Essays on

Procurement with Information Asymmetry

by

Bin Hu

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2011

Doctoral Committee:

Professor Izak Duenyas, Co-Chair
Associate Professor Damian R. Beil, Co-Chair
Professor Xiuli Chao
Associate Professor Uday Rajan

To my parents, whose love and caring are unconditional

ACKNOWLEDGEMENTS

First and foremost, I would like to express my gratitude towards my academic advisors and dissertation committee co-chairs, Professor Izak Duenyas and Professor Damian R. Beil, for their continual guidance during the past five years. They have devoted countless hours of their valuable time to helping my research progress; I cannot possibly complete my dissertation without their supervision. I would also like to thank my committee members, Professor Uday Rajan and Professor Xiuli Chao, for their constructive advices that help me improve the quality of my dissertation, as well as the anonymous associate editors and reviewers at *Management Science*, whose suggestions greatly contributed to the papers upon which this dissertation is based.

I would like to thank all faculty of the Department of Operations and Management Science. I have greatly benefited from them through both research seminars and informal conversations. I must also express special gratitude towards Professor Roman Kapuscinski and Professor Hyun-Soo Ahn, the Doctoral Program Coordinators, and Mr. Brian Jones, the Doctoral Program Assistant Director, for helping me with administrative matters over the years. The Ross School of Business, the Department of Operations and Management Science and Professor Izak Duenyas have financially supported me during various stages of my study, and I am grateful for their generosity. Finally, I thank my parents for their unconditional love and encouragement, and my girlfriend Yuanyuan Zeng for her constant support and caring.

I am solely responsible for all the mistakes, ambiguities and other shortcomings that remain in this dissertation.

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF APPENDICES

Appendix

CHAPTER I

Introduction

Sourcing, once seen as a tactical function of vertically integrated firms, has today become strategic for firms that now rely on extensive supply chains. In the US, manufacturers on average spend 40-60% of their revenue on procuring goods and services (U.S. Department of Commerce 2005). Vertical disintegration and specialization of firms lead to complex relationships within a supply chain; for example, dependence between firms across different tiers and competition among firms within each tier often coexist. In addition, firms generally harbor private information. For example, a supplier does not wish its cost data to be known by its customer (a buyer), and a buyer does not wish its demand data to be known by its supplier. The complex relationships and information asymmetry make firms' interactions highly strategic, and the recent rapid globalization of supply chains has magnified these effects. How should firms in supply chains of various structures make strategic procurement decisions in the presence of information asymmetry? The three essays in this dissertation study three specific problems on this topic. I will briefly introduce each essay below.

In the essay "Does Pooling Component Demands when Sourcing Lead to Higher Profits?" (Chapter II), I consider whether pooling purchases for a component used in multiple products with uncertain demands results in increased profits for the buyer. Demand pooling is a prominent concept widely taught in Operations Management courses, and the received intuition indicates that demand pooling is beneficial for the buyer due to variability reduction. However, the standard pooling logic assumes the price of the purchased component is exogenous and fixed. In this essay, I consider a setting where the buyer purchases the component from a sole-source supplier who strategically designs optimal price-quantity contracts to extract profit from the buyer. The supplier's ability to extract profit is mitigated by the fact that the buyer is privileged with superior demand information (e.g., the buyer may have privately surveyed her customers to forecast demands). I show that pooling may actually result in *decreased* profits for the buyer facing a powerful, strategic supplier. One of the key insights is that the variability reduction obtained by pooling can sometimes harm the buyer because less demand variability makes it easier for the supplier to extract higher profits through optimal pricing (adapted to the demands). I characterize cases when pooling in the presence of a sole-source strategic supplier is disadvantageous, and also provide insights into when it is still advantageous.

In the essay "Simple Auctions for Supply Contracts" (Chapter III), I study a simple, practically implementable optimal procurement mechanism for a newsvendor-like problem where the buyer's (newsvendor's) purchase price of the supplies is not fixed, but determined through interactions with candidate suppliers, and the suppliers' production costs are their private information. Previous literature has studied the buyer's optimal mechanisms in this setting, notably Chen (2007). Chen showed that the buyer's optimal procurement mechanism can be implemented by the buyer offering the suppliers a revenue function (specifying a payment for each quantity the buyer may purchase) and then auctioning off the supply contract with the specified revenue function to the highest bidding supplier. However, auctioning of supply contracts with a specified revenue function is seldom observed in practice, and suppliers would have to perform complex calculations in order to bid effectively in such an auction. In this essay, I show that the optimal mechanism can be implemented by a simple modified version of the standard open-descending auction for a fixed-quantity contract. What distinguishes this mechanism is its simplicity and familiarity for the suppliers — open-descending price auctions are ubiquitous in practice, and the suppliers' decision making in this mechanism is almost trivial. The fact that the suppliers can use very simple strategies in a familiar environment increases the chance that suppliers are willing to participate in such a mechanism. I further show that this simple mechanism can be generalized to ex ante asymmetric suppliers and a class of non-linear production costs, whereas Chen (2007) treated the case with ex ante symmetric suppliers with linear production costs.

The majority of procurement research focuses on interactions between buyers and their direct suppliers. Expanding this scope to include another tier of the supply chain, in the essay "Price-Quoting Strategies of a Tier-Two Supplier" (Chapter IV), I study the price-quoting strategies used by a tier-two supplier, whose tier-one customers compete for an OEM's indivisible contract. At most one of the tier-two supplier's quotes will ultimately result in downstream contracting and hence produce revenue for her, and when the tier-one suppliers' costs of fulfilling the OEM's contract are too high, the OEM may not award the contract to either tier-one supplier, in which case the tier-two supplier cannot earn any revenue. I characterize the tier-two supplier's optimal price-quoting strategies and show that she will use one of two possible types of strategies, with her choice depending on the tier-one suppliers' profit potentials: *secure*, whereby she will always have business; or *risky*, whereby she may not have business. Addressing potential fairness concerns, I also study price-quoting strategies in which all tier-one suppliers receive equal quotes. Finally, I show that a tier-two supplier's optimal mechanism resembles auctioning a single quote among the tier-one suppliers.

For clarity, in this dissertation principals are referred to as "she", and agents as "he".

CHAPTER II

Does Pooling Component Demands when Sourcing Lead to Higher Profits?

2.1 Introduction

There exists a large research literature on inventory pooling, dating back to the seminal paper by Eppen (1979). In its most basic form, inventory pooling refers to satisfying several demand streams from a common inventory. In Eppen's canonical example, a firm with several locations (a steel wholesaler with several satellite warehouses) considers satisfying different locations' demands from a central stock. More generally, the inventory pooling concept applies to areas such as component commonality, whereby the same component is used for multiple products, rather than a specific, non-interchangeable component being used for each product.

A central theme of the inventory pooling literature is that "statistical economies of scale" is a key benefit that makes pooling attractive for firms. As Eppen (1979) shows, for the simple case in which each location's demand is normally and symmetrically distributed, the amount of cost reduction enjoyed by the firm is positive and increases linearly in the standard deviation of each location's demand. In short, pooling enables variability reduction which results in higher profits. This powerful intuition is widely taught in operations management courses when examining pooling.

In this chapter I ask if and how the above intuition will carry over to the following problem. Consider two products a and b with independent demands, sold by OEMs A and B, respectively. Both products require a common component c from a sole-source supplier, and OEMs A and B independently approach the supplier and make separate agreements to purchase the component. Alternatively, instead of OEMs A and B, suppose one OEM D makes both products a and b (with the same demands as before), and obtains the component c from the sole-source supplier. Would OEM D experience higher profits than A and B combined because of the reduction in demand variability compared to the two separate OEMs?

I became interested in this problem after interacting with a large manufacturer on various supply chain issues. This manufacturer has several divisions, some of which utilize common components. I learned that in some cases these divisions bought the components from the same sole-source supplier. Additionally I noted that for some such components, the divisions purchased the components and satisfied demands independently (like OEMs A and B), while for other components, the divisions combined the component purchases and inventories (like OEM D). Initially, I thought that these different ways of treating various components stemmed from a lack of consistency rather than being driven by strategic purposes. I thought that the second ("OEM D") approach would be better in all cases because of the variability reduction. However, I was intrigued by the fact that these components were sole-sourced, which led me to think more carefully about the problem. This chapter shows that my initial supposition based on adapting received wisdom to this situation with a sole-source supplier was in fact incorrect, and that the decisions involved can be much more subtle and interesting.

In answering this question, classic inventory pooling theory would suggest that OEM D's profit is higher than the sum of OEMs A and B's profits. For instance, assuming demands for both products are normal and i.i.d., the classic insight is that the profit advantage enjoyed by OEM D would be positive, and linearly increasing in the standard deviation of the demand streams. Yet underlying this insight is a tacit assumption that the per item cost of the component c is the same, whether purchased by OEMs A, B or D. While this might appropriately model a situation in which the component is a commodity, in the case I observed, as well as what I model, the component was being bought from a sole supplier. Sole suppliers are common in many industries (for example, The Economist (2009) points out that 90% of the micro-motors used to adjust the rear-view mirrors in cars are made by Mabuchi, and TEL makes 80% of the etchers used in making LCD panels. In these cases even though other suppliers do exist, they generally could not provide the same level of quality or performance, therefore a buyer wanting to use a high quality product really has only one choice). The core of the problem I address boils down to the following question: Does presenting a less variable demand stream to a powerful sole supplier necessarily result in higher profits?

In stark contrast to the usual "statistical economies of scale" benefits of inventory pooling, I find that the presence of a powerful and strategic sole supplier can actually result in situations where OEMs A and B (with unpooled demands for c) may have higher profits in total than OEM D. That is, I find that a strategic supplier can reverse the common wisdom about the benefit of variability reduction through pooling. One of my key insights is that, in the presence of a powerful strategic supplier who can set prices for the component based on her expectations about the demands, the variability reduction obtained by pooling can sometimes turn into a disadvantage. This is because less demand variability can make it easier for the supplier to extract higher profits through optimal contract design (adapted to the demand). In this chapter, I characterize cases when pooling in the presence of a sole strategic supplier is disadvantageous, and also provide insights into when it is still advantageous.

The remainder of this chapter flows as follows: §2.2 provides a literature review.

My model is introduced in §2.3. §2.4 provides an analysis of which state of the world (pooled or unpooled) yields higher profits for the purchasing OEM(s). §2.5 concludes the chapter.

2.2 Literature Review

My research is closely related to two streams of existing literature. The first stream is component commonality and inventory pooling. Stemming from the seminal paper by Eppen (1979), a very large literature explores statistical economies of scale (the reduction of uncertainty upon merging multiple stochastic demand streams), and shows that under various settings, component commonality, storage centralization, and inventory sharing can reduce operations costs. Examples from this literature include Eppen and Schrage (1981), Gerchak and He (2003), Benjaafar et al. (2005), and more recently Hanany and Gerchak (2008). However, one tacit assumption in this literature is that the purchase price of the good discussed is exogenous, and usually the supplier is not modeled. This would be a reasonable assumption if the good is a commodity, but less so when one supplier is the only source of supply (e.g., because the supplier owns a patented technology). In this chapter I assume the component can only be purchased from a sole-source supplier who strategically takes into account the operational structure of the buyer(s), which leads to the second related stream of literature — procurement contract design. Based on the principal-agent model (e.g. Laffont and Martimort (2002)), this literature analyzes how a powerful, profit-maximizing member of the supply chain (the principal, may be a buyer or a supplier) should optimally design procurement contracts for the other members (agents) who possess private information. The principal-agent model captures the general practice of tailoring a contract to a specific buyer or supplier (instead of relying on one-size-fits-all contracting), and is a canonical modeling construct. In the operations management literature, several papers take the perspective of a buyer and find the optimal supplier selection and contracting mechanism, when the buyer faces operational issues like the need for fast delivery (Cachon and Zhang (2006)), random demand (Chen (2007)), or uncertain supplier qualifications (Wan and Beil (2009)). I analyze a different operational issue, namely component pooling, and study a supplier who designs contracts. Other operations papers have examined the supplier's perspective (although absent the component pooling issue which is the crux of my research), with early references including Corbett and de Groote (2000), and Ha (2001).

Relatively distantly related is a literature on the analysis of group purchasing, particularly coalition forming and stability issues, for example Hartman and Dror (2003). This literature usually assumes that either each buyer faces uncertain demand and group purchasing benefits the buyers due to statistical economies of scale, or the supplier announces a price schedule that offers greater discounts for larger purchase quantities, then analyzes how to allocate the benefit from group purchasing to form a stable coalition. In both cases the supplier is not acting strategically, and the pre-assumed benefit of group purchasing is a premise for the analysis. In contrast, I model a strategic supplier and ask whether there is always benefit from pooling purchases. Thus the core research problems are very different.

2.3 Model and Preliminaries

I consider a stylized model in which an original equipment manufacturer (OEM) approaches the sole supplier of a component to ask for pricing. The OEM needs this component to manufacture an end product i (e.g., product a or b that I introduced in §2.1). Assume each end product requires a single component. For tractability, I make the simplifying assumption that product i's mean demand μ^i can only take one of n known values μ^i_θ, where θ denotes one of n different demand types. For example, the demand type θ could be "high" or "low", with corresponding mean demands $\mu^i_{\text{high}} = 20,000$ and $\mu^i_{\text{low}} = 10,000$, in which case $n = 2$. Prior to approaching the supplier for the component, the OEM learns his demand type (e.g., whether demand is going to be high or low). However, actual demand is likely to differ from its expected mean when it is realized; to model this aspect I assume that product i of demand type θ will actually experience a realized demand $\mu_\theta^i + e^i$ where e^i is a random forecast error and is assumed to have mean zero, pdf f^i and cdf F^i. I assume e^i is independent of the product's demand type θ. Therefore, even though the OEM can find out his mean future demand μ_θ^i when he learns his demand type, the actual demand $\mu_\theta^i + e^i$ remains a random variable to the OEM until it is realized.

The supply chain is decentralized (the supplier and the OEM are independent decision makers) and I assume the supplier and the OEM each seek to maximize their own expected profits. The component that the OEM needs is made by only this supplier; given the market power of the supplier, she can offer the OEM take-it-or-leave-it contracts. In such a setting, if the supplier knew exactly the demand type (i.e., the mean demand) of the product sold by the OEM, she could extract all expected supply chain profits, always leaving zero profits for the OEM (or leaving the OEM just his reservation profit that would induce him to participate — which I assume to be zero without loss of generality). Of course, this would then trivialize any comparisons of the OEMs' profits. However, OEMs are usually privileged with better information about their end product demands. To model this aspect, I assume that the demand type that the OEM finds out is his private information; the supplier has a prior belief on the distribution of the product's demand types but does not know the actual demand type. I assume that the supplier's prior on product i's demand types is that with probability p_θ^i, the demand type will be θ (i.e., the mean demand will be μ_θ^i). Clearly, the forecast error e^i is a random variable for the supplier just as it is for the OEM. This means that the actual demand is uncertain to both the OEM and the supplier, but the demand type (i.e., the forecasted mean demand) is known only to the OEM. Therefore, although both parties face demand uncertainties, the

OEM has strictly more information than the supplier.

The timing of the events in the model is as follows:

Stage 1: The OEM learns his demand type.

Stage 2: The supplier offers the OEM a menu of contracts consisting of quantity-payment pairs (Q_θ, t_θ), each meant for a potential type θ. That is, the OEM can buy Q_θ units at total cost t_θ. The OEM decides which (if any) quantity-payment pair to choose.

Stage 3: Once the OEM has chosen the contract, the supplier produces the agreed-upon quantity at per unit cost c, delivers the units to the OEM, and receives the corresponding payment.

Stage 4: The OEM processes the components into finished goods. The demand for the finished goods is then realized. The OEM receives revenue r for each unit of demand he satisfies. I assume that $r > c$ and define $q \doteq c/r$ ($c = qr$). Unsatisfied demand is lost and excess inventory has no salvage value.

The supplier designs and uses an optimal menu of contracts in Stage 2; doing so maximizes her expected profit among all mechanisms that, ultimately, result in quantity and money being exchanged between her and the OEM. This setup implies that the supplier has market power, i.e., there are other OEMs so even a pooled OEM (like OEM D in my motivating example) is not a monopsony. Due to the revelation principle, the supplier can limit the search for an optimal menu of contracts to those that induce the OEM to choose the contract designed for his demand type. Formally, my model is

$$\max_{Q_\theta, t_\theta} \quad E_\Theta[t_\Theta - cQ_\Theta] \tag{2.1a}$$

$$\text{s.t.} \quad E_e[r \min\{\mu_\theta + e, Q_\theta\}] - t_\theta \geq 0, \forall\ \theta \tag{2.1b}$$

$$E_e[r \min\{\mu_\theta + e, Q_\theta\}] - t_\theta \geq E_e[r \min\{\mu_\theta + e, Q_{\theta'}\}] - t_{\theta'}, \forall\ \theta' \neq \theta \tag{2.1c}$$

where random variable Θ reflects the supplier's prior on the OEM's types. In the above formulation (and those to follow where it does not cause confusion), I suppress superscript "i" for readability. (2.1b) is the *participation constraint (PC)* with reservation profit set to zero, and (2.1c) is the *incentive compatibility constraint (IC)* which ensures that the OEM chooses the contract designed for his demand type. Following convention, I define *information rent* π_θ as the expected profit of an OEM having demand type θ who chooses the contract designed for his demand type:

$$\pi_\theta \doteq E_e[r \min\{\mu_\theta + e, Q_\theta\}] - t_\theta.$$

For convenience and without loss of generality, henceforth I denote a contract by a quantity-information rent pair (Q_θ, π_θ) rather than a quantity-payment pair. Using this notation and explicitly writing out the expectations, (2.1a)-(2.1c) can be shown to be equivalent to

$$\max_{Q_\theta, \pi_\theta} \quad \sum_\theta p_\theta \left[r \left((1-q)Q_\theta - \int_{-\infty}^{Q_\theta - \mu_\theta} F(x)\mathrm{d}x \right) - \pi_\theta \right] \tag{2.2a}$$

$$\text{s.t.} \quad \pi_\theta \geq 0, \forall\, \theta \tag{2.2b}$$

$$\pi_\theta \geq \pi_{\theta'} + r \int_{Q_{\theta'} - \mu_\theta}^{Q_{\theta'} - \mu_{\theta'}} F(x)\mathrm{d}x, \forall\, \theta' \neq \theta. \tag{2.2c}$$

The formulation can be further simplified, once one notices that my model satisfies the *Spence-Mirrlees property* (Laffont and Martimort (2002), p.53), namely the OEM's *marginal rate of substitution* $\frac{\partial \pi_\theta / \partial Q_\theta}{\partial \pi_\theta / \partial t_\theta}$ is monotonic in type θ. With this property, all the constraints can be substituted with the lowest type's participation constraint, local downward incentive constraints, and *monotonicity constraints (MC)* (contract quantity non-decreasing in type). Suppose the possible demand types faced by an OEM are such that $\mu_{\theta_1} \leq \mu_{\theta_1} \leq \cdots \leq \mu_{\theta_n}$; then (2.2a)-(2.2c) are equivalent to

$$\max_{Q_\theta, \pi_\theta} \quad \sum_{\theta=\theta_j} p_\theta \left[r \left((1-q)Q_\theta - \int_{-\infty}^{Q_\theta - \mu_\theta} F(x)\mathrm{d}x \right) - \pi_\theta \right] \tag{2.3a}$$

$$\text{s.t.} \quad \pi_{\theta_1} \geq 1 \tag{2.3b}$$

$$\pi_{\theta_{j+1}} \geq \pi_{\theta_j} + r \int_{Q_{\theta_j} - \mu_{\theta_{j+1}}}^{Q_{\theta_j} - \mu_{\theta_j}} F(x)\mathrm{d}x, \quad j = 1, \ldots, n-1 \tag{2.3c}$$

$$Q_{\theta_{j+1}} \geq Q_{\theta_j}, \quad j = 1, \ldots, n-1. \tag{2.3d}$$

Notice that Equations (2.3a)-(2.3d) describe the problem that the strategic supplier solves to derive an optimal menu of quantity-payment pairs (represented by equivalent quantity-information rent pairs in the above formulation) to offer to the OEM. Since there is information asymmetry in my setting, the supplier cannot obtain *all* of the supply chain profit, but has to leave some information rent to the OEM which constitutes the OEM's profit. The fundamental question that I am trying to answer is if pooling or not pooling demands results in higher total profits (information rents) for the OEMs. Next, I obtain results that characterize the information rent structure in preparation for exploring this question.

2.3.1 Solving for Information Rent

By the analysis on page 43 of Laffont and Martimort (2002), I know participation and incentive compatibility constraints (2.3b) and (2.3c) are binding at optimality. Applying this insight, the objective function (2.3a) can be recast as a function solely of Q_θ. Relaxing the MCs (2.3d) for now (I will test these constraints ex post, and in the case of them being violated, revise my solution), the first-order condition (FOC) for (2.3a) as a function of Q_{θ_j} can be written as:

$$
\begin{aligned}
&\Pr(\Theta = \theta_j)(1-q) = \Pr(\Theta \geq \theta_j)F(Q_{\theta_j} - \mu_{\theta_j}) - \Pr(\Theta \geq \theta_{j+1})F(Q_{\theta_j} - \mu_{\theta_{j+1}}) \\
&\iff \lambda(\theta_j)(1-q) = F(Q_{\theta_j} - \mu_{\theta_j}) - (1-\lambda(\theta_j))F(Q_{\theta_j} - \mu_{\theta_{j+1}}), \quad j < n, \qquad (2.4)
\end{aligned}
$$

where

$$\lambda(\theta_j) \doteq \frac{\Pr(\Theta = \theta_j)}{\Pr(\Theta \geq \theta_j)},$$

and the information rents can be derived recursively as:

$$\pi_1 = 0, \quad \pi_{\theta_{j+1}} = \pi_{\theta_j} + r\int_{Q_{\theta_j}-\mu_{\theta_{j+1}}}^{Q_{\theta_j}-\mu_{\theta_j}} F(x)\mathrm{d}x. \qquad (2.5)$$

For concision, I use $\Theta \geq \theta_j$ to denote $\Theta \in \{\bar{\theta} | \mu_{\bar{\theta}} \geq \mu_{\theta_j}\}$. As an example, if Θ can take two values, high or low, then $\Pr(\Theta \geq \text{low}) = \Pr(\Theta = \text{low or high})$ and $\Pr(\Theta \geq \text{high}) = \Pr(\Theta = \text{high})$.

Equation (2.3a), which represents the supplier's objective function, is generally not concave in its decision variable Q_{θ_j}. Consequently, solving FOC (2.4) cannot always guarantee a global maximizer. In the following proposition I provide a set of *sufficient* conditions that guarantee (2.4) has a unique solution and it is the global maximizer.

Proposition 2.1. *Equation* (2.4) *has a unique solution and the solution is a global maximizer of the supplier's expected profit, if the forecast error e's pdf f satisfies the following conditions:*

1. $f(-x) = f(x)$;

2. $f(x)$ *is continuous in* x;

3. $f(x_2) \leq f(x_1)$ *for all* $x_2 > x_1 \geq 0$;

4. *For all* $\delta > 0$ *and* $x \geq 0$, $f(x+\delta)/f(x)$ *is non-increasing in* x.

The conditions in Proposition 2.1 ensure that e is reasonably well-behaved. The first three conditions require that e has a symmetric, continuous and unimodal pdf. The fourth condition has the intuitive meaning that the pdf must be sufficiently smooth. This condition will be violated, for example, if the pdf is piecewise linear and alternates between being flat and steep. Considering that e is a demand forecast error, the conditions in Proposition 2.1 are quite natural and mild. In fact, it is trivial to test that many common distributions including uniform, triangular and normal satisfy all four conditions. For the rest of this chapter, I assume these conditions are satisfied, therefore Equation (2.4) determines the unique global maximizer.

The fundamental problem that I am addressing is the following: If two OEMs making two separate products a and b that use the same component independently go through the purchasing process described earlier in this section, would they actually experience higher profits than one OEM that makes both products and pools the demands for the common component? To gain insights into this fundamental problem using the simplest possible setting, I assume that the mean demand for products a and b can be of two possible types: Product i ($i = a, b$) has either high mean demand μ_h^i, or low mean demand μ_l^i. Thus, when OEM A approaches the supplier, he is offered a menu of two contracts, (Q_h^a, π_h^a) and (Q_l^a, π_l^a). Similarly, when OEM B approaches the supplier he is offered menu (Q_h^b, π_h^b) and (Q_l^b, π_l^b). I define $\delta^i \doteq \mu_h^i - \mu_l^i$ and assume $\delta^i \geq 0$. Unlike OEMs A and B who each sell a single product, OEM D sells both products. For consistency, I assume that OEM D will have demand type hh, hl, lh or ll, corresponding to a total mean demand for products a and b of $\mu_{hh}^D, \mu_{hl}^D, \mu_{lh}^D$, or μ_{ll}^D. (For example, type hl means product a's mean demand is μ_h^a and product b's mean demand is μ_l^b, and thus $\mu_{hl}^D = \mu_h^a + \mu_l^b$.) Accordingly, OEM D will be offered a menu of four contracts: (Q_{hh}^D, π_{hh}^D), (Q_{hl}^D, π_{hl}^D), (Q_{lh}^D, π_{lh}^D) and (Q_{ll}^D, π_{ll}^D). Without loss of generality, I assume $\delta^a \geq \delta^b$, and as a result $\mu_{hh}^D \geq \mu_{hl}^D \geq \mu_{lh}^D \geq \mu_{ll}^D$. In the rest of the chapter, I calculate the total information rent obtained by OEMs A and B each selling a single product, and compare it to the information rent obtained by OEM D selling both products, to understand if and when OEM D is better off than OEMs A and B combined.

When OEMs A and B separately approach the supplier, she designs an optimal menu of contracts for each OEM. For each product i, the supplier's problem is a two-type version of (2.3a)-(2.3d):

$$\begin{aligned} \max_{Q_\theta^i, \pi_\theta^i} \quad & \sum_{\theta=h,l} p_\theta^i \left[r \left((1-q)Q_\theta^i - \int_{-\infty}^{Q_\theta^i - \mu_\theta^i} F^i(x)\mathrm{d}x \right) - \pi_\theta^i \right] \\ \text{s.t.} \quad & \pi_l^i \geq 0, \quad \pi_h^i \geq \pi_l^i + r \int_{Q_l^i - \mu_h^i}^{Q_l^i - \mu_l^i} F^i(x)\mathrm{d}x \\ & Q_h^i \geq Q_l^i. \end{aligned} \tag{2.6}$$

Ignoring the MC (2.6), the FOC solution to the above problem is

$$p_l^i(1-q) = F^i(Q_l^i - \mu_l^i) - p_h^i F^i(Q_l^i - \mu_h^i), \quad Q_h^i = \mu_h^i + (F^i)^{-1}(1-q).$$

It is trivial to test that this solution satisfies (2.6), thus this is the optimal solution. At this solution, the information rents are

$$\pi_l^i = 0, \quad \pi_h^i = r \int_{Q_l^i - \mu_h^i}^{Q_l^i - \mu_l^i} F^i(x)\mathrm{d}x.$$

On the other hand, for OEM D, the supplier's problem is a four-type version of (2.3a)-(2.3d). D's demand type $\theta^a\theta^b$ has prior probability $p^D_{\theta^a\theta^b} = p^a_{\theta^a} p^b_{\theta^b}$ and mean demand $\mu^D_{\theta^a\theta^b} = \mu^a_{\theta^a} + \mu^b_{\theta^b}$. Assume the forecast error $e^D = e^a + e^b$ has pdf f^D and cdf F^D. The supplier's problem is

$$\max_{Q_\theta^D, \pi_\theta^D} \quad \sum_{\theta=hh,hl,lh,ll} p_\theta^D \left[r \left((1-q)Q_\theta^D - \int_{-\infty}^{Q_\theta^D - \mu_\theta^D} F^D(x)\mathrm{d}x \right) - \pi_\theta^D \right] \tag{2.7a}$$

$$\text{s.t.} \quad \pi_{ll}^D \geq 0 \tag{2.7b}$$

$$\pi_{lh}^D \geq \pi_{ll}^D + r \int_{Q_{ll}^D - \mu_{lh}^D}^{Q_{ll}^D - \mu_{ll}^D} F^D(x)\mathrm{d}x \tag{2.7c}$$

$$\pi_{hl}^D \geq \pi_{lh}^D + r \int_{Q_{lh}^D - \mu_{hl}^D}^{Q_{lh}^D - \mu_{lh}^D} F^D(x)\mathrm{d}x \tag{2.7d}$$

$$\pi_{hh}^D \geq \pi_{hl}^D + r \int_{Q_{hl}^D - \mu_{hh}^D}^{Q_{hl}^D - \mu_{hl}^D} F^D(x)\mathrm{d}x \tag{2.7e}$$

$$Q_{hh}^D \geq Q_{hl}^D \geq Q_{lh}^D \geq Q_{ll}^D. \tag{2.7f}$$

Ignoring the MCs (2.7f), the FOC solution of (2.7a)-(2.7e) is

$$p_{ll}^D(1-q) = F^D(Q_{ll}^D - \mu_{ll}^D) - (p_{lh}^D + p_{hl}^D + p_{hh}^D)F^D(Q_{ll}^D - \mu_{lh}^D) \tag{2.8}$$

$$p_{lh}^D(1-q) = (p_{lh}^D + p_{hl}^D + p_{hh}^D)F^D(Q_{lh}^D - \mu_{lh}^D) - (p_{hl}^D + p_{hh}^D)F^D(Q_{lh}^D - \mu_{hl}^D) \tag{2.9}$$

$$p_{hl}^D(1-q) = (p_{hl}^D + p_{hh}^D)F^D(Q_{hl}^D - \mu_{hl}^D) - p_{hh}^D F^D(Q_{hl}^D - \mu_{hh}^D) \tag{2.10}$$

$$Q_{hh}^D = \mu_{hh}^D + (F^D)^{-1}(1-q). \tag{2.11}$$

At this solution, the information rents are

$$\pi_{ll}^D = 0$$

$$\pi_{lh}^D = r \int_{Q_{ll}^D - \mu_{lh}^D}^{Q_{ll}^D - \mu_{ll}^D} F^D(x)\mathrm{d}x \tag{2.12}$$

$$\pi_{hl}^D = \pi_{lh}^D + r \int_{Q_{lh}^D - \mu_{hl}^D}^{Q_{lh}^D - \mu_{lh}^D} F^D(x)\mathrm{d}x$$

$$\pi_{hh}^D = \pi_{hl}^D + r \int_{Q_{hl}^D - \mu_{hh}^D}^{Q_{hl}^D - \mu_{hl}^D} F^D(x)\mathrm{d}x. \tag{2.13}$$

If the solution satisfies (2.7f), then it is indeed an optimal solution. Otherwise, if the MCs (2.7f) are not satisfied, it is necessary to revise the solution.

Proposition 2.2. *The contract quantities determined by* (2.8)-(2.11) *can only violate* $Q_{ll}^D \leq Q_{lh}^D$ *or* $Q_{lh}^D \leq Q_{hl}^D$*, and not both. When the solution of* (2.8)-(2.11) *violates* $Q_{ll}^D \leq Q_{lh}^D$*, replace* (2.8)-(2.9) *with*

$$\begin{aligned}(p_{ll}^D + p_{lh}^D)(1-q) &= F^D(\underline{Q} - \mu_{ll}^D) - (p_{hl}^D + p_{hh}^D)F^D(\underline{Q} - \mu_{hl}^D), \\ Q_{ll}^D &= Q_{lh}^D = \underline{Q}\end{aligned} \tag{2.14}$$

and keep (2.10) *and* (2.11) *unchanged, and the revised FOCs determine the optimal solution. When the solution of* (2.8)-(2.11) *violates* $Q_{lh}^D \leq Q_{hl}^D$*, replace* (2.9)-(2.10) *with*

$$\begin{aligned}(p_{lh}^D + p_{hl}^D)(1-q) &= (p_{lh}^D + p_{hl}^D + p_{hh}^D)F^D(\overline{Q} - \mu_{lh}^D) - p_{hh}^D F^D(\overline{Q} - \mu_{hh}^D), \\ Q_{lh}^D &= Q_{hl}^D = \overline{Q}\end{aligned} \tag{2.15}$$

and keep (2.8) *and* (2.11) *unchanged, and the revised FOCs determine the optimal solution.*

In the above revised solutions, two types are offered the same contract (e.g., $Q_{ll}^D = Q_{lh}^D = \underline{Q}$); this situation is referred to as *bunching*. If the solution of (2.8)-(2.11) violates a monotonicity constraint, an optimal solution can be obtained by forcing equality of the violated constraint and ignoring the rest of the MCs (they will always be satisfied), then re-deriving the FOC. The revised FOC (2.14) is actually the sum of (2.8) and (2.9) with $\underline{Q} \doteq Q_{ll}^D = Q_{lh}^D$, and (2.15) is the sum of (2.9) and (2.10) with $\overline{Q} \doteq Q_{lh}^D = Q_{hl}^D$. Note that (2.14) and (2.15) can still be represented by Equation (2.4) (with $\lambda(\theta_j)$ replaced by

$$\underline{\lambda} \doteq \frac{p_{ll}^D + p_{lh}^D}{p_{ll}^D + p_{lh}^D + p_{hl}^D + p_{hh}^D}$$

corresponding to $\underline{Q}$, and $\lambda(\theta_j)$ replaced by

$$\overline{\lambda} \doteq \frac{p_{lh}^D + p_{hl}^D}{p_{lh}^D + p_{hl}^D + p_{hh}^D}$$

corresponding to $\overline{Q}$), thus properties derived from (2.4) will remain true in both cases of bunching.

In conclusion, to solve for information rents, I ignore the MCs and assume all PCs and ICs are binding, and then solve the FOCs. For individual OEMs A and B, the solution is optimal. For OEM D, I must test whether the monotonicity constraints are satisfied. If they are satisfied, the solution is optimal. Otherwise, I need to solve the revised FOCs to obtain the optimal solution.

2.3.2 Three Key Drivers of Information Rent

The above analysis reveals that the information rent for any demand type can be expressed as the sum of several incremental information rents, where each incremental information rent can be characterized by Equations (2.4) and (2.5). Therefore, understanding the properties of the incremental information rents as governed by (2.4) and (2.5) is crucial for comparing pooling versus non-pooling profits. The three lemmas below identify three key drivers of incremental information rents. In what follows, θ and θ' always refer to a type and its adjacent higher type (assuming θ is not the highest type).

Lemma 2.1 (Type Rareness)**.** *Incremental information rent $\pi_{\theta'} - \pi_\theta$ is increasing in $\lambda(\theta)$.*[1]

Recall that $\lambda(\theta)$ was defined as $\frac{\Pr(\Theta=\theta)}{\Pr(\Theta\geq\theta)}$. The intuition behind Lemma 2.1 is easily seen for an unpooled OEM, say A, having high-type mean demand. In this case, $\lambda(l)$ equals p_l^a, and Lemma 2.1 states that π_h^a is increasing in p_l^a, which means it

[1]When bunching occurs, $\lambda(\theta)$ should be understood as $\underline{\lambda}$ or $\overline{\lambda}$.

is decreasing in p_h^a. To understand this, note that the larger p_h^a is, the more the supplier anticipates that OEM A has high-type mean demand, and consequently the lower information rent the high-type OEM A will receive. In the extreme case, if the supplier knows OEM A's mean demand is of high type with certainty ($p_h^a = 1$), then the OEM A would get no information rent at all. Similarly, the rarer the high-type OEM A is, the more information rent he receives. For OEM D, who can have mean demands of four different types, $\lambda(\theta)$ reflects how likely the type θ is, compared only within the set of types θ and higher. (The lower types do not matter because I am only considering the *incremental* information rent.) In summary, Lemma 2.1 characterizes the impact of type rareness on incremental information rents.

I define $\delta = \mu_{\theta'} - \mu_\theta$ to be the "gap" between the type-θ' and type-θ mean demands (recall that θ is the lower type of the two). Lemma 2.2 shows that the incremental information rent depends on the mean demands only through their gap.

Lemma 2.2 (Gap between Types). *Incremental information rent $\pi_{\theta'} - \pi_\theta$ as a function of μ_θ and $\mu_{\theta'}$ is determined only by their difference: $\delta = \mu_{\theta'} - \mu_\theta$. Furthermore, $\pi_{\theta'} - \pi_\theta < r\delta$.*

The supplier's contracts provide the OEM with information rent in order to ensure that the OEM picks the contract meant for his demand type, and not the contract meant for a lower demand type. Lemma 2.2 shows that the incremental information rent does not depend on individual mean demands, but only on the gap between them. Additionally, the incremental information rent is bounded by the largest possible revenue difference of the two types, i.e., $r\delta$. This means that when there is almost no gap between the mean demands of two adjacent types ($\delta \to 0$), the two adjacent types are almost identical and so there is little incremental information rent.

Another important element in the model is the demand forecast error. To be able to quantify the effect of forecast error variability, I first define variability in a family of distributions.

Definition 2.1 (Rescaling)**.** Suppose F is the cdf of a zero-mean random variable. For any constant $\gamma > 0$, define cdf $F_{(\gamma)}$ as $F_{(\gamma)}(x) \doteq F(x/\gamma)$.

$F_{(\gamma)}$ is a γ rescaling of F. $\{F_{(\gamma)}, \gamma > 0\}$ could be seen as a family of random variables stemming from F. One could easily verify that the variance of $F_{(\gamma)}$ is γ^2 times that of F. Thus when $\gamma > 1$ the variance of $F_{(\gamma)}$ is greater than that of F.

Lemma 2.3 (Demand Variability)**.** *Suppose $\gamma > 1$, and replace F by $F_{(\gamma)}$ in* (2.4) *and* (2.5)*. Then the incremental information rent $\pi_{\theta'} - \pi_\theta$ is increasing in γ.*

Lemma 2.3 implies that the incremental information rent is actually *increasing* in demand variability, within the same family of forecast error distributions. Notice that information rent stems from the supplier's uncertainty about the OEM's demand. Therefore, it is understandable that the OEM's profit increases in his demand variability. Recall that OEM D pools the demand for products a and b, and therefore faces reduced variability for his component demand. The interesting question then is whether the reduced variability could actually lead to OEM D receiving lower information rent than OEMs A and B.

2.4 Pooling versus Non-Pooling

In this section I compare OEM D's information rent to the total information rents of OEMs A and B. I have seen that incremental information rents are determined by contract quantities (see (2.5)), which are in turn determined implicitly by FOC (2.4). For OEM D, whose demand for component c can have four different types, calculating information rent involves adding multiple layers of incremental information rents. The complex multi-layer structure of D's information rent makes it difficult to directly compare the information rents in closed form. Thus I use the insights of Lemmas 2.1-2.3 to facilitate comparisons. To ensure that the forecast error distribution for D's component c demand (i.e., the forecast error distribution for products a

and b combined) is tractable, I will assume that the forecast error e^i for individual product i is a normal random variable $N(0, \sigma^i)$.[2] (I assume that forecast errors are small relative to mean demand, so that the probability of having negative demand is negligible.)

I will now focus on when OEM D receives lower information rent than OEMs A and B. I will study this analytically for each demand type faced by OEM D. (Of course, a comparison before the OEMs learn their demand types could also be done utilizing my analysis for each demand type and the prior distribution of the types. Since this extra step only involves taking a weighted average of information rents and yields similar insights to those I present below, it is omitted for brevity.) I will also numerically illustrate what parameters lead to OEM D or OEMs A and B receiving higher information rent. My primary method of presenting numerical results will be showing the "*pooling*" and "*non-pooling*" regions — denoted with P and N respectively — where OEM D's information rent is *greater* or *smaller* than that of OEMs A and B combined. I plot these regions in a two-dimensional box of δ^b versus $\delta^a - \delta^b$ (recall that I assumed $\delta^a \geq \delta^b$), and in different plots I vary either σ^i or p^i_h. If not indicated otherwise, the default parameters are: $p^i_h = 0.5$, $q(= c/r) = 0.2$, and $\sigma^i = 1$, $i = a, b$. Due to Lemma 2.2, the values of μ^i_θ are irrelevant except for $\delta^i = \mu^i_h - \mu^i_l$, thus I do not assume any value for μ^i_θ (δ^i are indicated at the axes).

I first take a look at the comparison for demand type lh.

Theorem 2.1. *Assume demand type is lh.*

1. *With sufficiently small σ^a, OEM D receives lower information rent than A and B combined.*

2. *Suppose OEM D receives lower information rent than A and B combined. Then*

[2] I have numerically tested error distributions other than normal (e.g., uniform and triangular) and verified that the qualitative insights I will obtain in this section using normally distributed forecast errors remain valid.

as σ^a decreases, or p_h^a increases, OEM D receives even lower information rent and is still outperformed by A and B.

To understand my result, first notice that when OEM A has low demand and OEM B has high demand, the combined information rents of A and B is just the information rent of B, because the low-type OEM A earns no information rent. The theorem's first part establishes the existence of cases where OEM D earns less profit than OEMs A and B combined. Compared to OEM B, what OEM D loses by having pooled demand is his position in the type hierarchy: OEM D of type lh has the second-lowest demand type, so the supplier does not have to grant him significant information rent, while OEM B's type is h and therefore he will be granted significant information rent. (Recall that calculating information rent involves adding multiple layers of incremental information rents, so higher types earn more information rent.) On the other hand, OEM D has higher demand variability $\sqrt{(\sigma^a)^2 + (\sigma^b)^2}$ than OEM B's σ^b, and I know (from Lemma 2.3) that higher demand variability improves information rent. However, when σ^a is small, OEM D's gain from the increased demand variability is minimal because $\sqrt{(\sigma^a)^2 + (\sigma^b)^2}$ is not much higher than σ^b. Therefore, when σ^a is small enough, OEM D will make less profit compared to OEMs A and B combined. This is the intuition behind the first part of Theorem 2.1.

Now I turn to the second part of the theorem. Since OEM A always gets zero information rent, lowering σ^a or increasing p_h^a has no impact on the combined information rent of OEMs A and B. However, lowering σ^a or increasing p_h^a does reduce the information rent of OEM D, who has pooled demand for products a and b. The reason is that information rent depends on demand variability (Lemma 2.3) and type rareness (Lemma 2.1), and OEM D — through pooling demand for products a and b — is negatively affected by either lowering σ^a or increasing p_h^a.

Figure 2.1 illustrates my result. Comparing the panels from left to right, one can see that the region in which A and B outperform D (the non-pooled region, denoted

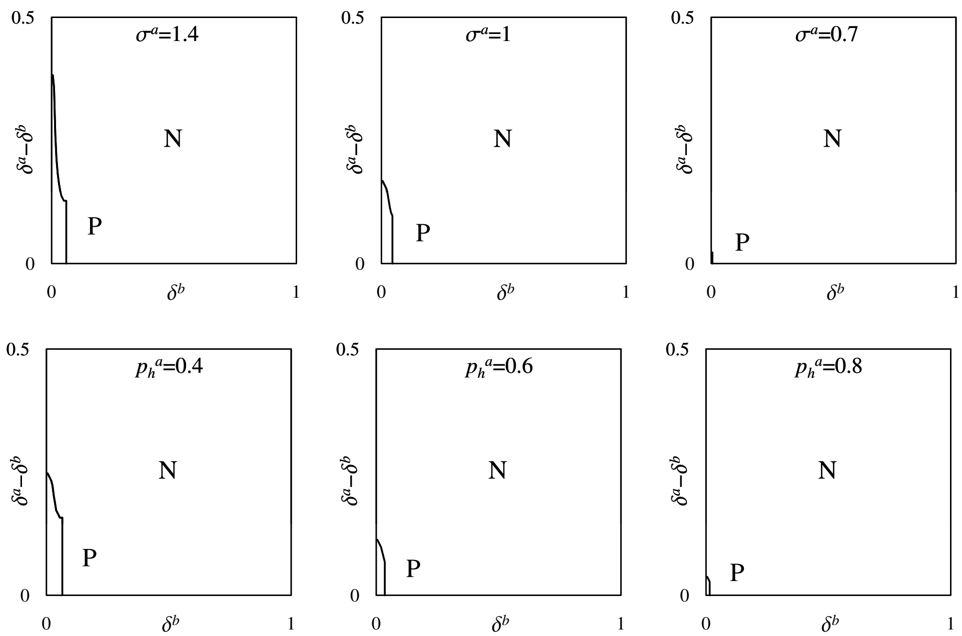

Figure 2.1: Pooling and non-pooling regions of demand type lh, varying σ^a and p_h^a

with N) grows as the demand variability for a, σ^a, decreases, and as the probability of product a having high-type demand, p_h^a, increases. Furthermore, Figure 2.1 shows that the regions in which non-pooling is optimal are fairly large and not limited to very low values of σ^a.

It is interesting to compare the intuition behind Theorem 2.1 to the traditional pooling intuition. In traditional inventory theory, pooling yields higher profits because of reduced demand variability. On the other hand, in my setting variability drives information rent and therefore reducing it can reduce the OEMs' profits. However, similar to the traditional setting, in my setting pooling can also be beneficial, although not for the usual reason. In my setting pooling can be beneficial because not all unpooled firms can make good use of their demand variability to generate information rents (e.g., firms with higher mean demands are capable of generating more information rents than firms with lower mean demands). Thus, pooling which makes use of the demand variability that would be "wasted" if used by a stand-alone OEM who is unable to earn much information rent on his own can be beneficial and result in higher profits.

Thus far I have shown that in the presence of a strategic supplier, an OEM D having demand type lh can earn less profit than OEMs A and B with unpooled demands. In fact this is not unique to demand type lh, but can also occur for type hl. The intuition behind these results is fairly similar, and is explained below.

Theorem 2.2. *Assume demand type is hl.*

1. *When σ^b is sufficiently small and δ^b is sufficiently close to δ^a, OEMs D receives lower information rent than A and B combined.*

2. *Suppose OEM D receives lower information rent than A and B combined. Then as σ^b decreases, OEM D receives even lower information rent and is still outperformed by A and B. Similarly, increasing p_h^b results in OEM D still receiving lower information rent than A and B combined, provided δ^a and δ^b are sufficiently close.*

The intuition behind this theorem is similar to Theorem 2.1's. The first part establishes the existence of cases where OEM D can be outperformed by OEMs A and B for demand type hl. With type hl, OEM A receives information rents whereas OEM B does not. For OEM D, the following tradeoff arises. On one hand, OEM A is of the highest type whereas OEM D is not, so OEM D is lower in the type hierarchy, which reduces his information rent. (In fact, when δ^b is sufficiently close to δ^a, OEM D of type hl is offered the same contract as type lh; therefore OEM D's loss of information rent due to lower type hierarchy is significant.) On the other hand, OEM D can make use of the demand variability of product b, which OEM B could not make use of. (Again, recall that higher demand variability results in more information rent by Lemma 2.3.) However, when product b's demand variability is low, OEM D's gain in information rent due to product b's demand variability is negligible, and as a result, OEM D ends up with lower profit than OEMs A and B combined. For the second part, notice that when OEM B has low-type demand and thus earns no information rent, lowering σ^b or increasing p_h^b has no impact on the combined information rent of A and B. Therefore, the result follows by noting that OEM D's information rent decreases when σ^b is decreased (by Lemma 2.3) or p_h^b is increased (by Lemma 2.1).

Figure 2.2 illustrates this result. Comparing the panels from left to right, the non-pooling region in which OEMs A and B outperform OEM D grows as the demand variability for b, σ^b, decreases, or the probability that product b has high demand, p_h^b, increases. Figure 2.2 again shows that the regions in which non-pooling is optimal are fairly large and not limited to very low values of σ^b. Notice that for any given σ^b and δ^b, as $\delta^a - \delta^b$ approaches zero, non-pooling is more likely to be preferred. However, the first two panels reveal that even when $\delta^a - \delta^b$ is zero, there are cases where OEM D earns higher profits. This occurs when σ^b is high, thus revealing that gaining sufficiently high demand variability from product b can indeed compensate for OEM D's decrease in type hierarchy when pooling. Once again, notice that pooling is beneficial when doing so utilizes *more* demand variability for generating information rent.

Concluding the results for types lh and hl, I make the following observation. Facing a powerful strategic supplier, the only source of profit for the OEMs is information rent. When not pooling purchases, only the high-type OEM gets information rent. When pooling purchases, OEM D has a hierarchy disadvantage compared with a high-type OEM A or B since D does not have the highest type. This can potentially be compensated by the fact that OEM D can make use of demand variability for both products, while for types lh or hl, the low-type OEM A or B cannot. Once again, I would like to draw the reader's attention to my finding that the presence of a strategic supplier can reverse the common wisdom about the benefit of variability reduction through pooling. In my setting, reduced variability harms the OEM, but pooling that utilizes more demand variability for generating information rent can actually be

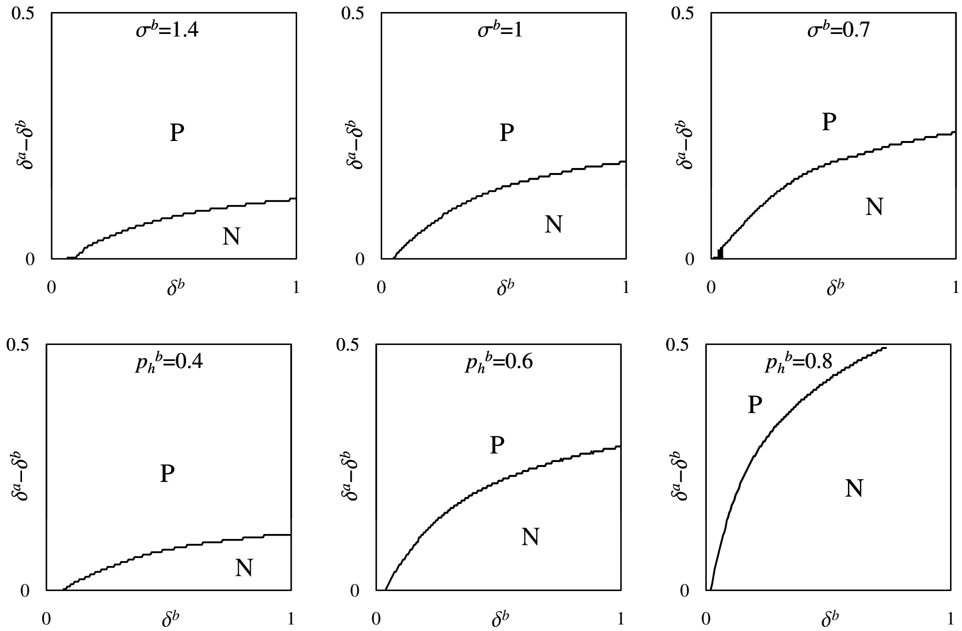

Figure 2.2: Pooling and non-pooling regions of demand type hl, varying σ^b and p_h^b

beneficial. This is counter to the received intuition about pooling.

But what happens when both products a and b have high mean demands? In this case, OEMs A, B, and D all have the highest type, so the tradeoff of hierarchy disadvantage versus increased variability that I set up above does not apply. Interestingly, situations where pooling results in lower profits can be easily found in this case as well.

Theorem 2.3. *Suppose demands for products a and b are symmetric: $\delta^a = \delta^b = \delta$, $p_h^a = p_h^b = p_h$, $\sigma^a = \sigma^b = \sigma$, $p_h > 0.15$, and σ is sufficiently large. If OEM D of type hh receives lower information rent than A and B combined, then as p_h increases, OEM D receives even lower information rent and is still outperformed by A and B.*

With demand type hh, both products' demand variability will generate information rent when OEMs A and B purchase separately, but for OEM D the variability is reduced upon pooling due to statistical economies of scale. This gives the OEMs A and B a variability advantage over OEM D, which gets greater as σ gets larger. On the other hand, OEM D's type-hh demand has the highest rank of four possible types

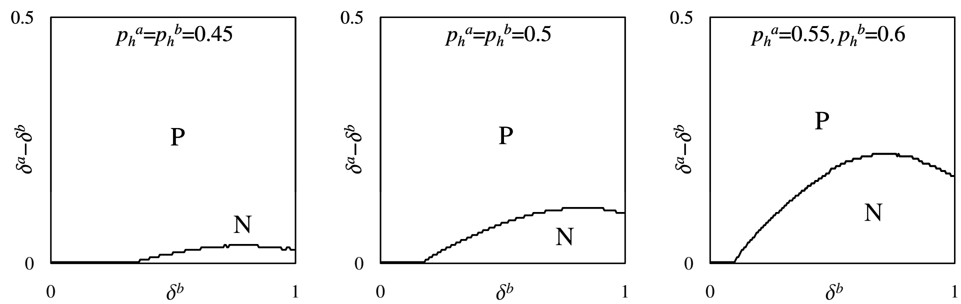

Figure 2.3: Pooling and non-pooling regions of demand type hh, for varying levels of p_h^i (in this figure $\sigma^i = 3$)

in the demand type hierarchy, while A and B's type-h demands are only the higher of two possible demand types. As a result, D gains an advantage in type hierarchy. This however is only a significant advantage when p_h is small (e.g., $p_h = 0.05$) because then p_h^2 would be *very* small and, by Lemma 2.1, the rareness of type-hh demand leads to a large information rent for OEM D of type hh. As p_h increases, the hh type becomes less rare, and this coupled with the decrease in variability from pooling ends up making OEM D worse off compared to OEMs A and B. Figure 2.3 illustrates this trend. Reading the panels from left to right, as p_h increases, the non-pooling region in which OEMs A and B outperform D grows. (Notice that although the sufficient conditions of Theorem 2.3 require $p_h^a = p_h^b$, it is easy to generate examples where OEMs A and B dominate OEM D when $p_h^a = p_h^b$ is violated, as can be seen on the last panel of Figure 2.3.)

Although in the above results I focused on cases when OEMs A and B outperform OEM D, the intuition I built can also be used to predict when the opposite happens. For example, consider the case where product b's demand has small gap but huge variability. On his own, OEM B would earn little information rent (e.g., *zero* information rent if the gap is zero, per Lemma 2.2). In this case OEM B's demand variability is "wasted". However, OEM D can combine the gap from product a with the variability from product b to generate significant information rent. This is captured in the theorem below.

Theorem 2.4. *Fixing all other parameters, as δ^b becomes sufficiently small, OEM D of types hl and hh receive higher information rent than A and B combined.*

Notice that Theorem 2.4 does not consider type lh. In this case, OEM A receives no information rent due to having low mean demand, and when δ^b is small, OEMs B and D both receive only negligible information rents. Therefore the profit comparisons for type lh are trivial. Theorem 2.4 is clearly demonstrated in Figures 2.2 and 2.3 by the fact that the pooling regions occur near the left edge (where $\delta^b \approx 0$). I again point out that, while Theorem 2.4 describes a case of my model where pooling is indeed beneficial, the reason is completely different from the standard pooling logic: OEM D receives higher profit because compared to OEM B, he can better utilize product b's demand variability to generate information rent.

So far I have focused on the OEMs' profits. Besides profits, the existing pooling literature has also studied how inventories (purchase quantities) change upon pooling. Below, I briefly discuss how quantities are affected by pooling in my model.

In a traditional pooling model (where the good is purchased from a commodity market at a fixed price c), how inventories change upon pooling is primarily driven by the *critical ratio* $(r-c)/r$ ($= 1 - q$ in my model), where r is the buyer's revenue from selling one unit of the good. When the demand pdf is symmetric about its mean (as is in the seminal paper Eppen (1979)), total inventory decreases upon pooling if $(r-c)/r > 0.5$, and increases if $(r-c)/r < 0.5$. The intuition is very simple: When the critical ratio is high (low), lost sales are more (less) expensive than leftover inventory, thus it is optimal to overstock (understock). Pooling reduces demand uncertainty, so the level of overstock (understock) to achieve the critical ratio is also reduced. This translates into decreased (increased) inventory when $(r-c)/r > (<) 0.5$. Note that the assumption of symmetric demand pdf is crucial; for example, Yang and Schrage (2009) shows that with a right-skewed demand pdf, "inventory anomaly" can occur, namely inventory can increase upon pooling even when $(r-c)/r > 0.5$. Interestingly, even

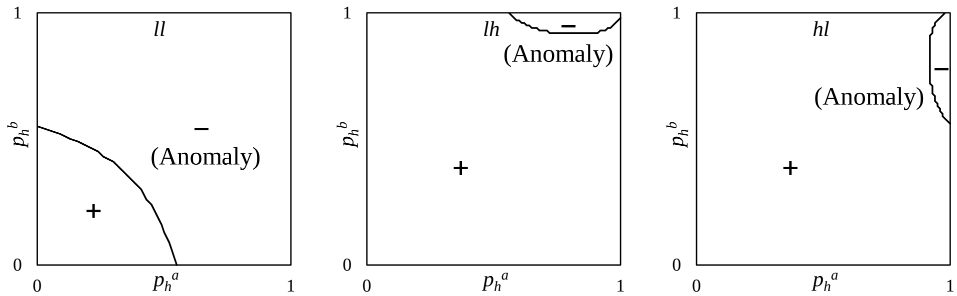

Figure 2.4: $Q^D_{\theta^a\theta^b} - Q^a_{\theta^a} - Q^b_{\theta^b}$ for types ll, lh and hl, and $q = 0.6$

when assuming symmetric (normal) demand pdfs in my model, “inventory anomaly” can still occur, thanks to the presence of information asymmetry. The next figure provides examples. In these examples, I set $q = 0.6$, $\sigma^a = \sigma^b = 1$, $\delta^a = \delta^b = 1$, and plot $Q^D_{\theta^a\theta^b} - Q^a_{\theta^a} - Q^b_{\theta^b}$ for varying p^a_h and p^b_h, and types ll, lh and hl. (There is no anomaly for type hh, as is explained in the next theorem.) In the traditional pooling model, setting $q = 0.6$ (corresponding to a critical ratio of 0.4) would result in positive values of $Q^D_{\theta^a\theta^b} - Q^a_{\theta^a} - Q^b_{\theta^b}$ (increased inventory upon pooling), but this is not always the case in my model.

The reason for the “anomaly” in my model is information asymmetry. It is well-known that information asymmetry in principal-agent models results in *downward distortion*, namely all types except the highest one purchase lower than first-best (newsvendor) quantities. Furthermore, one can show (see Lemma 2.1’s proof in Appendix A) that the rarer the type, the greater the quantity distortion for that type. Pooling changes the type distribution and affects the level of downward distortion, which can lead to the anomaly. For example, in the first panel (type ll), when p^a_l and p^b_l are small (say 0.3), OEM D of type ll is much rarer than OEMs A and B of type l ($p^D_{ll} = 0.09$). Consequently, OEM D experiences much stronger downward distortion in purchase quantity than OEMs A and B, leading to the anomaly. Similar observations can be made in the other two figures.

On the other hand, when p^a_l and p^b_l are sufficiently large, downward distortion in purchase quantity is weak. I also know that the highest type never experiences downward distortion. In these cases the purchase quantities behave as in the traditional pooling literature. The theorem below states conditions for the purchase quantities to behave as in traditional pooling models despite the presence of the strategic supplier and information asymmetry when the critical ratio $(r-c)/r > 0.5$. (The result for the case where $(r-c)/r < 0.5$ is similar and omitted for brevity).

Theorem 2.5. *Assume the critical ratio $(r-c)/r > 0.5$. Then for any OEM A's type θ^a and OEM B's type θ^b, when p_l^a and p_l^b are sufficiently close to 1, $Q^a_{\theta^a}+Q^b_{\theta^b} > Q^D_{\theta^a\theta^b}$. In particular, for demand type hh, $Q_h^a + Q_h^b > Q_{hh}^D$ if $(r-c)/r > 0.5$.*

The above discussion shows that in my model the traditional pooling intuition and information asymmetry both influence the behavior of purchase quantities upon pooling. When downward distortion is weak, the purchase quantities behave much like in a traditional model. On the other hand, information asymmetry and downward distortion can lead to "inventory anomaly", namely inventory moves in the opposite direction of what traditional pooling intuition suggests. This means I have identified information asymmetry as another possible cause of inventory anomaly, besides skewed demand distributions as previously identified in Yang and Schrage (2009).

2.5 Concluding Discussion

In this chapter, I consider whether pooling purchases for a component used in multiple products with uncertain demands results in increased profits for the buyer. Received intuition on pooling indicates that pooling demands benefits the buyer by reducing demand variability. However, my setting and the traditional pooling literature have a fundamental difference: I consider a strategic supplier who tries to maximize her profit by strategically pricing supplies based on what she anticipates about her customers' demands, whereas the traditional literature usually does not model a strategic supplier and the buyers effectively purchase from a commodity market at a fixed price. In my setting, pooling that significantly reduces demand variability can actually result in reduced profits for the buyers because of reduced information rent, which is not considered in the traditional pooling literature. I additionally show that the analysis is complex and subtle because the comparison does not only depend on the variability but also the "rank order" of the buyers' demands. I also show that information asymmetry and the resulting downward distortion can be a cause of "inventory anomaly", which is different from the cause previously identified in the literature.

The purpose of this chapter is to show that received wisdom on pooling can fail when the implicit assumption of an exogenous and fixed price is violated. I show this using a standard principal-agent model. Therefore, my current knowledge is limited to perfect competition amongst suppliers (e.g., commodity purchases) where the standard pooling logic applies, and the situation considered in this chapter where a sole powerful supplier is the only source of purchase for the buyers. However, there exist many industries in between where there may be multiple suppliers (e.g., duopolies) with partial power, or sole suppliers that are not able to make take-it-or-leave-it offers. Further research should explore when pooling demand will or will not be beneficial in these environments; because pooling is such a canonical concept in operations management, extending my knowledge to cover these cases as well would be an interesting avenue for future work. Nevertheless, I expect that my main insights would still apply: To the extent that the buyer's profit relies on superior information about his demand, demand pooling can be unattractive for reducing this informational advantage.

This chapter studies the effect of pooling on the OEMs' overall sales revenues and procurement costs when meeting demand for a product, without considering the higher-level decision of whether the OEMs should or should not pool their purchases. Whether the OEMs should pool their purchases is a very important question, however answering it requires the consideration of additional factors, such as the investment cost of instituting an inventory pooling system, the additional shipping cost when inventory is physically pooled (in a centralized storage facility), or the additional transshipping cost when inventory is virtually pooled (e.g., each OEM has his own storage facility, but an IT system enables inventory information sharing and transshipment across the OEMs' facilities can be made when necessary), etc. In this chapter I focus on the more fundamental question of whether pooling purchases always brings benefits to the OEMs, and leave the higher-level question of whether the OEMs should or should not pool purchases to future work. However, obviously, my analysis of sales revenues and procurement costs in the presence of pooling can aid managers seeking to make a decision about whether or not to institute pooling infrastructure in their organizations, and builds the foundation for further research on this topic.

CHAPTER III

Simple Auctions for Supply Contracts

3.1 Introduction

In this chapter, I consider a newsvendor-like problem where one buyer (a newsvendor) faces n candidate suppliers who hold private information about their production costs. The buyer needs to purchase goods from one or more suppliers before she can use the goods to meet a random demand. Unsatisfied demand is lost, and unsold inventory is discarded. If the buyer always purchases the good at an exogenous and fixed price (for example, from a commodity market), she only needs to determine the optimal purchase quantity, and the problem becomes the classical newsvendor problem. However, because each supplier's production cost is his private information, the buyer's purchase price of the good is not fixed, but will be determined through interaction with the suppliers. The buyer needs to design a sourcing mechanism to determine the purchase quantity and price, and which supplier(s) to purchase from.

The ex ante symmetric and linear cost version of this problem, namely when all suppliers' costs are linear in production quantity and their unit costs are identically distributed, has been the subject of several research papers, most recently and notably Chen (2007). In his paper Chen shows that the following *supply contract auction* is an optimal sourcing mechanism for the buyer: The buyer announces a supply contract that specifies her payments for all purchase quantities, and auctions the contract among the suppliers. The winning supplier pays an upfront fee to the buyer, then chooses to deliver any quantity of his choosing to the buyer and collects his payment according to the contract. Chen notes that this mechanism fits well with slotting allowance practice, where a supplier pays the retailer an upfront fee, then determines how much inventory to ship and display on the shelf. Chen also cites another optimal mechanism called the *quantity auction* proposed in Dasgupta and Spulber (1990), where the buyer announces a supply contract and makes the suppliers bid quantities they are willing to deliver under this contract in the "sealed-bid high-quantity" format. Chen notes that the supply contract auction he proposes has two advantages over the quantity auction. The first advantage is that the quantity auction must be carried out in the sealed-bid high-quantity format, while thanks to the revenue equivalence theorem, the supply contract auction can use several commonly seen formats, where suppliers bid prices not quantities, which is more akin to practice. The second advantage is that the contract used in the quantity auction depends on the number of participating suppliers while the contract used in the supply contract auction does not, so the buyer can run the supply contract auction without knowing the exact number of participating suppliers. In his paper Chen assumes that all the suppliers are ex ante symmetric, and have linear production costs.

While both the supply contract auction and the quantity auction are theoretically optimal, they are not the most familiar and simple mechanisms for the suppliers. Also, while slotting allowances are common in practice, the supply contract auction mechanism described by Chen is much less observed in practice, especially outside retail. In comparison, for example, the commonly used open-descending auction for a fixed-quantity contract has a very simple and easily understandable structure, and is widely observed in a variety of sourcing situations (Jap 2007). One reason for its popularity is that the requirement on the participating suppliers' decision-making sophistication is minimal — a supplier only needs to compare the current auction price with his own cost. In fact, Chen points out that such an auction is simpler than the supply contract auction and the quantity auction (Chen (2007), p.1563). However, the open-descending auction for a fixed-quantity contract is not an optimal mechanism for the buyer. Therefore, it would be ideal if the buyer can use an optimal mechanism that is as similar to an open-descending auction for a fixed-quantity contract as possible, and adds little complexity to the suppliers' decision-making process.

In this chapter, I show that a variation of the standard open-descending auction is an optimal mechanism for the buyer; I name it the *modified open-descending auction*. The buyer announces that the mechanism will consist of two stages. In Stage 1, the buyer will run a standard open-descending auction for an initial fixed quantity. In Stage 2, the winning supplier from Stage 1 will receive one additional offer from the buyer to supply more units at unit prices no higher than the auction's ending price. The timeline of the mechanism is as follows:

1. All suppliers participate in an open-descending auction for a fixed-quantity contract and one supplier emerges as the winner.

2. The buyer informs the winning supplier how much she is willing to pay for each additional unit the supplier chooses to deliver.

3. The winning supplier delivers the guaranteed initial quantity, plus any additional quantity of his choosing beyond the initial quantity.

4. The buyer's demand is realized, unsatisfied demand is lost, and unsold inventory is discarded.

What distinguishes the modified open-descending auction from the supply contract auction and the quantity auction is its familiarity to the suppliers. Stage 1 of my mechanism is just a standard open-descending auction. Stage 2 of my mechanism would also be familiar to the suppliers, because it is common in many industrial settings that suppliers will bid for an initial quantity and become a preferred supplier if chosen, with the expectation that they are to lower their prices in the future if they want more business. For example, in a large conglomerate that I worked with, the sourcing staff are expected to achieve a target price deflation over time for the goods they source from the same supplier. The other major advantage of the modified open-descending auction is that it requires a much less supplier sophistication. As I will later show in §3.2, in both stages of my mechanism, the only computation that a supplier needs to perform is comparing his own cost with another number. Furthermore, only one supplier will ever enter Stage 2. By contrast, in the quantity auction and the supply contract auction, all suppliers must perform calculations using the potentially complicated payment schedules provided by the buyer before determining their best bidding strategies. These benefits of my mechanism are important, because in practice suppliers are much more willing to participate in a mechanism that they find more familiar and simpler.

Finally, with minor modifications, my proposed mechanism remains optimal when the suppliers' production costs are concave in quantity and ex ante asymmetric, while Chen's paper assumes ex ante symmetric and linear production costs. Therefore, the modified open-descending auction is more practically implementable, and yet less restrictive.

3.2 Base Model: Linear Symmetric Costs

To be consistent with Chen (2007), I first assume linear symmetric production costs in my base model. Suppose a buyer needs to purchase a good from one or more of n candidate suppliers to satisfy an uncertain future demand. Unsatisfied demand is lost, and unsold inventory is discarded. Assume the buyer's *expected* revenue $R(Q)$ from stocking Q units is non-negative, increasing, and concave in Q.[1]

[1]One example of the expected revenue function that satisfies these requirements is the classical newsvendor's expected revenue function $R(Q) = pE_D[\min\{Q, D\}]$, where p is the unit retail price, and D is the uncertain demand. In this chapter I am not restricted to the classical newsvendor's

Define $r(k) \doteq R(k) - R(k-1)$ to be the buyer's expected marginal revenue from the k^{th} unit, which is non-increasing. Assume supplier i's unit production cost c_i is his private information, but it is common knowledge that each c_i is an independent draw from a random variable $C \in [\underline{c}, \bar{c}]$ with cdf F and pdf f. These are the exact same assumptions made in Chen (2007). My goal is to design a simple and practically implementable contracting mechanism that maximizes the buyer's expected profit.

To focus on the modified open-descending auction, I relegate the optimal mechanism design analysis to the Appendix, and jump directly to describing the mechanism. Define *virtual cost function* $\psi(c) \doteq c + \frac{F(c)}{f(c)}$ and assume ψ is increasing in c.[2] Define *initial quantity*

$$Q_0 \doteq \max\{Q \in \mathbb{Z} | r(Q) \geq \psi(\bar{c})\}$$

and assume $r(1) \geq \psi(\bar{c})$, so that $Q_0 \geq 1$.[3]

Definition 3.1 (Modified Open-Descending Auction)**.** The buyer first announces the following rules to all candidate suppliers:

- The mechanism has two stages. Stage 1 is an open-descending auction for a contract to supply Q_0 units. The auction price p will start at an initial price set by the buyer and thereafter will always equal the current lowest bid. Any supplier can always bid a price lower than the current auction price, which updates the auction price. The auction ends when no supplier is willing to bid any lower, and the final lowest bidder is the winner. Suppose the auction ends at price p_0. The auction's winner has the obligation to deliver Q_0 units to the buyer and the buyer will pay p_0 for each of the Q_0 units.

setup, but allow any general expected revenue functions that satisfy these requirements.

[2]This condition is often assumed in the economics literature, and is satisfied by many commonly seen distributions; see Footnote 2 on p.1564 of Chen (2007).

[3]This condition, used here and in a similar way in §3.3, guarantees that the buyer will always purchase. It only serves to simplify the analysis and is non-essential. When this condition is violated, I can redefine $Q_0 \equiv 1$ and run the modified open-descending auction similarly but with a starting price $\psi^{-1}(r(1))$; there is a possibility that all suppliers will drop out immediately after the auction starts, in which case the buyer chooses not to purchase at all.

- In Stage 2, the buyer will make one additional offer to the winning supplier from Stage 1 that specifies the prices the buyer will pay for each additional unit (beyond Q_0) that the supplier chooses to deliver. The exact prices will be specified in Stage 2, however no price for any additional unit will be higher than p_0. The winning supplier can choose to deliver any additional quantity (including zero).

After announcing the rules, the buyer executes the mechanism as follows.

Stage 1 The buyer runs an open-descending auction for a contract to supply Q_0 units, starting at price $\bar{c}$.

Stage 2 Suppose Stage 1's auction concludes at price p_0 and supplier i is the winner. The buyer then offers to pay $\min\{p_0, \psi^{-1}(r(k+Q_0))\}$ for the k^{th} additional unit beyond Q_0 that supplier i chooses to deliver, where $k \leq r^{-1}(\psi(\underline{c})) - Q_0$.

Put simply, the modified open-descending auction takes the following format: The buyer uses an open-descending auction to allocate a fixed-quantity contract. After the auction, the buyer makes one additional offer to the winning supplier to purchase more units at discounted prices.

Theorem 3.1. *The modified open-descending auction is an optimal mechanism for the buyer. In Stage 1, each supplier i's dominant bidding strategy is to keep lowering his bid until auction price p reaches c_i or the auction ends (when other suppliers are not willing to bid lower), whichever happens first. In Stage 2, supplier i (given he is the winner in Stage 1) will deliver an additional $q_0 \doteq \max\{q \in \mathbb{Z} | r(q + Q_0) > \psi(c_i)\}$ units to the buyer.*

Having described the modified open-descending auction and the suppliers' dominant strategies, below I highlight a few features of the mechanism to show why I believe it achieves my goal of designing a simple and practically implementable optimal mechanism.

First, as stated in §3.1, the mechanism is familiar to suppliers. Moreover, Theorem 3.1 shows that each supplier i's dominant bidding strategy in Stage 1's auction is identical to that in a standard open-descending auction (i.e., keep bidding lower until the auction price p reaches his cost of fulfilling the contract c_i). Each supplier i's computation is trivial: He only needs to compare his unit cost c_i with the current auction price p. Finally, in Stage 2 the winning supplier will receive an additional offer that takes the form of a (weakly) decreasing sequence of payments for each additional unit delivered. Therefore, to determine his optimal additional delivery, supplier i again only needs to compare his unit cost c_i with the decreasing sequence of payments to find out the last unit that he can earn a profit on, and deliver this additional quantity.

I want to point out that in judging whether a mechanism is "simple and practically implementable", I primarily consider how easy it is for the *suppliers* to understand its procedures, and to determine their best bidding strategies in the mechanism. A mechanism that is very simple from the suppliers' perspective may require computations from the buyer that are not as simple. For example, the modified open-descending auction appears to the suppliers as just an auction for a fixed-quantity contract starting from a reserve price, plus an additional offer in the form of a decreasing sequence of payments, and the suppliers' decision-making in this mechanism is almost trivial. However, from the buyer's perspective, computing the initial quantity, the reserve price, and the additional price offers all require some sophistication. Nevertheless, while a savvy buyer can make all efforts to design an optimal mechanism, she cannot necessarily force the suppliers to participate in the mechanism. Therefore, one very important practical consideration in designing a mechanism should be that it is simple and familiar to the suppliers, because suppliers may refuse to participate in a mechanism they find unfamiliar or confusing.

In Stage 1 of my mechanism, the buyer uses an open-descending auction for a fixed-quantity contract to find the lowest-cost supplier. It is well known that when auctioning a fixed-quantity contract, auction formats such as sealed first-price, Vickrey (sealed second-price) and reverse Dutch (open-ascending)[4] auctions all result in the same expected profit for the buyer as an open-descending auction (the *revenue equivalence theorem*). Thus it is natural to ask if in my mechanism the buyer could use any of these auction formats in Stage 1, i.e., if the term "open-descending" in Definition 3.1 could be replaced by "sealed first-price", "Vickrey", or "reverse Dutch". As it turns out, the revenue equivalence theorem no longer holds true in my setting: One *cannot* replace the open-descending auction in Stage 1 of my mechanism by any of these commonly used auction formats without compromising the mechanism's optimality.

Proposition 3.1. *In my mechanism, the open-descending auction in Stage 1 cannot be replaced by any auction format that results in full revelation of the lowest-cost supplier's private cost information to the buyer, without compromising the mechanism's optimality. In particular, it cannot be replaced by a sealed first-price, Vickrey, or reverse Dutch auction without compromising the mechanism's optimality.*

By optimal mechanism design theory, in any optimal mechanism the buyer *must* purchase the total quantity that makes marginal revenue equal the lowest marginal *virtual* cost. However, if in Stage 1 the buyer learns the winning supplier's cost information, then in Stage 2 she will offer to pay just above the supplier's cost for additional units, which will result in the total purchase quantity being greater than that of an optimal mechanism. Therefore, such a mechanism can never be optimal.

The fact that a sealed first-price, Vickrey, or reverse Dutch auction in Stage 1 will fully reveal the winning supplier's cost information to the buyer, however, is

[4]A reverse Dutch (open-ascending) auction for supplying Q_0 units runs as follows. The auction price p slowly increases from zero over time. When a supplier first indicates that he accepts the current price p_0, the auction ends and the supplier will deliver Q_0 units and receive p_0 for each of the Q_0 units from the buyer.

not obvious: The buyer only observes the suppliers' bids, which may or may not equal their actual costs. Nevertheless, it turns out that the buyer can infer the winning supplier's true cost from his bid and his equilibrium bidding strategy (see Proposition 3.1's proof for details). Therefore, replacing the open-descending auction in Stage 1 by a sealed first-price, Vickrey or reverse Dutch auction will cause sub-optimality. By contrast, with the open-descending auction the buyer can only learn an *upper bound* on the winning supplier's cost, not the exact cost.

As stated in Definition 3.1, in announcing her mechanism the buyer does not specify what the offer will be in Stage 2, and her freedom of offering any prices of her choosing (no higher than p_0) in Stage 2 will lead to sub-optimal purchase quantities when sealed first-price, Vickrey or reverse Dutch auctions are used in Stage 1. Of course, this problem could potentially be avoided if before Stage 1's auction the buyer removes her freedom by announcing the complete pricing recipe of how Stage 2's price offers will be computed as a function of the auction's outcome. However, announcing a pricing recipe — a multi-dimensional function that maps every possible auction outcome to a list of prices — is against my intention to design an optimal mechanism that is as simple and familiar as possible. Therefore, the attempt to replace the open-descending auction in my mechanism with another commonly used auction format will lead to either a sub-optimal mechanism, or a much more complicated mechanism. This speaks to the unique role that the open-descending auction plays in my simple optimal mechanism.

3.3 Extensions

In this section I extend the mechanism to allow concave and ex ante asymmetric production costs, neither of which is addressed in Chen (2007).

3.3.1 Ex Ante Symmetric Concave Costs

Production costs which are concave in quantity are commonly used to capture economies of scale widely observed in many industrial settings. Common causes for economies of scale include dilution of a large fixed setup cost (e.g., an initial machinery investment), bulk purchase discounts for materials, and learning curve effect. The modified open-descending auction can be easily extended to allow a class of concave production costs.

Assume supplier i's total cost for producing Q units is $H(Q)c_i$, where $H(Q)$ is an increasing and concave function common to all suppliers, satisfying $H(0) = 0$ and $H(\infty) = \infty$. I call c_i the *base cost* of supplier i. This cost model captures the case where the economies of scale is common in the industry, and suppliers differ from each other only in the base cost. This would be a reasonable assumption, for example, in a labor-intensive industry where the technology and learning curve effect of all suppliers are similar, and labor cost drives the production cost. In this case, base cost c_i can represent the hourly labor rate at the supplier. Define $h(k) \doteq H(k) - H(k-1)$, so $h(k)c_i$ is supplier i's marginal production cost for the k^{th} unit. Obviously $h(k)$ is non-increasing in k, and $H(Q) = \sum_{k=1}^{Q} h(k)$. Further assume $r(Q)/h(Q)$ is decreasing in Q and approaches 0 as $Q \to \infty$. Intuitively, this means the marginal revenue is diminishing faster than the marginal production cost. This assumption prevents the optimal production from becoming infinite. Base margin c_i is supplier i's private information, and is an independent draw from a random variable $C \in [\underline{c}, \bar{c}]$ with cdf F and pdf f. Assume increasing virtual cost function $\psi(c) \doteq c + \frac{F(c)}{f(c)}$. Define initial quantity

$$Q_0 \doteq \max\{Q \in \mathbb{Z} | r(Q)/h(Q) \geq \psi(\bar{c})\}$$

and assume $r(1)/h(1) \geq \psi(\bar{c})$, so that $Q_0 \geq 1$.

Definition 3.2 (Modified Open-Descending Auction for Symmetric Concave Costs)**.**

The buyer first announces the following rules to all candidate suppliers:

- The mechanism has two stages. Stage 1 is an open-descending auction for a contract to supply Q_0 units. The auction price P will start at an initial price set by the buyer and thereafter will always equal the current lowest bid. Any supplier can always bid a price lower than the current auction price, which updates the auction price. The auction ends when no supplier is willing to bid any lower, and the final lowest bidder is the winner. Suppose the auction ends at price P_0. The auction's winner has the obligation to deliver Q_0 units to the buyer and the buyer will pay P_0 for the Q_0 units.

- In Stage 2, the buyer will make one additional offer to the winning supplier from Stage 1 that specifies the prices the buyer will pay for each additional unit (beyond Q_0) that the supplier chooses to deliver. The exact prices will be specified in Stage 2, however the price for the k^{th} additional unit will be no higher than $h(k + Q_0)P_0/H(Q_0)$. The winning supplier can choose to deliver any additional quantity (including zero).

After announcing the rules, the buyer executes the mechanism as follows.

Stage 1 The buyer runs an open-descending auction for a contract to supply Q_0 units, starting at price $H(Q_0)\bar{c}$.

Stage 2 Suppose Stage 1's auction concludes at price P_0 and supplier i is the winner. The buyer then offers to pay $h(k + Q_0)\min\{P_0/H(Q_0), \psi^{-1}(r(k + Q_0)/h(k + Q_0))\}$ for the k^{th} additional unit beyond Q_0 that supplier i chooses to deliver, where $k \leq \max\{l|r(l + Q_0)/h(l + Q_0) > \psi(\underline{c})\}$.

Theorem 3.2. *The modified open-descending auction for symmetric concave costs is an optimal mechanism for the buyer. In Stage 1, each supplier i's dominant bidding strategy is to keep lowering his bid until auction price P reaches $H(Q_0)c_i$ or the* *auction ends (when other suppliers are not willing to bid lower), whichever happens first. In Stage 2, supplier i (given he is the winner) will deliver an additional* $q_0 \doteq \max\{q \in \mathbb{Z} | r(q+Q_0)/h(q+Q_0) > \psi(c_i)\}$ *units to the buyer.*

One can easily see that the symmetric linear costs problem (Definition 3.1 and Theorem 3.1) is a special case of the symmetric concave costs problem (Definition 3.2 and Theorem 3.2), with $h(k)$ replaced by 1 (equivalently, with $H(Q)$ replaced by Q). With symmetric concave costs, the price cap on the additional offer that the buyer announces before the auction is no longer a constant p_0 as in the base model, but is a decreasing curve $h(k+Q_0)P_0/H(Q_0)$ for the k^{th} additional unit the winning supplier chooses to deliver. The shape of the price cap curve matches the suppliers' diminishing marginal cost curve $h(\cdot)$.

Since I have described how to extend the modified open-descending auction to allow concave production costs, a natural question would be whether the mechanism can be extended to also allow *convex* production costs. Unfortunately, the answer is negative. The reason my mechanism works in linear and concave cost settings is that if a supplier does not win the auction in Stage 1, he will have no business, thus he has to bid down to his true cost in the auction. In a convex cost setting, it can be optimal for the buyer to multi-source to avoid climbing too high on a single supplier's increasing marginal cost curve, and the buyer does not know a priori how many suppliers she should source from. This makes it difficult to run a simple auction in Stage 1. Dasgupta and Spulber (1990) studies the optimal mechanism with strictly convex costs, however the proposed implementation of the optimal mechanism is very complex.

It is worth noting that a recent paper, Tunca and Wu (2009), studies simplifying complex mechanisms for strictly convex production costs using a two-stage approach. The author are motivated by the observation that running a complex mechanism such as in Dasgupta and Spulber (1990) with many suppliers can be impractical.

To address this concern, they propose a two-stage mechanism, in which the buyer first uses a preliminary auction to select a small number of suppliers as preferred candidates, then runs a complex mechanism (such as the mechanism proposed by Dasgupta and Spulber) among this reduced set of preferred candidates. This mechanism is generally *not optimal*: With convex production costs, it may be optimal for the buyer to purchase from *all* suppliers, so pre-selecting a subset of suppliers before learning their full cost information means in some cases the final purchasing decision will not be optimal. However, the authors point out that at the cost of optimality, the buyer has a more practically implementable mechanism (compared with an optimal mechanism such as the mechanism proposed by Dasgupta and Spulber). To reiterate, the two-stage mechanism proposed in Tunca and Wu (2009) achieves simplicity by using pre-screening to reduce the set of candidate suppliers at the expense of optimality in their convex cost setting. My modified open-descending auction also uses pre-screening to reduce the set of candidate suppliers (actually to the greatest extent possible, i.e., to a single supplier), but thanks to my mechanism's design and my concave (including linear) cost setting, I am able to maintain optimality.

3.3.2 Ex Ante Asymmetric Concave Costs

The modified open-descending auction can also be extended to allow ex ante asymmetric costs. In this section, I assume each supplier i's base cost c_i is drawn from an independent random variable $C_i \in [\underline{c}_i, \bar{c}_i]$ with cdf F_i and pdf f_i. Define virtual cost function $\psi_i(c) \doteq c + \frac{F_i(c)}{f_i(c)}$ and assume ψ_i is increasing in c. Define initial quantity

$$Q_0 \doteq \max\{Q \in \mathbb{Z} | r(Q)/h(Q) \geq \psi_i(\bar{c}_i), \forall i\}$$

and assume $r(1)/h(1) \geq \psi_i(\bar{c}_i),\ \forall i$, so that $Q_0 \geq 1$.

It is well known that with ex ante asymmetric bidders, an optimal mechanism biases towards higher-cost bidders, and against lower-cost bidders to intensify competition. Thus, I propose the following *modified biased open-descending auction*, where the suppliers' bids are treated by supplier-specific *biasing functions* before being submitted in the auction. (Biased auctions have been studied in the literature; see Rezende (2009) for a recent example.)

Definition 3.3 (Modified Biased Open-Descending Auction)**.** The buyer first announces the following rules to all candidate suppliers:

- The mechanism has two stages. Stage 1 is a biased open-descending auction for a contract to supply Q_0 units. The rules of the auction are announced as follows. Before the auction, the buyer will inform each supplier i of his biasing function $b_i(\cdot)$, and supplier i's bid P_i will be treated as price $b_i(P_i)$ in the auction. The auction price P will start at an initial price set by the buyer and thereafter will always equal the current lowest price. Any supplier can always submit a bid that lowers the current auction price. The auction ends when no supplier is willing to further lower the auction price, and the final lowest bidder is the winner. Suppose the auction ends at price P_0. The auction's winner, say supplier i, has the obligation to deliver Q_0 units to the buyer and the buyer will pay $b_i^{-1}(P_0)$ for the Q_0 units.

- In Stage 2, the buyer will make one additional offer to the winning supplier from Stage 1 that specifies the prices the buyer will pay for each additional unit (beyond Q_0) that the supplier chooses to deliver. The exact prices will be specified in Stage 2, however the price for the k^{th} additional unit will be no higher than $h(k+Q_0)\psi_i^{-1}(P_0/H(Q_0))$. The winning supplier can choose to deliver any additional quantity (including zero).

After announcing the rules, the buyer executes the mechanism as follows.

Stage 1 The buyer first informs each supplier i of the biasing function $b_i(\cdot) = H(Q_0)\psi_i(\cdot/H(Q_0))$, then runs an open-descending auction for a contract to supply Q_0 units, starting at price $H(Q_0)\max\{\psi_i(\bar{c}_i)\}$.

Stage 2 Suppose Stage 1's auction concludes at price P_0 and supplier i is the winner. The buyer then offers to pay $h(k+Q_0)\psi_i^{-1}(\min\{P_0/H(Q_0), r(k+Q_0)/h(k+Q_0)\})$ for the k^{th} additional unit beyond Q_0 that supplier i chooses to deliver, where $k \leq \max\{l | r(l+Q_0)/h(l+Q_0) > \psi_i(c_i)\}$.

Theorem 3.3. *The modified biased open-descending auction is an optimal mechanism for the buyer. In Stage 1, each supplier i's dominant strategy is to keep lowering his bid until auction price P reaches $b_i(H(Q_0)c_i)$ or the auction ends (when other suppliers are not willing to bid lower), whichever happens first. In Stage 2, supplier i (given he is the winner) will deliver an additional $q_0 \doteq \max\{q \in \mathbb{Z} | r(q+Q_0)/h(q+Q_0) > \psi_i(c_i)\}$ units to the buyer.*

One can again easily see that the symmetric concave costs problem (Definition 3.2 and Theorem 3.2) is a special case of the asymmetric concave costs problem (Definition 3.3 and Theorem 3.3), with all $\psi_i(\cdot)$ replaced by a common $\psi(\cdot)$. With asymmetric concave costs, the price cap curve on the additional offer $h(k+Q_0)\psi_i^{-1}(P_0/H(Q_0))$ that the buyer announces before the auction is supplier-specific. In addition, each supplier's bid must be treated by his individual biasing function b_i before being submitted to the auction. The use of biasing functions adds complexity to the mechanism, but biasing is unavoidable whenever an optimal mechanism with ex ante asymmetric agents is to be implemented.

3.4 Concluding Discussion

The classical newsvendor model and its variations are among the most important in operations management, witnessed by the numerous research papers published about them, as well as their prominence in business education and their applications in practice. The vast majority of newsvendor models in the literature assume that the newsvendor purchases the good at an exogenous and fixed price, and selects the optimal purchase quantity. In many real life situations, however, the newsvendor has to source the good from among multiple competing suppliers. In these situations, the competition among the suppliers drives the purchase price, the purchase price determines the buyer's purchase quantity, and the purchase quantity will in return affect the suppliers' competition. What is the buyer's optimal strategy facing these situations? For this problem, Chen (2007) made a notable contribution by developing an optimal sourcing mechanism for the buyer in which suppliers bid and pay for the right to supply any quantity of their choosing to the buyer. This is somewhat equivalent to the buyer "selling the business" to the supplier and has some parallels in the retail industry, for example slotting allowances as pointed out by Chen. However, "selling the business" whereby suppliers bid to purchase a payment schedule from the buyer is not as widely observed in practice as having the suppliers compete in a traditional procurement auction whereby suppliers bid prices to supply a guaranteed quantity.

In this chapter, I show that a simple variation of the standard open-descending auction for a fixed-quantity contract can implement the optimal mechanism for this problem. This is attractive because many suppliers are familiar with standard open-descending auctions. The suppliers' decision making in my modified open-descending auction is also extremely simple, which adds to its practical appeal. Speaking to the unique role that the open-descending auction plays in my setting, I find that one cannot replace the open-descending auction by other commonly used auction formats such as sealed first-price, Vickrey or reverse Dutch auctions, without compromising the mechanism's optimality, or adding complexity. Finally, the modified open-descending auction can be easily extended to allow concave production costs and ex ante asymmetric suppliers, showing great flexibility.

CHAPTER IV

Price-Quoting Strategies of a Tier-Two Supplier

4.1 Introduction

Business-to-business transactions are important: According to Kshetri and Dholakia (2002), the global value of goods and services traded among businesses is estimated to exceed US$60 trillion annually. This important area is addressed by a growing operations management literature on procurement, including work on reverse, or procurement, auctions, which have become commonplace in practice (Jap 2007). In such studies the auctioneer (often an OEM) and the bidders (the OEM's immediate suppliers, also called *tier-one suppliers*) are modeled as strategic decision makers, but the suppliers' costs are typically assumed to be exogenous. However, in reality these suppliers' costs are often influenced by many factors, including their own internal production costs, as well as their costs of sourcing from upstream suppliers. Although existing models capture cases where these upstream suppliers sell commodities and have little pricing power, when an upstream supplier is powerful she too should be considered as a strategic decision maker whose pricing decisions affect the downstream suppliers' costs and thus the outcome of the OEM's auction. To take a first step at analyzing this important but under-studied issue, in this chapter I focus on the strategic pricing decisions of a tier-two supplier whose downstream customers (tier-one suppliers) will compete in an OEM's reverse auction.

I became interested in this problem after encountering this type of setting in practice. I observed a situation in which a Fortune 50 OEM, who was considering whether it would be cost-effective to construct additional office space, ran a reverse auction to solicit bids from general contractors. For a particular key element of the construction, these general contractors were relying on the same specialized subcontractor who had quoted them prices. Obviously, the subcontractor faced a strategic problem of how to quote prices to the general contractors, but it was not clear to me what strategy should be employed, nor was there any existing research addressing this issue that I am aware of.

Several factors make understanding the tier-two supplier's price-quoting strategy an interesting and difficult problem, and to examine it I employ a stylized model capturing the following salient features of the underlying supply-chain situation:

First, I initially assume that all the tier-one suppliers depend on one single tier-two supplier for a critical part of their product/service. In my example above the tier-two supplier had specialization in a particular type of construction. More generally, it is not rare for many tier-one suppliers to rely on a tier-two supplier that specializes in a particular type of component or service — for example, The Economist (2009) points out that 90% of the micro-motors used to adjust the rear-view mirrors in cars are made by Mabuchi, and TEL makes 80% of the etchers used in making LCD panels. My initial assumption that all tier-one customers depend on the same tier-two supplier is a simplification of this reality, for expositional purposes. In §4.6 I will relax this simplifying assumption and show that the main insights from the base model still hold true.

Second, the tier-one customers are in competition for an indivisible contract from the OEM, modeled as an open-descending reverse auction held by the OEM. Naturally, a tier-one supplier will only execute a quote (purchase from the tier-two supplier at the quote price) if he has an order from the OEM. This alongside the first feature leads to the following implication in the supply chain: When her customers are competing for the same contract, the tier-two supplier can give each customer a quote, but anticipates that only one customer can possibly win the OEM's contract and subsequently purchase from her at the quoted price.

Third, the tier-two supplier's quotes directly affect her customers' costs. I model this as the tier-one supplier's costs for fulfilling the OEM's contract consisting of the cost of inputs from the tier-two supplier plus their internal cost of processing the inputs into final products. The tier-two supplier has an incentive to quote high prices to maximize revenue, but of course the OEM would walk away and not award the contract if the tier-one suppliers' costs end up being exorbitant. I capture the latter reality with a walk-away price, or reserve price, in the OEM's auction.

Fourth, firms typically closely guard their cost information to protect profits, so in reality the tier-two supplier does not perfectly know her customers' true internal processing costs. Of course, if the tier-two supplier somehow knew the tier-one suppliers' processing costs, she could pick out the most cost-efficient supplier and give this supplier a price quote that consumes as much of his profit margin as possible, leaving only minimal profit for him. In contrast, when faced with uncertainty about the tier-one suppliers' costs, the tier-two supplier benefits from giving quotes to multiple tier-one customers because doing so affords her multiple shots at the OEM's contract. I model this using the canonical asymmetric information setting from economics, whereby each tier-one supplier knows exactly his own internal processing cost, but the tier-two supplier only has an estimate.

These factors conspire to make the tier-two supplier's quoting problem interesting and difficult. The tier-two supplier has an incentive to offer quotes to multiple tier-one customers, but only has imperfect information about how aggressive she should be and which quote (if any) will ultimately bring her revenue. Particularly, she faces a complicating situation due to downstream competition — if she decides to tip the scales in favor of one customer (offers him a low quote that is likely to give him an overall cost advantage), this would come at the expense of making the other customers — those with higher quotes that would deliver more revenue to the tier-two supplier — less likely to win the auction. These tensions lead to several questions. Does the tier-two supplier always want to provide equal quotes to ex ante identical tier-one suppliers? In general, what does the tier-two supplier's quoting strategy look like? I address these questions in §4.4, where I find that the tier-two supplier would indeed offer non-identical quotes, particularly when the value of the underlying OEM contract is high.

There are two implications in the forgoing discussion. The first implication is that the supplier may offer her customers different (unequal) quotes. There is evidence that this is done in practice and does survive legal challenges. These legal challenges attempt to link unequal quotes to price discrimination, which the *Robinson-Patman Act* (RPA) forbids. In a recent example of this, Michael Foods (a manufacturer of egg and potato products) offered quotes for ingredients to two food service providers, Sodexo and Feesers, who then bid against each other for a food service contract at a downstream institution. Michael Foods offered Sodexo a quote lower than the quote it offered Feesers, and consequently Feesers sued alleging an RPA violation. While successful in a district court, this claim was overturned by the Third Circuit court (*Feesers, Inc. v. Michael Foods, Inc.*, Jan. 7, 2010), who cited RPA's "two purchaser" requirement: For RPA to apply, at least two sales must be alleged to different purchasers at different prices, meaning mere *offers to sell* are not sufficient. Citing a series of similar court decisions in several competitive bidding settings, Stoll and Goldfein (2007) note that courts generally find in favor of suppliers who offer their customers different price quotes where only the winning bidder actually purchases from the supplier, because the two purchaser requirement is not met. Nonetheless, they also point out that in at least one instance, a circuit court has ignored the two purchaser requirement and applied the RPA in a competitive bidding setting. The preponderance of evidence suggests that unequal quotes generally survive legal challenges, but Stoll and Goldfein (2007) point out that the Supreme Court has yet to definitively rule on whether RPA can be applied to the competitive bidding setting, and is not likely to do so for at least several years.

The legal challenges make it clear that tier-one customers may find unequal quotes unfair, as a customer may resent being forced to accept a higher quote if others are given lower quotes for the exact same good. Thus a tier-two supplier might wish to avoid unequal quotes solely to preserve customer goodwill. Addressing this possibility, in §4.5.1 I extend my analysis to consider the case where the tier-two supplier always provides equal quotes to its customers, and study how the identical quote compares to non-identical quotes. I find that the restriction to identical quotes generally hurts the tier-two supplier, unless its customers are ex ante symmetric and the underlying value of the OEM's contract is relatively low.

The second implication is that, the tier-two supplier quotes prices depending on her estimates of the tier-one suppliers' costs. If the tier-two supplier can somehow learn her customers' private cost information, she could potentially resolve her concern of getting too greedy with her quotes and accidentally ending up empty-handed. This is because better information can reduce/eliminate uncertainty over her customers' costs, enabling her to better target quotes without pushing the tier-one suppliers' costs past the OEM's walk-away reserve price. Furthermore, better information may allow the tier-two supplier to identify and back only the most efficient tier-one customer. While there is a clear benefit to soliciting information, the tier-two supplier cannot naively ask her customers for their costs — each tier-one supplier would tend to claim to have high internal processing cost, hoping that the tier-two supplier would quote him a lower price. These tensions lead to several questions: How should the tier-two supplier best solicit cost information from her customers? Does soliciting the cost information really resolve the aforementioned concerns? I explore these issues in §4.5.2, where I show that the tier-two supplier's optimal mechanism resembles auctioning off a single quote among the tier-one customers. The tier-two supplier has her customers compete for the opportunity to receive a single quote before they bid for the OEM's business, and in so doing the tier-two supplier increases the chance her quote results in revenue (leads to business with the OEM).

Before introducing modeling details in §4.3, in the next section I briefly review related literature.

4.2 Literature Review

This chapter deals with procurement and sourcing. There is a long line of research on this topic. A major branch of procurement literature this research fits into deals with competitive bidding. In these settings multiple suppliers compete in an auction for a contract from a buyer (see Elmaghraby (2000) for a survey on this topic; recent examples in the operations management literature include Chen et al. (2005), Chen et al. (2008), Li and Scheller-Wolf (2010), etc.). In this vast literature the bidding is virtually always analyzed at just a single supply-chain interface, namely the multiple suppliers seeking the contract and the buyer offering the contract are the only players considered. In this chapter, I expand upon this scope by also considering the actions of a tier-two supplier who is situated one tier below the interface at which the auction for the contract occurs. To the best of my knowledge, the only other paper to consider competitive bidding within a multiple-tier supply chain is Lovejoy (2010). Lovejoy considers a supply chain-formation problem in a supply chain of multiple tiers, from a commodity market to a final buyer (similar to the OEM in this chapter). In each tier there may be one or more potential suppliers, but only one will emerge as active. It is assumed that the supply cost of the commodity market and the final buyer's purchase price are exogenous, and the costs of the suppliers at all tiers are publicly known.

Lovejoy defines the "balanced principal solution" as a prediction of the supply chain formation and profit distribution outcome, whereby each supplier bids a price to each buyer and vice versa, then the mutually preferred pair form a supply chain and share the margin. A key difference between Lovejoy (2010) and this chapter is that the former assumes all supply chain members have complete and perfect cost information of all other parties, while imperfect information and the resulting information asymmetry and uncertainty play a pivotal role in this chapter.

A major part of my model and results revolve around the possibility of the tier-two supplier quoting different prices to her customers. Thus I also want to compare my findings to the vast literature on price discrimination (Stole (2007) offers an excellent review). In this literature, price discrimination generally occurs either between imperfect substitute goods, or between separated markets/segments. When there is perfect competition (all customers have full access to perfect substitute goods), no price difference is sustainable. Therefore, the source of price discrimination in the price discrimination literature is imperfect competition. In my setting, however, price discrimination arises despite the fact that the tier-one suppliers' offerings are perfect substitutes, and the OEM is free to choose any one supplier. This is because the source of my price discrimination is asymmetric information (tier-one suppliers are privileged with better information about their costs than the tier-two supplier). With asymmetric information, the tier-two supplier quotes different prices in order to manage the trade-off of her potential revenue versus the risk of not having an order. Thus, while similar in appearance to the traditional price discrimination which arises to take advantage of imperfect competition, the price discrimination in my setting arises for a totally different reason, namely to manage uncertainty caused by asymmetric information.

More generally, the main novelty of this research is that it adopts the perspective of a tier-two supplier. While there have been many studies of buyers' procurement auctions and suppliers' bidding strategies therein, mine is the first to study the price-quoting decisions of a supplier whose *customers* will be competing with each other in a downstream auction.

4.3 Base Model

4.3.1 Supply Chain Structure

I model a three-layer supply chain. The top layer is an original equipment manufacturer (OEM) who wishes to auction off an indivisible contract for the provision of goods or services. I refer to the suppliers who compete in the OEM's sourcing auction as *tier-one suppliers*. For expositional purposes I assume that there are two tier-one suppliers competing for the OEM's business, and denote them as TO_1 and TO_2 (TO stands for *tier-one*). The tier-one suppliers do not supply the good or service entirely by themselves; they require inputs from an outside source before they can produce the good or supply the service requested by the OEM. I model this outside source as a tier-two supplier, TT (standing for *tier-two*), who supplies TO_1 and TO_2 with a critical component or service. This stylized model captures a variety of situations, ranging from manufacturing (where TT supplies a critical component) to food service (where TT Michael Foods provides inputs to downstream TOs Sodexo and Feesers). For consistency I will refer to TT as providing a component used in production by TOs. The focus of this chapter is on how TT should price her component for customers TO_1 and TO_2. (A more general model with any number of TOs and also participants in the OEM's auction who do not depend on TT for a component is examined in §4.6.)

4.3.2 OEM's Auction

The OEM uses an open-descending auction with a reserve price. Such an auction is simple to describe to bidders and is widely used in practice (Jap 2007): The auction begins at the publicly announced reserve price r, and participants alternately bid the price down, until no one is willing to bid any lower. The last remaining bidder wins the contract and is paid the auction's ending price. If no one is willing to match the starting price r, no contract is awarded. The reserve price sets a ceiling on the amount the OEM is willing to pay for the contract, capturing the OEM's alternatives to contracting with a supplier. For example, when purchasing goods, the OEM would forgo the purchase if exorbitant procurement costs would make acquiring the goods unprofitable.

During the auction, tier-one suppliers TO_1, TO_2 compete on contract price. TO_i's cost to fulfill the OEM's contract is composed of two parts $x_i + y_i$, where x_i is the cost of purchasing the component produced by TT, and y_i is the cost of processing the component into the final product, shipping it to the OEM, etc. For simplicity I refer to y_i as processing cost. I assume the TOs are rational and seek to maximize their expected profits, as is standard in the auction-theoretic literature. If TO_i wins the OEM's auction at price p, his profit equals $p - x_i - y_i$; if he loses the auction his profit is zero. Note that in the OEM's auction, each TO_i finds it a dominant strategy to lower his bid until either he wins the auction, or the price drops below his total cost $x_i + y_i$. Therefore, TO_1 will win the contract if and only if $x_1 + y_1 < x_2 + y_2$ and $x_1 + y_1 < r$. Similarly, TO_2 will win the contract if and only if $x_2 + y_2 < x_1 + y_1$ and $x_2 + y_2 < r$. When $x_i + y_i > r$, $i = 1, 2$, the OEM will not award the contract.

4.3.3 TT's Problem

Like the TOs, TT is assumed to be a rational expected profit maximizer. The goal of this chapter is to study TT's price-quoting strategy that maximizes her expected profit. Let s_i denote TT's cost of supplying TO_i with the component needed for the OEM's contract, and normalize the components' alternative value for TT to be zero. I allow $s_1 \neq s_2$ even though the components delivered to each TO_i are identical, to address the possibly heterogeneous additional costs associated with each TO_i such as shipping costs. TO_i's processing cost y_i is his private information, and is a realization of random variable Y_i which has a commonly known, positive and finite pdf over a closed interval. Assume the Y_i's are independently distributed. The sequence of events begins with each TO_i soliciting from TT a quote x_i for a supply of the component. TO_i has the power to decide whether to execute the quote he receives, and obviously would only want to do so if he wins the OEM's contract. With their respective quotes in hand, TO_1 and TO_2 then compete in the OEM's auction. If TO_i wins the auction, TT will incur s_i to supply the component at the quoted price x_i; if neither TO_i is able to meet reserve price r, no deliveries or payments are made.

In what follows, for convenience I define TO_i's *realized base margin* to be $z_i \doteq r - y_i$, which is distributed according to $Z_i \doteq r - Y_i$. To ensure the typical auction-theoretic property that Z_i's failure rate is increasing (IFR), I assume that the Z_i's (or equivalently, Y_i's) have log-concave probability densities.[1] The realized base margin z_i is the highest revenue that TT can possibly achieve by selling to TO_i.

Thus far I have not imposed any assumptions on how TT will deal with the TOs regarding supplying the component. A common and easily implemented approach is that TT simply quotes a fixed price to each TO_i. Of course, in doing so TT would strategically account for the distributions of the base margins Z_i, $i = 1, 2$. I refer to this approach as *Quoting Prices (QP)*. Note that I did not rule out the possibility that TT provides different price quotes to its customers. As mentioned in §4.1, this

[1] Corollary 3 in Bagnoli and Bergstrom (2005) ensures that assuming Z_i's having log-concave probability densities leads to IFR of Z_i's. Corollary 5 in the same paper shows that assuming Z_i's having log-concave probability densities is equivalent to assuming so for Y_i's. Finally, the same paper points out that many common distributions have log-concave probability densities, including uniform, normal, logistic, and exponential distributions.

is indeed done in practice. While the TO_i's may perceive dissimilar quotes as unfair, the preponderance of court decisions suggests little if any legal ground for opposing such practices. Written formally, TT's problem is

$$\begin{aligned} \max \quad & P(x_1, x_2), && (4.1) \\ \text{where} \quad & P(x_1, x_2) \doteq \sum_{i=1,2} (x_i - s_i) \Pr\{TO_i \text{ wins the contract}\}, \\ & \qquad = \sum_{i=1,2} (x_i - s_i) \Pr(x_i + Y_i < r,\ x_i + Y_i < x_j + Y_j,\ j \neq i), \\ & \qquad = \sum_{i=1,2} (x_i - s_i) \Pr(Z_i > x_i,\ Z_i - x_i > Z_j - x_j,\ j \neq i). && (4.2) \end{aligned}$$

I call a pair of quotes (x_1, x_2) a *quoting strategy* and denote the pair as a whole by $X = (x_1, x_2)$. In particular, the *optimal quoting strategy*, namely the maximizing solution to problem (4.1), is denoted by $X^* = (x_1^*, x_2^*)$. It is clear that by a simple change of variables argument[2] I can transform the model into an equivalent one with $s_i = 0$. Therefore, without loss of generality, for concision I will hereafter assume $s_i = 0$. In subsequent sections I will study two other quoting approaches, but to avoid confusion I delay their formalization until I am ready to discuss them in §4.5.

4.4 Quoting Prices

Quoting prices to two competing tier-one suppliers can lead to quite complex trade-offs. For example, with two tier-one suppliers, TT can provide different prices to the two TOs. Suppose Z_i is uniformly distributed over $[a_i, a_i + 2]$, and a_1 and a_2 take the value of either 2 or 9 in the four possible combinations in Figure 4.1. At points A and D where $(a_1, a_2) = (2, 9)$ and $(9, 2)$, since TO_1 and TO_2 have very dissimilar base margin distributions, one would expect that TT will quote different prices to them. However, given that the two TOs are ex ante symmetric at points

[2] Define $\tilde{x}_i \doteq x_i - s_i, \tilde{Y}_i \doteq Y_i + s_i, \tilde{y}_i \doteq y_i + s_i$. Consequently, $\tilde{Z}_i = Z_i - s_i$ and $\tilde{s}_i = 0$.

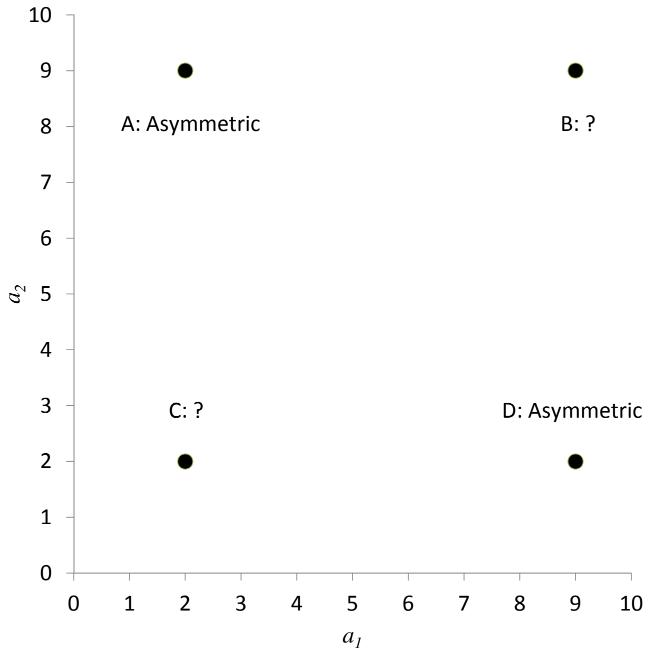

Figure 4.1: Symmetric or asymmetric quotes?

B and C where $(a_1, a_2) = (2, 2)$ and $(9, 9)$, would TT ever want to quote different prices? The answer is not immediately clear. These questions will be explored in this section.

4.4.1 Intuition

Before analytically examining TT's optimal quotes, I would like to discuss the subtleties involved using a simple example. Suppose two tier-one suppliers TO_1 and TO_2 are present for quotes. Assume both Z_i's have uniform distribution over $[2, 4]$. Since TT's quotes can be different for TO_1 and TO_2, for illustrative purposes I fix the quote to TO_2 at $x_2 = 2.5$ and examine what happens as I increase x_1, the quote to TO_1. As seen in Figure 4.2, the trade-offs are quite complex. As x_1 increases, the dark-colored lines represent TO_1's chance of winning the OEM's contract and the payoff that TT will receive if TO_1 wins, and light-colored lines represent the same for TO_2. Naturally, the payoff from TO_1, x_1, increases, and TO_1's chance of winning

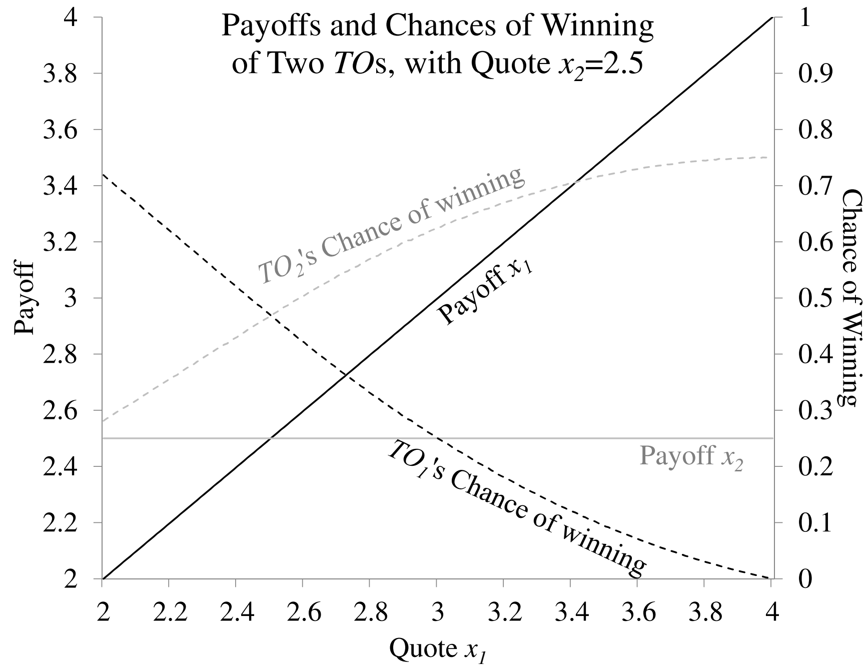

Figure 4.2: Payoffs and chances of winning for varying x_1, with $x_2 = 2.5$.

decreases. However, there is more: As x_1 increases, even though x_2 is fixed, the chance that TO_2 wins increases because TO_1 becomes less competitive. TT's profit is greatly affected by the competition amongst TT's own customers. For example, suppose TT quotes $x_1 = 3.6$ to TO_1. Although x_1 is much higher than x_2 so TT prefers TO_1 to win the contract, in reality TO_1's chance of winning is only about 7%, while TO_2's chance of winning is much higher at over 70%. This results in an interesting effect for TT: Whichever TO she prefers to win, because she quotes a high price to him, actually competes unfavorably against the other TO for the very same reason! Such subtleties make it difficult to intuitively see when, if ever, and why, a high-low quote combination for ex ante symmetric TOs as in my example may be a good idea. To answer this question, I need to resort to analytical tools.

4.4.2 Analytical Results

Recall that TT's problem is described by (4.1) and (4.2), and I assume $s_i = 0$ without loss of generality. First I provide a preliminary result related to the support of Z_i, denoted by $[a_i, a_i + h_i]$. Proof of this and subsequent results are provided in Appendix C.

Proposition 4.1. *It is never optimal to quote x_i out of the support of base margin Z_i, $i = 1, 2$.*

This is an intuitive result considering that TT will not give up guaranteed profit potential by quoting below a_i, but neither would TO_i execute a quote yielding him negative "profit" as would happen when TT quotes above a_i+h_i. Because the optimal x_i always stays in $[a_i, a_i + h_i]$, I only need to characterize x_i in this interval. I next identify a special category of quoting strategies.

Definition 4.1. A quoting strategy $X^s = \{x_1^s, x_2^s\}$ is said to be secure if $x_i = a_i$ for some i. A quoting strategy $X^r = \{x_1^r, x_2^r\}$ is said to be risky if $x_i > a_i$, $i = 1, 2$.

I call a strategy *secure* when with this strategy at least one TO_i will certainly meet the reserve price r, and thus TT can *secure* her business. With a secure strategy, $\Pr\{\text{both } TO\text{s losing}\} = 0$. On the other hand, with a *risky* strategy I have $\Pr\{\text{both } TO\text{s losing}\} > 0$, namely there is positive probability that TT will not transact.

Intuitively, when using a secure strategy, TT only needs to secure the business with one TO_i. The next proposition formalizes this idea and characterizes the *optimal secure strategy*, namely the one that generates the highest expected profit for TT among all secure strategies.

Proposition 4.2. *An optimal secure strategy X^{s*} will have only one $x_i^{s*} = a_i$. For the other $j \neq i$ it must be true that $x_j^{s*} > a_j$. In addition, either $x_j^{s*} > a_i$, or $x_j^{s*} = a_j + h_j$.*

The intuition behind Proposition 4.2 is straightforward. When quoting $x_i^{s*} = a_i$ to TO_i, TT is already guaranteed payoff a_i even before quoting to TO_j. Therefore, it would only make sense to provide a quote to TO_j that can possibly generate higher payoff to TT, i.e., $x_j^{s*} > a_i$. The only exception is when $a_i > a_j + h_j$ and no meaningful quote to TO_j ($a_j \leq x_j < a_j + h_j$) can match a_i; in this case TT's best strategy is to ensure TO_j never meets the reserve price by quoting $x_j^{s*} = a_j + h_j$ to him.

Developing upon this intuition, I show a special property of the optimal secure strategy. Assume Z_i is replaced by $\hat{Z}_i \doteq Z_i + a$ in (4.2), $i = 1, 2$ (i.e., shift both Z_i's by a). I denote the optimal secure strategy for problem (4.1) by $X^{s*} = \{x_1^{s*}, x_2^{s*}\}$, and denote the optimal secure strategy after shifting Z_i's by $\hat{X}^{s*} = \{\hat{x}_1^{s*}, \hat{x}_2^{s*}\}$.

Proposition 4.3. *For all a, $\hat{x}_i^{s*} = x_i^{s*} + a$, $i = 1, 2$. In other words, the optimal secure strategy remains fixed relative to the support of Z_i's as both Z_i's are shifted by the same amount.*

According to Proposition 4.2, with a secure strategy where $x_i^s = a_i$, TT locks in a guaranteed payoff a_i, then gambles with TO_j for additional profit, or eliminates TO_j if he is too inefficient. When quoting for TO_j, TT only cares about balancing the *additional* profit $x_j^s - x_i^s$, and the chance of getting it. Neither is affected when both Z_i's are shifted by the same amount. Thus the optimal secure strategy X^{s*} remains fixed relative to the support of Z_i's, invariant in a.

A secure strategy guarantees TT at least a minimal payoff. The following theorem, which studies how TT's quoting strategy changes as both TOs' base margins increase, characterizes when TT finds it optimal to use a secure strategy, and when she instead finds it optimal to use a risky strategy.

Theorem 4.1. *Replace Z_i by $\hat{Z}_i \doteq Z_i + a$ in (4.2) and consider the resulting optimal strategy. There exists a threshold $T_{sec} < \infty$ such that when $a < T_{sec}$ the optimal strategy is risky and when $a > T_{sec}$ the optimal strategy is secure.*

This theorem establishes a threshold in a for the transition of the optimal strategy between secure and risky. The result is important because it describes how the tier-two supplier changes her behavior as she faces tier-one suppliers with higher base margins (e.g., because the tier-one suppliers become more efficient and cut their costs, or because the OEM announces a higher reserve price). As TOs' base margins increase, TT will eventually want to use a secure strategy. This makes intuitive sense: The business becomes more and more lucrative, and TT wants to make sure she at least obtains the business. On the other hand, if the base margins are fairly low, losing the business is not as damaging, and TT may want to take a risk and gamble with higher quotes to the TOs. Doing so may lead to loss of business, but can provide higher payoffs if business is won.

I continue to use the numerical example in Figure 4.1 to demonstrate this behavior (see Figure 4.3). Recall that $Z_i \sim U[a_i, a_i + 2]$. I plot TT's strategy (*risky* or *secure*) as a_1 and a_2 take values from 0 to 10. Theorem 4.1 indicates that starting anywhere in Figure 4.3, moving northeast along a 45° line will eventually lead to using a secure strategy (and never risky again).

I am now ready to answer an important question asked at the beginning of §4.4 about Figure 4.1, namely how TT would quote when facing ex ante identical TOs. Theorem 4.1 finds that with ex ante identical TOs, at sufficiently high a_i (e.g., point B $(a_1, a_2) = (9, 9)$), TT will use a secure quoting strategy, which (by Proposition 4.2) is always asymmetric. I want to point out that, although TT provides asymmetric quotes at point B $(9, 9)$ as well as points A $(2, 9)$ and C $(9, 2)$, the reasons are different. Quoting asymmetric prices at points A $(2, 9)$ and D $(9, 2)$ is expected as the TOs are ex ante asymmetric. At point B $(9, 9)$, although the TOs are ex ante identical, TT views them differently: She uses one to lock in the business, and gambles with the other. It is the different roles TT wants them to play that leads to asymmetric quotes to symmetric TOs.

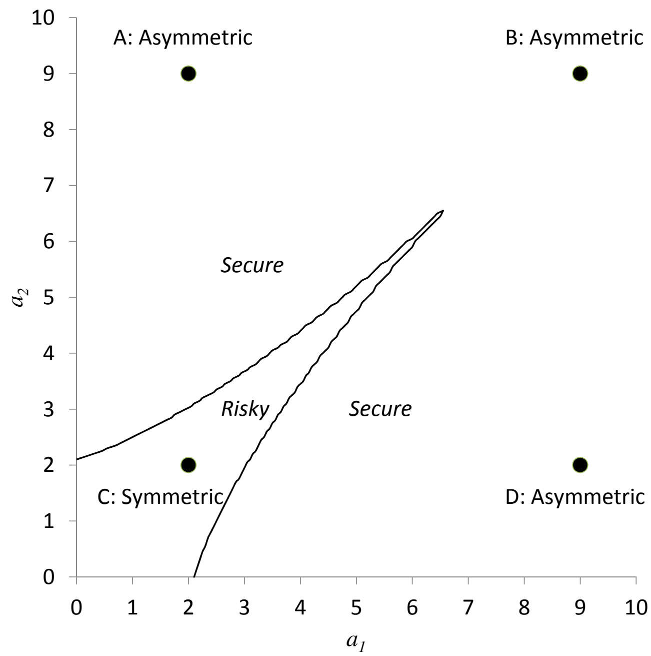

Figure 4.3: Regions of risky and secure strategies, where $Z_i \sim U[a_i, a_i + 2]$.

The different roles TT has her customers play in a secure strategy lead to the following interesting dynamic: TT puts up with the fact that the TO with the lower quote is most likely to win in the OEM's auction, leaving less chance for TT to get the higher quote. Despite this fact, Theorem 4.1 shows that TT will desire to treat her two customers differently and induce this situation when the profit potentials are high enough and locking in a sizeable profit is paramount. Finally, I also want to point out that when using a secure strategy, TT may still get the higher quote so there is usually still uncertainty in TT's payoff.

On the other hand, when a secure strategy is not as attractive, I may expect that TT would want to treat her two customers more equally. Indeed, the next theorem shows that when the two TOs are ex ante symmetric, if the profit potential is sufficiently low, the optimal strategy will be risky and symmetric, and thus provides both customers with equal quotes.

Theorem 4.2. *Assume Z_1 and Z_2 are i.i.d., replace Z_i by $\hat{Z}_i \doteq Z_i + a$ in* (4.2),

$i = 1, 2$, and consider the optimal quotes (x_1^*, x_2^*). There exists some $T_{sym} \leq T_{sec}$ such that when $a < T_{sym}$ the optimal strategy is risky and symmetric.

With Theorem 4.2 I can finish answering the question raised earlier about Figure 4.1, namely how TT would quote when facing ex ante identical TOs: With ex ante identical TOs, at sufficiently low a_i (e.g., point C $(2, 2)$), TT will quote equal prices. The reason is that, when the profit potential is sufficiently low, TT does not care simply about securing the minimal possible profit, but seeks to maximize her expected profit in light of the risk of getting nothing. This is achieved with a symmetric quoting strategy.

In conclusion, TT's price-quoting strategy is greatly affected by the profit potentials of the TOs. When the profit potentials are high, TT gives one TO a low quote and the other TO a higher quote, using the TOs for different strategic purposes. This is the consequence of TT's desire to lock in a sure-fire payoff, but this results in even ex ante symmetric customers being treated differently. On the other hand, when the tier-one profit potentials are low, TT uses a different approach that does not guarantee her a payoff, but takes the risk of not getting any business in exchange for potentially higher payoffs. Since neither TO is given a sure-fire low quote, there is no clear advantage granted to either customer. In such a case, TT treats her customers more equally as they serve similar strategic purposes for her. A numerical example of the optimal price quotes will be later provided in Figure 4.4.

4.5 Alternative Approaches

Thus far I have studied how TT should best provide quotes to her customers (the TOs). In doing so I allowed TT to provide different quotes to different TOs. However, there can be situations where the TOs would perceive different quotes as being unfair. Such a concern may prompt TT to restrict herself to only offering identical quotes to her customers (imposing the constraint $x_1 \equiv x_2$). I refer to this approach as *Quoting Equal Prices (QEP)*. Intuitively, this case is more likely to occur when TT does not have much supply chain power.

At the other extreme, if TT is unconcerned about fairness issues, and she has the power to set forth any rules of her choosing in supplying the component to the TOs, she could use an *Optimal Mechanism (OM)* to maximize her expected profit. As optimality rather than simplicity is the primary concern, an optimal mechanism may transcend quoting prices only based on priors about TOs' costs, and involve more elaborate procedures such as soliciting cost information from the TOs. In this section I study these two alternative approaches.

4.5.1 Quoting Equal Prices

With the QEP approach, TT's problem is written formally as

$$\begin{aligned} \max \quad & P(x) & (4.3) \\ \text{where} \quad & P(x) \doteq x \sum_{i=1,2} \Pr\{TO_i \text{ wins the contract}\}, \\ & \quad\quad = x \Pr(x < \max\{Z_1, Z_2\}). & (4.4) \end{aligned}$$

The QEP approach is the same as QP except for the constraint to provide equal quotes to the TOs. Having learned that the optimal QP quotes can be different even to ex ante symmetric TOs, it is natural to ask in what way the constraint of using equal quotes affects TT's strategies. I take a number of steps to answer this question.

First, as a natural extension of Proposition 4.1, I have the following preliminary result.

Proposition 4.4. *Suppose TO_i's base margin Z_i has support $[a_i, a_i + h_i]$. Then it is never optimal to quote x out of* $[\max\{a_1, a_2\}, \max\{a_1 + h_1, a_2 + h_2\}]$.

Then I show the optimal quote is described by the following condition.

Proposition 4.5. *The optimal QEP quote x^* must satisfy*

$$x^* = \frac{1 - F_1(x^*)F_2(x^*)}{f_1(x^*)F_2(x^*) + F_1(x^*)f_2(x^*)}.$$

Furthermore, when Z_1 and Z_2 are i.i.d. with cdf F and pdf f, the unique optimal QEP quote x^* is determined by

$$x^* = \frac{1 - F(x^*)^2}{2F(x^*)f(x^*)}.$$

Proposition 4.5 allows me to characterize the behavior of the optimal QEP quote with symmetric TOs.

Theorem 4.3. *Assume Z_1 and Z_2 are i.i.d. with support $[0, h]$, replace Z_i by $\hat{Z}_i \doteq Z_i + a$ in (4.4), $i = 1, 2$ and consider the optimal QEP quote x^*. As a increases, x^* must be greater than a, but $x^* - a$ will asymptotically converge to 0. In other words, facing symmetric TOs, the optimal QEP quote asymptotically converges to, but never reaches, a secure quote.*

Corollary 4.1. *Assume Z_1 and Z_2 are i.i.d., replace Z_i by $\hat{Z}_i \doteq Z_i + a$ in (4.2) and (4.4), $i = 1, 2$, and consider the optimal QP quotes (x_1^*, x_2^*) and the optimal QEP quote x^*. When a is sufficiently small, I have $x^* = x_1^* = x_2^*$, and when a is sufficiently large, I have $x_1^* < x^* < x_2^*$ (assume without loss of generality that $x_1^* \leq x_2^*$).*

Comparing Theorem 4.3 to Theorem 4.1 reveals an important insight: While the optimal QP strategy may be risky or secure (Theorem 4.1), the optimal QEP strategy with ex ante symmetric TOs is *always* risky (Theorem 4.3). This distinction makes clear the significance of the capability to discriminate, as the usage of secure strategies depends on it. In addition, Corollary 4.1 establishes how the optimal QEP quote compares to the optimal QP quotes. As illustrated in Figure 4.4, which continues the

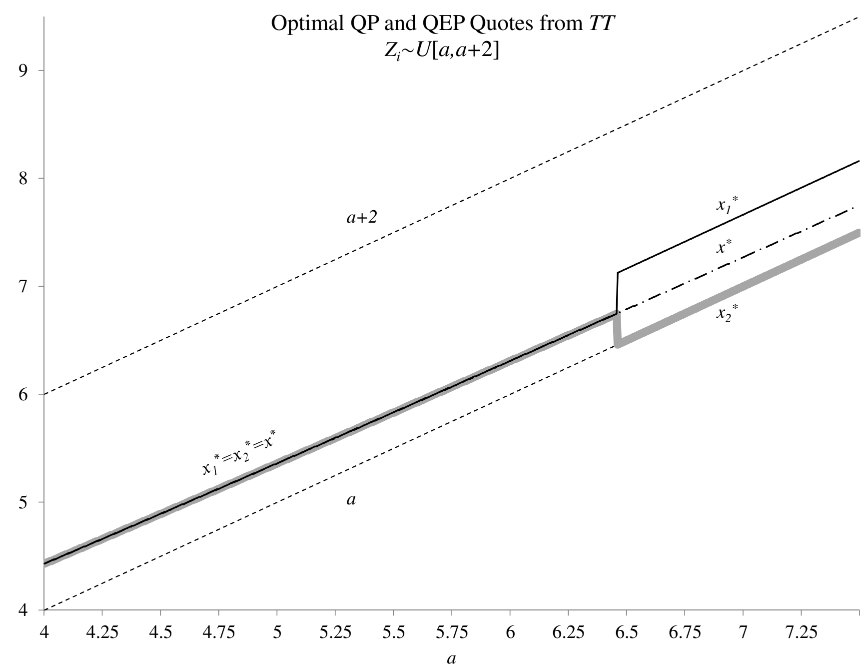

Figure 4.4: Optimal QP and QEP strategies, with symmetric uniform costs

numerical example of Figure 4.1 ($Z_i \sim U[a, a+2]$), when the profit potential is low (a is small) the optimal QEP quote coincides with the equal optimal QP quotes, and when the profit potential is high (a is large) the optimal QEP quote lies between the unequal optimal QP quotes. This addresses the question posed at the beginning of the subsection: With symmetric TOs, the constraint of using equal quotes has a negative impact on TT's expected profit only when the TOs' profit potentials are high. As an aside, Figure 4.4 also illustrates several results from Section 4.4; in this example I have $T_{sym} = T_{sec} \approx 6.4$ (Theorems 4.1 and 4.2), and x_1^* and x_2^* move in parallel as a increases once the optimal strategy becomes secure at $a = T_{sec}$ (Proposition 4.3).

4.5.2 Optimal Mechanism

In §4.4 I established how TT should optimally quote prices to the TOs based on her priors about TOs' processing costs. These quotes, however, could potentially lead to non-transaction (if she quotes above both TOs' profit margins), or transaction with money left on the table (if she quotes too low). TT may mitigate such concerns if she could find ways of soliciting TOs' private cost information. However, soliciting accurate information is not easy. For example, if TT simply asks the TOs "what is the highest quote you can accept (what is your base margin z_i)", each TO will claim to have a very low base margin in the hopes of receiving a lower quote from TT. The asymmetry of information clearly puts TT at a disadvantage. On the other hand, although both TOs are potential customers of TT, eventually at most one TO will get the OEM's contract. Therefore TT might counter against the TOs' incentives to underreport their base margins by playing one TO against another when she determines their quotes.

There are of course infinitely many ways in which TT could solicit, and predicate her actions upon, information from the TOs. The challenge is to find the way that generates the greatest expected profit for TT. To tackle this challenge, I utilize *optimal mechanism design* theory. A generic mechanism can be described as a set of "rules" that operate on cost signals provided by the TOs. I will formalize this below, but before doing so I make an observation. Canonical mechanism design analysis, Myerson (1981), designs a seller's optimal mechanism where the mechanism is comprised of *allocation* and *payment* rules, and after the seller receives cost signals the mechanism determines who receives the item and who pays how much. In my setting, however, TT does not have the power to allocate the OEM's contract; the decision of which TO (if any) gets the contract is determined by the OEM's auction. Instead of directly controlling the contract allocation, TT can indirectly affect the contract allocation via the quotes she provides the TOs. To capture this fact, in my setting I replace allocation rules with quoting rules, which alongside the (upfront) payment rules form my mechanism. The resulting mechanism subsumes a range of possibilities, from pure non-contingent upfront purchasing as in Myerson (1981) (positive upfront payments and zero quotes), to payments and purchasing purely contingent upon the OEM's auction (zero upfront payments and positive quotes), and any mixture in-between.

Having described that my mechanism involves cost reports, upfront payments, and quotes, I now formalize these concepts. TT announces a mechanism — a set of rules

$$\{p_1(\check{z}_1, \check{z}_2), p_2(\check{z}_1, \check{z}_2), x_1(\check{z}_1, \check{z}_2), x_2(\check{z}_1, \check{z}_2)\}$$

that map the TOs reported base margins $\check{z}_i$ into upfront payments $p_i \geq 0$ and quotes $x_i \geq 0$ to TO_i, $i = 1, 2$. Each TO chooses whether to participate, and if so, reports his base margin as $\check{z}_i$ (not necessarily equal to his true base margin z_i). With the TOs' reports in hand, TT then announces the payment $p_i(\check{z}_1, \check{z}_2)$ that TO_i must pay TT upfront (before participating in the OEM's auction), and quotes the price $x_i(\check{z}_1, \check{z}_2)$ that TO_i must pay TT if he chooses to order the component from TT (of course, a rational TO_i would order only if he wins the OEM's contract). In a sense, the upfront payment can be interpreted as a fee that a TO pays for the right to purchase the component if he wins the OEM's contract. Thus, this setup is similar to the capacity reservation and execution type of supply contracts, which are prevalent in practice. The TOs' strategic behaviors in the mechanism are predicted by the Bayesian-Nash equilibrium concept, and TT's expected profit is derived from the TOs' equilibrium strategies.

The goal of mechanism design is to find the mechanism which maximizes TT's expected profit. By the *revelation principle*, it suffices to search for the optimal mechanism only among *direct-revelation mechanisms*, namely those in which each TO willingly reports his true base margin in equilibrium. Consequently, TT's (simplified) mechanism design problem is

$$\max_{p_1,p_2,x_1,x_2} \quad \sum_{i=1,2} [p_i(\check{z}_1, \check{z}_2) + x_i(\check{z}_1, \check{z}_2) \Pr\{TO_i \text{ wins the contract}\}] \tag{4.5}$$

$$\text{s.t.} \quad v_i(z_i, z_i) \geq 0, \tag{4.6}$$

$$v_i(z_i, z_i) \geq v_i(z_i, \check{z}_i), \ \forall z_i, \check{z}_i, \tag{4.7}$$

where $v_i(z_i, \check{z}_i)$, defined as

$$E_{Z_j}[(z_i - x_i(\check{z}_i, Z_j) - (Z_j - x_j(\check{z}_i, Z_j))^+)^+ \Pr\{Z_j - x_j(\check{z}_i, Z_j) < z_i - x_i(\check{z}_i, Z_j)\} - p_i(\check{z}_i, Z_j)],$$

is TO_i's expected profit under mechanism $\{p_1(\check{z}_1, \check{z}_2), p_2(\check{z}_1, \check{z}_2), x_1(\check{z}_1, \check{z}_2), x_2(\check{z}_1, \check{z}_2)\}$, given that TO_i's true base margin is z_i, he reports base margin $\check{z}_i$, and TO_j, $j \neq i$ reports truthfully $\check{z}_j = z_j$. In the formulation, (4.6) and (4.7) ensure, respectively, that TT's mechanism is indeed *individually rational* (each TO_i willingly participates in the mechanism) and *incentive compatible* (each TO_i finds truthfulness optimal). With the range of search significantly reduced by the revelation principle, I can now characterize the optimal mechanism. Define the *virtual base margin* as $\psi_i(z_i) \doteq z_i - \frac{1-F_i(z_i)}{f_i(z_i)}$, where F_i and f_i are the cdf and pdf, respectively, of TO_i's base margin $Z_i \doteq r - Y_i$.[3] The following theorem describes the optimal mechanism.

Theorem 4.4. *An optimal mechanism* $\{p_1^*(\check{z}_1, \check{z}_2), p_2^*(\check{z}_1, \check{z}_2), x_1^*(\check{z}_1, \check{z}_2), x_2^*(\check{z}_1, \check{z}_2)\}$ *is as follows*[4]*: If* $\psi_i(\check{z}_i) > \max\{\psi_j(\check{z}_j), 0\}$*, then* $p_i^* + x_i^* = \max\{\psi_i^{-1}(\psi_j(\check{z}_j)), \psi_i^{-1}(0)\}$*,* $p_j^* = 0$*, and* $x_j^* > \check{z}_j$*; if* $\max\{\psi_1(\check{z}_1), \psi_2(\check{z}_2)\} < 0$ *then* $p_i^* = 0$ *and* $x_i^* > \check{z}_i$*,* $i = 1, 2$*.*

Theorem 4.4 reveals three key aspects of the optimal mechanism. First, TT chooses to "back" at most one (or possibly neither) TO_i in the OEM's auction, meaning she provides him with a quote low enough to guarantee that this TO_i can

[3]Since by assumption Z_i's probability density is log-concave, the results of Bagnoli and Bergstrom (2005) can be used to show that $\psi_i(\cdot)$ must be an increasing function.

[4]For readability I suppress the arguments $(\check{z}_1, \check{z}_2)$ when writing p^* and x^*.

meet the OEM's reserve price. Any TO that is not "backed" instead receives a very high quote that effectively prices them out of the running in the OEM's auction. Note that how much TT can get a TO to pay her depends on how much payment this TO can expect to get from the OEM. Once the TOs reveal their true costs, TT can control which TO_i will win the OEM's contract. Eliminating the other TO_j will maximize how much TO_i can get from the OEM, and in return, how much TT can get from TO_i.

Second, the optimal mechanism can be implemented by pure quotes, wherein the upfront payments are always zero. Notice that after soliciting cost information, TT knows which TO is going to win the contract. Because the outcome of the OEM's auction is no longer uncertain, contingent payments are interchangeable with upfront payments; in particular, the optimal mechanism can be achieved with quotes only, a favorable result considering the simplicity of using just quotes. However, I will need both upfront payments and price quotes to describe the optimal mechanism when I extend my model to more general settings in §4.6.

My third observation is based on the fact that under the optimal mechanism, TT's equilibrium realized profit equals $\max\{\psi_1(z_1), \psi_2(z_2), 0\}$, where $\psi_i(z_i) = z_i - \frac{1-F_i(z_i)}{f_i(z_i)}$ is TT's profit if she "backs" TO_i. Note that TT does not take all of TO_i's base margin z_i, but leaves him a profit of $\frac{1-F_i(z_i)}{f_i(z_i)}$, which increases in z_i. This setup is necessary to counter the TOs' natural tendency to underreport their base margins. It explains my third observation: When TT compares the TOs to determine which one to back, the comparison is done on virtual base margin rather than actual base margin. Of course, as is well known in the mechanism design literature, with ex ante symmetric tier-one suppliers, the one who reports the highest base margin will be chosen, while when the tier-one suppliers are ex ante asymmetric, this may not always be the case.

The distinction between the optimal mechanism and the QP approach is evident. With the QP approach, TT is willing to quote prices to both TOs, thus effectively taking two shots at success in the OEM's auction. Doing so is optimal for TT because at the time she provides quotes she is uncertain about which TO will succeed. However, when TT can make her quotes contingent upon information reported by her customers, she uses this opportunity to identify which TO is more profitable to back exclusively in the OEM's auction. This confirms my speculation made at the outset of this subsection that soliciting information could be superior to passively accommodating uncertainty.

Although Theorem 4.4 found the optimal mechanism, I have thus far been silent on how this mechanism might be implemented by TT. The optimal mechanism can actually be implemented as an auction. Instead of directly providing quotes to both TOs, TT has the TOs bid against each other to receive a quote. When the TOs are ex ante symmetric (Z_1 and Z_2 have identical distributions), the optimal mechanism can be implemented as the following open-ascending auction. Before the auction, TT sets the reserve price at $\psi_1^{-1}(0)$ ($= \psi_2^{-1}(0)$), and announces that only the TO willing to accept the highest price will receive a quote. Starting from the reserve price, the TOs take turns bidding higher prices they are willing to accept until one drops out, at which point the remaining TO is the auction winner and is provided a quote from TT equal to his final bid (the auction's ending price). If no TO is willing to match the reserve price, TT will not provide any quotes. When the TOs are ex ante asymmetric, the optimal mechanism can be implemented similarly, but with one key change: TT uses a biasing rule when comparing bids to decide the auction winner, and when computing the winner's quote price. This biasing rule intensifies competition by favoring the weaker TO (the one with a lower ex ante base margin). (For brevity I omit the details of running a biased auction; interested readers are referred to Duenyas et al. (2011) which describes running such auctions in detail.)

A key takeaway from this subsection is that the optimal mechanism resembles an auction for a quote. In practice, auctions for quotes are not commonplace; what is more common (as is seen in the court cases referenced in §4.1) is for the tier-two supplier to simply quote prices to their customers. Thus, for now the analyses of Sections 4.4 and 4.5.1 appear to more closely resemble how tier-two suppliers act in practice. However, since auctions for firm contracts are becoming prevalent in supply chains, it is plausible that an auction for a quote could be deployed by a tier-two supplier, especially as I have shown it to be the optimal mechanism. Also, although a tier-two supplier using an auction to select an exclusive customer is seldom observed, it is not uncommon to see a supplier causes a bidding war among its customers who want to buy its whole business to obtain exclusive access to an important technology it possesses. While the scenario of buyers bidding for the whole business of the supplier is different from my setting, the motivation behind it is actually very similar, namely the supplier can achieve maximal value by making the buyers compete for exclusive access to the important technology.

4.6 Extensions

Thus far in this chapter, to make the presentation simple, I assumed that only two TOs participate in the OEM's auction, and both depend on TT for the crucial component. Of course, in many cases in industry there are more than two TOs competing in the OEM's auction. Furthermore, it is possible that while some TOs rely on components from TT, others (e.g., a TO with its own internal component production capability) might not. This section examines such possibilities.

The base model I introduced in §4.3 has a three-tier structure: TT quotes prices for the component to TO_1 and TO_2, and TO_1 and TO_2 compete for the OEM's contract. In this section's more general model, the basic structure remains unchanged, but instead of two TOs, I assume there are n TOs (TO_i, $i = 1, ..., n$) who request price quotes from TT before they compete for the OEM's contract. In addition, I also assume there are m *outside competitors* CP_j, $j = 1, ..., m$ participating alongside the

n TOs in the OEM's auction, where the CPs do not require inputs of the component from TT. (When $n = 2$ and $m = 0$, the general model reduces to the base model.) Exactly as in Section 4.3, TO_i's cost to fulfill the OEM's contract is $x_i + y_i$ where x_i is the cost of purchasing the component from TT, and y_i is the cost of processing the component into the final product; y_i is TO_i's private information but it is common knowledge that y_i is a realization of random variable Y_i. In contrast, CP_j's *total* cost c_j includes its cost of the component and cost of processing it into the final product. c_j is a realization of random variable C_j whose distribution is common knowledge, and to ensure typical auction-theoretic properties (IFR), I assume that C_j has a positive, log-concave probability density; see Footnote 1. I assume all Y_i and C_j are independently distributed. Again, the OEM's reserve price, announced prior to the auction, is denoted by r. Note that under this setting, if a TO_i is to win the OEM's contract, it means $x_i + y_i < x_j + y_j$, $\forall j \neq i$ (TO_i has the lowest total cost among TOs), $x_i + y_i < r$ (TO_i can meet the reserve price), and $x_i + y_i < c_j$, $\forall j$ (TO_i has lower total cost than all CPs).

I find it convenient to define $R \doteq \min\{r, C_1, \ldots, C_m\}$. R represents a "sufficient statistic" capturing the OEM's reserve price and the outside competitors that all TOs face in the auction. Namely, to win the OEM's contract, TO_i's total cost must be the lowest among all TOs *and* lower than R. Thus, TO_i's base margin in the general model is given by $R - Y_i$, replacing $r - Y_i$ in the base model. With TO_i's base margin defined as $Z_i \doteq R - Y_i$ and with Z_i's support denoted by $[a_i, a_i + h_i]$, I now extend my analytical approach for the base model. A critical, complicating factor of course is that R is a random variable, whereas r was a constant.

As before, I consider three different approaches: Quoting Prices (QP), Quoting Equal Prices (QEP), and Optimal Mechanism (OM).

4.6.1 Quoting Prices

TT's problem is

$$\begin{aligned}
\max \quad & P(x_1, ..., x_n) \\
\text{where} \quad & P(x_1, ..., x_n) \doteq \sum_{i=1}^{n} x_i \Pr\{TO_i \text{ wins the contract}\} \\
& = \sum_{i=1}^{n} x_i \Pr(Z_i \geq x_i,\ Z_i - x_i > Z_j - x_j,\ j \neq i). \qquad (4.8)
\end{aligned}$$

Despite the increased complexity in TT's problem, I nonetheless show that many key results from §4.4 remain.

Theorem 4.5. *Propositions 4.1, 4.2, 4.3 and Theorem 4.1 still hold true for the general model.*

Therefore, my main insight holds true in the general model: TT uses two different types of strategies, risky or secure, and the TOs' profit potentials determine which will be used — a risky strategy when the profit potentials are low, and a secure strategy when they are high.

4.6.2 Quoting Equal Prices

Now I impose the identical quotes restriction. Denote TT's (identical) quote to all TOs by x. TT's problem is

$$\begin{aligned}
\max \quad & P(x) \\
\text{where} \quad & P(x) = x \Pr\{\text{Some } TO_i \text{ wins the contract}\} \\
& = x \Pr(x < \max_{i=1,\ldots,n} \{Z_i\}).
\end{aligned}$$

Section 4.5.1's main results continue to hold:

Theorem 4.6. *Propositions 4.4 and 4.5, and Theorem 4.3, hold for the general model, where I make the following changes to the results' statements to accommodate the general model: In Proposition 4.5 the condition*

$$x^* = \frac{1 - F_1(x^*)F_2(x^*)}{f_1(x^*)F_2(x^*) + F_1(x^*)f_2(x^*)}$$

is replaced by

$$x^* = \frac{1 - F(x^*)}{f(x^*)}$$

where $F(x)$ and $f(x)$ are the cdf and pdf of the random variable $\max_{i=1,\ldots,n}\{Z_i\}$, respectively; and in Proposition 4.5 and Theorem 4.3 the statement "Z_1 and Z_2 are i.i.d." is replaced by "the Y_is are i.i.d. and R has a log-concave pdf".

Therefore, when using the identical quotes approach, the insight that secure strategies are never optimal is preserved with the general model.

4.6.3 Optimal Mechanism

With the general model, a mechanism can be described with a set of rules

$$\{p_1(\check{y}_1, ..., \check{y}_n), ..., p_n(\check{y}_1, ..., \check{y}_n), x_1(\check{y}_1, ..., \check{y}_n), ..., x_n(\check{y}_1, ..., \check{y}_n)\}$$

where $p_i \geq 0$ is the transfer payment TO_i must pay TT upfront (before participating in the OEM's auction), and $x_i \geq 0$ is TT's price quote to TO_i, which he pays if and only if he places an order for the component. Both p_i and x_i are functions of all TO_is' reported processing costs $\check{y}_1, \ldots, \check{y}_n$ (where $\check{y}_i$ does not necessarily equal TO_i's real cost y_i).[5] Let the upper bound of R's support be denoted by u, let Y_i's cdf and pdf be denoted by G_i and pdf g_i, respectively, and define *virtual cost function*

[5] In §4.5.2, because r is known there is no difference between asking TOs to report base margin $\check{z}_i$ or cost $\check{y}_i$. In the general model however, since the TOs do not know R before the auction, it is only possible to ask them to report their costs as I do here.

$\omega_i(y) = y + \frac{G_i(y)}{g_i(y)}$.[6]

Theorem 4.7. *An optimal mechanism*

$$\{p_1^*(\check{y}_1, ..., \check{y}_n), ..., p_n^*(\check{y}_1, ..., \check{y}_n), x_1^*(\check{y}_1, ..., \check{y}_n), ..., x_n^*(\check{y}_1, ..., \check{y}_n)\}$$

is as follows[7]*: If* $\omega_i(\check{y}_i) \leq \min_{j \neq i}\{\omega_j(\check{y}_j), u\}$*, then* $p_i^* = E_R[\mathbf{1}_{\{\omega_i(\check{y}_i)<R\}}(R - \omega_i(\check{y}_i) + \check{y}_i - \omega_i^{-1}(\min_{j \neq i}\{\omega_j(\check{y}_j), u\}))]$ *and* $x_i^* = \frac{G_i(\check{y}_i)}{g_i(\check{y}_i)}$*, and* $\forall j \neq i$*,* $p_j^* = 0$ *and* $x_j^* > u - \check{y}_i$*; if* $\min_i\{\omega_i(\check{y}_i)\} > u$*, then* $p_i^* = 0$ *and* $x_i^* > u - \check{y}_i$*,* $\forall i$*.*

The optimal mechanism for the general model is a generalization of that for the base model. The main difference is that price quotes (which are paid only upon ordering) *and* upfront payments are required in the general model and cannot be interchanged, whereas they are interchangeable in the base model. What leads to this distinction is that in the base model the reserve price r is a constant, while its counterpart in the general model R is a random variable. With an uncertain R, the outcome of the OEM's auction is no longer fully predictable even when TT has all TOs' cost information, and thus contingent and upfront payments are not interchangeable. Nonetheless, the optimal mechanism's basic structure remains: TT backs only one TO, and the optimal mechanism can still be implemented via an auction-type approach in a similar spirit to that described in §4.5.2 (details are omitted for brevity).

4.7 Concluding Discussion

In this chapter I consider a tier-two supplier's optimal price-quoting strategy in the following model: The tier-two supplier gives price quotes to her customers, who are tier-one suppliers competing for an OEM's contract. Because (i) the tier-two

[6]Since I assume Y_i has log-concave probability density, due to Bagnoli and Bergstrom (2005) I know $\omega_i(\cdot)$ must be an increasing function.

[7]For readability I suppress the arguments $(\check{y}_1, ..., \check{y}_n)$ when writing p^* and x^*.

supplier's price quote will only realize as revenue when her customer places an order, and (ii) at most one of the competing tier-one suppliers will get the contract and hence need to place an order, this situation has an interesting and complicating implication: Even if the tier-two supplier can give price quotes to multiple customers, eventually at most one quote can results in revenue, and this may or may not be the quote that the tier-two supplier would most like to fulfill, i.e., the most lucrative quote. To the best of my knowledge, despite the widespread existence of multi-tier supply chains, my research is the first to identify and study the tier-two supplier's price-quoting problem. By studying decision-making at the second tier of the supply chain, I complement the extant procurement literature that focuses primarily on transactions within a single supply-chain interface (e.g., a buyer and her immediate suppliers).

In analyzing the problem, I identify two types of strategies deployed by the tier-two supplier: secure, where at least one of the tier-two supplier's price quotes will be exercised (i.e., the customer will place an order); and risky, where there is positive probability that none of the quotes will be exercised. I find that the optimal strategy is risky when the tier-one suppliers' profit potentials are low, and secure when the profit potentials are high. With a secure strategy, the tier-two supplier uses her customers for different strategic purposes (even when they are ex ante symmetric): She provides one customer a low quote to guarantee winning the business, and gambles with the others' quotes to obtain possible high profits, but at lower probabilities.

In analyzing the tier-two supplier's problem I allow her to strategically provide non-identical quotes to her customers. Non-identical quotes for the same item are seen in practice (cf. paragraph 9, §4.1), and my work has the potential to inform how to make the best price-quoting decisions and why. However, there is also evidence from practice that tier-one suppliers might perceive non-identical quotes as unfair. This motivates me to also study the tier-two supplier's strategy when she constrains herself to provide equal quotes to her customers. I find that with this constraint, the tier-two supplier facing ex ante identical tier-one customers will never use a secure strategy, and that such a constraint negatively affects her expected profit only when her customers' profit potentials are high.

Finally, while the results above are motivated by the existing practice of tier-two suppliers offering price quotes to their customers, it is of interest to ask what among all mechanisms would be theoretically optimal for the tier-two supplier. To answer this question, I employ mechanism design theory. The identified optimal mechanism has the following structure: The tier-two supplier solicits cost information from her customers, and with this information she strategically backs only one of the tier-one suppliers. This structure enables the optimal mechanism to be implemented by auctioning off a single quote among the tier-one suppliers. Although to my knowledge auctions for price quotes are not commonly deployed by tier-two suppliers, my analysis points to clear benefits of using an auction in such a scenario, and thus has the potential to inspire the use of such approaches in practice.

This research should be of interest to tier-two suppliers seeking to make better pricing decisions. I also hope that the research will spur further research into pricing decisions at various tiers of the supply chain, an important but under-studied area of procurement.

CHAPTER V

Conclusion

The three essays in this dissertation study three specific procurement-related problems in the presence of information asymmetry.

In the essay "Does Pooling Component Demands when Sourcing Lead to Higher Profits?" (Chapter II), I find that pooling purchases from a strategic sole supplier of a component used in multiple products with uncertain demands can in fact reduce buyers' profits because of reduced information rent, which is not considered in the traditional pooling literature. One important insight from this essay is that pooling which reduces demand uncertainty faced by the buyers can also reduce their informational advantage over the supplier, enabling the supplier to extract more of the buyers' profits. This insight highlights the sole supplier's strategic behaviors as a dominant factor in determining the buyers' profits, and warns procurement managers that naively applying the received wisdom about pooling — which ignores the supplier's possible strategic responses — may lead to an unwanted outcome.

The essay "Simple Auctions for Supply Contracts" (Chapter III) shows that a buyer facing several candidate suppliers who have private cost information can use a simple modified open-descending auction as an optimal mechanism. The major advantages of this mechanism over the mechanisms previously proposed in the literature are that this mechanism is very familiar to the suppliers, and the suppliers only need to use extremely simple strategies in this mechanism. The fact that the suppliers can use very simple strategies in a familiar environment means that they are more apt to participate in such a mechanism. The modified open-descending auction can also be extended to allow concave production costs and ex ante asymmetric suppliers. All these features suggest that the mechanism has great potential for practical use.

The essay "Price-Quoting Strategies of a Tier-Two Supplier" (Chapter IV) is one of the first research works to study the price-quoting strategies of a tier-two supplier, whose customers (the tier-one suppliers) will compete downstream for an indivisible contract. The tier-two supplier's price-quoting problem is complicated by the fact that at most one quote will generate revenue, and she does not know which one will, if any. I show that the tier-two supplier's optimal strategy has the following structure: If the business is lucrative enough, she uses a secure strategy whereby at least one of the tier-two suppliers' price quotes will be exercised. Otherwise, she uses a risky strategy whereby there is positive probability that none of the quotes will be exercised. In addition, a secure strategy always consists of a higher quote for one customer and a lower quote for the other(s), even when the customers are ex ante symmetric. This means asymmetric quotes can be optimal for ex ante symmetric customers. The intuition is that the tier-two supplier can use her customers for different strategic purposes: She provides one customer a low quote to guarantee winning the business, and gambles with the others' quotes to obtain possible high profits, but at lower probabilities. I also show that if the tier-two supplier is not limited by giving out fixed-price quotes, then she can maximize her expected profit by making her customers compete for exclusive access to her product (auctioning a single quote). This essay complements the extant procurement literature that focuses primarily on transactions at a single supply-chain interface (e.g., a buyer and its immediate suppliers), and should be of interest to tier-two suppliers seeking to make better pricing decisions.

This dissertation has the potential to help procurement managers understand certain business situations more clearly and make better decisions. In particular, it highlights the impact of information asymmetry on procurement, and suggests strategies to tackle resulting challenges. More broadly, this dissertation adds to the nascent and growing Operations Management scholarship on procurement, and contributes to the field's general understanding of supply chain management.

APPENDICES

APPENDIX A

Proofs of Chapter II

Proof of Proposition 2.1. Let the first derivative of the supplier's objective function with respect to Q_{θ_j} be denoted by

$$FD_{\theta_j}(Q_{\theta_j}) \doteq r\Pr(\Theta \geq \theta_j)[\lambda(\theta_j)(1-q) - F(Q_{\theta_j} - \mu_{\theta_j}) + (1-\lambda(\theta_j))F(Q_{\theta_j} - \mu_{\theta_{j+1}})],$$

and the second derivative be denoted by

$$SD_{\theta_j}(Q_{\theta_j}) \doteq r\Pr(\Theta \geq \theta_j)[(1-\lambda(\theta_j))f(Q_{\theta_j} - \mu_{\theta_{j+1}}) - f(Q_{\theta_j} - \mu_{\theta_j})].$$

It is straightforward to observe that $FD_{\theta_j}(Q_{\theta_j})$ is positive when Q_{θ_j} is sufficiently small (so that $F(Q_{\theta_j} - \mu_{\theta_j}) = 0$), and negative when Q_{θ_j} is sufficiently large (so that $F(Q_{\theta_j} - \mu_{\theta_{j+1}}) = 1$). Assume f has support over $(-s, s)$ and satisfies all four conditions in Proposition 2.1. Next, for each of three possible cases I will show that $FD_{\theta_j}(Q_{\theta_j})$ is first positive then negative, thus the supplier's objective function is unimodal and the unique solution to (2.4) characterizes the global optimal solution.

Case 1: $\mu_{\theta_{j+1}} - \mu_{\theta_j} \geq 2s$ $(\mu_{\theta_j} - s < \mu_{\theta_j} + s \leq \mu_{\theta_{j+1}} - s < \mu_{\theta_{j+1}} + s)$.

In this case, it is clear that $FD_{\theta_j}(Q_{\theta_j})$ is decreasing over $(\mu_{\theta_j} - s, \mu_{\theta_j} + s)$, constant over $[\mu_{\theta_j} + s, \mu_{\theta_{j+1}} - s]$, and increasing over $(\mu_{\theta_{j+1}} - s, \mu_{\theta_{j+1}} + s)$. Yet

I already know $FD_{\theta_j}(Q_{\theta_j}) > 0$ when $Q_{\theta_j} = \mu_{\theta_j} - s$ and $FD_{\theta_j}(Q_{\theta_j}) < 0$ when $Q_{\theta_j} = \mu_{\theta_{j+1}} + s$. Thus it is clear that $FD_{\theta_j}(Q_{\theta_j})$ is first positive then negative.

Case 2: $s < \mu_{\theta_{j+1}} - \mu_{\theta_j} < 2s$ $(\mu_{\theta_j} - s < \mu_{\theta_j} < \mu_{\theta_{j+1}} - s < \mu_{\theta_j} + s < \mu_{\theta_{j+1}} < \mu_{\theta_{j+1}} + s)$.

In this case, it is clear that $FD_{\theta_j}(Q_{\theta_j})$ is decreasing over $(\mu_{\theta_j} - s, \mu_{\theta_{j+1}} - s)$, and increasing over $(\mu_{\theta_j} + s, \mu_{\theta_{j+1}} + s)$. Between $\mu_{\theta_{j+1}} - s$ and $\mu_{\theta_j} + s$, the second derivative $SD_{\theta_j}(Q_{\theta_j})$ equals

$$r \Pr(\Theta \geq \theta_j)[(1 - \lambda(\theta_j))f(Q_{\theta_j} - \mu_{\theta_{j+1}}) - f(Q_{\theta_j} - \mu_{\theta_j})].$$

By Condition 3, $f(Q_{\theta_j} - \mu_{\theta_{j+1}})$ is non-decreasing and $f(Q_{\theta_j} - \mu_{\theta_j})$ is non-increasing, therefore $SD_{\theta_j}(Q_{\theta_j})$ is non-decreasing. Because of this, I know $FD_{\theta_j}(Q_{\theta_j})$ is convex over $(\mu_{\theta_{j+1}} - s, \mu_{\theta_j} + s)$. If $FD_{\theta_j}(Q_{\theta_j}) < 0$ at $\mu_{\theta_{j+1}} - s$, then by convexity of $FD_{\theta_j}(Q_{\theta_j})$, $FD_{\theta_j}(Q_{\theta_j})$ must be negative over $(\mu_{\theta_{j+1}} - s, \mu_{\theta_j} + s)$. If $FD_{\theta_j}(Q_{\theta_j}) > 0$ at $\mu_{\theta_{j+1}} - s$, then by convexity of $FD_{\theta_j}(Q_{\theta_j})$, $FD_{\theta_j}(Q_{\theta_j})$ crosses zero only once. In both cases, it is clear that $FD_{\theta_j}(Q_{\theta_j})$ is first positive then negative.

Case 3: $\mu_{\theta_{j+1}} - \mu_{\theta_j} \leq s$ $(\mu_{\theta_j} - s < \mu_{\theta_{j+1}} - s \leq \mu_{\theta_j} < \mu_{\theta_{j+1}} \leq \mu_{\theta_j} + s < \mu_{\theta_{j+1}} + s)$.

In this case, it is clear that $FD_{\theta_j}(Q_{\theta_j})$ is decreasing over $(\mu_{\theta_j} - s, \mu_{\theta_j})$, convex over $(\mu_{\theta_j}, \mu_{\theta_{j+1}})$ (similar to Case 2), and increasing over $(\mu_{\theta_j} + s, \mu_{\theta_{j+1}} + s)$. I now characterize $FD_{\theta_j}(Q_{\theta_j})$ over $(\mu_{\theta_{j+1}}, \mu_{\theta_j} + s)$. Recall that the second derivative $SD_{\theta_j}(Q_{\theta_j})$ equals

$$r \Pr(\Theta \geq \theta_j)[(1 - \lambda(\theta_j))f(Q_{\theta_j} - \mu_{\theta_{j+1}}) - f(Q_{\theta_j} - \mu_{\theta_j})].$$

By Condition 4, $f(Q_{\theta_j} - \mu_{\theta_j})/f(Q_{\theta_j} - \mu_{\theta_{j+1}})$ is non-increasing in Q_{θ_j}. Thus I know over $(\mu_{\theta_{j+1}}, \mu_{\theta_j} + s)$, $SD_{\theta_j}(Q_{\theta_j})$ cannot first be positive then negative; it may only be always positive, or always negative, or first negative then positive.

When $SD_{\theta_j}(Q_{\theta_j})$ is always positive over $(\mu_{\theta_{j+1}}, \mu_{\theta_j}+s)$, $FD_{\theta_j}(Q_{\theta_j})$ is increasing over $(\mu_{\theta_{j+1}}, \mu_{\theta_{j+1}}+s)$. Following the same proof as in Case 2 I know $FD_{\theta_j}(Q_{\theta_j})$ must be first positive then negative.

When $SD_{\theta_j}(Q_{\theta_j})$ is negative over $(\mu_{\theta_{j+1}}, t)$ and positive over $[t, \mu_{\theta_j}+s)$ where $\mu_{\theta_{j+1}} < t \leq \mu_{\theta_j}+s$, I know the right-derivative of $FD_{\theta_j}(Q_{\theta_j})$ at $\mu_{\theta_{j+1}}$ is negative. Condition 2 then guarantees that the left-derivative of $FD_{\theta_j}(Q_{\theta_j})$ at $\mu_{\theta_{j+1}}$ is also negative. Combining this and the fact that $FD_{\theta_j}(Q_{\theta_j})$ is convex over $(\mu_{\theta_j}, \mu_{\theta_{j+1}})$, I know $FD_{\theta_j}(Q_{\theta_j})$ must also be decreasing over $(\mu_{\theta_j}, \mu_{\theta_{j+1}})$. As a result, $FD_{\theta_j}(Q_{\theta_j})$ is decreasing over $(\mu_{\theta_j}-s, t)$ and increasing over $(t, \mu_{\theta_{j+1}}+s)$, so it is first positive then negative. □

Proof of Proposition 2.2. To prove that the violation of the MCs can only occur in the two cases mentioned in the proposition, I need to show the solution to (2.8)-(2.11) (ignoring MCs (2.7f)) always satisfies $Q^D_{ll} < Q^D_{hl} < Q^D_{hh}$.

As discussed in §2.3.1, the participation and incentive compatibility constraints are always binding at optimality. First plug the binding constraints (2.7b)-(2.7e) into (2.7a), then take the first derivative of the objective function with respect to Q^D_{hl}. Doing so yields

$$FD_{hl}(Q) \doteq \left.\frac{\partial[(2.7\mathrm{a})]}{\partial Q^D_{hl}}\right|_{Q^D_{hl}=Q} = p^D_{hl}(1-q) - (p^D_{hl}+p^D_{hh})F^D(Q-\mu^D_{hl}) + p^D_{hh}F^D(Q-\mu^D_{hh}).$$

Due to Proposition 2.1, I know that $FD_{hl}(Q) = 0$ always has an interior solution. Recall that from (2.11) I have $Q^D_{hh} = \mu^D_{hh} + (F^D)^{-1}(1-q)$. If I evaluate $FD_{hl}(Q)$ at $Q = Q^D_{hh} = \mu^D_{hh} + (F^D)^{-1}(1-q)$, I can easily see that $FD_{hl}(Q^D_{hh}) < 0$. This means the solution to $FD_{hl}(Q) = 0$, Q^D_{hl}, is smaller than Q^D_{hh}.

I next show $Q^D_{ll} < Q^D_{hl}$. Notice that

$$\mu^D_{lh} - \mu^D_{ll} = (\mu^a_l + \mu^b_h) - (\mu^a_l + \mu^b_l) = (\mu^a_h + \mu^b_h) - (\mu^a_h + \mu^b_l) = \mu^D_{hh} - \mu^D_{hl}.$$

Define $\delta \doteq \mu^D_{lh} - \mu^D_{ll} = \mu^D_{hh} - \mu^D_{hl}$. Using Equations (2.8), (2.10), (2.12) and (2.13), and applying the definitions of δ and γ, (2.8) and (2.10) can be rewritten as

$$\lambda(\theta)(1-q) = F(s) - (1-\lambda(\theta))F(s-\delta), \tag{A.1}$$

and (2.12) and (2.13) can be rewritten as

$$\pi_{\theta'} - \pi_\theta = r\int_{s-\delta}^{s} F(x)\mathrm{d}x, \tag{A.2}$$

for types $(\theta, \theta') = (ll, lh)$ and (hl, hh), where $s = Q_\theta - \mu_\theta$, $\lambda(ll) = p^D_{ll}$ for (2.8), and $\lambda(hl) = p^D_{hl}/(p^D_{hl} + p^D_{hh})$ for (2.10). Taking the total derivative of (A.1) with respect to $\lambda(\theta)$ yields

$$\begin{aligned} 1-q &= \frac{\mathrm{d}s}{\mathrm{d}\lambda(\theta)}[f(s) - (1-\lambda(\theta))f(s-\delta)] + F(s-\delta) \\ \Longrightarrow \frac{\mathrm{d}s}{\mathrm{d}\lambda(\theta)} &= \frac{1-q-F(s-\delta)}{f(s)-(1-\lambda(\theta))f(s-\delta)}. \end{aligned} \tag{A.3}$$

Because the objective function is concave at the global optimal solution (by Proposition 2.1), I know the denominator of (A.3) is positive. The numerator is also positive, because Equation (A.1) and the fact that $F(s) > F(s-\delta)$ together imply

$$\lambda(\theta)(1-q) = F(s) - (1-\lambda(\theta))F(s-\delta) > \lambda(\theta)F(s-\delta).$$

Therefore, I know s is increasing in $\lambda(\theta)$. Notice that

$$\lambda(ll) = p^D_{ll} = p^a_l p^b_l < \lambda(hl) = \frac{p^D_{hl}}{p^D_{hl} + p^D_{hh}} = \frac{p^a_h p^b_l}{p^a_h(p^b_l + p^b_h)} = p^b_l.$$

As a result, $Q^D_{ll} - \mu^D_{ll} = Q^D_{ll} - \mu^a_l - \mu^b_l < Q^D_{hl} - \mu^a_h - \mu^b_l = Q^D_{hl} - \mu^D_{hl}$, which immediately yields $Q^D_{ll} < Q^D_{hl}$.

Ultimately I wish to solve optimization problem (2.7a)-(2.7f). Thus far I have examined the solution obtained with MCs (2.7f) ignored. If the solution satisfies (2.7f), then it is the optimal solution to (2.7a)-(2.7f). If it violates (2.7f), which as argued above can only be true if $Q^D_{ll} > Q^D_{lh}$ or $Q^D_{lh} > Q^D_{hl}$, then I must revise my solution. I now show how to revise (2.8)-(2.10) to obtain the true optimal solution. I show this for the case when $Q^D_{ll} \leq Q^D_{lh}$ is violated; the other case is similar.

Suppose when ignoring the MCs the FOC solution is such that $Q^D_{ll} > Q^D_{lh}$. Consider a function $I(x,y) = G(x) + H(y)$ such that both G and H are unimodal, and suppose the x_0 that maximizes G and the y_0 that maximizes H are such that $x_0 > y_0$. If I impose the requirement that $x \leq y$ and find the optimal $(x^*, y^*) \doteq \{(x,y)|x \leq y,\ G(x)+H(y) \geq G(x')+H(y'),\ \forall\ x' \leq y'\}$, it is easy to see (by a contradiction argument) that $y_0 \leq x^* = y^* \leq x_0$. Since by Proposition 2.1 my maximization problem is decomposable, and unimodal in both Q^D_{ll} and Q^D_{lh}, this result applies to my problem. I thus know that adding back the violated constraint $Q^D_{ll} \leq Q^D_{lh}$ will result in a new optimal solution for which $Q^D_{ll} = Q^D_{lh} = \underline{Q}$ for some $\underline{Q}$. Differentiating the objective function with this equality enforced and setting the derivative to zero yields FOC (2.14). Furthermore, since Equation (2.14) is a specific instance of Equation (2.4), I can apply Proposition 2.1 to conclude that the FOC solution $\underline{Q}$ will be the unique global maximizer of the supplier's objective function. $\square$

Proof of Lemma 2.1. By (2.5),

$$\pi_{\theta'} - \pi_\theta = r \int_{Q_\theta - \mu_{\theta'}}^{Q_\theta - \mu_\theta} F(x)\mathrm{d}x,$$

so to show $\pi_{\theta'} - \pi_\theta$ increases in $\lambda(\theta)$, it is sufficient to show Q_θ increases in $\lambda(\theta)$. I focus on type ll; arguments for the other types are similar. Depending on whether the MC for type ll in (2.7f) is slack or binding, the equation that determines Q_{ll} is either (2.8) or (2.14), respectively. I denote the solution to FOC (2.8) by $Q^*(2.8)$, and the solution to FOC (2.14) by $Q^*(2.14)$. Similarly denote the solution to FOC (2.9) by $Q^*(2.9)$.

Suppose when $\lambda(ll) = \lambda_1$, $Q^*(2.8) < Q^*(2.9)$ (so the MC for type ll in (2.7f) is slack and $Q_{ll}^D = Q^*(2.8)$). To prove Lemma 2.1 for this case, I need to show three things: (i) While the MC for type ll remains slack, when $\lambda(ll)$ increases, $Q^*(2.8)$ increases. (ii) If increasing $\lambda(ll)$ to some point $\lambda(ll) = \lambda_2$ causes the MC for type ll to become binding, then I switch my consideration to the parameter $\underline{\lambda}$ and need to show that when $\underline{\lambda}$ increases, $Q^*(2.14)$ increases. (iii) At the "switch point" $\lambda(ll) = \lambda_2$ between the non-bunching and bunching cases, I have continuity, namely $Q_{ll}^D = Q^*(2.8) = Q^*(2.14)$.

(i)-(ii) follow from the fact that both (2.8) and (2.14) fit the structure of (2.4), so as argued following Equation (A.1), I know $Q^*(2.8)$ and $Q^*(2.14)$ are increasing in $\lambda(ll)$ and $\underline{\lambda}$, respectively. I next prove (iii). $Q^*(2.8)$ increases in $\lambda(ll)$. Since $Q^*(2.9)$ remains constant (since I am increasing $\lambda(ll)$ in a way that keeps $\lambda(lh)$ fixed), if type ll's MC becomes binding at $\lambda(ll) = \lambda_2$, then $Q^*(2.8) = Q^*(2.9)$ at $\lambda(ll) = \lambda_2$. Furthermore, as argued following Equation (A.1), I know that $Q^*(2.8) \le Q^*(2.14) \le Q^*(2.9)$, which in turn implies that $Q_{ll}^D = Q^*(2.8) = Q^*(2.14)$ at $\lambda(ll) = \lambda_2$.

The above arguments proved the lemma when increasing $\lambda(ll)$ took it from non-bunching to bunching. The case where increasing $\underline{\lambda}$ takes it from bunching to non-bunching can be treated analogously. Iteratively applying these results yields the lemma over any interval of $\lambda(ll)$ or $\underline{\lambda}$. $\square$

Proof of Lemma 2.2. Regardless of whether there is bunching, the equations characterizing the optimal solution and information rent always fit the structure of (2.4) and (2.5), respectively. Therefore I only need to prove Lemma 2.2 for the general equations (2.4) and (2.5).

Consider two pairs of mean demands $\mu_\theta, \mu_{\theta'}$ and $\mu'_\theta, \mu'_{\theta'}$ such that $\mu_{\theta'} - \mu_\theta = \mu'_{\theta'} - \mu'_\theta = \delta$. Notice that (2.4) and (2.5), with either pair of mean demands, can be rewritten as (A.1) and (A.2) where $s = Q_\theta - \mu_\theta$ or $s = Q'_\theta - \mu'_\theta$. Therefore $\mu_\theta, \mu_{\theta'}$ and

$\mu'_{\theta}, \mu'_{\theta'}$ yield the same information rent $\pi_{\theta'} - \pi_{\theta}$. This shows that μ_{θ} and $\mu_{\theta'}$ determine information rent only through $\delta = \mu_{\theta'} - \mu_{\theta}$. The inequality $\pi_{\theta'} - \pi_{\theta} \leq r\delta$ follows because $F(x) \leq 1$ for all x. $\square$

Proof of Lemma 2.3. I prove the result for type $\theta = ll$ only; the proofs for other types are analogous. The proof's structure is similar to that of Lemma 2.1's proof. Suppose when $\gamma = \gamma_1$, $Q^*(2.8) < Q^*(2.9)$ (so the MC for type ll in (2.7f) is slack, meaning $Q^D_{ll} = Q^*(2.8)$). To prove Lemma 2.3 for this case, I need to show three things. (i) While the MC for type ll remains slack, when γ increases $\pi_{lh} - \pi_{ll}$ also increases. (ii) If increasing γ to some point $\gamma = \gamma_1$ causes the MC for type ll to become binding, then I switch my consideration to the bunching case meaning $Q^D_{ll} = Q^*(2.9)$, and I need to show that when γ continues to increase, $\pi_{lh} - \pi_{ll}$ increases. (iii) At the "switch point" $\gamma = \gamma_1$ between the non-bunching and bunching cases, I have continuity, namely $Q^D_{ll} = Q^*(2.8) = Q^*(2.14)$ (which implies that the information rent gap $\pi_{lh} - \pi_{ll}$ is also continuous at the switch point).

I begin by addressing (iii). The arguments are precisely the same as that used to prove point (iii) in the proof of Lemma 2.1, except with γ's in the role of λ's. Therefore, it suffices to show (i)-(ii).

Notice that the incremental information rent $\pi_{lh} - \pi_{ll}$ is determined by (2.12) together with (2.8) or (2.14) (corresponding to non-bunching or bunching, respectively). In both cases the incremental information rent can be expressed in the following general way, where θ and θ' play the roles of ll and lh, respectively:

$$\lambda(\theta)(1-q) = F_{(\gamma)}(\tilde{s}) - (1-\lambda(\theta))F_{(\gamma)}(\tilde{s}-\delta), \tag{A.4}$$

$$\tilde{\pi}_{\theta'} - \tilde{\pi}_{\theta} = r\int_{\tilde{s}-\delta}^{\tilde{s}} F_{(\gamma)}(x)\mathrm{d}x, \tag{A.5}$$

where $\delta \doteq \mu_{\theta'} - \mu_{\theta}$. Suppose $\gamma > 1$. To show that the incremental information rent increases in γ, it suffices to show that $\tilde{\pi}_{\theta'} - \tilde{\pi}_{\theta}$ is greater than $\pi_{\theta'} - \pi_{\theta}$, which I will define to be the incremental information rent when $\gamma = 1$; namely

$$\pi_{\theta'} - \pi_\theta = r \int_{s-\delta}^{s} F_{(1)}(x)\mathrm{d}x$$

where

$$\lambda(\theta)(1-q) = F_{(1)}(s) - (1-\lambda(\theta))F_{(1)}(s-\delta).$$

To this end, let s' be the root of the following equation:

$$\lambda(\theta)(1-q) = F_{(\gamma)}(s') - (1-\lambda(\theta))F_{(\gamma)}(s'-\gamma\delta), \tag{A.6}$$

and notice that $s' = \gamma s$. I now show that $s' \leq \tilde{s}$. Because s' is the solution to (A.6), I have

$$F_{(\gamma)}(s') - (1-\lambda(\theta))F_{(\gamma)}(s'-\delta) \leq F_{(\gamma)}(s') - (1-\lambda(\theta))F_{(\gamma)}(s'-\gamma\delta) = \lambda(\theta)(1-q).$$

By the proof of Proposition 2.1, I know that when $\tilde{s}$ increases, the RHS of (A.4) is first smaller then greater than $\lambda(\theta)(1-q)$. Therefore, I know the solution to (A.4), $\tilde{s}$, must be greater than s'.

Now I show $\tilde{\pi}_{\theta'} - \tilde{\pi}_\theta \geq \pi_{\theta'} - \pi_\theta$. In fact, since $\tilde{s} \geq s'$, I have

$$\begin{aligned}
\tilde{\pi}_{\theta'} - \tilde{\pi}_\theta &= r\int_{\tilde{s}-\delta}^{\tilde{s}} F_{(\gamma)}(x)\mathrm{d}x \geq r\int_{s'-\delta}^{s'} F_{(\gamma)}(x)\mathrm{d}x = r\int_0^{\delta} F_{(\gamma)}(s'-t)\mathrm{d}t \\
&= \frac{1}{\gamma} r \int_0^{\gamma\delta} F_{(\gamma)}(s'-k/\gamma)\mathrm{d}k \geq \frac{1}{\gamma} r \int_0^{\gamma\delta} F_{(\gamma)}(s'-k)\mathrm{d}k \\
&= \frac{1}{\gamma} r \int_{s'-\gamma\delta}^{s'} F_{(\gamma)}(y)\mathrm{d} = \frac{1}{\gamma}\gamma(\pi_{\theta'} - \pi_\theta) = \pi_{\theta'} - \pi_\theta,
\end{aligned}$$

where $k = \gamma t$ and $y = s' - k$. This proves (i) and (ii).

The above arguments prove the lemma when increasing γ took it from non-bunching to bunching. The opposite case where increasing γ takes it from bunching to non-bunching can be treated analogously. Iteratively applying these results yields the lemma over any interval of γ. $\square$

Proof of Theorem 2.1. Before proving the theorem I introduce a technique that will appear multiple times in this and later proofs. I first note that the incremental information rent $\pi_{\theta'} - \pi_\theta$ is continuous in $\lambda(\theta)$, σ, and δ^θ. (Continuity in $\lambda(\theta)$ and γ was shown in the course of proving Lemmas 2.1 and 2.3. Continuity in δ^θ comes from similar arguments.)

By the definition of continuity, when I establish a property about the size of the incremental information rent at a certain point in the parameter space, it will also hold in a sufficiently small neighborhood around that point. With this technique, I can show that my desired properties hold in a neighborhood of a particular point by showing they hold *at* the particular point.

I begin with the first part of the theorem. When $\sigma^a = 0$, $\sqrt{(\sigma^a)^2 + (\sigma^b)^2} = \sigma^b$. Notice that OEM B's demand of type h has gap δ^b, variability σ^b, and $\lambda(l) = p_l^b$. In comparison, OEM D's demand of type lh has gap δ^b, variability $\sqrt{(\sigma^a)^2 + (\sigma^b)^2} = \sigma^b$ and $\lambda(ll) = p_l^a p_l^b < p_l^b = \lambda(l)$. Thus OEM B's type-h demand has the same gap and variability as OEM D's type-lh demand, but has higher λ. Due to Lemma 2.1, I know $\pi_h^b > \pi_{lh}^D$, thus OEMs A and B have higher combined information rent than D. By continuity, this is also true when σ^a is sufficiently small.

I now prove the theorem's second part. When σ^a decreases, OEM B's information rent will not change, nor will the low-type OEM A's information rent (it is constant at zero). Thus the combined information rents of A and B do not change. However, the information rent of D decreases because when σ^a decreases, the variability of the combined demand $\sqrt{(\sigma^a)^2 + (\sigma^b)^2}$ also decreases and, by Lemma 2.3, so does D's information rent. Similarly, when p_h^a increases, OEMs A and B's information rents are not affected, but $\lambda(ll) = p_l^a p_l^b$ decreases and (by Lemma 2.1) this causes D's information rent to decrease. $\square$

Proof of Theorem 2.2. I begin with the theorem's first part. When $\delta^a = \delta^b$, types lh and hl are indistinguishable, so bunching occurs between them, and consequently $\pi^D_{hl} = \pi^D_{lh}$. Therefore the comparison of π^a_h versus π^D_{hl} is essentially the comparison of π^a_h versus π^D_{lh}. On the other hand, when $\sigma^b = 0$, $\sqrt{(\sigma^a)^2 + (\sigma^b)^2} = \sigma^a$. Notice that OEM A with type-h demand has variability σ^a, gap δ^a and $\lambda(l) = p^a_l$. In comparison, OEM D of type-lh demand has the same variability $\sqrt{(\sigma^a)^2 + (\sigma^b)^2} = \sigma^a$, gap $\delta^b = \delta^a$ and lower $\lambda(ll) = p^a_l p^b_l < p^a_l = \lambda(l)$. Due to Lemma 2.1 I know $\pi^a_h > \pi^D_{lh}$, thus OEMs A and B have higher combined information rent than D. By continuity this is also true when σ^b is sufficiently small and δ^a is sufficiently close to δ^b.

The second part's proof is similar to that of Theorem 2.1. As σ^b decreases, the information rent of OEM A's type-h demand does not change, but OEM D's information rent decreases. Also, as p^b_h increases, π^a does not change, but π^D_{lh} decreases since $\lambda(ll) = p^a_l p^b_l$ decreases. □

Proof of Theorem 2.3. The theorem assumes σ is sufficiently large. However, due to the scalability of the information rents in δ and σ, I can replace this condition by δ is sufficiently small. To see why, note that if I choose scalar $\kappa > 0$ and set

$$\pi_{\theta'} - \pi_\theta = r \int_{s-\delta}^{s} F_{(\sigma)}(x)\mathrm{d}x$$

where

$$\lambda(\theta)(1-q) = F_{(\sigma)}(s) - (1-\lambda(\theta))F_{(\sigma)}(s-\delta)$$

and

$$\pi'_{\theta'} - \pi'_\theta = r \int_{s'-\kappa\delta}^{s'} F_{(\kappa\sigma)}(x)\mathrm{d}x$$

where

$$\lambda(\theta)(1-q) = F_{(\kappa\sigma)}(s') - (1-\lambda(\theta))F_{(\kappa\sigma)}(s'-\kappa\delta),$$

then $s' = \kappa s$, and hence $\pi'_{\theta'} - \pi'_\theta = \kappa(\pi_{\theta'} - \pi_\theta)$. Thus, to study the relative sizes of OEM A's, B's and D's information rents with σ large, I can analyze the related scaled-down setting (where δ is small) which reduces the magnitude of the information rents but leaves their relative sizes unchanged. Consequently, for the remainder of the proof I will instead use the condition $\delta \to 0$.

When the demands for a and b are symmetric, bunching occurs between types lh and hl. Thus to show that OEMs A and B continue to outperform D when p_l decreases, I need to show that $\pi_h^a + \pi_h^b$'s derivative with respect to p_l is smaller than $(\pi_{hh}^D - \pi_{lh}^D) + \pi_{lh}^D$'s derivative with respect to p_l. I now establish a property that facilitates my comparison of the derivatives.

Due to Proposition 2.2's proof,

$$\frac{\mathrm{d}Q_\theta}{\mathrm{d}\lambda(\theta)} = \frac{1 - q - F(Q_\theta - \mu_{\theta'})}{f(Q_\theta - \mu_\theta) - (1 - \lambda(\theta))f(Q_\theta - \mu_{\theta'})}. \tag{A.7}$$

My first step is to show that

$$\frac{\mathrm{d}Q_\theta}{\mathrm{d}\lambda(\theta)} \to \delta/\lambda^2(\theta) \quad \text{as} \quad \delta \to 0. \tag{A.8}$$

I begin by noting that a Taylor expansion of F about the point $Q_\theta - \mu_{\theta'}$, together with the fact that $\mu_{\theta'} = \mu_\theta + \delta$, yields

$$F(Q_\theta - \mu_\theta) = F(Q_\theta - \mu_{\theta'}) + \delta f(Q_\theta - \mu_{\theta'}) + o(\delta). \tag{A.9}$$

Recall FOC (2.4):

$$\lambda(\theta)(1 - q) = F(Q_\theta - \mu_\theta) - (1 - \lambda(\theta))F(Q_\theta - \mu_{\theta'}).$$

With (A.9), this implies that as $\delta \to 0$ the RHS of (2.4) approaches $\lambda(\theta)F(Q_\theta - \mu_{\theta'}) +$

$\delta f(Q_\theta - \mu_{\theta'})$, hence

$$1 - q - F(Q_\theta - \mu_{\theta'}) \to \frac{\delta f(Q_\theta - \mu_{\theta'})}{\lambda(\theta)} \quad \text{as} \quad \delta \to 0. \tag{A.10}$$

Since the RHS of (A.10) approaches zero as $\delta \to 0$, I also have

$$Q_\theta - \mu_{\theta'} \to F^{-1}(1-q) \quad \text{as} \quad \delta \to 0. \tag{A.11}$$

By Lipschitz continuity of f, combining (A.11) and (A.10) yields

$$1 - q - F(Q_\theta - \mu_{\theta'}) \to \frac{\delta f(F^{-1}(1-q))}{\lambda(\theta)} \quad \text{as} \quad \delta \to 0. \tag{A.12}$$

I now apply these insights to the RHS of (A.7). Since $\mu_\theta + \delta = \mu_{\theta'}$, (A.11) implies that the denominator of (A.7)'s RHS approaches $\lambda(\theta) f(F^{-1}(1-q))$ as $\delta \to 0$. Additionally, (A.12) implies that the numerator of (A.7)'s RHS approaches $\delta f(F^{-1}(1-q))/\lambda(\theta)$ as $\delta \to 0$. Upon simplifying, I recover (A.8).

Furthermore, $\frac{\mathrm{d}(\pi_{\theta'} - \pi_\theta)}{\mathrm{d}Q_\theta} = r(F(Q_\theta - \mu_\theta) - F(Q_\theta - \mu_{\theta'}))$ (by Equation (2.5)) $\to r\delta f(F^{-1}(1-q))$ as $\delta \to 0$ (by combining (A.9) and (A.11)). Combining this and (A.8) yields

$$\frac{\mathrm{d}(\pi_{\theta'} - \pi_\theta)}{\mathrm{d}\lambda(\theta)} = \frac{\mathrm{d}(\pi_{\theta'} - \pi_\theta)}{\mathrm{d}Q_\theta} \frac{\mathrm{d}Q_\theta}{\mathrm{d}\lambda(\theta)} \to r\delta^2 f(F^{-1}(1-q))/\lambda^2(\theta) \quad \text{as} \quad \delta \to 0. \tag{A.13}$$

Let $F_{(\sigma)}$ denote the error distribution of e^i. Then $F_{(\sqrt{2}\sigma)}$ denotes the error distribution of e^D. Generally speaking, for any $\kappa > 0$, since $F_{(\kappa\sigma)}(\kappa s) = F_{(\sigma)}(s)$, I know

$$\left.\frac{\mathrm{d}}{\mathrm{d}x} F_{(\kappa\sigma)}(x)\right|_{x=\kappa s} = \left.\frac{1}{\kappa} \frac{\mathrm{d}}{\mathrm{d}x} F_{(\sigma)}(x)\right|_{x=s}$$

and

$$F^{-1}_{(\kappa\sigma)}(1-q) = \kappa F^{-1}_{(\sigma)}(1-q).$$

Combining these equalities yields

$$f_{(\kappa\sigma)}(F^{-1}_{(\kappa\sigma)}(1-q)) = \frac{1}{\kappa} f_{(\sigma)}(F^{-1}_{(\sigma)}(1-q)). \tag{A.14}$$

I now use (A.13) and (A.14) to prove the theorem. By (A.13),

$$\frac{\mathrm{d}\pi^a_h}{\mathrm{d}\lambda(l)} = \frac{\mathrm{d}\pi^b_h}{\mathrm{d}\lambda(l)} \to r\delta^2 f_{(\sigma)}(F^{-1}_{(\sigma)}(1-q))/\lambda^2(l),$$

and by (A.14),

$$\frac{\mathrm{d}\pi^D_{lh}}{\mathrm{d}\lambda(ll)} \to r\delta^2 f_{(\sqrt{2}\sigma)}(F^{-1}_{(\sqrt{2}\sigma)}(1-q))/\lambda^2(ll) = \frac{1}{\sqrt{2}} r\delta^2 f_{(\sigma)}(F^{-1}_{(\sigma)}(1-q))/\lambda^2(ll).$$

Similarly,

$$\frac{\mathrm{d}(\pi^D_{hh} - \pi^D_{lh})}{\mathrm{d}\overline{\lambda}} \to \frac{1}{\sqrt{2}} r\delta^2 f_{(\sigma)}(F^{-1}_{(\sigma)}(1-q))/\overline{\lambda}^2.$$

($\overline{\lambda}$ is used because types lh and hl are bunched.)

Note that $\lambda(l) = p_l$ for both OEMs A and B, and $\lambda(ll) = p_l^2$ and $\overline{\lambda} = 2p_l/(1+p_l)$ for OEM D. Therefore, by the chain rule, to show

$$\frac{\mathrm{d}(\pi^a_h + \pi^b_h)}{\mathrm{d}p_l} < \frac{\mathrm{d}[(\pi^D_{hh} - \pi^D_{hl}) + \pi^D_{hl}]}{\mathrm{d}p_l},$$

I need only show

$$2\sqrt{2}/p_l^2 < 1/p_l^4 \cdot 2p_l + (1+p_l)^2/(4p_l^2) \cdot 2/(1+p_l)^2.$$

The latter can be verified to be true when $p_l < 0.85$. □

Proof of Theorem 2.4. I first show that bunching between types ll and lh will occur if δ^b is sufficiently small. When $\delta^b \to 0$, (2.8) converges to $1 - q = F^D(Q^D_{ll} - \mu^D_{ll})$, so its solution $Q^*(2.8)$ converges to $\mu_{ll}^D + (F^D)^{-1}(1-q)$. Furthermore, since

$$\begin{aligned} p_{lh}^D(1-q) &= (p_{lh}^D + p_{hl}^D + p_{hh}^D)F^D(Q_{lh}^D - \mu_{lh}^D) - (p_{hl}^D + p_{hh}^D)F^D(Q_{lh}^D - \mu_{hl}^D) \\ &> (p_{lh}^D + p_{hl}^D + p_{hh}^D)F^D(Q_{lh}^D - \mu_{lh}^D) - (p_{hl}^D + p_{hh}^D)F^D(Q_{lh}^D - \mu_{lh}^D) = p_{lh}^D F^D(Q_{lh}^D - \mu_{lh}^D), \end{aligned}$$

I have $Q^*(2.9) < \mu_{lh}^D + (F^D)^{-1}(1-q)$. Notice that $Q^*(2.9) - [\mu_{lh}^D + (F^D)^{-1}(1-q)]$ only depends on μ_{lh}^D and μ_{hl}^D through the gap $\mu_{hl}^D - \mu_{lh}^D = \delta^a - \delta^b$, which converges to δ^a as $\delta^b \to 0$. Thus, as $\delta^b \to 0$, $Q^*(2.9) - [\mu_{lh}^D + (F^D)^{-1}(1-q)]$ approaches a negative constant. Therefore, for sufficiently small δ^b, $Q^*(2.8) > Q^*(2.9)$, violating the MC. This means bunching occurs between types ll and lh as $\delta^b \to 0$.

As in (A.1)-(A.2), the information rent of OEM i ($i = a$ or b) is described by

$$\pi_h^i = r\int_{s^i-\delta^i}^{s^i} F^i(x)\mathrm{d}x$$

and

$$p_l^i(1-q) = F^i(s^i) - p_h^i F^i(s^i - \delta^i),$$

where $s^i \doteq Q_l^i - \mu_l^i$. When bunching occurs between types ll and lh, the information rent of OEM D is described by

$$\begin{aligned} \pi_{lh}^D &= r\int_{\underline{s}-\delta^b}^{\underline{s}} F^D(x)\mathrm{d}x \\ \pi_{hl}^D &= r\int_{\underline{s}-\delta^a}^{\underline{s}} F^D(x)\mathrm{d}x \\ \pi_{hh}^D &= r\int_{\underline{s}-\delta^a}^{\underline{s}} F^D(x)\mathrm{d}x + r\int_{s_{hl}-\delta^b}^{s_{hl}} F^D(x)\mathrm{d}x \\ (p_{ll}^D + p_{lh}^D)(1-q) &= F^D(\underline{s}) - (p_{hl}^D + p_{hh}^D)F^D(\underline{s} - \delta^a) \qquad \text{(A.15)} \\ p_{hl}^D(1-q) &= (p_{hl}^D + p_{hh}^D)F^D(s_{hl}) - p_{hh}^D F^D(s_{hl} - \delta^b), \qquad \text{(A.16)} \end{aligned}$$

where $\underline{s} \doteq \underline{Q} - \mu_{ll}^D$ (recall that bunching occurs between types ll and lh) and $s_{hl} \doteq$

$Q^D_{hl} - \mu^D_{hl}$. Since $p^D_{\theta^a\theta^b} = p^a_{\theta^a} p^b_{\theta^b}$, I can further simplify (A.15) and (A.16) as

$$p^a_l(1-q) = F^D(\underline{s}) - p^a_h F^D(\underline{s} - \delta^a)$$

and

$$p^b_l(1-q) = F^D(s_{hl}) - p^b_h F^D(s_{hl} - \delta^b).$$

The similar structure of OEMs A, B and D's information rent is now evident.

To prove the theorem for OEM D of type hl, I must show that $\pi^D_{hl} > \pi^a_h$ (recall that $\pi^b_l = 0$). Observe that the equations determining π^D_{hl} and π^a_h are identical except that the former involves F^D while the latter involves F^a. Since F^D has higher variability, from Lemma 2.3 I know that $\pi^D_{hl} > \pi^a_h$. Hence, to prove the theorem for OEM D of type hh, it is sufficient to show that $\pi^D_{hh} - \pi^D_{hl} > \pi^b_h$. Again, the equations determining $\pi^D_{hh} - \pi^D_{hl}$ and π^b_h are identical except that the former involves F^D while the latter involves F^b. The result follows from Lemma 2.3. □

Proof of Theorem 2.5. Proposition 2.1's proof showed that $FD_{\theta_j}(Q_{\theta_j})$ is first positive then negative. From the expression for $FD_{\theta_j}(Q_{\theta_j})$ it then follows that Q_{θ_j} is decreasing in $\mu_{\theta_{j+1}}$. Because $\mu_{\theta_j} \le \mu_{\theta_{j+1}} < \infty$, I can obtain upper and lower bounds on Q_{θ_j} by solving (2.4) with $\mu_{\theta_{j+1}}$ replaced by μ_{θ_j} and ∞, respectively. Doing so establishes that

$$\mu_{\theta_j} + F^{(-1)}(\lambda(\theta_j)(1-q)) < Q_{\theta_j} \le \mu_{\theta_j} + F^{(-1)}(1-q).$$

The theorem assumes $(r-c)/r > 0.5$ (equivalently, $q < 0.5$), and p^a_l and p^b_l are sufficiently close to 1. The above bounds ensure that

$$Q^i_{\theta^i} > \mu^i_{\theta^i} + (F^i)^{(-1)}(\lambda(\theta^i)(1-q)), \quad i = a, b,$$

and

$$Q^D_{\theta^a\theta^b} \le \mu^a_{\theta^a} + \mu^b_{\theta^b} + (F^D)^{(-1)}(1-q).$$

Therefore, to show $Q^a_{\theta^a} + Q^b_{\theta^b} > Q^D_{\theta^a\theta^b}$, it is sufficient to show

$$(F^a)^{(-1)}(\lambda(\theta^a)(1-q)) + (F^b)^{(-1)}(\lambda(\theta^b)(1-q)) > (F^D)^{(-1)}(1-q).$$

If $\lambda(\theta^a)$ and $\lambda(\theta^b)$ are sufficiently close to 1,

$$(F^a)^{(-1)}(\lambda(\theta^a)(1-q)) + (F^b)^{(-1)}(\lambda(\theta^b)(1-q))$$

will be sufficiently close to

$$(F^a)^{(-1)}(1-q) + (F^b)^{(-1)}(1-q).$$

Because $F^i(x) = \Phi_{(\sigma^i)}(x)$, I have

$$(F^a)^{(-1)}(1-q)+(F^b)^{(-1)}(1-q) = \Phi^{-1}_{(\sigma^a+\sigma^b)}(1-q) > \Phi^{-1}_{(\sqrt{(\sigma^a)^2+(\sigma^b)^2})}(1-q) = (F^D)^{(-1)}(1-q).$$

Therefore, the result follows if $\lambda(\theta^a)$ and $\lambda(\theta^b)$ are sufficiently close to 1.

Recall that when $\theta^i = l$, $\lambda(\theta^i) = p^i_l$, which is close to 1 by assumption. When $\theta^i = h$, by definition $\lambda(\theta^i) \equiv 1$. Therefore, the result follows. $\square$

APPENDIX B

Proofs of Chapter III

The Appendix includes the optimal mechanism design analysis and the proofs. I mentioned that Theorem 3.1 is a special case of Theorem 3.2 (with $h(k)$ replaced by 1, or equivalently, $H(Q)$ replaced by Q), which is in turn a special case of Theorem 3.3 (with all $\psi_i(\cdot)$ replaced by a common $\psi(\cdot)$). Therefore, it suffices to analyze the model with asymmetric and concave production costs as described in §3.3.2, and show proof for Theorem 3.3. I will first find the optimal direct-revelation mechanism, then show the modified biased open-descending auction always leads to the same outcome (purchase quantity and payment) as the optimal direct-revelation mechanism.

A mechanism can be described as a set of binding rules announced by the buyer stating how purchase quantities and payments will be determined based on the suppliers' bids. My goal is to design the purchase quantity and price rules that maximize the buyer's expected profit. Due to the *revelation principle*, it suffices to search for the optimal mechanism within direct-revelation mechanisms, namely the ones that require the suppliers to report their base costs and map their reported base costs to purchase decisions, while providing the incentive for the suppliers to report their base costs truthfully. For brevity, I use $-i$ to represent $j, \forall j \neq i$ in subscripts. For supplier i, denote his reported base cost with s_i, the purchase quantity from and payment to him with $Q_i(s_i, s_{-i})$ and $M_i(s_i, s_{-i})$. The expected payment to supplier i when he reports s_i, given that all other suppliers report truthfully, equals

$$m_i(s_i) \doteq E_{C_{-i}}[M_i(s_i, C_{-i})].$$

Also define

$$q_i(s_i) \doteq E_{C_{-i}}[H(Q_i(s_i, C_{-i}))].$$

Note that because H is increasing, if Q_i is decreasing in s_i, so is q_i. One can then see that supplier i's expected profit when reporting s_i, given all other suppliers report truthfully, equals

$$u_i(s_i) \doteq m_i(s_i) - c_i q_i(s_i).$$

The buyer's mechanism design problem can be written as

$$\begin{aligned} \max_{Q_i, M_i} \quad & E\left[R\left(\sum_{i=1}^{n} Q_i(C_i, C_{-i})\right)\right] - \sum_{i=1}^{n} E_{C_i}[m_i(C_i)] \\ \text{s.t.} \quad & Q_i \in \mathbb{Z}^+,\ u_i(c_i) \geq 0,\ u_i(c_i) \geq u_i(s_i),\ \forall s_i \neq c_i,\ i = 1, ..., n. \end{aligned} \tag{B.1}$$

By the standard analysis on pp 64-67 of Krishna (2002), the incentive compatibility constraint $u_i(c_i) \geq u_i(s_i)$ can be replaced with

$$u_i(c_i) = u_i(\bar{c}_i) + \int_{c_i}^{\bar{c}_i} q_i(t_i) \mathrm{d}t_i$$

and $q_i(\cdot)$ being non-increasing. In addition, since $u_i(c_i) \geq u_i(\bar{c}_i)$ under these two conditions, $u_i(c_i) \geq 0$ can also be replaced with $u_i(\bar{c}_i) \geq 0$. Combining

$$u_i(c_i) = u_i(\bar{c}_i) + \int_{c_i}^{\bar{c}_i} q_i(t_i) \mathrm{d}t_i$$

and

$$u_i(c_i) = m_i(c_i) - c_i q_i(c_i)$$

yields

$$m_i(c_i) = u_i(\bar{c}_i) + c_i q_i(c_i) + \int_{c_i}^{\bar{c}_i} q_i(t_i) \mathrm{d}t_i.$$

As a result,

$$\begin{aligned}
E_{C_i}[m_i(C_i)] =& \int_{\underline{c}_i}^{\bar{c}_i} m_i(c_i) f_i(c_i) \mathrm{d}c_i \\
=& u_i(\bar{c}_i) + \int_{\underline{c}_i}^{\bar{c}_i} c_i q_i(c_i) f_i(c_i) \mathrm{d}c_i + \int_{\underline{c}_i}^{\bar{c}_i} \int_{c_i}^{\bar{c}_i} q_i(t_i) f_i(c_i) \mathrm{d}c_i \mathrm{d}t_i \\
=& u_i(\bar{c}_i) + \int_{\underline{c}_i}^{\bar{c}_i} c_i q_i(c_i) f_i(c_i) \mathrm{d}c_i + \int_{\underline{c}_i}^{\bar{c}_i} \int_{\underline{c}_i}^{t_i} f_i(c_i) \mathrm{d}c_i q_i(t_i) \mathrm{d}t_i \\
=& u_i(\bar{c}_i) + \int_{\underline{c}_i}^{\bar{c}_i} \left(c_i + \frac{F_i(c_i)}{f_i(c_i)} \right) q_i(c_i) f_i(c_i) \mathrm{d}c_i \\
=& u_i(\bar{c}_i) + E\left[\left(C_i + \frac{F_i(C_i)}{f_i(C_i)} \right) H(Q_i(C_i, C_{-i})) \right].
\end{aligned}$$

Define virtual cost function $\psi_i(c) \doteq c + \frac{F_i(c)}{f_i(c)}$ and assume $\psi_i(\cdot)$ is non-decreasing for all i. Then the mechanism design problem becomes

$$\begin{aligned}
\max_{Q_i, M_i} \quad & E\left[R\left(\sum_{i=1}^n Q_i(C_i, C_{-i}) \right) - \sum_{i=1}^n \psi_i(C_i) H(Q_i(C_i, C_{-i})) \right] - \sum_{i=1}^n u_i(\bar{c}_i) \qquad \text{(B.2)} \\
\text{s.t.} \quad & Q_i \in \mathbb{Z}^+,\ u_i(\bar{c}_i) \geq 0,\ q_i(\cdot) \text{ non-increasing},\ i = 1, ..., n.
\end{aligned}$$

First observe that the optimal mechanism must have $u_i(\bar{c}_i) = 0$ for all i. Next, recall that Q_i non-increasing in c_i implies $q_i(\cdot)$ non-increasing. Therefore, if I find a purchase quantity rule Q_i that maximizes

$$R\left(\sum_{i=1}^n Q_i(c_i, c_{-i}) \right) - \sum_{i=1}^n \psi_i(c_i) H(Q_i(c_i, c_{-i}))$$

for every realization $\{c_i\}_{i=1}^n$ and is non-increasing in c_i, then Q_i will be the optimal purchase quantity rule. The next lemma states that to maximize

$$R\left(\sum_{i=1}^{n} Q_i(c_i, c_{-i})\right) - \sum_{i=1}^{n} \psi_i(c_i) H(Q_i(c_i, c_{-i})),$$

the buyer should purchase from the lowest-virtual cost supplier i quantity Q such that marginal revenue $r(Q)$ equals marginal virtual cost $\psi_i(c_i)h(Q)$. (The buyer compares suppliers based on virtual costs, rather than actual costs, in order to account for information rents; this is a well-known result in mechanism design.) Obviously, the resulting purchase quantity rule Q_i^* is non-increasing in c_i, therefore it is the optimal purchase quantity rule.

Lemma 2.1. *Denote with subscript* (i) *the ordered virtual costs:* $\psi_{(1)}(c_{(1)}) \leq ... \leq \psi_{(n)}(c_{(n)})$. *The following purchase quantity rule is optimal:* $Q^*_{(1)}(c_{(1)}, c_{-(1)}) = \max\{Q \in \mathbb{Z} | r(Q) \geq \psi_{(1)}(c_{(1)})h(Q)\}$*;* $Q^*_{(i)} \equiv 0$*,* $\forall i \geq 2$*.*

Proof of Lemma 2.1. First I show that it is optimal to only purchase from supplier (1). Suppose instead the buyer purchases $Q_{(i)} > 0$ and $Q_{(j)} > 0$. Since $\psi_{(i)}(C_{(i)}) \geq \psi_{(1)}(C_{(1)})$ for all $i \geq 2$, and due to the concavity of $H(\cdot)$,

$$\begin{aligned} &H(Q_{(i)}) + H(Q_{(j)}) = H(Q_{(i)}) + [H(Q_{(j)}) - H(0)] \\ \geq &H(Q_{(i)}) + [H(Q_{(i)} + Q_{(j)}) - H(Q_{(i)})] = H(Q_{(i)} + Q_{(j)}), \end{aligned}$$

I know

$$\psi_{(i)}(c_{(i)})H(Q_{(i)}) + \psi_{(j)}(c_{(j)})H(Q_{(j)}) \geq \psi_{(1)}(c_{(1)})H(Q_{(i)} + Q_{(j)}).$$

Thus it is optimal to only purchase from supplier (1). Then maximizing the integrand

$$R\left(\sum_{i=1}^{n} Q_i(c_i, c_{-i})\right) - \sum_{i=1}^{n} \psi_i(c_i) H(Q_i(c_i, c_{-i}))$$

is equivalent to maximizing

$$R\left(Q_{(1)}(c_{(1)}, c_{-(1)})\right) - \psi_{(1)}(c_{(1)})H(Q_{(1)}(c_{(1)}, c_{-(1)})),$$

which is achieved at

$$Q^*_{(1)}(c_{(1)}, c_{-(1)}) = \max\{Q \in \mathbb{Z} | r(Q) \geq \psi_{(1)}(c_{(1)})h(Q)\}.$$

Finally, $Q^*_{(1)}$ is increasing in $c_{(1)}$ because $\psi_{(i)}(\cdot)$ is assumed non-decreasing and $r(\cdot)/h(\cdot)$ is assumed decreasing. $\square$

Using Q^*_i from Lemma 2.1, the corresponding optimal payment rule can be written as

$$M^*_i(c_i, c_{-i}) = c_i H(Q^*_i(c_i, c_{-i})) + \int_{c_i}^{\bar{c}_i} H(Q^*_i(t_i, c_{-i}))\mathrm{d}t_i,$$

which will guarantee the required relationship

$$m^*_i(c_i) = c_i q^*_i(c_i) + \int_{c_i}^{\bar{c}_i} q^*_i(t_i)\mathrm{d}t_i.$$

The next lemma provides a more convenient expression of the optimal allocation and payment rules, which will be useful in proving Theorem 3.3.

Lemma 2.2. *Define*

$$P_i(c_{-i}, k) \doteq \max\{c | \psi_i(c) \leq \psi_j(c_j),\ \forall j \neq i,\ \psi_i(c) \leq r(k)/h(k)\},$$

$$W_i(c_i, c_{-i}, k) \doteq \mathbf{1}_{\{c_i \leq P_i(c_{-i}, k)\}},$$

$$N_i(c_i, c_{-i}, k) \doteq h(k)P_i(c_{-i}, k)W_i(c_i, c_{-i}, k).$$

Then

$$Q^*_i(c_i, c_{-i}) = \sum_{k=1}^{\infty} W_i(c_i, c_{-i}, k)$$

and

$$M_i^*(c_i, c_{-i}) = \sum_{k=1}^{\infty} N_i(c_i, c_{-i}, k).$$

Lemma 2.2 expresses the optimal allocation and payment rules from a per-unit perspective. $P_i(c_{-i}, k)$ denotes a threshold of supplier i's base cost c_i; the buyer will purchase the k^{th} unit from supplier i if his base cost c_i is below the threshold. $W_i(c_i, c_{-i}, k)$ is an indicator function of whether the buyer will purchase the k^{th} unit from supplier i, thus summing up $W_i(c_i, c_{-i}, k)$ for all k yields the optimal purchase quantity from supplier i. Finally, the definition of $N_i(c_i, c_{-i}, k)$ states that the buyer will pay exactly $h(k)$ times the threshold $P_i(c_{-i}, k)$ for the k^{th} unit (given that she will purchase this unit from supplier i), thus summing up $N_i(c_i, c_{-i}, k)$ for all k yields the optimal payment to supplier i.

Proof of Lemma 2.2. First I show that

$$Q_i^*(c_i, c_{-i}) = \sum_{k=1}^{\infty} W_i(c_i, c_{-i}, k).$$

Since $r(k)/h(k)$ is decreasing in k, $W_i(c_i, c_{-i}, k)$ is decreasing in k, namely once the buyer does not purchase a k^{th} unit from supplier i, she will not purchase any further units from supplier i. Thus the buyer only purchases from one supplier, and then by Lemma 2.1 it is clear that

$$Q_i^*(c_i, c_{-i}) = \sum_{k=1}^{\infty} W_i(c_i, c_{-i}, k).$$

Next I show

$$\sum_{k=1}^{\infty} N_i(c_i, c_{-i}, k) = M_i^*(c_i, c_{-i}).$$

Since

$$\int_{c_i}^{\bar{c}_i} W_i(t_i, c_{-i}, k)\mathrm{d}t_i = P_i(c_{-i}, k) - c_i,$$

$N_i(c_i, c_{-i}, k)$ can be rewritten as

$$W_i(c_i, c_{-i}, k) \left[h(k)c_i + \int_{c_i}^{\bar{c}_i} h(k) W_i(t_i, c_{-i}, k) \mathrm{d}t_i \right].$$

Summing this over all k yields

$$\sum_{k=1}^{\infty} N_i(c_i, c_{-i}, k) = c_i H \left(\sum_{k=1}^{\infty} W_i(c_i, c_{-i}, k) \right) + \int_{c_i}^{\bar{c}_i} H \left(\sum_{k=1}^{\infty} W_i(t_i, c_{-i}, k) \right) \mathrm{d}t_i$$

$$= c_i H(Q_i^*(c_i, c_{-i})) + \int_{c_i}^{\bar{c}_i} H(Q_i^*(t_i, c_{-i})) \mathrm{d}t_i,$$

which equals $M_i^*(c_i, c_{-i})$. □

Now I make use of Lemma 2.2 to prove Theorem 3.3.

Proof of Theorem 3.3. First I show that continuing to bid while $b_i^{-1}(P) > H(Q_0)c_i$ is supplier i's dominant strategy: When $b_i^{-1}(P) > H(Q_0)c_i$, if the auction ends and supplier i is the winner, he can just choose to deliver the guaranteed Q_0 units and secure a positive profit $b_i^{-1}(P) - H(Q_0)c_i$. On the other hand, if he loses the auction, he has profit 0. Thus he will keep bidding. When $b_i^{-1}(P) < H(Q_0)c_i$, if the auction ends and supplier i is the winner, he incurs a negative profit $b_i^{-1}(P) - H(Q_0)c_i$. Because the buyer pays no more than $h(k+Q_0)\psi_i^{-1}(P/H(Q_0))$ for the k^{th} additional unit, supplying additional units could only worsen the supplier's already negative profit (notice that $b_i^{-1}(P) < H(Q_0)c_i \Leftrightarrow h(k+Q_0)\psi_i^{-1}(P/H(Q_0)) < h(k+Q_0)c_i$). On the other hand, if he drops out, he has profit 0. Thus he will drop out.

Next I show when all suppliers' strategy is to keep bidding lower as long as $b_i^{-1}(P) > H(Q_0)c_i$, the purchase quantity and payment of the modified open-descending auction is identical to those described in Lemmas 2.1 and 2.2. Since the auction's starting price $H(Q_0)\max\{\psi_i(\bar{c}_i)\} \geq H(Q_0)\psi_i(c_i) = b_i(H(Q_0)c_i)$, $\forall i$, all suppliers can meet the reserve and the auction ends when the second-lowest-cost supplier drops out. Utilizing the ordered virtual costs notation, the winner of the auction will be supplier (1) with the lowest virtual cost $\psi_{(1)}(c_{(1)})$, and the auction's ending price P_0 will equal $H(Q_0)\psi_{(2)}(c_{(2)})$. Supplier (1) will first deliver Q_0 units, then will deliver the k^{th} additional unit as long as

$$h(k+Q_0)\psi_{(1)}^{-1}(\min\{\psi_{(2)}(c_{(2)}), r(k+Q_0)/h(k+Q_0)\}) \geq h(k+Q_0)c_{(1)}.$$

It is not difficult to see that

$$\begin{aligned}&\psi_{(1)}^{-1}(\min\{\psi_{(2)}(c_{(2)}), r(k+Q_0)/h(k+Q_0)\})\\ =&\sup\{c|\psi_{(1)}(c) \leq \psi_{(i)}(c_{(i)}),\ \forall i \geq 2,\ \psi_{(1)}(c) \leq r(k+Q_0)/h(k+Q_0)\}.\end{aligned}$$

The right-hand side expression equals $P_{(1)}(c_{-(1)}, k+Q_0)$ defined in Lemma 2.2. Thus supplier (1)'s total delivery quantity under the modified open-descending auction can be written as $Q = \max\{q \in \mathbb{Z}|c_{(1)} \leq P_{(1)}(c_{-(1)}, q)\}$, which equals Q_i^* by Lemma 2.2. Additionally, it is easy to verify that in the modified biased open-descending auction, the buyer pays $h(q)P_{(1)}(c_{-(1)}, q)$ to supplier (1) for the q^{th} unit shipped (this expression also applies to the initial Q_0 units), which equals $N_{(1)}(c_{(1)}, c_{-(1)}, q)$ defined in Lemma 2.2. Hence the modified open-descending auction is an optimal mechanism. □

Proof of Proposition 3.1. By the revelation principle, for any mechanism there exists an outcome-equivalent direct mechanism. Among direct mechanisms, the buyer's expected profit maximization problem (B.1) can be simplified into problem (B.2). By Lemma 2.1, in the symmetric linear costs case as I consider here, the optimal purchase quantity equals $Q^* = \max\{Q \in \mathbb{Z}|r(Q) \geq \psi(c_{(1)})\}$ where supplier (1) is the supplier with the lowest unit cost. Note that this optimal purchase quantity is *necessary* for the mechanism's optimality. However, if the auction in Stage 1 reveals $c_{(1)}$ to the buyer, then in Stage 2 the buyer will offer to pay $c_{(1)} + \varepsilon$ for each additional unit delivered, up to $r^{-1}(c_{(1)}) - Q_0$ units, where ε is an arbitrarily small positive number.

Receiving this offer, supplier (1) will deliver a total $\max\{Q \in \mathbb{Z} | r(Q) \geq c_{(1)}\}$ units, which is greater than Q^*. Such a mechanism cannot be optimal.

Next, suppose that I can replace the open-descending auction in the mechanism by a sealed first-price, Vickrey or reverse Dutch auction and retain optimality. I show a contradiction. Notice that each of these latter three auction formats, a supplier's bidding strategy can be described as a simple bidding function $\beta_i(c_i)$ that only depends on the supplier's cost c_i. Since I assume an optimal mechanism, the lowest-cost supplier must *always* win Stage 1's auction. This is because (as I mentioned above) any optimal mechanism must ensure that the buyer purchases from the lowest-cost supplier, and in the mechanism the buyer will only purchase from Stage 1's auction winner. The requirement that the lowest-cost supplier must always win Stage 1's auction implies $c_i < c_j \Leftrightarrow \beta_i(c_i) < \beta_j(c_j), \forall i \neq j$, which in turn implies the following properties of β_i: All suppliers use the same bidding function: $\beta_i(\cdot) \equiv \beta(\cdot)$ for all i; $\beta(\cdot)$ is deterministic; $\beta(\cdot)$ is monotonic. The three properties imply $\beta_i(\cdot) \equiv \beta(\cdot)$ is a one-to-one mapping between supplier i's unit cost c_i and his bid β_i in the auction. Hence, at the conclusion of the auction the buyer can infer the winning supplier i's unit cost $c_i = \beta^{-1}(\beta_i)$ from his bid β_i. By the earlier discussion, such a mechanism cannot be optimal. $\square$

APPENDIX C

Proofs of Chapter IV

Note. Although Propositions 4.1 through 4.4 and Theorem 4.1 are stated for two TOs with independent Z_i's, their proofs here are provided for any number of TOs and without assuming Z_i's are independent. This will facilitate proofs of Theorems 4.5 and 4.6.

Proof of Proposition 4.1. Recall that Z_i has positive pdf over $[a_i, a_i + h_i]$. I take the example of TO_1 and show that any quoting strategy $(x_1, ..., x_n)$ with $x_1 > a_1 + h_1$ or $x_1 < a_1$ is weakly dominated by a quoting strategy with $x_1 \in [a_1, a_1 + h_1]$.

First I show quoting strategy $(x_1, ..., x_n)$ with $x_1 > a_1 + h_1$ is weakly dominated by $(a_1 + h_1, x_2, ..., x_n)$. This is straightforward: The former strategy renders TO_1 unable to meet the OEM's reserve price, and so does the latter. Therefore, they lead to the same level of expected profit for TT.

Next I show that, for any i, any strategy with $x_i < a_i$ is weakly dominated by one where $x_i \geq a_i$. Assuming without loss of generality $a_1 - x_1 \geq a_i - x_i$, $\forall i$ and $a_1 - x_1 = t > 0$, it suffices to show that quoting strategy $(x_1, ..., x_n)$ is weakly dominated by $(a_1, x_2 + t, ..., x_n + t)$, called the revised strategy.

Revised strategy $(a_1, x_2 + t, ..., x_n + t)$ is generated by increasing each quote in $(x_1, ..., x_n)$ by $t = a_1 - x_1$. Notice that doing so will not affect the TOs' sample path (realizations of Z_i's) cost comparison in the OEM's auction; namely, the most cost-efficient TO will remain so after the change. On the other hand, notice that at least one TO can meet the OEM's reserve price with both strategies $(x_1, ..., x_n)$ and $(a_1, x_2+t, ..., x_n+t)$, meaning the most cost-efficient TO will always meet the reserve price. As a result, the revised strategy does not change the winning TO. However, because quotes to all TOs increased by t, whichever TO wins will, under the revised strategy, yield an additional profit of t for TT, meaning TT's expected profit with $(a_1, x_2+t, ..., x_n+t)$ is higher than $(x_1, ..., x_n)$ by $t > 0$. $\square$

Proof of Proposition 4.2. I begin by stating a simple fact useful in showing the proposition: TT's profit is improved by increasing the *lowest* quote(s) of a secure strategy until either reaching the next lowest quote, or reaching the point beyond which the strategy would become risky. This is easy to see by a sample-path argument. When increasing the lowest quote(s), the sample paths where these TOs would win the OEM's contract either remain so, or become ones where the other TOs would win, both of which mean higher profit for TT. (Note that the key assumption that the quoting strategy remains secure ensures these sample-paths do not become ones where no TO could meet the reserve price.)

Without loss of generality, I consider a secure strategy $X^s = (x_1^s, ..., x_n^s)$ where $x_1^s = a_1 \geq a_i$ for all i such that $x_i^s = a_i$. I take three steps to show the proposition. First, repeatedly using the above argument, I can improve TT's expected profit by increasing any $x_j^s < a_1$ to a_1 (the strategy remains secure because I still have $x_1^s = a_1$). Denote the resulting strategy by $X'^s = (x_1'^s, ..., x_n'^s)$. Now, $x_1'^s = a_1$ is the lowest quote and is a secure quote. Without loss of generality, assume $a_1 = x_1'^s = ... = x_j'^s < x_{j+1}'^s \leq ... \leq x_n'^s$. Second, simultaneously increasing $x_i'^s$, $i = 2, ..., j$ to $x_{j+1}'^s$ improves TT's expected profit. I use a sample-path argument to show this. When increasing $x_i'^s$, $i = 2, ..., j$ to $x_{j+1}'^s$, the sample paths where these TOs would win the OEM's contract become ones where TO_1 wins, or one of TO_k, $k \geq 2$ wins. The former sample paths do not affect TT's profit, and the latter sample paths will improve TT's profit. Denote the resulting secure strategy by $X''^s = (x_1''^s, ..., x_n''^s)$, where $x_1''^s = a_1$, and $x_i''^s > \max\{a_1, a_i\}$, $i = 2, ..., n$. Finally, if $a_i + h_i < a_1$, then I can decrease $x_i''^s$ to $a_i + h_i$ without reducing TT's expected profit — this is part of Proposition 4.1's proof.

As shown in the above three steps, any secure strategy X^s is dominated by another secure strategy that satisfies the conditions in Proposition 4.2. Therefore these conditions are necessary for optimality. $\square$

Proof of Proposition 4.3. For any secure strategy X^s, define $\hat{X}^s$ by $\hat{x}_i^s \doteq x_i^s + a$, $\forall i$. Since

$$\begin{aligned} P(\hat{X}^s) =& a + a_1 + \sum_{i=2}^{n} (\hat{x}_i^s - a - a_1) \Pr\{TO_i \text{ winning}\} \\ =& a + a_1 + \sum_{i=2}^{n} (x_i^s - a_1) \Pr\{TO_i \text{ winning}\}, \end{aligned} \tag{C.1}$$

the optimal secure strategy $\hat{X}^{s*}$ that maximizes $P(\hat{X}^{s*})$ is obviously invariant in a and is characterized by $\hat{x}_i^{s*} = x_i^{s*} + a$, $\forall i$. $\square$

Proof of Theorem 4.1. First I prove a lemma that will be used in this proof.

Lemma 3.1. *Fix a secure quoting strategy X^s and a risky quoting strategy X^r. Replace Z_i by $\hat{Z}_i \doteq Z_i + a$ in (4.2) and consider the secure strategy $\hat{X}^s \doteq X^s + a = (x_1^s + a, ..., x_n^s + a)$ and the risky strategy $\hat{X}^r \doteq X^r + a = (x_1^r + a, ..., x_n^r + a)$. The derivatives of TT's expected profit $P(\hat{X}^s)$ and $P(\hat{X}^r)$ with respect to a are constants and satisfy*

$$\frac{\mathrm{d}}{\mathrm{d}a} P(\hat{X}^s) \equiv 1, \quad \frac{\mathrm{d}}{\mathrm{d}a} P(\hat{X}^r) \equiv 1 - \Pr\{\textit{All TOs losing}\} < 1.$$

Proof of Lemma 3.1. First consider the secure strategy. Assume X^s has $x_1^s = a_1$.

Since the expected profit can be written as in (C.1), it is immediately seen that

$$\frac{\mathrm{d}}{\mathrm{d}a} P(\hat{X}^s) \equiv 1.$$

Next consider the risky strategy. Because

$$P(\hat{X}^r) = \sum_{i=1}^{n} (a + x_i^r) \Pr\{TO_i \text{ winning}\},$$

I know

$$\frac{\mathrm{d}}{\mathrm{d}a} P(\hat{X}^r) \equiv 1 - \Pr\{\text{All } TO\text{s losing}\} < 1.$$

□

Take the optimal secure strategy X^{s*} and any risky strategy X^r. Replace Z_i by $\hat{Z}_i \doteq Z_i + a$ in (4.2), and consider the optimal secure strategy $\hat{X}^{s*} \doteq X^{s*} + a$ and the risky strategy $\hat{X}^r \doteq X^r + a$. Due to Lemma 3.1,

$$\frac{\mathrm{d}}{\mathrm{d}a} P(\hat{X}^{s*}) - \frac{\mathrm{d}}{\mathrm{d}a} P(\hat{X}^r) > 0$$

is a constant. As a result, there exists a finite $a^*_{X^r}$ such that $P(\hat{X}^{s*}) > P(\hat{X}^r)$ when $a > a^*_{X^r}$. Define $T_{sec} = \sup\{a^*_{X^r}\}$ for all X^r. Then when $a > T_{sec}$, the optimal secure strategy $\hat{X}^{s*}$ generates higher expected profit for TT than *all* risky strategies, thus is the optimal strategy.

Next I show T_{sec} is finite. Denote $\hat{Z}_i$'s pdf by $\hat{f}_i$ and recall that $\hat{f}_i$ is positive and finite over $[a + a_i, a + a_i + h]$. Consider a family of strategies

$$\hat{X}^\delta \doteq (\hat{x}_1 = a + a_1 + \delta, \hat{x}_2(a + a_1 + \delta)^*, ..., \hat{x}_n(a + a_1 + \delta)^*)$$

where $\hat{x}_i(a + a_1 + \delta)^*$ is the optimal $\hat{x}_i$ given $\hat{x}_1 = a + a_1 + \delta$. I first show that there exists $T_1 < \infty$, $\varepsilon > 0$, such that $a > T_1$ and $0 < \delta < \varepsilon$ imply that TT prefers using

$\hat{X}^0$ to using $\hat{X}^\delta$, that is, $P(\hat{X}^0) \geq P(\hat{X}^\delta)$. Since by definition

$$P(\hat{X}^0) \geq P(a + a_1, \hat{x}_2(a + a_1 + \delta)^*, ..., \hat{x}_n(a + a_1 + \delta)^*).$$

it suffices to show

$$P(a + a_1, \hat{x}_2(a + a_1 + \delta)^*, ..., \hat{x}_n(a + a_1 + \delta)^*) - P(\hat{X}^\delta) \geq 0.$$

Notice that $\hat{X}^\delta$ can be thought of as strategy

$$(\hat{x}_1 = a + a_1, \hat{x}_2(a + a_1 + \delta)^*, ..., \hat{x}_n(a + a_1 + \delta)^*)$$

with the quote $\hat{x}_1 = a + a_1$ increased by δ. For small δ, increasing δ has three effects on TT's profit. First, with some positive probability $p_0 > 0$ invariant in a, TT may no longer have business because no TO meets the reserve price, which leads to a loss of *at least* $a + a_1$. Second, TT will receive δ additional profit when TO_1 wins the contract, whose probability is *at most* $1 - \delta \hat{f}_1(a + a_1) + o(\delta)$. Third, the chance that another TO_i wins may increase by *at most* $\delta \hat{f}_1(a + a_1) + o(\delta)$, which brings in an increased profit of *at most* $\max\limits_{2 \leq i \leq n}\{a + a_i + h_i - a - a_1\}$. Combining the three effects, I can upper bound the profit impact of increasing $\hat{x}_1$ by δ:

$$\begin{aligned}&P(a + a_1, \hat{x}_2(a + a_1 + \delta)^*, ..., \hat{x}_n(a + a_1 + \delta)^*) - P(\hat{X}^\delta) \\ \geq &(a + a_1)p_0 - \delta(1 - \delta \hat{f}_1(a + a_1)) - \max_{2 \leq i \leq n}\{a_i + h_i - a_1\}\delta \hat{f}_1(a + a_1) + o(\delta).\end{aligned}$$

It is obvious that there exist $T_1 < \infty$, $\varepsilon > 0$ such that when $a > T_1$ and $0 < \delta < \varepsilon$, the right hand side is positive, and thus TT prefers using $\hat{X}^0$ to using $\hat{X}^\delta$. In other words, for a greater than a finite T_1, a secure strategy is preferred to all risky strategies that have $a + a_1 < \hat{x}_1 < a + a_1 + \varepsilon$. Repeating the proof for all $\hat{x}_i$, I know that for a greater than a finite T_1, the optimal strategy among all strategies that have $\hat{x}_i < a + a_i + \varepsilon$ for some i must be secure.

Next I show that there exists a $T_2 < \infty$ such that for a greater than T_2 the optimal strategy is secure among all strategies that have $\hat{x}_i > a + a_i + \varepsilon$ for all i. To show this, assume $\hat{X}^*$ is the optimal strategy among those that satisfy $\hat{x}_i > a + a_i + \varepsilon$ for all i. For all a, TT's profit with $\hat{X}^*$ is upper-bounded by $(a + \max_i\{a_i + h_i\})(1 - \Pr\{\text{all } TO\text{s losing}|\hat{x}_i = a + a_i + \varepsilon\})$. In comparison, TT's profit with any secure strategies is lower-bounded by $a + \min_i\{a_i\}$. Obviously, $\Pr\{\text{all } TO\text{s losing}|\hat{x}_i = a + a_i + \varepsilon\}$ is a positive constant invariant in a. Therefore, there exists some finite T_2 such that as $a > T_2$, $\hat{X}^*$ is dominated by a secure strategy.

Combining the above results, I know $T_{sec} \leq \max\{T_1, T_2\} < \infty$, namely T_{sec} is finite. □

Proof of Theorem 4.2. Denote the i.i.d. Z_i's (common) cdf by F. Fix a quoting strategy $X = (x_1, x_2)$, and consider the strategy $\hat{X} = (\hat{x}_1, \hat{x}_2) = (x_1 + a, x_2 + a)$. TT's expected profit under strategy $\hat{X}$ equals

$$P(\hat{x}_1, \hat{x}_2) = (a + x_1)\int_{x_1}^{\infty} F(z_1 - x_1 + x_2)\mathrm{d}F(z_1) + (a + x_2)\int_{x_2}^{\infty} F(z_2 - x_2 + x_1)\mathrm{d}F(z_2),$$

and

$$\begin{aligned}
&\frac{\partial P(\hat{x}_1, \hat{x}_2)}{\partial x_1} - \frac{\partial P(\hat{x}_1, \hat{x}_2)}{\partial x_2} \\
=&2(x_2 - x_1)\int_0^{\infty} f(x_1 + t)f(x_2 + t)\mathrm{d}t + (a + x_2)f(x_2)F(x_1) - (a + x_1)f(x_1)F(x_2) \\
&- \int_0^{\infty}(f(x_1 + t)F(x_2 + t) - F(x_1 + t)f(x_2 + t))\mathrm{d}t.
\end{aligned}$$

Furthermore,

$$\frac{\mathrm{d}}{\mathrm{d}a}\left[\frac{\partial P(\hat{x}_1, \hat{x}_2)}{\partial x_1} - \frac{\partial P(\hat{x}_1, \hat{x}_2)}{\partial x_2}\right] = f(x_2)F(x_1) - f(x_1)F(x_2). \tag{C.2}$$

Notice that $1/\sqrt{2}$ times (C.2) equals the constant rate of change of $P(\hat{x}_1, \hat{x}_2)$'s $(1, -1)$ directional derivative in a. Obviously, (C.2) is anti-symmetric in x_1 and x_2, meaning switching the indices will change its sign. Therefore, if I can show for $x_1 < x_2$ it is negative, then when a is sufficiently small (possibly negative), $P(\hat{x}_1, \hat{x}_2)$'s $(1, -1)$ directional derivative will be positive when $x_1 < x_2$, and the $(-1, 1)$ directional derivative will be positive when $x_1 > x_2$, implying the optimal quotes must be symmetric. Below I show (C.2) is indeed negative for $x_1 < x_2$. Observe that

$$\text{(C.2)} < 0 \text{ for } x_1 < x_2 \Longleftrightarrow f(x)/F(x) \text{ is decreasing in } x,$$

which is in turn implied by the assumption that Y_i has log-concave pdf.

$T_{sym} \leq T_{sec}$ is implied by Proposition 4.2, which requires an optimal secure strategy to never be symmetric (implying that an optimal strategy that is symmetric must be risky and can only occur with $a \leq T_{sec}$). □

Proof of Proposition 4.4. Quoting $x < \max_i\{a_i\}$ is strictly dominated by quoting $\max_i\{a_i\}$, because $\max_i\{a_i\}$ is a higher quote and yet guarantees at least one TO will meet reserve price r. Quoting $x > \max_i\{a_i + h_i\}$ is weakly dominated by quoting $\max_i\{a_i + h_i\}$, because both quotes guarantee TO_i cannot meet the reserve price. □

Proof of Proposition 4.5. The necessary condition is a first-order condition. To show the condition is also sufficient when the Z_i's are i.i.d., I only need to show $\max\{Z_1, Z_2\} = r - \min\{Y_1, Y_2\}$ is IFR, so that its failure rate function

$$\frac{2F(x)f(x)}{1 - F(x)^2}$$

is increasing. Since Y_i is assumed to have a log-concave density, by the results of Huang and Ghosh (1982), $\min\{Y_1, Y_2\}$ also has a log-concave density; then by the results of Bagnoli and Bergstrom (2005), $\max\{Z_1, Z_2\} = r - \min\{Y_1, Y_2\}$ is IFR. □

Proof of Theorem 4.3. Denote the cdf and pdf of $\hat{Z}_i$ by $\hat{F}$, $\hat{f}$, respectively. Notice that $\hat{Z}_i$ has support $[a, a+h]$. From Proposition 4.5 I know the optimal QEP quote x^* is characterized by

$$x^* = \frac{1-\hat{F}^2(x^*)}{2\hat{F}(x^*)\hat{f}(x^*)}.$$

Since as $x \to a$,

$$\frac{1-\hat{F}^2(x)}{2\hat{F}(x)\hat{f}(x)} \to \infty > x,$$

I know $x^* > a$ must always be true.

To show the optimal quote converges to a secure quote $x^* = a$, I take any $\varepsilon > 0$ and consider a quote $x = a + \eta$ where $\eta > \varepsilon$. TT's expected profit from this quote is upper-bounded by $(a+h)(1-\hat{F}^2(a+\varepsilon))$, which is obviously dominated by a, TT's profit from quoting $x = a$, when

$$a > h(1-\hat{F}^2(a+\varepsilon))/\hat{F}^2(a+\varepsilon).$$

Therefore, for any $\varepsilon > 0$, quoting $x > a+\varepsilon$ cannot be optimal when a is sufficiently large, meaning x^* asymptotically converges to a secure quote $x = a$. □

Proof of Corollary 4.1. When a is sufficiently small, I know the optimal QP quotes are symmetric (Theorem 4.2), which must also be the optimal QEP quote. When a is sufficiently large, I know that $x_1^* = a$ and $x_2^* - a$ is invariant in a (Theorem 4.1, Proposition 4.3), but x^* asymptotically converges to a (Theorem 4.3), so it must be true that $x_1^* < x^* < x_2^*$. □

Proof of Theorem 4.4. Obviously, the mechanism design problem in Theorem 4.4 is a special case of the problem in Theorem 4.7, with $R \equiv r$ and $n = 2$. Assuming Theorem 4.7 is true (the proof is provided later), I only need to show the mechanism of Theorem 4.7 reduces to the mechanism of Theorem 4.4.

Now assume $R \equiv r$ and $n = 2$. By definition, $z_i + y_i = Z_i + Y_i \equiv r,\ i = 1, 2$, and $u \equiv r$. Considering the relationship between Z_i and Y_i, it is not difficult to verify the following equation: $\psi_i(z_i) + \omega_i(y_i) = r$. Theorem 4.7 states that the supplier i (if any) satisfying $\omega_i(\check{y}_i) \le \min_{j\neq i}\{\omega_j(\check{y}_j), u\}$ is charged upfront payment

$$p_i^* = E_R[\mathbf{1}_{\{\omega_i(\check{y}_i)<R\}}(R - \omega_i(\check{y}_i) + \check{y}_i - \omega_i^{-1}(\min_{j\neq i}\{\omega_j(\check{y}_j), u\}))]$$

and quoted price

$$x_i^* = \frac{G_i(\check{y}_i)}{g_i(\check{y}_i)}.$$

Notice that when $R \equiv r$ and $n = 2$,

$$\omega_i(\check{y}_i) \le \min_{j\neq i}\{\omega_j(\check{y}_j), u\} \iff \psi_i(\check{z}_i) \ge \max\{\psi_j(\check{z}_j), 0\},$$

and for supplier i that satisfies the above condition,

$$\begin{aligned} p_i^* &= E_R[\mathbf{1}_{\{\omega_i(\check{y}_i)<R\}}(R - \omega_i(\check{y}_i) + \check{y}_i - \omega_i^{-1}(\min_{j\neq i}\{\omega_j(\check{y}_j), u\}))] \\ &= \psi_i^{-1}(\max\{\psi_j(\check{z}_j), 0\}) - \frac{1 - F_i(\check{z}_i)}{f_i(\check{z}_i)}, \\ x_i^* &= \frac{G_i(\check{y}_i)}{g_i(\check{y}_i)} = \omega_i(\check{y}_i) - \check{y}_i = r - \psi_i(\check{z}_i) - r + \check{z}_i = \frac{1 - F_i(\check{z}_i)}{f_i(\check{z}_i)}. \end{aligned}$$

Furthermore, since $R \equiv r$ and TT backs only one TO_i in the auction, I know the OEM will pay r to TO_i for the contract and TO_i will transfer a total amount of $p_i + x_i$ to TT, implying $r \ge p_i + x_i$. Alternatively, any upfront payment $p_i' \ge 0$ and price quote $x_i' \ge 0$ satisfying $p_i' + x_i' = p_i + x_i$ would ensure that TO_i meets the reserve price (because $r \ge p_i + x_i$), and transfer the same amount to TT. Therefore, I know that any $p_i^* \ge 0$ and $x_i^* \ge 0$ satisfying $p_i^* + x_i^* = \max\{\psi_i^{-1}(\psi_j(\check{z}_j)), \psi_i^{-1}(0)\}$ form an optimal mechanism. This concludes the proof. □

Proof of Theorem 4.5. The proofs of Propositions 4.1 through Proposition 4.3 and Theorem 4.1 were provided for the generalized model (notice that none of the proofs require the Z_i's to be independently distributed). □

Proof of Theorem 4.6. The proof of Proposition 4.4 was provided for the generalized model.

To show Proposition 4.5 extends to the general model, I first need to show $\max_i\{Z_i\} = R - \min_i\{Y_i\}$ is IFR. Since Y_i and R are assumed to have log-concave density, by the results of Huang and Ghosh (1982), $\min_i\{Y_i\}$ also has log-concave densities; then by the results of Bagnoli and Bergstrom (2005) and Theorem 3.2 in Barlow, Marshall and Proschan (1963), $\max_i\{Z_i\} = R - \min_i\{Y_i\}$ is IFR.

To show the next step, I denote the cdf, pdf and support of R by W, w and $[\underline{R}, \overline{R}]$, and those of $Y \doteq \min_i\{Y_i\}$ by G, g, and $[\underline{Y}, \overline{Y}]$, respectively. I need to show $\max_i\{Z_i\} = R - Y$ has zero probability density at left end point of its support: $f(\underline{R} - \overline{Y}) = 0$ (so $x^* = \frac{1-F(x^*)}{f(x^*)}$ cannot hold at this point). By definition,

$$F(z) = \int_{\underline{Y}}^{\overline{Y}} g(y)W(y+z)\mathrm{d}y, \quad f(z) = F'(z) = \int_{\underline{Y}}^{\overline{Y}} g(y)w(y+z)\mathrm{d}y.$$

Therefore,

$$f(\underline{R} - \overline{Y}) = \int_{\underline{Y}}^{\overline{Y}} g(y)w(y + \underline{R} - \overline{Y})\mathrm{d}y.$$

Notice that $\forall y \in [\underline{Y}, \overline{Y}]$, $y + \underline{R} - \overline{Y} \leq \underline{R}$, and consequently $f(\underline{R} - \overline{Y}) = 0$. This concludes the extension of the proof. $\square$

Proof of Theorem 4.7. For concision I use subscript $-i$ to denote subscript j, $\forall j \neq i$. Suppose Z_i has support $[a_i, b_i]$. For TO_i with cost y_i and report $\check{y}_i$, if I assume all other TO_j's report their true costs $\check{y}_j = y_j$, then TO_i's expected profit is

$$\begin{aligned} v_i(\check{y}_i, y_i) &= E_{R,Y_{-i}}[(\min\{R, Y_{-i} + x_{-i}(\check{y}_i, Y_{-i})\} - y_i - x_i(\check{y}_i, Y_{-i}))^+ - p_i(\check{y}_i, Y_{-i})], \\ \frac{\partial v_i}{\partial y_i} &= -E_{R,Y_{-i}}[\mathbf{1}_{\{\min\{R, Y_{-i}+x_{-i}(\check{y}_i, Y_{-i})\} - y_i - x_i(\check{y}_i, Y_{-i})>0\}}]. \end{aligned}$$

A direct-revelation mechanism must satisfy $\max_{\check{y}_i} v_i(\check{y}_i, y_i) = v_i(y_i, y_i)$. By the envelope theorem:

$$v_i(y_i, y_i) = v_i(b_i, b_i) + \int_{y_i}^{b_i} E_{R,Y_{-i}}[\mathbf{1}_{\{\min\{R, Y_{-i}+x_{-i}(z,Y_{-i})\}-z-x_i(z,Y_{-i})>0\}}]\mathrm{d}z,$$

and

$$\begin{aligned} &E_{Y_{-i}}[p_i(y_i, Y_{-i})] \\ =&E_{R,Y_{-i}}[(\min\{R, Y_{-i} + x_{-i}(y_i, Y_{-i})\} - y_i - x_i(y_i, Y_{-i}))^+] - v_i(y_i, y_i). \end{aligned} \tag{C.3}$$

So TT's expected profit equals

$$\begin{aligned} P \doteq& \sum_{i=1}^{n} E_{Y_i}[E_{R,Y_{-i}}[p_i(Y_i, Y_{-i}) + x_i(Y_i, Y_{-i})\mathbf{1}_{\{\min\{R,Y_{-i}+x_{-i}(Y_i,Y_{-i})\}-Y_i-x_i(Y_i,Y_{-i})>0\}}]] \\ =& \sum_{i=1}^{n} \left\{ \int_{a_i}^{b_i} g_i(z)[E_{R,Y_{-i}} \left[(\min\{R, Y_{-i} + x_{-i}(z, Y_{-i})\} - z - x_i(z, Y_{-i}))^+ \right] \right. \\ &- \int_z^{b_i} E_{R,Y_{-i}}[\mathbf{1}_{\{\min\{R,Y_{-i}+x_{-i}(w,Y_{-i})\}-w-x_i(w,Y_{-i})>0\}}]\mathrm{d}w \\ &+ \left. \left. E_{R,Y_{-i}}[(x_i(z, Y_{-i}))\mathbf{1}_{\{\min\{R,Y_{-i}+x_{-i}(z,Y_{-i})\}-z-x_i(z,Y_{-i})>0\}}] \right] \mathrm{d}z - v_i(b_i, b_i) \right\}. \end{aligned}$$

Notice that

$$\begin{aligned} &\int_{a_i}^{b_i} g_i(z) \int_z^{b_i} E_{R,Y_{-i}}[\mathbf{1}_{\{\min\{R,Y_{-i}+x_{-i}(w,Y_{-i})\}-w-x_i(w,Y_{-i})>0\}}]\mathrm{d}w\mathrm{d}z \\ =& \int_{a_i}^{b_i} \int_{a_i}^{w} g_i(z) E_{R,Y_{-i}}[\mathbf{1}_{\{\min\{R,Y_{-i}+x_{-i}(w,Y_{-i})\}-w-x_i(w,Y_{-i})>0\}}]\mathrm{d}z\mathrm{d}w \\ =& \int_{a_i}^{b_i} G_i(w) E_{R,Y_{-i}}[\mathbf{1}_{\{\min\{R,Y_{-i}+x_{-i}(w,Y_{-i})\}-w-x_i(w,Y_{-i})>0\}}]\mathrm{d}w. \end{aligned}$$

Plugging it back in, I have

$$
\begin{aligned}
P =& \sum_{i=1}^{n} \left\{ \int_{a_i}^{b_i} g_i(z) \left[E_{R,Y_{-i}} \left[(\min\{R, Y_{-i} + x_{-i}(z, Y_{-i})\} - z - x_i(z, Y_{-i}))^+ \right.\right.\right. \\
& - \frac{G_i(z)}{g_i(z)} \mathbf{1}_{\{\min\{R, Y_{-i}+x_{-i}(z,Y_{-i})\} - z - x_i(z,Y_{-i}) > 0\}} \\
& \left.\left.\left. + \ x_i(z, Y_{-i}) \mathbf{1}_{\{\min\{R, Y_{-i}+x_{-i}(z,Y_{-i})\} - z - x_i(z,Y_{-i}) > 0\}} \right]\right] \mathrm{d}z - v_i(b_i, b_i) \right\} \\
=& \sum_{i=1}^{n} \left\{ \int_{a_i}^{b_i} g_i(z) \left[E_{R,Y_{-i}} \left[\mathbf{1}_{\{\min\{R, Y_{-i}+x_{-i}(z,Y_{-i})\} - z - x_i(z,Y_{-i}) > 0\}} \right.\right.\right. \\
& \left.\left.\left. \left(\min\{R, Y_{-i} + x_{-i}(z, Y_{-i})\} - z - x_i(z, Y_{-i}) - \frac{G_i(z)}{g_i(z)} + x_i(z, Y_{-i}) \right) \right]\right] \mathrm{d}z - v_i(b_i, b_i) \right\} \\
=& E_{R,Y} \left[\sum_{i=1}^{n} \left\{ \mathbf{1}_{\{\min\{R, Y_{-i}+x_{-i}(Y_i,Y_{-i})\} - Y_i - x_i(Y_i,Y_{-i}) > 0\}} \right.\right. \\
& \left.\left. \left(\min\{R, Y_{-i} + x_{-i}(Y_i, Y_{-i})\} - Y_i - \frac{G_i(Y_i)}{g_i(Y_i)} \right) - v_i(b_i, b_i) \right\} \right].
\end{aligned}
$$

To maximize P, I set $v_i(b_i, b_i) = 0$. Observe that

$$\min\{R, y_{-i} + x_{-i}(y_i, y_{-i})\} - y_i - x_i(y_i, y_{-i}) > 0$$

can be true for at most one i. Thus, if

$$\min\{R, y_{-i} + x_{-i}(y_i, y_{-i})\} - y_i - \frac{G_i(y_i)}{g_i(y_i)}$$

takes value

$$\left(R - y_i - \frac{G_i(y_i)}{g_i(y_i)} \right)^+$$

whenever

$$y_i - \frac{G_i(y_i)}{g_i(y_i)} < y_j - \frac{G_j(y_j)}{g_j(y_j)}, \quad \forall j \neq i,$$

then the sum must be maximized. This is achieved by setting

$$x_i(y_i, y_{-i}) = \frac{G_i(y_i)}{g_i(y_i)}$$

for the TO_i with the smallest $\omega_i(y_i)$ if $\omega_i(y_i) < u$ (otherwise $x_i(y_i, y_{-i}) > u - y_i$), and $x_{-i}(y_i, y_{-i}) > u - y_{-i}$ (eliminate all other TOs with forbiddingly high quotes). This is the optimal quoting rule. With the optimal quotes, TT's expected profit can be written as

$$P = E_{R,Y}[(R - \min\{\omega_i(Y_i)\})^+].$$

Now consider the payment rule. The envelope theorem implies (C.3). With the optimal quotes, (C.3) can be rewritten as

$$\begin{aligned} &E_{Y_{-i}}[p_i(y_i, Y_{-i})] \\ =&E_{R,Y_{-i}}\left[\mathbf{1}_{\{\omega_i(y_i)<\min\{R,\omega_{-i}(Y_{-i})\}\}}(R-\omega_i(y_i)) - \int_{y_i}^{b_i} \mathbf{1}_{\{\omega_i(z)<\min\{R,\omega_{-i}(Y_{-i})\}\}}\mathrm{d}z\right]. \end{aligned} \tag{C.4}$$

Define the following payment rule:

$$p_i(y_i, y_{-i}) = E_R\left[\mathbf{1}_{\{\omega_i(y_i)<\min\{R,\omega_{-i}(y_{-i})\}\}}(R-\omega_i(y_i)) - \int_{y_i}^{b_i} \mathbf{1}_{\{\omega_i(z)<\min\{R,\omega_{-i}(y_{-i})\}\}}\mathrm{d}z\right].$$

Obviously, this payment rule ensures that (C.4) always holds. Assuming $\omega_i(y_i) < \omega_j(y_j),\ \forall j \neq i$, this payment rule is equivalent to

$$\begin{aligned} p_j(y_i, y_{-i}) =&0,\ j \neq i, \\ p_i(y_i, y_{-i}) =&E_R[(R-\omega_i(y_i))^+] - \int_{y_i}^{\omega_i^{-1}(\min\{\omega_{-i}(y_{-i})\})} \Pr(R > \omega_i(z))\mathrm{d}z \\ =&E_R[\mathbf{1}_{\{\omega_i(y_i)<R\}}(R-\omega_i(y_i) + y_i - \omega_i^{-1}(\min\{\omega_{-i}(y_{-i}), u\}))]. \end{aligned}$$

This is the optimal payment rule, and coupled with the optimal quote rule I have an optimal mechanism.

Finally, I must confirm incentive compatibility and individual rationality of the above mechanism. Under this mechanism, TO_i's expected profit, reporting cost $\check{y}_i$, equals

$$\begin{aligned}v_i(\check{y}_i, y_i) =& E_{R,Y_{-i}}[\mathbf{1}_{\{\omega_i(\check{y}_i)<\min\{R,\omega_{-i}(Y_{-i})\}\}}[(R-\omega_i(\check{y}_i)+\check{y}_i-y_i)^+ \\ & -(R-\omega_i(\check{y}_i)+\check{y}_i-\omega_i^{-1}(\min\{\omega_{-i}(Y_{-i}),u\}))].\end{aligned}$$

When $\check{y}_i = y_i$, obviously $v_i(y_i, y_i) \geq 0$, so the participation constraint is satisfied.

When $\check{y}_i > y_i$, on the sample paths of $\{\omega_i(\check{y}_i) < R\}$, I must have $R-\omega_i(\check{y}_i)+\check{y}_i-y_i > 0$, and $(R-\omega_i(\check{y}_i)+\check{y}_i-y_i)^+ - (R-\omega_i(\check{y}_i)+\check{y}_i-\omega_i^{-1}(\min\{\omega_{-i}(Y_{-i}),u\})) = \omega_i^{-1}(\min\{\omega_{-i}(Y_{-i}),u\})-y_i > 0$ does not depend on $\check{y}_i$. However, by reporting higher than true cost $\check{y}_i > y_i$, TO_i receives this positive profit on fewer sample paths. Thus TO_i prefers reporting truthfully, $\check{y}_i = y_i$.

When $\check{y}_i < y_i$, on the sample paths of $\{\omega_i(\check{y}_i) < \omega_i(y_i) < \min\{\omega_{-i}(Y_{-i}),u\}\}$, since $R-\omega_i(\check{y}_i)+\check{y}_i-y_i < R-\omega_i(\check{y}_i)$, TO_i's profit is no higher than reporting truthfully. On the sample paths of $\{\omega_i(\check{y}_i) < \min\{\omega_{-i}(Y_{-i}),u\} < \omega_i(y_i)\}$, when $R-\omega_i(\check{y}_i)+\check{y}_i-y_i > 0$, the profit is $\omega_i^{-1}(\min\{\omega_{-i}(Y_{-i}),u\})-y_i < 0$; when $R-\omega_i(\check{y}_i)+\check{y}_i-y_i < 0$, the expected profit is

$$\begin{aligned}& -E_R[R-\omega_i(\check{y}_i)+\check{y}_i-\omega_i^{-1}(\min\{\omega_{-i}(Y_{-i}),u\})] \\ =& -\int_{\check{y}_i}^{\omega_i^{-1}(\min\{\omega_{-i}(Y_{-i}),u\})} \Pr(R > \omega_i(z))\mathrm{d}z < 0.\end{aligned}$$

In both cases, these sample paths would have reduced TO_i's expected profit. Thus TO_i again prefers reporting truthfully, $\check{y}_i = y_i$.

The above analysis shows that $v_i(\check{y}_i, y_i) \leq v_i(y_i, y_i)$ is always true, so the incentive constraints are satisfied. This concludes the proof. □

BIBLIOGRAPHY

Bagnoli, M., T. Bergstrom. 2005. Log-concave probability and its applications. *Economic Theory* **26** 445-469.

Barlow, R. E., A. W. Marshall, F. Proschan. 1963. Properties of probability distributions with monotone hazard rate. *Annals of Mathematical Statistics* **34** 375-389.

Benjaafar, S., W. L. Cooper, J. S. Kim. 2005. On the Benefits of Pooling in Production-Inventory Systems. *Management Science* **51** 548-565.

Cachon, G. P., F. Zhang. 2006. Procuring fast delivery: sole sourcing with information asymmetry. *Management Science* **52** 881-896.

Chen, F. 2007. Auctioning supply contracts. *Management Science* **53** 1562-1576.

Chen, R. R., R. O. Roundy, R. Q. Zhang, G. Janakiraman. 2005. Efficient auction mechanisms for supply chain procurement. *Management Science* **51** 467-482.

Chen, Y., S. Seshadri, E. Zemel. 2008. Sourcing through auctions and audits. *Production and Operations Management* **17** 121-138.

Corbett, C. J., X. de Groote. 2000. A supplier's optimal quantity discount policy under asymmetric information. *Management Science* **46** 444-450.

Dasgupta, S., D. F. Spulber. 1990. Managing procurement auctions. *Information Economics and Policy* **4** 5-29.

Duenyas, I., B. Hu, D. R. Beil. 2011. Simple auctions for supply contracts. *Working paper.*

Elmaghraby, W. J. 2000. Supply contract competition and sourcing policies. *Manufacturing and Service Operations Management* **2** 350-371.

Eppen, G. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science* **25** 498-501.

Eppen, G., L. Schrage. 1981. Centralized ordering policies in a multi-warehouse system with lead times and random demand. *TIMS Studies in the Management Sciences. Vol. 16, Multi-Level Production/Inventory Control Systems, Theory and Practice.* North-Holland, Amsterdam, Netherlands.

Feesers, Inc. v. Michael Foods, Inc., 591 F.3d 191 (3d Cir. Jan. 7, 2010)

Hanany, E., Y. Gerchak. 2008. Nash bargaining over allocations in inventory pooling contracts. *Naval Research Logistics* **55** 541-550.

Gerchak, Y., Q. He. 2003. On the relation between the benefits of risk pooling and the variability of demand. *IIE Transactions* **35** 1027-1031.

Ha, A. Y. 2001. Supplier-buyer contracting: asymmetric cost information and cutoff level policy for buyer participation. *Naval Research Logistics* **48** 41-64.

Hartman, B. C., M. Dror. 2003. Optimizing centralized inventory operations in a cooperative game theory setting. *IIE Transactions* **35** 243-257.

Huang, J. S., M. Ghosh. 1982. A note on strong unimodality of order statistics. *Journal of the American Statistical Association* **77** 929-930.

Jap, S. D. 2007. The impact of online reverse auction design on buyer-supplier relationships. *Journal of Marketing* **71** 146-159.

Krishna, V. 2002. *Auction Theory.* Elsevier, San Diego, CA

Kshetri, N., N. Dholakia. 2002. Determinants of the global diffusion of B2B e-commerce. *Electronic Markets* **12** 120-129.

Laffont, J., D. Martimort. 2002. *Theory of Incentives: the Principal-Agent Model.* Princeton University Press, Princeton, NJ.

Li, C., A. Scheller-Wolf. 2010. Push or pull? Auctioning supply contracts. *Production and Operations Management*, forthcoming.

Lovejoy, W. 2010. Bargaining chains. *Management Science* **56** 2282-2301.

Myerson, R. 1981. Optimal auction design. *Mathematics of Operations Research* **6** 58-73.

Rezende, L. 2009. Biased procurement auctions. *Economics Theory* **38** 169-185.

Stole, L. A. 2007. Price discrimination and competition. *Handbook of Industrial Organization Vol. 3* 2223-2286. Elsevier, Oxford, UK.

Stoll, N., S. Goldfein. 2007. Robinson-Patman Act: take two. *New York Law Journal* **November 20, 2007**.

The Economist. 2009. Japan's technology champions: invisible but indispensable. *The Economist* **November 7, 2009** 64-66.

Tunca, T. I., Q. Wu. 2009. Multiple sourcing and procurement process selection with bidding events. *Management Science* **55** 763-780.

U.S. Department of Commerce. 2005. Statistics for industry groups and industries: 2004. *Annual Survey of Manufactures* **December 2005**.

Wan, Z., D. R. Beil. 2009. RFQ auctions with supplier qualification screening. *Operations Research* **57** 934-949.

Yang, H., L. Schrage. 2009. Conditions that cause risk pooling to increase inventory. *European Journal of Operational Research* **192** 837-851.