogeneity. We conclude in Section 5.

2 The Model

We use the language of procurement markets to lay out our model, as our baseline trading protocol captures procurement markets particularly well, although various other scenarios, including sellers' advertising their goods and firms' posting vacancies, can also be accommodated.

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 the procurers whose value of the project is not greater than v. The set $\mathcal{V}$ has a finite minimal element $\underline{v}$ and a maximal element $\overline{v}$, where $0 < \underline{v} \leq \overline{v}$.

[7]This result is new by itself. Inefficiencies in finite markets with price posting (no bidding) are shown in Galenianos, Kircher and Virág (2011). For reserve price posting (with subsequent bidding), Julien, Kennes and King (2002) find constrained efficiency in the case of two procurers. In our Online Appendix, we prove that the efficiency result does not extend beyond the two-procurer case: whenever there are more than two procurers, commitment always leads to inefficiencies.

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 one project, at an opportunity cost c of forgoing other work. For simplicity we start here with the case of homogeneous contractors with common $c < \underline{v}$.[8]

All agents 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 (normalized) utility 0.

Market interaction. Agents match according to the following sequence. *Communication*: each procurer first publicly announces a cheap-talk message. Assume, without loss of generality, that the set of feasible messages, denoted by $\mathcal{M} \in \mathbb{R}$, coincides with the set $\mathcal{V}$. *Search*: after observing all announced messages, each contractor selects a procurer.[9] *Trading*: each contractor then submits a quote specifying his asking price to the selected procurer, not knowing who else is competing for the job. Subsequently, each procurer decides which contractor, if any, to hire at the quoted price.

The overall market interaction closely resembles those of existing models of competitive search and competing auctions. The main difference is that announcements in existing models are binding obligations that change the nature of the subsequent trading protocol. A posted price leaves no room for price adjustment, while a binding reserve price changes the types of bids that are allowed in the subsequent interaction. In our model, messages do not affect the subsequent trading game, even though they might change players' strategies within that game. Notice 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, the wage range often announced in labor markets, or verbal descriptions of the value of the job.[10]

In principle, any trading protocol can be employed at the trading stage. Messages are cheap talk as long as the trading game does not depend on them. However, the trading protocol affects players' incentives and, therefore, the functioning of the market. We focus on the above trading protocol for three reasons. First, it is the protocol that is widely adopted in real markets,

[8] Contractor heterogeneity is important in procurement markets and arguably the reason why auctions are commonly adopted. We allow for contractor heterogeneity in Section 4.3. Whereas the analysis becomes significantly more complicated, the main logic goes through unchanged.

[9] Contractors are not allowed to contact multiple procurers. Not only is this a standard assumption in the literature, but also does it reflect capacity constraints facing contractors in reality. Small contractors may lack the capability to commit to more than one project or to prepare more than one bid, as acknowledged, e.g., in the survey on EU public procurements by Piga and Zanza (2004). Theoretically, multiple contacts (applications) are considered in Albrecht, Gautier and Vroman (2006), Galenianos and Kircher (2009), and Kircher (2009). While we expect our key insights to carry over, it goes beyond the scope of this paper.

[10] In the model, there is a one-to-one mapping between the support of the distribution of wages facing a given procurer and his type. Therefore, announcing this wage range reveals the type.

in particular, in procurement markets. Second, it delivers our main result regarding the full potential of cheap-talk communication (Theorem 1) in a particularly simple way. Finally, as shown in Theorem 2, it serves as a benchmark that allows us to classify other trading protocols. In other words, it enables us to evaluate the performance of other trading protocols without repeating the analysis for each trading protocol.

Strategies and equilibrium. We impose anonymity on agents' strategies. This is in line with the spirit of the literature on competitive search and competing auctions and allows us to model trading frictions inherent in many markets.[11] Precisely, we restrict attention to the equilibria in which agents of each type play an identical strategy. Procurers' communication strategies are represented by a function $m : \mathcal{V} \to \mathcal{M}$, where $m(v)$ denotes the message sent by type v procurers. All contractors use an identical search and bidding strategy: A search strategy is a function $P : \mathcal{M} \to [0,1]$, where $P(m)$ is the probability that each contractor selects a procurer who announced a message weakly below m. A bidding strategy is a function $H(\cdot;\cdot) : \mathbb{R}_+ \times \mathcal{M} \to [0,1]$, where $H(b;m)$ is the probability that a contractor who selected a procurer with message m quotes a price weakly below b. Each procurer's optimal hiring strategy is straightforward, so its formal description is omitted: each procurer hires a contractor who offers the lowest price, provided that the price is not greater than her own value of the project.

In the current large market setting, it is necessary to specify the matching function. We adopt the canonical "urn-ball" matching technology: Given procurers' communication strategies $m(\cdot)$ and contractors' search strategies $P(\cdot)$, the ratio of contractors to procurers at message m, usually called the queue length, is given by

$$\lambda(m) = \frac{\beta dP(m)}{\int_{\{v\in\mathcal{V}:m(v)=m\}} dF(v)}. \tag{1}$$

Given a message, there is no further information transmission, and contractors approach procurers randomly. In a large market, this generates the familiar urn-ball matching process: the probability that exactly $n \in \mathcal{N}_0$ contractors select a procurer is given by $\lambda(m)^n e^{-\lambda(m)}/n!$. Conditional on the event that a contractor selects a procurer, the probability that n other contractors select the same procurer is also given by $\lambda(m)^n e^{-\lambda(m)}/n!$. While, as in the literature, these matching probabilities are directly imposed on the model, they can be formally derived by considering finite markets and letting the size of the market go to infinity. See Burdett, Shi and Wright (2001) for an excellent discussion, and Section 4.2 for a formal derivation within our framework.

[11] See Burdett, Shi and Wright (2001) for an excellent discussion on the issue. In procurement markets, collusion and anti-trust concerns are important, but they are beyond the scope of this paper. The assumption here is that antitrust is imposed and agents must decide in isolation, which is intended to be captured through the anonymity assumption.

An equilibrium is then a set of strategies $(m(\cdot), P(\cdot), H(\cdot;\cdot))$, together with contractors' beliefs about the distribution of types associated with each message, such that given $(m(\cdot), P(\cdot), H(\cdot;\cdot))$ and the resulting $\lambda(\cdot)$ according to (1), it is optimal for procurers to communicate according to $m(\cdot)$ and for contractors to search and bid according to $m(\cdot)$ and $H(\cdot;\cdot)$, respectively, and contractors' beliefs are derived by Bayes' rule whenever possible.[12] For most parts of the paper, we focus on fully revealing equilibria, in which contractors' beliefs about the distribution of types for each message are trivial and, therefore, can be suppressed. A fully revealing equilibrium is an equilibrium in which $m(\cdot)$ is injective. Without loss of generality, we restrict attention to the fully revealing equilibrium in which each procurer truthfully announces her own value, that is, $m(v) = v$ for all $v \in \mathcal{V}$.

3 Main Results

In this section, we present our two main results and prove them.

3.1 Fully Revealing Equilibrium in the Baseline Model

Our first result concerns the existence and efficiency of the fully revealing equilibrium in the baseline model where the first-price auction serves as the trading protocol. We compare the equilibrium outcome to the optimal solution of a planner who wishes to maximize gross surplus, knowing procurers' types (or dictating truthful revelation $m(v) = v$) but being constrained to assign a common search strategy $P(\cdot)$ to all contractors. In other words, we consider the planner who faces the same anonymity restrictions as in the decentralized equilibrium and can achieve trading probabilities only as explained in connection with (1). We say that an equilibrium is constrained efficient if it yields the same outcome as the planner's solution.

Theorem 1 *In the baseline model, the fully revealing equilibrium always exists and is constrained efficient.*

The theorem consists of two parts, the existence of the fully revealing equilibrium and the constrained efficiency of the equilibrium. We first characterize contractors' equilibrium bidding strategies and use the characterization to prove each claim of the theorem.

3.1.1 Preliminary: Contractors' Equilibrium Bidding Strategies

Suppose contractors believe that a procurer's true value of the project is equal to v and the associated queue length for the procurer is given by λ. Denote by $\tilde{H}(b; v, \lambda)$ the probability that

[12] As is typically the case in cheap-talk models, additional restrictions on beliefs and queue lengths outside of the support of equilibrium strategies are not necessary at this point.

a contractor quotes a price weakly below b to the procurer. In addition, denote by $U(v,\lambda)$ the expected payoff of a contractor who selects the procurer and by $V(v',v,\lambda)$ the expected payoff of a procurer whose true value is v' but is believed to be v by contractors. The following lemma provides closed-form solutions for all three functions for the fully-revealing equilibrium.

Lemma 1 *(1) Contractor's expected payoff:* $U(v,\lambda)=e^{-\lambda}(v-c)$.
(2) Equilibrium bidding strategies: For each $b\in[c+U(v,\lambda),v]$, $\tilde{H}(b;v,\lambda)$ *is given by the value that satisfies*

$$e^{-\lambda\tilde{H}(b;v,\lambda)}(b-c)=U(v,\lambda). \tag{2}$$

(3) Procurers' expected payoffs: For each $v'\in\mathcal{V}$,

$$V(v',v,\lambda)=\int_{c+U(v,\lambda)}^{\min\{v',v\}}(v'-b)d\left(1-e^{-\lambda\tilde{H}(b;v,\lambda)}\right).$$

In particular, $V(v,v,\lambda)=\left(1-e^{-\lambda}-\lambda e^{-\lambda}\right)(v-c)$.

Proof. A key observation is that $\tilde{H}(\cdot;v,\lambda)$ has no atom in its support: Suppose b is in the support. A contractor can secure payoff $e^{-\lambda}(v-c)>0$ by quoting v, because he faces no competition (no other contractor selects the same procurer) with probability $e^{-\lambda}$. This implies that $b-c\geq e^{-\lambda}(v-c)>0$, that is, b must be bounded away from c. If there is an atom at b, then a contractor can obtain a strictly higher expected payoff by quoting a price slightly below b, because his price would decrease slightly, while his winning probability would jump. This contradicts the fact that b is an optimal price for the contractor.

Given this observation, (1) follows from the fact that the contractor who quotes the highest price wins the project if and only if there is no competitor (which occurs with probability $e^{-\lambda}$), and thus his price must be optimal conditional on him being the only quote provider to the procurer (and, therefore, be equal to v). (2) follows from the fact that in equilibrium a contractor must be indifferent over all prices in the support of his mixed bidding strategy and a contractor who bids b wins the project with probability $e^{-\lambda\tilde{H}(b;v,\lambda)}$. Finally, (3) follows from the fact that a procurer accepts the lowest price if and only if it is below her true value of the project v' (thus, $\min\{v',v\}$) and the probability that the lowest quote to the procurer is not greater than b is equal to $1-e^{-\lambda\tilde{H}(b;v,\lambda)}$. ■

One can intuitively understand the payoffs through the revenue equivalence between first- and second-price auctions.[13] Suppose a procurer with project value v runs a second-price procurement auction with reserve price equal to v. Then, a contractor would get the entire surplus if and only if there is no other competitor, while the procurer would get the entire surplus if and only if there

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

are at least two contractors. Therefore, a contractor's expected payoff is equal to $e^{-\lambda}(v-c)$, while the procurer's expected payoff is equal to $(1-e^{-\lambda}-\lambda e^{-\lambda})(v-c)$.

3.1.2 Constrained Efficiency of the Fully Revealing Equilibrium

We now solve for contractors' equilibrium search strategies in the fully revealing equilibrium. The result is then used to establish the constrained efficiency of the fully revealing equilibrium.

Suppose contractors believe that all procurers announced their true values. Given other contractors' search strategies $P(\cdot)$, or equivalently, via Equation (1), queue lengths $\lambda(\cdot)$, an optimal search strategy for an individual contractor is to select a procurer from whom he can obtain the highest expected payoff. Formally, let $u \equiv \max_{v\in\mathcal{V}} U(v,\lambda(v)) = e^{-\lambda(v)}(v-c)$. Following the literature, we refer to this utility u as *market utility*. Then, a contractor's optimal strategy is to select a procurer with message (value) v only if $e^{-\lambda(v)}(v-c) \geq u$. In other words, the queue length $\lambda(v)$ for a procurer can be positive only if $e^{-\lambda(v)}(v-c) = u$.

It follows that the equilibrium queue lengths $\lambda(\cdot)$ (equivalently, contractors' equilibrium search strategies) are characterized by the following two conditions:

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

The latter condition states that the total number of quotes made in this market (left-hand side) must coincide with the total number of contractors β: $\lambda(v)$ is the average number of quotes a type v procurer receives, while $dF(v)$ is the measure of type v procurers. Combining the two conditions,

$$\int_{\mathcal{V}} \lambda(v)dF(v) = \int_{c+u}^{\overline{v}} log\left(\frac{v-c}{u}\right) dF(v) = \beta. \tag{3}$$

Since the left-hand side is strictly decreasing in u, there exists a unique value of u that satisfies this equation. Given the uniqueness of u, $\lambda(v)$ is also uniquely determined for all $v \in \mathcal{V}$. The equilibrium is then fully determined by $m(v) = v$ (and contractors' associated beliefs), $P(\cdot)$ obtained from $\lambda(\cdot)$ via (1), and $H(b;v) = \tilde{H}(b;v,\lambda(v))$ as in Lemma 1.

To evaluate the extent of efficiency of this fully revealing equilibrium, consider the social planner's problem, described at the beginning of this section. His objective is to design a common search strategy $P(\cdot)$ for contractors so as to maximize gross surplus. Again, via (1), this is equivalent to finding the optimal queue lengths $\lambda(\cdot)$. Given the queue length $\lambda(v)$, the probability that a type v procurer is selected by at least one contractor and completes her project is equal to $1-e^{-\lambda(v)}$. Therefore, the problem facing the social planner can be expressed as follows:

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

Letting μ be the Lagrangian multiplier for the constraint, $(\lambda(v), \mu)$ is the solution if and only if[14]

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

Notice that the conditions are identical to those for contractors' equilibrium search strategies. By the same argument as above, there is a unique pair $(\mu,\ \lambda(\cdot))$ that satisfy the conditions. Therefore, the fully revealing equilibrium outcome coincides with the social planner's optimal solution.

To see the underlying reason behind this efficiency result, recall that in a two-sided market environment as ours, contractors behave socially efficiently if they receive exactly their marginal contribution to the value of the local interaction among the single procurer and the contractors that selected her. In our model, this means that contractors' equilibrium search strategies would be efficient if each contractor receives the entire surplus when he is the only contractor in a meeting, while receives 0 payoff when there are other contractors. This arises precisely when payoffs are determined at the contractor-optimal point in the (ex-post) core of the local interaction, which can be implemented with the second-price auction with reserve price equal to the procurer's true value.[15] 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.

3.1.3 Existence of the Fully Revealing Equilibrium

We now consider procurers' incentives at the communication stage and prove that the fully revealing equilibrium always exists.

Suppose a type v' procurer unilaterally deviates and announces v. The procurer's basic trade-off is between queue length (the number of contractors she can attract) and price (how much each contractor quotes): If $v > v'$, then relatively more contractors would select the procurer, while each of them would quote a relatively higher price. The opposite happens for $v < v'$. The crux of our proof is to show that this trade-off is exactly balanced when the price is determined according to the first-price auction.

Consider the probability that the procurer hires a contractor at a price below $b \leq \min\{v', v\}$. Since the queue length for a firm with message m is $\lambda(m)$ and each contractor offers a price below b with probability $H(b; m) = \tilde{H}(b; m, \lambda(m))$, the probability is equal to $1 - e^{-\lambda(v')H(b;v')}$ if the procurer announces v' and $1 - e^{-\lambda(v)H(b;v)}$ if the procurer announces v. These two probabilities

[14]The objective function is strictly concave, while the constraint is linear. Therefore, the first-order condition is both necessary and sufficient.

[15]Gautier and Holzner (2013) recently generalize this insight into the case of multiple applications. With multiple applications, one side's optimal points in the ex-post core depend on the network structure associated with each meeting and, therefore, are much harder to determine. In addition, a more general allocation mechanism than the second-price auction is necessary to implement the efficient outcome.

are identical even though the queue lengths and bidding strategies differ: By the equilibrium conditions for contractors' search strategies, as long as the queue lengths are strictly positive, a contractor must be indifferent between selecting a type v' procurer and a type v procurer. In addition, conditional on selecting a particular procurer, a contractor must be indifferent over all prices in the support of their bidding strategies. Therefore, for any price below $\min\{v', v\}$, a contractor must have the same winning probability whether he selects a type v' procurer or a type v procurer. A contractor who quotes b wins if and only if no other contractor quotes less than b to the same procurer. Therefore, his winning probability is equal to $e^{-\lambda(m)H(b;m)}$. It follows that in equilibrium the products of the queue length and the probability of bidding a price below b are identical, i.e., $\lambda(v')H(b; v') = \lambda(v)H(b; v)$.[16]

The above result implies that a procurer's deviation affects her expected payoff only through the changes in the probabilities that she hires a contractor at a price between v' and v. If $v < v'$, then the deviating procurer loses the opportunity to hire a contractor at a price between v and v'. Therefore, the downward deviation is strictly unprofitable. If $v > v'$, then the deviating procurer now receives price quotes between v' and v with a positive probability. However, the procurer has no incentive to accept such a price, and thus those price quotes are irrelevant to the procurer's expected payoff. Consequently, the upward deviation would not change the procurer's expected payoff.[17] To sum up, each procurer strictly prefers truth-telling to downward deviations and weakly to upward deviations.

3.2 Alternative Trading Protocols

Our second result concerns the effects of adopting an alternative trading protocol. We consider the same market game as described in Section 2, except that now we replace the first-price auction with some other trading protocol. We examine what trading protocols can induce efficient market outcomes as in our baseline model.

For efficiency, two necessary properties of the market interaction are immediate. The equilibrium must induce full revelation of types: otherwise different types would obtain the same

[16]This can be derived directly from (2) in Lemma 1 and the fact that contractors' expected payoffs at both messages (the right-hand side in (2)) must be equalized in order for contractors to select both messages.

[17]Formally, combining Lemma 1 with the fact that $\lambda(m)H(b; m)$ is independent of m,

$$V(v', v, \lambda(v)) = \int_{c+u}^{\min\{v',v\}} (v' - b) d\left(1 - e^{-\lambda(v)H(b;v)}\right) = \int_{c+u}^{\min\{v',v\}} (v' - b) d\left(1 - e^{-\lambda(v')H(b;v')}\right).$$

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

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

Obviously, this expression is strictly positive if $v < v'$, while equal to 0 if $v > v'$.

queue length, which cannot be efficient. In addition, contractors' expected payoffs in the fully revealing equilibrium must be constantly proportional to those in our baseline model. They must be exactly identical if one would allow endogenous entry.[18] Note that this property concerns only agents' *interim* expected payoffs, not their *ex post* expected payoffs. The latter may depend on other variables, such as the realized number of contractors for each procurer, and need not coincide with those from the first-price auction. The only binding restriction is that contractors' expected payoffs conditional on v and the planner's queue length $\lambda(v)$ are identical to those from the first-price auction. From procurers' perspective, this means that the equilibrium outcome must be revenue-equivalent to that of the first-price auction.

Given these necessary restrictions, the central question is under what conditions a fully revealing equilibrium exists. In a fully revealing equilibrium, contractors' beliefs about each procurer's type are degenerate: if a procurer sends message v, then contractors assign probability 1 to the event that his type is v. Revenue equivalence implies that the same queue length as in our baseline model, $\lambda(v)$, is assigned to each type v procurer. Each trading protocol then generates a unique mapping from v to a distribution of prices. Denote by $L(b;v)$ the probability that the best price quote to a type v procurer is less than or equal to b. For example, in our baseline model $L(b;v) = 1 - e^{-\lambda \tilde{H}(b;v,\lambda(v))}$. Since the procurer obviously cannot hire if she is not selected by any contractor, $\lim_{b\to\infty} L(b;v) \leq 1 - e^{-\lambda}$.

For procurers' truth-telling incentives, consider a type v' procurer's deviation to v (by sending message v). While in principle the trading prices for the deviant type could vary in many ways from those of the true type, many insights can be gained under the following restriction on the trading protocol: assume that deviant type v' faces the same distribution of prices, $L(\cdot;v)$, as a (truth-telling) type v procurer. This means that the distribution of prices for a procurer depends only on contractors' (degenerate) beliefs about the procurer's type, not on the procurer's true type. The procurer still retains the option to reject a price, but cannot influence the distribution of prices. This restriction clearly holds in our baseline model. It also seems reasonable in many other settings, because we only impose it when contractors have degenerate beliefs: even though a deviant might want to convince the contractors that she has a different type, that is not possible under degenerate beliefs. Below we provide some trading protocols that satisfy this assumption.

Obviously, a type v' procurer would accept only a price below v'. This implies that the problem reduces to whether a type v' procurer would prefer $L(\cdot;v')$ to $L(\cdot;v)$ over the interval $(-\infty, v')$, that is, $\int_{-\infty}^{v'}(v'-b)dL(b;v') \geq \int_{-\infty}^{v'}(v'-b)dL(b;v)$. The following theorem establishes that the existence of the fully revealing equilibrium boils down to the level of price dispersion resulting from each trading protocol (at full type revelation and optimal queue lengths), and the first-price auction in the baseline model provides a tight benchmark for the level of price

[18] Entry can be on either side: procurers might enter at some fixed cost, or contractors could have a search cost. In either case it is easy to show that efficient entry arises if and only if the expected payoffs are as in the baseline model.

dispersion. Since efficiency requires revenue equivalence, we measure price dispersion in terms of mean-preserving spreads and say that a mean-preserving spread is strict if the distribution puts strictly more weight above (below) any particular cutoff to the right (left) of the mean.[19]

Theorem 2 *Consider an alternative trading protocol. If the distribution of prices for each procurer resulting from this protocol is a mean-preserving contraction (strict mean-preserving spread) of that from the first-price auction, then the fully revealing equilibrium always (never) exists. In the case of a strict mean-preserving contraction, all procurers' truth-telling incentives are strict in the fully revealing equilibrium.*

Proof. Denote by $W(v', v)$ the expected payoff of a type v' procurer when contractors believe that his value is v. In addition, for notational simplicity, let $L^*(b; v) \equiv 1 - e^{-\lambda \tilde{H}(b;v,\lambda(v))}$ denote the distribution of prices from our baseline setting.

Suppose for each $v \in \mathcal{V}$, $L(\cdot; v)$ is a strict mean-preserving spread of $L^*(\cdot; v)$. If $v > v'$, then

$$W(v', v) = \int_{-\infty}^{v'} (v' - b) dL(b; v) > \int_{-\infty}^{v'} (v' - b) dL^*(b; v) = V(v, v', \lambda(v)) = V(v', v', \lambda(v')) = W(v', v').$$

The first inequality is due to the property of mean-preserving spreads: Since $L(\cdot; v)$ is a strict mean-preserving spread of $L^*(\cdot; v)$,

$$\int_{-\infty}^{v'} (v' - b)(dL(b; v) - dL^*(b; v)) = \int_{-\infty}^{v'} (L(b; v) - L^*(b; v)) \, db > 0.$$

Notice that a strict mean-preserving spread is necessary to ensure the strict inequality. If the inequality holds weakly, then the result depends on the locations of the points at which the inequality holds strictly. The second-last equality uses the procurer's indifference over upward deviations in the baseline model, while the last equality utilizes the revenue equivalence.

The same argument can be used to establish that if $L(\cdot; v)$ is a mean-preserving contraction of $L^*(\cdot; v)$, then $W(v', v) \leq W(v', v')$, with strict inequality in the case of strict mean-preserving contractions. It remains to show that downward deviations are also unprofitable in the case of mean-preserving contractions. Mean-preserving contractions imply that if $v < v'$, then a procurer never faces a price above $v (< v')$ even with the alternative trading protocol. It then follows that

$$W(v', v) = V(v', v, \lambda(v)) < V(v', v', \lambda(v')) = W(v', v',).$$

[19]Formally, a distribution function F_A is a mean-preserving spread of another distribution function F_B if they have the same mean and $\int_{-\infty}^{x} (F_A(t) - F_B(t)) dt \geq 0$ for all x, with strict inequality for some x. We say that F_A is a strict mean-preserving spread of F_B if the latter property holds with strict inequality for any x in the union of the supports of F_A and F_B, except at the minimum of the union. We also define a strict mean-preserving contraction in an analogous manner. The strict notion requires that F_A is more dispersed than F_B through the support. In other words, it excludes the case where F_A differs from F_B only locally (in particular, only around the mean).

■

To intuitively understand the result, 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.[20] On the other hand, each contractor quotes a relatively higher price. If a procurer must hire a contractor after the deviation, then, as shown in the literature on competing auctions and competitive search, the latter negative effect outweighs the former positive effect, and thus the procurer would strictly prefer truth-telling. If communication is cheap talk, then a procurer can refuse to hire any contractor. This reduces the latter negative effect, which makes the overall effect ambiguous. Since a procurer rejects only a price above his own value, the decrease in the latter effect is larger, the more frequently does the deviating procurer face prices above his value. This implies that the more dispersed the distribution of prices is, the more likely is it that an upward deviation is profitable. As shown in Section 3.1.3, if the first-price auction is employed as the trading protocol, then the two effects become exactly balanced. It follows that the profitability of an upward deviation with an alternative trading protocol (and, therefore, the existence of the fully revealing equilibrium) ultimately depends on whether the distribution of prices induced by the protocol is more or less dispersed than the one by the first-price auction.

As a concrete example, consider the case where procurers still run first-price auctions, but contractors observe the number of other contractors for the same procurer before they place bids. In this case, in the fully revealing equilibrium a contractor offers v to a type v procurer if he is the only quote provider, while the price is driven down to c if there are at least two contractors. Notice that these are the same ex post outcomes as when procurers run second-price auctions with reserve price equal to their values. For simplicity we call this outcome "the second-price auction." As is well-known in auction theory (see, e.g., Proposition 2.4 in Krishna, 2009), the second-price auction spreads prices more than the first-price auction. Theorem 2 then implies that the fully revealing equilibrium does not exist with this trading protocol. Intuitively, when a type v' procurer deviates to $v(> v')$, she clearly benefits from attracting relatively more contractors, as it would increase her chance of hiring a contractor at c. However, she does not face any real deviation cost. When there is only one contractor, he would offer v, instead of v'. The increase in the price, however, does not affect the procurer's expected payoff, because she would simply reject the price.[21] In fact, the same unraveling applies to any informative equilibrium. It can be shown that there does not exist any informative equilibrium with the

[20] A procurer never faces a price above his value after a downward deviation. Therefore, the problem described in this paragraph is irrelevant to downward deviations. For the existence of the fully revealing equilibrium, as shown in the proof of Theorem 2, it suffices to check the (un)profitability of upward deviations.

[21] Formally, when the queue length for a procurer is λ, the probability that at least two contractors select the procurer is given by $1 - e^{-\lambda} - \lambda e^{-\lambda}$. Since $\lambda(v) > \lambda(v')$,

$$L(c; v) = 1 - e^{-\lambda(v)} - \lambda(v)e^{-\lambda(v)} > L(c; v') = 1 - e^{-\lambda(v')} - \lambda(v')e^{-\lambda(v')}.$$

second-price auction.

As for an opposite example yielding a contraction of prices, consider a bilateral random-proposer bargaining protocol, as in Grossman and Perry (1986) and Gul and Sonnenschein (1988), and recently adopted within a similar search context by Menzio (2007). Precisely, consider the following trading protocol. Once contractors choose procurers, one contractor is randomly selected among those who chose the same procurer, and the other contractors are discarded. Then, the contractor and the procurer engage in an infinite-horizon bilateral bargaining game in which the procurer (contractor) is selected as a proposer with probability θ (probability $1-\theta$) in each period and they apply common discount factor δ. Under the refinements in beliefs proposed in these aforementioned papers, in the limit as δ tends to 1, the contractor obtains exactly a share $1-\theta$ of the value of the lowest procurer type in the support of his beliefs. In the fully revealing equilibrium, his beliefs are degenerate, and thus he obtains $(1-\theta)(v-c)$ from a type v procurer. Note that a deviant procurer who lies about his type can do no more than either also give this amount to the contractor or reach no agreement at all (i.e., reject the price). The reason is that his type is outside of the support of beliefs of the contractor, and so he cannot convince the contractor to lower his demand. While this fits the framework we are analyzing, it is clear that this bargaining protocol is generically not revenue-equivalent to the first-price auction and therefore does not induce efficient outcomes.

For the revenue equivalence, a type v procurer must obtain a share θ such that his expected total payoff $(1-e^{-\lambda})\theta(v-c)$ equals the expected payoff from the first-price auction, $(1-e^{-\lambda}-\theta e^{-\lambda})(v-c)$. Then, it is necessary that his share θ depends on the market condition reflected in the queue length. In particular, he would need to obtain $\theta(\lambda)=(1-e^{-\lambda}-\lambda e^{-\lambda})/(1-e^{-\lambda})$, where $\theta(\lambda)$ coincides with the probability that at least one more contractor selects the procurer conditional on the event that she is selected by one contractor. This expression is strictly increasing in λ, which captures the intuitive idea that a procurer's bargaining power increases as more contractors wish to trade with her. The previous trading protocol does not achieve such outcomes as other contractors have already been discarded and, therefore, their presence is meaningless at the time of bargaining. When other contractors matter, a trading share $\theta(\lambda)$ can arise for many reasons, and one can construct simple bargaining games that yield exactly such an outcome.[22] With this modification, the trading protocol becomes revenue-equivalent to

Therefore,

$$W(v',v)=L(c;v)(v'-c)>L(c;v')(v'-c)+(L(v';v')-L(c;v'))(v'-v')=W(v',v').$$

[22]Various bargaining games might lead to the outcome. One very stylized but easy-to-explain example is the following: Each procurer initially bargains with one of her contractors. Initially neither of them knows whether there is any other contractor who also selected the procurer. The procurer starts by making an offer to the contractor. If the contractor accepts, he gets the proposed transfer and executes the work. If he rejects and there is no other contractor, he makes a take-it-or-leave it offer to the procurer in return. If there are other

the first-price auction. Since the price distribution is contracted to a single point, by Theorem 2, the efficient fully revealing equilibrium exists and provides strict incentives for all procurers to truthfully reveal their types.

4 Discussion

In this section, we provide further characterizations for our baseline model.

4.1 Other Equilibria and Equilibrium Selection

As in other communication games, our game also has other equilibria than the fully revealing equilibrium. To begin with, there always exists a babbling equilibrium: If all procurers babble (uniformly randomize over all messages), then contractors would ignore all messages and randomly select a procurer. This, in turn, makes all procurers indifferent over all messages. The following proposition characterizes the set of all partitional equilibria in which the set of procurer types that play an identical communication strategy is convex in $\mathcal{Y}$.[23]

Proposition 1 *Suppose that $(1 - F(v))(v - c)$ is strictly quasi-concave. Then, for any $\tilde{v} \in \mathcal{V}$, there is an equilibrium in which all procurers below $\tilde{v}$ fully reveal their types, while the other procurers pool all together. There does not exist any other partitional equilibrium that yields a different outcome from any of these.*

Proof. See the appendix. ■

This is a consequence of the fact that each procurer has a strict disincentive to deviate downward, but is indifferent over all upward deviations. A procurer below $\tilde{v}$ has no incentive to deviate for the same reason as in the fully revealing equilibrium. A procurer above $\tilde{v}$ can deviate only to below $\tilde{y}$, but the deviation is not profitable, again, as in the fully revealing equilibrium. The interval above $\tilde{v}$ cannot be segmented, because if so the highest procurer in the lower segment would deviate.

---

contractors available for the procurer, then one of them is randomly selected and makes a take-it-or-leave-it offer to the procurer. The initial contractor can always reject any offer and make a counter-offer when he is the only contractor. Since the (conditional) probability that there is no other contractor is given by $\lambda e^{-\lambda}/(1-e^{-\lambda})$, he can ensure at least $\lambda e^{-\lambda}/(1-e^{-\lambda})(v-c)$ from this game. It follows that the outcome of this bargaining game is that the procurer offers payoff $\lambda e^{-\lambda}/(1-e^{-\lambda})(v-c)$ (i.e., price $c+\lambda e^{-\lambda}/(1-e^{-\lambda})(v-c)$) to the initial contractor, who then accepts it. Note that each contractor's expected payoff is equal to the probability that he is selected as the initial contractor ($(1-e^{-\lambda})/\lambda$) times his expected payoff when selected as the initial contractor, which coincides with $e^{-\lambda}(v-c)$. In addition, any deviant procurer cannot reduce the price below $c + \lambda e^{-\lambda}/(1-e^{-\lambda})(v-c)$, as contractors have degenerate beliefs and expect to get at least this amount. This is clearly an example of mean-preserving contractions in Theorem 2.

[23] There may exist non-partitional equilibria, even though we are not aware of any. We are also not aware of a way to systematically characterize such equilibria.

Nevertheless, the fully revealing equilibrium is more appealing than the others at least for two reasons. First, it yields the socially efficient outcome and, therefore, would be a natural focal point. Second, it is the only equilibrium that satisfies NITS (no incentive to separate) suggested by Chen, Kartik and Sobel (2008). We show that if the condition is applied to the highest procurer type (that is, if we require the highest procurer type not to have an incentive to deviate),[24] then the fully revealing equilibrium uniquely satisfies the condition. We further show that the fully revealing equilibrium can be selected through one of the main motivations for NITS, without relying on the choice of a type to apply the condition.[25]

Proposition 2 *Among all the partitional equilibria, only the fully revealing equilibrium satisfies NITS that is applied to the highest procurer type. If procurers face costs of lying, then the fully revealing equilibrium is the unique equilibrium that satisfies a natural monotonicity condition (the higher a procurer's type is, the higher message he sends) or the intuitive criterion.*

Proof. See the appendix. ■

4.2 Finite Markets

In the literature, there is a long tradition to build up the large economy with a continuum of agents as the limit of games with finite numbers of agents. This is even more relevant in our baseline setting as procurers' incentives for truth-telling are weak. We show that procurers' incentives are strict in finite markets, and thus the fully revealing equilibrium in the baseline model has a desirable robustness property. Moreover, most other insights from the large market setting arise essentially unchanged in the finite market setting, albeit at the cost of a more restrictive type space and of additional notational complexity.

4.2.1 Setup

Now suppose there are N procurers and M contractors, where M is the smallest integer such that $M \geq \beta N$. In order to obtain sharp results, assume that the set of possible procurer types $\mathcal{V}$ is also finite: $\mathcal{V} \equiv \{v^1, ..., v^I\}$.[26] Without loss of generality, assume that $v^i < v^{i+1}$ for each $i \leq I-1$. For notational simplicity, denote by f^i the probability that each procurer draws project

[24] NITS concerns the incentive of the *lowest* sender type to separate from other types in the standard cheap-talk game. In our model where agents' payoffs through communication are endogenously determined, it is not clear which type is the appropriate *lowest* type.

[25] A procurer's indifference over upward deviations prevents us from applying another prominent equilibrium selection criterion, "neologism-proofness" by Farrell (1993). The upward incentive is strict with an alternative trading protocol that generates less price dispersion than the first-price auction (see Theorem 2) or in finite markets (see Section 4.2). In those cases, it is straightforward to show that the fully revealing equilibrium is also the unique equilibrium that satisfies neologism proofness.

[26] This restriction on the type space allows us to obtain clear-cut results on the fully revealing equilibrium even in finite markets. For the reason clarified shortly, if the set $\mathcal{V}$ is continuous, we would have to look for partitional

value v^i. We also denote by v_n procurer n's value of the project, by $\vec{v}$ a profile of all procurers' values, and by $\vec{v}_{-n}$ a profile of all procurers' values save procurer n's. All other specifications of the model are identical to the ones in Section 2.

As before, we impose anonymity on agents' strategies and focus on the equilibria in which agents of the same type play an identical strategy. Procurers' communication strategies are represented by a function $m : \mathcal{V} \to \mathcal{V}$ where $m(v)$ denotes the message announced by a procurer with project value v. Contractors' search strategies are a function $q : \mathcal{V}^N \times \{1, ..., N\} \to [0,1]$ where $q(\vec{v}, n)$ denotes the probability that each contractor selects procurer n when the announced value profile is $\vec{v}$. Lastly, contractors' bidding strategies are a function $H : \mathbb{R}_+ \times \mathcal{V}^N \times \{1, ..., N\} \to [0, 1]$ where $H(b; \vec{v}, n)$ denotes the probability that a contractor who selected procurer n quotes a price less than or equal to b. For notational simplicity, we simply denote $q(\vec{v}, n)$ by $q_n(\vec{v})$ and $H(b; \vec{v}, n)$ by $H_n(b; \vec{v})$, respectively. As in the large market setting, we study the fully revealing (perfect Bayesian) equilibrium in which $m(v) = v$ for all $v \in \mathcal{V}$.

The following lemma is a finite-market counterpart to Lemma 1. We denote by $U_n(\vec{v})$ the expected payoff of a contractor who selects procurer n when the announced value profile is $\vec{v}$. The proof is essentially identical to that of Lemma 1 and, therefore, omitted.

Lemma 2 *(1) Contractors' expected payoff:* $U_n(\vec{v}) = (1 - q_n(\vec{v}))^{M-1}(v_n - c)$.
(2) Equilibrium bidding strategies: For each $b \in [c + U_n(\vec{v}), v_n]$, $H_n(b; \vec{v})$ *is given by the value that satisfies*

$$(1 - q_n(\vec{v})H_n(b; \vec{v}))^{M-1}(b - c) = U_n(\vec{v}).$$

4.2.2 Efficiency of the Fully Revealing Equilibrium

We now characterize contractors' equilibrium search strategies for a given value profile $\vec{v} \in \mathcal{V}^N$ and prove that the fully revealing equilibrium is constrained efficient even in finite markets.

Following the same reasoning as in Section 3.1.2, the equilibrium conditions for contractors' search strategies $q(\cdot, \cdot)$ are given by

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

where $u(\vec{v}) \equiv \max_n U_n(\vec{v})$. Rewriting the condition yields

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

equilibria as in, for example, Crawford and Sobel (1982). It would still be possible to construct a sequence of partitional equilibria that converges to the fully revealing equilibrium, in the sense that the size of each partition element shrinks to zero, but the 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.

Summing over n and using the fact that $\sum_n q_n(\vec{v}) = 1$ result in

$$N - 1 = \sum_n \min\left\{ \left(\frac{u(\vec{v})}{v_n - c} \right)^{\frac{1}{M-1}}, 1 \right\}. \tag{4}$$

From this equation, it follows that for each $\vec{v} \in \mathcal{V}^N$, contractors' market utility is uniquely determined. The uniqueness of $q_n(\vec{v})$ follows from the uniqueness of $u(\vec{v})$.

The finite-market analogue to the constrained social planner's problem is as follows:

$$\max_{q(\cdot,\cdot)} \sum_n (1 - (1 - q_n(\vec{v}))^M)(v_n - c) \text{ subject to } \sum_n q_n(\vec{v}) \leq 1.$$

Letting μM be the multiplier to the constraint, $q(\vec{v}, \cdot)$ is the solution if and only if

$$(1 - q_n(\vec{v}))^{M-1}(v_n - c) \leq \mu, = \mu \text{ if } q_n(\vec{v}) > 0, \text{ and } \sum_n q_n(\vec{v}) = 1.$$

As in large markets, this condition coincides with the one for the fully revealing equilibrium. Therefore, again, the fully revealing equilibrium is constrained efficient.

Proposition 3 *In finite markets, the fully revealing equilibrium is always constrained efficient.*

This result implies that in finite markets, our cheap-talk model can outperform the canonical (reserve) price posting model, unlike in large markets where both models achieve constrained efficiency. In finite markets with heterogeneity, posting models are typically not constrained efficient: Galenianos, Kircher and Virág (2011) provide a general inefficiency result for the price posting (directed search) model. For the reserve price posting (competing auctions) model, Julien, Kennes and King (2002) find efficiency for the case of two procurers. However, in our Online Appendix, we prove that if there are at least three procurers in the market and at least two of them are heterogeneous (i.e., there is procurer heterogeneity), then the market outcome fails to achieve constrained efficiency. In other words, the efficiency result by Julien, Kennes and King (2002) does not extend beyond the case of two procurers. Constrained efficiency of our model arises 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.

4.2.3 Limit Market Outcomes

We now show that the finite market outcomes indeed converge to the large market outcome in Section 3.1 as the market size (N) grows infinity.

Denote by $\tilde{q}^i_N$ the probability that each contractor selects a type v^i procurer and by $\tilde{u}_N$ contractors' market utility. Both $\tilde{q}^i_N$ and $\tilde{u}_N$ depend on the realizations of procurers' types and, therefore, are random variables. In order to characterize the limits of the random variables, suppose that for each N there are Nf^i number of type v^i procurers in the market.[27] Let q_N^i and u_N be the values that satisfy the equilibrium conditions for contractors' search strategies, that is,

$$\left(1 - \frac{q_N^i}{Nf^i}\right)^{\beta N - 1} (v^i - c) \leq u_N, = u_N \text{ if } q_N^i > 0, \text{ and } \sum_i q_N^i = 1.$$

As shown above, there exists a unique vector $(q_N^1, ..., q_N^I, u_N)$ that satisfies the conditions. As N tends to infinity, the deterministic sequence converges to a vector $(q^1, ..., q^I, u)$ that satisfies

$$e^{-\beta q^i/f^i}(v^i - c) \leq u, = u \text{ if } q^i > 0, \text{ and } \sum_i q^i = 1. \tag{5}$$

In order to show that the random vector $(\tilde{q}_N^1, ..., \tilde{q}_N^I, \tilde{u}_N)$ also converges to $(q^1, ..., q^I, u)$, it suffices to invoke the central limit theorem: for large N, the distribution of the number of type v^i procurers approximates the normal distribution with mean Nf^i and variance $f^i(1 - f^i)/N$. Since the variance disappears in the limit as N tends to infinity, it follows that $(\tilde{q}_N^1, ..., \tilde{q}_N^I, \tilde{u}_N)$ converges to the same limit as $(q_N^1, ..., q_N^I, u_N)$.

To show that the limit market outcome coincides with the large market outcome in Section 3.1, observe that (5) is equivalent to

$$\beta q^i = \max\left\{\log\left(\frac{v^i - c}{u}\right), 0\right\} f^i, \text{ and } \sum_i q^i = 1.$$

Combining the conditions yields

$$\beta = \sum_i \max\left\{\log\left(\frac{v^i - c}{u}\right), 0\right\} f^i = \int_{v > c+u} \log\left(\frac{v - c}{u}\right) dF(v).$$

This condition coincides with (3), which is the equilibrium condition for contractors' market utility in large markets. In other words, contractors' market utility in the limit market is identical to the one in the large market. It is then straightforward that the queue length of contractors for a type v procurer also converges to $\lambda(v^i) = \frac{\beta q^i}{f^i}$ as N tends to infinity.

4.2.4 Existence of the Fully Revealing Equilibrium

We now consider procurers' incentives and examine whether they would be willing to reveal their true types.

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 select

[27] We note that Nf^i needs not be an integer.

the procurer with a different probability ($q_n(v', \vec{v}_{-n})$ instead of $q_n(\vec{v})$) and quote different prices to her ($H_n(\cdot; v', \vec{v}_{-n})$ instead of $H_n(\cdot; \vec{v})$). Second, unlike in large markets, the deviation changes contractors' market utility from $u(\vec{v})$ 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 (unconditional) probability that procurer n receives at least one quote less than or equal to $b \leq \min\{v, v'\}$. When procurer n announces $\tilde{v}$, the probability is equal to

$$1 - (1 - q_n(\tilde{v}, \vec{v}_{-n})H_n(b; \tilde{v}, \vec{v}_{-n}))^M,$$

because $q_n(\tilde{v}, \vec{v}_{-n})H_n(b; \tilde{v}, \vec{v}_{-n})$ is the probability that a contractor selects the procurer and quotes a price below b, and there are M contractors in the market. Considering only the non-trivial case where $q_n(\tilde{v}, \vec{v}_{-n}) > 0$, Lemma 2 and the equilibrium conditions for contractors' search strategies imply

$$(1 - q_n(\tilde{v}, \vec{v}_{-n})H_n(b; \tilde{v}, \vec{v}_{-n}))^{M-1}(b - c) = U_n(\tilde{v}, \vec{v}_{-n}) = u(\tilde{v}, \vec{v}_{-n}).$$

It follows that procurer n's probability of getting the lowest quote below b is equal to

$$1 - \left(\frac{u(\tilde{v}, \vec{v}_{-n})}{b - c}\right)^{\frac{M}{M-1}}.$$

Notice that the probability would not be affected by procurer n's announcement $\tilde{v}$ if $u(\tilde{v}, \vec{v}_{-n})$ were independent of $\tilde{v}$. 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 (i.e., $\tilde{v}$ is the maximal quote made to the procurer), this makes a procurer strictly prefer truth-telling to downward deviations, while indifferent over upward deviations.

In finite markets, contractors' market utility $u(\tilde{v}, \vec{v}_{-n})$ does depend on procurer n's announcement. From (4), it is clear that $u(\tilde{v}, \vec{v}_{-n})$ is increasing in $\tilde{v}$. This is fairly intuitive, because a higher v means more surplus available in the market. Procurer n's probability of getting the lowest quote below b is strictly decreasing in contractors' market utility $u(\tilde{v}, \vec{v}_{-n})$. This is also intuitive, because a contractor would quote a higher price when he has a higher reservation value. Together, these imply that ceteris paribus, a procurer prefers announcing a lower message, so as to drive down contractors' market utility. Clearly, this makes downward deviations more attractable, while upward deviations less attractable.

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, but 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, measured by the number of procurers in the market 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, as shown above, each procurer's market power shrinks as the market size increases and becomes negligible in the limit. In other words, contractors' market utility is independent of a procurer's announcement in the limit as N tends to infinity, and thus the market utility effect is arbitrarily small in sufficiently large finite markets. We conclude that a procurer strictly prefers truth-telling to downward deviations, provided that the market size is sufficiently large.

Proposition 4 *In finite markets, the fully revealing equilibrium exists if the market size (N) is sufficiently large. In the equilibrium, a procurer strictly prefers truth-telling to both upward and downward deviations.*

4.3 Contractor Heterogeneity

We now generalize our analysis of the baseline model into the case of two-sided heterogeneity. The assumption of homogeneous contractors significantly simplified the notation and analysis. However, not only is it not sufficiently realistic, but also does it not allow us to explore the full extent to which communication facilitates trade. In price posting models, it is well-known that price discrimination is crucial for efficiency (see, in particular, Shi, 2002; Shimer, 2005; Peters, 2010). It allows procurers to prioritize more efficient contractors, thereby reducing the risk of mismatches. In other words, market outcomes are more efficient with price discrimination than without it. In our model, this raises a question of whether market communication could also subsume the role of price discrimination, which can be answered only in a setting where contractors are heterogeneous.

We consider the two most common specifications where contractors differ either in terms of their opportunity cost of working (see, e.g., Peters, 1997b; Mortensen and Wright, 2002) or in terms of observable skills that add the same value across firms (see, e.g., Shi, 2002; Shimer, 2005). We focus on the central ideas, relegating a formal analysis to the appendix.

4.3.1 Heterogeneous Opportunity Costs of Working

Setup. Suppose each contractor independently and identically draws his opportunity cost of working from the set $\mathcal{C} \equiv [\underline{c}, \overline{c}]$ according to the distribution function G with continuous and everywhere-positive density g at the beginning of the game. This is private information for each contractor. In addition, to ease the exposition and notation, assume that the set of procurers' project values $\mathcal{V}$ is also convex (i.e., $\mathcal{V} \equiv [\underline{v}, \overline{v}]$) and the distribution function F also admits a continuous and everywhere-positive density f. To avoid triviality, also assume that $\underline{c} < \underline{v}$ and $\overline{c} < \overline{v}$. All the other configurations of the physical environment are identical to those in Section 2.

It is convenient to define the following queue lengths. Denote by $\lambda(v,c)$ the queue length of type c contractors for a type v procurer. In addition, denote by $\Lambda(v,c)$ the queue length of contractors whose types are below c for a type v procurer. Formally, $\Lambda(v,c) \equiv \int_{\underline{c}}^{c} \lambda(v,x)dx$.

Efficiency of the fully revealing equilibrium. Consider the same communication game as in Section 2, and suppose procurers' types were fully revealed. Since revenue equivalence continues to hold, suppose each procurer runs the second-price auction with reserve price equal to her value. Then, a type c contractor is hired by a type v procurer if and only if there is no other contractor whose type is below c and, therefore, with probability $e^{-\Lambda(v,c)}$. The price paid to the contractor is determined by the second lowest contractor type among those who selected the same procurer. The second lowest contractor type is below c' if and only if at least one contractor below c' selects the procurer and, therefore, with probability $1 - e^{-\Lambda(v,c')}$. It follows that the expected payoff of a type c contractor who selects a type v procurer is equal to

$$e^{-\Lambda(y,c)}(v-c) - \int_{c}^{\overline{c}} (v-c')d\left(1 - e^{-\Lambda(v,c')}\right). \tag{6}$$

Observe that this is identical to the marginal social contribution of a type c contractor who selects a type v procurer. The first term is the expected social surplus he creates by selecting a type v procurer. The second term is the negative externality by the contractor to the contractors above his type: If a type c contractor did not select the procurer, then the procurer could have hired a contractor whose type is above c. It follows that contractors would search efficiently and the fully revealing equilibrium is constrained efficient.

Existence of the fully revealing equilibrium. Denote by $u(c)$ a type c contractor's market utility. In the fully revealing equilibrium, $u(c)$ coincides with (6) for any v such that $\lambda(y,c) > 0$. Differentiating (6) with respect to c yields

$$u'(c) = -e^{-\Lambda(v,c)}.$$

Notice that $e^{-\Lambda(v,c)}$ is the probability that a type c contractor who selects a type v procurer wins the project. Since $u'(c)$ is obviously independent of v, this implies that a type c contractor's winning probability must be independent of a procurer's type. This, in turn, implies that with the first-price auction, a type c contractor's price quote also should not vary according to a procurer's type. In other words, each contractor quotes the same price to all procurers. Let $w(c)$ denote the price quote by a type c contractor. It is clear that $w(\cdot)$ is continuous because of incentive compatibility, and strictly increasing because a lower cost contractor would quote a lower price due to the usual single crossing property.

To prove the existence of the fully revealing equilibrium, suppose a type v' procurer deviated and announced v, instead of v'. The procurer hires a contractor at a price below $b \leq \min\{v, v'\}$ if and only if at least one contractor who quotes less than b selects her, whose probability is $1 - e^{-\Lambda(v,c)}$ where c is such that $w(c) = w$. As shown above, this probability is independent of a procurer's announcement. Therefore, as in the homogeneous contractors case, a procurer's deviation does not affect her probability of hiring a contractor at each price. Given this, it is straightforward to show that, again as in the homogeneous contractors case, a procurer strictly prefers truth-telling to downward deviations, while is indifferent over all upward deviations.[28]

4.3.2 Heterogeneous Skills

Setup. Now suppose contractors have the same opportunity cost of working, but differ in terms of their observable skills. For simplicity, normalized c to 0. Each contractor's skill, denoted by s, is independently and identically drawn from the set $\mathcal{S} \equiv [\underline{s}, \overline{s}]$ according to the distribution function G at the beginning of the game. Skills are observable by the procurer. Assume, again, that the two distribution functions F and G admit continuous and everywhere-positive densities f and g, respectively. If a type s contractor is hired by a type v procurer at a price p, then the contractor receives utility p, while the procurer obtains utility $v + s - p$.[29] Since the analysis is essentially identical to the previous one, we illustrate only the necessary changes for this specification.

Efficiency of the fully revealing equilibrium. As before, suppose procurers' types were fully revealed and each procurer runs the second-price auction. In the current specification, a procurer's reserve price depends on a contractor's type. Specifically, a type v procurer's reserve price for a type s contractor is equal to $v + s$. Denote by $\lambda(v, s)$ the queue length of type s contractors for a type v procurer, and let $\Lambda(v, s) \equiv \int_s^{\overline{s}} \lambda(v, x)dx$ denote the queue length of the contractors above s for a type v procurer. Since a type s contractor is hired only when no other contractor above s selects the same procurer and his price is determined by the next most skilled

[28]For this result, it is necessary that the highest price quote to a type v procurer is equal to v. This follows from the fact that the highest contractor type who selects the procurer wins if and only if there is no other contractor, and thus his quote must be optimal conditional on him being the only quote provider to the procurer.

[29]Our analysis is restricted to the case where there is no complementarity between procurers' and contractors' types. The most general form of this specification would allow complementarity as well as two-sided heterogeneity, as in Shimer (2005). We do know that the fully revealing equilibrium continues to be efficient in such a more general setting, but do not know whether the fully revealing equilibrium would exist in our baseline model.

contractor, his expected payoff by selecting a type v procurer is equal to

$$e^{-\Lambda(v,s)}(v+s) - \int_{\underline{s}}^{s} (v+s')d\left(1-e^{-\Lambda(v,s')}\right).$$

This payoff is, again, identical to the contractor's marginal social contribution, with the first term representing the expected social surplus he creates by selecting a type v procurer and the second term the negative externality to the contractors below his type. All contractors search efficiently and the fully revealing equilibrium is constrained efficient.

Existence of the fully revealing equilibrium. Letting $u(s)$ be a type s contractor's market utility, as before, we get

$$u'(s) = e^{-\Lambda(v,s)}.$$

From here, again, we find that a contractor's winning probability is independent of a procurer's type and, therefore, with the first-price auction, each contractor quotes the same price to all procurers. Then, by the same reasoning as before, the fully revealing equilibrium exists because a procurer strictly prefers truth-telling to downward deviations and weakly to upward deviations.

5 Concluding Remarks

We consider a two-sided matching environment and show that cheap-talk communication can be perfectly informative and induce constrained efficient market outcomes. The choice of the trading protocol is crucial for the results. We show that constrained efficiency can be obtained if and only if the trading protocol adopted is revenue-equivalent to the first-price auction and does not disperse prices more than the first-price auction.[30] We also illustrate that our insights hold even in finite markets and with two-sided heterogeneity. Our results identify the settings in which cheap talk can serve as an efficient competitive instrument, in the sense that the central insights from the literature on competing auctions and competitive search continue to hold unaltered even without ex ante price commitment. They also shed light on an active market design question of whether reserve prices should be made public or not.

Although we consider only one round of market interaction, we conjecture that our results would carry over to repeated settings such as those analyzed in Peters (1991), and even to settings with business-cycle shocks and on-the-job-search as in Menzio and Shi (2011). In such a setting, if parties can write binding contracts about future actions once they meet, contractors essentially ask for parts of the ensuing surplus, and the rest of our results would apply. For

[30] To our knowledge, this role of price dispersion is new to the literature. Price dispersion (i.e., the failure of the law of one price) has received much attention for a long time. However, the literature has been mainly interested in explaining or generating price dispersion itself (see, e.g., Burdett and Judd, 1983).

other environments, we have less clear intuition. In particular, when contractors' types interact non-additively with procurers' types, the analysis becomes more intricate. While our exact technique (based on identical conditional distributions of bids below one's announcement) no longer applies, some of our basic insights are likely to carry over. We leave these and similar explorations in broader competitive market games for future work.

Appendix

Proof of Proposition 1. We first show that for each $\tilde{v} \in \mathcal{V}$, it is an equilibrium that all procurers below $\tilde{v}$ reveal their types, while all other procurers pool together. We prove later that there does not exist any other partitional equilibrium.

Contractors' equilibrium bidding strategies. It suffices to characterize contractors' equilibrium bidding strategies for a procurer above $\tilde{v}$. As in the proof of Lemma 1, a contractor's expected payoff is equal to the probability that he is the only quote provider $e^{-\lambda}$ times his expected utility when he knows that no other contractor selected the same procurer. For the latter part, define

$$v^*(\tilde{v}) \equiv argmax_{v \geq \tilde{v}} \int_v^{\overline{v}} (v-c)\frac{dF(v)}{1-F(\tilde{v})} = \frac{1-F(v)}{1-F(\tilde{v})}(v-c).$$

The strict quasi-concavity assumption guarantees that $v^*(\tilde{v})$ is uniquely determined. Then, a contractor's expected payoff is equal to

$$e^{-\lambda}\frac{1-F(v^*(\tilde{v}))}{1-F(\tilde{v})}(v^*(\tilde{v})-c).$$

Denote by $\tilde{H}(b;[\tilde{v},\overline{v}],\lambda)$ the probability that a contractor quotes a price weakly below b to a procurer in $[\tilde{v},\overline{v}]$. Then, for any b in the support of $\tilde{H}(\cdot;[\tilde{v},\overline{v}],\lambda)$, as in the proof of Lemma 1, $\tilde{H}(b;[\tilde{v},\overline{v}],\lambda)$ can be obtained from a contractor's indifference between $v^*(\tilde{v})$ and b.

Contractors' equilibrium search strategies. Let $u(\tilde{v})$ denote contractors' market utility in the prescribed strategy profile. In addition, abusing notation slightly, denote by $\lambda([\tilde{v},\overline{v}])$ the queue length for a procurer above $\tilde{v}$. Similarly to Section 3.1.2, the equilibrium queue lengths must satisfy

$$e^{-\lambda(v)}(v-c) \leq u(\tilde{v}), = u(\tilde{v}) \text{ if } \lambda(v) > 0, \forall v < \tilde{v},$$

$$e^{-\lambda([\tilde{v},\overline{v}])}\frac{1-F(v^*(\tilde{v}))}{1-F(\tilde{v})}(v^*(\tilde{v})-c) \leq u(\tilde{v}), = u(\tilde{v}) \text{ if } \lambda([\tilde{v},\overline{v}]) > 0,$$

and

$$\int_{\underline{v}}^{\tilde{v}} \lambda(v)dF(v) + \int_{\tilde{v}}^{\overline{v}} \lambda([\tilde{v},\overline{v}])dF(v) = \beta.$$

Let $\overline{u}$ be contractors' market utility in the fully revealing equilibrium. In addition, denote by $\underline{u}$ contractors' market utility in the babbling equilibrium, which can be obtained from the fact that $\tilde{v} = \underline{v}$ and $\lambda([\tilde{v},\overline{v}]) = \beta$ in the equilibrium. It is then straightforward to show that for each $\tilde{v} \in [\underline{v},\overline{v}]$, there exists a unique value of $u(\tilde{v}) \in [\underline{v},\overline{u}]$ that satisfies the above equilibrium. As can be easily expected, $u(\tilde{v})$ is increasing in $\tilde{v}$, that is, the more procurer types are revealed, the higher utility do contractors obtain.

Procurers' incentives at the communication stage. Now we prove that no procurer has an incentive to deviate in the prescribed strategy profile.

First, consider a type v' procurer where $v' < \tilde{v}$. The procurer has no incentive to deviate to any other $v < \tilde{v}$ for the same reason as in the fully revealing equilibrium. Suppose he deviates to above $\tilde{v}$. Only the quotes below $v' < \tilde{v}$ are relevant to his expected payoff, but, again due to the equilibrium conditions for contractors' search and bidding strategies, the probabilities of her getting those quotes are independent of his deviation. Consequently, she is indifferent between truth-telling and deviating to above $\tilde{v}$.

Now suppose $v' \geq \tilde{v}$, and consider a type v' procurer's deviation to $v < \tilde{v}$. As in the fully revealing equilibrium, her probability of hiring a contractor at a price below v does not change, but she no longer receives the quotes between $[v, \min\{v', v^*(\tilde{v})\}]$. Therefore, she has a strict disincentive to deviate below $\tilde{v}$.

Uniqueness. Suppose there exists another form of partitional equilibrium. In the equilibrium, there must exist v' and v'' such that the two types announce an identical message, and $v' < v'' < \overline{v}$. Without loss of generality, assume that v' and v'' are the minimal and maximal values in the partition element, respectively. In the equilibrium, a type v'' procurer necessarily has an incentive to deviate upward. If she deviates upward, then for any price below v'', she faces the same probability of hiring a contractor as when she truthfully reveals her type. When pooling the types in $[v', v'']$ (playing the prescribed strategy), her probability of hiring a contractor at a price below v' is the same as when she reveals her type, but her probabilities for prices between v' and v'' become strictly smaller: For any $b \in (v', v'']$,

$$e^{-\lambda([v',v''])\tilde{H}(b;[v',v''],\lambda([v',v'']))}\frac{F(v'') - F(b)}{F(v'') - F(v')}(b - c) = u = e^{-\lambda(v'')\tilde{H}(b;v'',\lambda(v''))}(b - c).$$

Since $\lambda([v', v''])\tilde{H}(b; [v', v''], \lambda([v', v''])) < \lambda(v'')\tilde{H}(b; v'', \lambda(v''))$ for any $b \in (v', v''] \cap \mathcal{Y}$, a type v'' procurer obtains strictly less than when she reveals her value (or deviates upward). ■

Proof of Proposition 2. The result that only the fully revealing equilibrium satisfies NITS that is applied to the highest procurer type follows immediately from the uniqueness proof of Proposition 1: If $\tilde{v} < \overline{v}$, then a type $\overline{v}$ would strictly prefer deviating upward and truthfully revealing her type.

Now suppose procurers face costs of lying, as in Kartik (2009). Specifically, take a continuous function $C : \mathcal{V} \times \mathcal{V} \to \mathbb{R}$ such that for any $v \in \mathcal{V}$, $C(m, v)$ is strictly decreasing in m if $m < v$ and strictly increasing in m if $m > v$. Then assume that a type v procurer suffers from lying cost $C(m, y)$ if she announces m. Since lying costs are concerned only with procurers' expected payoffs at the communication stage, no previous characterization, other than the existence of equilibria, is affected by the introduction of lying costs. Under this specification, we show that the fully revealing equilibrium is the unique equilibrium that satisfies either a plausible monotonicity condition or the intuition criterion (Cho and Kreps, 1987).

Message monotonicity. Suppose, as in Chen, Kartik and Sobel (2008) and Kartik (2009), we restrict attention to message-monotone equilibria in which a procurer with a higher value announces a weakly higher message. In any such equilibrium, a type $\overline{v}$ firm strictly prefers announcing $\overline{v}$ to any other message, because it would maximize its expected profit from the labor market (again, see the uniqueness proof of Proposition 1) as well as minimize the cost of lying. This immediately implies that only the fully revealing equilibrium satisfies the monotonicity condition.

Intuitive criterion. Fix a partitional equilibrium and let $\tilde{v}$ be the cutoff procurer type for the equilibrium (see Proposition 1). In the presence of lying costs, all the procurers in $[\tilde{v}, \overline{v}]$ must pool on message $\tilde{v}$: If they use multiple messages then, due to lying costs, different procurer types would self-select into different messages and, therefore, cannot induce the same outcome for different messages. In addition, if they pool on a message strictly above $\tilde{v}$, then the types sufficiently close to $\tilde{y}$ would prefer announcing $\tilde{y}$ (or any message below, but sufficiently close to, $\tilde{y}$) to pooling on the message.

Now consider a message $v > \tilde{v}$. Since each procurer is indifferent over all upward deviations in the absence of lying cost, message v is equilibrium dominated for all types below $\tilde{v} + \epsilon$ for some $\epsilon > 0$. On the other hand, if contractors believe that the message comes from the procurer types strictly above $\tilde{v}$, all the procurers above $\tilde{v}$ strictly prefer v to $\tilde{v}$. Therefore, message v is not equilibrium dominated for those firms. Applying the intuition criterion, contractors must believe that the messages above $\tilde{v}$ come from the procurer types above $\tilde{v}$. It is then straightforward that only the fully revealing equilibrium satisfies the intuitive criterion. ■

Formal Analysis of Contractor Heterogeneity

Here, we provide a formal analysis of the baseline model with contractor heterogeneity, which was introduced in Section 4.3. Since the same analysis applies to both specifications, we focus on the first specification where contractors differ in their opportunity cost of working.

Strategies and equilibrium. As in the homogeneous contractors case, we impose anonymity on agents' strategies. Procurers' strategies are a function $m : \mathcal{V} \to \mathcal{M}$ where $m(v)$ denotes the message sent by type v procurers. Contractors' search strategies are represented by a function $P : \mathcal{M} \times \mathcal{C} \to [0, 1]$ where $P(m, c)$ denotes the probability that a type c contractor selects a procurer who announces a message weakly below m. Contractors' (pure) bidding strategies are a function $b : \mathcal{M} \times \mathcal{C} \to \mathbb{R}_+$ where $b(m, c)$ denotes the price type c contractors quote to a procurer with message m. Procurers' optimal hiring strategies are trivial, so we omit their formal descriptions. The formal definition for the fully revealing equilibrium is essentially identical to that for the homogeneous contractors case and, therefore, is omitted as well.

Given procurers' communication strategies $m(v) = v$ and contractors' search strategies $P(\cdot, \cdot)$, the queue length of type c contractors for a type v procurer is given by

$$\lambda(v, c) = \frac{\beta g(c) dP(v, c)}{f(v)}. \tag{7}$$

In addition, let $\Lambda(v, c) = \int_{\underline{c}}^{c} \lambda(v, x) dx$.

Contractors' equilibrium bidding strategies. We utilize revenue equivalence to characterize contractors' equilibrium bidding strategies. Suppose a type v procurer runs a second-price auction with reserve price equal to v. As usual, bidding his cost is a weakly dominant strategy for each contractor. Then, a type c contractor obtains $v - c$ if no other contractor selects the procurer (ignoring contractors above v without loss of generality), whose probability is equal to $e^{-\Lambda(v,v)}$. If other contractors also select the procurer, then let c' be the lowest type among the contractors. Then, a type c contractor's payoff is $\max\{c' - c, 0\}$. Since the probability that c' is less than x is equal to $1 - e^{-\Lambda(v,x)}$, a type c contractor's expected payoff is equal to

$$U(v,c) \equiv e^{-\Lambda(v,v)}(v-c) + \int_c^{\overline{c}} (c'-c) d\left(1 - e^{-\Lambda(v,c')}\right).$$

Arranging the terms yields

$$U(v,c) = e^{-\Lambda(v,c)}(v-c) - \int_c^{\overline{c}} (v-c') d\left(1 - e^{-\Lambda(v,c')}\right). \tag{8}$$

The revenue equivalence implies that even with a first-price auction, a type c contractor's expected payoff must be as in (8). To deduce $b(v,c)$, recall that the winning probability of a type c contractor is equal to $e^{-\Lambda(v,c)}$ in the second-price auction. Since the first-price auction must give him the same winning probability, it follows that

$$e^{-\Lambda(v,c)}(b(v,c) - c) = U(v,c).$$

Efficiency of the fully revealing equilibrium. Let $u(c) \equiv \max_v U(v,c)$. Then, by the same argument as in Section 3.1.2, the equilibrium queue lengths $\lambda(\cdot,\cdot)$ are characterized by

$$u(c) \geq U(v,c) = e^{-\Lambda(v,c)}(v-c) - \int_c^{\overline{c}} (v-c') d\left(1 - e^{-\Lambda(v,c')}\right), \tag{9}$$

with equality holding if $\lambda(v,c) > 0$, and

$$\int_{\underline{v}}^{\overline{v}} \lambda(v,c) f(v) dv \leq \beta g(c), \forall c \in \mathcal{C}. \tag{10}$$

To establish the constrained efficiency of this market outcome, as in the homogeneous contractors case, consider a social planner who designs contractors' search strategies so as to maximize output in the market. Since it is optimal for each procurer to hire the contractor with the lowest cost, the social planner faces the following maximization problem:

$$\max P(\cdot,\cdot) \int_{\underline{v}}^{\overline{v}} \left(\int_{\underline{c}}^{\overline{c}} (v-c) d\left(1 - e^{-\Lambda(v,c)}\right) \right) dF(v).$$

Using (7) and the fact that $P(\overline{v}, c) \leq 1$ for any $c \in \mathcal{C}$, the problem can be rewritten as follows:

$$\max_{\lambda(\cdot,\cdot)} \int_{\underline{v}}^{\overline{v}} \left(\int_{\underline{c}}^{\overline{c}} (v-c) d\left(1 - e^{-\Lambda(v,c)}\right) \right) dF(v)$$

subject to

$$\int_{\underline{v}}^{\overline{v}} \lambda(v,c)dF(v) \leq \beta g(c), \forall c \in \mathcal{C}.$$

The objective function is concave, while the constraints are linear. Therefore, letting $\mu(c)$ be the multiplier on the constraint for c, $(\lambda(\cdot,\cdot),\mu(\cdot))$ is the solution if and only if

$$\mu(c) \geq e^{-\Lambda(v,c)}(v-c) - \int_{c}^{\overline{c}} (v-c')d\left(1-e^{-\Lambda(v,c')}\right),$$

with equality holding if $\lambda(v,c) > 0$, and

$$\int_{\underline{v}}^{\overline{v}} \lambda(v,c)dF(v) \leq \beta g(c), \forall c \in \mathcal{C}.$$

Notice that these optimality conditions coincide with those for contractors' equilibrium search strategies, (9) and (10). It follows that the market outcome is constrained efficient.

Existence of the fully revealing equilibrium. A proof for the existence of the fully revealing equilibrium is already in Section 4.3. Let us only mention that the (almost-everywhere) differentiability of $u(\cdot)$ follows from the fact that $u(\cdot)$ is continuous and strictly decreasing, which in turn derives from the usual incentive compatibility argument that each type can mimic any other type.

References

Acemoğlu, Daron and Robert Shimer, "Efficient unemployment insurance," *Journal of Political Economy*, 1999a, *107* (5), 893–928.

_ **and** _ , "Holdup and efficiency with search frictions," *International Economic Review*, 1999b, *40* (4), 827–849.

Albrecht, James, Pieter A. Gautier, and Susan Vroman, "Equilibrium directed search with multiple applications," *Review of Economic Studies*, 2006, *73* (4), 869–891.

_ **,** _ **, and** _ , "Directed search in the housing market," *mimeo*, 2013.

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.

Basov, Suren, Ian King, and Lawrence Uren, "The employed, the unemployed, and the unemployable: directed search with worker heterogeneity," *mimeo*, 2011.

Burdett, Kenneth and Kenneth L. Judd, "Equilibrium price dispersion," *Econometrica*, 1983, *51* (4), 955–969.

— , **Shouyong Shi, and Randall Wright**, "Pricing and matching with frictions," *Journal of Political Economy*, 2001, *109* (5), 1060–1085.

Chen, Ying, Navin Kartik, and Joel Sobel, "Selecting cheap-talk equilibria," *Econometrica*, 2008, *76* (1), 117–136.

Cho, In-Koo and David M Kreps, "Signaling games and stable equilibria," *The Quarterly Journal of Economics*, 1987, *102* (2), 179–222.

Crawford, Vincent P. and Joel Sobel, "Strategic information transmission," *Econometrica*, 1982, *50* (6), 1431–1451.

Delacroix, Alain and Shouyong Shi, "Directed search on the job and the wage ladder," *International Economic Review*, 2006, *47* (2), 651–699.

Farmer, Amy and Victor Stango, "Ask prices, offers, and time to sale in an online exchange," *Economic Inquiry*, 2004, *42* (1), 14–28.

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.

Fernández-Blanco, Javier, "Labor market equilibrium with rehiring," *International Economic Review*, 2013, *54* (3), 885–914.

Galenianos, Manolis and Philipp Kircher, "Directed search with multiple job applications," *Journal of Economic Theory*, 2009, *114* (2), 445–471.

— , — , **and Gábor Virág**, "Market power and efficiency in a search model," *International Economic Review*, 2011, *52* (1), 85–104.

Gautier, Pieter A. and Christian L. Holzner, "Simultaneous search and network efficiency," *mimeo*, 2013.

Geromichalos, Athanasios, "Directed search and optimal production," *Journal of Economic Theory*, 2012, *147* (6), 2303–2331.

Grossman, Sanford J. and Motty Perry, "Sequential bargaining under asymmetric information," *Journal of Economic Theory*, 1986, *39* (1), 120–154.

Guerrieri, Veronica, "Inefficient unemployment dynamics under asymmetric information," *Journal of Political Economy*, 2008, *116* (4), 667–708.

— **and Robert Shimer**, "Dynamic adverse selection: a theory of illiquidity, fire sales, and flight to quality," *mimeo*, 2013.

— , — , **and Randall Wright**, "Adverse selection in cometitive search equilibrium," *Econometrica*, 2010, *78* (6), 1823–1862.

Gul, Faruk and Hugo Sonnenschein, "On delay in bargaining with one-sided uncertainty," *Econometrica*, 1988, *56* (3), 601–611.

Horowitz, Joel L., "The role of the list price in housing markets: theory and an econometric model," *Journal of Applied Econometrics*, 1992, *7* (2), 115–129.

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.

Kartik, Navin, "Strategic communication with lying costs," *Review of Economic Studies*, 2009, *76* (4), 1359–1395.

Kircher, Philipp, "Efficiency of simultaneous search," *Journal of Political Economy*, 2009, *117* (5), 861–913.

Krishna, Vijay, *Auction Theory*, Academic Press, San Diego, CA, 2009.

Lester, Benjamin, "Information and prices with capacity constraints," *American Economic Review*, 2011, *101* (4), 1591–1600.

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.

_ and Shouyong Shi, "Block recursive equilibria for stochastic models of search on the job," *Journal of Economic Theory*, 2010, *145* (4), 1453–1494.

_ and _ , "Efficient search on the job and the business cycle," *Journal of Political Economy*, 2011, *119* (3), 468–510.

Merlo, Antonio and François Ortalo-Magné, "Bargaining over residential real estate: evidence from England," *Journal of Urban Economics*, 2004, *56* (2), 192–216.

Moen, Espen R., "Competitive search equilibrium," *Journal of Political Economy*, 1997, *105* (2), 385–411.

Montgomery, James D., "Wage dispersion and interindustry wage differentials," *The Quarterly Journal of Economics*, 1991, *106* (1), 163–179.

Mortensen, Dale T. and Randall Wright, "Competitive pricing and efficiency in search equilibrium," *International Economic Review*, 2002, *43* (1), 1–20.

Peters, Michael, "Bertrand equilibrium with capacity constraints and restricted mobility," *Econometrica*, 1984, *52* (5), 1117–1127.

_ , "Ex ante price offers in matching games non-steady states," *Econometrica*, 1991, *59* (5), 1425–1454.

_ , "On the equivalence of Walrasian and non-Walrasian equilibria in contract markets: the case of complete contracts," *Review of Economic Studies*, 1997a, *64* (2), 241–264.

_ , "A competitive distribution of auctions," *Review of Economic Studies*, 1997b, *64* (1), 97–123.

_ , "Limits of exact equilibria for capacity constrained sellers with costly search," *Journal of Economic Theory*, 2000, *95* (2), 139–168.

_ , "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 Comission, 2011.

Rocheteau, Guillaume and Randall Wright, "Money in search equilibrium, in competitive equilibrium, and in competitive search equilibrium," *Econometrica*, 2005, *73* (1), 175–202.

Rogerson, Richard, Robert Shimer, and Randall Wright, "Search-theoretic models of the labor market: survey," *Journal of Economic Literature*, 2005, *43* (4), 959–988.

Rudanko, Leena, "Labor market dynamics under long-term wage contracting," *Journal of Monetary Economics*, 2009, *56* (2), 170–183.

_ , "Aggregate and idiosyncratic risk in a frictional labor market," *American Economic Review*, 2011, *101* (6), 2823–2843.

Shi, Shouyong, "A directed search model of inequality with heterogeneous skills and skill-biased technology," *Review of Economic Studies*, 2002, *69* (2), 467–491.

_ , "Directed search for equilibrium wage-tenure contracts," *Econometrica*, 2009, *77* (2), 561–584.

Shimer, Robert, "The assignment of workers to jobs in an economy with coordination frictions," *Journal of Political Economy*, 2005, *113* (5), 996–1025.

Virág, Gábor, "Competing auctions: finite markets and convergence," *Theoretical Economics*, 2010, *5* (2), 241–274.