Simultaneously Managing Procurement Costs and Risks

by

Zhixi Wan

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2009

Doctoral Committee:

Assistant Professor Damian R. Beil, Chair
Professor Izak Duenyas
Assistant Professor Volodymyr O. Babich
Associate Professor Uday Rajan

To my dear parents

Acknowledgements

I would like to first thank my advisor, Professor Damian Beil, for his guidance and great help during the past six years. It is him whom I have been keeping learning from. I also would like to sincerely thank my committee members, Professor Izak Dueyas, Professor Volodymyr Babich, and Professor Uday Rajan, for their constructive advice that has improved the quality of my dissertation.

My dissertation chapters are based on three research papers. Two of them (related to Chapters 2 and 3) are coauthored with Professor Damian Beil and the third one (related to Chapter 4) is coauthored with Professor Damian Beil and Professor Elena Katok (who is in Pennsylvania State University). I sincerely thank them for their hard work. Working with them has become one of the most enjoyable memory during the last six years.

I really appreciate all the opportunities that The Ross School of Business and The University of Michigan have offered me. A lot of thanks to the Operations and Management Science department, who admitted me in 2003, and to all the faculty members (during 2003 to 2009), who have educated me in various ways. Special thanks to my two Ph.D. Coordinators, Professor Roman Kapuscinski and Professor Hyun-Soo Ahn, who have been very helpful to me since my first day in the department.

All errors in this dissertation are mine.

Table of Contents

List of Figures

Chapter 1

Introduction

Transcending its more tactical beginnings in vertically integrated firms, today the procurement function serves as a vital gatekeeper to low-cost and low-risk inputs for a firm. My research focuses on theoretical and practical insights for strategic procurement. Globalization has made supply chains longer and more heterogeneous, making procurement increasingly complex. Strategic procurement involves identifying operationally capable suppliers, while also managing total-cost risks. The procurement function must also negotiate terms and prices with suppliers, involving economic concepts of incentives and private information. Reflecting these multifaceted challenges, my dissertation research combines operational processes such as supplier qualification screening with economic processes such as auctions and mechanism design for price discovery. This dissertation consists of three essays on two types of procurement risks: Two essays focus on supplier non-performance risks and the third one examines regional procurement cost risks.

A. Supplier Non-performance Risks

The procurement function must negotiate reasonable prices with suppliers and — equally important — ensure that contracts are awarded only to qualified suppliers. Contracting with unqualified suppliers can have dire consequences for the buyer, illustrated most recently by widespread recalls of products manufactured abroad . In the essay "RFQ Auctions with Supplier Qualification Screening" (Chapter 2), I use an optimal mechanism analysis to study how a buyer can best use a reverse auction in combination with supplier qualification screening processes to determine which qualified supplier will be awarded a contract. The main takeaway for procurement managers is that the standard industrial practice of fully qualifying all suppliers before the auction can be improved upon by judiciously delaying all or part of the qualification screening process until after the auction. To my knowledge, this is the first paper that models and analyzes an auction setting in which contracting is contingent upon suppliers' passing costly qualification screening, a feature that is virtually ubiquitous in practice. The second essay, "Procurement Auctions with an Incumbent and Partially Qualified Entrant" (Chapter 3), motivated by observations at a Fortune 100 manufacturer I interacted with, extends the main insights from the first essay to a setting where a buyer conducts an auction with her incumbent supplier and an entrant supplier. While the incumbent is known to be qualified for contract award, the entrant can get the contract only if he passes the buyer's qualification screening. My analysis finds that if the buyer delays the entrant's qualification screening until after the auction, the incumbent will strategically drop out of the auction early to forestall a bidding war, depending on his cost, degree of risk aversion, and the likelihood that the entrant would pass qualification screening. This work will contribute to the auction literature as the first study with bidders possessing asymmetric probabilities of passing qualification screening, a feature common in practice.

B. Regional Cost Risks

Buyers practicing global sourcing are increasingly aware of "non-price costs" covering logistics, shipping insurance and commissions . Such "non-price costs" are subject to cost shocks (e.g., port strikes, regulatory changes) in the region where suppliers are located. The third essay, "Bargaining Power and Supply Base Diversification" (Chapter 4), examines the following strategic decision of the buyer: Facing potential regional cost shocks, should the supply base include similar suppliers (selected from the same geographic region) or diversified suppliers (selected from different geographic regions)? Idealistically, diversifying hedges the buyer's risk by improving the chances that she can transact with a low-regional-cost supplier.

However, my study finds that diversification provides regional cost hedging at the peril of losing so much price parity across suppliers that suppliers with a regional cost advantage can opportunistically take a windfall profit by escalating the contract price. As a result, the buyer's optimal supply base design depends on her degree of bargaining clout, i.e., her ability to impose auction mechanisms to curtail suppliers' windfall profit-taking. The main message is that the more bargaining power the buyer has, the more she prefers a diversified supply base. To my knowledge, this work is the first study of supply base design to mitigate regional cost shocks. It contributes to the literature especially in that it exemplifies a principle that is often neglected in the operations literature, namely, a firm needs to carefully evaluate its bargaining (channel) clout when choosing its operational strategies.

Chapter 2

RFQ Auctions with Supply Qualification Screening

2.1. Introduction

The average U.S. manufacturer spends 40-60% of its revenue income to purchase goods and services (U.S. Department of Commerce 2005). Vital for most companies, the procurement function must negotiate reasonable prices with suppliers and — equally important — it must make reasonable efforts to ensure that contracts are made with *qualified* suppliers who are indeed able to fulfill the contract.

Contracting with unqualified suppliers can have dire consequences for the buyer. Numerous media reports recently exposing product safety issues in the U.S. have traced problems to suppliers failing to meet a buyer's requirements, resulting in dangerous lead paint in toys (Spencer and Casey 2007), tires lacking proper safety features (Welch 2007), and pet food containing noxious chemicals (Myers 2007). Recalls of faulty products produced by noncompliant suppliers have cost downstream firms millions of dollars and inflicted untold damage on their reputations. Menu Foods' market capitalization was slashed in half soon after the firm recalled over 60 million packages of dog and cat food in March 2007; contaminants were found to have originated in a supplier's raw ingredients (Myers 2007). Welch (2007) reports that New Jersey-based tire importer Foreign Tire Sales traced its consumers' complaints of faulty tires to an unauthorized design change made by its supplier, whose design engineer decided to omit gum strips, apparently unaware of their role in preventing tread separation. A surprised Foreign Tire Sales was forced to recall half a million tires, and may go bankrupt as a result. Vulnerabilities to supplier non-performance deepen for lengthly global supply chains, while price pressures and a multitude of global supply options compounds the procurement manager's challenge of discerning who is an able supplier and who is a charlatan.

To verify a supplier's qualification and thereby reduce the likelihood of non-performance, the procurement function must spend time and money vetting suppliers with qualification screening. Screening often involves references checks, financial status checks, surge capacity verification, and even site visits to supplier production facilities. Supplier employees might be interviewed and assessed, for example, to ensure that they understand all the engineering and safety requirements of the product in their charge. Checks can even reach back into the supplier's supply chain; to avoid mistakes like those uncovered in the recent pet food contamination scandal, Procter & Gamble now vets its suppliers' raw ingredients sources to ensure that they are reputable (Myers 2007).

Typically, qualification screening precedes price negotiations with suppliers. This is particularly common when buyers use qualification screening before a Request For Quotes (RFQ) auction for a well-defined good or service, a prototypical procurement setting on which this chapter focuses. (For details on procurement auctions in practice, see, for example, Jap 2003.) To convince the suppliers to bid aggressively, the buyer touts the fact that only the lowest price will win the contract when all participants in the auction are absolutely qualified to win the contract.

Yet committing to award business to the lowest bidder requires the buyer to spend significant time and resources screening *all* suppliers entering the auction. This chapter analyzes how much costly supplier qualification screening should be performed before the auction. At one extreme the buyer uses pre-qualification only, in which she fully screens all suppliers entering the auction and commits to awarding the contract directly to the lowest bidder, so suppliers in an open-descending auction (were this the format used) would bid down to their true cost (but, they may not have to), and the buyer directly awards the contract to the lowest bidder. At the other extreme the buyer uses post-qualification only, in which she screens suppliers only after seeing their bids in the auction; instead of screening all suppliers, the buyer homes in on the most promising bidders and screens them in sequence until finding one who is qualified. But, this comes with a tradeoff: during the post-auction screening process, a bid will have to be discarded if the bidder is found to be unqualified for the contract, hence the winner will be the lowest *qualified* bidder, if any. In this chapter, these two extremes and mixtures of the two in which suppliers are partially screened before the auction are considered.

Our study appears to be the first to model and analyze an auction setting in which contracting is contingent upon passing costly supplier qualification screening, a common feature of RFQs in practice. The central tradeoff of the research problem is this: while delaying some or all qualification screening until after the auction saves the buyer qualification screening costs, doing so increases contract payments and also risks non-transaction (turning to a costly option outside of the auction, such as internal production) if all suppliers in the auction turn out to be unqualified. By optimally balancing these tradeoffs, the buyer can significantly reduce total procurement cost (qualification cost plus contract payment) by judiciously delaying all or part of supplier qualification screening until after the auction, provided the qualification screening is not too costly (prompting defection to the outside option) nor too cheap (making total pre-qualification the best option). Our analysis endogenizes the level of pre- and post-qualification as well as the negotiation mechanism chosen by the buyer.

Section 3.2 provides a literature review, and §3.3 introduces the model and assumptions. Through a mechanism design approach, the optimal auction and post-qualification structure is derived in §2.4.2. Behavior of the optimal balance between pre- and post-qualification is characterized in §2.4.3. Section 2.5 provides numerical illustrations, and §2.6 discusses practical considerations and extensions. Proofs are provided in §2.8.

2.2. Literature Review

Many practical decisions faced by procurement managers have been addressed by academic research: contract type, such as fixed price versus cost plus; negotiation framework such as auction versus face-to-face; and competition type such as sole versus dual sourcing. See Elmaghraby (2000) for a detailed survey on procurement studies in economics and operations management. The present chapter focuses on a procurement situation in which a buyer holds an auction to award a contract to one of several potential suppliers. There is a sizeable literature on auctions — books by Krishna (2002) and Milgrom (2004) provide excellent treatments and detailed references — but only a handful of such studies include the processes which occur before and after an auction. Typically these processes seek to mitigate the risk of consummating a transaction in which one or more parties does not obtain what it expected, called *non-performance* risk. Note that non-performance could describe, for example, an item falling short of the winner's expectations, or a winner failing his obligations to the auctioneer. The supplier qualification process central to this chapter — and to our knowledge novel in the literature — is one such process meant to mitigate the risk of supplier non-performance in a procurement auction.

One type of non-performance in forward auctions is the winner's failure to pay. Papers dealing with this issue have looked at the ability of bidders to borrow money and the ensuing possibility of broke winners (Zheng 2001), and the use of deposits or fees forfeited to the auctioneer in the event a winning bid is reneged (Rothkopf 1991, Waehrer 1995). In a procurement auction context, surety bonds (analogous to a bid deposit) to partially offset non-performance costs are examined by Calveras et al. (2004), while Braynov and Sandholm (2003) study a buyer who is unable to directly verify the "trustworthiness" of suppliers, but knowing the form of the suppliers' cost functions can design bidding options which cause each supplier to reveal themselves as either a high or low trustworthy type, allowing the buyer to estimate the expected utility of contracting with that supplier. In contrast, this research assumes that the buyer verifies (at a cost) that a supplier is qualified up to some threshold prior to contracting (she will not contract with unqualified types). Practitioners we have spoken with use qualification and surety bonds in tandem, the former to proactively avoid problems (the focus and main contribution area of this chapter), the latter to partially recoup costs if problems arise.

A second type of non-performance is misevaluation of the item. For example, costly bid preparation (or costly entry, or due diligence) plays a central role in forward auctions for non-standard, complex items such as an entire company or its assets. To encourage participation in auctions where bidders trade off their bid preparation costs (possibly millions of dollars) against their anticipated likelihood of winning the item, Ye (2007) suggests inviting only bidders whose bid in an initial, assumed costless round of bidding signals that they stand a good chance of winning the item. While Ye finds that screening out low value bidders can promote competition in a complex item auction by limiting bidders' unnecessary bid preparation costs, this chapter examines screening costs borne by the auctioneer (buyer) and find that screening out unqualified suppliers promotes competition by increasing the likelihood that each bid in an RFQ auction will be qualified and therefore eligible for contracting.

In our context of relatively well-specified RFQ auctions this chapter assumes that a supplier's bid indicates the value offered to the buyer should that supplier be deemed qualified. This allows the buyer to delay all or part of the qualification process until after the auction, at which point — with bids in hand — she can home in on the suppliers offering the highest value (who may or may not turn out to be qualified). Other auction theoretic papers have focused on situations where it is costly for the auctioneer to estimate even the value offered in the suppliers' bids. In such situations, the buyer could employ a sequential search model whereby suppliers are communicated with and their bids evaluated until either finding a supplier whose cost is sufficiently low or exhausting the supply pool (McAfee and McMillan 1988). In an effort to explain unconsummated Request For Proposals auctions documented by Snir and Hitt (2003), Carr (2003) models an auction for professional services where proposals (bids) are difficult to compare; in his model, faced with high evaluation costs after the auction, the auctioneer might simply forgo evaluating *any* proposals in favor of an outside option. The buyer sometimes turns to an outside option in our model, but for different reasons — either due to reserve prices, or after disqualifying all suppliers invited to the auction.

Methodologically, our study is also related to the screening literature, in particular studies such as Feinberg and Huber (1996) which assume that some form of screening can be performed cheaply (e.g., bids can be observed) relative to more costly forms of screening (e.g., qualification). In our context, partial qualification screening impacts the extent to which bidders compete in the auction, by creating randomness in the number of qualified bidders in the auction. For auctions with an uncertain number of bidders, optimal mechanisms and equilibrium bid functions have been respectively derived by McAfee and McMillan (1987) and Harstad et al. (1990), but neither studies qualification processes or the attendant possibility of having to turn to an outside option.

Our study is related in spirit to multiple dimensional auctions (Che 1993, Beil and Wein 2003, Chen 2007) in that both seek to take non-price factors into account. Although the goals of a multiple dimension auction and qualification processes are related, they are distinct: multiple dimension auctions serve as tools to better express the value of non-price abilities of suppliers, such as quality. Qualification processes seek to verify the ability of a supplier to deliver on the promises expressed by his bid, be they promises on price or any other dimensions.

2.3. Model

Consider a risk-neutral, cost-minimizing buyer seeking to award an indivisible contract to a *qualified* supplier. A supplier is called qualified if the buyer is willing to transact with the supplier without performing additional due diligence to verify this supplier satisfies all pre-award requirements. These requirements vary widely in practice depending on the buyer's needs and the contract type (see Leenders and Fearon 1997 for a discussion of purchasing processes). The constituent requirements themselves exhibit varying degrees of standardization. Among the more standard requirements are the need to verify the supplier's reputability (e.g., through published ratings) and ability to ramp up production. Not all requirements are so transparent, as qualification can encompass relational aspects that are difficult to codify; e.g., a just in time manufacturer we spoke with visits supplier management to ensure they and the supplier "see eye to eye on lean principles" before awarding contracts. These verification processes take time and can be costly, particularly if involving visits to distant supplier facilities. As is common in industry, this verification is referred to as the *qualification process*, or the act of *qualifying* a supplier. A supplier is described as *qualified* once he successfully passes the qualification process.[1]

To begin formalizing the model, we imagine a continuum of qualification requirements, with zero representing requirements that every supplier satisfies and one representing requirements that virtually no supplier satisfies. We let q_0 be the buyer's *qualification threshold*, a scalar between zero and one representing the pre-award requirements. For each supplier i, we will define his *qualification level* q_i as the maximum qualification threshold that supplier i can pass. Due to opaque requirements set by the buyer, such as needing to see eye to eye on lean principles, or to create rapport with the buyer's internal customer (e.g., the engineering department), supplier i does not precisely know its true qualification level q_i, but the buyer and supplier share the common belief that q_i is distributed according to distribution H on domain $[0,1]$. This setup could model, for instance, a buyer deciding to outsource a portion of its production currently done in-house, facing new suppliers she knows little about and who in turn know little about her (possibly idiosyncratic) qualification requirements. The strictness of the buyer's pre-award requirements is captured by $1 - H(q_0)$, the probability that $q_i \geq q_0$.

The buyer and each supplier responding to the RFQ are equally unsure of the supplier's qualification until costly qualification verification is undertaken by the buyer. If $q_i \geq q_0$,

[1] Because the entire notion of being "qualified" is based on the buyer's specific requirements for the contract up for bid, being deemed qualified does not necessarily imply absolutely zero risk of non-performance. For instance, when purchasing a simple, cheap, and non-critical indirect good (such as office supplies), a buyer might be satisfied with relatively light screening of the supplier's status and abilities prior to contracting.

the buyer's qualification process on supplier i would reveal that supplier i is qualified. The *qualification cost* the buyer would incur to do so is denoted by K, the total cost to the buyer of verifying that an individual supplier meets all requirements to be deemed qualified. For example, K may include the cost of purchasing and testing supplier products, travel to supplier facilities abroad, etc. For simplicity, this chapter assumes that K is the same for all suppliers. This assumption is most appropriate when suppliers are similar, at least in terms of the cost drivers of qualification, such as distance from the buyer or number of units that the buyer must purchase to run a test sample. A more general model would allow all the qualification verification cost to differ for different suppliers. This ex ante asymmetric suppliers extension is left for future research.

On the other hand, if $q_i < q_0$, supplier i would be rejected during the buyer's qualification process after failing to meet a requirement. In this latter case, how much cost would the buyer incur? Assuming that the requirements are nested (passing a larger threshold implies passing a smaller threshold, but not vice-versa), the cost strictly increases with q_i and approaches K as q_i approaches q_0. To streamline the exposition, this chapter further assumes that the cost is linear and normalized such that a threshold of zero costs zero to verify, implying that weeding out an unqualified supplier i costs the buyer $\frac{q_i}{q_0}K$. This is without loss of generality, because any nonlinear and strictly increasing cost function of q_i over $[0, q_0]$ can be renormalized to be linear by redefining q_i and renormalizing the distribution H. Note that in this dissertation the terms "increase" and "decrease" are used in their weak sense.

Each supplier i privately knows his cost to fulfill the contract, x_i, which he observes perfectly prior to the auction. Cost x_i is distributed according to a commonly known distribution F with density f on domain $[0, 1]$, which for simplicity is assumed statistically independent of other suppliers' costs (§2.6.1 relaxes this assumption) and i's own qualification level. In reality, a supplier who is likely to be qualified might be expected to have relatively high costs if qualification requires that costly spare capacity be kept on hand for an ability to ramp up production in the face of surge orders from the buyer; on the other hand, lower costs might prevail if qualification includes lean principles which impart efficiency. Such issues are left for future research. For simplicity this chapter will assume that all distributions have positive and continuous densities over their domains.

The buyer utilizes an auction to extract private cost information from the risk-neutral suppliers, who — as is standard in the auction literature — are assumed to be fully rational players following a Bayesian Nash bidding equilibrium in which bidding is assumed to be costless. To capture practical demands on the buyer's resources such as auction participant training and technical support, this chapter assume that the buyer will have a finite number of bidders in the auction, denoted by n, which is treated as fixed and exogenous in our analysis.

We now describe the three-stage procurement process the buyer uses for supplier qualification screening (qualification checks) and supplier cost screening (competitive auction).

Pre-qualification stage. The buyer announces n *pre-qualification thresholds* $\underline{q_i} \leq q_0$ ($i = 1, \ldots n$), where $\underline{q_i}$ represents the requirements the i^{th} auction participant must meet to participate in the auction. Each pre-qualification threshold $\underline{q_i}$ represents part ($\underline{q_i} < q_0$) or all ($\underline{q_i} = q_0$) of the requirements that a supplier must satisfy in order for the buyer to consider them qualified for contract award. The pre-qualification thresholds are not assumed to be identical, and can be different for different auction participants (the less general but perhaps more equitable case of symmetric pre-qualification thresholds is analyzed in §2.4.4). The buyer faces a pool of ex ante symmetric suppliers. To find its i^{th} auction participant, the buyer verifies requirements of supply pool members one at a time until finding one who passes pre-qualification threshold $\underline{q_i}$. (For simplicity this chapter assume an infinite supplier pool.) After doing so for each $i = 1, \ldots, n$, the buyer has n pre-qualified suppliers, where the i^{th} participant (say bidder i) has qualification level $q_i \geq \underline{q_i}$.

Auction stage. The n pre-qualified bidders participate in a price-only auction. When bidding, bidder i, who passed pre-qualification level $\underline{q_i}$, is estimated to be qualified with probability

$$\beta_i = \frac{1 - H(q_0)}{1 - H(\underline{q_i})} .$$

This probability, referred as bidder i's *qualification probability*, is strictly increasing in $\underline{q_i}$. In particular, β_i equals its upper bound 1 when $\underline{q_i} = q_0$, and equals its lower bound $\underline{\beta} \triangleq 1 - H(q_0)$ when $\underline{q_i} = 0$. In the sequel, the β_i's figure prominently in the analysis and are, for convenience, directly used as the buyer's pre-qualification threshold decisions. This is without loss of generality: since H is strictly increasing, there is a one-to-one correspondence between $\underline{q_i}$ and β_i, and we define function $q : [\underline{\beta}, 1] \to [0, q_0]$ such that

$$\underline{q_i} = q(\beta_i) \triangleq H^{-1}[1 - \frac{1 - H(q_0)}{\beta_i}] .$$

We primarily use notation $q(\beta_i)$ in lieu of $\underline{q_i}$. We use the bold notation $\boldsymbol{\beta}$ for the vector of qualification probabilities $(\beta_1, \beta_2, ..., \beta_n)$, and use $\underline{\boldsymbol{\beta}}$ for the vector having $\beta_i = \underline{\beta}$ for all $i = 1, \ldots n$. Unlike a traditional auction where bidding is directly followed by a contract award decision, this auction yields a *post-qualification sequence*, which specifies (based on auction bids and bidder qualification probabilities) an ordered subset of bidders who will be post-qualified in the next and final stage of procurement, the post-qualification stage.

Post-qualification stage. After the auction the buyer post-qualifies suppliers up to qualification threshold q_0 according to the post-qualification sequence, until finding a qualified bidder to contract with, or disqualifying all bidders included in the post-qualification sequence. In the latter case, the buyer turns to her "outside option" at a cost of C_o. Cost C_o describes an option outside the auction, but could, for example, correspond to in-house production.

The buyer faces two key decisions in the above procurement process: what pre-qualification thresholds to set; and what auction and post-qualification mechanism to use to find (if possible) a desirable, qualified bidder to contract with. Our analysis first finds the optimal auction and post-qualification mechanism, and then works backwards to characterize the optimal pre-qualification thresholds. Before moving to this analysis, we conclude this section by discussing the main tradeoffs involved.

The choice of pre-qualification thresholds affects the balance between qualification costs on the one hand, and contract payment on the other. If $\boldsymbol{\beta} = (1, \ldots, 1)$, the buyer performs all due diligence prior to the auction — the situation tacitly assumed in traditional auction theory. In such a case the buyer incurs qualification costs for *at least* n (possibly more if some are rejected en-route to being qualified) suppliers prior to the auction, and awards the contract directly to the most attractive bidder after the auction. On the other hand, $\boldsymbol{\beta} = \underline{\boldsymbol{\beta}}$ models postponement of all due diligence; after the auction, the buyer only pays qualification costs until finding the first qualified bidder or turning to her outside option. Clearly the expected total cost of qualifying suppliers is greater in the case where $\boldsymbol{\beta} = (1, \ldots, 1)$. But the buyer also must consider the expected costs of contracting and non-transaction, for which the cost relationship can be reversed — consider the following toy example with two bidders showing why more pre-qualification reduces the expected costs of contract payment and non-transaction.

Suppose that bidder 1's and bidder 2's true costs are \$100,000 and \$125,000, respectively, and the auction reserve price is \$150,000. For simplicity, suppose that the auction mechanism induces truthful bidding. If bidder 1 has been fully qualified ($\underline{q_1} = q_0$), he is the obvious candidate for contract award. However, to ensure truthful bidding, the mechanism's award and payment rules must be designed taking bidder incentives into account. An interesting dynamic arises regarding payment to bidder 1. Should the mechanism pay bidder 1 precisely \$125,000, the cost of the losing bidder? If $\underline{q_2} = q_0$, the answer — according to the incentive compatibility of the Vickrey auction — is yes, which is based on the buyer's next best alternative being bidder 2. However, things are different if $\underline{q_2} < q_0$. Intuitively the mechanism must compensate bidder 1 for the fact that bidder 2 might be unqualified for the contract, and hence possibly a non-viable alternative for the buyer. Thus, the fact that the buyer delayed qualification on bidder 2 means that she must pay bidder 1 more than \$125,000 for the contract. On the other hand, if bidder 2's true cost had been less, say \$80,000 instead of \$125,000, the buyer might find it worthwhile to first post-qualify bidder 2, which may or may not turn out favorably for the buyer — during post-qualification she might discover that bidder 2 is unqualified ($\underline{q_2} \le q_2 < q_0$), at which point she has wasted post-qualification money on bidder 2 and her next best option will be bidder 1, whose payment is then the reserve price. The optimal mechanism analysis of the next section determines how the buyer should run an auction and post-qualification process when suppliers are possibly unqualified, where clearly the buyer's expected payment increases with the amount of qualification that is delayed until after the auction.

2.4. Analysis

In this section we first fix the qualification screening strategy (equivalently, fix the qualification probabilities $\beta_1, \ldots, \beta_n$) and compute the buyer's total (contracting plus qualification) expected costs. We begin by deriving $\overline{\mathrm{PRE}}$, the expected pre-qualification cost, in §2.4.1. In §2.4.2 we derive $\overline{\mathrm{PAY}}$, $\overline{\mathrm{POST}}$, and $\overline{\mathrm{NT}}$, respectively the expected payment to the auction winner, the expected post-qualification cost, and the expected non-transaction cost. These costs are derived via an optimal mechanism analysis. Having derived the total expected cost as a function of the qualification screening policy, we then characterize the optimal qualification screening strategy in §2.4.3. Following this general analysis, §2.4.4 discusses the special (but more "fair") case of symmetric pre-qualification thresholds.

2.4.1 Expected Pre-Qualification Cost

As explained in §3.3, we assume an infinite supply pool from which the buyer samples until finding n suppliers that pass their respective pre-qualification thresholds. The buyer begins by sampling suppliers one at a time until finding a supplier whose qualification level is at least $q(\beta_1)$. This "successful" supplier is admitted to the auction as bidder 1. This process is then repeated for $i = 2, \ldots, n$. The ordering of the β_i's is not important, because "failures", i.e., suppliers found to have a qualification level below $q(\beta_i)$, are permanently discarded for being unqualified for the contract, and the supply pool population is infinite and samples (supplier qualification levels) are independent. The expected pre-qualification cost is comprised of n pre-qualification "successes" plus geometrically distributed numbers of pre-qualification "failures" before each success. Because the qualification cost is linear in the amount of qualification performed, for each success the buyer pays $\frac{q(\beta_i)}{q_0}K$ to pre-qualify a supplier with qualification level $y \geq q(\beta_i)$ up to the pre-qualification threshold $q(\beta_i)$. For each failed pre-qualification on a supplier with qualification level $y < q(\beta_i)$, the buyer pays $\frac{y}{q_0}K$. Since y is a random variable and meets or exceeds the pre-qualification level $q(\beta_i)$ with probability $1 - H(q(\beta_i))$, the buyer expects to pay

$$\overline{\text{PRE}} = \sum_{i=1}^{n} \{ \frac{q(\beta_i)}{q_0} K + \left[\frac{1}{1 - H(q(\beta_i))} - 1 \right] \int_{y=0}^{q(\beta_i)} \frac{y}{q_0 H(q(\beta_i))} K dH(y) \}. \tag{2.1}$$

2.4.2 Expected Costs in Auction and Post-Qualification

The auction and post-qualification stages together comprise a mechanism of awarding a contract to one of n bidders or to the buyer's outside option. The buyer designs this mechanism to minimize her expected auction and post-qualification cost $\overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}$. The mechanism extracts private cost information from the suppliers, as is common in traditional mechanism design, but also involves post-qualification screening, which is not traditional in the mechanism design literature. The auction and post-qualification mechanism has three components: a set of possible messages (or "bids") for each bidder; a rule that describes how the buyer sequences the bidders for post-qualification; and a payment rule that maps bids to an amount transferred between the buyer and bidders. The post-qualification sequencing rule is the analogue of the allocation rule in standard mechanism design problems absent post-qualification; it captures the fact that the allocation decision itself is not entirely in the buyer's hands, as post-qualification outcomes (the realization of q_i's) affect whether or not a bidder will actually be qualified for contract award.

Both rules (sequencing and payment) are functions of the messages sent by bidders. Thanks to the revelation principle (e.g., Gibbard 1973, Green and Laffont 1977, Dasgupta et al. 1979), given any mechanism and an equilibrium for that mechanism, there exists an outcome equivalent "direct" mechanism in which it is an equilibrium for each bidder to bid his true cost. This allows us to restrict our optimal mechanism search to direct mechanisms; in what follows the set of messages sent is assumed without loss of generality to be just the bidder cost vector $\boldsymbol{x} \triangleq (x_1, x_2, \ldots, x_n)$. The remainder of this subsection formalizes the mechanism design problem.

We will let $\mathrm{Q}(\boldsymbol{x})$ denote the post-qualification sequencing rule, an ordered subset of $\{1, \ldots, n\}$ representing the bidders that will be post-qualified in a determined sequence. For example, a post-qualification sequence $(3, 1, 5, 2)$ for $n = 5$ means that the buyer would post-qualify bidder 3, bidder 1, bidder 5, and bidder 2 in sequence until finding the first qualified bidder or disqualifying all these four bidders, without post-qualifying bidder 4 at all. For a given vector of costs $\boldsymbol{x}$, using Q and $\boldsymbol{\beta}$ we can compute Δ_i the probability that bidder i wins the contract as

$$\Delta_i(\boldsymbol{x}) = \beta_i \prod_{\substack{j \text{ ranks ahead} \\ \text{of } i \text{ in } \mathrm{Q}(\boldsymbol{x})}} (1 - \beta_j), \text{ if } i \in \mathrm{Q}(\boldsymbol{x}); \quad \Delta_i(\boldsymbol{x}) = 0, \text{ if } i \notin \mathrm{Q}(\boldsymbol{x}). \tag{2.2}$$

This assumes that the buyer uses a deterministic sequencing rule, although this is without loss of generality; see the proof of Proposition 1. The mechanism's payment rule, which we will denote as $\mathrm{M}(\boldsymbol{x}) \in \mathbb{R}^n$, is the expected monetary transfer from the buyer to the bidders.

Of course, in order to constitute a viable direct mechanism, Q and M must satisfy constraints ensuring both incentive compatibility (truthful bidding in equilibrium) and individual rationality (bidders expect non-negative profits by participating). To write down these constraints for bidder i, first define $\boldsymbol{x}_{-i} = (x_1, x_2, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$, and let $\mathbf{F}_{-i}$ be the joint distribution of $\boldsymbol{x}_{-i}$. Define $\mathrm{m}_i(z_i) \triangleq \int_{\boldsymbol{x}_{-i}} \mathrm{M}_i(z_i, \boldsymbol{x}_{-i}) d\mathbf{F}_{-i}(\boldsymbol{x}_{-i})$ to be the expected payment to bidder i when i reports cost z_i and all others report their true costs, and likewise define $\delta_i(z_i) \triangleq \int_{\boldsymbol{x}_{-i}} \Delta_i(z_i, \boldsymbol{x}_{-i}) d\mathbf{F}_{-i}(\boldsymbol{x}_{-i})$ to be the probability bidder i is awarded the contract when i reports z_i to the auctioneer and all others report their true costs. Under truthful bidding, if all bidders $j \neq i$ report their true costs, bidder i's expected profit from reporting z_i is given by $\mathrm{m}_i(z_i) - \delta_i(z_i)x_i$. Thus, incentive compatibility and individual rationality can be expressed as

$$U_i(x_i) \triangleq \mathrm{m}_i(x_i) - \delta_i(x_i)x_i = \max_{z_i \in [0,1]} \mathrm{m}_i(z_i) - \delta_i(z_i)x_i, \qquad (2.3)$$

$$U_i(x_i) \geq 0. \qquad (2.4)$$

$U_i(x_i)$ is the expected profit of bidder i in equilibrium. Constraint (2.3) is called the *incentive compatibility* constraint, and constraint (2.4) is called the *incentive compatibility* constraint. For Q and M satisfying (2.3)-(2.4), we now compute $\overline{\mathrm{NT}}$, $\overline{\mathrm{POST}}$, and $\overline{\mathrm{PAY}}$.

Expected non-transaction cost. If no bidder wins the contract, the buyer must turn to the outside option at cost C_o. Letting $\mathbf{F}$ denote the joint distribution of $\boldsymbol{x}$, from (2.2) it is easy to see that $\overline{\mathrm{NT}} = \int_{\boldsymbol{x}} [1 - \sum_{i=1}^{n} \Delta_i(\boldsymbol{x})] C_o d\mathbf{F}(\boldsymbol{x})$.

Expected post-qualification cost. Because $\Delta_i(\boldsymbol{x})$ is the probability that i wins the contract, $\Delta_i(\boldsymbol{x})/\beta_i$ must be the probability that bidder i undergoes post-qualification. Taking the expected cost of post-qualifying bidder i, and dividing this cost by β_i and calling the result $a(\beta_i)$, yields

$$a(\beta_i) \triangleq \frac{1}{1 - H(q_0)} \int_{q_i = q(\beta_i)}^{q_0} \frac{q_i - q(\beta_i)}{q_0} K dH(q_i) + \frac{q_0 - q(\beta_i)}{q_0} K. \qquad (2.5)$$

Thus, the expected post-qualification cost is $\overline{\mathrm{POST}} = \int_{\boldsymbol{x}} \sum_{i=1}^{n} \Delta_i(\boldsymbol{x}) a(\beta_i) d\mathbf{F}(\boldsymbol{x})$.

Expected payment. Finding the optimal mechanism is aided by expressing the buyer's expected payment (up to an additive constant) as a function of award probabilities (as is standard in mechanism design analyses). Applying the envelope theorem (see Milgrom

2004, p67) to (2.3) implies that $U_i'(x_i) = -\delta_i(x_i)$ for $x_i \in [0,1]$. Treating this as a differential equation and using the expected profit at $x_i = 1$ as an integration constant yields an equation for bidder i's equilibrium expected profit:

$$\mathrm{m}_i(x_i) - \delta_i(x_i)x_i = \mathrm{m}_i(1) - \delta_i(1) + \int_{x_i}^{1} \delta_i(z_i)dz_i. \tag{2.6}$$

Solving for m_i, the buyer's expected payment to bidder i given the reported cost x_i, and then integrating over $\boldsymbol{x}$ and summing over n yields

$$\overline{\mathrm{PAY}} \;=\; \sum_{i=1}^{n}[\mathrm{m}_i(1) - \delta_i(1)] + \sum_{i=1}^{n}\int_{\boldsymbol{x}} [\Delta_i(\boldsymbol{x})\psi(x_i)]\, d\mathbf{F}(\boldsymbol{x}), \qquad \text{where } \psi(x_i) \triangleq x_i + \frac{F(x_i)}{f(x_i)}. \tag{2.7}$$

The value $\psi(x_i)$, which equals bidder i's true cost, x_i, plus $\frac{F(x_i)}{f(x_i)}$, the informational rent accruing to bidder i's private knowledge of his cost x_i, is commonly referred to as i's *virtual cost* in the mechanism design literature.

Mechanism design program. If (2.6) holds, it is easy to check that incentive compatibility constraint (2.3) is guaranteed if the post-qualification sequencing rule is such that $\delta_i(z_i)$ is decreasing in $z_i \in [0,1]$. (See, for example, Myerson 1981 Lemma 2 for a similar result.) Furthermore, from (2.6), we have that individual rationality constraint (2.4) holds as long as the payment and post-qualification sequencing rules satisfy $\mathrm{m}_i(1) - \delta_i(1) \geq 0$. Clearly, equation (2.7) implies that a cost-minimizing buyer will set $\mathrm{m}_i(1) - \delta_i(1) = 0$. Thus, to find the cost-minimizing direct mechanism $(\mathrm{Q}^*, \mathrm{M}^*)$ employed by the buyer in the auction and post-qualification stage, it is sufficient to solve the following program,

$$\min_{\mathrm{Q,M}} \quad \underbrace{\sum_{i=1}^{n}\int_{\boldsymbol{x}} \Delta_i(\boldsymbol{x})\,[\psi(x_i) + a(\beta_i) - C_o]\, d\mathbf{F}(\boldsymbol{x}) + C_o}_{\overline{\mathrm{PAY}} \;+\; \overline{\mathrm{POST}} \;+\; \overline{\mathrm{NT}}} \tag{2.8a}$$

$$\text{Subject to} \quad \mathrm{m}_i(x_i) \;=\; \delta_i(x_i)x_i + \int_{x_i}^{1} \delta_i(z_i)dz_i \;\; \forall i, \tag{2.8b}$$

and verify that at this solution $\delta_i(z_i)$ is decreasing in $z_i \in [0, 1]$ for all i. Program (2.8) ignores the pre-qualification cost, which is considered sunk by the auction and post-qualification mechanism. Note that if the buyer selects the pre-qualification thresholds to be $\underline{q_i} = q_0$ for all i — i.e., she performs all qualification screening prior to the auction — then $\beta_i = 1$ for all i and the above program devolves into the standard optimal auction analysis of Myerson (1981).

The Optimal Auction and Post-Qualification Mechanism

We now find an optimal solution $(\mathrm{Q}^*, \mathrm{M}^*)$ to program (2.8). To ensure that $\delta_i(z_i)$ is decreasing, we will assume that the virtual cost $\psi(x_i)$ is strictly increasing in the true cost x_i (see the proof of Proposition 1). This standard, technical assumption is satisfied, for example, if F is logconcave; see Bagnoli and Bergstrom (2005) for details about logconcave functions, which include uniform, normal, logistic and exponential distributions.

Proposition 1 *An optimal direct, individually rational, and incentive compatible auction and post-qualification mechanism* (Q^*, M^*) *that minimizes* $\overline{PAY} + \overline{POST} + \overline{NT}$ *is as follows.*

- *Set* $Q^*(\boldsymbol{x})$ *such that the buyer only post-qualifies those bidders whose* x_i *and* β_i *are such that* $\psi(x_i) + a(\beta_i) \leq C_o$*, and sequence all such bidders for post-qualification according to ascending value of* $\psi(x_i) + a(\beta_i)$*, breaking ties randomly.*

- *Set* $M^*(\boldsymbol{x})$ *such that the buyer pays the contract winner, i.e., the first bidder (if any) deemed qualified by post-qualification, a payment of*

$$x_i + \int_{x_i}^{1} \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i, \tag{2.9}$$

  and pays nothing to all other bidders, where $\Delta_i^*(\boldsymbol{x})$ *is calculated via equation (2.2) given* $Q^*(\boldsymbol{x})$*.*

Note that Proposition 1 implies that the buyer should set bidder-specific reserve prices r_i, $i = 1 \ldots n$, such that $r_i = \max\{\min\{\psi^{-1}[C_o - a(\beta_i)], 1\}, 0\}$, and promise not to post-qualify any bidder whose bid falls beyond the reserve price for that bidder. When bidders have equal β_i's (symmetric pre-qualification), a common reserve price for all bidders is optimal; see §2.4.4.

Adjusted virtual costs account for post-qualification. Proposition 1 indicates that the buyer should sequence bidders for post-qualification according to *adjusted virtual cost,* which we define to be $\psi(x_i) + a(\beta_i)$. The added adjustment $a(\beta_i)$ captures both the cost of post-qualifying bidder i and the risk that bidder i fails post-qualification. Intuitively, this adjustment should be lower if i already survived very strict pre-qualification, and indeed $a(\beta_i)$ does decrease in β_i. Thus, the optimal sequencing rule takes into account not only virtual cost, but also favors bidders who are more likely to be qualified by applying adjustments when sequencing the bids for post-qualification. If the buyer uses an identical pre-qualification hurdle for all bidders (identical β_i's), the optimal sequencing rule simply post-qualifies bidders according to cost x_i (bidders' identical $a(\beta_i)$'s wash out, and ψ is increasing); see §2.4.4, where the case of symmetric pre-qualification is discussed.

Payment to winner depends on competitors' virtual costs *and* qualification probabilities. Proposition 1 implies that to induce all bidders to bid their true costs, the buyer finds it optimal to commit to pay zero to all bidders save the contract winner, say bidder i, who is paid his cost x_i plus a markup $\int_{x_i}^{1} \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i$. This payment can be rewritten (see the proof of Proposition 1) as

$$\sum_{j=1}^{t} \psi^{-1}[\psi(x_{i_j}) + a(\beta_{i_j}) - a(\beta_i)]\beta_{i_j} \prod_{k=1}^{j-1}(1-\beta_{i_k}) + r_i \prod_{k=1}^{t}(1-\beta_{i_k}) \text{ (if } t \geq 1) \text{ or } r_i \text{ (if } t = 0), \quad (2.10)$$

where i_j $(j = 1, ..., t)$ is the index of the j^{th} bidder out of t bidders sequenced after bidder i for post-qualification according to the optimal sequencing rule such that $\psi(x_{i_j}) + a(\beta_{i_j}) \leq$

$\psi(r_i)+a(\beta_i)$. When the buyer reaches bidder i in the post-qualification stage, she has at least $t+1$ options available to her besides bidder i, namely the t bidders described above, plus the outside option. Bidder i could have increased his bid and still been preferred by the buyer to each of these $t+1$ options. When computing the payment to bidder i, the buyer rewards bidder i for this. In particular, the value $\psi^{-1}[\psi(x_{i_j}) + a(\beta_{i_j}) - a(\beta_i)]$ is exactly the amount that bidder i could have bid and still been preferred to bidder i_j. This value is weighted by $\beta_{i_j} \prod_{k=1}^{j-1}(1 - \beta_{i_k})$, exactly the probability that the buyer would end up contracting with bidder i_j if the buyer were to, upon reaching bidder i in the qualification list, forgo qualifying bidder i. Similarly, the reserve price r_i is the maximum amount bidder i could have bid and still been preferred to the outside option, which is weighted by $\prod_{k=1}^{t}(1 - \beta_{i_k})$ (if $t \geq 1$) or 1 (if $t = 0$), exactly the probability that the buyer would end up using the outside option if the buyer were to, upon reaching bidder i in the qualification list, forgo qualifying bidder i. As a check, when $\beta_i = 1$ for all i the payment reduces to the smallest rejected bid x_{i_1} (or the reserve price, whichever is smaller), which is exactly the payment rule of the traditional second-price auction with a reserve price.

Optimal Mechanism Implementations

Auction formats differ in the amount of visibility bidders have on competitors' bids. At one extreme is the sealed-bid format, where bidders submit a single bid known only to the bidder and auctioneer. At the other extreme is an open-bid format, where bidders can see competing bids and respond by updating their own bid. In the following, we show that the optimal auction can be implemented in either of these two format types.

Sealed-bid format. As is straightforward to see given Proposition 1, the buyer can conduct a sealed-bid auction after directly announcing the optimal sequencing rule as specified in Proposition 1 and the payment rule per (2.10) above; this will lead to an equilibrium where all bidders truthfully bid their costs.

Open-bid format. In our open-bid format, bidding proceeds as follows. The buyer starts the auction by setting the current bid at C_o. All those bidders who wish to remain in the auction at the current bid signal their willingness to do so. The buyer then lowers the current bid by some decrement, and all those bidders still willing to remain in the auction at this new current bid again signal their willingness to do so. We assume this decrement is arbitrarily small, ensuring that bidders can bid down to any value they wish to. For simplicity, we assume that re-entry is disallowed: once a bidder fails to signal willingness to remain in the auction, they cannot resume bidding at a lower price (the equilibrium remains if this assumption is relaxed). We let b_i denote bidder i's dropout bid, that is, bidder i remains in the auction up to bid b_i, but drops out when the buyer lowers the current bid below b_i. The buyer keeps lowering the current bid until at most one bidder remains in the auction. After the auction, the buyer discards bidders who did not participate in the auction (i.e., those who never signalled a willingness to remain in the auction), and post-qualifies the remaining bidders in order of ascending dropout bids (the b_i's), breaking ties randomly.

This setup resembles a traditional clock auction (e.g., Ausubel and Cramton 2006), with the crucial difference that bids can be interpreted as bids for positions in the post-qualification sequence. The dropout bids are used to compute the contract payment in the following manner. Suppose that bidder i is awarded the contract, that is, i is the first bidder that successfully passes post-qualification. Let i_j and t be as defined immediately following (2.10). The payment awarded to bidder i is precisely that given in (2.10), with one small change: b_{i_j} is used in lieu of $\psi(x_{i_j})+a(\beta_{i_j})$ within the bracketed term of the first summation.

Because this is an open-bid format, during the auction, in addition to the current bid, bidders are able to see competing bids remaining in the auction. In equilibrium, each bidder i will simply remain in the auction unless the current bid dips below his true adjusted virtual cost $\psi(x_i)+a(\beta_i)$, at which point he drops out. It is easy to check that in equilibrium bidders are sequenced and paid exactly as Proposition 1 specifies. In summary, we have

Proposition 2 *Under the open-auction format described above, the strategy of dropping out* *of the auction at one's true adjusted virtual cost constitutes a Bayesian Nash equilibrium under which the optimal mechanism from Proposition 1 is implemented.*

2.4.3 Optimal Qualification Screening Strategy

The buyer seeks the overall optimal mechanism that minimizes the expected total cost $\overline{\text{PRE}} + \overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}$. For a given set of pre-qualification thresholds, the previous subsection derived the buyer's optimal auction and post-qualification mechanism, minimizing $\overline{\text{PAY}}+\overline{\text{POST}}+\overline{\text{NT}}$. Hence, it remains only to proceed one layer up and investigate the optimal selection of the pre-qualification thresholds (equivalently β_i's). The buyer faces four practical concerns: uncertain supplier costs, uncertain supplier qualification levels, costs associated to verifying the qualification of a supplier, and an outside option cost. Comparisons of the latter two costs determine the buyer's optimal pre-qualification thresholds, for given uncertainties over supplier costs and qualification levels. We first identify the central tradeoff involved in the qualification strategy decision, followed by Proposition 3 which characterizes the buyer's optimal qualification screening strategy via three switching thresholds for qualification cost.

Central tradeoff of qualification screening decision. Under the optimal auction and post-qualification mechanisms described in §2.4.2, the more qualification due diligence the buyer performs before the auction, the more she pays before the auction but the less she pays after the auction. That is,

$$\overline{\text{PRE}} \text{ increases in } \beta_i, \text{ while } \overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}} \text{ decreases in } \beta_i, \text{ for all } i = 1, \ldots, n. \quad (2.11)$$

(This result is proved formally in §2.8.) Therefore, in an optimal qualification screening strategy the buyer trades off the pre-qualification cost against the total post-auction cost. Hereafter we call this the *central tradeoff* of the qualification decision.

Optimal qualification screening strategy thresholds. Although the central tradeoff reveals the buyer's key consideration when deciding the pre-qualification thresholds, solving for the optimal $\boldsymbol{\beta}^*$ which minimizes $\overline{\text{PRE}} + \overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}$ for fixed K and C_o is precluded by the complexity of the total cost expression (see, for example, equation (2.10)). Instead, we wish to describe the optimal qualification screening strategy more qualitatively by characterizing when various possible types of qualification strategies will be used. For any fixed outside option cost C_o, we seek to show that there exist three qualification cost thresholds defined as follows. Let K^{nt} be the smallest qualification cost such that, for any qualification cost $K \geq K^{nt}$, forgoing the auction in favor of the outside option is optimal. Let K^{post} be the smallest qualification cost such that, for any qualification cost K such that $K^{post} \leq K < K^{nt}$, delaying all qualification screening until after the auction (referred to as *post-only* for shorthand) is optimal, that is, $\beta_i = \underline{\beta}$ for all i. Let K^{pre} be the largest qualification cost such that, for any qualification cost $K \leq K^{pre}$, doing all qualification screening before the auction (referred to as *pre-only* for shorthand) is optimal, that is, $\beta_i = 1$ for all i.

Proposition 3 *Fix the outside option cost C_o. Under the buyer's optimal (total-cost minimizing) qualification screening strategy, K^{nt}, K^{post}, and K^{pre} all exist, are positive, and are finite. In other words,*
(i) pre-qualification only, that is $\beta_i = 1$ for all i, is optimal if $K \leq K^{pre}$,
(ii) a mix of pre- and post-qualification is optimal only if $K^{pre} < K < K^{post}$,
(iii) post-qualification only, that is $\beta_i = \underline{\beta}$ for all i, is optimal if $K^{post} \leq K < K^{nt}$, and
(iv) it is optimal to forgo the auction in favor of the outside option if and only if $K^{nt} \leq K$.

Moreover, the thresholds K^{nt}, K^{post}, and K^{pre} all increase in the outside option C_o.

Part (i) of Proposition 3 shows that there exists a threshold K^{pre}, separating the decision between complete versus partial pre-qualification. Keeping the outside option cost constant, if the qualification cost is below this threshold, the buyer prefers to completely pre-qualify suppliers before the auction; furthermore, any increase in the outside option cost while holding the qualification cost fixed will result in the buyer still preferring complete pre-qualification. Part (iii) of Proposition 3 shows that a similar threshold K^{post} exists between the partial pre-qualification versus the post-qualification only decision. Parts (i), (ii), and (iii) of the proposition together indicate that, keeping the outside option cost fixed, as qualification cost increases from zero the buyer's decision shifts from complete pre-qualification, to a mixture of pre and post-qualification, and then to the complete post-qualification decision. Finally, the buyer prefers to forgo the auction altogether if the cost of qualification is too high, as shown in Part (iv) of Proposition 3.

Proposition 3 indicates that despite the risk of non-transaction and larger contract payment, the buyer sometimes finds it profitable to postpone some or all qualification screening until after the auction. The decision of how much qualification to postpone depends on the qualification cost, as the buyer has more incentive to risk non-transaction and higher payments with qualification postponement if doing so avoids high qualification expenses before the auction.[2]

2.4.4 Symmetric Pre-Qualification Levels

While it can be in the buyer's interests to set unequal β_i's for different bidders when $K^{pre} < K < K^{post}$, doing so may raise fairness concerns about the procurement process. Under the optimal mechanisms described in §2.4.2, bidders who are more thoroughly pre-qualified are advantaged: first, they are more likely to win the contract, since they are more likely to be ranked at the top of the post-qualification sequence given that their virtual costs require less upward adjustment to capture post-qualification costs (i.e., $a(\beta_i)$ decreases in β_i); second, they will be paid more upon winning the contract, given that the payment to any winning bidder i per formula (2.10) increases in β_i. In a symmetric environment with no a priori differences among suppliers, the buyer may wish to pre-qualify suppliers equally, even though this might technically not be optimal, in order to avoid ill-will created by arbitrarily

[2] While our analyses assume that the outside option cost is finite, if C_o were instead allowed to be infinite (to capture a case in which the buyer absolutely must transact), it is straightforward to see that the buyer would find it optimal to fully qualify at least one bidder in the auction, in order to eliminate the possibility of non-transaction.

choosing one supplier to be more pre-qualified than the others.

Symmetric pre-qualification also simplifies the buyer's optimal auction and post-qualification mechanism. When bidders are equally pre-qualified, the optimal mechanisms of §2.4.2 use a common reserve price for all bidders. If β is the probability that any bidder in the auction is truly qualified, this optimal reserve price is $r = \max\{\min\{\psi^{-1}[C_o - a(\beta)], 1\}, 0\}$. Moreover, because a common qualification cost adjustment $a(\beta)$ is applied to all bidders, the optimal post-qualification sequencing rule reduces to simply ranking bidders according to ascending virtual costs, and in turn this is equivalent to simply ranking bidders according to true costs (by the virtual cost function increasing in true cost). In the following, we discuss a modified version of a standard sealed-bid first-price auction, where the contract winner is paid what he bids.

Sealed-bid first-qualified-price auction. When bidders are equally pre-qualified, the buyer can conduct a sealed-bid "first-qualified-price" auction to implement the optimal auction and post-qualification mechanism. In such an auction, the buyer sequences all bidders (who bid below the common reserve price r) according to ascending bid values, and awards the contract to the first qualified bidder (if any) with a payment equaling the bidder's bid.

Proposition 4 *Under symmetric pre-qualification, the sealed-bid first-qualified-price auction with reserve price* $r = \max\{\min\{\psi^{-1}[C_o - a(\beta)], 1\}, 0\}$ *implements the optimal mechanism from Proposition 1. In this auction, it is a Bayesian Nash equilibrium for a bidder with cost* $x_i \leq r$ *to bid* $x_i + \int_{x_i}^{r} [\frac{1-\beta F(z_i)}{1-\beta F(x_i)}]^{n-1} dz_i$.

It is intuitive that bidders react to the competitors' qualification probability β in equilibrium. This reaction is captured by the size of their price markup: for a fixed reserve price r (which is set by the buyer), the markup $\int_{x_i}^{r} [\frac{1-\beta F(z_i)}{1-\beta F(x_i)}]^{n-1} dz_i$ decreases in β.

Threshold policy for symmetric pre-qualification levels. With a restriction to symmetric pre-qualification levels, Proposition 3 still holds. We omit the proof of this; however, the intuition is that the central tradeoff (captured by equation (4.5)) remains with this restriction.

2.5. Numerical Illustrations

In this section we illustrate the optimal qualification screening strategy and show cost savings for a straightforward case in which the cost distribution is uniform and the qualification level distribution has decreasing density. This models a setting in which supplier cost types are evenly dispersed and the buyer has an increasing marginal cost to screen out each additional percent of unqualified suppliers. We normalize the cost distribution F to $U[\$500{,}000, \$1{,}000{,}000]$ to reflect dollar values, and take $H(q) = \sqrt{q}$ for tractability. We suppose that the buyer would like to hold an auction with three bidders, and that the buyer has a strict qualification requirement ($q_0 = 0.80$) such that only about eleven percent of all suppliers are truly qualified ($\underline{\beta} = 1 - H(q_0) = 11\%$, e.g., this could model the case in which foreign suppliers are extremely unlikely to meet rigorous qualification requirements set by the buyer). We say that an *asymmetric* strategy allows bidders' qualification probabilities to differ, while a *symmetric* strategy requires that bidders have equal qualification probabilities.

2.5.1 Optimal Qualification Probabilities

Asymmetric strategy. Figure 2.1(a) illustrates Proposition 3: three thresholds K^{pre}, K^{post}, and K^{nt} exist and are increasing in the outside option cost C_o. This figure depicts the optimal qualification screening strategy: pre-only is optimal in region A, a mix of pre- and post-qualification is optimal in region B, post-only is optimal in region C, and forgoing the auction is optimal in region D.

To further examine the optimal asymmetric qualification screening strategy, we consider a specific outside option cost of \$1,200,000, which is 20% higher than the worst possible supplier cost (e.g., this could model procurement with low cost foreign suppliers). Figure 2.1(b) illustrates the switching behavior of the optimal asymmetric qualification screening strategy for this outside option cost. When $\$17{,}500 < K < \$30{,}500$, the optimal qualification screening strategy entails partially pre-qualifying one bidder but fully pre-qualifying the

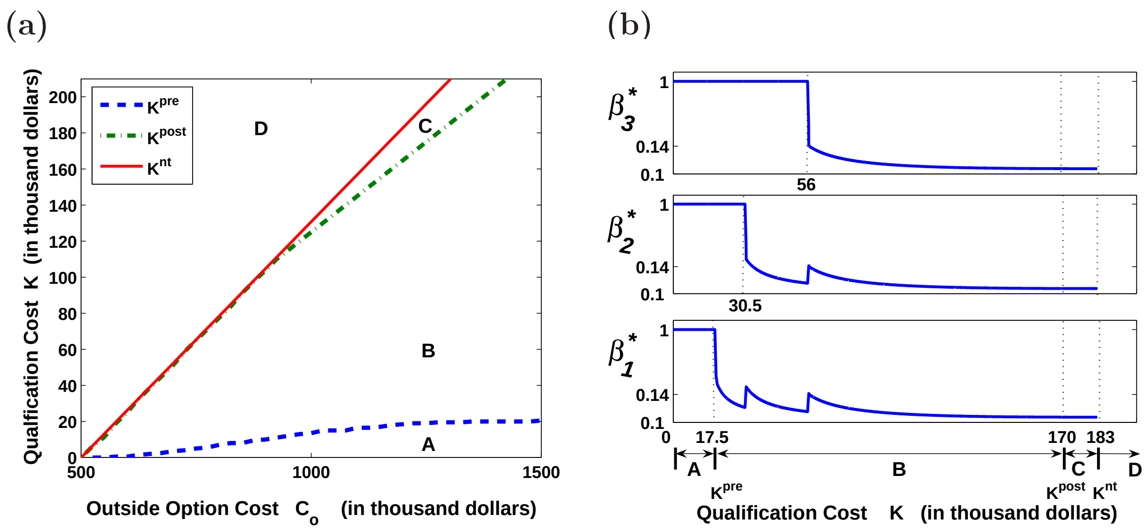

Figure 2.1: Illustration of regions where it is optimal to use pre-only (A); a mix of pre- and post-qualification (B); post-only (C); or forgo the auction (D). Plots assume $n = 3$ bidders, qualification threshold $q_0 = 0.80$, and supplier cost and qualification level distributions $F \sim U[\$500{,}000, \$1{,}000{,}000]$ and $H(q) = \sqrt{q}$, respectively. Panel (a) depicts regions in the (K, C_o) plane, and panel (b) shows optimal bidder qualification probabilities assuming the outside option cost is fixed at $C_o = \$1{,}200{,}000$. In panel (b), $K^{pre} = \$17{,}500$, $K^{post} = \$170{,}000$, $K^{nt} = \$183{,}000$; furthermore, the buyer prefers to fully pre-qualify three, two, or one bidder(s) if the qualification cost is between zero and \$17,500, between \$17,500 and \$30,500, or between \$30,500 and \$56,000, respectively.

other two. Such a strategy balances the pre-qualification cost with the total post-auction cost, given that it saves pre-qualification cost (by postponing some qualification due diligence for one bidder) while guaranteeing that the buyer finds a qualified bidder to contract with (by having two bidders fully pre-qualified). As the qualification cost increases, it is optimal to postpone more due diligence to save pre-qualification cost upfront. This explains why the pre-qualification level for the partially pre-qualified bidder, β_1^*, decreases as K increases from \$17,500 to \$30,500. However, when K exceeds \$30,500, having two bidders fully pre-qualified is too expensive to be optimal. Instead, it is optimal to fully pre-qualify only one bidder and partially pre-qualify the other two when $\$30{,}500 < K < \$56{,}000$, and it is optimal to partially pre-qualify all three bidders when $\$56{,}000 < K < \$170{,}000$.

Symmetric strategy. Figure 2.2(a) plots the optimal bidder qualification probability

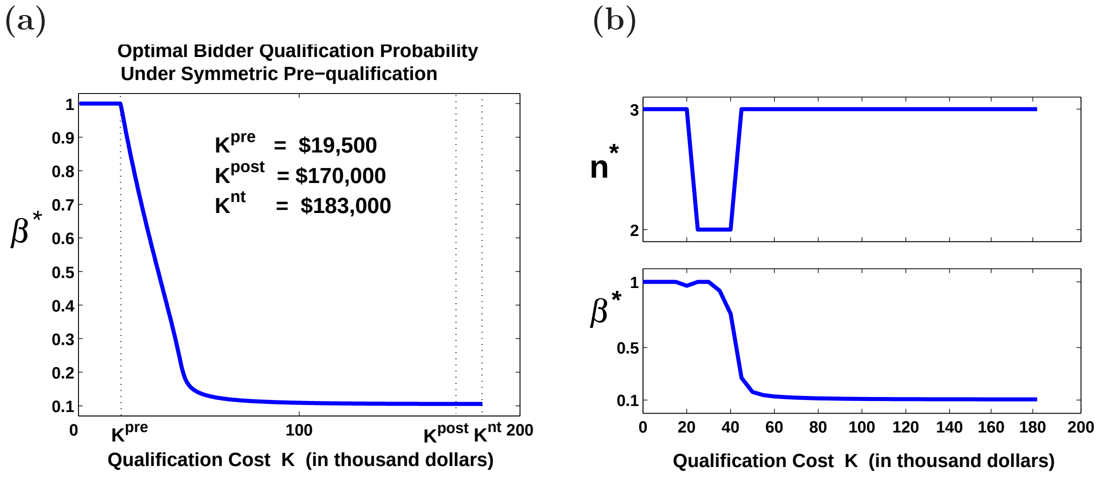

Figure 2.2: Panel (a) plots the optimal symmetric bidder qualification probability β^* for $n = 3$ bidders. Panel (b) plots the optimal number of bidders under symmetric qualification, $n^* \leq n = 3$, and the optimal symmetric qualification probability, β^*, for n^* bidders. All graphs are plotted against qualification cost, and assume outside option cost $C_o = \$1{,}200{,}000$, qualification requirement $q_0 = 0.8$, and supplier cost and qualification level distributions $F \sim U[\$500{,}000, \$1{,}000{,}000]$ and $H(q) = \sqrt{q}$, respectively.

under symmetric qualification (β^*). Threshold K^{pre} in Figure 2.2(a) ($19,500) is larger than that in Figure 2.1(b) ($17,500) and the other two thresholds K^{post} and K^{nt} equal those in Figure 2.1(b), since the optimal qualification strategy without restriction is not symmetric for $\$17{,}500 < K < \$19{,}500$ while the optimal qualification strategy without restriction is symmetric for K near K^{post} ($170,000) and K^{nt} ($183,000). Figure 2.2(a) assumes the buyer invites $n = 3$ bidders to the auction. With asymmetric pre-qualification thresholds, the buyer always finds it optimal to invite up to n bidders; however, it turns out that a buyer who is restricted to use symmetric pre-qualification might choose to invite $n^* < n$ bidders to her auction. Figure 2.2(b) describes such a case: when $\$20{,}000 < K < \$44{,}000$, the optimal number of bidders under symmetric pre-qualification, $n^* = 2$, is less than $n = 3$. In other words, fully pre-qualifying two bidders is desirable, but fully pre-qualifying a third bidder does not sufficiently reduce expected total costs to warrant the additional pre-qualification costs spent on the third bidder. Without the symmetric pre-qualification restriction, the buyer could simply use post-qualification on the third bidder; however, the buyer forgoes the third bidder if this bidder cannot be invited without being pre-qualified to the same level as the other two bidders, as is the case when restricting to symmetric pre-qualification. In summary, inviting fewer bidders is a realistic way in which a buyer might cope with a symmetric pre-qualification strategy restriction.

2.5.2 Cost Comparisons

Optimal asymmetric strategy versus pre-only. Fully pre-qualifying all bidders before the auction (pre-only) is generally standard in practice. Let

$$\begin{array}{c}\text{pre-only strategy}\\ \text{optimality gap}\end{array} = 1 - \frac{\text{total cost with optimal asymmetric qualification screening}}{\text{total cost with pre-only}}. \tag{2.12}$$

In Figure 2.3(a), in addition to the previous scenario where the buyer has a strict qualification requirement ($q_0 = 0.80$), we also examine the optimality gap for a scenario where the buyer has a more lenient qualification requirement ($q_0 = 0.20$). The maximal optimality gap is around 16% when $q_0 = 0.80$, and 30% when $q_0 = 0.20$. For both qualification requirement scenarios, the optimality gap peaks when qualification screening becomes too costly to make the auction with pre-only worthwhile. To the left of this peak, the rate of cost savings is small when K is small because the pre-only strategy is (or is close to) optimal. To the right of this peak, the rate diminishes as K increases to K^{nt} because the optimal qualification strategy's total cost increases to C_o while the pre-only strategy's cost is fixed at C_o for K exceeding \$40,500 (when $q_0 = 0.8$), or \$120,000 (when $q_0 = 0.2$).

When the qualification cost is not too expensive (\$17,500 $< K <$ \$40,500), the optimality gap is greater under the stricter qualification requirement. For example, if supplier qualification costs \$40,000 (perhaps \$20,000 is spent to purchase and test supplier products, \$15,000 to send buyer employees to inspect supplier facilities abroad, and \$5,000 on time-intensive meetings with stakeholders throughout the buyer's company), about 15.4% of total procurement costs are saved by postponing qualification checks when $q_0 = 0.80$, compared

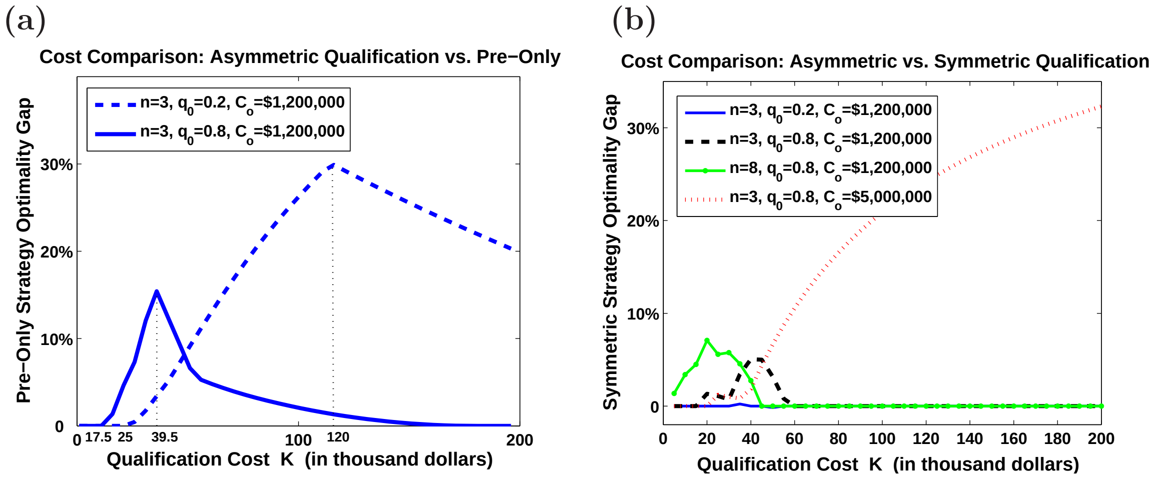

Figure 2.3: Panel (a) compares the optimal asymmetric and pre-only qualification strategies' costs (per equation (2.12)). Panel (b) compares the optimal asymmetric and optimal symmetric qualification strategies' costs (per equation (2.13)). All graphs are plotted against qualification cost, and assume supplier cost and qualification level distributions $F \sim U[\$500{,}000, \$1{,}000{,}000]$ and $H(q) = \sqrt{q}$, respectively.

to about 3.4% when $q_0 = 0.20$. When $q_0 = 0.8$, only about eleven percent of all suppliers are truly qualified, making the full pre-qualification policy very costly. Consequently, optimally postponing a portion of qualification can be very beneficial.

A more lenient qualification requirement of $q_0 = 0.20$ makes pre-qualification more affordable and hence mitigates the relative benefits of aggressive post-qualification. However, under the more lenient qualification requirement, the affordability of pre-only allows the relative cost savings of the optimal qualification strategy to be sustained for larger qualification costs K, eventually reaching 30%. The relative cost savings are more quickly capped in the stricter requirement case, because pre-qualification becomes prohibitively expensive for the buyer, who abandons the auction and resorts to the outside option at a smaller K.

Symmetric versus asymmetric strategies. Given that a symmetric pre-qualification strategy is "fairer" and "simpler", a buyer might prefer it if it can well approximate the optimal asymmetric strategy. Let

$$\begin{array}{c}\text{symmetric strategy}\\ \text{optimality gap}\end{array} = 1 - \frac{\text{total cost with optimal asymmetric strategy}}{\text{total cost with optimal symmetric strategy}}. \quad (2.13)$$

From Figure 2.3(b), we see that the optimal symmetric strategy approximates the optimal asymmetric strategy well when the outside option cost C_o, the number of bidders n, and the qualification requirement q_0 are small.[3] In the range of qualification cost K considered, this rate is no more than 0.2% when C_o = \$1,200,000, $n = 3$, and $q_0 = 0.2$. The maximal rate observed is about 5% $\sim$ 7% when q_0 increases to 0.8. However, the symmetric strategy approximates the optimal asymmetric strategy poorly when the outside option cost is large and qualification is expensive but not prohibitively so (making an auction still worthwhile). Intuitively, a buyer will always find it prudent to heavily use pre-qualification to avoid the downside risk of an onerous outside option. However, a buyer who can employ asymmetric pre-qualification can also enjoy the upside benefit of potentially reduced contracting costs from speculatively inviting "long-shots" to the auction, bidders admitted with much or all of their qualification screening postponed. For example, the optimality gap can reach about 32% when C_o increases to \$5,000,000 in Figure 2.3(b). Therefore, in a scenario where the buyer cannot produce in-house and the contract is extremely important — for example, flu vaccine procurement by a government, the buyer could be much better off with the flexibility to apply asymmetric pre-qualification thresholds to bidders, to simultaneously avoid the downside risk of non-transaction while still enjoying speculative benefits from casting a wider net for bids.

In summary, regardless of whether the buyer uses asymmetric or symmetric pre-qualification, these results suggest that the buyer should seriously consider postponing some or all of the supplier qualification process, especially when the outside option cost is high, and either the qualification requirement is strict and the qualification cost is moderate, or the qualification

[3]Note that in calculating the optimality gap in Figure 2.3(b) we have included the symmetric strategy's flexibility to optimally invite fewer than n bidders.

requirement is lenient but qualification is expensive.

2.6. Practical Considerations and Extensions

2.6.1 A More General Supplier Cost Model

Our analyses up to this point have assumed that suppliers can privately and perfectly observe their own costs prior to the auction and that these costs are statistically independent. However, one can easily imagine a situation in which suppliers' costs are related due to similar or shared cost drivers. Moreover, suppliers might form a rough estimate of their cost on their own, and utilize information from competing suppliers' bids during the auction to update this estimate. In this subsection we discuss how §4.4's results can be adapted to a more general cost model that allows for these possibilities, culminating in Proposition 5.

Suppose that each supplier i possesses private information about his idiosyncratic status, for example, production technology, inventory level, order backlog, etc. which is denoted by $\theta_i \in [0,1]$. Supplier i's cost to fulfill the contract is then modelled by a function $x_i(\boldsymbol{\theta},\ \xi) \triangleq X(\theta_i, c(\boldsymbol{\theta}), \xi)$, where $\boldsymbol{\theta} \triangleq (\theta_1, \ldots, \theta_n)$ represents the collection of all bidders' private information, $c(\boldsymbol{\theta})$ characterizes common cost factors that depend on $\boldsymbol{\theta}$, and ξ characterizes exogenous and publicly observable factors that could impact costs directly (via X) and/or indirectly through affecting the distribution of bidders' private information (the θ_i's). We assume that x_i is strictly increasing in θ_i.

As an illustration, consider the following simple model in which costs are correlated and suppliers adjust their cost estimates after seeing competitors' bids. Let $X(\theta_i, c(\boldsymbol{\theta}), \xi) = \theta_i + \sum_j e(\theta_j) + \xi$, where $e : [0,1] \to \mathbb{R}$ is an increasing function. This model is essentially that which appears in Myerson (1981) without the term ξ. The summation $\sum_j e(\theta_j)$ could represent cost estimate revisions based on anticipated demand in the supply base for specialized components, with suppliers' private information (the θ_i's) incorporating their own anticipated component demand based on their on-hand component inventory and current order backlog. In a similar vein, ξ could reflect prevailing market prices for commodity inputs whose prices are unaffected by the suppliers' aggregate usage, being only a negligible fraction of the commodities' overall demand. Although this illustrative example includes the exogenous factor ξ in a simple additive fashion, the general cost model can handle richer structures; for example, suppliers' costs can depend on ξ in a non-additive way in order to model heterogeneous efficiencies in using commodity inputs. We now return to describing the general model.

So that suppliers are ex ante symmetric, we assume that $c(\boldsymbol{\theta})$ is exchangeable in all elements of $\boldsymbol{\theta}$, and θ_i has a commonly known conditional distribution $F(\cdot|\xi)$ with a positive and continuous density $f(\cdot|\xi)$. Furthermore, we assume that θ_i's are conditionally independent given any fixed publicly observable ξ. Hereafter, for notational simplicity, we will suppress the variable ξ when writing F, f, X and x_i's. By treating ξ as implicit, the general supplier cost formulation introduced two paragraphs above reduces to the formulation of correlated bidder costs in Branco (1997).

With this cost model, our analysis of the optimal auction and post-qualification mechanism needs one key modification, which is to redefine the virtual cost as a function of the collection of all bidders' private information and the exogenous factor, that is,

$$\hat{\psi}_i(\boldsymbol{\theta}) \triangleq x_i(\boldsymbol{\theta}) + \frac{F(\theta_i)}{f(\theta_i)} \frac{\partial x_i(\boldsymbol{\theta})}{\partial \theta_i}.$$

In §2.8, we derive the optimal mechanism assuming that (i) $\frac{\partial \hat{\psi}_i(\boldsymbol{\theta})}{\partial \theta_i} > 0$ for all i at all $\boldsymbol{\theta}$; and (ii) $\frac{\partial \hat{\psi}_i(\boldsymbol{\theta})}{\partial \theta_i} > \frac{\partial \hat{\psi}_j(\boldsymbol{\theta})}{\partial \theta_i}$ for all i and $j \neq i$ at all $\boldsymbol{\theta}$. While (i) ensures that bidder i's virtual cost $\hat{\psi}_i(\boldsymbol{\theta})$ strictly increases in his type, (ii) ensures that, although increases in i's type can also increase the virtual costs of $j \neq i$, such increases are smaller than the increase caused to i's own virtual cost. Condition (i) is satisfied, for example, if $x_i(\boldsymbol{\theta})$ is increasing and convex in θ_i, and $\frac{F(\theta_i)}{f(\theta_i)}$ is increasing in θ_i. Condition (ii) is then satisfied, for example, if $\frac{\partial x_i(\boldsymbol{\theta})}{\partial \theta_i} > \frac{\partial x_j(\boldsymbol{\theta})}{\partial \theta_i}$ and $\frac{\partial^2 x_i(\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \leq 0$. The insight is that under conditions (i) and (ii) the redefined virtual cost $\hat{\psi}_i(\boldsymbol{\theta})$ under the general cost model inherits a key property of the virtual cost $\psi(x_i)$ under the independent private cost model, that is, any increase (decrease) of bidder i's report to the buyer encourages the buyer to demote (promote) bidder i in the post-qualification sequence.

Under conditions (i) and (ii) on the general cost model, it remains optimal to sequence bidders for post-qualification according to ascending value of their adjusted virtual costs $\hat{\psi}_i(\boldsymbol{\theta}) + a(\beta_i)$, as we show in Proposition 5 below. As before, the contract winner is the first bidder (if any) deemed qualified by post-qualification. Furthermore, echoing equation (2.10), incentive compatibility can again be ensured by paying the contract winner, say bidder i, according to the amounts i could have increased his bid by and still been preferred to the t bidders sequenced after him for post-qualification, and the reserve price. Yet there is one slight difference when actually computing the payment under the general cost model: Because an increase in i's bid can affect other bidders' virtual costs, when computing the payment for bidder i, the buyer must take into account how the post-qualification sequence of the t bidders would have been reshuffled if i had increased his bid.

We now define notation needed to generalize equation (2.10). Let $\boldsymbol{\theta}_{-i}$ be the vector $\boldsymbol{\theta}$ excluding the i^{th} element, and define $r_i = \max\{\min\{z, 1\}, 0\}$ where z is such that $\hat{\psi}_i(z, \boldsymbol{\theta}_{-i}) + a(\beta_i) = C_o$. Let z_{i_j} be such that $\hat{\psi}_i(z_{i_j}, \boldsymbol{\theta}_{-i}) + a(\beta_i) = \hat{\psi}_{i_j}(z_{i_j}, \boldsymbol{\theta}_{-i}) + a(\beta_{i_j})$ and $z_{i_1} \leq z_{i_2} \leq \cdots \leq z_{i_t}$, where i_j $(j = 1, ..., t)$ are the indices of the t bidders such that $\hat{\psi}_i(\boldsymbol{\theta}) + a(\beta_i) \leq \hat{\psi}_{i_j}(\boldsymbol{\theta}) + a(\beta_{i_j})$ and $z_{i_j} \in [\theta_i, r_i]$. Note that z_{i_j} and r_i, representing the maximal amount that bidder i could have bid and still been preferred to bidder i_j and the outside option, respectively, are well defined due to conditions (i) and (ii).

Proposition 5 *Under the general cost model described above, the optimal auction and post-qualification mechanism proceeds as follows: only post-qualify a bidder i if $\hat{\psi}_i(\boldsymbol{\theta}) + a(\beta_i) \leq C_o$; sequence all such bidders for post-qualification according to ascending value of $\hat{\psi}_i(\boldsymbol{\theta}) + a(\beta_i)$, breaking ties randomly; award the contract to the first bidder i (if any) deemed qualified by post-qualification with a payment of*

$$\sum_{j=1}^{t-1} x_i(z_{i_j}, \boldsymbol{\theta}_{-i}) \beta_{i_j} \prod_{k=1}^{j-1} (1 - \beta_{i_k}) + x_i(r_i, \boldsymbol{\theta}_{-i}) \prod_{k=1}^{t} (1 - \beta_{i_k}) \ \text{(if } t \geq 1\text{) or } x_i(r_i, \boldsymbol{\theta}_{-i}) \ \text{(if } t = 0\text{)},$$

and pay nothing to all other bidders. Additionally, Proposition 3 characterizing the optimal *qualification strategy continues to hold.*

The key reason that Proposition 3 continues to hold under the general cost model is that the central tradeoff described in equation (4.5) does not rely on the format of the bidder's cost function x_i. In summary, Proposition 5 implies that the main results of §4.4 — the optimal mechanism (Proposition 1) and qualification thresholds (Proposition 3) — continue to hold under the general cost model described above, albeit with a bit more care needed when calculating the contract winner's payment in order to account for correlation among bidders' costs.

2.6.2 Value of Credible Reserve Price

Our optimal mechanism derivation in §2.4.2 assumed that the buyer could credibly commit to not awarding the contract to any bidder bidding above the reserve price set for him. As Milgrom (1987) points out, an auctioneer who cannot credibly commit to throwing away bids between the reserve price and the auctioneer's own valuation is disadvantaged: the auctioneer cannot achieve the optimal ex-ante expected profits because bidders will ignore the announced reserve price.

Post-qualification of bidders places additional importance on reserve price credibility. Proposition 1 characterizes the optimal reserve price for bidder i as $r_i = \max\{\min\{\psi^{-1}(b_i^{\max}), 1\}, 0\}$, where $b_i^{\max} \triangleq C_o - a(\beta_i)$ is the maximum bid from bidder i that the buyer would find profitable to post-qualify. Since $\psi(x) > x$, we can have $r_i < b_i^{\max}$; i.e., in the optimal mechanism the buyer promises to ignore bidder i's bid if it is in $(r_i, b_i^{\max}]$, even though she actually would find it profitable to post-qualify such a bid after the auction.

If this promise is uncredible, the buyer is forced to set $b_i^{\max}$ as the reserve price for bidder i and cannot apply an optimal mechanism (optimal reserve prices). This is particularly problematic when $b_i^{\max}$ is greater than 1, the upper bound of supplier cost. If all other bidders fail post-qualification and bidder i is the last remaining bidder, bidder i can make a take-it-or-leave-it offer of $b_i^{\max}$ to the buyer. For example, suppose that all bidders have a symmetric qualification probability equal to β. Per Proposition 4, a fully credible buyer would find it optimal to run a sealed-bid first-qualified-price auction with an optimal reserve price r (recall $1 \geq r$). If the buyer's lack of credibility forces her to instead use $b^{\max} = C_o - a(\beta) > 1$ as a reserve price, a cost type 1 bidder i expects to earn $(b^{\max} - 1)\beta(1-\beta)^{n-1}$ from the auction. In contrast, note that when $\beta = 1$ for all bidders, as is assumed in classical auction theory, cost type 1 bidders expect to earn *zero* profits in the auction provided the auction has at least two bidders.

Thus, two factors inflate the costs of an uncredible buyer in our setting: forgone "price discrimination" opportunities, as seen in classical auction theory; and the "being held hostage by the last remaining bidder" effect, which allows even a worst-type bidder to expect positive profits from the auction and is a consequence of post-qualification which to our knowledge is new to the literature. Continuing the symmetric β's example above, when $C_o - a(\beta) > 1 = r$ only the "being held hostage" effect exists; when $1 \geq C_o - a(\beta) > r$, only the "price discrimination" effect exists; and when $C_o - a(\beta) > 1 > r$ both effects exist.

These observations suggest that post-qualification should be used carefully when the buyer has little or no negotiating clout with suppliers and must rely solely on competition among suppliers for price concessions. In such situations a supplier can command a very high price in a one-on-one negotiation with the buyer, a damaging scenario for the buyer that is risked by post-qualification. Intuitively, this can push the buyer to employ more stringent pre-qualification screening. On the other hand, the buyer is between a rock and a hard place if qualification screening is prohibitively expensive, in which case she might find it too costly to pre-qualify extensively enough to avoid getting squeezed by the suppliers. For a simple illustration of both of these possibilities, consider the setup of Figure 2.2. For qualification cost \$30,000, the optimal symmetric qualification probability for a credible buyer is 0.66, which is lower than 0.93, the symmetric qualification probability that can be (using numerical analysis) shown to be optimal for an uncredible buyer (who uses a sealed-bid first-qualified-price auction with reserve price $C_o - a(\beta)$). However, at a higher qualification cost of \$45,000 the opposite happens: the credible buyer's optimal symmetric qualification probability of 0.27 exceeds 0.11, the optimal probability for the uncredible buyer. While we will leave a more complete exploration of credibility and qualification screening interactions to future work, the above suggests that despite the importance of credibility in avoiding cost inflation under post-qualification, even an uncredible buyer might prefer heavy use of post-qualification if qualification is very expensive.

2.7. Conclusions

When issuing an RFQ for competitive bid, finding a supplier truly qualified to fulfill the contract is often as important as price concerns. Costly supplier qualification processes are virtually ubiquitous in industry to help buyers proactively avoid problems and expenses associated with supplier non-performance, e.g., recalls and product liability issues. This chapter explicitly models and suggests optimal policies for both the supplier qualification and competitive price negotiation processes together, and to our knowledge is the first study of optimal supplier qualification processes in the operations management and auction theoretic literatures. To save on total supplier qualification and contracting costs we allow the buyer to delay all or part of the qualification process until after the competitive price negotiation (an auction) and then home in on the lowest (virtual adjusted cost) bidders. While delaying qualification is not to our knowledge common practice in industry, our study provides a mathematical framework which suggest such a post-qualification stage can indeed be beneficial.

In particular, we find that the pre-qualification only strategy is optimal solely when supplier qualification is relatively cheap. Because postponing qualification means that some attractive bids in the auction may be disqualified, it makes sense to completely pre-qualify suppliers if doing so is cheap. However, for moderate sized qualification costs the buyer can do much better if some costly qualification is delayed until after the auction, because reduced qualification costs with judicious post-qualification can more than offset expected increases in the contracting costs (as determined by our auction theoretic analysis). Figure 2.3(a) shows total (qualification plus procurement) cost savings of about $3\% \sim 15\%$ for a contract worth \$1.2M to the buyer when qualifying a supplier costs \$40,000, supplier cost types are evenly dispersed, and the qualification level distribution has a decreasing density. More generally, Proposition 3 partitions the two-dimensional qualification cost (K) and outside option cost (C_o) plane into regions where either traditional pre-qualification only, our novel post-qualification only, or our novel mix of the two are optimal (illustrated in Figure 2.1).

While operations management analyses such as ours may merely galvanize a reconsideration of current procurement policies, as supply chains lengthen and supply sources become globalized and more varied, the increase in potential new suppliers and the growing number of RFQ events could make the standard pre-qualification only strategy prohibitive for resource-constrained procurement departments that cannot possibly fully pre-qualify all suppliers invited to all bidding events. Post-qualification might eventually be used out of sheer necessity to accommodate constrained qualification resources, but fortunately our study shows that post-qualification can be part of an optimally balanced supplier qualification strategy even without such resource constraints.

Our study is built on classical auction theory, and extends the auction theory literature by developing an optimal auction and post-qualification mechanism. We characterize the optimal mechanism in Proposition 1 and propose both sealed-bid and open-bid formats to implement it, including Proposition 4 showing that a simple variant of a standard sealed-bid auction — a sealed-bid, first-qualified-price auction — implements the optimal mechanism when all bidders are pre-qualified up to the same level.

The spirit of our results are broad in the sense that they characterize the tradeoff between costly pre-qualification and increased likelihood that an attractive bid will be tenable (the bidder is truly qualified for the contract). For general auction formats not studied in the present chapter, we expect that our main insight that post-qualification can be an effective cost-reduction strategy for the buyer will continue to hold. However, for auction mechanisms with extremely rich strategy spaces, post-qualification can add additional complexities to the bidding equilibrium analysis. For instance, were post-qualification used in conjunction with a standard reverse English auction, bidders may find it optimal to drop out of the auction before reaching their true costs in a way that intricately depends on all previous bids in the auction up to that point. Introducing post-qualification in various auction formats presents many opportunities for research on competitive bidding.

Section 2.6 discusses several possible extensions to our model, which while fairly general does make some important assumptions in order to keep the analyses focused and tractable. Other extensions are also possible. The present chapter shows that the buyer can profitably postpone qualification and employ asymmetric pre-qualification policies, even under our assumption that suppliers are ex ante symmetric. The ex ante asymmetric suppliers case, in which H, K, and F are allowed to be supplier-specific, is an interesting possible direction for future work; while the analysis would be more complex, we suspect that the main insights of this chapter would be preserved, as ex ante supplier asymmetry would likely make our asymmetric pre-qualification screening policies even more fitting.

In this chapter we also assume that suppliers are privately informed about their costs, and the buyer and each supplier are equally unsure of the supplier's likelihood of meeting the buyer's qualification requirements. One interesting and challenging direction for future work is the multidimensional asymmetric information case in which superior prior knowledge of the qualification level is held on the supplier side. The ensuing optimal procurement mechanism analysis would involve information asymmetry that is multidimensional, a setting which is generally seen (e.g., Rochet and Stole 2003) as a current challenge in mechanism design research.

Our analysis assumes an infinite supplier pool, which ensures that the buyer is able to fill the auction. However, in practice the number of suppliers who show interest in the RFQ is finite and could be small for specialized purchases. With a finite supplier pool, the number of suppliers passing pre-qualification and entering the auction could be non-deterministic. Presumably the buyer could dynamically adjust the pre-qualification threshold per the remaining supplier pool size. We leave such extensions to our future work.

2.8. Proofs

2.8.1 Proof of Proposition 1

An optimal Q should minimize (2.8a), which we rewrite as $\int_{\boldsymbol{x}} S(Q) d\mathbf{F}_{-i}(\boldsymbol{x})$, where

$$S(\mathrm{Q}) = C_o + \sum_{i=1}^{n} \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]. \tag{2.14}$$

We first show that an optimal sequencing rule $\mathrm{Q}^*(\boldsymbol{x})$ should not include any i such that $\psi(x_i) + a(\beta_i) > C_o$, and then show that it should include all i such that $\psi(x_i) + a(\beta_i) \le C_o$.

Let $\mathrm{Q}^*(\boldsymbol{x}) = (i_1, \dots, i_L)$, where $L \le n$ is the number of elements in the sequence. Suppose there exists $i_k \in \{i_1, \dots, i_L\}$ such that $\psi(x_{i_k}) + a(\beta_{i_k}) > C_o$, and without loss of generality suppose i_k is the largest such element, so that $j > k$ implies $\psi(x_{i_j}) + a(\beta_{i_j}) \le C_o$. Consider $\hat{\mathrm{Q}}(\boldsymbol{x}) = (i_1, \dots, i_{k-1}, i_{k+1}, \dots, i_L)$, which is obtained by removing i_k from $\mathrm{Q}^*(\boldsymbol{x})$ without changing the sequence of other elements. Using equation (2.2), calculate $\Delta_i^*(\boldsymbol{x})$ and $\hat{\Delta}_i(\boldsymbol{x})$, the winning probabilities derived from the squences Q^* and $\hat{\mathrm{Q}}$. It turns out that $\Delta_i^*(\boldsymbol{x}) = \hat{\Delta}_i(\boldsymbol{x})$ for $i = i_1, \dots, i_{k-1}$ while $\Delta_i^*(\boldsymbol{x}) < \hat{\Delta}_i(\boldsymbol{x}) = \frac{\Delta_i^*(\boldsymbol{x})}{1-\beta_{i_k}}$ for $i = i_{k+1}, \dots, i_L$. This implies $S(\hat{\mathrm{Q}}) - S(\mathrm{Q}^*) < 0$ (as can be easily checked), which contradicts the assumption that Q^* minimizes (2.14). This contradiction indicates that $\mathrm{Q}^*(\boldsymbol{x})$ should not include any i such that $\psi(x_i) + a(\beta_i) > C_o$.

If $\mathrm{Q}^*(\boldsymbol{x}) = (i_1, \dots, i_L)$ does not include an element $i_k \in \{1, \dots, n\} \backslash \{i_1, \dots, i_L\}$ such that $\psi(x_{i_k}) + a(\beta_{i_k}) \le C_o$, we consider $\hat{\mathrm{Q}}(\boldsymbol{x}) = (i_1, \dots, i_L, i_k)$, which is obtained by appending i_k to the end of $\mathrm{Q}^*(\boldsymbol{x})$ without changing the sequence of the other elements. Calculating $\Delta_i^*(\boldsymbol{x})$ and $\hat{\Delta}_i(\boldsymbol{x})$ via equation (2.2) reveals that $\Delta_i^*(\boldsymbol{x}) = \hat{\Delta}_i(\boldsymbol{x})$ for $i = i_1, \dots, i_L$ while $\hat{\Delta}_{i_k}(\boldsymbol{x}) \ge 0$. This implies $S(\hat{\mathrm{Q}}) - S(\mathrm{Q}^*) = \hat{\Delta}_{i_k}(\boldsymbol{x})[\psi(x_{i_k}) + a(\beta_{i_k}) - C_o] \le 0$. Thus, $\hat{Q}$ performs as good or better than Q^*, and we conclude that $\mathrm{Q}^*(\boldsymbol{x})$ should include all i such that $\psi(x_i) + a(\beta_i) \le C_o$.

We next show that any optimal $\mathrm{Q}^*(\boldsymbol{x})$ should sequence the elements in ascending value of $\psi(x_i) + a(\beta_i)$. Otherwise, if $\mathrm{Q}^*(\boldsymbol{x}) = (i_1, \dots, i_L)$ includes an element $i_k \in \{i_1, \dots, i_L\}$ such that $\psi(x_{i_k}) + a(\beta_{i_k}) > \psi(x_{i_{k+1}}) + a(\beta_{i_{k+1}})$, we consider $\hat{\mathrm{Q}}(\boldsymbol{x}) = (i_1, \dots, i_{k+1}, i_k \dots, i_L)$, which is obtained by switching i_k and i_{k+1} without changing the sequence of the other elements. Calculating $\Delta_i^*(\boldsymbol{x})$ and $\hat{\Delta}_i(\boldsymbol{x})$ via equation (2.2) reveals that $\Delta_i^*(\boldsymbol{x}) = \hat{\Delta}_i(\boldsymbol{x})$ for $i = i_1, \ldots, i_{k-1}, i_{k+2} \ldots, i_L$ while $\Delta_{i_k}^*(\boldsymbol{x}) = \beta_{i_k} \prod_{j=1}^{k-1}(1-\beta_{i_j})$, $\Delta_{i_{k+1}}^*(\boldsymbol{x}) = \beta_{i_{k+1}}(1-\beta_{i_k}) \prod_{j=1}^{k-1}(1-\beta_{i_j})$, $\hat{\Delta}_{i_k}(\boldsymbol{x}) = \beta_{i_k}(1-\beta_{i_{k+1}}) \prod_{j=1}^{k-1}(1-\beta_{i_j})$, and $\hat{\Delta}_{i_{k+1}}(\boldsymbol{x}) = \beta_{i_{k+1}} \prod_{j=1}^{k-1}(1-\beta_{i_j})$. This implies $S(\hat{\mathrm{Q}}) - S(\mathrm{Q}^*) = \beta_{i_{k+1}}\beta_{i_k} \prod_{j=1}^{k-1}(1-\beta_{i_j})\{[\psi(x_{i_{k+1}}) + a(\beta_{i_{k+1}}) - C_o] - [\psi(x_{i_k}) + a(\beta_{i_k}) - C_o]\} < 0$, which contradicts the assumption that Q^* minimizes (2.14). This contradiction indicates that $\mathrm{Q}^*(\boldsymbol{x})$ should sequence the elements in ascending value of $\psi(x_i) + a(\beta_i)$.

Given that $\mathrm{Q}^*(\boldsymbol{x})$ sequences the elements in ascending value of $\psi(x_i) + a(\beta_i)$ and that $\psi(x_i)$ increases in x_i, $\Delta_i(\boldsymbol{x})$ is decreasing in x_i and hence $\delta_i(x_i)$ is also decreasing in x_i. To find a payment rule $\mathrm{M}^*(\boldsymbol{x})$ satisfying (2.8b) with the proposed $\mathrm{Q}^*(\boldsymbol{x})$ so that all bidders bid their true costs in equilibrium, we consider

$$\mathrm{M}_i^*(\boldsymbol{x}) = \Delta_i^*(\boldsymbol{x})x_i + \int_{x_i}^1 \Delta_i^*(z_i, \boldsymbol{x}_{-i})dz_i. \tag{2.15}$$

This payment rule satisfies (2.8b), as (2.8b) can be obtained by integrating both sides of (2.15) over $\boldsymbol{x}_{-i}$ with distribution $\mathbf{F}_{-i}$. Given (2.15), if $i \notin \mathrm{Q}^*(\boldsymbol{x})$, we have $\Delta_i^*(\boldsymbol{x}) = 0$ and hence $\mathrm{M}_i^*(\boldsymbol{x}) = 0$; if $i \in \mathrm{Q}^*(\boldsymbol{x})$, we have $\Delta_i^*(\boldsymbol{x}) > 0$ and hence $\mathrm{M}_i^*(\boldsymbol{x}) = \Delta_i^*(\boldsymbol{x})[x_i + \int_{x_i}^1 \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i]$.

Finally, we justify our restriction to Q's that deterministically (as opposed to randomly) map $\boldsymbol{x}$ to a post-qualification sequence. Consider a $\hat{\mathrm{Q}}$ that randomizes over post-qualification sequences. Note that all the arguments in §2.4.2 to derive program (2.8) continue to hold if we replace Q by $\hat{\mathrm{Q}}$ and change equation (2.2) to

$$\Delta_i(\boldsymbol{x}) = E_{\hat{\mathrm{Q}}}\left[\begin{cases} \beta_i \prod_{j \text{ ranks ahead of } i \in \hat{\mathrm{Q}}(\boldsymbol{x})} (1-\beta_j), & \text{if } i \in \hat{\mathrm{Q}}(\boldsymbol{x}) \\ 0, & \text{if } i \notin \hat{\mathrm{Q}}(\boldsymbol{x}) \end{cases}\right].$$

(The derivations of $\overline{\mathrm{PAY}}$ go through because the probability of winning with a given report enters bidder i's payoff function linearly.) Finally, because under $\hat{\mathrm{Q}}$ the buyer's expected payoff function (2.8a) is a weighted average of payoffs under the post-qualification sequences that are possible under $\hat{\mathrm{Q}}$, she can do no better than simply using with probability 1 the best post-qualification sequence Q^* as defined above.

We now derive equation (2.7). Let i_j $(j = 1, ..., t)$ be the index of the j^{th} bidder sequenced after bidder i for post-qualification according to the optimal sequencing rule such that $\psi(x_{i_j}) + a(\beta_{i_j}) \leq \psi(r_i) + a(\beta_i)$, where t is the number of such bidders, and $r_i = \max\{\min\{\psi^{-1}[C_o - a(\beta_i)], 1\}, 0\}$ denotes the reserve price for bidder i. Note that

$$\begin{aligned}
&x_i + \int_{x_i}^{1} \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i \\
&= x_i + \int_{x_i}^{\psi^{-1}[\psi(x_{i_1})+a(\beta_{i_1})-a(\beta_i)]} \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i \\
&\quad + \sum_{j=1}^{t-1} \int_{\psi^{-1}[\psi(x_{i_j})+a(\beta_{i_j})-a(\beta_i)]}^{\psi^{-1}[\psi(x_{i_{j+1}})+a(\beta_{i_{j+1}})-a(\beta_i)]} \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i \\
&\quad + \int_{\psi^{-1}[\psi(x_{i_t})+a(\beta_{i_t})-a(\beta_i)]}^{r_i} \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i + \int_{r_i}^{1} \frac{\Delta_i^*(z_i, \boldsymbol{x}_{-i})}{\Delta_i^*(\boldsymbol{x})} dz_i, \\
&= x_i + \int_{x_i}^{\psi^{-1}[\psi(x_{i_1})+a(\beta_{i_1})-a(\beta_i)]} 1 \cdot dz_i + \sum_{j=1}^{t-1} \int_{\psi^{-1}[\psi(x_{i_j})+a(\beta_{i_j})-a(\beta_i)]}^{\psi^{-1}[\psi(x_{i_{j+1}})+a(\beta_{i_{j+1}})-a(\beta_i)]} \prod_{k=1}^{j} (1 - \beta_{i_k}) dz_i \\
&\quad + \int_{\psi^{-1}[\psi(x_{i_t})+a(\beta_{i_t})-a(\beta_i)]}^{r_i} \prod_{k=1}^{t} (1 - \beta_{i_k}) dz_i + \int_{r_i}^{1} 0 \cdot dz_i,
\end{aligned}$$

$$= \sum_{j=1}^{t} \psi^{-1}[\psi(x_{i_j}) + a(\beta_{i_j}) - a(\beta_i)]\beta_{i_j} \prod_{k=1}^{j-1} (1 - \beta_{i_k}) + r_i \prod_{k=1}^{t} (1 - \beta_{i_k}), \text{ if } t \geq 1; \text{ or } = r_i, \text{ if } t = 0.$$

2.8.2 Proof of Proposition 2

Suppose all bidders $j \neq i$ drop out of the auction at their true adjusted virtual cost, and i is active at current bid b. Without loss generality, the sequence of bidders who dropped out by bid b are denoted by $1, 2, ..., i(b)$, where bidder 1 dropped out earliest and bidder $i(b)$ dropped out latest. The set of bidders besides bidder i who have not dropped out by bid b is denoted by $\zeta(b)$. Since $b_j = \psi(x_j) + a(\beta_j)$ for all $j \in \{1, 2, ..., i(b)\}$, the expected profit from dropping out at b is given by

$$\prod_{l\in\zeta(b)}(1-\beta_l)\Big\{\sum_{j=1}^{i(b)}\psi^{-1}[b_j-a(\beta_i)]\beta_j\prod_{k=j+1}^{i(b)}(1-\beta_k)+r_i\prod_{k=1}^{i(b)}(1-\beta_k)-x_i\Big\}$$
$$=\sum_{j=1}^{i(b)}\{\psi^{-1}[b_j-a(\beta_i)]-x_i\}\beta_j\prod_{l\in\zeta(b_j)}(1-\beta_l)+(r_i-x_i)\prod_{l\in\zeta(C_o)}(1-\beta_l),$$

which decreases in b when $b>\psi(x_i)+a(\beta_i)$ and increases in b when $b<\psi(x_i)+a(\beta_i)$, since $\psi^{-1}[b_j-a(\beta_i)]\geq x_i$ for all $j\in\{1,2,...,i(\psi(x_i)+a(\beta_i))\}$, $\psi^{-1}[b_j-a(\beta_i)]\leq x_i$ for all $j\in\zeta(\psi(x_i)+a(\beta_i))$, and the weights $\beta_j\prod_{l\in\zeta(b_j)}(1-\beta_l)$ and the term $(r_i-x_i)\prod_{l\in\zeta(C_o)}(1-\beta_l)$ do not change with b. Therefore, it is optimal for bidder i to remain in the auction unless the current bid dips below his true adjusted virtual cost $\psi(x_i)+a(\beta_i)$, at which point he drops out, given that all other bidders do the same.

2.8.3 Proof of Equation (4.5)

Given (2.1), it is easy to see that $\overline{\text{PRE}}$ strictly increases in β_i for all i, since $q(\beta_i)$ strictly increases in β_i (recall that H has a continuous and positive density).

From the proof of Proposition 1, the expected total auction and post-qualification cost under the optimal mechanism is

$$\overline{\text{PAY}}+\overline{\text{POST}}+\overline{\text{NT}}=\int_{\boldsymbol{x}}\left\{C_o+\sum_{i=1}^{n}\Delta_i(\boldsymbol{x})[\psi(x_i)+a(\beta_i)-C_o]\right\}d\mathbf{F}(\boldsymbol{x}),$$

where $\Delta_i(\boldsymbol{x})$ is bidder i's probability of winning derived from the optimal sequencing rule via equation (2.2) given $(\beta_1,\ldots,\beta_n)$ and a realization of $\boldsymbol{x}$. Consider $(\hat{\beta}_1,\ldots,\hat{\beta}_n)$ such that $\hat{\beta}_k>\beta_k$ for some $k\in\{1,\ldots,n\}$ and $\hat{\beta}_i=\beta_i$ for all $i\neq k$, and let $\hat{\Delta}_i(\boldsymbol{x})$ be i's probability of winning corresponding to the realization of $\boldsymbol{x}$ and the optimal sequencing rule given $(\hat{\beta}_1,\ldots,\hat{\beta}_n)$. We show that $\overline{\text{PAY}}+\overline{\text{POST}}+\overline{\text{NT}}$ decreases in β_k by showing $\overline{\text{PAY}}+\overline{\text{POST}}+\overline{\text{NT}}$

with $(\beta_1, \dots, \beta_n)$ is no less than that with $(\hat{\beta}_1, \dots, \hat{\beta}_n)$. In particular,

$$\begin{aligned} C_o + \sum_{i=1}^{n} \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o] &\geq C_o + \sum_{i=1}^{n} \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\hat{\beta}_i) - C_o] \\ &\geq C_o + \sum_{i=1}^{n} \hat{\Delta}_i(\boldsymbol{x})[\psi(x_i) + a(\hat{\beta}_i) - C_o]. \end{aligned}$$

The first inequality (which is strict when $\Delta_k(\boldsymbol{x}) > 0$) results from the fact that $a(\hat{\beta}_i) = a(\beta_i)$ for all $i \neq k$, while $a(\beta_k) > a(\hat{\beta}_k)$ since $a(\beta_k)$ strictly decreases in β_k (by equation (2.5) and the fact that $q(\beta_k)$ strictly increases in β_k). The second inequality results from the fact that the $\hat{\Delta}_i(\boldsymbol{x})$'s are derived from the optimal sequencing rule which minimizes the expected total auction and post-qualification cost given $\boldsymbol{x}$ and $(\hat{\beta}_1, \dots, \hat{\beta}_n)$. As the first inequality is strict when $\Delta_k(\boldsymbol{x}) > 0$, we conclude that $\overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}$ strictly decreases in β_k if $\int_{\boldsymbol{x}} 1_{\{\psi(x_k)+a(\beta_k) \leq C_o\}} d\mathbf{F}(\boldsymbol{x}) > 0$.

2.8.4 Proof of Proposition 3

In the following, Step 1 derives the unique $K^{nt} < \infty$ and shows the existence of K^{post}; Step 2 shows the existence of $K^{pre} > 0$. As $K^{pre} \leq K^{post} < K^{nt}$ is straightforward by definition, we then conclude that $0 < K^{pre} \leq K^{post} < K^{nt} < \infty$. Finally, Step 3 shows that K^{nt} is increasing in C_o, and that the largest of all such thresholds K^{pre} and the smallest of all such thresholds K^{post} are both increasing in C_o.

Step 1. Given that $a(\underline{\beta}) = \frac{1}{1-H(q_0)} \int_{q_i=0}^{q_0} \frac{q_i}{q_0} K dH(q_i) + K$ is linearly increasing in K, it is easy to check that $C_o - a(\underline{\beta}) = 0$ only when K equals $C_o / [\frac{1}{1-H(q_0)} \int_{q_i=0}^{q_0} \frac{q_i}{q_0} dH(q_i) + 1] < \infty$. We show that this value, denoted as K^{nt} hereafter, is exactly the unique threshold defined in the proposition, such that running the auction with post-only is better than the outside option whenever $K < K^{nt}$, while it is optimal to forgo the auction in favor of the outside option whenever $K > K^{nt}$. We first prove the following lemma, which helps establish K^{nt} as well as K^{post}.

Lemma 1 *It is optimal to run the auction with post-only when K is close enough to K^{nt} from below.*

Proof. Let $\overline{\text{TOTAL}} \triangleq \overline{\text{PRE}} + \overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}$ denote the total ex ante expected cost of the buyer. Let $\overline{\text{TOTAL}}|_{K,\,\boldsymbol{\beta}}$ denote the expected total cost with K and pre-qualification probabilities $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_n)$, given a fixed C_o. We can construct an $\varepsilon > 0$ such that $\overline{\text{TOTAL}}|_{K,\,\boldsymbol{\beta}} > \overline{\text{TOTAL}}|_{K,\,\underline{\boldsymbol{\beta}}}$, $\forall K \in [K^{nt} - \varepsilon, K^{nt}]$, $\forall \boldsymbol{\beta} \succ \underline{\boldsymbol{\beta}}$, where $\boldsymbol{\beta} \succ \underline{\boldsymbol{\beta}}$ means $\beta_j \geq \underline{\beta}\ \forall\ j$ and $\beta_i > \underline{\beta}$ for some i. The construction is based on two facts.

Fact (i): $\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K^{nt},\,\underline{\boldsymbol{\beta}}} > 0$ for all i. It suffices to show $\frac{\partial [\overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}]}{\partial \beta_i}|_{K^{nt},\,\underline{\boldsymbol{\beta}}} = 0$, given that $\frac{\partial \overline{\text{PRE}}}{\partial \beta_i}|_{K^{nt},\,\underline{\boldsymbol{\beta}}} > 0$ by the proof of equation (4.5) (Section 2.8.3).

$$
\begin{aligned}
\frac{\partial[\overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}]}{\partial \beta_i}|_{K^{nt},\,\underline{\boldsymbol{\beta}}} \;&=\; \lim_{\beta_i \downarrow \underline{\beta}} \frac{\partial \int_{\boldsymbol{x}} \{C_o + \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]\}\, d\mathbf{F}(\boldsymbol{x})}{\partial \beta_i}|_{K^{nt}}, \\
= \lim_{\beta_i \downarrow \underline{\beta}} & \frac{\partial \int_0^{\psi^{-1}[C_o - a(\beta_i)]} \beta_i[\psi(x_i) + a(\beta_i) - C_o] dF(x_i)}{\partial \beta_i}|_{K^{nt}}, \\
= \lim_{\beta_i \downarrow \underline{\beta}} & \int_0^{\psi^{-1}[C_o - a(\beta_i)]} \frac{\partial}{\partial \beta_i}\{\beta_i[\psi(x_i) + a(\beta_i) - C_o]\} dF(x_i)|_{K^{nt}} \\
+ \lim_{\beta_i \downarrow \underline{\beta}} & \frac{\partial \psi^{-1}[C_o - a(\beta_i)]}{\partial \beta_i} f(\psi^{-1}(C_o - a(\beta_i)))\beta_i[\psi(\psi^{-1}(C_o - a(\beta_i))) + a(\beta_i) - C_o]|_{K^{nt}} = 0.
\end{aligned}
$$

The first equality holds because $\Delta_j(\boldsymbol{x}) = 0$ for all $j \neq i$ when $K = K^{nt}$ and $\beta_j = \underline{\beta}$; the second equality holds because $\Delta_i(\boldsymbol{x}) = 0$ when $\psi(x_i) + a(\beta_i) > C_o$ and $\Delta_i(\boldsymbol{x}) = \beta_i$ when $\psi(x_i) + a(\beta_i) \leq C_o$; the third equality is by Leibniz's integral rule; and the fourth equality holds because both terms before the fourth equality sign equal zero.

Fact (ii): $\overline{\text{TOTAL}}|_{K^{nt},\,\boldsymbol{\beta}} > \overline{\text{PRE}}|_{K^{nt},\,\boldsymbol{\beta}} + \overline{\text{POST}}|_{K^{nt},\,\boldsymbol{\beta}} + \overline{\text{NT}}|_{K^{nt},\,\boldsymbol{\beta}}$

$$
\geq \int_{\boldsymbol{x}} \sum_{i=1}^{n} \Delta_i(\boldsymbol{x}) C_o d\mathbf{F}(\boldsymbol{x}) + \int_{\boldsymbol{x}} [1 - \sum_{i=1}^{n} \Delta_i(\boldsymbol{x})] C_o d\mathbf{F}(\boldsymbol{x}) = C_o = \overline{\text{TOTAL}}|_{K^{nt},\,\underline{\boldsymbol{\beta}}}.
$$

The first inequality follows from $\overline{\text{PAY}}|_{K^{nt},\,\boldsymbol{\beta}} > 0$, which holds due to Proposition 1 and because the probability of contracting with a qualified bidder is positive when $K = K^{nt}$ and $\boldsymbol{\beta} \succ \underline{\boldsymbol{\beta}}$ (since $f(x) > 0$ for all $x \in [0, 1]$). The final equality follows because $\overline{\text{PRE}}|_{K^{nt},\,\underline{\boldsymbol{\beta}}} = 0$

by equation (2.1) and $[\overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}}]|_{K^{nt},\, \underline{\boldsymbol{\beta}}} = C_o$ by Proposition 1. The second inequality is from the expression for $\overline{\text{NT}}|_{K^{nt},\, \boldsymbol{\beta}}$ on page 18, and $\overline{\text{PRE}}|_{K^{nt},\, \boldsymbol{\beta}} + \overline{\text{POST}}|_{K^{nt},\, \boldsymbol{\beta}} \geq \int_{\boldsymbol{x}} \sum_{i=1}^{n} \Delta_i(\boldsymbol{x}) C_o d\mathbf{F}(\boldsymbol{x})$, which is true by the expression for $\overline{\text{POST}}|_{K^{nt},\, \boldsymbol{\beta}}$ on page 18 and

$$\begin{aligned}\overline{\text{PRE}}|_{K^{nt},\, \boldsymbol{\beta}} &= \sum_{i=1}^{n} \{ \frac{q(\beta_i) K^{nt}}{q_0} + \left[\frac{1}{1 - H(q(\beta_i))} - 1 \right] \int_{y=0}^{q(\beta_i)} \frac{y K^{nt} dH(y)}{q_0 H(q(\beta_i))} \} \quad \text{by (2.1)}, \\ &= \sum_{i=1}^{n} \beta_i [a(\underline{\beta}) - a(\beta_i)]|_{K=K^{nt}} \geq \int_{\boldsymbol{x}} \sum_{i=1}^{n} \Delta_i(\boldsymbol{x}) [C_o - a(\beta_i)] d\mathbf{F}(\boldsymbol{x})|_{K=K^{nt}}.\end{aligned}$$

We can now construct the $\varepsilon > 0$ as follows. We can write $\overline{\text{PAY}} + \overline{\text{POST}} + \overline{\text{NT}} =$

$$\sum_{i=1}^{n} \int_{\{x_i | \psi(x_i) + a(\beta_i) \leq C_o\}} \underbrace{\beta_i \prod_{j \neq i} \left[1 - \beta_j F(\psi^{-1}[\psi(x_i) + a(\beta_i) - a(\beta_j)])\right]}_{\text{probability that } i \text{wins contract, given } x_i} (\psi(x_i) + a(\beta_i)) dF(x_i)$$
$$+ C_o \underbrace{\prod_{j=1}^{n} \left[1 - \beta_j F(\psi^{-1}[C_o - a(\beta_j)])\right]}_{\text{probability that no bidder wins the contract}} . \quad (2.16)$$

With (2.16) and the expression (2.1) for $\overline{\text{PRE}}$, we can conclude that $\overline{\text{TOTAL}}$ is continuously differentiable in β_i, a fact which is useful in finding ϵ.

Given that $\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K^{nt}}$ is continuous in $\boldsymbol{\beta}$ and positive at $\underline{\boldsymbol{\beta}}$ by Fact (i), there exists $\delta > 0$ such that $\eta^{(1)} \triangleq \min\{\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K^{nt},\, \boldsymbol{\beta}} : \forall i, \forall \boldsymbol{\beta} \in [\underline{\beta}, \underline{\beta} + \delta]^n\} > 0$. Given that $\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K, \boldsymbol{\beta}}$ is continuous in $(K, \boldsymbol{\beta})$ and hence uniformly continuous on $[0, K^{nt}] \times [\underline{\beta}, \underline{\beta} + \delta]^n$, there exists an $\varepsilon^{(1)} > 0$ such that $|\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K, \boldsymbol{\beta}} - \frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K^{nt}, \boldsymbol{\beta}}| < \frac{\eta^{(1)}}{2}$ for all $K \in [K^{nt} - \varepsilon^{(1)}, K^{nt}]$ and all $\boldsymbol{\beta} \in [\underline{\beta}, \underline{\beta} + \delta]^n$. Therefore, $\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K, \boldsymbol{\beta}} > 0$ for all $K \in [K^{nt} - \varepsilon^{(1)}, K^{nt}]$ and all $\boldsymbol{\beta} \in [\underline{\beta}, \underline{\beta} + \delta]^n$, and hence $\overline{\text{TOTAL}}|_{K,\, \boldsymbol{\beta}} > \overline{\text{TOTAL}}|_{K,\, \underline{\boldsymbol{\beta}}}$ for all $K \in [K^{nt} - \varepsilon^{(1)}, K^{nt}]$ and all $\boldsymbol{\beta} \in [\underline{\beta}, \underline{\beta} + \delta]^n \backslash \{\underline{\boldsymbol{\beta}}\}$.

Next, we address the result over the complement set $[\underline{\beta}, 1]^n \backslash [\underline{\beta}, \underline{\beta} + \delta]^n$; to enable our argument, we will actually prove the result over a slightly larger set, $[\underline{\beta}, 1]^n \backslash [\underline{\beta}, \underline{\beta} + \delta)^n$, which is compact. Given that $\overline{\text{TOTAL}}|_{K^{nt},\, \boldsymbol{\beta}} - \overline{\text{TOTAL}}|_{K^{nt},\, \underline{\boldsymbol{\beta}}}$ is positive by Fact (ii) and continuous in $\boldsymbol{\beta}$ on $[\underline{\beta}, 1]^n \backslash [\underline{\beta}, \underline{\beta} + \delta)^n$, there exists $\eta^{(2)} \triangleq \min\{\overline{\text{TOTAL}}|_{K^{nt},\, \boldsymbol{\beta}} - \overline{\text{TOTAL}}|_{K^{nt},\, \underline{\boldsymbol{\beta}}} : \boldsymbol{\beta} \in$

$[\underline{\beta}, 1]^n \backslash [\underline{\beta}, \underline{\beta} + \delta)^n\} > 0$. Given that $\overline{\text{TOTAL}}|_{K,\, \boldsymbol{\beta}} - \overline{\text{TOTAL}}|_{K,\, \underline{\boldsymbol{\beta}}}$ is continuous in $(K, \boldsymbol{\beta})$ and hence uniformly continuous on $[0, K^{nt}] \times \{[\underline{\beta}, 1]^n \backslash [\underline{\beta}, \underline{\beta} + \delta)^n\}$, there exists an $\varepsilon^{(2)} > 0$ such that $|[\overline{\text{TOTAL}}|_{K,\, \boldsymbol{\beta}} - \overline{\text{TOTAL}}|_{K,\, \underline{\boldsymbol{\beta}}}] - [\overline{\text{TOTAL}}|_{K^{nt},\, \boldsymbol{\beta}} - \overline{\text{TOTAL}}|_{K^{nt},\, \underline{\boldsymbol{\beta}}}]| < \frac{\eta^{(2)}}{2}$ for all $K \in [K^{nt} - \varepsilon^{(2)}, K^{nt}]$ and $\boldsymbol{\beta} \in [\underline{\beta}, 1]^n \backslash [\underline{\beta}, \underline{\beta} + \delta)^n$. It is easy to check that $\overline{\text{TOTAL}}|_{K,\, \boldsymbol{\beta}} - \overline{\text{TOTAL}}|_{K,\, \underline{\boldsymbol{\beta}}} > 0$ for all $K \in [K^{nt} - \varepsilon^{(2)}, K^{nt}]$ and $\boldsymbol{\beta} \in [\underline{\beta}, 1]^n \backslash [\underline{\beta}, \underline{\beta} + \delta)^n$.

Therefore, $\overline{\text{TOTAL}}|_{K,\, \boldsymbol{\beta}} > \overline{\text{TOTAL}}|_{K,\, \underline{\boldsymbol{\beta}}}$ for all $K \in [K^{nt} - \varepsilon, K^{nt}]$ and all $\boldsymbol{\beta} \succ \underline{\boldsymbol{\beta}}$, where $\varepsilon = \min\{\varepsilon^{(1)}, \varepsilon^{(2)}\}$. That is, post-only is optimal for $K \in [K^{nt} - \varepsilon, K^{nt}]$. ■

We now show that K^{nt} is indeed the threshold sought. We can write $\overline{\text{TOTAL}} = \overline{\text{PRE}} + C_o + \int_{\boldsymbol{x}} \{\sum_{i=1}^n \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]\}\, d\mathbf{F}(\boldsymbol{x})$. For $\boldsymbol{x}$ such that $\frac{\partial \Delta_i(\boldsymbol{x})}{\partial K} = 0$, we have that the derivative of $\Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]$ with respect to K is $\Delta_i(\boldsymbol{x}) \frac{\partial a(\beta_i)}{\partial K}$, which is nonnegative and strictly positive if $\Delta_i(\boldsymbol{x}) > 0$ and $\boldsymbol{\beta} \neq (1, \ldots, 1)$. Furthermore, because $\Delta_i(\boldsymbol{x})$ only changes at a finite number of K's for which the buyer is indifferent between the post-qualification sequence of one or more bidders, we have that $\Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]$ is continuous in K. Together with $\frac{\partial \overline{\text{PRE}}}{\partial K}$ nonnegative and strictly positive if $\boldsymbol{\beta} \succ \underline{\boldsymbol{\beta}}$, we conclude that $\overline{\text{TOTAL}}|_{K,\boldsymbol{\beta}}$ increases in K and increases strictly for $K < K^{nt}$ (for which the probability of award to a bidder is non-negative).

Running the auction with post-only is better than the outside option when $K < K^{nt}$ because $\overline{\text{TOTAL}}|_{K,\, \underline{\boldsymbol{\beta}}} < \overline{\text{TOTAL}}|_{K^{nt},\, \underline{\boldsymbol{\beta}}} = C_o$, where the inequality is due to the fact that the expected total cost strictly increases in $K < K^{nt}$, and the equality was established in Fact (ii) in the proof of Lemma 1.

It is optimal to forgo the auction in favor of the outside option when $K > K^{nt}$ because $\overline{\text{TOTAL}}|_{K,\, \boldsymbol{\beta}} \geq \overline{\text{TOTAL}}|_{K^{nt},\, \boldsymbol{\beta}} \geq C_o$, where the first inequality holds because expected total cost increases in K, and the second inequality follows from arguments analogous to those used to establish Fact (ii) in the proof of Lemma 1.

To complete Step 1, note that Lemma 1 automatically implies the existence of K^{post}.

Step 2. The existence of K^{pre} is due to the existence of K^{nt} and the fact that it is always optimal to do pre-only when $K = 0$, true because $\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}|_{K=0} < 0$. We omit the proof of this, which is straightforward since when $K = 0$ increasing β_i towards 1 is costless and only increases the chance of transacting with an attractive bid. Given that $\frac{\partial \overline{\text{TOTAL}}}{\partial \beta_i}$ is uniformly continuous in (K, β_i) on $[0, K^{nt}] \times [\underline{\beta}, 1]$, $\overline{\text{TOTAL}}$ strictly decreases in β_i on the full support of $[\underline{\beta}, 1]$ when K is positive and close enough to zero. This implies that $K^{pre} > 0$.

Step 3. By the definition of K^{nt}, we can conclude that K^{nt} is increasing in C_o. The following lemma implies that the largest threshold K^{pre} defined above and the smallest threshold K^{post} defined above are both increasing in C_o.

Lemma 2 *Suppose $\hat{C}_o > \tilde{C}_o > 0$. Let $\hat{\boldsymbol{\beta}}^* = (\hat{\beta}_1^*, \ldots, \hat{\beta}_n^*)$ and $\tilde{\boldsymbol{\beta}}^* = (\tilde{\beta}_1^*, \ldots, \tilde{\beta}_n^*)$ be the optimal pre-qualification probabilities corresponding to $\hat{C}_o$ and $\tilde{C}_o$, respectively for a fixed $K < K^{nt}|_{\tilde{C}_o}$. It is impossible that $\tilde{\boldsymbol{\beta}}^* \succ \hat{\boldsymbol{\beta}}^*$.*

Proof. First we show that if $\hat{C}_o > \tilde{C}_o$ and $\hat{\beta}_i < \tilde{\beta}_i$, then for fixed $K < K^{nt}|_{\tilde{C}_o}$ and $(\beta_1, \ldots, \beta_{i-1}, \beta_{i+1}, \ldots, \beta_n)$,

$$\overline{\text{TOTAL}}|_{\hat{C}_o, \tilde{\beta}_i} - \overline{\text{TOTAL}}|_{\tilde{C}_o, \tilde{\beta}_i} < \overline{\text{TOTAL}}|_{\hat{C}_o, \hat{\beta}_i} - \overline{\text{TOTAL}}|_{\tilde{C}_o, \hat{\beta}_i}. \tag{2.17}$$

Recall that $\overline{\text{TOTAL}} = \overline{\text{PRE}} + C_o + \int_{\boldsymbol{x}} \{\sum_{i=1}^n \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]\} d\mathbf{F}(\boldsymbol{x})$. As C_o increases, $\overline{\text{PRE}}$ remains constant. For $\boldsymbol{x}$ such that $\frac{\partial \Delta_i(\boldsymbol{x})}{\partial C_o} = 0$, we have that the derivative of $C_o + \sum_{i=1}^n \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]$ with respect to C_o is $1 - \sum_{i=1}^n \Delta_i$. Furthermore, because $\Delta_i(\boldsymbol{x})$ only changes at a finite number of C_o's for which the buyer is indifferent between the post-qualification sequence of one or more bidders, we have that $C_o + \sum_{i=1}^n \Delta_i(\boldsymbol{x})[\psi(x_i) + a(\beta_i) - C_o]$ is continuous in C_o, and piecewise linearly increases with a slope $1 - \sum_{i=1}^n \Delta_i(\boldsymbol{x})$.

Given fixed $\boldsymbol{x}$, K, $\tilde{C}_o$, $\tilde{\beta}_i$, $\hat{\beta}_i$, and $(\beta_1, \ldots, \beta_{i-1}, \beta_{i+1}, \ldots, \beta_n)$, both sides of inequality (2.17) (evaluated at $\boldsymbol{x}$ rather than in expectation) are continuous and piecewise linearly increasing in $\hat{C}_o$. Therefore, it suffices to prove inequality (2.17) by showing, for fixed $\boldsymbol{x}$, the left hand side always has a larger slope than the right hand side. This is true as $1 - \sum_{i=1}^n \Delta_i(\boldsymbol{x}) = \prod_{\{k \in \{1,\ldots,n\} | \ \psi(x_k) + a(\beta_k) \leq C_o\}} (1 - \beta_k)$ is decreasing in β_i. As the slope $1 - \sum_{i=1}^n \Delta_i(\boldsymbol{x})$ is strictly decreasing in β_i if $\psi(x_i) + a(\beta_i) \leq C_o$, and $K < K^{nt}|_{\tilde{C}_o}$ implies

$\int_{\boldsymbol{x}} 1_{\{\psi(x_i)+a(\beta_i)\leq C_o\}} d\mathbf{F}(\boldsymbol{x}) > 0$ for all $C_o \in [\tilde{C}_o, \hat{C}_o]$ and $\beta_i \in [\hat{\beta}_i, \tilde{\beta}_i]$, we have that the inequality in (2.17) is strict.

By definition of $\hat{\boldsymbol{\beta}}^*$ and $\tilde{\boldsymbol{\beta}}^*$, it must be true that $\overline{\text{TOTAL}}|_{\hat{C}_o,\hat{\boldsymbol{\beta}}^*} \leq \overline{\text{TOTAL}}|_{\hat{C}_o,\tilde{\boldsymbol{\beta}}^*}$, and $\overline{\text{TOTAL}}|_{\tilde{C}_o,\tilde{\boldsymbol{\beta}}^*} \leq \overline{\text{TOTAL}}|_{\tilde{C}_o,\hat{\boldsymbol{\beta}}^*}$, which together imply that

$$\overline{\text{TOTAL}}|_{\hat{C}_o,\hat{\boldsymbol{\beta}}^*} - \overline{\text{TOTAL}}|_{\tilde{C}_o,\hat{\boldsymbol{\beta}}^*} \leq \overline{\text{TOTAL}}|_{\hat{C}_o,\tilde{\boldsymbol{\beta}}^*} - \overline{\text{TOTAL}}|_{\tilde{C}_o,\tilde{\boldsymbol{\beta}}^*}.$$

It is impossible that $(\tilde{\beta}_1^*, \ldots, \tilde{\beta}_n^*) \succ (\hat{\beta}_1^*, \ldots, \hat{\beta}_n^*)$, given that inequality (2.17) holds. ■

2.8.5 Proof of Proposition 4

Proof of the equilibrium bidding strategy is similar to the classical proof of the symmetric first-price forward auction bidding strategy; readers can refer to Krishna (2002), pp16–19, aware of the fact that a bidder bidding z_i will win the contract (be the lowest *qualified* bidder) with probability $\beta[1-\beta F(z_i)]^{n-1}$. With this bidding strategy, it is straightforward to verify that the sealed-bid first-qualified-price auction with reserve price r satisfies the conditions of Proposition 1.

2.8.6 Proof of Proposition 5

With the general cost model, our analysis of the optimal mechanism involves changes as follows. The sequencing rule Q, the payment rule M, and the winning probabilities Δ_i's should be redefined as functions of the collection of private information $\boldsymbol{\theta}$. Let $\mathbf{F}_{-i}$ and $\mathbf{F}$ be the joint distributions of $\boldsymbol{\theta}_{-i}$ and $\boldsymbol{\theta}$, respectively. Let $\chi_i(z_i, \theta_i) \triangleq \int_{\boldsymbol{\theta}_{-i}} \Delta_i(z_i, \boldsymbol{\theta}_{-i}) x_i(\boldsymbol{\theta}) d\mathbf{F}_{-i}(\boldsymbol{\theta}_{-i})$ be the expected cost bidder i incurs in the contract allocation when i reports his private information as z_i and all others report their true private information. Let $\mathrm{m}_i(z_i) \triangleq \int_{\boldsymbol{\theta}_{-i}} \mathrm{M}_i(z_i, \boldsymbol{\theta}_{-i}) d\mathbf{F}_{-i}(\boldsymbol{\theta}_{-i})$ be the expected payment to bidder i when i reports his private information as z_i and all others report their true private information. Equations (2.3) and

(2.4) become

$$
\begin{aligned}
U_i(\theta_i) &\triangleq \mathrm{m}_i(\theta_i) - \chi_i(\theta_i, \theta_i) = \max_{z_i \in [0,1]} \mathrm{m}_i(z_i) - \chi_i(z_i, \theta_i), && (2.18)\\
U_i(\theta_i) &\geq 0. && (2.19)
\end{aligned}
$$

Since the envelope theorem implies $U_i'(\theta_i) = -\frac{\partial \chi_i(\theta_i, \hat{\theta}_i)}{\partial \hat{\theta}_i}|_{\hat{\theta}_i = \theta_i}$, equation (2.6) becomes

$$
\mathrm{m}_i(\theta_i) - \chi_i(\theta_i, \theta_i) = \mathrm{m}_i(1) - \chi_i(1,1) + \int_{\theta_i}^{1} \frac{\partial \chi_i(z_i, \hat{\theta}_i)}{\partial \hat{\theta}_i}|_{\hat{\theta}_i = z_i} dz_i. \qquad (2.20)
$$

Consequently, equation (2.7) changes to

$$
\overline{\mathrm{PAY}} = \sum_{i=1}^{n} [\mathrm{m}_i(1) - \chi_i(1,1)] + \sum_{i=1}^{n} \int_{\boldsymbol{\theta}} \left[\Delta_i(\boldsymbol{\theta}) \hat{\psi}_i(\boldsymbol{\theta}) \right] d\mathbf{F}(\boldsymbol{\theta}), \qquad (2.21)
$$

where

$$
\hat{\psi}_i(\boldsymbol{\theta}) = x_i(\boldsymbol{\theta}) + \frac{F(\theta_i)}{f(\theta_i)} \frac{\partial x_i(\boldsymbol{\theta})}{\partial \theta_i}.
$$

If (2.20) holds, incentive compatibility constraint (2.18) is guaranteed if the post-qualification sequencing rule is such that $\Delta_i(\boldsymbol{\theta})$ is decreasing in θ_i. (For the moment, we postpone the proof of this.) Furthermore, from (2.20), we have that individual rationality constraint (2.19) holds as long as the payment and post-qualification sequencing rules satisfy $\mathrm{m}_i(1) - \chi_i(1,1) \geq 0$. Clearly, equation (2.21) implies that a cost-minimizing buyer will set $\mathrm{m}_i(1) - \chi_i(1,1) = 0$. Because $\overline{\mathrm{POST}}$ and $\overline{\mathrm{NT}}$ are computed as in §2.4.2, to find the cost-minimizing direct mechanism $(\mathrm{Q}^*, \mathrm{M}^*)$ employed by the buyer in the auction and post-qualification stage, it suffices to solve the following program,

$$
\min_{\mathrm{Q,M}} \quad \sum_{i=1}^{n} \int_{\boldsymbol{\theta}} \Delta_i(\boldsymbol{\theta}) \left[\hat{\psi}_i(\boldsymbol{\theta}) + a(\beta_i) - C_o \right] d\mathbf{F}(\boldsymbol{\theta}) + C_o \qquad (2.22a)
$$

$$\text{Subject to } \mathrm{m}_i(\theta_i) = \chi_i(\theta_i, \theta_i) + \int_{\theta_i}^{1} \frac{\partial \chi_i(z_i, \hat{\theta}_i)}{\partial \hat{\theta}_i}|_{\hat{\theta}_i = z_i} dz_i \quad \forall i, \tag{2.22b}$$

and verify that at the solution $\Delta_i(\boldsymbol{\theta})$ decreases in $\theta_i \in [0,1]$. We now explain why (2.22b) (i.e, (2.20)) and that $\Delta_i(\boldsymbol{\theta})$ decreases in θ_i, together imply that (2.18) holds. First, define $g(z_i, \theta_i) \triangleq \mathrm{m}_i(z_i) - \chi_i(z_i, \theta_i)$. For a given function v, define $D_k v(r,t)$ to be the derivative of v with respect to its k^{th} argument, evaluated at (r,t). For example, $D_1 x_i(\theta_i, \boldsymbol{\theta}_{-i}) = \frac{\partial x_i(z_i, \boldsymbol{\theta}_{-i})}{\partial z_i}|_{z_i = \theta_i}$. We have

$$D_1 D_2 g(z_i, \theta_i) = -\frac{\partial}{\partial z_i} \int_{\boldsymbol{\theta}_{-i}} \Delta_i(z_i, \boldsymbol{\theta}_{-i}) D_1 x_i(\theta_i, \boldsymbol{\theta}_{-i}) d\mathbf{F}_{-i}(\boldsymbol{\theta}_{-i}) \geq 0, \tag{2.23}$$

where the inequality holds since Δ_i decreases in θ_i, and x_i increases in θ_i. Furthermore,

$$\begin{aligned} D_1 g(z_i, \hat{\theta}_i) &= D_1 \chi_i(z_i, z_i) - D_1 \chi_i(z_i, \hat{\theta}_i) \qquad \text{by using (2.22b) to express } m_i(z_i), \\ &= \frac{\partial}{\partial z_i} \int_{\boldsymbol{\theta}_{-i}} \Delta_i(z_i, \boldsymbol{\theta}_{-i}) [x_i(z_i, \boldsymbol{\theta}_{-i}) - x_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})] d\mathbf{F}_{-i}(\boldsymbol{\theta}_{-i}). \end{aligned} \tag{2.24}$$

We now show that $D_1 g(z_i, \hat{\theta}_i)$ increases in $\hat{\theta}_i$. For $\theta_i < \hat{\theta}_i$, (2.24) implies

$$D_1 g(z_i, \hat{\theta}_i) - D_1 g(z_i, \theta_i) = \frac{\partial}{\partial z_i} \int_{\boldsymbol{\theta}_{-i}} \Delta_i(z_i, \boldsymbol{\theta}_{-i}) [x_i(\theta_i, \boldsymbol{\theta}_{-i}) - x_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})] d\mathbf{F}_{-i}(\boldsymbol{\theta}_{-i}) \geq 0,$$

where the inequality holds because Δ_i decreases in θ_i, and x_i increases in θ_i. Thus, $\theta_i < \hat{\theta}_i$ and $D_1 g(z_i, \theta_i) = 0$ together imply $D_1 g(z_i, \hat{\theta}_i) \geq 0$, while $\theta_i > \hat{\theta}_i$ and $D_1 g(z_i, \theta_i) = 0$ together imply $D_1 g(z_i, \hat{\theta}_i) \leq 0$. This together with (2.23) establishes that g satisfies the smooth single crossing differences property (see Milgrom 2004 p101). Furthermore, $g(\theta_i, \theta_i) = \int_{\theta_i}^{1} D_2 g(z_i, z_i) dz_i$ by (2.22b). Hence, by Theorem 4.2 of Milgrom (2004) (where, in Milgrom's notation, we have $t = \theta_i$ and set $\bar{x}(t) \triangleq t = \theta_i$), we have that $\theta_i = \arg\max_{z_i \in [0,1]} g(z_i, \theta_i)$, that is, incentive compatibility constraint (2.18) holds.

Having established the mechanism design program (2.22), it is straightforward to see that Proposition 1 holds with $\psi(x_i)$ replaced by $\hat{\psi}_i(\boldsymbol{\theta})$, remembering that Q, M, and Δ

are now functions of $\boldsymbol{\theta}$. Finally, as they do not rely on the format of the bidder's cost function x_i, proofs of the central tradeoff (equation (4.5)) and Proposition 3 continue to hold under the general cost model. The original proofs can be duplicated with notation changes, remembering that $\hat{\psi}_i$ is now a function of $\boldsymbol{\theta}$.

Chapter 3

Procurement Auctions with An Incumbent and Partially Qualified Entrant

3.1. Introduction

We study a procurement problem for a buyer who has an expiring contract with her incumbent supplier. Instead of directly renewing the contract with the incumbent, the buyer, approached by a new entrant supplier, wishes to conduct an open-descending procurement auction between the incumbent and the entrant, seeking either price concessions from the incumbent or a better price from the entrant. While the contract price is a concern for the buyer, she will not contract with the entrant unless the entrant is verified to be fully *qualified* for the business. As is common in industry, we will refer to performing *qualification screening* on a supplier as the act of verifying that a supplier is indeed able to comply with all the contract specifications (e.g., on product, delivery, packaging, etc.) with a reasonable degree of certainty. The qualification process is *costly*, which can include testing the entrant's products, visiting the entrant's production facilities, verifying the entrant's surge capacity availability, auditing the entrant's financial status, etc.

The buyer faces a strategic decision regarding the timing of performing qualification screening on the entrant. The buyer can choose to screen the entrant before conducting the auction, called "pre-qualification." If the entrant successfully passes pre-qualification an auction is held in which the low bid wins the contract. However, pre-qualification may backfire on the buyer if the entrant fails pre-qualification and must be discarded — in such a case the buyer not only wastes the qualification cost but also loses the opportunity to run the auction, forcing her to renew the incumbent's contract without any reduction in price.

Alternatively, the buyer can choose to delay screening the entrant until after the auction, at which point she screens the entrant only if the entrant wins the auction, called "pre-qualification." In this case the incumbent knows that he could lose the auction but still win the contract if the entrant fails post-qualification. We study the following **research questions:**

1. What is the incumbent's optimal bidding strategy under post-qualification? Will the incumbent boycott the auction (drop out at the reserve price), and if not, how aggressively will he bid?

2. How does the answer to question 1 depend on the probability that the entrant is truly qualified, the buyer's qualification cost, and the auction reserve price?

3. Under what circumstances will the buyer prefer to use post-qualification?

In answering research question 1, we find that under post-qualification the incumbent deploys one of three types of strategies: "boycott the auction," "test-the-water," and "bid-to-win." Under the boycott strategy, the incumbent drops out of the auction at the reserve price, and simply hopes that the entrant fails post-qualification. This strategy is used when the incumbent knows he is unlikely to beat the entrant on price alone, i.e., when the incumbent's cost is quite high. When the incumbent's true cost is moderate he uses the test-the-water strategy: he bids against the entrant in the hopes of clinching the contract on price alone, but does so only half-heartedly — if the entrant stays in the auction long enough, eventually the incumbent will abandon the effort and drop out before reaching his true cost. Only when the incumbent is certain he can beat the entrant on price alone will he deploy the bid-to-win strategy in which he lowers his bid until the entrant drops out. This novel result holds under quite general assumptions, and its predictions are markedly different from the canonical open-descending auction analysis (e.g., Krishna 2002 Chapter 2) in which there is no post-qualification stage and therefore all bidders have a dominant strategy to bid down to their true cost.

In addressing research question 2, we find that the incumbent bids less aggressively when the reserve price is large, since he finds boycotting the auction more tempting. Thus, high profit potential for the incumbent may, ironically, cause it to short-circuit competition rather than compete harder to retain the contract. Continuing to address research question 2, as the buyer's qualification cost increases, we find two countervailing effects on the incumbent's bidding strategy. On one hand, a higher qualification cost implies the buyer adds a higher "switching cost" to the entrant's bid, making it easier for the incumbent to beat the entrant on price alone as the qualification cost increases; however, on the other hand, a higher qualification cost makes the low-cost incumbent types who use the bid-to-win strategy effectively drop out at a higher price as the lowest possible entrant cost is higher. Thus, the incumbent may bid more, or less aggressively, when the buyer's cost of qualifying the entrant is higher, depending on which of two countervailing effects prevails. Analogously, as the entrant's probability of surviving qualification increases, the incumbent bids more, or less aggressively, depending on which of two effect prevails, and in general the effect is non-monotonic. On one hand, the entrant is more likely to survive post-qualification, encouraging the incumbent to try to win the auction. On the other hand, the buyer is more willing to post-qualify the entrant (who is more likely to survive the post-qualification) so the "switching cost" shrinks, scaring the incumbent away from trying to compete on price. The managerial implication is again that when the incumbent can lose the auction but win the contract, a tougher competitor — namely an entrant that is more likely to be qualified — might actually forestall competition.

Turning to research question 3, we prove that the buyer's decision follows a threshold: If the cost of qualifying the entrant becomes large enough, eventually the buyer will prefer the post-qualification strategy.

Key features of our novel incumbent bidding equilibrium are qualitatively consistent with the empirical findings. Zhong (2007) analyzes incumbent and entrant behavior in multi-item procurement auctions held by a large high-tech company. Relative to entrants, she finds that incumbents using what she calls "early- and mid-evaluator" strategies (roughly akin to our boycott and test-the-water strategies) have higher final bids, while incumbents using more aggressive strategies tend to have lower final bids. Similar to our equilibrium analysis of a stylized and simplified setting, this empirical data suggests incumbents seem to choose between timid testing and all-out competing for the contract. Moreover, Zhong (2007) finds that incumbents very often win the contract without being the lowest bidder, which is also consistent with our stylized post-qualification model.

The next section reviews related literature, followed by a discussion of the model in §3.3. Section 3.4 provides theoretical analyses of the model. Concluding remarks are provided in §3.5. Proofs of propositions are in §3.6.

3.2. Literature Review

Elmaghraby (2000) provides a detailed review of work on procurement in the operations and economics literature. Many such papers, including ours, apply auctions as the means of price discovery during the procurement process. Books by Krishna (2002) and Milgrom (2004) provide excellent treatments and detailed references on auctions.

This chapter studies how supplier qualification screening manifests itself in the auction bidding behavior of entrants and incumbents. Supplier qualification screening is a process by which the buyer performs due diligence to avoid consummating a transaction with a supplier who will not fulfill its obligations. As such, our work is related to other auction papers examining measures taken to redress non-performance by a counterparty. In the context of procurement auctions, Calveras et al. (2004) study how the buyer can require surety bonds to partially offset its cost if the supplier does not perform. In a forward auction context, Rothkopf (1991) and Waehrer (1995) study the use of deposits that are forfeited to the auctioneer in the event a winning bid is reneged. However, unlike surety bonds or bid deposits, which are measures imposed on bidders to reactively recoup losses in the event of non-performance, supplier qualification screening is a proactive measure employed upfront (at a cost to the buyer) to reduce the risk of supplier non-performance.

In this chapter, the incumbent, who is known to the buyer, has already passed qualification screening. The entrant, who is unknown to the buyer, has not yet been qualified. The existing literature has studied different features which make incumbent and entrant unalike. Zhong (2007) is an empirical analysis of incumbent and entrant behavior in procurement auctions, and hypothesizes that the differences between incumbent and entrant bidding behavior she observes could be due to non-price factors (such as quality) that differ between them. More generically, incumbent or entrant status can be used as motivation for studying bidders with asymmetric cost distributions; see Chapter 4.3 of Krishna (2002). In our analysis, in addition to the asymmetry over qualification, we allow the incumbent and entrant costs to follow different distributions.

While supplier qualification is common in practice, surprisingly little has been written about it in the procurement auction literature. To our knowledge, only one other paper studies supplier qualification in the context of procurement auctions. The paper (Chapter 2 of this dissertation), Wan and Beil (2008), focuses on the buyer's optimal auction design problem and studies how to optimally combine two separate phases of supplier qualification screening when suppliers are ex ante symmetric. In contrast, the present chapter focuses the suppliers' strategic bidding under the more common open-descending auction format, and studies the buyer's qualification timing decision when facing a finite supply pool (in this case, consisting of two suppliers who, as an incumbent and entrant, are ex ante asymmetric). In summary, the present chapter contributes new theory to the nascent auctions with qualification screening literature and provides the first experimental analysis of auctions with a possibly unqualified bidder.

3.3. Model

We consider a procurement manager, or *buyer*, who seeks to award a single, indivisible contract for goods or services. The buyer already has a pre-existing *incumbent* supplier, denoted by i, who currently performs the contract. As is common in practice, we assume that the contract covers a finite period of time (e.g., one to two years), after which point it must be reinstated. To this end we assume that the buyer becomes aware of an *entrant*, denoted by e, a new supplier who approaches the buyer seeking out new business. The buyer is interested in leveraging supply-side competition for the contract by conducting an auction in which she solicits competing bids from both the incumbent and the entrant. We let R denote the price the incumbent currently charges the buyer for the contract; thus, the incumbent's true cost to perform the contract, denoted by x_i, is assumed to be at most R. We assume that x_i is distributed according to a c.d.f. F_i (with p.d.f. $f_i > 0$) on the support $[l, R]$ where $l < 1 \leq R$, and that the entrant's true cost x_e follows a c.d.f F_e (with p.d.f. $f_e > 0$) on the support $[0, 1]$. We assume that x_i and x_e are privately known and independently distributed, and the distributions F_i and F_e are common knowledge. We assume that both suppliers seek to maximize their expected utility. We let $U(\cdot)$ denote the incumbent's utility function. Thus, the utility of an incumbent with true cost x_i is $U(p - x_i)$ if he wins the contract and receives payment p from the buyer, or is $U(0)$ if he does not win the contract. We assume that $U(\cdot)$ is concave (i.e., the incumbent is risk-neutral or risk-averse) and $\frac{U'(p-x_i)}{U(p-x_i)-U(0)}$ goes to zero as $p - x_i$ goes to infinity. In our model setting, the entrant will have a dominant bidding strategy (see §3.4); thus, we do not explicitly specify the entrant's utility function. Our theoretical analyses of the suppliers' strategic bidding behavior utilizes the Bayesian Nash equilibrium concept, which is standard in the auction literature.

Due to its incumbency status, the incumbent is already qualified for the contract; due to opaque requirements set by the buyer, we assume that both the buyer and the suppliers only know that the probability that the entrant is indeed qualified equals $0 < \beta < 1$, the entrant's *qualification probability*. For instance, β close to one corresponds to very light qualification checks that any entrant supplier is very likely to pass, while β close to zero corresponds to very strict qualification requirements that relatively few entrant suppliers would be able to pass. Qualification screening checks can be costly, involving tests of supplier products, trips to the supplier's production facilities, etc. We let $K \geq 0$ denote the *qualification cost*, that is, the cost that would be incurred by the buyer to verify whether the entrant is, or is not, qualified for the contract.

The buyer seeks to minimize her expected total procurement cost, that is, the contract price plus any supplier qualification costs. The buyer strategically chooses either "post-qualification" or "pre-qualification", describing the timing of performing qualification screening on the entrant.

Post-qualification. Under post-qualification, the buyer directly conducts a typical reverse clock auction between the entrant and incumbent, but *without attempting to qualify the entrant ahead of time*; see Ausubel and Cramton (2006) for discussion about clock auctions in practice. For simplicity, we assume the auction kicks off with a calling price p equal to R and the calling price p continuously drops as the auction progresses. The auction ends when either or both bidders drop out. Suppose the auction ends at a calling price $p = b$. If it was the entrant that dropped out first, the incumbent wins the contract and gets paid b; otherwise, the buyer performs qualification screening on the entrant, and awards the contract to the entrant with a payment $b - \frac{K}{\beta}$ if the entrant passes, but contracts with the incumbent and pays the incumbent b if the entrant fails. The buyer can reasonably implement such payment rules when she can commit to making take-or-leave-it offers to suppliers. In particular, she takes the incumbent's bid b as a signal that his cost is lower than b so she is surely informed that the incumbent will accept the payment offer if she credibly commit to pay the incumbent no more than b. By subtracting $\frac{K}{\beta}$ when computing the entrant's contract payment, the buyer accounts for the need to post-qualify the entrant. The buyer essentially runs a total-cost auction, where she computes the total cost bid from a supplier to be the supplier's price offer plus a markup to account for qualification expenses. This markup is assumed to be zero for the incumbent, capturing the fact that the incumbent is already qualified for the contract. For the entrant, the markup is equal to $\frac{K}{\beta}$ which accounts for the fact that prior to accepting an entrant's bid, the buyer would have to perform K dollars worth of qualification checks on the entrant, which the entrant would pass with probability β. Note that the buyer would be indifferent between post-qualifying a price offer of $b - \frac{K}{\beta}$ from the entrant and directly accepting a price offer of b from the incumbent, since

$$K + (b - \frac{K}{\beta})\beta + b(1 - \beta) = b.$$

The left hand side equals the expected total procurement cost if the entrant is post-qualified: Checking the entrant's qualification status costs the buyer K, the entrant passes post-qualification with probability β, and the incumbent is awarded the contract at price b if the entrant fails post-qualification. The right hand side equals the total procurement cost if the contract is directly awarded to the incumbent (who is already qualified) at price b. This markup $\frac{K}{\beta}$ can be thought of as a "switching cost" related to the need to perform costly qualification screening on the entrant. Intuitively, as the cost of qualification (K) increases or the entrant's qualification probability (β) decreases, the entrant becomes less attractive to the buyer, which is reflected by a larger markup $\frac{K}{\beta}$. In effect, the markup shifts the entrant's cost distribution to the right, making the entrant less competitive. Because the entrant's true cost x_e is distributed between zero and one, the *effective cost* (i.e., true cost plus the markup) of the entrant is distributed between $[\frac{K}{\beta}, 1 + \frac{K}{\beta}]$. Of course, additional switching costs that are unrelated to qualification — such as the need to change order processing procedures — could also be incorporated into the model by simply shifting F_e to the right. We assume $R > \frac{K}{\beta}$; otherwise, no entrant cost type would ever win a positive profit.

Pre-qualification. Under pre-qualification, the buyer pays K to screen the incumbent *before* the auction. With probability $1 - \beta$, the entrant is found to be unqualified and is discarded, and without any competitive threat to the incumbent the contract is defacto renewed with the incumbent at prevailing price R. This captures a situation in which the buyer must rely on supplier competition for price concessions. However, with probability

β the pre-qualification establishes that the entrant is qualified, at which point an auction is conducted between the entrant and the incumbent. The auction details are exactly as before, save the need to use post-qualification: Whichever bidder drops out first loses (ties are broken randomly), and the other wins and is paid the loser's dropout bid. In other words, the auction is essentially the well-known open-descending reverse English auction.

3.4. Theoretical Analyses

If the buyer has chosen pre-qualification and the entrant has passed pre-qualification, it is apparent that both suppliers have a dominant strategy to bid down to their true costs before dropping out. If the buyer has chosen post-qualification, it is apparent that the entrant must win the auction in order to possibly win the contract; therefore, the entrant's dominant strategy is to bid down to $x_e + \frac{K}{\beta}$ before dropping out. However, if the buyer chooses post-qualification, it is less apparent how exactly the incumbent will bid against a partially qualified entrant because *the incumbent can lose the auction but still win the contract* if the entrant wins the auction but fails post-qualification screening — an important feature of the auction with post-qualification screening. In such a case, the incumbent is paid his drop out bid. Thus, when $\beta < 1$ the incumbent has an incentive to hold back on bidding by dropping out of the auction before reaching his true cost x_i, in order to preserve his profit margin. In contrast, if the entrant were already fully qualified ($\beta = 1$), the incumbent would have a dominant strategy to bid down to its true cost x_i before dropping out; this is the case that is assumed in classical auction theory (all bidders fully qualified, hence the auction winner is automatically the contract winner). In the following §3.4.1, we analyze the incumbent's Bayesian Nash equilibrium bidding behavior in the auction with post-qualification, addressing *research questions 1 and 2*.

In comparing the post- and pre-qualification strategies, the buyer faces the following tradeoffs. If she uses post-qualification, she avoids wasting money qualifying an entrant whose price in the auction might not turn out to be competitive, and also avoids losing the opportunity to run an auction in case the entrant is actually unqualified. However, post-qualification causes the incumbent to hold back on bidding in the auction. Pre-qualification can induce more aggressive bidding by the incumbent, but pre-qualification backfires on the buyer if the entrant fails pre-qualification and must be discarded. Thus, the buyer's decision depends on how the incumbent will bid in an auction with post-qualification. If the buyer thinks the incumbent will bid very aggressively even if the entrant might be unqualified, post-qualification can be an attractive strategy. On the other hand, if the incumbent will only bid aggressively if the buyer can tout the fact that the entrant is fully qualified and only the low bid will win the contract, the buyer may be forced to use pre-qualification. In the following §3.4.2, we characterize theoretical predictions for the buyer's optimal strategy, addressing *research question 3*.

3.4.1 Incumbent Bidding Strategy Under Post-Qualification

Proposition 6 *For an incumbent with cost x_i, there exists a static optimal bid-down-to level $\underline{p}(x_i) \in [\frac{K}{\beta}, R]$ such that the incumbent should stay in the auction if the auction price is above $\underline{p}(x_i)$ and drop out of the auction at the price $\underline{p}(x_i)$ if the entrant has not dropped out. There exist two thresholds x_B and x_W such that*

- $x_W \leq x_B$, $x_B < R$, *and* $x_W \leq \frac{K}{\beta}$,
- $\underline{p}(x_i) = R$ *if and only if* $x_i \geq x_B$,
- $\underline{p}(x_i) = \frac{K}{\beta}$ *if and only if* $x_i \leq x_W$,
- $x_i < \underline{p}(x_i)$, $\frac{K}{\beta} < \underline{p}(x_i) < R$, *and* $\underline{p}(x_i)$ *strictly increases in* x_i *if* $x_W < x_i < x_B$.

In words, Proposition 6 says that the incumbent uses three types of strategies, depending on its cost x_i: boycott the auction; participate in the auction until price drops too low, then drop out; participate in the auction until the entrant drops out. The incumbent uses the first strategy when his cost is so high (i.e., $x_i \geq x_W$) that he is unlikely to beat the entrant on price alone. In such cases, the incumbent prefers to retain his profit margin by boycotting the auction (i.e., $\underline{p}(x_i) = R$) and simply hopes the entrant is disqualified. On the other hand, the incumbent uses the third strategy when he is absolutely certain that he can defeat the entrant on price alone. To see this, note that this strategy is only used when the incumbent's cost is sufficiently small, that is, $x_i \leq x_W$; furthermore, we prove that the threshold x_W is always less than or equal to $\frac{K}{\beta}$, which in turn is no greater than the entrant's total cost $x_e + \frac{K}{\beta}$. In such cases, the incumbent prefers to win the auction to avoid any chance of losing the contract (i.e., $\underline{p}(x_i) = \frac{K}{\beta}$). The second strategy is a mixture of these approaches: when the incumbent thinks he stands a reasonable chance of defeating the entrant on price alone (i.e., $x_W < x_i < x_B$), he prefers to "test the waters" by staying in the auction, at least initially. He hopes the entrant will drop out quickly, but if the auction price gets too low the incumbent abandons the effort and drops out of the auction before the price dropping to his true cost (i.e., $\underline{p}(x_i) > x_i$). Rather than seeking to beat the entrant on price, he changes tactics and instead hopes the entrant will be disqualified.

Proposition 7 *The optimal bid-down-to level $\underline{p}(x_i)$ is increasing in R for all x_i; however, $\underline{p}(x_i)$ is generally not monotone in K or β.*

Intuitively, a higher current contract price R implies a larger potential profit margin available to the incumbent, making him more willing to attempt to retain such profit margin by dropping out the auction early or even boycotting the auction. However, the qualification cost K and the qualification probability β in general do not monotonically affect the incumbent's bidding. Technically, this is because the incumbent's marginal benefit from a lower bid is not monotone in K or β. More intuitively, an increase of K or β can yield two competing effects on the incumbent's bidding tradeoffs. To see this, as K increases, on one hand, the incumbent is encouraged to use a lower bid (so as to win the contract by winning the auction directly) because the entrant is less competitive in cost (due to a larger markup $\frac{K}{\beta}$) and hence the incumbent finds it easier to beat the entrant on price alone; however, on the other hand, as K increases, a low-cost incumbent who uses a bid-to-win strategy will effectively drop out at a higher bid $\frac{K}{\beta}$. Analogously, as β increases, on one hand, the incumbent is encouraged to try and win the auction because the entrant is more likely to survive post-qualification screening (which means that it is less attractive for the incumbent to give up winning the auction and bet on the entrant's failing post-qualification screening); however, on the other hand, the incumbent is discouraged to try and win the auction because the entrant is more competitive in cost (due to a lower markup $\frac{K}{\beta}$) and hence the incumbent finds it more difficult to beat the entrant on price alone but more attractive to simply let the entrant win the auction rather than risk his profit margin by bidding low.

Special case 1: risk-neutral incumbent

To illustrate Propositions 6-7, in the following we consider a risk-neutral incumbent, i.e., $U(p - x_i) = p - x_i$, and assume the entrant's cost distribution is uniformly distributed, that is, $F_e \sim U[0, 1]$. For given K, β, and R, let $\hat{x}(K, \beta, R)$ be the $x_i \in (-\infty, \frac{2\beta-1}{\beta} + \frac{K}{\beta})$ solving the following equation:

$$0 = \begin{cases} \frac{(\beta x_i - K)^2}{2(2\beta-1)} - \beta x_i + \frac{K}{\beta} + \frac{1}{2} - (1-\beta)R & \text{, if } \frac{K}{\beta} < x_i \leq \frac{2\beta-1}{\beta} + \frac{K}{\beta}; \\ -\beta x_i + \frac{K}{\beta} + \frac{1}{2} - (1-\beta)R & \text{, if } x_i \leq \frac{K}{\beta}. \end{cases}$$

Note that $\hat{x}(K, \beta, R)$ is well defined because the right hand side of the equation is continuous and convex on $(-\infty, \frac{2\beta-1}{\beta} + \frac{K}{\beta}]$, is negative at $x_i = \frac{2\beta-1}{\beta} + \frac{K}{\beta}$ and goes to positive infinity as x_i goes to negative infinity.

Proposition 8 *Assume $F_e \sim U[0, 1]$. For a risk-neutral incumbent, the optimal bid-down-to level $\underline{p}(x_i) = \frac{\beta x_i}{2\beta-1} - \frac{(1-\beta)K}{\beta(2\beta-1)}$ when $x_W < x_i < x_B$. The thresholds x_B and x_W, depending on R, K, and β, are given in the following table.*

	$0 < \beta \leq \frac{1}{2}$		$\frac{1}{2} < \beta < 1$	
	$R \leq 1 + \frac{K}{\beta}$	$R > 1 + \frac{K}{\beta}$	$R \leq 1 + \frac{K}{\beta}$	$R > 1 + \frac{K}{\beta}$
x_B	$\frac{K}{2\beta^2} - \frac{R}{2\beta} + R$	$\frac{K}{\beta^2} + \frac{1}{2\beta} - \frac{R}{\beta} + R$	$\frac{2\beta-1}{\beta}R + \frac{(1-\beta)K}{\beta^2}$	$\hat{x}(K, \beta, R)$
x_W	$\frac{K}{2\beta^2} - \frac{R}{2\beta} + R$	$\frac{K}{\beta^2} + \frac{1}{2\beta} - \frac{R}{\beta} + R$	$\frac{K}{\beta}$	$\min\{\hat{x}(K, \beta, R), \frac{K}{\beta}\}$

Proposition 8 illustrates Proposition 6 by presenting the exact expressions of the optimal "test-the-water" level and the thresholds x_W and x_B. Demonstrations of the incumbent's

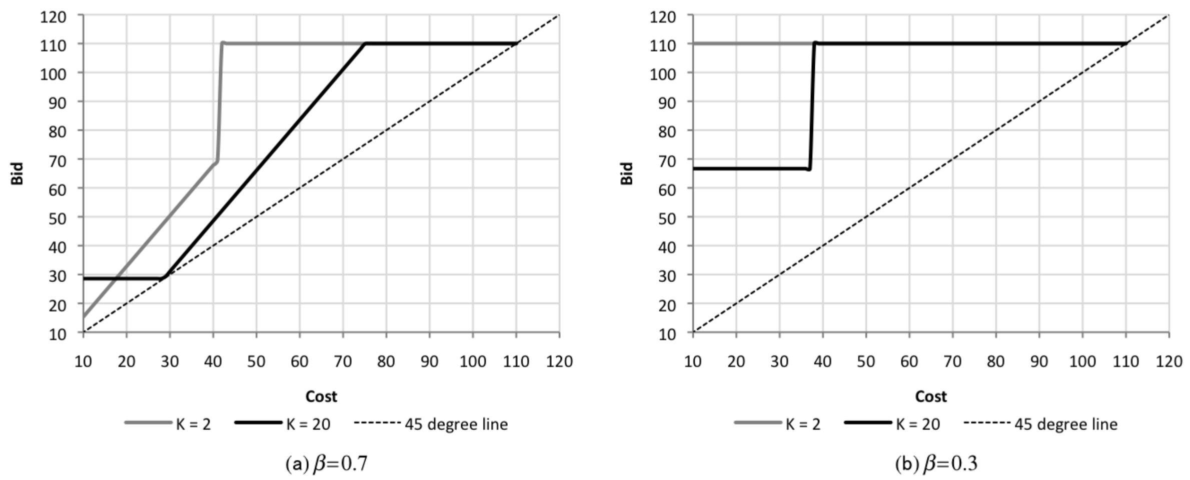

Figure 3.1: A risk-neutral incumbent's bidding functions when $K = 2$ and $K = 20$. Assume $F_i \sim U[0.1, 1.1]$, $F_e \sim U[0, 1]$, and $\beta = 0.7$ (left panel) or $\beta = 0.3$ (right panel).

optimal bidding functions are provided in Figure 3.1. It is easy to see that the optimal bid-down-to level $\underline{p}(x_i)$ is never lower than and increases in the incumbent's cost x_i. It is interesting to notice that the "test-the-water" strategy is employed only when $\frac{1}{2} < \beta < 1$. The intuition is that, as β decreases and approaches $\frac{1}{2}$ from above, more and more incumbent types with medium cost find the entrant less credible a threat and hence would choose boycotting over testing-the-water.

The exact expressions of the bidding functions provided by Proposition 8 enable us to further explore the effect of the qualification cost K and the entrant's qualification probability β on the aggressiveness of the incumbent's bidding behavior. An incumbent can win a contract in two different ways: i) win the auction outright, which happens when $\underline{p}(x_i) \leq x_e + \frac{K}{\beta}$; ii) lose the auction but win the contract because the entrant fails post-qualification, which happens only when $x_e + \frac{K}{\beta} \leq \underline{p}(x_i)$. Only in case (i) does the incumbent's aggressive bidding actually drive the entrant to drop out and lose the auction. Thus, in exploring how is the aggressiveness of the incumbent's bidding affected by the qualification cost K and the entrant's qualification probability β, we choose to examine how K and β affect *the probability that the incumbent wins the auction*, which is defined to be the probability that

$\underline{p}(x_i) \leq x_e + \frac{K}{\beta}$, corresponding to the case i) discussed as above.

Proposition 9 *Assume $F_e \sim U[0,1], F_i \sim U[l,R]$ and the incumbent is risk-neutral. The probability that the incumbent wins the auction increases in K; it increases in β if K is close to zero, and it decreases in β if K is large and close to βR.*

Proposition 9 shows that, although the incumbent's optimal bid-down-to level is not monotone in K, the probability that the incumbent wins the auction outright increases in K. In other words, on average, the incumbent bids more aggressively when K is larger. Intuitively, this is because the entrant's effective cost is larger. However, the effect of β on the probability that the incumbent wins the auction depends on the value of K. When K is small, the effect that the entrant's cost is less advantaged with larger β is dominated; therefore, the incumbent is encouraged to bid more aggressively, because the entrant is more likely to be truly qualified and hence it is less likely to win the contract if the entrant wins the auction. However, when K is large, the effect that the entrant's cost is less advantaged with larger β dominates; therefore, the incumbent is encouraged to bid less aggressively if β is larger (and hence the entrant's effective cost is lower).

Special case 2: incumbent with constant absolute risk aversion

Assume the incumbent has a constant absolute risk aversion (CARA) with a utility function $U(p - x_i) = -\frac{1}{\gamma}e^{-\gamma(p-x_i)}$. To facilitate characterization of the incumbent's bidding strategy, we introduce some notation. For given $\gamma > 0, K \geq 0$, $0 < \beta < 1$, and $y \geq \frac{K}{\beta}$, define $\overline{x}(y) \triangleq y - \frac{1}{\gamma}\ln[1 + \gamma\frac{(1-\beta)}{\beta}(y - \frac{K}{\beta})]$ and $\underline{x}(y) \triangleq y - \frac{1}{\gamma}\ln[\frac{e^{\gamma(y-\frac{K}{\beta})}-1-\gamma(1-\beta)(y-\frac{K}{\beta})}{\gamma(y\beta-K)}]$. Consider a term $e^{\gamma(y-x_i)} - 1 - \gamma(y - \frac{K}{\beta})\frac{1-\beta}{\beta}$, which is convex in y and approaches $+\infty$ as y approaches either $-\infty$ or $+\infty$. This term, when $x_i > \frac{1-2\beta}{\gamma(1-\beta)} - \frac{1}{\gamma}\ln(\frac{1-\beta}{\beta}) + \frac{K}{\beta}$, has two solutions on $y \in (-\infty, +\infty)$, of which we let $\overline{p}(x_i)$ denote the larger one. Let $\tilde{x}$ be the x_i solving $x_i = \underline{x}(\overline{p}(x_i))$.

Proposition 10 *Assume $F_e \sim U[0,1]$ and $U(p - x_i) = -\frac{1}{\gamma}e^{-\gamma(p-x_i)}$. The optimal bid-down-to level $\underline{p}(x_i) = \overline{p}(x_i)$ when $x_W < x_i < x_B$. When $1 \leq R \leq 1 + \frac{K}{\beta}$, the thresholds x_B and x_W,*

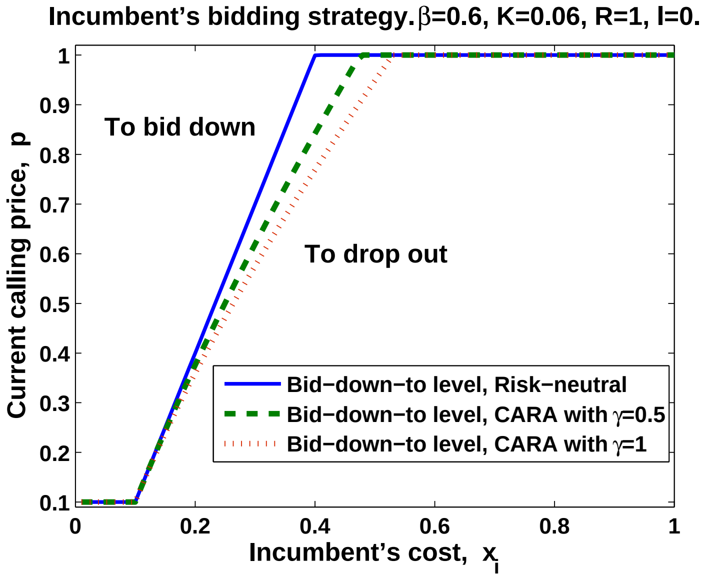

Figure 3.2: Incumbent's optimal bidding strategy given his cost x_i, the current calling price p, and his risk preference. $K = 0.06$, $\beta = 0.6$, $l = 0$, $R = 1$.

depending on R, K, and β, are given in the following table.

	$0 < \beta < \frac{1}{2}$		$\frac{1}{2} \leq \beta < 1$
	$R \leq 1 + \frac{K}{\beta}$; $\overline{x}(R) \leq \underline{x}(R)$	$R \leq 1 + \frac{K}{\beta}$; $\overline{x}(R) > \underline{x}(R)$	$R \leq 1 + \frac{K}{\beta}$
$x_B(R)$	$\underline{x}(R)$	$\overline{x}(R)$	$\overline{x}(R)$
$x_W(R)$	$\underline{x}(R)$	$\tilde{x}$	$\frac{K}{\beta}$

When $R > 1+\frac{K}{\beta}$, *the thresholds* $x_B(R) = \min\{x_B(1+\frac{K}{\beta}), \hat{x}(R)\}$ *and* $x_W(R) = \min\{x_W(1+\frac{K}{\beta}), \hat{x}(R)\}$, *where* $\hat{x}(R)$ *is the unique* x_i *such that* $\int_{\hat{p}(x_i)}^{R} \{\beta[U(y-x_i) - U(0)]f_e(y - \frac{K}{\beta}) - (1-\beta)F_e(y - \frac{K}{\beta})\frac{d\ U(y-x_i)}{dy}\}dy = 0$, *and* $\hat{p}(x_i)$ *is the bid-down-to level when* $R = 1 + \frac{K}{\beta}$.

3.4.2 Buyer's Optimal Qualification Screening Strategy

As discussed in §3.3, under the pre-qualification strategy, the buyer spends qualification cost K on qualifying the entrant, which yields one of two outcomes: with probability β the entrant is found to be qualified and a typical reverse English auction is subsequently run with two fully qualified bidders thus resulting in an expected contract payment of $E\max\{x_i, x_e\}$; and with probability $1-\beta$ the entrant is found to be unqualified and is discarded, and consequently the contract is renewed with the incumbent at price R. In summary, under pre-qualification the buyer's expected total (payment plus qualification) cost is

$$\beta E \max\{x_i, x_e\} + (1-\beta)R + K. \tag{3.1}$$

Per our analysis in §3.4.1, if the buyer uses the post-qualification strategy, an incumbent's bidding strategy can be described as a bid-down-to level $\underline{p}(x_i)$. The entrant wins the auction if $\underline{p}(x_i) > \min\{x_e + \frac{K}{\beta}, R\}$; if so, the buyer incurs a qualification cost K to vet the entrant and pays $\underline{p}(x_i) - \frac{K}{\beta}$ to the entrant if the entrant survives post-qualification (which happens with probability β), but pays $\underline{p}(x_i)$ to the incumbent if the entrant fails post-qualification (which happens with probability $1-\beta$). Otherwise, if $\underline{p}(x_i) \leq \min\{x_e + \frac{K}{\beta}, R\}$, the incumbent wins the auction, and thus keeps the contract with a payment from the buyer equal to either the entrant's dropout bid or the reserve price (whichever is smaller), $\min\{x_e + \frac{K}{\beta}, R\}$. Therefore, under the post-qualification strategy, the buyer's expected total cost is

$$\begin{aligned} & E \max\{\min\{x_e + \frac{K}{\beta}, R\}, K + \beta[\underline{p}(x_i) - \frac{K}{\beta}] + (1-\beta)\underline{p}(x_i)\}, \\ = \quad & E \max\{\min\{x_e + \frac{K}{\beta}, R\}, \underline{p}(x_i)\}. \end{aligned} \tag{3.2}$$

The buyer finds the optimal qualification strategy by comparing (3.1) with (3.2). The following proposition says that the buyer prefers post-qualification screening if the qualification cost K is large enough.

Proposition 11 *Given any F_i, F_e, $U(\cdot)$, $\beta \in (0,1)$, l and R, there exists a threshold $\underline{K}$ such that it is optimal for the buyer to choose post-qualification if $K > \underline{K}$.*

The intuition to Proposition 11 is that when the qualification cost is high enough post-qualification screening is preferred because it helps the buyer avoid wasting money qualifying an entrant whose price in the auction might not turn out to be competitive. Although Proposition 11 shows that the buyer's optimal strategy can be characterized by a threshold

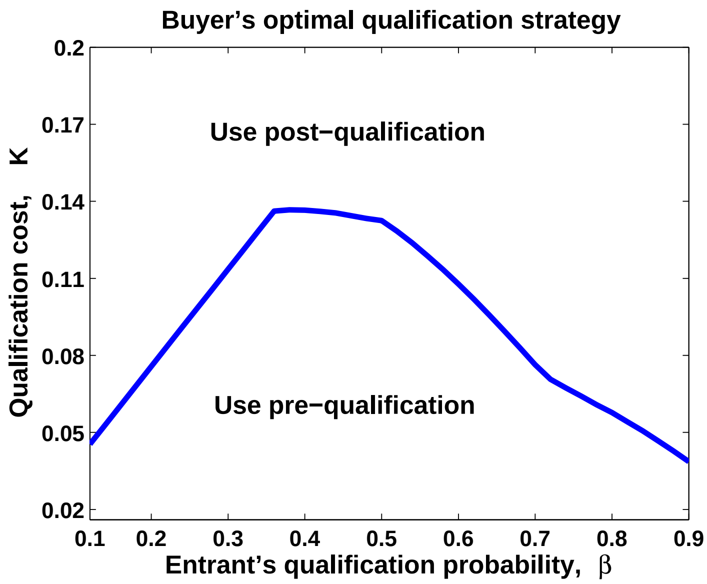

Figure 3.3: Buyer's optimal qualification strategy given the qualification cost K, the entrant's qualification probability β. Plot assumes $F_e \sim U[0,1]$, $F_i \sim [0.1, 1.1]$, and that the incumbent is risk-neutral.

of qualification cost K, there in general does not exist a similar threshold of qualification probability β; for example, a buyer can prefer post-qualification either when β is small enough or large enough, as we will show in the following Figure 3.3. This is because the incumbent's bidding behavior under post-qualification is in general not monotone in β (Proposition 7) and hence the buyer's expected total cost under post-qualification is in general not monotone in β (per (3.2)), although the buyer's expected total cost under pre-qualification is monotone in β (per (3.1)).

To illustrate Proposition 11, we assume $l = 0.1$, $R = 1.1$, $F_e \sim U[0,1]$, $F_i \sim U[0.1, 1.1]$, and that the incumbent is risk-neutral. Figure 3.3 characterizes the buyer's optimal qualification strategy given various qualification costs K and entrant qualification probabilities β. In the figure, the hill-shaped line corresponds to the cases where the buyer is indifferent between the pre-qualification and post-qualification strategies and it divides the plane into upper and lower parts, such that it is optimal for the buyer to choose the post-qualification strategy in the upper part, but to choose the pre-qualification strategy in the lower part. Figure 3.3, on one hand, illustrates Proposition 11: The buyer prefers pre-qualification only when the qualification cost K is small (making pre-qualification cheap); on the other hand, it shows the non-monotone effect of qualification probability β: The buyer prefers pre-qualification when β is not too small (in which case pre-qualification is very likely to disqualify the bidder), nor too large (in which case the incumbent is inclined to bid aggressively even under post-qualification because he knows the entrant would stand only a small chance of failing post-qualification).

3.5. Conclusions

We consider a buyer who uses a procurement auction to structure contract negotiation between an incumbent and an entrant supplier, seeking either price concessions from the incumbent or a low-price contract from the entrant. To make sure that the entrant is a truly qualified supplier for the contract, the buyer needs to perform qualification screening process on the entrant – a process that is costly and time-consuming for the buyer. The buyer is faced a strategic decision on the timing of screening the entrant, choosing either screening the entrant before holding the auction, so-called pre-qualification screening, or delaying the qualification screening until after the auction only if the entrant bids wins the auction, so-called post-qualification screening. The buyer has to incur the qualification screening cost upfront and may be backfired if the entrant fails the screening process — in such case, the buyer has to discard the entrant and forego the opportunity to run the auction for a low contract price. Whereas the buyer can avoid wasting qualification screening cost and losing the opportunity to run an auction by delaying the entrant's screening until after the auction (i.e., by using post-qualification), the buyer needs to take into account the incumbent's strategic hold-back in bidding against a partially qualified entrant because in such case the incumbent can possibly win the contract even if loses the auction — this happens if the entrant fails the post-qualification process and the buyer comes back to the incumbent for contract renewal.

We analytically examine the incumbent's bidding strategy against an entrant who is partially qualified, i.e., who can survive the qualification screening process with probability less than one. We find that the incumbent may use three types of bidding strategies, depending on his cost. On one extreme, when the incumbent's cost is very high and hence it is very unlikely to beat the entrant on price alone, the incumbent finds it optimal to drop out of the auction as soon as the auction starts, i.e., to effectively boycott the auction so as to retain a profit margin as large as possible and simply hope the entrant would fail the post-qualification screening. On the other extreme, when the incumbent's cost is very low, the buyer has a large profit margin to retain and possess a for-sure winning opportunity if he will to take. In such cases, he prefers to guarantee winning the contract to directly win the auction, and hence finds it optimal to stay in the auction until the entrant drops out. Between the two extremes are the cases where the incumbent has medium cost and reasonable chance to beat the entrant on price. In such cases, the incumbent finds it optimal to test the water: He stays in the auction for a while, seeing if he can directly win the auction (i.e., the entrant drops out the auction), but he drops out before the auction price gets too low so as to retain a reasonable profit margin. The incumbent's bidding behavior is also affected by the auction reserve price, the buyer's qualification screening cost, and the probability that the entrant can survive the qualification screening process. In general, the incumbent bids less aggressively with a higher reserve price, but may bid less or more aggressively with higher screening cost or qualification probability.

With characterizations of the incumbent's bidding strategy under post-qualification screening, we examine the buyer's optimal qualification screening strategy. We find that the buyer prefers pre-qualification only when the qualification cost is small because the buyer is pushed away from pre-qualification by high qualification screening in avoidance of wasting money on an unqualified entrant. On the other hand, we find the non-monotone effect of the qualification probability: The buyer prefers pre-qualification when the probability is not too small nor too large. Intuitively, a small qualification probability implies that had the buyer chosen pre-qualification screening, it would be very likely that the entrant is unqualified and she would have to waste the qualification cost and the auction opportunity; a large qualification probability implies that the incumbent will be inclined to bid aggressively even under post-qualification because he knows the entrant would stand only a small chance of failing post-qualification.

3.6. Proofs

3.6.1 Proof of Proposition 6

Existence of static bid-down-to level

The optimal strategy of the incumbent is characterized by a static bid-down-to level because we can show that dynamically choosing between drop-out and stay-in as time goes forward yields the same outcome as if the buyer has chosen a fixed bid-down-to level at the outset of the auction.

Let $Y(p)$ be the maximum utility the incumbent can expect to gain by optimally choosing his strategy given the current calling price p. At any calling price p, the incumbent can choose to drop out or stay in. If drops out at p, he obtains an expected utility given by

$$(1-\beta)U(p-x_i)+\beta U(0); \tag{3.3}$$

however, if he chooses to remain in the auction the incumbent's maximum expected utility is given by

$$\sup_{\max\{x_i,\frac{K}{\beta}\}\le t<p} \int_t^p U(y-x_i)g(y|p)dy + G(t|p)Y(t), \tag{3.4}$$

where $g(y|p)=\frac{f_e(y-\frac{K}{\beta})}{F_e(p-\frac{K}{\beta})}$ is the density of conditional probability that the entrant drops out at y given that the entrant has not dropped out by p and $G(t|p)=[1-\int_t^p g(y|p)dy]=\frac{F_e(y-\frac{K}{\beta})}{F_e(p-\frac{K}{\beta})}$. Therefore, we have

$$Y(p)=\max\{(3.3),(3.4)\}.$$

In words, the above states that the incumbent's optimal bidding strategy (drop-out decision) can be characterized as a stopping problem over the (continuous) state p, the calling price in the auction. At state p, the incumbent makes the following tradeoff calculation. He first calculates his expected utility of dropping out of the auction (equation (3.3)), which is a function of the calling price, his true cost x_i, and the probability that the entrant would fail post-qualification, $1-\beta$. Second, he calculates the expected utility of remaining in the auction, in the hopes of beating the entrant and winning the contract outright (equation (3.4). This expected utility depends on the price at which the incumbent predicts the entrant will drop out (the variable y in equation (3.4)), which in turn depends on g, itself a function of the entrant's cost distribution (F_e), markup ($\frac{K}{\beta}$), and the current auction calling price. If the expected utility of dropping out of the auction is greater than that of staying in, the incumbent drops out.

In the following we show that (3.4) equals

$$\sup_{\max\{x_i,\frac{K}{\beta}\}\leq t<p} \int_t^p U(y-x_i)g(y|p)dy + G(t|p)[(1-\beta)U(t-x_i)+\beta U(0)]. \tag{3.5}$$

Let $V_1(p) = \sup_{\max\{x_i,\frac{K}{\beta}\}\leq t<p} v_1(p,t)$ and $V_2(p) = \sup_{\max\{x_i,\frac{K}{\beta}\}\leq t<p} v_2(p,t)$, where

$$\begin{aligned} v_1(p,t) &= \int_t^p U(y-x_i)g(y|p)dy + G(t|p)Y(t), \\ Y(t) &= \max\{(1-\beta)U(t-x_i)+\beta U(0), V_1(t)\} \end{aligned}$$

and

$$v_2(p,t) = \int_t^p U(y-x_i)g(y|p)dy + G(t|p)[(1-\beta)U(t-x_i)+\beta U(0)].$$

Given that $Y(t) \geq [(1-\beta)U(t-x_i)+\beta U(0)]$ by definition, we have $V_1(p) \geq V_2(p)$. Thus, it suffices to show $V_1(p) \leq V_2(p)$.

- If $V_1(p) \leq [(1-\beta)U(p-x_i)+\beta U(0)]$, then $V_1(p) \leq V_2(p)$ because $[(1-\beta)U(p-x_i)+\beta U(0)] = v_2(p,p) \leq V_2(p)$.

- If $V_1(p) > [(1-\beta)U(p-x_i)+\beta U(0)]$, then there exists $t_{\inf}(p) \in [\max\{x_i, \frac{K}{\beta}\}, p)$ such that

$$t_{\inf}(p) = \inf \underset{\max\{x_i, \frac{K}{\beta}\} \leq t < p}{\arg\max} v_1(p,t).$$

We establish $V_1(p) \leq V_2(p)$ by showing $V_2(p) \geq v_2(p, t_{\inf}(p)) = v_1(p, t_{\inf}(p)) = V_1(p)$, where the inequality is by the definition of $V_2(p)$, the last equality is by the definition of $t_{\inf}(p)$, and the first equality is true because $Y(t_{\inf}(p)) = [(1-\beta)U(t_{\inf}(p) - x_i) + \beta U(0)]$. That $Y(t_{\inf}(p)) = [(1-\beta)U(t_{\inf}(p) - x_i) + \beta U(0)]$ is proved by contradiction. Suppose otherwise, i.e., $Y(t_{\inf}(p)) > [(1-\beta)U(t_{\inf}(p) - x_i) + \beta U(0)]$. This implies $\max_{\max\{x_i, \frac{K}{\beta}\} \leq t < t_{\inf}(p)} v_1(t_{\inf}(p), t)$ attains at some $\hat{t} \in \max\{x_i, \frac{K}{\beta}\} \leq t < t_{\inf}(p)$, i.e.,

$$Y(t_{\inf}(p)) = v_1(t_{\inf}(p), \hat{t}). \tag{3.6}$$

Note that

$$\begin{aligned}
\max_{\max\{x_i, \frac{K}{\beta}\} \leq t < p} v_1(p,t) &= v_1(p, t_{\inf}(p)) \\
&= \int_{t_{\inf}(p)}^{p} U(y - x_i) g(y|p) dy + G(t_{\inf}(p)|p) Y(t_{\inf}(p)) \\
&= \int_{t_{\inf}(p)}^{p} U(y - x_i) g(y|p) dy + G(t_{\inf}(p)|p) v_1(t_{\inf}(p), \hat{t}) \\
&= \int_{\hat{t}}^{p} U(y - x_i) g(y|p) dy + G(\hat{t}|p) Y(\hat{t}) = v_1(p, \hat{t}),
\end{aligned}$$

where the first equality is given by the definition of $t_{\inf}(p)$ and $t_{\inf}(p) < p$, the second equality is by the definition of $v_1(p, t_{\inf}(p))$, the third equality is due to equation (3.6), and the last two equalities follow algebra. However, that $v_1(p, \hat{t}) = v_1(p, t_{\inf}(p))$ and $\hat{t} < t_{\inf}(p)$ contradicts that $t_{\inf}(p)$ is defined as the infimum.

The fact that (3.4) equals (3.5) implies that the optimal dynamic bidding strategy yields the same payoff to the bidder as using the static bid-down-to level.

Optimal static bid-down-to level

When the incumbent chooses a fixed bid-down-to level at the outset of the auction, the existence of the optimal bid-down-to level is guaranteed because the the incumbent optimizes a continuous objective function over a compact set. In particular, the static optimal bid-down-to level $\underline{p}(x_i)$ is such that $\underline{p}(x_i) = \underset{t\in[\max\{x_i,\frac{K}{\beta}\},R]}{\arg\max} \Pi(t)$, where

$$\begin{aligned}\Pi(t) \triangleq \quad & \int_t^R U(y-x_i)f_e(y-\frac{K}{\beta})dy \\ & +F_e(t-\frac{K}{\beta})[(1-\beta)U(t-x_i)+\beta U(0)] + U(R-x_i)[1-F_e(R-\frac{K}{\beta})].\end{aligned} \tag{3.7}$$

The decision set is $t \in [\max\{x_i, \frac{K}{\beta}\}, R]$ because the auction price can never be beyond $[\frac{K}{\beta}, R]$ and bidding below the true cost x_i can never be profitable for the incumbent. The incumbent's expected utility as a function of the chosen bid-down-to level t, $\Pi(t)$, has the first term $\int_t^R U(y-x_i)f_e(y-\frac{K}{\beta})dy$ corresponding to the cases in which the incumbent wins the auction outright because the entrant drops out at $y \in (t, R)$, the second term $F_e(t-\frac{K}{\beta})[(1-\beta)U(t-x_i)+\beta U(0)]$ corresponding to the cases in which the entrant wins the auction, and the last term $U(R-x_i)[1-F_e(R-\frac{K}{\beta})]$ corresponding to the cases in which the entrant loses because his effective cost is above the reserve price.

Existence of thresholds x_W and x_B

Note that

$$\begin{aligned}\frac{d\Pi(t)}{dt} &= F_e(t-\frac{K}{\beta})(1-\beta)U'(t-x_i) - \beta f_e(t-\frac{K}{\beta})[U(t-x_i)-U(0)] \\ &= F_e(t-\frac{K}{\beta})(1-\beta)U'(t-x_i)\left[1-\frac{\beta}{1-\beta}\frac{f_e(t-\frac{K}{\beta})}{F_e(t-\frac{K}{\beta})}\frac{U(t-x_i)-U(0)}{U'(t-x_i)}\right],\end{aligned} \tag{3.8}$$

which strictly increases in x_i (i.e., we have $\frac{\partial^2\Pi(t)}{\partial x_i \partial t} > 0$) because $U'(t-x_i)$ increases in x_i

(i.e., because $U(\cdot)$ is concave) and $[U(t-x_i)-U(0)]$ strictly decreases in x_i. The fact that $\frac{d\Pi(t)}{dt}$ strictly increases in x_i implies that $\underline{p}(x_i)$ increases in x_i and moreover strictly increases when $\max\{x_i, \frac{K}{\beta}\} < \underline{p}(x_i) < R$. To see this, suppose $\underline{p}(x_i)$ does not increase in x_i, then it must exist $x_i^{(1)} < x_i^{(2)}$ such that $\underline{p}(x_i^{(1)}) > \underline{p}(x_i^{(2)})$. With a bit abuse of notation, let $\Pi(t; x_i)$ denote the incumbent's expected utility if the incumbent's cost is x_i and it chooses t as the bid-down-to level. On one hand, by definition of $\underline{p}(x_i^{(1)})$ and $\underline{p}(x_i^{(2)})$, we have

$$\Pi(\underline{p}(x_i^{(1)}); x_i^{(1)}) - \Pi(\underline{p}(x_i^{(2)}); x_i^{(1)}) \geq 0, \text{ and } \Pi(\underline{p}(x_i^{(2)}); x_i^{(2)}) - \Pi(\underline{p}(x_i^{(1)}); x_i^{(2)}) \geq 0. \tag{3.9}$$

On the other hand, we notice that

$$\begin{aligned} &\Pi(\underline{p}(x_i^{(1)}); x_i^{(1)}) - \Pi(\underline{p}(x_i^{(2)}); x_i^{(1)}) + \Pi(\underline{p}(x_i^{(2)}); x_i^{(2)}) - \Pi(\underline{p}(x_i^{(1)}); x_i^{(2)}) \\ &= \int_{\underline{p}(x_i^{(2)})}^{\underline{p}(x_i^{(1)})} -\frac{d\Pi(t; x_i^{(1)})}{dt} dt + \int_{\underline{p}(x_i^{(2)})}^{\underline{p}(x_i^{(1)})} \frac{d\Pi(t; x_i^{(2)})}{dt} dt > 0 \end{aligned} \tag{3.10}$$

where the inequality holds because $\underline{p}(x_i^{(1)}) > \underline{p}(x_i^{(2)})$ and $\frac{d\Pi(t;x_i^{(2)})}{dt} > \frac{d\Pi(t;x_i^{(1)})}{dt}$. However, (3.9) contradicts with (3.10), which implies that $\underline{p}(x_i)$ must increase in x_i. Moreover, if $\max\{x_i^{(1)}, \frac{K}{\beta}\} < \underline{p}(x_i^{(1)}) < R$, then it must be $\frac{d\Pi(t;x_i^{(1)})}{dt}|_{t=\underline{p}(x_i^{(1)})} = 0$; this implies that $\underline{p}(x_i^{(1)}) < \underline{p}(x_i^{(2)})$ — Otherwise, if $\underline{p}(x_i^{(1)}) = \underline{p}(x_i^{(2)})$, we have $\frac{d\Pi(t;x_i^{(2)})}{dt}|_{t=\underline{p}(x_i^{(2)})} > \frac{d\Pi(t;x_i^{(1)})}{dt}|_{t=\underline{p}(x_i^{(2)})} = \frac{d\Pi(t;x_i^{(1)})}{dt}|_{t=\underline{p}(x_i^{(1)})} = 0$, which contracts to the optimality of $\underline{p}(x_i^{(2)})$.

Given any $x_i \in (\frac{K}{\beta}, R)$, we have $\frac{d\Pi(t)}{dt}|_{t\downarrow x_i} = F_e(x_i - \frac{K}{\beta})(1-\beta)U'(0) > 0$, which has two implications. First, $\underline{p}(x_i) > x_i$ for all $x_i \in (\frac{K}{\beta}, R)$. Second, $\frac{d\Pi(t)}{dt} > 0$ for all $t \in [x_i, R]$ if $x_i < R$ is close enough to R, which in turn implies that $\underline{p}(x_i) = R$ if $x_i < R$ is close enough to R. The fact that $\underline{p}(x_i)$ is increasing implies that there exists a unique threshold $x_B < R$ such that $\underline{p}(x_i) = R$ if and only if $x_i \geq x_B$.

If $U(\cdot)$ is such that $\frac{U'(t)}{U(t)-U(0)}$ goes to zero as t goes to infinity, then per (3.8) we have that $\frac{d\Pi(t)}{dt}$ is negative for all $t \in [\frac{K}{\beta}, R]$ when x_i is small enough. In other words, there exists a largest threshold x_W such that $\overline{p}(x_i) = \max\{x_i, \frac{K}{\beta}\}$ for all $x_i \leq x_W$. The fact that $\underline{p}(x_i) > x_i$

for all $x_i \in (\frac{K}{\beta}, R)$ implies that $x_W \leq \frac{K}{\beta}$ and hence $\overline{p}(x_i) = \frac{K}{\beta}$ for all $x_i \leq x_W$. The fact that $\underline{p}(x_i)$ is increasing implies $x_W \leq x_R < R$ and that $\overline{p}(x_i) = \frac{K}{\beta}$ if and only if $x_i \leq x_W$.

When $x_W < x_i < x_B$, we have $\frac{K}{\beta} < \underline{p}(x_i) < R$, which is implied by the existence and the uniqueness of the thresholds x_W and x_B, and we have $\underline{p}(x_i) > x_i$, which holds for $x_i > \frac{K}{\beta}$ because we have proved that $\underline{p}(x_i) > x_i$ for all $x_i \in (\frac{K}{\beta}, R)$ and holds for $x_W < x_i \leq \frac{K}{\beta}$ (when $x_W < \frac{K}{\beta}$) because $\underline{p}(x_i) > \frac{K}{\beta} \geq x_i$. Finally, that $\underline{p}(x_i)$ strictly increases when $x_W < x_i < x_B$ is implied by $\frac{d^2\Pi(t)}{dx_i dt} > 0$, which was proved earlier on.

3.6.2 Proof of Proposition 7

Consider any $R^{(1)} < R^{(2)}$ and any $t^{(1)} < t^{(2)} < R^{(1)}$. For a given x_i, when $R = R^{(1)}$ we have $\Pi(t^{(2)}) - \Pi(t^{(1)}) = \int_{t^{(1)}}^{t^{(2)}} \frac{d\Pi(t)}{dt} dt$, which does not change as $R^{(1)}$ increases because $\frac{d\Pi(t)}{dt}$ does not change with R per (3.8). This implies that the optimal bid-down-to levels when $R = R^{(1)}$ and $R^{(2)}$, denoted by $\underline{p}(x_i)|_{R=R^{(1)}}$ and $\underline{p}(x_i)|_{R=R^{(2)}}$, respectively, should be such that either $\underline{p}(x_i)|_{R=R^{(2)}} = \underline{p}(x_i)|_{R=R^{(1)}}$ or $\underline{p}(x_i)|_{R=R^{(2)}} > R^{(1)} > \underline{p}(x_i)|_{R=R^{(1)}}$. Namely, the optimal bid-down-to level increases in R.

3.6.3 Proof of Proposition 8

When $U(t - x_i) = t - x_i$ and $F_e \sim U[0, 1]$, per (3.7) and (3.8), we have for $t \in [\frac{K}{\beta}, R]$

$$
\begin{aligned}
\Pi(t) &= \int_t^R (y - x_i)dy + (t - \frac{K}{\beta})(1-\beta)(t - x_i) + (R - x_i)(1 + \frac{K}{\beta} - R), \\
&\qquad\qquad \text{if } 1 + \frac{K}{\beta} \geq R, \\
&= \int_t^{1+\frac{K}{\beta}} (y - x_i)dy + (t - \frac{K}{\beta})(1-\beta)(t - x_i), \text{ if } 1 + \frac{K}{\beta} < R \text{ and } t \leq 1 + \frac{K}{\beta}, \\
&= (1-\beta)(t - x_i), \text{ if } 1 + \frac{K}{\beta} < R \text{ and } t > 1 + \frac{K}{\beta},
\end{aligned}
$$

$$\text{and } \frac{d\Pi(t)}{dt} = \beta x_i + (1-2\beta)t - \frac{K(1-\beta)}{\beta},$$
$$\text{if } 1+\frac{K}{\beta} \geq R, \text{ or if } 1+\frac{K}{\beta} < R \text{ and } t \leq 1+\frac{K}{\beta},$$
$$= 1-\beta, \text{ if } 1+\frac{K}{\beta} < R \text{ and } t > 1+\frac{K}{\beta}.$$

Let $t^*(x_i) \equiv \frac{\beta x_i}{2\beta-1} - \frac{(1-\beta)K}{\beta(2\beta-1)}$, that is, $\frac{d\Pi(t)}{dt}|_{t=t^*(x_i)} = 0$ if $1+\frac{K}{\beta} \geq R,$ or if $1+\frac{K}{\beta} < R$ and $t \leq 1+\frac{K}{\beta}$.

Cases with $R \leq 1+\frac{K}{\beta}$. Note that $\Pi(t)$ is convex in t when $0 < \beta < \frac{1}{2}$, is linear in t when $\beta = \frac{1}{2}$, and is concave in t when $\frac{1}{2} < \beta < 1$.

- When $0 < \beta \leq \frac{1}{2}$, the convexity of $\Pi(t)$ implies that the optimal solution $\underline{p}(x_i)$ equals either $t = \max\{x_i, \frac{K}{\beta}\}$ or $t = R$; this together with the fact that $\underline{p}(x_i) > x_i$ for all $x_i \in (\frac{K}{\beta}, R)$ (which was proved in the proof of Proposition 6) further imply that the optimal solution $\underline{p}(x_i)$ equals either $t = \frac{K}{\beta}$ or $t = R$. Note that $\Pi(R) = (R - \frac{K}{\beta})(1-\beta)(R-x_i) + (R-x_i)(1+\frac{K}{\beta} - R)$ and $\Pi(\frac{K}{\beta}) = \int_{\frac{K}{\beta}}^{R}(y-x_i)dy + (R-x_i)(1+\frac{K}{\beta}-R)$. It is easy to check that $\Pi(R) > \Pi(\frac{K}{\beta})$ if and only if $x_i > \frac{K}{2\beta^2} - \frac{R}{2\beta} + R$. Namely, $x_W = x_B = \frac{K}{2\beta^2} - \frac{R}{2\beta} + R$.

- When $\frac{1}{2} < \beta < 1$, the concavity of $\Pi(t)$ implies that $\underline{p}(x_i) = \max\{\frac{K}{\beta}, \min\{t^*(x_i), R\}\}$. That is, $\underline{p}(x_i) = \frac{K}{\beta}$ if $x_i \leq \frac{K}{\beta}$ (because $t^*(x_i) \leq \frac{K}{\beta}$ when $x_i \leq \frac{K}{\beta}$), $\underline{p}(x_i) = R$ if $x_i \geq \frac{2\beta-1}{\beta}R + \frac{(1-\beta)K}{\beta^2}$ (because $t^*(x_i) \geq R$ when $x_i \geq \frac{2\beta-1}{\beta}R + \frac{(1-\beta)K}{\beta^2}$), and $\underline{p}(x_i) = t^*(x_i)$ if $\frac{K}{\beta} < x_i < \frac{2\beta-1}{\beta}R + \frac{(1-\beta)K}{\beta^2}$. Namely, $x_W = \frac{K}{\beta}$ and $x_B = \frac{2\beta-1}{\beta}R + \frac{(1-\beta)K}{\beta^2}$.

Cases with $R > 1+\frac{K}{\beta}$.

- When $0 < \beta \leq \frac{1}{2}$, the convexity of $\Pi(t)$ over $t \in [\frac{K}{\beta}, 1+\frac{K}{\beta}]$ and the fact that $\Pi(t)$ increases when $t \in [1+\frac{K}{\beta}, R]$ together imply that $\Pi(t)$ is quasiconvex, and hence imply that the optimal solution $\underline{p}(x_i)$ equals either $t = \max\{x_i, \frac{K}{\beta}\}$ or $t = R$. Again, the fact that $\underline{p}(x_i) > x_i$ for all $x_i \in (\frac{K}{\beta}, R)$ (which was proved in the proof of Proposition 6) further imply that the optimal solution $\underline{p}(x_i)$ equals either $t = \frac{K}{\beta}$ or $t = R$. It is easy to check that $\Pi(R) = (1-\beta)(R-x_i)$ and $\Pi(\frac{K}{\beta}) = \frac{1}{2} + \frac{K}{\beta} - x_i$, and that $\Pi(R) > \Pi(\frac{K}{\beta})$ if and only if $x_i > \frac{K}{2\beta^2} + \frac{1}{2\beta} - \frac{R}{\beta} + R$. Namely, $x_W = x_B = \frac{K}{2\beta^2} + \frac{1}{2\beta} - \frac{R}{\beta} + R$.

- When $\frac{1}{2} < \beta < 1$, the concavity of $\Pi(t)$ over $t \in [\frac{K}{\beta}, 1 + \frac{K}{\beta}]$ and the fact that $\Pi(t)$ increases when $t \in [1 + \frac{K}{\beta}, R]$ together imply that $\underline{p}(x_i)$ equals either R or $\hat{t}(x_i) \equiv \max\{\frac{K}{\beta}, \min\{t^*(x_i), 1 + \frac{K}{\beta}\}\}$. Note that $\Pi(R) = (1-\beta)(R - x_i)$ and $\Pi(\hat{t}(x_i)) = \int_{\hat{t}(x_i)}^{1+\frac{K}{\beta}} (y - x_i)dy + [\hat{t}(x_i) - \frac{K}{\beta}](1-\beta)[\hat{t}(x_i) - x_i]$. Thus, for $\frac{2\beta-1}{\beta} + \frac{K}{\beta} \le x_i < R$, we have $\hat{t}(x_i) = 1 + \frac{K}{\beta}$ and hence $\Pi(\hat{t}(x_i)) = (1-\beta)[1 + \frac{K}{\beta} - x_i] < \Pi(R)$; for $x_i < \frac{K}{\beta}$, we have $\hat{t}(x_i) = \frac{K}{\beta}$ and hence $\Pi(x_i) = \int_{\frac{K}{\beta}}^{1+\frac{K}{\beta}} (y - x_i)dy = \frac{1}{2} + \frac{K}{\beta} - x_i$, which is less than $\Pi(R)$ for x_i small enough. Therefore, the continuity of $\Pi(\hat{t}(x_i))$ implies that there exists a threshold $\hat{x}_i < \frac{2\beta-1}{\beta} + \frac{K}{\beta}$ such that $\Pi(\hat{x}_i) = \Pi(R)$ and $\Pi(\hat{t}(x_i)) < \Pi(R)$ if and only if $x_i < \hat{x}_i$; namely, $x_B = \hat{x}_i$. In particular, $\hat{x}_i$ solves

$$\int_{\hat{t}(\hat{x}_i)}^{1+\frac{K}{\beta}} (y - \hat{x}_i)dy + [\hat{t}(\hat{x}_i) - \frac{K}{\beta}](1-\beta)[\hat{t}(\hat{x}_i) - \hat{x}_i] = \Pi(\hat{x}_i) = \Pi(R) = (1-\beta)(R - \hat{x}_i);$$

using simplification, we have for given K, β, and R, $x_B = \hat{x}(K, \beta, R)$ be the $x_i \in (-\infty, \frac{2\beta-1}{\beta} + \frac{K}{\beta})$ solving the following equation:

$$0 = \begin{cases} \frac{(\beta x_i - K)^2}{2(2\beta-1)} - \beta x_i + \frac{K}{\beta} + \frac{1}{2} - (1-\beta)R & , \text{if } \frac{K}{\beta} < x_i \le \frac{2\beta-1}{\beta} + \frac{K}{\beta}; \\ -\beta x_i + \frac{K}{\beta} + \frac{1}{2} - (1-\beta)R & , \text{if } x_i \le \frac{K}{\beta}. \end{cases}$$

Finally, if $\hat{x}(K, \beta, R) > \frac{K}{\beta}$, we have $\underline{p}(x_i) = \hat{t}(x_i)$ for $x_i < \hat{x}(K, \beta, R)$ (i.e., $\underline{p}(x_i) = t^*(x_i)$ for $\frac{K}{\beta} < x_i < \hat{x}(K, \beta, R)$ and $\underline{p}(x_i) = t^*(x_i)$ for $x_i \le \frac{K}{\beta}$), which implies that $x_W = \frac{K}{\beta}$; if $\hat{x}(K, \beta, R) \le \frac{K}{\beta}$, we have $\underline{p}(x_i) = \frac{K}{\beta}$, which implies that $x_W = x_B$. To summarize, $x_W = \min\{\frac{K}{\beta}, x_B\}$.

3.6.4 Proof of Proposition 9

The probability that the incumbent wins the auction outright equals

$$1 - Prob\left[x_e + \frac{K}{\beta} \leq \underline{p}(x_i)\right] = \int_l^R \max\{1 + \frac{K}{\beta} - \underline{p}(x_i), 0\}dx_i.$$

It increases in K because $1 + \frac{K}{\beta} - \underline{p}(x_i)$ increase in K per Proposition 8.

We then study the effect of β. The probability is in general not monotone in β since $1 + \frac{K}{\beta} - \underline{p}(x_i)$ is general non-monotone in β. To prove the proposition, we first consider the case with $K = 0$.

- When $0 < \beta \leq \frac{1}{2}$, $Prob\left[x_e + \frac{K}{\beta} \leq \underline{p}(x_i)\right]$ equals $\frac{(R-x_B)\min\{R,1\}}{R-l}$. It decreases in β because x_B increases in β. To see this, $\frac{dx_B}{d\beta} = \frac{R}{2\beta^2}$ (if $R \leq 1$) or $-\frac{1}{2\beta^2} + \frac{R}{\beta^2}$ (if $R > 1$), which is positive. Thus, the winning probability increases in β.

- When $\frac{1}{2} < \beta < 1$, $R \leq 1$ and $l < 0$, $Prob\left[x_e + \frac{K}{\beta} \leq \underline{p}(x_i)\right]$ equals $\frac{R}{2(R-l)}(R + R - x_B)$, which decreases in β because $x_B = (2 - \frac{1}{\beta})R$ increases in β. Thus, the winning probability increases in β.

- When $\frac{1}{2} < \beta < 1$, $R \leq 1$ and $0 < l \leq x_B$, $Prob\left[x_e + \frac{K}{\beta} \leq \underline{p}(x_i)\right]$ equals $[\frac{R}{2}(R + R - x_B) - \frac{l^2}{2}\frac{\beta}{2\beta-1} = \frac{R^2}{2\beta} - \frac{l^2\beta}{2(2\beta-1)}]\frac{1}{R-l}$, which decreases in β, because its derivative with respect to β equals $[-\frac{R^2}{2\beta^2} + \frac{l^2}{2(2\beta-1)^2}]\frac{1}{R-l}$, which is non-positive because $l \leq x_B = \frac{2\beta-1}{\beta}R$. Thus, the winning probability increases in β.

- When $\frac{1}{2} < \beta < 1$, $R \leq 1$ and $l > x_B$, $Prob\left[x_e + \frac{K}{\beta} \leq \underline{p}(x_i)\right]$ equals R. It is constant as β changes.

- When $\frac{1}{2} < \beta < 1$, $R > 1$, $x_B \leq 0$, the winning probability equals $\max\{x_B - l, 0\}\frac{1}{R-l}$, which increases in β because x_B increases in β. To see that, per Proposition 8 we know if $x_B \leq 0$ then $x_B = \frac{1}{2\beta} - \frac{R}{\beta} + R$, with its first order derivative with respect to β equal to $\frac{R-0.5}{\beta^2} > 0$.

- When $\frac{1}{2} < \beta < 1$, $R > 1$, $x_B > 0$ and $l \leq 0$, the winning probability equals $[(x_B - l) - (x_B)^2 \frac{\beta}{2(2\beta-1)}]\frac{1}{R-l}$, equal to $[\frac{1}{2\beta} - \frac{R}{\beta} + R - l]\frac{1}{R-l}$ per Proposition 8, which increases in β.

- When $\frac{1}{2} < \beta < 1$, $R > 1$, $x_B > 0$ and $0 < l \leq x_B$, the winning probability equals $[(x_B - l) - (x_B)^2 \frac{\beta}{2(2\beta-1)} - \frac{l}{2}(1 + 1 - \frac{\beta l}{2\beta-1})]\frac{1}{R-l}$, equal to $[\frac{1}{2\beta} - \frac{R}{\beta} + R - 2l + \frac{\beta l^2}{2(2\beta-1)}]\frac{1}{R-l}$ (per Proposition 8), with the first order derivative with respect to β equal to $[\frac{2R-1}{2\beta^2} - \frac{l^2}{2(2\beta-1)^2}]\frac{1}{R-l}$, which is positive because $l \leq x_B < \frac{2\beta-1}{\beta}$ per Proposition 8. Hence, the probability increases in β.

- When $\frac{1}{2} < \beta < 1$, $R > 1$, $x_B > 0$ and $l > x_B$, the winning probability equals zero. Thus, it must increase in β.

We next consider K approaches βR.

- When $0 < \beta \leq \frac{1}{2}$, since $R \leq 1 + \frac{K}{\beta}$ we have $x_B = x_W = \frac{K}{2\beta^2} - \frac{R}{2\beta} + R$ and $Prob\left[x_e + \frac{K}{\beta} \leq \underline{p}(x_i)\right] = (R - \frac{K}{\beta})(R - x_B)\frac{1}{R-l}$, which increases in β because $\frac{d(K/\beta)}{d\beta} < 0$ and $\frac{dx_B}{d\beta} = \frac{1}{\beta^2}(\frac{R}{2} - \frac{K}{\beta}) < 0$ when K approaches βR. Thus, the winning probability decreases in β.

- When $\frac{1}{2} < \beta < 1$, since $R \leq 1 + \frac{K}{\beta}$ and $l < \frac{K}{\beta}$ we have $Prob\left[x_e + \frac{K}{\beta} \leq \underline{p}(x_i)\right] = (R - \frac{K}{\beta})^2 \frac{1}{2\beta}\frac{1}{R-l}$, which increases in β when K is close to βR, because its derivative with respective β equals $\frac{R-\frac{K}{\beta}}{\beta^2}(\frac{K}{\beta} - \frac{1}{2}(R - \frac{K}{\beta}))\frac{1}{R-l}$, which is positive if K is close enough to βR. Thus, the winning probability decreases in β as K approaches βR.

3.6.5 Proof of Proposition 10

We find it more convenient to characterize the optimal bidding strategy by studying the equivalent dynamic bidding problem. Note that equation (3.5) minus equation (3.3) is given by

$$\sup_{\max\{x_i, \frac{K}{\beta}\} \leq t < p} \int_t^p \{\beta[U(y - x_i) - U(0)]g(y|p) - (1-\beta)G(y|p)\frac{\partial U(y - x_i)}{\partial y}\}dy. \tag{3.11}$$

Given K, β, x_i, and p, define

$$DF(y|p) \triangleq \beta[U(y-x_i)-U(0)]g(y|p)-(1-\beta)G(y|p)\frac{\partial U(y-x_i)}{\partial y}. \tag{3.12}$$

Thus, for given F_e, K, β, and x_i, the incumbent's optimal dynamic bidding strategy at a current calling price p is given in the following lemma.

Lemma 3 *For given F_e, K, β, x_i, and p, it is optimal for the incumbent to*

- *Drop out at the current calling price p, if $\int_t^p DF(y|p)dy \leq 0$ for all $t \in [\max\{x_i, \frac{K}{\beta}\}, p]$, and to*
- *Stay in the auction, if $\int_t^p DF(y|p)dy > 0$ for some $t \in [\max\{x_i, \frac{K}{\beta}\}, p]$.*

We first study cases where $1 \leq R \leq 1+\frac{K}{\beta}$ and then study cases where $R > 1+\frac{K}{\beta}$.

Cases with $1 \leq R \leq 1+\frac{K}{\beta}$. In these cases, $g(y|p)=1$ and $G(y|p)=\frac{y-\frac{K}{\beta}}{p-\frac{K}{\beta}}$; applying equation (3.12) with $U(p-x_i)=-\frac{1}{\gamma}e^{-\gamma(p-x_i)}$, we get

$$DF(y|p) \quad = \quad \frac{\beta e^{-\gamma(y-x_i)}}{\gamma(p-\frac{K}{\beta})}[e^{\gamma(y-x_i)}-1-\gamma(y-\frac{K}{\beta})\frac{1-\beta}{\beta}]. \tag{3.13}$$

Fixing K, β, γ, y, and p, it is easy to check that $DF(y|p)$ is strictly decreasing in x_i and its sign is the same as that of $s(y,x_i) \triangleq e^{\gamma(y-x_i)}-1-\gamma(y-\frac{K}{\beta})\frac{1-\beta}{\beta}$. It is easy to check that $s(y,x_i)$ is strictly decreasing in x_i. Thus, fixing y, $s(y,x_i)$ (and hence $DF(y|p)$) is greater than, or equal to, or less than 0, if and only if x_i is less than, or equal to, or greater than $\overline{x}(y) \triangleq y-\frac{1}{\gamma}\ln[1+\gamma(y-\frac{K}{\beta})\frac{1-\beta}{\beta}]$, respectively. Furthermore, $s(y,x_i)$ is strictly convex in y, and thus has at most two roots. We will let $\bar{p}(x_i)$ and $\tilde{p}(x_i)$ denote the larger and smaller roots, respectively, if they exist.

Cases with $1 \leq R \leq 1+\frac{K}{\beta}$ and $\frac{1}{2} \leq \beta < 1$. For these cases, we find the following about the sign of $DF(y|p)$.

- If $x_i \leq \frac{K}{\beta}$, then $\max\{\frac{K}{\beta}, x_i\} = \frac{K}{\beta}$ and the interval of interest is $y \in [\frac{K}{\beta}, p]$. Since one can show that $DF(\frac{K}{\beta}|p) \geq 0$ and $\frac{\partial}{\partial y}DF(y|p)$ is non-negative at $y = \frac{K}{\beta}$, the strict concavity of $DF(y|p)$ in y implies that $DF(y|p) > 0$ for $y \in (\frac{K}{\beta}, p]$ (and hence it is optimal to bid until the entrant drops out or the price reaches $\frac{K}{\beta}$).

- If $x_i > \frac{K}{\beta}$, then $\max\{\frac{K}{\beta}, x_i\} = x_i$ and the interval of interest is $y \in [x_i, p]$. For a fixed $x_i \leq \frac{K}{\beta}$, $s(y, x_i)$ is convex in y and is negative at $y = x_i$ and goes to infinity as y goes to infinity. Thus, $\overline{p}(x_i)$ exists and is strictly greater than x_i. In other words, $DF(y|p)$ is negative, zero, or positive, if and only if $y \in [x_i, \overline{p}(x_i))$, $y = \overline{p}(x_i)$, or $y \in (\overline{p}(x_i), p]$, respectively. This implies that, if $x_i \geq \overline{x}(p)$ then $DF(y|p) \leq 0$ for $y \in [\frac{K}{\beta}, p]$ (and hence it is optimal to drop out at p), if $\frac{K}{\beta} < x_i < \overline{x}(p)$ then $DF(y|p) \leq 0$ for $y \in [\frac{K}{\beta}, \overline{p}(x_i)]$ and $DF(y|p) > 0$ for $y \in (\overline{p}(x_i), p]$ (and hence it is optimal to bid until the entrant drops out or the price reaches $\overline{p}(x_i)$).

Because the auction starts at reserve price R, we can describe the incumbent's strategy by examining the case $p = R$. One can check that $R > \frac{K}{\beta}$ and $\frac{1}{2} \leq \beta < 1$ together imply $\overline{x}(R) > \frac{K}{\beta}$. To summarize, when $1 \leq R \leq 1 + \frac{K}{\beta}$ and $\frac{1}{2} \leq \beta < 1$ the incumbent's strategy is the following:

- If $x_i \geq \overline{x}(R)$, boycott the auction by dropping out immediately at R;

- Otherwise, if $\frac{K}{\beta} < x_i < \overline{x}(R)$, bid until the entrant either drops out or the price reaches $\overline{p}(x_i)$;

- Otherwise, bid until the entrant drops out.

Cases with $1 \leq R \leq 1 + \frac{K}{\beta}$ and $0 < \beta < \frac{1}{2}$. For these cases, the sign of $DF(y|p)$ over the whole interval $y \in [\frac{K}{\beta}, p]$ falls into one of the following five patterns, illustrated by Figure 3.4, depending on the value of x_i. This is because, for fixed x_i, $s(y, x_i)$ is convex in y.

- Pattern 1: $DF(y|p) \leq 0$ for all $y \in [\frac{K}{\beta}, p]$. With pattern 1, it is optimal to drop out at p.

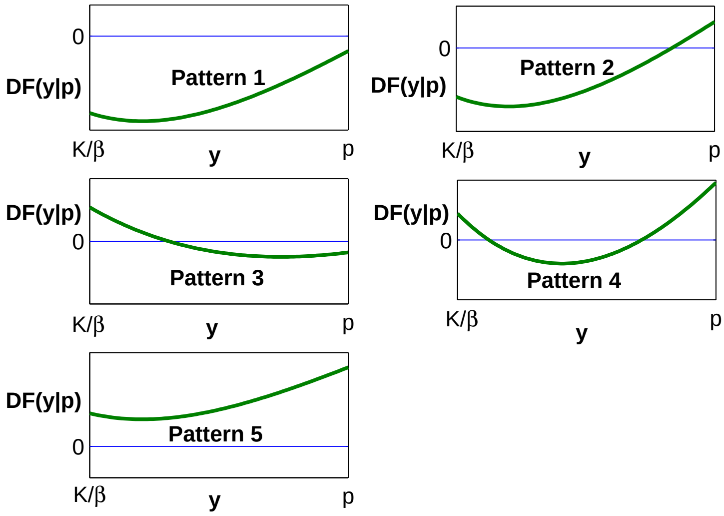

Figure 3.4: Five patterns of $DF(y|p)$.

- Pattern 2: $DF(y|p) \leq 0$ for all $y \in [\frac{K}{\beta}, \overline{p}(x_i)]$ and $DF(y|p) > 0$ for all $y \in (\overline{p}(x_i), p]$. With pattern 2, it is optimal to bid down to $\overline{p}(x_i)$.

- Pattern 3: $DF(y|p) > 0$ for all $y \in [\frac{K}{\beta}, \tilde{p}(x_i))$ and $DF(y|p) \leq 0$ for all $y \in [\tilde{p}(x_i), p]$. With pattern 3, it is optimal to bid until the entrant drops out if $\int_{\frac{K}{\beta}}^{p} DF(y|p)dy > 0$, or otherwise, it is optimal to drop out at p.

- Pattern 4: $DF(y|p) > 0$ for all $y \in [\frac{K}{\beta}, \tilde{p}(x_i)) \cup (\overline{p}(x_i), p]$ and $DF(y|p) < 0$ for all $y \in (\tilde{p}(x_i), \overline{p}(x_i))$. With pattern 4, it is optimal to bid until the entrant drops out if $\int_{\frac{K}{\beta}}^{\overline{p}(x_i)} DF(y|\overline{p}(x_i))dy > 0$, or otherwise, it is optimal to bid down to $\overline{p}(x_i)$.

- Pattern 5: $DF(y|p) \geq 0$ for all $y \in (\frac{K}{\beta}, p)$. With pattern 5, it is optimal to bid until the entrant drops out.

We now simplify the incumbent's optimal bidding strategy when $1 \leq R \leq 1 + \frac{K}{\beta}$ and $0 < \beta < \frac{1}{2}$. Let $\underline{x}(p) \triangleq p - \frac{1}{\gamma}\ln[\frac{e^{\gamma(p-\frac{K}{\beta})} - 1 - \gamma(1-\beta)(p-\frac{K}{\beta})}{\gamma(p\beta - K)}]$. It is easy to check that $\underline{x}(p)$ is the unique x_i such that $\int_{\frac{K}{\beta}}^{p} DF(y|p)dy = 0$. In other words, $\int_{\frac{K}{\beta}}^{p} DF(y|p)dy$ is positive, zero, or negative, if and only if x_i is less than, or equal to, or greater than $\underline{x}(p)$, respectively.

Case $\underline{x}(R) \geq \overline{x}(R)$. The pattern of the sign of $DF(y|p)$ must be pattern 3 when $x_i = \underline{x}(R)$. This is because $s(R, \underline{x}(R)) \leq s(R, \overline{x}(R)) = 0$ (given that $s(y, x_i)$ is strictly decreasing in x_i) and $s(\frac{K}{\beta}, \underline{x}(R)) > 0$ (given that $\int_{\frac{K}{\beta}}^{R} DF(y|R)dy = 0$ when $x_i = \underline{x}(R)$). Thus, if $x_i \geq \underline{x}(R)$ (and hence $\int_{\frac{K}{\beta}}^{R} DF(y|R)dy \leq 0$) the pattern of the sign of $DF(y|p)$ must be pattern 3 or pattern 1; this is because $DF(y|R)$ is strictly decreasing in x_i with y and R fixed. With either pattern 3 or pattern 1, it is optimal to drop out at R, given that Given that $\int_{\frac{K}{\beta}}^{R} DF(y|R)dy \leq 0$. Otherwise, if $x_i < \underline{x}(R)$ (and hence $\int_{\frac{K}{\beta}}^{R} DF(y|R)dy > 0$) the pattern of the sign of $DF(y|p)$ must be pattern 3, or pattern 4, or pattern 5; again this is because $DF(y|R)$ is strictly decreasing in x_i with y and R fixed. With either pattern 3, or pattern 4, or pattern 5, it is optimal to bid until the entrant drops out.

Case $\underline{x}(R) < \overline{x}(R)$. The pattern of the sign of $DF(y|p)$ must be pattern 2 or 4 when $x_i = \underline{x}(R)$. This is because $s(R, \underline{x}(R)) > s(R, \overline{x}(R)) = 0$ (given that $s(y, x_i)$ is strictly decreasing in x_i) and $\int_{\frac{K}{\beta}}^{R} DF(y|R)dy = 0$ when $x_i = \underline{x}(R)$. With either pattern 2 or 4, there exists $\overline{p}(\underline{x}(R)) \in (\frac{K}{\beta}, R)$ such that when $x_i = \underline{x}(R)$ we have $DF(y|R) = 0$ for $y = \overline{p}(\underline{x}(R))$ and $DF(y|R) > 0$ for all $y \in (\overline{p}(\underline{x}(R)), R]$. Given that $\int_{\frac{K}{\beta}}^{\overline{p}(\underline{x}(R))} DF(y|R)dy < \int_{\frac{K}{\beta}}^{R} DF(y|R)dy = 0$, there must exist a unique $x_i < \underline{x}(R)$ such that $\int_{\frac{K}{\beta}}^{\overline{p}(x_i)} DF(y|R)dy = 0$; this is because $DF(y|R)$ is strictly decreasing in x_i with y and R fixed and $DF(y|R)$ goes to infinity for all y as x_i goes to negative infinity. Denote such unique x_i by $\tilde{x}$. Thus, when $x_i \geq \overline{x}(R)$ (and hence $x_i > \underline{x}(R)$, i.e., $\int_{\frac{K}{\beta}}^{R} DF(y|R)dy < 0$) the pattern of the sign of $DF(y|p)$ must be pattern 3 or pattern 1, which implies that it is optimal to drop out at R; when $\tilde{x} < x_i < \overline{x}(R)$ (and hence $\int_{\frac{K}{\beta}}^{\overline{p}(x_i)} DF(y|\overline{p}(x_i))dy < 0$) the pattern of the sign of $DF(y|p)$ must be pattern 2 or pattern 4, which implies that it is optimal to bid down to $\overline{p}(x_i)$; when $x_i \leq \tilde{x}$ the pattern of the sign of $DF(y|p)$ must be pattern 4 or pattern 5, which implies that it is optimal to bid until the entrant drops out.

Cases with $R > 1 + \frac{K}{\beta}$. In these cases, $g(y|p) = 0$, $G(y|p) = 1$ when $y > 1 + \frac{K}{\beta}$ and $g(y|p) = 1$ and $G(y|p) = y - \frac{K}{\beta}$ when $y \leq 1 + \frac{K}{\beta}$; applying equation (3.12) with $U(p - x_i) = -\frac{1}{\gamma} e^{-\gamma(p - x_i)}$, we get

$$DF(y|p) = \begin{cases} -(1-\beta)e^{-\gamma(y-x_i)} & \text{, if } y > 1+\frac{K}{\beta}; \\ \frac{\beta e^{-\gamma(y-x_i)}}{\gamma}[e^{\gamma(y-x_i)} - 1 - \gamma(y-\frac{K}{\beta})\frac{1-\beta}{\beta}] & \text{, if } y \leq 1+\frac{K}{\beta}. \end{cases}$$

Let $\hat{p}(x_i)$ be the bid-down-to level when $R = 1+\frac{K}{\beta}$. The incumbent chooses to drop at R, if $\int_{\hat{p}(x_i)}^{R} DF(y|R)dy \leq 0$, or to bid down to $\hat{p}(x_i)$, otherwise. Given that $\hat{p}(x_i)$ increases in x_i, $DF(y|R)$ strictly decreases in x_i, and $\int_{\hat{p}(x_i)}^{R} DF(y|R)dy$ goes to infinity as x_i goes to negative infinity (since $\hat{p}(x_i)$ goes to negative infinity and $DF(y|R) > 0$ when x_i is small enough), $\int_{\hat{p}(x_i)}^{R} DF(y|R)dy$ strictly decreases in x_i and hence there must exist an $\hat{x}(R)$ such that $\int_{\hat{p}(x_i)}^{R} DF(y|R)dy \leq 0$ if and only if $x_i \geq \hat{x}(R)$. Therefore, when $R > 1+\frac{K}{\beta}$ the incumbent's strategy is the following:

- If $x_i \geq x_B$, drop out immediately at R, i.e., *boycott* the bidding process;
- Otherwise, if $x_W < x_i < x_B$, bid down to $\overline{p}(x_i)$;
- Otherwise, bid until the entrant drops out.

The thresholds $x_B = \min\{x_B(1+\frac{K}{\beta}), \hat{x}(R)\}$ and $x_W = \min\{x_W(1+\frac{K}{\beta}), \hat{x}(R)\}$, where $x_B(1+\frac{K}{\beta})$ and $x_W(1+\frac{K}{\beta})$ correspond to the two thresholds when $R = 1+\frac{K}{\beta}$.

3.6.6 Proof of Proposition 11

For any $0 < \beta < 1$, the expected total cost under pre-qualification linearly increases in qualification cost K (per (3.1)), whereas the expected total cost under post-qualification is bounded by R from above (per (3.2)). Thus, post-qualification yields a lower expected total cost when K is large enough; namely, it must exist a $\underline{K} \geq 0$ such that the buyer prefers post-qualification whenever $K > \underline{K}$.

Chapter 4

Bargaining Power and Supply Base Diversification

4.1. Introduction

It is common for buyers (procurement managers) responsible for procuring an item to identify a *supply base*, a group of qualified suppliers that are capable of producing the item. A supply base is a well-known tool for managing risks. For specialized items where availability is the main objective, buyers can place orders with multiple suppliers to manage non-delivery risks (e.g., Anupindi and Akella 1993). But, as is our focus in this paper, a supply base can also be a crucial strategic tool for purchasing commodity-type items where cost, not availability, is the central issue.

Buyers typically do not know the true costs of suppliers, who possess private information about their cost drivers (inventory level, capacity utilization, financial status, etc.). To find a low price, buyers increasingly employ procurement auctions aimed at price discovery (Jap 2003). As the practitioner survey Beall et al. (2003) page 49 points out, "If a qualified supply base is identified, and the market for a particular commodity/purchase family group changes rapidly, [procurement auctions] are an excellent tool to award business for short duration and re-auction regularly. For example, one company interviewed purchases highly engineered printed circuit boards quarterly through [procurement auctions]." In a procurement auction, competition between suppliers can come down to cents or fractions of a cent, yet these small differences can translate into millions of dollars of savings to the buyer given large volumes — a high tech firm we interacted with runs quarterly auctions in which commodity (cables, connectors, etc.) suppliers compete on unit prices in increments of one tenth of a cent.

When margins are razor-thin, factors such as transportation costs, taxes and commissions become non-negligible (Pederson 2004). Buyers are increasingly aware of the need to make sourcing decisions based on *total cost*, which from the buyer's perspective measures the total cost of procuring from the supplier. In addition to the supplier's price, total cost includes non-price costs such as duties, transportation costs, shipping insurance and commissions (Ariba 2005). In this paper we introduce the idea of strategic supply base design to mitigate total procurement cost shocks, and examine how the buyer's optimal supply base design is affected by the buyer's bargaining clout. We now motivate and introduce both these concepts.

Supply base design to mitigate cost shocks. The "non-price costs" associated with a supplier can be closely related to the supplier's geographic region, and thus subject to *cost shocks* affecting that region. For example, shipping costs associated with procuring from a supplier are largely affected by local logistics markets and regulations within the supplier's region, and can be dramatically increased by labor strikes or regulation changes. In February 2007, the CN Railway strike disabled almost three quarters of Canada's rail capacity, forcing companies such as Ford to look for much more expensive alternatives like truck freight for shipments from its Canadian suppliers. Seeking heightened security for the Olympics in the summer of 2008, the Chinese government forbade a wide range of hazardous materials at six major ports; affected buyers incurred significant rerouting costs. Other examples of regional cost shocks include ocean shipping insurance rates (which are based on geopolitical and geosecurity elements along shipping routes[1]).

Ideally, a buyer could respond to regional cost shocks by instantly augmenting her supply base with new suppliers from unaffected regions. However, for some buyers this can be impractical (for all but the most catastrophic scenarios), because finding and qualifying a

[1]A recent example is the thousand-percent increase in shipping insurance premiums for Asia to Europe ocean transport, as freighters funneling through the Suez canal face a gauntlet of pirates and kidnappers based in an increasingly destabilized Somalia (Costello 2008).

new supplier is usually time-consuming and costly. The process of vetting suppliers, called supplier qualification screening, typically involves reference checks, financial audits, site visits to supplier facilities abroad, approval and buy-in from the buyers' internal customers, etc; see Chapter 2 of this dissertation. At a Fortune 100 manufacturer we interacted with it takes an average of 8 to 26 weeks to find and qualify a new supplier — even for commodity parts.

Instead of frequently finding and qualifying totally new suppliers, buyers, including the large manufacturer we interacted with, build their supply base as a long-term strategic decision, and then frequently auction off short-term supply contracts among them to find the current lowest-total-cost supplier. For such buyers, therefore, an important strategic decision arises when forming their supply base: Facing potential regional cost shocks, should the buyer's supply base include *similar* suppliers (selected from the same region) or *diversified* suppliers (selected from different regions)?

Intuitively, geographically diversifying the supply base, i.e., selecting suppliers from different regions, can mitigate regional cost shocks. For example, once a prolonged labor strike at the ports in region A drives up the cost of transporting goods from the supplier in region A, a buyer who sources a large and expensive-to-transport component can avoid a high transportation cost if she has a supplier in an unaffected region B. However, a buyer seeking to minimize total procurement cost needs to take into account the impact of diversifying on her contract payment: Will the supplier in region B strategically mark up his price to make a windfall profit based on his cost advantage over the supplier in region A? If so, how should the buyer design her supply base in the first place to manage both regional-costs risks and supplier-windfall-profit-taking risks?

Bargaining power and supply base design decision. In our study the buyer's contract payment is determined through a competitive bidding process (i.e., auction). Thus, it is crucial to understand how the buyer's ability to design auctions (i.e., choose auction format and rules) should be taken into account when she designs her supply base. We term such ability the buyer's *bargaining power*. For forward auctions, Bulow and Klemperer (1996) point out that an auctioneer with no bargaining power can only run an English auction with no reserve price while an auctioneer with full bargaining power can utilize an optimal auction mechanism.

Similarly, in this paper, at one extreme we model a buyer with no bargaining power — such a buyer cannot make credible take-it-or-leave-it offers and must solely rely on supplier competition for price concessions, utilizing a simple reverse English auction with no reserve price. In such an auction, the lowest-total-cost supplier charges the buyer a price that is set according to second-lowest supplier's total-cost, creating the risk of severe windfall-profit taking. Returning to our example two paragraphs above, the supplier in region B could take windfall profits and consequently the buyer's total cost could be the total cost of the supplier in region A, which includes A's regional cost shock! Thus, the imperative to diversify the supply base (i.e., choose suppliers from different regions) is mitigated by the need for cost parity among suppliers. We find that the optimal amount of diversification depends on the total number of suppliers and the likelihood of regional cost shocks.

At the other extreme, we model a buyer having full bargaining power, who thus can design an optimal procurement mechanism within which suppliers compete for the buyer's business (e.g., could promise to bias against the supplier in B who has regional cost advantage). Between the two extremes there can be intermediate cases, where for example the buyer is unable to use an optimal mechanism but can commit to using a reserve price in a reverse English auction. We find that supplier cost parity is less crucial for buyers with more bargaining power — such buyers are better served by a diversified supply base — and the optimal supply-base-design strategy can depend on the distributions of supplier costs and regional cost shocks.

The next section reviews related literature, and §4.3 introduces the model and assumptions. Section 4.4 analyzes the buyer's optimal supply-base-design problem and focuses on the two cases: one in which the buyer has full bargaining power and uses the optimal mechanism, and the other in which the buyer has no bargaining power and uses a reverse English auction without reserve price. Section 4.5.1 extends the results to the cases with codependent regional costs, §4.5.2 extends the analysis to the cases with asymmetric regions, and §4.5.3 studies the case in which the buyer has intermediate bargaining power and uses a reserve price in reverse English auctions. We provide numerical illustrations of our results in §4.6 and conclude in §4.7. All proofs are provided in the electronic companion.

4.2. Literature Review

Our paper analytically studies how buyers should select suppliers to mitigate regional cost risks, and is thus related to the supply risk management literature. However, our paper differs from the majority of the literature in two main aspects. First, we focus on supply risks that can be modeled as "cost shocks," while the existing literature mainly focuses on catastrophic "supply shocks" that cause supply shortages. Such "supply shocks," more commonly referred to as supply disruptions, include natural disasters (fire, hurricane, earthquake, etc.), supplier bankruptcy, etc. Researchers have studied various mitigation and contingency strategies to manage supply disruption risks; readers are referred to Tomlin (2006), which categorizes these strategies as stockpiling, multi-sourcing, using backup options, managing demand, and others. Among these categories, multi-sourcing and using backup options are related to supply base design. Studies on multi-sourcing to mitigate supply disruptions typically focus on buyers' inventory management decisions (e.g., determining the optimal ordering quantity and split of quantities among suppliers) and model the impact of disruptions by various random yield models. Recent examples include Dada et al. (2007), Federgruen and Yang (2007, 2008), etc.; readers are referred to Tomlin (2006), which provides a detailed survey of early work of this stream. For work including backup options in the supply base, see, for example, Yang et al. (2008) and references therein.

Second, this paper studies price escalation risks (e.g., windfall-profit taking by suppliers), while the majority of supply risk management literature presumes exogenous contract prices (or unit procurement costs) and ignores suppliers' strategic pricing behavior. One exception is Babich et al. (2007), which endogenizes suppliers' pricing decisions in a multi-sourcing problem where a buyer allocates ordering quantities among suppliers with correlated default risks. They assume that suppliers have full information of competitors' costs, and show how suppliers' pricing decisions can be affected by their default risk correlations. In particular, they find that the buyer prefers suppliers with positively correlated default risks despite the loss of diversification benefits, because default risk correlation increases supplier competition. In our paper, which studies the supply base design problem in the presence of suppliers' regional cost risks, we model supplier competition via procurement auctions in which suppliers possess private cost information, and we show how supplier competition can be affected by correlations across suppliers' cost shocks. We find that the buyer's bargaining power dictates her preference for the supply base design, namely, a buyer with stronger bargaining power prefers a more diversified supply base, which effects less correlation across suppliers' cost shocks.

The term "bargaining power" is probably one of the most widely used but vaguely defined concepts in the literature of bargaining models. In the literature of bargaining games with complete information, the asymmetric Nash bargaining model (Roth 1979) "captures some imprecisely defined 'bargaining power' " (Binmore et al. 1986) by including weighting scalars in the calculation of utility products. However, the literature on bargaining games with incomplete information focuses on analyzing bargaining outcomes given different bargaining mechanisms (see Ausubel et al. (2002) for a detailed survey), without explicitly defining players' "bargaining power." In the present paper, we interpret the term "bargaining power" as the buyer's ability to impose an auction mechanism that she favors, an interpretation that can be traced to the prominent work of Bulow and Klemeperer (1996). In other words, we use the term "bargaining power" as a way to rank the auction mechanisms that we study in this paper.

Extensive work has examined procurement cost reduction via supply base competition. Elmaghraby (2000) provides a comprehensive survey of early work on competitive sourcing strategies including auctions; more recently, Grey et al. (2005) surveys the literature on e-marketplaces in relationship-based supply chains, and Elmaghraby (2007) surveys industry practices in designing and running auctions for e-sourcing events. Our paper considers a buyer who finds the lowest-price provider by auctioning off short-term supply contracts. In supply chain settings, Chen (2007) studies buyers' auctions of supply contracts, and Chen and Vulcano (2008) studies a supplier's auction to sell capacity and compares first- with second-price auction formats. Methodologically, our paper is related to the auction and mechanism design literature; readers are referred to the books by Krishna (2002) and Milgrom (2004), which provide excellent treatments and detailed references on auction theory.

4.3. Model and Preliminaries

4.3.1 Model

We study a stylized model in which a risk-neutral buyer selects a cohort of N qualified suppliers to form a supply base for a component. We allow N to be any integer no less than two; however, in practice it is not rare that a buyer forms a supply base for a component with two suppliers, especially when the buyer needs to manage procurement of many different kinds of components.[2] In period $t = 0$, the buyer designs the supply base. For simplicity, we assume that designing supply base is a one-time decision and no suppliers are removed from or added to the supply base after the establishment of the supply base. This models the cases in which frequently finding and qualifying new suppliers are not practical due to costly and time-consuming supplier qualification screening processes. To focus on the supply base diversification decision, we assume that N is exogenously given. Suppliers can be selected from different geographic *regions*. The buyer's decision variables are the number of regions to select suppliers from, R, and the number of suppliers to select from each region, denoted

[2]For such a buyer, having many suppliers for each component would result in a great number of suppliers overall. Due to the burden of supplier qualification processes — not only on the purchasing agents but also on internal customers such as engineering — it can be difficult to have a huge total number of suppliers. For instance, for one of the firms we interacted with, it was common to have just two suppliers for a component since this was the minimum number of suppliers required by company policy. As the buyer put it to us, with two suppliers, "the monkey was off the back of the internal customer (engineering)." Basically, the engineering department viewed the need to qualify multiple suppliers as an unwelcome responsibility, and consequently was reluctant to qualify more than the minimum (two).

by $n_1, n_2, \ldots, n_R$ for region 1, region 2 ..., and region R, respectively, where $\sum_{r=1}^{R} n_r = N$. We assume that there are at least N symmetric regions available and within each region up to N suppliers can be found. (We extend our results to the cases with asymmetric regions in §4.5.2. The analysis changes in a straightforward way if limited number of regions are available; see our discussion in §4.7.) Thus, the number of regions R can be any integer from 1 to N; in particular, $R = 1$ means selecting all suppliers from only one region, which we call the *pooling strategy*, and $R = N$ means selecting each supplier from a different region, which we call the *fully diversifying strategy*.

After establishing her supply base (finding and pre-qualifying the suppliers) in period $t = 0$, in each of the following periods $t = 1, \ldots, \infty$ the buyer runs an auction to award an indivisible short-term contract to one of the suppliers in the supply base. This setup is most appropriate when the buyer procures commodity parts from suppliers, who do not fully rely on the buyer's contract to keep afloat. To keep the analysis focused and tractable, we assume that the buyer does not store inventory and does not have in-house production, hence she must contract with one supplier in every period. This setup could model, for example, a buyer who produces high tech, short life-cycle products, relies on suppliers for key components, and holds quarterly supply auctions. When analyzing auction outcomes we assume that the suppliers are risk-neutral and fully rational players following a Bayesian Nash bidding equilibrium, as is standard in the auction literature.

Two types of costs are associated with each supplier $i = 1, \ldots, N$ in each period t. The first type of cost is an idiosyncratic production cost, $x_i^t \in [0, 1]$, which as typical in auction models is assumed to be independently and identically distributed across suppliers and periods according to a commonly known distribution F. Cost x_i^t represents supplier i's firm-specific and privately known cost of fulfilling the contract offered in period t, per supplier i's inventory level, capacity utilization, working capital position, debt status, etc. For simplicity we assume F has a positive and continuous density f and is stationary over time (this assumption can be relaxed; see our discussion in §4.7). As is standard in the auction theoretic literature, we also assume that $x + \frac{F(x)}{f(x)}$ is increasing in x — a technical assumption that ensures pure strategy implementation of the optimal mechanism described later in this section, which is satisfied, for example, if f is logconcave, including normal, logistic, and exponential distributions (Bagnoli and Bergstrom 2005).

The second type of cost is a region-specific cost, y_r^t, which represents (from the buyer's perspective) common costs affecting all suppliers in region r in period t. In our analysis, we assume such regional costs are not related to suppliers' production costs but are the additional procurement expenses the buyer incurs when doing business with a supplier in the region, for instance, logistics costs. (Our results easily extend to cases where regional factors also influence suppliers' production costs; see our discussion in §4.7.) We assume that y_r^t's are commonly observable at the outset of period t and are independently and identically distributed across regions and periods according to a commonly known distribution G with finite mean (i.e., $E[y_t^r] < \infty$). We discuss how our results extend to the cases where regional costs are possibly codependent in §4.5.1, and study the cases where regions can be asymmetric in terms of distributions of supplier costs and regional costs in §4.5.2. We let a_i^t denote the regional cost of supplier i in period t, that is, $a_i^t = y_r^t$ if supplier i is located in region r. For simplicity we assume G is stationary over time, although this too can be relaxed (see §4.7).

The buyer seeks to minimize her expected long-term total procurement cost. Let $\boldsymbol{x}^t \stackrel{def}{=} (x_1^t, x_2^t ..., x_N^t)$ denote the vector of realized supplier production costs in period t; let $\boldsymbol{y}^t \stackrel{def}{=} (y_1^t, y_2^t, ..., y_N^t)$ denote the vector of realized regional costs in period t; let $\boldsymbol{a}^t \stackrel{def}{=} (a_1^t, a_2^t, ..., a_N^t)$ denote the vector of realized regional costs of suppliers in period t. Thus, the supply base design problem can be formulated as follows:

$$\min_{n_1,...,n_R} \underset{(\boldsymbol{x}^t,\boldsymbol{y}^t)}{E} \left[\sum_{t=1}^{\infty} \beta^t \pi^{Mech}(\boldsymbol{x}^t; \boldsymbol{a}^t) \right] = \frac{\beta}{1-\beta} \underset{(\boldsymbol{x}^1,\boldsymbol{y}^1)}{E} \left[\pi^{Mech}(\boldsymbol{x}^1; \boldsymbol{a}^1) \right]$$

$$s.t. \quad R \in \{1, ..., N\}, \quad n_i \in \mathbb{N} \; \forall i \in \{1, ..., R\}, \text{ and } \; n_1 + ... + n_R = N,$$

where β is a discount factor and $\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a})$ is the buyer's period-t total procurement cost given the auction mechanism $Mech$. Since x_i^t's and a_i^t's are assumed to be identically distributed from period to period, the buyer's objective is simplified to minimizing the expected one-period total procurement cost. Therefore, we omit the superscript t for notational convenience in the rest of this paper. In §4.4, we focus on two auction mechanisms — the optimal mechanism (denoted by $Mech = OPT$) and the reverse English auction without reserve price (denoted by $Mech = RE$), representing the cases in which the buyer has full bargaining power and zero bargaining power, respectively (Bulow and Klemperer 1996). In §4.5.3, we further examine our main result (the effect of bargaining power on the buyer's optimal supply base design strategy) by studying the case in which the buyer has intermediate bargaining power and can impose a reserve price in a reverse English auction (denoted by $Mech = RER$). We describe these three auction mechanisms in §4.3.2.

The buyer's supply base design strategy affects her expected total procurement cost because different strategies yield different $\boldsymbol{a}^t$ given a realized $\boldsymbol{y}^t$. For example, in a four-supplier case ($N = 4$), if the buyer selects all four suppliers from region 1, (i.e., pooling), the suppliers' regional costs are $\boldsymbol{a} = (y_1, y_1, y_1, y_1)$, while if the buyer selects two suppliers from regions 1 and 2 each, the suppliers' regional costs are $\boldsymbol{a} = (y_1, y_1, y_2, y_2)$. The pooling strategy enables the buyer "win big" (i.e. secure a low regional cost no matter which supplier wins the contract) if region 1 happens to have a low regional cost, but it is clearly a very risky strategy — the buyer would have to "lose big" (i.e. suffer a high regional cost no matter which supplier wins the contract) if a large cost shock hits region 1. In contrast, a diversification strategy, say the two-region strategy, engenders regional cost disparities among suppliers and hence increases the likelihood for the buyer to access at least some suppliers from low-cost regions. But is this more temperate, diversified approach better than potentially winning big with a pooling strategy? As yet the buyer's preference for or against diversification is unclear, mainly because the buyer's contract price is determined through supplier competition (an auction), which would obviously be affected by the cost disparities introduced by diversification strategies. Thus, the buyer's optimal supply base design strategy would reasonably depend on the number of suppliers N, the cost distributions F and G, and the auction mechanism. In this paper, we characterize the buyer's optimal supply base design strategy and describe when and if her optimal strategy depends on her ability to choose an auction mechanism (per her bargaining power). To this end, we next formally describe the auction mechanisms we will examine.

4.3.2 Auction Mechanisms

Optimal mechanism (OPT). When the buyer has full bargaining power, she can offer suppliers a join-or-leave-it mechanism such that all suppliers will participate and meanwhile the buyer's expected total procurement cost is minimized. We refer to such a mechanism as the *optimal mechanism* (OPT). Let $\psi(x_i) \stackrel{def}{=} x_i + \frac{F(x_i)}{f(x_i)}$, which is commonly referred to as supplier i's *virtual cost* in the mechanism design literature, and let $\psi(x_i)+a_i$ denote supplier i's *adjusted virtual cost*, that is, supplier i's virtual cost adjusted by the additive regional cost a_i. In equilibrium, the optimal mechanism awards the contract to supplier j having the lowest adjusted virtual cost, i.e., $j = \underset{i=1\ldots N}{\arg\min}\{\psi(x_i)+a_i\}$, breaking ties evenly, and pays the contract winner $\min\{\psi^{-1}[\psi(x_{j_1})+a_{j_1}-a_j],1\}$, where $j_1 = \underset{i=1\ldots N,\ i\neq j}{\arg\min}\{\psi(x_i)+a_i\}$ is the losing supplier with the lowest adjusted virtual cost. The payment is truncated from above by an optimal reserve price of 1. Because the buyer must contract with a supplier, if the buyer uses a reserve price, it is always optimal to set it at the worst possible supplier cost type, i.e., at 1. The optimal awarding and payment rules can be derived by straightforward adaptation of Myerson (1981) (Rezende 2009 also points this out). To implement this optimal auction, we now propose a modified reverse clock auction, in which bidding proceeds as follows. The auction begins at calling price $\psi(1)+\max_{i=1\ldots N}\{a_i\}$, and continuously drops. Each bidder signals their willingness to stay in the auction or drop out, and the auction ends when at most one bidder remains in the auction. Let p be the calling price when the auction ends. The last bidder remaining in the auction, say bidder j, wins and is paid $\min\{\psi^{-1}(p-a_j),1\}$; ties are broken evenly.

Proposition 12 *The optimal mechanism can be implemented by the modified reverse clock* *auction described above. Furthermore, in such an auction bidders have a dominant strategy of staying in the auction until the calling price reaches their true adjusted virtual cost.*

Reverse English auction without/with reserve (RE/RER). In the case where the buyer does not have any bargaining power, she can only demand price concessions on the basis of competing offers from suppliers and cannot credibly impose a reserve price. Thus, the contract award and payment decisions can be modeled as outcomes of a reverse English total-cost auction without a reserve price (RE). In such an auction, a supplier's regional cost is added to his price bid to yield a *total-cost* bid. For instance, a supplier with a regional cost of \$8,000 who bids \$100,000 for the contract has a total-cost bid of \$108,000. During the auction, suppliers can see the current lowest total-cost bid, and can respond by lowering their own price bid if they are not the current lowest total-cost bidder. In such an auction, it is a weakly dominant strategy for a supplier to bid down to his true total cost $x_i + a_i$ (although he may not have to); see, for example, Maskin and Riley (2000). Thus, the auction ends when the second-lowest total-cost supplier drops out of the auction, and the lowest total-cost supplier wins the contract and is paid the second-lowest total cost minus the winner's regional cost. We also study cases where the buyer has some bargaining power and can impose the optimal reserve price of 1 in a reverse English auction (RER). In such a case, it remains optimal for bidders to bid down to their true costs before dropping out.

Under the three mechanisms, the buyer's expected total procurement cost can be written as expectations of order statistics as follows:

$$\underset{(\boldsymbol{x},\boldsymbol{y})}{E}\left[\pi^{OPT}(\boldsymbol{x},\boldsymbol{a})\right] = \underset{(\boldsymbol{x},\boldsymbol{y})}{E}\left[\min_{i=1,\ldots,N}\{\psi(x_i)+a_i\}\right]; \tag{4.1a}$$

$$\underset{(\boldsymbol{x},\boldsymbol{y})}{E}\left[\pi^{RE}(\boldsymbol{x},\boldsymbol{a})\right] = \underset{(\boldsymbol{x},\boldsymbol{y})}{E}\left[\underset{i=1,\ldots,N}{\text{second min}}\{x_i+a_i\}\right]; \tag{4.1b}$$

$$\underset{(\boldsymbol{x},\boldsymbol{y})}{E}\left[\pi^{RER}(\boldsymbol{x},\boldsymbol{a})\right] = \underset{(\boldsymbol{x},\boldsymbol{y})}{E}\left[\underset{i=1,\ldots,N}{\text{second min}}\{x_i+a_i, 1+a_i\}\right], \tag{4.1c}$$

where "second min$\{\cdot\}$" denotes the second-lowest value in the set. Throughout the paper, let $X_{k:N}$ and $Y_{k:N}$ denote the k^{th}-lowest order statistic out of N independent and identical random draws from distributions F and G, respectively, let $\overline{X}_{k:N}$ and $\overline{Y}_{k:N}$ denote their expectations, respectively, let $\mathbb{I}_{\{A\}}$ denote the indicator function of event A, and let $\vee$ and $\wedge$ denote the componentwise maximum and minimum operators, respectively.

4.4. Analysis and Results

To evaluate the buyer's expected total procurement cost under different diversification strategies, we need to compute the expectation of asymmetrically distributed random variables as shown by equations (4.1). However, this would be technically intractable mainly because closed-form expressions for expectations of order statistics are generally restricted to identically and independently distributed random variables following a handful of distributions such as power-function family distributions and exponential distribution. Our problem is even more challenging in that the expected total procurement cost, first, takes ex ante expectation over $\boldsymbol{x}$ given realized $\boldsymbol{a}$, involving order statistics of random variables from asymmetric distributions, and then, it takes ex ante expectation over $\boldsymbol{a}$, which involves elements that can exhibit various correlation depending on the supply base design strategy.

Thus, in order to have a hope of tackling the challenging problem of optimal supply base design, we need to exploit whatever structure can be found in the problem. We accomplish this by undertaking an iterative analysis of the buyer's diversification tradeoff, introduced next.

4.4.1 Diversification Tradeoff

Suppose the buyer compares an R-region diversification strategy $(\hat{n}_1, \hat{n}_2, \ldots, \hat{n}_R)$ with the $(R+1)$-region strategy $(\hat{n}_1, \hat{n}_2, \ldots, \hat{n}_{R-1}, \tilde{n}_R, \tilde{n}_{R+1})$ such that $\hat{n}_R = \tilde{n}_R + \tilde{n}_{R+1}$. Let $\hat{\boldsymbol{a}}$ be the vector of suppliers' regional costs under the R-region strategy and let $\tilde{\boldsymbol{a}}$ denote the vector of suppliers' regional costs under the $(R+1)$-region strategy. Given that all suppliers have independent and identical cost distributions, the difference between the two strategies comes entirely from the suppliers' regional costs $\hat{\boldsymbol{a}}$ and $\tilde{\boldsymbol{a}}$. We use a sample-path analysis as follows. On a sample path with given regional costs $\boldsymbol{y}$, the suppliers' regional costs $\hat{\boldsymbol{a}}$ and $\tilde{\boldsymbol{a}}$ can only differ from each other in the last $\tilde{n}_R$ elements. In particular, when region R experiences a larger cost shock than region $(R+1)$ does, i.e., $y_R > y_{R+1}$, switching to the $(R+1)$-region strategy would have saved the buyer money, resulting in a *diversifying upside*. Conversely, when region R experiences a smaller cost shock than region $(R+1)$ does, i.e., $y_R < y_{R+1}$, switching to the $(R+1)$-region strategy would have resulted in an expected disbenefit for the buyer, the *diversifying downside*. To facilitate expressing suppliers' regional costs under these two strategies, we let

$$\begin{aligned}
\boldsymbol{a^{hh}} &\overset{def}{=} (a_1, \ldots, a_{N-n_R}, \underbrace{y_R \vee y_{R+1}, \ldots, y_R \vee y_{R+1}}_{\tilde{n}_R \text{ elements}}, \underbrace{y_R \vee y_{R+1}, \ldots, y_R \vee y_{R+1}}_{\tilde{n}_{R+1} \text{ elements}}), \\
\boldsymbol{a^{hl}} &\overset{def}{=} (a_1, \ldots, a_{N-n_R}, \underbrace{y_R \vee y_{R+1}, \ldots, y_R \vee y_{R+1}}_{\tilde{n}_R \text{ elements}}, \underbrace{y_R \wedge y_{R+1}, \ldots, y_R \wedge y_{R+1}}_{\tilde{n}_{R+1} \text{ elements}}), \\
\boldsymbol{a^{lh}} &\overset{def}{=} (a_1, \ldots, a_{N-n_R}, \underbrace{y_R \wedge y_{R+1}, \ldots, y_R \wedge y_{R+1}}_{\tilde{n}_R \text{ elements}}, \underbrace{y_R \vee y_{R+1}, \ldots, y_R \vee y_{R+1}}_{\tilde{n}_{R+1} \text{ elements}}), \text{ and} \\
\boldsymbol{a^{ll}} &\overset{def}{=} (a_1, \ldots, a_{N-n_R}, \underbrace{y_R \wedge y_{R+1}, \ldots, y_R \wedge y_{R+1}}_{\tilde{n}_R \text{ elements}}, \underbrace{y_R \wedge y_{R+1}, \ldots, y_R \wedge y_{R+1}}_{\tilde{n}_{R+1} \text{ elements}}).
\end{aligned}$$

In other words, $\boldsymbol{a^{hh}}$ denotes the vector $\hat{\boldsymbol{a}}$ when $y_R \geq y_{R+1}$, $\boldsymbol{a^{hl}}$ denotes the vector $\tilde{\boldsymbol{a}}$ when $y_R \geq y_{R+1}$, $\boldsymbol{a^{ll}}$ denotes the vector $\hat{\boldsymbol{a}}$ when $y_R < y_{R+1}$, and $\boldsymbol{a^{lh}}$ denotes the vector $\tilde{\boldsymbol{a}}$ when $y_R < y_{R+1}$. Thus, for a given $Mech$ and realized $(\boldsymbol{x}, \boldsymbol{y})$ we can write

$$\begin{aligned}
\text{diversifying upside} &\overset{def}{=} \left[\pi^{Mech}(\boldsymbol{x}, \hat{\boldsymbol{a}}) - \pi^{Mech}(\boldsymbol{x}, \tilde{\boldsymbol{a}})\right] \mathbb{I}(y_R \geq y_{R+1}) \\
&= \left[\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{hh}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{hl}})\right] \mathbb{I}(y_R \geq y_{R+1}), \quad \text{and} \\
\text{diversifying downside} &\overset{def}{=} \left[\pi^{Mech}(\boldsymbol{x}, \tilde{\boldsymbol{a}}) - \pi^{Mech}(\boldsymbol{x}, \hat{\boldsymbol{a}})\right] \mathbb{I}(y_R < y_{R+1}) \\
&= \left[\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{lh}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{ll}})\right] \mathbb{I}(y_R \geq y_{R+1}).
\end{aligned}$$

When regions R and $R+1$ are symmetric, we have

$$\begin{aligned} \text{expected diversification upside} &= \frac{1}{2} \underset{(\boldsymbol{x},\boldsymbol{y})}{E} \left[\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{hh}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{hl}}) \right], \text{ and} \\ \text{expected diversification downside} &= \frac{1}{2} \underset{(\boldsymbol{x},\boldsymbol{y})}{E} \left[\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{lh}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{ll}}) \right]. \end{aligned}$$

Definition 1 *In comparing the R-region strategy with the $(R+1)$-region strategy, given the auction mechanism $Mech$ and the realized costs $(\boldsymbol{x}, \boldsymbol{y})$, we call*

$$\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{hh}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{hl}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{lh}}) + \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a^{ll}})$$

the diversification tradeoff of a buyer considering switching from the R-region strategy to the $(R+1)$-region strategy.

Apparently, the buyer prefers the $(R+1)$-region strategy if the expected diversification tradeoff is positive; she prefers the R-region strategy, otherwise. In general, however, the buyer's preference is not trivial because the diversification tradeoff on a sample path can be positive or negative, depending on $(\boldsymbol{x}, \boldsymbol{y})$, and hence the buyer's preference between a more and a less diversified supply base depends on the supplier production cost distribution F, the regional cost distribution G, and the auction mechanism.

However, noticing that $\boldsymbol{a^{hh}} = \boldsymbol{a^{hl}} \vee \boldsymbol{a^{lh}}$ and $\boldsymbol{a^{ll}} = \boldsymbol{a^{hl}} \wedge \boldsymbol{a^{lh}}$, we can prove that the diversification tradeoff is always (i.e., regardless of the realized costs $\boldsymbol{x}$ or $\boldsymbol{y}$) non-positive/non-negative when the buyer's per-period cost function $\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a})$ is submodular/supermodular in $\boldsymbol{a}$ for all $\boldsymbol{x}$, per the definitions of submodular and supermodular functions (see, e.g., p.43 of Topkis 1998). Formally, we have the following lemma.

Lemma 4 *If $\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a})$ is supermodular in $\boldsymbol{a}$ for all $\boldsymbol{x}$, then the buyer always prefers the $(R+1)$-region strategy to the R-region strategy, which in turn implies that the fully diversifying strategy is optimal. If $\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a})$ is submodular in $\boldsymbol{a}$ for any fixed $\boldsymbol{x}$, the converse is true, which in turn implies that the pooling strategy is optimal.*

In other words, Lemma 4 provides a tractable shortcut to the buyer's optimal supply base design problem: Instead of comparing diversification strategies after computing the buyer's expected total procurement cost under each possible supply base design strategy — which in general is technically intractable as we mentioned — we can possibly find the optimal strategy by examining the super- or submodularity of the per-period total cost function $\pi^{Mech}(\cdot)$. Using this approach, we thus next explore the optimal supply base design strategy and the effect of buyer's bargaining power.

4.4.2 Optimal Supply Base Design Strategy For A Buyer with Full Bargaining Power

Per (4.1a), we have $\pi^{OPT}(\boldsymbol{x}, \boldsymbol{a}) = \min_{i=1,\ldots,N}\{\psi(x_i) + a_i\}$, which implies that, for any vector of supplier production cost $\boldsymbol{x}$, and any two vectors of suppliers' regional costs $\boldsymbol{a}$ and $\boldsymbol{a'}$, it must be true that

$$\pi^{OPT}(\boldsymbol{x}, \boldsymbol{a}\wedge\boldsymbol{a'}) = \pi^{OPT}(\boldsymbol{x}, \boldsymbol{a})\wedge\pi^{OPT}(\boldsymbol{x}, \boldsymbol{a'}) \text{ and } \pi^{OPT}(\boldsymbol{x}, \boldsymbol{a}\vee\boldsymbol{a'}) \geq \pi^{OPT}(\boldsymbol{x}, \boldsymbol{a})\vee\pi^{OPT}(\boldsymbol{x}, \boldsymbol{a'}).$$

This in turn implies that $\pi^{OPT}(\boldsymbol{x}, \boldsymbol{a})$ is supermodular in $\boldsymbol{a}$ for any $\boldsymbol{x}$. Therefore, from Lemma 4, we obtain the optimal supply base design strategy under mechanism OPT, as stated in the following proposition.

Proposition 13 *For any F, G, or N, it is always optimal to fully diversify if mechanism OPT is used.*

Proposition 13 highlights a remarkably general result: Whenever the buyer has the power to use the optimal mechanism, it is optimal to fully diversify the supply base, *regardless* of the number of suppliers N, or the cost distributions F and G. This is because, although the buyer's expected total procurement cost under any supply base design strategy in general depends on N, F, and G, the diversification tradeoff (per Definition 1) of a buyer using OPT is non-negative for *all* $\boldsymbol{x}$, $\boldsymbol{y}$, and $R < N$.

To provide intuition for Proposition 13, we will use an example to illustrate why mechanism OPT allows the buyer to enjoy the benefits of diversification, and how OPT functions. Later, we will use these illustrations as a point of contrast to what happens when the buyer has zero bargaining power and uses mechanism RE. For simplicity, our example will assume that regional costs follow a two-point distribution, and are either high, y_H, or low, y_L. The optimal mechanism involves the virtual cost function $\psi(\cdot)$, and for convenience we assume suppliers' production costs are uniformly distributed, making the virtual cost function linear (namely, $\psi(x) = 2x$).

With this setup, we examine a two-supplier case for which the pooling strategy has two suppliers in region 1 and the diversifying strategy has suppliers 1 and 2 in regions 1 and 2, respectively. Figure 4.1(a) pictorially illustrates the diversification upside and diversification downside for a particular pair of supplier production cost realizations x_1 and x_2. Because the function of a reserve price is straightforward, the figure depicts cost realizations for which a supplier with a regional cost advantage will win and receive a payment set by his competitor's dropout bid rather than via the reserve price.[3] In this discussion we assume OPT is implemented with the auction format described in Proposition 12.

Example 1:

- The top panel of Figure 4.1(a) depicts the diversification upside, which occurs when $y_1 = y_H > y_L = y_2$. Had the pooling strategy been used, both suppliers would have the high regional cost, supplier 1 would drop out when the calling price reached his true adjusted virtual cost $\psi(x_1)+y_H$ and supplier 2 would win the auction and be paid $\psi^{-1}(\psi(x_1)+y_H-y_H) = x_1$, yielding a total procurement cost x_1+y_H to the buyer. In contrast, had the diversifying strategy been used, the buyer's supply base would have one supplier (supplier 2) with the low regional cost. In such a case, mechanism OPT would capture the cost reduction opportunity by awarding the contract to supplier 2 (so the buyer incurs a low regional cost) and paying supplier 2 $\psi^{-1}(\psi(x_1) + y_H -$

[3]In particular, Figure 4.1(a) assumes $x_1 > x_2$ and $\psi(x_1)+y_L < \psi(x_2)+y_H$, and $y_H - y_L < \max_{i=1,2}\{2-2x_i\}$.

$y_L) = x_1 + \frac{y_H - y_L}{2}$. Consequently, the buyer pockets a diversification upside equal to $(x_1 + y_H) - [(x_1 + \frac{y_H - y_L}{2}) + y_L] = \frac{y_H - y_L}{2}$.

- The lower panel of Figure 4.1(a) depicts the diversification downside, which occurs when $y_1 = y_L < y_H = y_2$. Had the pooling strategy been used, both suppliers would have low regional costs, and supplier 2 would win the auction and be paid $\psi^{-1}(\psi(x_1) + y_L - y_L) = x_1$, yielding a total procurement cost of $x_1 + y_L$ to the buyer. However, had the diversifying strategy been used, supplier 1 would be the only supplier with a low regional cost. Mechanism OPT would award the contract to the low-regional-cost supplier 1 and pay him $\psi^{-1}(\psi(x_2) + y_H - y_L) = x_2 + \frac{y_H - y_L}{2}$. Thus the diversification downside equals $[(x_2 + \frac{y_H - y_L}{2}) + y_L] - (x_1 + y_L) = \frac{y_H - y_L}{2} - (x_1 - x_2)$.

Note that in this example, the diversification upside exceeds the diversification downside. Because symmetry implies that either occurs with equal probability, this example confirms that for these realizations of x_1 and x_2 the buyer always benefits from diversifying.

While this example applied to a particular set of assumptions on F, G, N, and realizations of x_1 and x_2, the important takeaway is that mechanism OPT helps the buyer capture surplus from a supplier enjoying a regional cost advantage because OPT's rules bias against such a supplier. For example, consider the outcome for the upper-right part of Figure 4.1(a): Supplier 2 is paid $x_1 + \frac{y_H - y_L}{2}$, which is actually $\frac{y_H - y_L}{2}$ dollars less than the lowest total cost the buyer could possibly incur if she transacted with supplier 1. The buyer gets away with this by promising ex ante to compare suppliers' virtual costs, not actual costs, when determining the auction winner and payment. This biases against the advantaged supplier. In particular, when supplier 2 enjoys a regional cost advantage, he only wins the auction if

$$x_2 \in \{x_2 | \psi(x_2) + y_L \leq \psi(x_1) + y_H\} \subset \{x_2 | x_2 + y_L \leq x_1 + y_H\}. \tag{4.2}$$

In summary, mechanism OPT biases against advantaged suppliers in order to reduce their payment, and in doing so might impose an inefficient allocation (evidenced by the proper

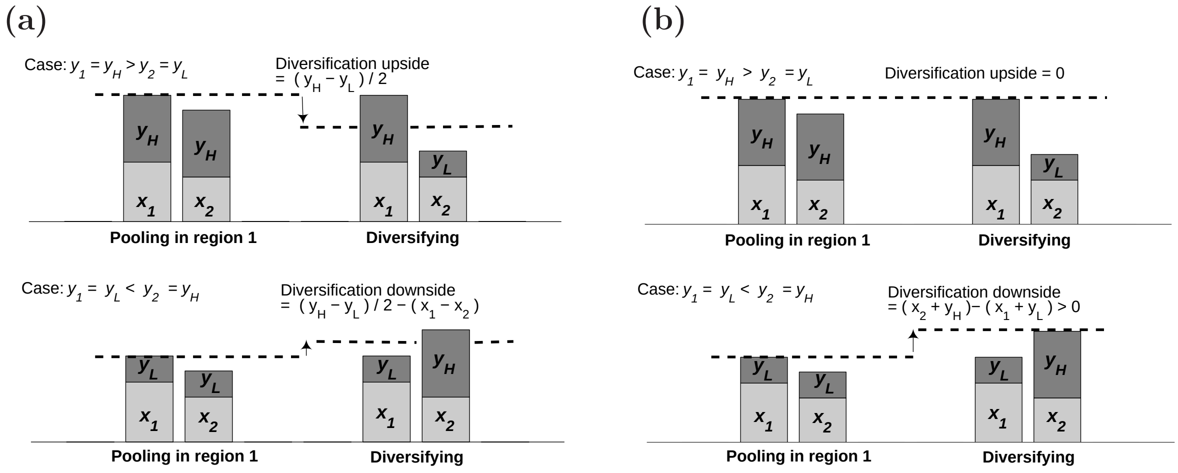

Figure 4.1: The diversification upside and downside. Panel (a) plots for mechanism OPT with $N = 2$ suppliers, assuming $F \sim U[0, 1]$; Panel (b) plots for mechanism RE with $N = 2$ suppliers.

subset relation in (4.2); see also MacAfee and McMillan 1989, Rezende 2009). The upshot is that, because diversifying engenders cost realizations in which suppliers can enjoy a regional cost advantage and mechanism OPT allows the buyer to capitalize on the resulting cost-saving opportunities, the buyer finds it optimal to fully diversify her supply base.

Despite being theoretically optimal, mechanism OPT may be difficult to implement in practice. First, it requires the buyer to impose take-it-or-leave-it allocation and payment rules that bias against suppliers with a cost advantage. The buyer may have a difficult time convincing suppliers to go along with such a scheme, who might not understand why they should be put at a disadvantage even though they have a low regional cost that is attractive for the buyer. In such a case, the optimal mechanism may be off the table and the buyer might have to employ another mechanism which does not require her to exert bargaining power over the suppliers. This motivates our analysis using mechanism RE in the next subsection.

4.4.3 Optimal Supply Base Design Strategy For A Buyer with Zero Bargaining Power

We now examine the cases in which the buyer uses a reverse English auction with no reserve price (zero bargaining power). We first study the case in which the buyer designs a supply base with two suppliers, and then study the general case in which the supply base consists of $N \geq 3$ suppliers.

Two suppliers

Per (4.1b), we have $\pi^{RE}(\boldsymbol{x}, \boldsymbol{a}) = \max\{x_1 + a_1, x_2 + a_2\}$ when $N = 2$, which implies that, for any vector of supplier production costs $\boldsymbol{x}$, and any two vectors of suppliers' regional costs $\boldsymbol{a}$ and $\boldsymbol{a'}$, it must be true that

$$\pi^{RE}(\boldsymbol{x}, \boldsymbol{a} \vee \boldsymbol{a'}) = \pi^{RE}(\boldsymbol{x}, \boldsymbol{a}) \vee \pi^{RE}(\boldsymbol{x}, \boldsymbol{a'}) \text{ and } \pi^{RE}(\boldsymbol{x}, \boldsymbol{a} \wedge \boldsymbol{a'}) \leq \pi^{RE}(\boldsymbol{x}, \boldsymbol{a}) \wedge \pi^{RE}(\boldsymbol{x}, \boldsymbol{a'}).$$

This in turn implies that $\pi^{RE}(\boldsymbol{x}, \boldsymbol{a})$ is submodular in $\boldsymbol{a}$ for all $\boldsymbol{x}$. Therefore, from Lemma 4, we obtain the optimal supply base design strategy under mechanism RE when $N = 2$, as stated in the following proposition.

Proposition 14 *With two suppliers, for any F or G, it is always optimal to pool if mechanism RE is used.*

Surprisingly, with two suppliers, Proposition 14 shows that, rather than diversifying the supply base, the buyer prefers to select the two suppliers from the same region if she has no bargaining power (i.e., uses reverse English auctions without reserve price). Perhaps more surprising, this preference persists for any supplier cost distribution and any regional cost distribution. Why is it never optimal to spread out the regional cost risk by diversification when the buyer uses mechanism RE, even if large supply shocks are very likely? This is because without a reserve price mechanism RE fully exposes the buyer to windfall profit-taking by the advantaged supplier with lower regional cost, who largely absorbs what would have been the buyer's upside benefit of diversifying. Such windfall profit-taking by the advantaged supplier is so severe that the buyer always has no diversification upside. We illustrate this with an example. As a point of contrast to mechanism OPT, we use the same setup as for Example 1, but apply mechanism RE instead of mechanism OPT. The diversification upside and downside are illustrated in Figure 1(b)'s top and bottom panels, respectively. [4]

Example 2:

- The top panel of Figure 1(b) illustrates the diversification upside, which occurs when $y_1 = y_H > y_L = y_2$. Had the pooling strategy been used, both suppliers would have high regional cost, and the lowest total-cost supplier (supplier 2) would win the auction and be paid supplier 1's total cost minus supplier 2's regional cost, i.e., $x_1 + y_H - y_H = x_1$. Thus, the buyer's total procurement cost would be the largest total cost, i.e., $x_1 + y_H$. In contrast, had the diversifying strategy been used, supplier 2 would have low regional cost. However, this does not mean the buyer will get any benefit from having such a low-regional-cost supplier. On the contrary, supplier 2, with lower total cost, would win the auction but charge price $x_1 + y_H - y_L$, matching supplier 1's total cost and yielding a total procurement cost $x_1 + y_H$ to the buyer. In other words, mechanism RE would allow the advantaged supplier 2 to fully absorb the diversification benefit, leaving zero diversification upside to the buyer.

- The lower panel of Figure 1(b) illustrates the diversification downside, which occurs when $y_1 = y_L < y_H = y_2$. Had the pooling strategy been used, both suppliers would have low regional costs, and supplier 2, with lower total cost, would win the auction and be paid supplier 1's total cost minus supplier 2's regional cost, i.e., $x_1 + y_L - y_L = x_1$. Thus, the buyer's total procurement cost would be the largest total cost, i.e., $x_1 + y_L$. In contrast, had the diversifying strategy been used, only supplier 1 would have low regional cost. Supplier 1 would charge price $x_2 + y_H - y_L$, matching supplier 2's total

[4] Figure 4.1(b) assumes $x_1 > x_2$ and $x_1 + y_L < x_2 + y_H$.

cost and yielding a total procurement cost of $x_2 + y_H$ to the buyer. In other words, mechanism RE would allow the advantaged supplier 1 to fully absorb the benefit of its regional cost advantage, saddling the buyer with a large diversification downside.

In summary, mechanism RE allows the supplier with lower regional cost to take so much windfall profit that diversifying yields no diversification benefit but effectively exposes the buyer to more regional cost risk. Consequently, the buyer with no bargaining power is worse off by using the diversifying strategy, and hence would rather pool the two suppliers in the same region.

As Propositions 13-14 reveal, the optimal supply base designs under OPT and RE are polar opposites. This encapsulates our fundamental message in this paper: Bargaining power is a key driver of supply base diversification decisions. The buyer's bargaining power (i.e., the auction mechanism she is able to deploy) determines how much diversification benefit she can pocket, which then informs her decision to diversify the supply base or not. A weak buyer who foresees not being able to pocket the benefits of diversification should take this into account and design her supply base with less diversification than she would if she held full bargaining power over suppliers. This key finding is confirmed in our following analysis of the case in which the buyer uses mechanism RE with $N \geq 3$ suppliers.

More than two suppliers

The goal of this subsubsection is to show that a buyer with zero bargaining power (RE) and three or more suppliers finds that, in many cases, it is suboptimal to fully diversify. This is in stark contrast to the strategy of always fully diversifying, which is optimal for a buyer with full bargaining power (OPT). Thus, this subsubsection reenforces the main message of the paper: The buyer's bargaining power can drastically affect her optimal supply base design.

With three or more suppliers, the optimal diversification strategy under mechanism RE turns out to be much more complicated than that with two suppliers, for two reasons. First, as discussed earlier, the buyer's total procurement cost $\pi^{RE}(\boldsymbol{x}, \boldsymbol{a})$ is the second-lowest order statistic of the possibly correlated total costs of N suppliers. Second, this total procurement cost can easily be shown to be neither submodular nor supermodular in $\boldsymbol{a}$ for all $\boldsymbol{x}$ — hence, the buyer's preference between the R-region strategy and the $(R+1)$-region strategy in general depends on N, R, and the cost distributions F and G.

To gain insight into the buyer's supply base design preference, we will characterize how the buyer's preference is affected by the shape and scale of the regional cost distribution G, given an N and supplier production cost distribution F. For any distribution G, we accomplish this by examining a family of models $\{(F, G^{(s)}), s \in \mathbb{R}\}$, where $G^{(s)}(y) \stackrel{def}{=} G(\frac{y}{s})$. Regional cost distribution $G^{(s)}$ has the same "shape" as G, but a different scale. We call s the scale parameter. Since we have normalized the range of F to the unit interval, this sequence of models captures an increasing variation of the regional cost distribution relative to that of the supplier production cost distribution.

Large regional costs drive preference away from pooling. Section 4.4.3 shows that with two suppliers a buyer using mechanism RE always finds it optimal to pool. Following that intuition, does the buyer always prefer to pool even with three or more suppliers? Here we show that the answer is "no." In particular, when the scale parameter s is large, regional cost variation dominates suppliers' production cost variation. This can model, for example, cases in which suppliers use standard production technology and the variability of production costs is negligible in comparison to that of regional costs — which could be driven by a variety of sources ranging from incremental transportation rate changes to catastrophic port strikes. When the regional cost variation is relatively large, pooling and "losing big" (as described on page 98) could be catastrophic for the buyer. (For example, if a strike hit the originating port in a region containing all N of her suppliers.) Intuitively the buyer might want to diversify her supply base, but then again we recall the $N = 2$ case for which we know that pooling is optimal due to supplier windfall-profit taking. How can the buyer benefit from diversifying when using mechanism RE? The key is that, with $N \geq 3$ suppliers, the buyer can partially diversify by grouping suppliers together into different regions. This curbs unilateral regional cost advantages and forestalls windfall-profit taking. Suppliers in a low-cost region will — in the course of competing for the buyer's business in the auction — transfer the surplus of their regional cost advantage to the buyer. This result is formalized in Proposition 15 below, which shows that, in fact, the buyer would always prefer to have at least $R = \lfloor \frac{N}{2} \rfloor$ regions (with at least two suppliers per region) to any less diversified strategy with $R < \lfloor \frac{N}{2} \rfloor$ regions. Therefore, we see that the number of suppliers can affect the supply base design strategy under mechanism RE, and the pooling strategy can dominated by the partially diversifying strategy.

Preference between full and partial pooling driven by shape of regional cost distribution. We now turn to the main goal of this subsubsection, which is to show that a buyer using mechanism RE with three or more suppliers need not find fully diversifying optimal. In particular, we show that the buyer prefers the partially diversifying strategy to the fully diversifying strategy when the scale parameter s is large and G "has a low-cost tail" (roughly speaking, this will be made more precise shortly). These conditions make windfall-profit taking a serious concern for the buyer. When s is big, regional costs largely determine the auction winner. Furthermore, when G has a low-cost tail there is more chance for low "outlier" regional costs. Hence, fully diversifying under these conditions is apt to backfire by yielding a winning supplier with a sizeable, unilateral regional cost advantage that he absorbs through windfall-profit taking: Even though the buyer has suppliers in N regions, she incurs costs as if she contracts with a supplier in the second-cheapest of N regions. On the other than, if the buyer chose to forestall windfall profit-taking by partially diversifying, then she would have two suppliers in each of $\lfloor \frac{N}{2} \rfloor$ regions, and incur costs associated with the cheapest of $\lfloor \frac{N}{2} \rfloor$ regions. After netting out the production cost, the buyer's preference between partial and full diversification depends on the relative sizes of $\overline{Y}_{2:N}$ and $\overline{Y}_{1:\lfloor \frac{N}{2} \rfloor}$, which in turn is driven by the shape of G.

The following proposition summarizes the results so far in this section.

Proposition 15 *When the scale parameter s of the regional cost distribution is large enough:*

- *The partially diversifying strategy that has $\lfloor \frac{N}{2} \rfloor$ regions with at least two suppliers each dominates any strategy with fewer regions;*
- *The buyer prefers the partially diversifying strategy to the fully diversifying strategy if $\overline{Y}_{2:N} > \overline{Y}_{1:\lfloor \frac{N}{2} \rfloor}$; vice versa.*

Given the above discussion, it is natural to ask what distributions G have a "low-cost tail" shape that causes the buyer to prefer partial diversification. More precisely, in light of Proposition 15 we can equivalently ask for what distributions G will $\overline{Y}_{2:N}$ exceed $\overline{Y}_{1:\lfloor \frac{N}{2} \rfloor}$. One can show (see the e-companion) that

$$\begin{aligned} \overline{Y}_{1:\lfloor \frac{N}{2} \rfloor} - \overline{Y}_{2:N} &= \int_{-\infty}^{\infty} \bar{G}^{\lfloor \frac{N}{2} \rfloor}(y)dy - \int_{-\infty}^{\infty} \left[N\bar{G}^{N-1}(y) - (N-1)\bar{G}^{N}(y)\right] dy \qquad (4.3) \\ &= \int_0^1 \left[z^{\lfloor \frac{N}{2} \rfloor} - Nz^{N-1} + (N-1)z^N\right] \frac{1}{g[\bar{G}^{-1}(z)]} dz, \end{aligned}$$

where $\bar{G}(y) \stackrel{def}{=} 1 - G(y)$ denotes the tail distribution of regional costs and $\bar{G}^{-1}$ is the inverse function of $\bar{G}$. Note that there exists some $0 < \underline{z} < 1$ such that $z^{\lfloor \frac{N}{2} \rfloor} - nz^{n-1} + (n-1)z^n$ is positive (negative) for all $0 < z < \underline{z}$ ($\underline{z} < z < 1$). Therefore, if the distribution G has most of its density piled close to the left end-point of the support (the low-cost end-point), then $g[\bar{G}^{-1}(z)]$ is large for z close to one but is small for z close to zero, and thus $\overline{Y}_{1:\lfloor \frac{N}{2} \rfloor} - \overline{Y}_{2:N}$ must be positive, which implies that fully diversifying is preferred. In contrast, if G has most of its density close to the right end-point (the high-cost end-point) of the support (and so has a "low-cost tail"), then $g(\bar{G}^{-1}(z))$ is large for z close to zero but is small for z close to one, and thus $\overline{Y}_{1:\lfloor \frac{N}{2} \rfloor} - \overline{Y}_{2:N}$ must be negative, which implies that the partially diversifying strategy is preferred. To further illustrate this (i.e., the shape of G drives the buyer's preference between the two strategies), we consider the power-function family $G^{(s)}(y) = (\frac{y}{s})^v$ with a shape parameter $v > 0$ on the support $[0, s]$, because such a distribution's density function takes various shapes according to the shape parameter $v > 0$: As illustrated in

Figure 4.3(a), the density distribution is very concentrated near the left endpoint when v is small; in contrast, the density distribution flattens out as v increases. Consistent to the above discussion on the shape of distribution G, we find that the buyer prefers fully diversifying when v is small but prefers the partially diversifying strategy otherwise.

Corollary 1 *Suppose* $G^{(s)}(y) = (\frac{y}{s})^v$. *When the scale parameter* s *is large enough, there exists a threshold* $\underline{v}$ *such that the buyer prefers fully diversifying if* $v < \underline{v}$*, and the partially diversifying if* $v > \underline{v}$.

The insight to Corollary 1 is consistent to that to Propositions 13 and 14. Whereas diversification gives the buyer better chance to access to a supplier with low regional cost, the supplier is likely to take a windfall profit by charging a high price and thus absorb the buyer's diversification benefit if the supplier's competitors all have higher regional costs and the buyer has little bargaining power against this advantaged supplier, leaving the buyer with only the diversification downside. Such an advantage supplier is likely to emerge and enjoy a large advantage over competitors unless the regional cost distribution G has density piled up at the low-cost end. That's why the buyer prefers the partially diversifying strategy under RE unless G has density piled up at the low-cost end, i.e., the shape parameter v is small if G is a power-function distribution. Intuitively, these insights would not be sensitive to the scale parameter s, i.e., the relative variation of the regional costs compared to that of the production costs; in fact, our numerical examinations in §4.6 will show that Proposition 15 and Corollary 1 hold even when s is small, i.e., when the regional cost variation is comparable to the production cost variation.

4.4.4 General Takeaways

Sections 4.4.2–4.4.3 suggest that the buyer should carefully evaluate her bargaining clout relative to the suppliers before deciding to diversify her supply base. Buyers with strong bargaining power always find it optimal to have just one supplier per region (fully diversifying); this is because diversifying mitigates the exposure to regional cost shocks, while such buyers can use reserve prices and biasing rules to prevent suppliers from absorbing the benefits of diversification and hence capture significant benefits from having a supplier in a low-cost region. In contrast, for buyers who have extremely little channel power and cannot prevent advantaged suppliers from making significant windfall profits, there is little benefit to diversifying and they should only add more regions to their portfolio if doing so is unlikely to forfeit the benefits of diversification to an advantaged supplier. Because a supplier with a sizeable regional cost advantage can largely absorb the resulting benefits if he is the only supplier in his region, the buyer can avoid the emergence of such an advantaged supplier by keeping two suppliers in each region (partially diversifying). In particular, for a buyer who uses a two-supplier supply base, there is little benefit to diversifying and she should always instead simply pool her risk by choosing both suppliers in a single region. While such a buyer is inevitably more vulnerable to cost shocks, by pooling she ensures greater cost parity between suppliers, which she needs to drive down suppliers' price bids. A buyer with weak bargaining power who uses more than two suppliers will diversify a two-supplier region into two single-supplier regions only if neither region is likely to contain an advantaged supplier that can absorb the benefits of the diversification. Because the benefit of this diversification accrues precisely when just one of the two regions experiences a large regional cost, this pushes the buyer to prefer diversifying only if, in such a case, she has supplier(s) in a third region which is unlikely to experience a large regional cost. Consequently, buyers with weak bargaining power prefer to fully diversify only if there are at least three suppliers in the supply base and the regional cost distribution does not have a left (low-cost) tail.

4.5. Extensions

4.5.1 Codependent Regional Costs

Our analyses in §4.4 assumed independent regional costs. However, one can easily imagine situations where regions exhibit vulnerability to common shock factors, such as the price of oil or global shipping volumes. This subsection thus examines the results of Propositions 13-15 for the case with ex ante symmetric and codependent regions. We first discuss how the result extend, and then conclude this subsection with a proposition formalizing this discussion.

Regional codependence does not affect buyer's preference for diversifying with OPT. Propositions 13 revealed that the buyer prefers to fully diversify no matter what the production cost and regional cost distributions are. The intuition is that the buyer diversifies in order to enjoy the upside benefit of finding suppliers with attractive regional costs. Surprisingly, this general result extends even when the regional costs can be codependent (technically, this follows from the fact that Lemma 4 holds even when regional costs are codependent). Thus, even if regions' costs are correlated, the buyer still places exactly one supplier in each region in order to diversify her regional cost risk as much as possible.

Regional codependence affects the buyer's preference under RE only when she has three or more suppliers. Because Lemma 4 remains valid in the presence of regional codependence, Proposition 14 extends to codependent regions and consequently a buyer using mechanism RE always prefers to forestall windfall-profit taking by pooling when she has just two suppliers. Diversifying leaves the buyer with downside risk but no upside benefit, and this remains true no matter how small the windfall-profit taking risk is: Regardless of how highly positively correlated the regional costs are, the buyer's preference for pooling is unchanged. On the other hand, when the buyer has three or more suppliers and uses mechanism RE, regional codependence can affect her supply base design preference. To see why, consider correlated, random regional costs $(y_1, y_2, \ldots, y_R)$ for which there exists a random state variable ξ having distribution $P(\cdot)$ such that the y_r's are independently distributed according to some distribution function $G(\cdot|\xi)$. (In the terminology of Shaked (1977), these random variables are "positive dependent by mixture.") Using §4.4.3's results for the case with identically and independently distributed regional costs (Proposition 15), we can see that the effect of regional codependence on the diversification strategy thus depends on both $P(\cdot)$ and the family of distributions $G(\cdot|\xi)$. (Unless of course the family $G(\cdot|\xi)$ is such that $\overline{Y}_{1:\lfloor \frac{N}{2} \rfloor}$ never exceeds $\overline{Y}_{2:N}$ [or vice-versa] — in such cases the codependence does not affect the buyer's preference between fully and partially diversifying.) For example, consider $G(y|\xi) = s^{-\xi}y^{\xi}$ with a given $s > 0$. In such a case, a small value of state variable ξ can model the case in which regional costs are mainly driven by global factors and exhibit small variance; in contrast, a large ξ can model the case in which regional costs are mainly driven by local factors and exhibit large variance. According to Corollary 1, it is clear that the buyer prefers partially diversifying if the state variable distribution $P(\xi)$ has enough of its density concentrated at large values of ξ, but would instead prefer partially pooling were the density concentrated primarily at small values of ξ.

The following proposition formalizes the discussion in this section.

Proposition 16 *With codependent, ex ante symmetric regions,*

- *Under mechanism OPT, the buyer's preference is robust to regional codependence and she always finds it optimal to fully diversify her supply base.*
- *Under mechanism RE,*
  - *with two suppliers the buyer's preference is robust to regional codependence and she always finds it optimal to pool her supply base;*
  - *with three or more suppliers, regional codependence can affect the buyer's preference between fully diversifying and partially diversifying.*

4.5.2 Asymmetric Regions

We have thus far assumed that all regions are symmetric; this assumption has enabled us to shed some light to the in general challenging problem of supply diversification and focus on the effect of the buyer's bargaining power. In this subsection, we examine the cases in which different types of regions may have different supplier production cost distributions and regional cost distributions. We show that Propositions 13 and 14 remain valid even when there are asymmetric regions, as long as each region type has enough "copies" of regions. Relaxing this assumption, however, may push the buyer away from fully diversifying; we provide an example at the end of this subsection.

We use an index k for a region type. For a type k region, let F_k denote the supplier production cost distribution and let G_k denote the regional cost distribution. For a supplier in a type k region, let H_k denote the total cost distribution and let $\hat{H}_k$ denote the adjusted virtual cost distribution. Namely, $H_k \stackrel{def}{=} F_k \oplus G_k$, where the convolution operator $\oplus$ is such that $H_k(z) = \int_{-\infty}^{\infty} G(z-x)dF(x)$; $\hat{H}_k \stackrel{def}{=} \hat{F}_k \oplus G_k$, where $\hat{F}_k(z) \stackrel{def}{=} F_k(\psi_i^{-1}(z))$ is the distribution of the virtual cost. We assume that suppliers' production costs and regional costs are all independent.

Proposition 17 *If mechanism OPT is used, it is always optimal to fully diversify even when regions are asymmetric. However, the optimal supply base design strategy may involve different region types.*

Our main insights remain valid under asymmetric regions: With full bargaining power (OPT), it is optimal to fully diversify the supply base. This is because any R-strategy having $n_r \geq 2$ suppliers in some region r is dominated by the $(R+1)$-strategy strategy having $n_r - 1$ suppliers in region r and one supplier in the region $(R+1)$ which is of the same type as region r. However, the way in which the buyer chooses to fully diversify the supply base may be more nuance than that in the case with symmetric regions. In particular, the buyer might use multiple types of regions when fully diversifying the supply base. For example, consider a two-supplier case with two types of regions, denoted by $k = 1, 2$. Suppose that $\hat{H}_1$ is the two-point distribution with probability mass q at zero and probability mass $1-q$ at one, and that $\hat{H}_2$ is the uniform distribution $U[0,1]$. It is easy to check the following are true. If the buyer uses two type 1 regions, the buyer's expected total cost equals $(1-q)^2$; if the buyer uses two type 2 regions, the buyer's expected total cost equals $\frac{1}{3}$; if the buyer uses one type 1 and one type 2 region, the buyer's expected total cost equals $\frac{(1-q)}{2}$. The "mixed diversification strategy" that uses both region types is optimal when $q > \frac{1}{3}$. Intuitively, in such cases type 2 regions are less "risky" but a second type 2 region has a decreasing effect on reducing the buyer's expected total cost. As a result, the buyer can be better off using one type 2 region as a "safety" and then gambling on a "bet" by having one type 1 region.

In summary, our main message of the paper remains under asymmetric regions: With more bargaining power, the buyer prefers a more diversified supply base. However, the ability to mix-and-match asymmetric regions adds complexity in determining how to best execute the fully diversifying strategy — a complex question that we leave to future work.

With asymmetric regions, Proposition 14 still holds. In other words, it is never optimal to select two suppliers from two asymmetric regions under mechanism RE. Intuitively, any two types of regions have a "better" type, and hence using two regions of different types is dominated by using two regions of the "better" type, which is in turn dominated by selecting two suppliers from one region of the "better" type according to Proposition 14. Let $\tilde{H}_k \stackrel{def}{=} F_k^2 \oplus G_k$; namely, it denotes the distribution of the buyer's total procurement cost if she has two suppliers in a type k region. We have the following proposition.

Proposition 18 *With two suppliers, it is always optimal to pool if mechanism RE is used, even when regions are asymmetric. In particular, it is optimal to choose a type k region such that $\int_{-\infty}^{\infty} z d\tilde{H}_k$ is minimized.*

Proposition 18 generates managerial insights for buyers who uses two suppliers and mechanism RE. Consider such a buyer who has access to suppliers in both a local region, denoted by $k = 1$, and a foreign region, denoted by $k = 2$. Suppose that F_2 is stochastically smaller than F_1, and that G_1 is stochastically smaller than G_2, which models that foreign suppliers have stochastically lower production costs but the buyer incurs higher transportation cost and risk if using foreign suppliers. According to Proposition 18, the buyer should choose between "go foreign" (i.e., selecting two supplier from the foreign region) or "stay local" (i.e., selecting two supplier from the local region) by taking a total-cost perspective, that is, comparing region types via the distribution $\tilde{H}_k$, but she should never do a "mix' (i.e., using one local supplier and one foreign supplier).

Apparently, Proposition 18 holds even if each region type has only one "copy" of region since the optimal strategy is after all to use only one region. However, if some "better" region type has limited "copies", instead of fully diversifying, a buyer using mechanism

OPT may select multiple suppliers from such a "better" region. Here is a simple example. Consider a buyer who wants to use two suppliers and can only access to two regions with the same production cost distribution $U[0,1]$ but different regional cost distributions G_1 and G_2, respectively. Suppose that G_k has probability mass p_k at 1 and probability mass $1-p_k$ at 0 and $p_2 > p_1$. It is easy to check that it is optimal to select both suppliers from region 1 if and only if $p_2 > \frac{p_1}{5-4p_1}$.

4.5.3 Optimal Supply Base Design Strategy For A Buyer with Intermediate Bargaining Power

Mechanisms OPT and RE represent the full and zero bargaining power cases, respectively. Using these two mechanisms, our comparisons of the buyer's diversification preference has revealed the following insight: Buyers with more bargaining power favor more diversification. In this subsection we examine whether this insight extends to a third auction mechanism, the reverse English auction with a reserve price (RER). Compared to the zero-bargaining power, mechanism RER adds the power to set a credible reserve price. Thus, our goal is to see whether adding the reserve price to the reverse English auction format will encourage, or discourage, diversification. To make the results comparable with Propositions 13-15, we adopt the same assumptions that suppliers' production costs and regional costs are all independent and regions are symmetric. In the following, we first show that, with two suppliers, the buyer can find it optimal to diversify if using mechanism RER, in contrast to that it is always optimal to pool if using mechanism RE (Proposition 14); we next discuss that, with $N \geq 3$ suppliers, when the scale parameter of the regional cost distribution is large, the buyer finds it always optimal to fully diversify if using mechanism RER, in contrast to that the buyer's preference toward fully diversifying depends on the shape of the regional cost distribution if using mechanism RE (Proposition 15).

With two suppliers, the sign of the buyer's diversification tradeoff (per Definition 1) under mechanism RER can be positive or negative, depending on the realized regional cost difference $|y_1 - y_2|$ and the larger production cost $X_{2:2}$, as stated in the following lemma.

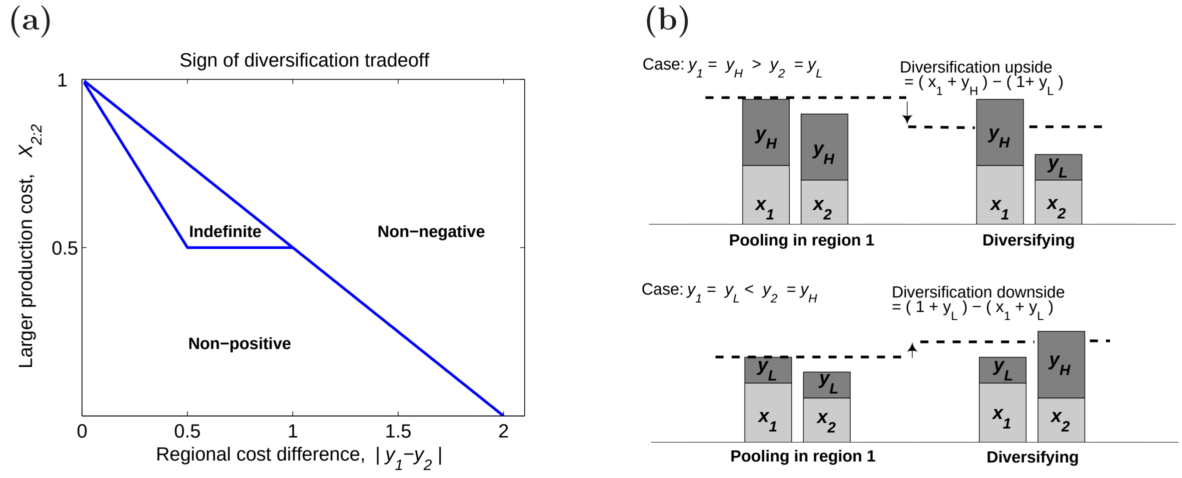

Figure 4.2: Effect of reserve price when $N = 2$. Panel (a) illustrates Lemma 5; panel (b) illustrates the effect of reserve price on the diversification tradeoff.

This is in stark contrast to the always non-positive diversification tradeoff under mechanism RE.

Lemma 5 *If* $|y_1 - y_2| \geq 2(1 - X_{2:2})$, *the buyer's diversification tradeoff (per Definition 1) is non-negative. If* $|y_1 - y_2| \leq 1 - X_{2:2}$, *or if* $|y_1 - y_2| \leq 2(1 - X_{2:2})$ *and* $X_{2:2} \leq \frac{1}{2}$, *the diversification tradeoff is non-positive.*

Figure 4.2(a) illustrates Lemma 5, and Figure 4.2(b) demonstrates the main message of Lemma 5: When the regional cost difference is large (e.g., $|y_1 - y_2| \geq 2 - 2X_{2:2}$), a reserve price effectively limits windfall profit-taking and allows the buyer to capture significant cost savings when sourcing from a low regional cost supplier. Consequently, diversifying makes the buyer better off under mechanism RER. To directly compare with the case under mechanism RE, Figure 4.2(b) shows the diversification upside and downside under mechanism RER for the same setting as in Figure 4.1(b).

Example 3:

As a direct comparison against what would happen under mechanism RE, Figure 4.2(b) shows the diversification upside and downside under mechanism RER for the same setting as in Figure 4.1(b).

- When $y_1 = y_H > y_L = y_2$, had the buyer diversified, although the advantaged supplier 2 would still win the auction (as in Figure 4.1(b)), the reserve price would cap its payment at 1, yielding a total procurement cost $1 + y_L$ to the buyer. Namely, the reserve price would increase the diversification upside to $(x_1 + y_H) - (1 + y_L)$, compared to zero in Figure 4.1(b).

- When $y_1 = y_L < y_H = y_2$, had the buyer diversified, although the advantaged supplier 1 would still win the auction (as in Figure 4.1(b)), the reserve price would again cap its payment at 1, yielding a total procurement cost $1 + y_L$ to the buyer. Namely, the reserve price would decrease the diversification downside to $(1 + y_L) - (x_1 + y_L)$, compared to $(x_2 + y_H) - (x_1 + y_L)$ in Figure 4.1(b).

As a result of effectively truncating large windfall profit opportunities, when the regional cost difference is large (e.g., $|y_1 - y_2| \geq 2 - 2X_{2:2}$), the buyer who can use a reserve price in a reverse English auction has a positive diversification tradeoff equal to $|y_1 - y_2| - (2 - 2X_{2:2})$. However, when the regional cost difference is small, the reserve price is inactivated — in such cases the auction payment is (as in RE) set purely by pricing competition. Thus, even with a reserve price, the buyer could have a negative diversification tradeoff.

Since the buyer prefers diversifying (pooling) if the expected diversification tradeoff is positive (negative), Lemma 5 implies that pooling is optimal only if the realized regional cost difference $|y_1 - y_2|$ and the larger production cost $X_{2:2}$ are both small enough. Formally, we have the following proposition.

Proposition 19 *With two suppliers, if mechanism RER is used,*

- *it is optimal to diversify if the expected regional cost difference $E[|y_1 - y_2|]$ is large enough, and*

- *it is optimal to pool if the probability that $|y_1 - y_2| > 1$ and the probability that $X_{2:2} > \frac{1}{2}$ are both small enough.*

We now turn to the case with three or more suppliers. First, note that when the scale parameter s of the regional cost distribution is large enough, mechanism RER is a good "facsimile" of mechanism OPT: In both cases the winner is likely to be determined by regional costs, while the payment is likely to be determined by the reserve price. Using this intuition, it is possible that when s is large it is optimal to fully diversify under mechanism RER. This contrasts with the buyer's preference under RE, which even with a large s favored partially pooling depending on the shape of the regional cost distribution, G.

4.6. Numerical Illustrations

The purpose of this section is threefold. First, we illustrate the buyer's optimal supply-base strategy under mechanisms OPT and RE in a two-supplier case (Propositions 13-14). Second, we illustrate that regional codependence does not change the optimal strategy under mechanisms OPT and RE in a two-supplier case (Proposition 16) but can affect the magnitude of the buyer's benefit from choosing the "right" strategy. Third, we illustrate the buyer's preference between the partially and the fully diversifying strategy when using mechanism RE and $N \geq 3$ suppliers. In particular, we show that Proposition 15 holds even when the scale parameter s is not very large, i.e., the variation of the regional cost distribution is comparable to that of the supplier production cost.

To measure the relative performance of different supply base design strategies in any given setting (with fixed number of suppliers N, cost distributions F and G, and the auction mechanism), we benchmark all supply base design strategies to the pooling strategy and define for each supply base design strategy

$$\textit{rate of cost improvement} = 1 - \frac{\text{expected total cost under the strategy}}{\text{expected total cost under the pooling strategy}}. \tag{4.4}$$

In this section, we assume the supplier production cost distribution $F \sim U[0,1]$. For the regional cost distribution we use the power-function family $G^{(s)}(y) = (\frac{y}{s})^v$ with a shape parameter v and a scale parameter s on the support $[0, s]$. Figure 4.3(a) illustrates the

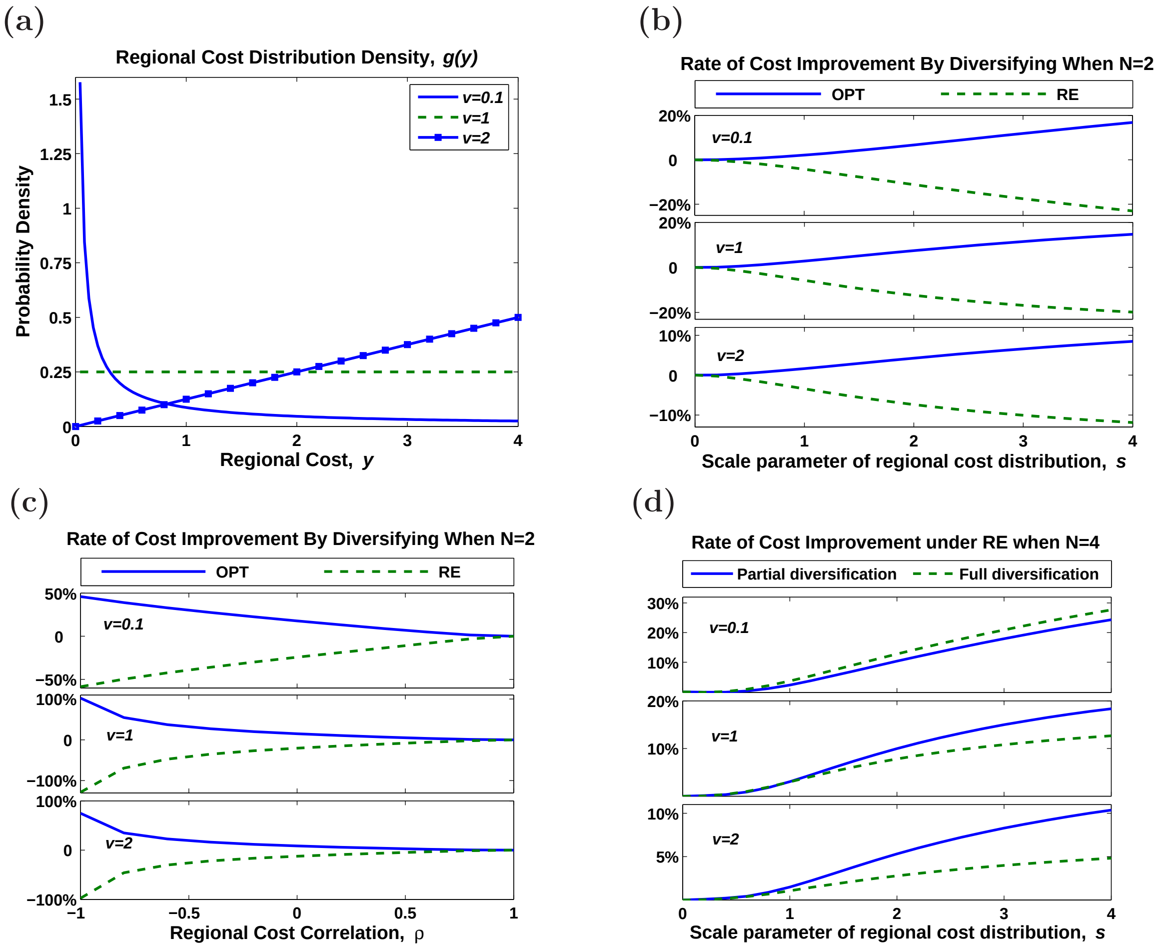

Figure 4.3: Power-function distribution's density and rates of cost improvement per equation (4.4). Panel (a) illustrates the power-function distribution's density, $g(y) = vs^{-v}y^{v-1}$, with scale parameter $s = 4$ and shape parameters $v = 0.1$, 1, 2. Panels (b) and (d) assume that suppliers' production costs follow $F \sim U[0, 1]$ and regional costs follow the power-function distribution with density $g(y) = vs^{-v}y^{v-1}$, $v = 0.1, 1, 2$. Panel (b) plots the rates of cost improvement by diversifying under mechanisms OPT and RE when $N = 2$ against the scale parameter $s = 0 \sim 4$. Panel (c) plots the rates of cost improvement by diversifying under mechanisms OPT and RE when $N = 2$ against the regional cost correlation $\rho = -1 \sim 1$. Panel (d) plots the cost improvement of the partial $(R = 2)$ and the full $(R = 4)$ diversification strategies when $N = 4$ against the scale parameter $s = 0 \sim 4$.

density functions of such distributions with a fixed shape parameter $s = 4$ and three shape parameters $v = 0.1, 1, 2$. As it shows, the power-function distributions have density very concentrated near the left endpoint of the support when v is small, i.e., $v = 0.1$; however, the shape of the density functions flatten out toward the right endpoint of the support when the shape parameter increases, i.e., $v = 1$ and $v = 2$.

Figure 4.3(b) plots the rates of cost improvement by the diversifying strategy under mechanism OPT and RE in a two-supplier case. In all three cases ($v = 0.1, 1, 2$), the rate is positive/negative under mechanism OPT/RE, confirming that diversifying is optimal under mechanism OPT (Proposition 13) and pooling is optimal under mechanism RE (Proposition 14). In all cases, the magnitude of the rate of cost improvement increases as the scale of the regional cost distribution increases. This confirms the intuition that the buyer can be significantly better off by optimizing the supply base design especially when the variation of regional costs is relative large compared to that of supplier production costs.

We examine the effect of regional cost codependence by studying a two-supplier case, in which two regions (regions 1 and 2) have correlated and identically distributed regional costs. In particular, we assume that region 1 has a random cost $y_1^c = \lambda y_1 + \sqrt{1-\lambda^2} y_2$ and region 2 has a random cost $y_2^c = \lambda y_2 + \sqrt{1-\lambda^2} y_1$, where $\lambda = -\frac{\sqrt{2}}{2} \sim \frac{\sqrt{2}}{2}$, and y_1 and y_2 are independent draws from the power-function distribution $G^{(s)}(y) = (\frac{y}{s})^v$ with $v = 0.1, 1, 2$ and $s = 4$. Namely, y_1^c and y_2^c are identically distributed and have correlation $\rho = 2\lambda\sqrt{1-\lambda^2} \in [-1, 1]$. For such a case, Figure 4.3(c) plots the rate of cost improvement by diversifying against the regional cost correlation ρ. It illustrates that in a two-supplier case the regional cost correlation does not change the buyer's preference between pooling and diversifying (Proposition 16): Diversifying is always optimal under mechanism OPT and pooling is always optimal under mechanism RE. Moreover, it also shows that the magnitude of the buyer's benefit from choosing the optimal strategy decreases with the regional cost correlation. Intuitively, the buyer is indifferent between pooling and diversifying when the regional costs are perfectly positively correlated — in such a case, regions are essentially the same region. However, choosing the optimal supply base design strategy is most beneficial when the regional costs are perfectly negatively correlated — the case in which regional cost disparity is most significant. A buyer who uses mechanism OPT can take advantage of such regional cost disparity and hence finds it most beneficial to diversify the supply base; in contrast, a buyer who uses mechanism RE expects most severe windfall profit-taking by advantaged suppliers and consequently finds it most beneficial to pool the two suppliers in the same region.

Figure 4.3(d) plots the rates of cost improvement of the partially diversifying strategy $R = 2$ and the fully diversifying strategy $R = 4$ for a setting in which the buyer uses mechanism RE and $N = 4$ suppliers. It clearly verifies Proposition 15: When the scale parameter s is large enough, first, the pooling strategy is dominated by the partially diversifying strategy; second, the buyer prefers fully diversifying when the shape parameter is small, i.e., $v = 0.1$, but prefers the partially diversifying strategy when the shape parameter is not small, i.e., $v = 1, 2$. Moreover, the figure indicates that Proposition 15 can hold even when the scale parameter $s \simeq 1$, i.e., when the variation of the regional cost distribution is comparable to that of the supplier production cost distribution. However, it also indicates that the performance difference between the partially and fully diversifying strategies is reasonably small when the scale parameter s is small.

4.7. Conclusions

A buyers' total procurement cost includes not only the contract payment to a supplier but also other costs (logistics costs, duties, etc.), which are often subject to various regional cost shocks such as labor strikes, regulation changes, and geopolitical events. To mitigate regional cost risks, a buyer seeking to minimize her total procurement cost can strategically reduce the cost correlation across suppliers by diversifying her supply base (i.e., choosing suppliers from different regions). However, in settings where the buyer's payment to her supplier is determined by a competitive bidding process (i.e., an auction), the buyer's upside benefit of diversification — having a significantly cost-advantaged supplier — can be undermined by this supplier's windfall profit-taking. This paper models the interaction between the buyer's supply-base-design strategy and the risks of windfall profit-taking by suppliers, and characterizes the optimal supply-base-design strategy under various auction mechanisms. To our knowledge, this paper is the first study of supply base design to mitigate regional cost shocks.

We find that the buyer needs to make a tradeoff between the benefit from diversifying her risk of exposure to cost shocks, and the risk that suppliers will absorb such benefits for themselves by taking windfall profits. The ability of suppliers to take windfall profits depends upon the buyer's bargaining power, that is, the buyer's ability to choose an auction mechanism to suppress supplier profits. In particular, at one extreme, when the buyer has full bargaining power and thus can impose the optimal mechanism (i.e., the optimal reserve price plus the optimal contract allocation rule that biases against cost-advantaged suppliers), windfall profit-taking is curbed and consequently the buyer finds it optimal to fully diversify her supply base (i.e., select each supplier from a different region). However, at the other extreme, when the buyer has no bargaining power and solely relies on supplier competition for price concessions (i.e., uses a reverse English auction with no reserve price), supplier windfall profit-taking can be severe and consequently the buyer diversifies less. With two suppliers she always finds it optimal to pool both suppliers in a single region. With more suppliers she prefers a blended strategy: She diversifies by using multiple regions, but keeps two suppliers per region to hedge her bets and eliminate the risk that any supplier possesses a unilateral regional cost advantage. We also study cases where the buyer has intermediate bargaining power and thus can impose a reserve price when using a reverse English auction. We find that imposing a reserve price allows the buyer to truncate large supplier profits, so when the cost shock size is large the buyer prefers fully diversifying when she can use a reserve price. This again highlights that buyers with strong bargaining power prefer to diversify more, while buyers with less bargaining power prefer to diversify less due to concerns about windfall profit-taking.

We find that introducing codependence across regional cost shocks generally leaves the buyer's diversification decision unchanged, and only affects the buyer's decision in cases where there are at least three suppliers and the buyer is susceptible to severe windfall profit-taking by advantaged suppliers (formats RE). In these cases, codependence can encourage or discourage fully diversifying, depending on whether or not it reduces the risk of an advantaged supplier emerging, that is, reduces the risk of windfall profit-taking.

Our study was motivated by focusing on shocks to the buyers' "non-price" costs (logistics costs, tariffs, etc.), but our results can easily be extended to cases where suppliers share regional cost drivers, such as costs associated with a small local labor force, regional energy market, or a common second-tier supply base. More precisely, all analyses in this paper follow if regional cost y_r is re-interpreted as a commonly known cost factor shared by all suppliers within region r.

Although we assume that the distributions capturing production costs and regional costs remain static over time, Propositions 13-14 (and their extensions Propositions 16-18) directly extend to cases where these distributions vary over time, given that these results hold regardless of production cost distribution F or regional cost distribution G. For Corollary 1, which says the optimal strategy depends on the shape of the regional cost distribution G, we suspect that the optimal supply-base design decision depends on the shape of the regional cost distribution G "on average", if G is time-variant.

We examined three auction mechanisms that are of both theoretically and practically important. Of course, buyers may use other auction mechanisms or unstructured bargaining processes — for example, the buyer may use first-price auctions, or may negotiate with the advantaged supplier. For such cases, we suspect that the key insight of our paper will continue to apply: The more bargaining clout the buyer has to control windfall-profit taking by cost-advantaged suppliers, the more she will prefer building a diversified supply base. Our results can also extend to the cases where the buyer has only a limited number of regions to choose suppliers from — suppose there are only $\overline{R} < N$ regions available — the buyer tends to use all $\overline{R}$ regions if fully diversifying is optimal in the unconstrained case, or tends to use $\min\{\overline{R}, \lfloor \frac{N}{2} \rfloor\}$ regions if the partially diversifying strategy is optimal in the unconstrained case. Some other extensions are possible, for example, imposing a fixed cost of using additional regions, or allowing asymmetry among regions. Our preliminary numerical studies suggest that these issues affect the buyer's supply base design decisions in quite straightforward ways, but we leave further examinations for future work.

Finally, to keep our analysis focused and tractable we ignore the buyer's inventory decisions. To the extent that the buyer can anticipate regional cost shocks, she may choose to speculatively purchase inventory to avoid future cost spikes, e.g., impending tariffs increases in a certain country. Interestingly, our analysis suggests that the usefulness of such a strategy depends on the buyer's bargaining clout. Speculative inventory might behoove a buyer with little bargaining clout, who might use it to help avoid paying windfall profits to cost advantaged suppliers. On the other hand, speculative inventory would likely be of much less benefit to a buyer with strong bargaining clout, who could contract with cost-advantaged suppliers without paying an undue price premium, thereby reducing the speculative benefits of holding inventory. We leave a detailed analysis of the interplay between inventory decisions, bargaining power and supply base design to our future work.

4.8. Proofs

4.8.1 Proof of Proposition 12

If a supplier drops out of the auction when the calling price is higher than his true adjusted virtual cost and there is at least one other supplier staying in the auction, the supplier loses the auction and gets zero profit. The buyer can possibly win the auction and get a positive profit by staying in the auction until the calling price reaches his true adjusted virtual cost. Thus, dropping out of the auction before the calling price reaches the true adjusted virtual cost is a dominated strategy.

If a supplier stays in the auction when the calling price falls below his true adjusted virtual cost, it is possible that he wins the auction. But he will get a negative profit because the payment is lower than his cost if the auction ends with a calling price lower than his true adjusted virtual cost. Thus, staying in the auction when the calling price falls below the true adjusted virtual cost is a dominated strategy.

Therefore, it is a dominant strategy for each supplier to stay in the auction until the calling price reaches his true adjusted virtual cost. Consequently, the optimal mechanism is implemented: The supplier with the lowest adjusted virtual cost will win the auction and is paid exactly as the optimal mechanism's payment rule specifies.

4.8.2 Proof of Lemma 4

The expected difference between the total procurement cost under the R-region strategy and that under the $(R+1)$-region strategy equals

$$\frac{1}{2} \underset{(\boldsymbol{x},\boldsymbol{y})}{E} \left[\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a}^{\boldsymbol{hh}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a}^{\boldsymbol{hl}}) + \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a}^{\boldsymbol{ll}}) - \pi^{Mech}(\boldsymbol{x}, \boldsymbol{a}^{\boldsymbol{lh}})\right]. \tag{4.5}$$

Note that $\boldsymbol{a}^{\boldsymbol{hh}} = \boldsymbol{a}^{\boldsymbol{hl}} \vee \boldsymbol{a}^{\boldsymbol{lh}}$ and $\boldsymbol{a}^{\boldsymbol{ll}} = \boldsymbol{a}^{\boldsymbol{hl}} \wedge \boldsymbol{a}^{\boldsymbol{lh}}$. Therefore, the lemma follows from the definitions of supermodular and submodular functions.

4.8.3 Proof of Proposition 13

The fact that $\pi^{OPT}(\boldsymbol{x}, \boldsymbol{a}) = \min_{i=1,\ldots,N}\{\psi(x_i)+a_i\}$ implies that, for any $\hat{\boldsymbol{a}}$ and $\tilde{\boldsymbol{a}}$, we have that $\pi^{OPT}(\boldsymbol{x}, \hat{\boldsymbol{a}} \wedge \tilde{\boldsymbol{a}}) = \pi^{OPT}(\boldsymbol{x}, \hat{\boldsymbol{a}}) \wedge \pi^{OPT}(\boldsymbol{x}, \tilde{\boldsymbol{a}})$ and $\pi^{OPT}(\boldsymbol{x}, \hat{\boldsymbol{a}} \vee \tilde{\boldsymbol{a}}) \geq \pi^{OPT}(\boldsymbol{x}, \hat{\boldsymbol{a}}) \vee \pi^{OPT}(\boldsymbol{x}, \tilde{\boldsymbol{a}})$. Thus, $\pi^{OPT}(\boldsymbol{x}, \boldsymbol{a})$ is supermodular in $\boldsymbol{a}$ for any $\boldsymbol{x}$. Therefore, the proposition follows from Lemma 4.

4.8.4 Proof of Proposition 14

When $N = 2$, we have $\pi^{RE}(\boldsymbol{x}, \boldsymbol{a}) = \max\{\psi(x_1) + a_1, \psi(x_2) + a_2\}$. It implies that, for any $\hat{\boldsymbol{a}}$ and $\tilde{\boldsymbol{a}}$, we have that $\pi^{RE}(\boldsymbol{x}, \hat{\boldsymbol{a}} \vee \tilde{\boldsymbol{a}}) = \pi^{RE}(\boldsymbol{x}, \hat{\boldsymbol{a}}) \vee \pi^{RE}(\boldsymbol{x}, \tilde{\boldsymbol{a}})$ and $\pi^{RE}(\boldsymbol{x}, \hat{\boldsymbol{a}} \wedge \tilde{\boldsymbol{a}}) \leq \pi^{RE}(\boldsymbol{x}, \hat{\boldsymbol{a}}) \wedge \pi^{RE}(\boldsymbol{x}, \tilde{\boldsymbol{a}})$. Thus, $\pi^{RE}(\boldsymbol{x}, \boldsymbol{a})$ is submodular in $\boldsymbol{a}$ for any $\boldsymbol{x}$. Therefore, the proposition follows from Lemma 4.

4.8.5 Proof of Proposition 15

It is equivalent to prove the proposition by considering a family of models $\{(F^{(s)}, G), s \in \mathbb{R}\}$, where $F^{(s)}(x) \stackrel{def}{=} F(sx)$. As s goes to infinity, the probability that production cost equals zero goes to one. Therefore, as s goes to infinity, the buyer's expected total procurement cost equals $\overline{Y}_{1:R}$ if she uses $R \leq \lfloor \frac{N}{2} \rfloor$ regions with at least two suppliers in each. The partially diversifying strategy that has $\lfloor \frac{N}{2} \rfloor$ regions with at least two suppliers each dominates any strategy with fewer regions, since $\overline{Y}_{1:R}$ decreases in R when $R \leq \lfloor \frac{N}{2} \rfloor$. Because the buyer's expected total procurement cost equals $\overline{Y}_{2:N}$ if she fully diversifies, she prefers the partially diversifying strategy to the fully diversifying strategy if $\overline{Y}_{2:N} > \overline{Y}_{1:\lfloor \frac{N}{2} \rfloor}$; vice versa.

4.8.6 Proof of Equation (4.3)

We have $\overline{Y}_{1:\lfloor \frac{N}{2} \rfloor} = \int_{-\infty}^{\infty} \bar{G}^{\lfloor \frac{N}{2} \rfloor}(y)dy$ because the tail probability $\Pr(Y_{1:R} > z) = \Pr(y_r > z, r = 1, \ldots, R) = \bar{G}^R(z)$. We have $\overline{Y}_{2:N} = \int_{-\infty}^{\infty} \left[N\bar{G}^{N-1}(y) - (N-1)\bar{G}^N(y)\right] dy$ because the tail probability $\Pr(Y_{2:R} > z) = \Pr(y_r > z, \text{ for all } r \in \{1, \ldots, R\}) + \Pr(y_{\hat{r}} \leq z \text{ for an } \hat{r} \in \{1, \ldots, R\}, y_r > z, \text{ for all } r \in \{1, \ldots, R\} - \{\hat{r}\}) = RG(z)\bar{G}^{R-1}(z) = R[\bar{G}^{R-1}(z) - \bar{G}^R(z)]$.

4.8.7 Proof of Corollary 1

When the regional costs are independent draws from a power-function distribution $G(y) = a^{-v}y^v$ with scale parameter $a > 0$ and shape parameter $v > 0$, we have $\overline{Y}_{1:R} = \frac{a\Gamma(R+1)\Gamma(1+1/v)}{\Gamma(R+1+1/v)}$ and $\overline{Y}_{2:R} = \frac{a\Gamma(R+1)\Gamma(2+1/v)}{\Gamma(R+1+1/v)}$; see Malik (1967). Thus, for $N \geq 4$ even, we have

$$\frac{\overline{Y}_{1:\frac{N}{2}}}{\overline{Y}_{2:N}} = \frac{(N+1/v)(N-1+1/v)\cdots(N/2+1+1/v)}{(1+1/v)N(N-1)\cdots(N/2+1)}.$$

We now show that there exists a threshold $\underline{v} > 0$ such that the above fraction is greater than 1 when $v < \underline{v}$ and less than 1 when $v > \underline{v}$. To see this, note that the numerator minus the denominator can be written as $-b_1v^{-1} + b_2v^{-2} + \ldots + b_{\frac{N}{2}}v^{-\frac{N}{2}}$ with $b_1, \ldots, b_{\frac{N}{2}} > 0$. Thus, the threshold $\underline{v}$ is the unique positive root of $b_1 = b_2v^{-1} + \ldots + b_{\frac{N}{2}}v^{-\frac{N}{2}+1}$. Note $\underline{v}$ is unique because $b_2v^{-1} + \ldots + b_{\frac{N}{2}}v^{-\frac{N}{2}+1}$ is strictly decreasing, approaches positive infinity as

v approaches zero, and approaches zero as v approaches positive infinity. For $N \geq 3$ odd, we can similarly prove that $\overline{Y}_{1:\frac{N-1}{2}}/\overline{Y}_{2:N}$ is greater (less) than 1 if v is greater (less) than a threshold $\underline{v} > 0$. Namely, the proposition holds.

4.8.8 Proof of Proposition 16

In presence of regional codependence, Propositions 13-14 still holds because the proof of Lemma 4 does not use the assumption that regions are independent. To see this, note that the sign of equation (4.5) is not affected by the distribution of $\boldsymbol{y}$ if $\pi^{Mech}(\boldsymbol{x}, \boldsymbol{a})$ is supermodular or submodular in $\boldsymbol{a}$ for all $\boldsymbol{x}$.

4.8.9 Proof of Proposition 17

When the buyer uses mechanism OPT, it is optimal to select all suppliers from different regions even when regions are asymmetric, because any R-strategy having $n_r \geq 2$ suppliers in some region r is dominated by the $(R+1)$-strategy strategy having $n_r - 1$ suppliers in region r and one supplier in the region $(R+1)$ which is of the same type as region r. This is true because the proof of Lemma 4 is still valid given that the distributions of $\boldsymbol{x}$ and $\boldsymbol{y}$ are the same under both the R-strategy and the $(R+1)$-strategy strategy.

4.8.10 Proof of Proposition 18

Note that the buyer's total procurement cost equals $\max\{x_1 + y_1, x_2 + y_2\}$ under mechanism RE. Thus, if the buyer has one supplier in region 1 and one supplier in region 2, the expected total cost equals $\int z dH_1(z)H_2(z)$; if the buyer uses two copies of region k, $k = 1, 2$, the expected total cost equals $\int z dH_k(z)H_k(z)$. The diversification strategy that has one supplier in region 1 and one supplier in region 2 is dominated by either or both of the diversification strategies that use two copies of region k, $k = 1, 2$, because

$$\begin{aligned}
& 2\int_{-\infty}^{\infty} z dH_1(z)H_2(z) - \int_{-\infty}^{\infty} z dH_1(z)H_1(z) - \int_{-\infty}^{\infty} z dH_2(z)H_2(z) \\
= & -\int_{-\infty}^{\infty} z d[H_1(z) - H_2(z)]^2 = \int_{-\infty}^{\infty} [H_1(z) - H_2(z)]^2 dz > 0,
\end{aligned}$$

where the last equality uses integration by parts and the fact that $H_1(-\infty) = H_2(-\infty) = 0$ and $H_1(\infty) = H_2(\infty) = 1$. The proposition follows because the buyer prefers pooling to using two symmetric regions.

4.8.11 Proof of Lemma 5

With $N = 2$ suppliers, the buyer's total cost under RER equals

$$\pi^{RER}(\boldsymbol{x}, \boldsymbol{a}) = [(x_1 + a_1) \vee (x_2 + a_2)] \wedge (1 + a_1) \wedge (1 + a_2).$$

$$\begin{aligned}
\text{Thus,} \quad \pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{hh}}) &= x_1 \vee x_2 + y_1 \vee y_2, \\
\pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{ll}}) &= x_1 \vee x_2 + y_1 \wedge y_2, \\
\pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{hl}}) &= 1 \wedge [(x_1 + |y_1 - y_2|) \vee x_2] + y_1 \wedge y_2, \\
\pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{hl}}) &= 1 \wedge [(x_2 + |y_1 - y_2|) \vee x_1] + y_1 \wedge y_2.
\end{aligned}$$

$$\begin{aligned}
&\text{Hence, } \pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{hh}}) - \pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{hl}}) + \pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{ll}}) - \pi^{RER}(\boldsymbol{x}, \boldsymbol{a^{lh}}) \\
&= 2(x_1 \vee x_2) + |y_1 - y_2| - \{1 \wedge [(x_1 + |y_1 - y_2|) \vee x_2]\} - \{1 \wedge [(x_2 + |y_1 - y_2|) \vee x_1]\}. \qquad (4.6)
\end{aligned}$$

Assuming $x_1 \geq x_2$ without loss of generality, equation (4.6) equals

- $0 \cdot \mathbb{I}(|y_1 - y_2| \in [0, x_1 - x_2]) + (x_1 - x_2 - |y_1 - y_2|) \cdot \mathbb{I}(|y_1 - y_2| \in (x_1 - x_2, 1 - x_1]) + (2x_1 - x_2 - 1) \cdot \mathbb{I}(|y_1 - y_2| \in (1 - x_1, 1 - x_2]) + (2x_1 + |y_1 - y_2| - 2) \cdot \mathbb{I}(|y_1 - y_2| \in (1 - x_2, \infty))$, if $2x_1 - x_2 - 1 \leq 0$;
- $0 \cdot \mathbb{I}(|y_1 - y_2| \in [0, 1 - x_1]) + (x_1 + |y_1 - y_2| - 1) \cdot \mathbb{I}(|y_1 - y_2| \in (1 - x_1, x_1 - x_2]) + (2x_1 - x_2 - 1) \cdot \mathbb{I}(|y_1 - y_2| \in (x_1 - x_2, 1 - x_2]) + (2x_1 + |y_1 - y_2| - 2) \cdot \mathbb{I}(|y_1 - y_2| \in (1 - x_2, \infty))$, if $2x_1 - x_2 - 1 > 0$.

This implies that equation (4.6) is non-positive if $|y_1 - y_2| < (1 - x_1)$, is non-negative if $|y_1 - y_2| \geq 2(1 - x_1)$, and is non-positive if $2x_1 - x_2 - 1 \leq 0$ (or, sufficiently, if $x_1 \leq \frac{1}{2}$) and

$|y_1 - y_2| < 2(1 - x_1)$. The lemma thus follows.

4.8.12 Proof of Proposition 19

Assuming $x_1 \geq x_2$ without loss of generality, equation (4.6) is greater than $(2x_1 - x_2 - 1) \cdot \mathbb{I}(|y_1 - y_2| \leq 2 - 2x_1) + [|y_1 - y_2| - (2 - 2x_1)] \cdot \mathbb{I}(|y_1 - y_2| > 2 - 2x_1)$ when $2x_1 - x_2 - 1 < 0$, and it is greater than $[|y_1 - y_2| - (2 - 2x_1)] \cdot \mathbb{I}(|y_1 - y_2| > 2 - 2x_1)$ when $2x_1 - x_2 - 1 \geq 0$. Thus, the expectation of equation (4.6) over the distributions of $\boldsymbol{x}$ and $|y_1 - y_2|$ is greater than $\min\{2x_1 - x_2 - 1, 0\} \Pr(|y_1 - y_2| \leq 2 - 2x_1) + E[|y_1 - y_2| - (2 - 2x_1)||y_1 - y_2| > (2 - 2x_1)] \Pr(|y_1 - y_2| > 2 - 2x_1)$, which is greater than $E[|y_1 - y_2|] - 1 - (2 - 2x_1) \geq E[|y_1 - y_2|] - 3$. This implies that it is optimal to diversify if $E[|y_1 - y_2|] \geq 3$.

Equation (4.6) is always less than $|y_1 - y_2|$, and is less than $(2x_1 - 1) \cdot \mathbb{I}(1 - x_1 \leq |y_1 - y_2| \leq 1)$ when $x_1 < \frac{1}{2}$. Thus, the expectation of equation (4.6) over the distributions of $\boldsymbol{x}$ and $|y_1 - y_2|$ is less than $E\left[(2x_1 - 1) \cdot \mathbb{I}(1 - x_1 \leq |y_1 - y_2| \leq 1) \mid x_1 < \frac{1}{2}, |y_1 - y_2| < 1\right] \cdot Prob(x_1 < \frac{1}{2}, |y_1 - y_2| < 1) + E[|y_1 - y_2|] \cdot [1 - Prob(x_1 < \frac{1}{2}, |y_1 - y_2| < 1)]$. Since $E[|y_1 - y_2|]$ is finite (because $E[|y_1|] = E[|y_2|] < \infty$) and $E\left[(2x_1 - 1) \cdot \mathbb{I}(1 - x_1 \leq |y_1 - y_2| \leq 1) \mid x_1 < \frac{1}{2}, |y_1 - y_2| < 1\right] < 0$, the expectation of equation (4.6) over the distributions of $\boldsymbol{x}$ and $|y_1 - y_2|$ is negative if $Prob(x_1 < \frac{1}{2}, |y_1 - y_2| < 1)$ is large enough. The lemma thus follows.

Chapter 5

Conclusion

This dissertation consists of three essays on operational strategies to manage procurement costs and risks, focusing on two types of risks concerning corporate procurement managers, namely, supplier non-performance risk and procurement cost risk. To manage the supplier non-performance risk, a buyer can conduct supplier qualification screening processes and only award supply contracts to those qualified suppliers who can pass the qualification screening processes. To mitigate the procurement cost risks, a buyer can diversify her supply base by including pre-qualified suppliers from different geographic regions in the supply base.

Firms' operational strategies to manage procurement risks can induce strategic responses of their supply chain patenters, which could then affect the firms' total procurement cost equations and hence the buyer's optimal operational strategy.

When the buyer conducts a procurement auction in combination with the supplier qualification screening process to award a new contract to a group of new suppliers, the suppliers may strategically make their bidding strategy contingent on their believes about their competitors' probabilities of being truly qualified, which consequently affects the buyer's expected contract payment to winning suppliers. Since the suppliers' believes about competitors' qualification probabilities are affected by the timing that the buyer performs qualification screening process, the buyer needs to optimize the timing so as to minimize her expected total procurement cost. To explore the buyer's optimal qualification screening strategy, an analytic framework is developed in which the suppliers strategic bidding behavior is captured in a mechanism design approach. The analysis shows that the buyer's optimal qualification screening timing depends on both the buyer's qualification screening cost and the buyer's cost to fulfill the contract outside of the auction: it is optimal to delay more qualification screening until after the auction if the qualification screening cost is higher or the outside option cost is lower.

Analogously, when the buyer conducts a procurement auction between an incumbent supplier and an entrant, the timing of qualification screening the entrant would affect the incumbent's bidding strategy. The incumbent in an open-descending auction setting chooses three strategies — boycotting, testing-the-water, and bidding-to-win — depending on his cost. Such strategic bidding behavior can also be affected by the buyer's qualification cost and the entrant's qualification probability. Thus, a buyer seeking total procurement cost minimization needs to tradeoff between "scaring the incumbent for price concessions" and the incumbent's strategic holding back in bidding. The research shows that the buyer should delay screening the entrant until after the auction (to save qualification cost and retain the opportunity of running the auction) unless the qualification cost is very cheap and the entrant's qualification probability is medium.

The buyer also needs to take into account of suppliers' strategic response when diversifying her supply base. This is because diversification amplifies cost asymmetry among suppliers and could enable suppliers with cost advantage to use such advantage against the interests of the buyer — by charging a high contract price and taking windfall profit out of the buyer's diversification benefit — especially when the contract price is mainly determined through the suppliers' competition. Such windfall-profit taking can be so significant that it would surprisingly be better for the buyer not to diversify the supply base in a typical setting where the buyer uses two suppliers and runs an open descending reverse English auction. With more suppliers and regions, such a buyer can contain the advantaged suppliers' windfall profit-taking by keeping two competing suppliers within the same region, i.e., using a partial diversification strategy.

References

[Anupindi and Akella (1993)] Anupindi, R. and R. Akella. 1993. Diversification under supply uncertainty. *Management Science* **39**, 944-963.

[Ariba (2005)] Ariba Global Services Organization. 2005. Supply watch. Quarter **3**.

[Ausubel et al. (2002)] Ausubel, L. M., P. Cramton, and R. J. Deneckere. 2002. Bargaining with incomplete information. In *Handbook of Game Theory with Economic Applications*, **3**, Aumann, R. J. and S. Hart (eds.), Elsevier Science Publishers.

[Ausubel and Cramton (2006)] Ausubel, L. M. and P. Cramton. 2006. Dynamic Auctions in Procurement. In *Handbook of Procurement*, N. Dimitri, G. Piga and G. Spagnolo (eds.), Cambridge University Press.

[Babich et al. (2007)] Babich, V., A. N. Burnetas, P. H. Ritchken. 2007. Competition and diversification effects in supply chains with supplier default risk. *Manufacturing & Service Operations Management* **9**, 123-146.

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

[Beal et al. (2003)] Beall, S., C. Carter, P. L. Carter, T. Germer, T. Hendrick, S. D. Jap, L. Kaufmann, D. Maciejewski, R. Monczka, K. Petersen. 2003. The role of reverse auctions in strategic sourcing. Center for Advanced Procurement Studies Research.

[Beil and Wein (2003)] Beil, D. R., L. M. Wein. 2003. An inverse-optimization-based auction mechanism to support a multiattribute RFQ process. *Management Sci.* **49** 1529–1545.

[Binmore et al. (1986)] Binmore, K., A. Rubinstein, and A. Wolinsky. 1986. The Nash bargaining solution in economic modelling. *Rand Journal of Economics*, **17**, 176-188.

[Branco (1997)] Branco, F. 1997. The design of multidimensional auctions. *The RAND Journal of Economics* **28** 63–81.

[Braynov and Sandholm (2003)] Braynov, S. and T. Sandholm. 2003. Auctions with untrustworthy bidders. In *Proceedings of the IEEE Conference on Electronic Commerce.* Newport Beach, CA, June 24–27.

[Bulow and Klemperer (1996)] Bulow, J. and P. Klemperer. 1996. Auctions versus negotiations. *The American Economic Review* **86**, 180-194.

[Calveras et al. (2004)] Calveras, A., J.-J. Ganuza and E. Hauk. 2004. Wild bids. Gambling for resurrection in procurement contracts. *Journal of Regulatory Economics* **26** 41–68.

[Carr (2003)] Carr, S. M. 2003. Note on online auctions with costly bid evaluation. *Management Sci.* **49** 1521–1528.

[Che (1993)] Che, Y.-K. 1993. Design competition through multidimensional auctions. *RAND Journal of Economics* **24** 668–680.

[Chen (2007)] Chen, F. 2007. Auctioning supply contracts. *Management Sci.* **53** 1562–1576.

[Chen and Vulcano (2008)] Chen, Y. and G. Vulcano. 2008. Effects of information disclosure under first- and second-price auctions in a supply chain setting. *Manufacturing & Service Operations Management* **11**, 299-316.

[Costello (2008)] Costello, M. 2008. Shipping insurance cost soars with piracy surge off Somalia. *The Times*, September 11, 2008.

[Dada et al. (2007)] Dada, M., N. C. Petruzzi, and L. B. Schwarz. 2007. A newsvendor 's problem when suppliers are unreliable. *Manufacturing & Service Operations Management* **9**, 9-32.

[Dasgupta et al. (1979)] Dasgupta, P., P. Hammond and E. Maskin. 1979. The implementation of social choice rules: Some general results on incentive compatibility. *Review of Economic Studies* **46** 185–216.

[Elmaghraby (2000)] Elmaghraby, W. J. 2000. Supply contract competition and sourcing policies. *Manufacturing & Service Operations Management* **2** 350–371.

[Elmaghraby (2007)] Elmaghraby, W. 2007. Auctions within e-sourcing events. *Production and Operations Management* **16**, 409-422.

[Federgruen and Yang (2007)] Federgruen, A. and N. Yang. 2007. Optimal supply diversification under general supply risks. To appear in *Operations Research.*

[Federgruen and Yang (2008)] Federgruen, A. and N. Yang. 2008. Selecting a portfolio of suppliers under demand and supply risks. *Operations Research* **56** 916-936.

[Feinberg and Huber (1996)] Feinberg, F. M. and J. Huber. 1996. A theory of cutoff formation under imperfect information. *Management Sci.* **42** 65–84.

[Fudenberg and Tirole (1991)] Fudenberg, D. and J. Tirole. 1991. *Game Theory.* The MIT Press.

[Gibbard (1973)] Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* **41** 587–601.

[Green and Laffont (1977)] Green, J. and J. J. Laffont. 1977. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica* **45** 427–438.

[Grey et al. (2005)] Grey, W., T. Olavson, and D. Shi. 2005. The role of e-marketplaces in relationship-based supply chains: A survey. *IBM Systems Journal* **44**, 109-123.

[Harstad et al. (1990)] Harstad, R. M., J. H. Kagel and D. Levin. 1990. Equilibrium bid functions for auctions with an uncertain number of bidders. *Economics Letters* **33** 35–40.

[Hewitt and Savage (1955)] Hewitt, E. and L.J. Savage. 1955. Symmetric measures on Cartesian products. *Trans. Amer. Math. Soc.* **80**, 470-501.

[Jap (2003)] Jap, S. D. 2003. An Exploratory study of the introduction of online reverse auctions. *Journal of Marketing* **67** 96–107.

[Krishna (2002)] Krishna, V. 2002. *Auction Theory.* Academic Press.

[Leenders and Fearon (1997)] Leenders, M. R. and H. E. Fearon. 1997. *Purchasing and Supply Management.* 11th edition. Irwin.

[Malik (1967)] Malik, H. J. 1967. Exact moments of order statistics from a power-function distribution. *Skand. Aktuarietidskr.* **50**, 64 - 69.

[Maskin and Riley (2000)] Maskin, E. and J. Riley. 2000. Asymmetric auctions. *The Review of Economic Studies* **67**, 413-438.

[McAfee and McMillan (1987)] McAfee, R. P. and J. McMillan. 1987. Auctions with a stochastic number of bidders. *Journal of Economic Theory* **25** 1–19.

[McAfee and McMillan (1988)] McAfee, R. P. and J. McMillan. 1988. Search mechanisms. *Journal of Economic Theory* **44** 99–123.

[McAfee and McMillan (1989)] McAfee, R. P. and J. McMillan. 1989. Government procurement and international trade. *Journal of International Economics* **26**, 291-308.

[Milgrom (1987)] Milgrom, P. 1987. Auction theory. T. Bewley, eds.*Advances in Economic Theory Fifth World Congress.* Cambridge University Press, 1–32.

[Milgrom (2004)] Milgrom, P. 2004. *Putting Auction Theory to Work.* Cambridge University Press.

[Myers (2007)] Myers, R. 2007. Food fights. *CFO Magazine* **1 June**.

[Myerson (1981)] Myerson, R. B. 1981. Optimal auction design. *Math. Oper. Res.* **6** 58–73.

[Nash (1950)] Nash, J. F. 1950. The bargaining problem. *Econometrica*, **18**, 155-162.

[Pedersen (2004)] Pedersen, A. 2004. Putting a price tag on an international procurement office. *Purchasing* Oct. 7.

[Rezende (2008)] Rezende, L. 2008. Biased procurement auctions. To appear in *Economic Theory.*

[Rochet and Stole (2003)] Rochet, J. C. and L. Stole. 2003. The economics of multidimensional screening. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress.*

[Roth (1979)] Roth, A. 1979. *Axiomatic Models of Bargaining.* Berlin: Springer-Verlag.

[Rothkopf (1991)] Rothkopf, M. H. 1991. On auctions with withdrawable winning bids. *Marketing Sci.* **10** 40–57.

[Shaked (1977)] Shaked, M. 1977. A concept of positive dependence for exchangeable random variables. *The Annals of Statistics.* **5**, 505-515.

[Snir and Hitt (2003)] Snir, E. M. and L. M. Hitt. 2003. Costly bidding in online markets for IT services. *Management Sci.* **49** 1504–1520.

[Spencer and Casey (2007)] Spencer, J. and N. Casey. 2007. Toy recall shows challenge China poses to partners. *The Wall Street Journal.* 3 Aug, pg. A1.

[Strausz (2006)] Strausz, R. 2006. Deterministic versus stochastic mechanisms in principal-agent models. *Journal of Economic Theory* **128** 306–314.

[Tomlin (2006)] Tomlin, B. 2006. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science* **52**, 639-657.

[Topkis (1998)] Topkis, D. M. 1998. *Supermodularity and Complementarity.* MIT Press, Cambridge, Massachusetts.

[U.S. Department of Commerce (2005)] U.S. Department of Commerce, Bureau of the Census. 2005. Statistics for Industry Groups and Industries: 2004, Annual Survey of Manufactures. M04(AS)-1, pp.1–155, and Appendix. *http://www.census.gov/prod/2005pubs/am0431gs1.pdf*

[Waehrer (1995)] Waehrer, K. 1995. A model of auction contracts with liquidated damages. *Journal of Economic Theory* **67** 531-555.

[Wan and Beil (2008)] Wan, Z. and D. R. Beil. 2008. RFQ auctions with supplier qualification screening. To appear in *Operations Research.*

[Welch (2007)] Welch, D. 2007. Made in China: Faulty tires how a Chinese supplier's bad decision turned into one importer's worst nightmare, and may mean the end of his business. *http://www.businessweek.com/bwdaily/dnflash/content/jul2007/db20070711_487422.htm.*

[Yang et al. (2009)] Yang, Z., G. Aydin, V. Babich, and D. R. Beil. 2008. Supply disruptions, asymmetric information and a backup production option. *Management Science* **55** 192–209.

[Ye (2007)] Ye, L. 2007. Indicative bidding and a theory of two-stage auctions. *Games and Economic Behavior* **58** 181–207.

[Zheng (2001)] Zheng, C. Z. 2001. High bids and broke winners. *Journal of Economic Theory* **100** 129–171.

[Zhong (2007)] Zhong, F. 2007. Empirical analysis of online procurement auctions: business value, bidding behavior, learning and incumbent effect. Ph.D. Thesis at Georgia Institute of Technology.