Specification Vagueness and Supply Quality Risk

Yimin Wang

W. P. Carey School of Business, Department of Supply Chain Management, Arizona State University, Tempe, Arizona

Received 5 January 2012; revised 16 January 2013; accepted 24 January 2013
DOI 10.1002/nav.21530
Published online 6 March 2013 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: Specifying quality requirement is integral to any sourcing relationship, but vague and ambiguous specifications can often be observed in practice, especially when a buyer is in the initial stage of sourcing a new product. In this research, we study a supplier's production incentives under vague or exact quality specifications. We prove that a vague specification may in fact motivate the supplier to increase its quantity provision, resulting in a higher delivery quality. Vague quality specification can therefore be advantageous for a buyer to screen potential suppliers with an initial test order, and then rely on the received quality level to set more concrete quality guidelines. There is a degree, though, to which vague quality specification can be effective, as too much vagueness may decrease the supplier's quantity provision and hence the expected delivery quality. © 2013 Wiley Periodicals, Inc. Naval Research Logistics 60: 222–236, 2013

Keywords: quality control; specification vagueness; supply quality risk

1. INTRODUCTION

Specifying quality requirements is an important consideration in almost all business transactions. Although firms strive to set comprehensive and exact quality requirements, incomplete and vague requirements are prevalent in the business world. Williamson [42] proclaims that "all contracts are incomplete and imperfect documents" and that firms often prefer private ordering (i.e., voluntary agreements) to court system to resolve such issues, Tirole [41] holds a similar sentiment by pointing out that "many contracts are vague or silent on a number of key features." Recently, Ariely [2] also argues that incomplete and vague requirements can in fact be beneficial to firms. Anecdotal evidence in different industries suggests that firms may use incomplete or vague requirements intentionally.

> "Honda, often, especially at the beginning of a new product line, makes things vague on what the expectations are in terms of quality. For instance, they do not spell out the exact level of drag of a hinge as it closes and opens." A manager at Honda later explained that this initial vagueness actually helps them "push new standards or increased quality level". Choi and Hong [9].

In contrast, our discussion with former process engineers in the industrial diamonds industry suggests that the prevalence of vague quality requirements (for tooling equipment) is largely unintentional, partly because they do not know in advance what quality standards are most appropriate. Such vagueness in quality requirements sometimes does result in a contentious supplier relationship when the supplier's products are rejected due to "realized" quality standards. The suppliers may tolerate such vague quality requirements if the residual value of sustaining the relationship with the buyer is significant, for example, a promise of more future businesses. Similar observations can also be found in the textile industry, where vague quality requirements are often observed due to the subtlety and complexity of fabric production process.[1]

> "A majority of the quality problems in the textile industry can be traced back to yarn quality parameters such as evenness, imperfections, count variations, hairiness and strength. Yet, often companies dealing with yarns or yarn-based products may not understand enough about yarn and fiber quality to define their requirements clearly. This vague situation creates great economic inefficiencies in sourcing: the quality levels are either too low or too high." (Uster Technologies [1]).

Regardless of whether firms set vague quality requirements intentionally or not, the broad issue of specifying

Correspondence to: Y. Wang (yimin_wang@asu.edu)

[1] In a multitier textile processing supply chain, however, the buyer may not be in a strong legal position to reject the supplier's quality provision. In such cases, vague quality requirements are often subsumed by industry standards.

quality requirements is an important area of study. This is especially the case when quality requirements interact with supplier's incentives to improve delivery quality. Given that firms almost always want suppliers to exert effort to improve quality, we believe it both important and relevant to identify alternative levers that firms may use to incentivize suppliers to improve delivery quality. In this article, we focus on how vagueness in quality requirements influences the supplier's delivered quality level, that is, the expected quality level of those units actually delivered by the supplier.

To set our work on a concrete footing, it is instructive to consider how one might apply our work in the initial stage of a supplier selection process. Suppose a buyer wishes to source a particular component or product from some potential suppliers whose cost structures and/or process capabilities are not completely transparent to the buyer. Instead of investing significant time and effort upfront to profile each supplier and then write a comprehensive contract, the buyer may simply approach these suppliers with some vague quality requirements for a small test order. The buyer may then examine which supplier achieves the highest delivered quality level ex post, and subsequently work with this particular supplier for regular orders with a more precise set of quality requirements.

From a practice point of view, the above approach applies to situations where the specifiability of goods is low, that is "items that do not have clearly defined attributes that competing suppliers can translate into unambiguous specifications." (Beall et al. [5]). In fact, ambiguous specifications are frequently encountered in different purchasing contexts across many different industries. A comprehensive study by Center for Advanced Purchasing Studies, for example, noted that "one of the most common supplier complaints concerning e-RAs [reverse auctions] was the lack of clear specifications and rules from the buying firm." (Beall et al. [5]) Our modeling framework is therefore particularly appropriate in the RFxs stage, which is often considered a critical step in evaluating potential suppliers (Monczka et al. [30]).[2]

The above framework also applies even if the buyer perfectly knows its desired requirements a priori. By not revealing its true requirements to suppliers, the buyer may create sufficient incentives for some suppliers to exceed its original expected delivery quality level. It is worth pointing out that we are not the first to observe that vague requirements may positively incentivize suppliers. Gal-Or et al. [16], for example, study whether a buyer should reveal its true requirements (valuation) to suppliers, but they focus on whether the buyer should approach suppliers sequentially or simultaneously to incentivize suppliers to reduce their prices.

In cases where the buyer truly does not know its requirements in advance, Terwiesch and Loch [40] and Wolinsky [43] characterize how vagueness in buyer's requirements incentivizes (or disincentivizes) the supplier's effort in providing customized products to the buyer. In particular, Terwiesch and Loch [40] study the optimal number of prototypes the supplier should build and how much it should charge the buyer, whereas Wolinsky [43] focus on the supplier's optimal consultation fee when the buyer performs a sequential search. Our work complements the above research by focusing on how vagueness in quality requirements may motivate the supplier to improve the expected delivery quality via production lot-sizing decisions.

1.1. Research Perspective

Throughout the article, we adopt the convention that a higher quality is more desirable to the buyer, and the term quality may either refer to conformance quality (for example, tolerance level) or performance quality (for example, smoothness of ball bearings). We focus on situations where quality dimension is known and measurable, but the buyer's quality specification may be vague. This could be the case, for example, when an exact specification is possible but is complex to stipulate and costly to manage. We do not study scenarios where the quality itself cannot be measured or the buyer is unsure of what quality dimensions to measure. Such scenarios can be found in the noncontractibility literature, and we refer the interested reader to Baiman et al. [3] and references therein for more details.

In this research, we do not study supplier-buyer interactions from a game-theoretic approach; instead, we focus on the effect of specification vagueness on supplier's quality provision. The supplier may meet the buyer's quality specifications (either exact or vague) in several different ways: (1) improve the capability of its production process, (2) set quantity provisions by producing more and only deliver the fraction of output with better quality, or (3) perform post production repairs and rework to bring the quality up to the target standard. Our focus is the second approach, that is quantity provisions. If the supplier's production process is constrained by technology limitations, such as in automobile, semiconductor, industrial diamonds, and biochemical industries, then the first approach can be costly. On the other hand, repair and rework may not be feasible for certain products: repair and rework on industrial pipe flanges, for example, are often prohibited because they may weaken the structural integrity of the flange.

[2] RFxs refer to several variants of supplier screening approach, such as RFP (request for proposal), RFQ (request for quotation), or RFS (request for sample). An RFP "may only detail the end use or performance characteristics of the needed item and ask the supplier to propose specifically how they would satisfy these needs" (Beall et al. [5]). Our model can be seen as a variant of RFP (with sample production request) or RFS (with ambiguous specifications).

Quantity provision is a common practice used in many different industries when there exist quality variations, and it is often referred to alternatively as reject allowance (Levitan [27]), overplanning in material requirements planning (MRP) with uncertain quality (Murthy and Ma [32]), hedging or yield factor in MRP/MRPII (Mula [31]), and it is also somewhat related to MLPO, that is, multiple lot-sizing production to order (Grosfeld-Nir and Gerchak [17]). Note that in practice firms are often reticent about quality problems (and hence quantity provision strategies), perhaps due to brand image or competitive concerns. Nevertheless, there is evidence that suppliers do use quantity provisions to cope with quality variations (ERG [14]), and that suppliers' quality provision decisions often do respond to perceived quality requirements (BCG [4]).

In closing, we note that in theory one can always raise the target quality level with exact specifications, such as imposing tighter tolerance level, but this may not be easily implemented without knowing precisely the supplier's detailed cost structure and production capabilities. On the other hand, one can also pay for higher quality, that is, pay more per unit for higher quality levels. This is certainly a reasonable approach, albeit at a higher procurement cost. In contrast, our primary focus is to study ceteris paribus, that is, with same procurement cost, the effect of specification vagueness on expected delivery quality levels.

1.2. Literature

Although not a focus here, production capability improvements and repair/rework are well-studied in the literature and they form two important classes of quality control problem in many industries.

Reyniers and Tapiero [36, 37] explore various operational levers, for example, pricing, penalties for defects, and warranty terms, that a buyer can use to induce the supplier to improve the quality of its production process. Sheopuri and Zemel [38] also consider a similar problem, but they focus on remedial actions by the buyers, that is, individual versus class actions. From a somewhat different angle, Hwang et al. [22] consider the effect of two operational levers, appraisal and certification, that a buyer might adopt to improve supply quality. They argue that appraisal (inspection) may not be desirable because it can cause the supplier to perform excessive or unwanted preemptive inspections that increase the supply chain's cost. A common assumption made in the above research is that the supplier's quality is perfectly contractible and that the buyer knows the supplier's "cost of quality," that is, the effort required to reach a certain quality level. In contrast, Baiman et al. [3], Lim [28], and Kaya and Özer [24] study a buyer's contract design problem when the supplier's "cost of quality" is not observable by the buyer. This stream of literature mainly focuses on contract design, for example, price rebate, inspection, and sales price commitment, to improve the supply quality.

The existing research on repair and rework mainly focuses on the optimal inspection and repair policies in a dynamic production system. This stream of literature often considers the inspection policy jointly with procurement decisions in a production and inventory system, see for example Britney [7], Lee and Rosenblatt [26], Peters et al. [34], Ou and Wein [33], and Chen et al. [8]. They study the joint replenishment and quality screening policy when there exist lot-by-lot quality variations. We refer the interested reader to Yao and Zheng [45] for a comprehensive treatment of optimal inspection policies in a dynamic inventory-production system. The random yield and uncertain supply literature, for example, Dada et al. [10], Dong and Tomlin [11], Kazaz and Webster [25], Liu et al. [29], Swinney and Netessine [39], and Yano and Lee[44], is also related to supply quality problem, but the main thrust of that literature is finding optimal procurement and inventory policies, assuming that the yield or supply distribution is exogenous.

More generally, our work is also related to the incomplete contract literature. Some key motivations for the incomplete contract literature include unforeseen contingencies and transaction cost, such that the buyer and the supplier may not be able to incorporate all future contingencies into the contract or the cost of writing a completely contract and enforcing it through the court is prohibitively high. Some influential earlier work on incomplete contract includes Grossman and Hart [18] (share of property rights), Hart and Moore [20, 21] (renegotiation and limited property rights), and Williamson [42] (decision rights), among others. We refer the interested reader to Tirole [41] for a detailed critique of the incomplete contract literature.

One could view our setting as a special case of an incomplete contract in which the supplier does not have renegotiation power ex post and the buyer holds all decision rights. Our work, however, differs fundamentally from the incomplete contract literature in terms of economic rationale: we allow the buyer to institute "incomplete contract" without invocation of transaction costs concerns. In other words, the buyer could choose to use an incomplete contract even if all future contingencies can be exhausted (and hence incorporated into the contract) without significant transaction costs. As a side note, our work is also tangentially related to the fuzzy decision making literature where fuzzy concepts are modeled by a deterministic membership function, see Bellman and Zadeh [6].

Our research complements the existing literature by studying how specification vagueness influences the supplier's quantity provision and the consequent expected delivery quality. As alluded to in the introduction, one key advantage of vague specification is that the buyer does not have to incur upfront cost to learn precisely about the supplier's "cost of quality" to incentivize the supplier to increase quantity provision and the expected delivery quality. This is particularly useful when the buyer is in the initial stage of supplier selection process, in which important information, such as supplier's cost and capability, may not be readily available.

We prove that under certain conditions the buyer may be better off with a vague quality specification. We find that all else being equal a vague specification is more attractive if the supplier has a less reliable production process, whereas an exact specification is more attractive if the supplier has a more reliable process. One important implication of our findings is that contracting complexity in supply quality may not always be a problem: the buyer may in fact receive a higher delivery quality with vague quality specifications, even if an exact specification is possible. Our results also make a weaker requirement on the buyer's knowledge in supplier's precise cost structure and production capability, which facilitates the buyer's initial supplier selection process. By introducing a "proneness to low quality" index in the supplier's production capability (quality), the buyer may use some benchmark index to determine whether it is desirable to use vague specification. Our findings also indicate that there is a degree to which specification vagueness is beneficial, too much vagueness can decrease the supplier's quantity provision and hence decrease the expected delivery quality.

The rest of this article is organized as follows. We describe the model in Section 2 and introduce some preliminary analysis in Section 3. The effect of specification vagueness is analyzed in Section 4. We conclude in Section 5. All proofs are contained in Appendix.

2. MODEL

Consider a supplier that needs to deliver q amount of output to a buyer. The supplier's production process is imperfect, that is, when the supplier launches a production run, the quality of its output can be described by a certain class of distribution functions. The exact form (e.g., parameters) of the realized quality distribution can vary depending on random factors (denote as ϵ), such as temperature, humidity, chemical impurity, and material variations, that may influence the supplier's production process.

Let $F_\epsilon(x)$ describe the fraction of output that has a quality level less than x, conditional on the realized random factors ϵ. Define $\overline{F}_\epsilon(x) = 1 - F_\epsilon(x)$. If the supplier launches a production lot size of Q, the amount of output that exceeds quality level x is $Q\overline{F}_\epsilon(x)$. This implies that the fraction of the output that exceeds a certain quality level is invariant to the production lot size, although the amount of the output that exceeds a certain quality level is increasing in the lot size Q. The above random quality model is a direct extension of the commonly used stochastic proportional random yield model, which "applies to circumstances where yield losses occur because of limited capabilities of the production system." (Yano and Lee [44])

2.1. Quality Expectations

The buyer has certain expectations on the minimum quality level l and it will only accept the fraction of the supplier's delivery with quality level greater than l. The buyer may use an exact quality specification by setting a fixed minimum quality level $l = \underline{l}$ such that it will only accept the fraction of the supplier's delivery with quality level $x \geq \underline{l}$ and reject the rest. Alternatively, the buyer may use a vague quality specification by being somewhat vague about the exact values of the minimum quality level l, and, in this case, the buyer may reject a portion of the supplier's delivery even if the quality of that portion exceeds $\underline{l}$, and vice versa. Our adoption of minimum acceptance quality level is driven by the notion that a higher quality is more desirable, which is advocated by quality management pioneers including Taguchi and Deming and widely recognized in practice (Evans and Lindsay [15]). Our adoption of the minimum quality level, however, does not imply one-sided specification only, see Remark 1 below.

REMARK 1: One needs not to construe the minimum quality level l as one-sided specification only: it can also represent two-sided specification such as a tolerance interval. Suppose that an exact specification for a hinge's damping ratio is 1.200 ± 0.020, then one can define $x = 0.040 - |\zeta - 1.200|$ (where ζ is the actual damping ratio) and set $\underline{l} = 0.020$ such that the buyer will only accept the fraction of delivery with quality level $x \geq \underline{l}$, that is, $\zeta \in 1.180$–1.220. In contrast, with vague specification l can be greater or less than 0.020 such that the buyer may accept deliveries within a wider or tighter tolerance interval.

REMARK 2: Conceptually, the minimum quality level l can also capture uncertain, two-sided target specifications. Suppose the buyer is unsure about a hinge's ideal damping ratio, but it can tell (on seeing the delivery) whether the actual damping ratio is close to *ideal*, where the buyer may use crude measures such as likert scale (e.g., "1. not close at all", "2. slight close," "3. somewhat close", "4. very close," and "5. extremely close") to describe the quality distance between actual and ideal target specifications. Let x and l denote (the inverse of) the actual and specified quality distances (e.g., numerical values associated with likert scale), respectively, then only those delivery with quality level $x \geq l$ will be accepted. With exact specification, for example, the buyer may set a fixed $\underline{l} = 4$ such that only those delivery with quality distance "very close" or "extremely close" will be accepted. On the other hand, the buyer may be vague about

l such that it may accept "somewhat close" and above, but may also accept "extremely close" only.[3]

For vague specification, the minimum quality level l is uncertain, and, in this case, let $G(\cdot)$ denote the distribution of l. Note that $G(\cdot)$ can also describe the exact specification as a special case: $G(l) = 1$ for any $l \geq \underline{l}$ and $G(l) = 0$ for $l < \underline{l}$. To make $G(\cdot)$ more concrete, we illustrate a particular functional form below. Such a functional form, however, is not required for subsequent analysis. Let

$$G(l) = \frac{1}{2}\left(1 + \operatorname{erf}\left[\frac{l - \underline{l}}{\sigma\sqrt{2}}\right]\right), \qquad (1)$$

where $\operatorname{erf}(\cdot)$ is the standard error function and σ captures the vagueness of the quality specification. Note that $\lim_{\sigma\to 0} G(l) = H_{\underline{l}}(l) = H(l - \underline{l})$, where $H(\cdot)$ is the Heaviside step function (assuming $H(0) = 1$). This corresponds to exact specification where the minimum quality level is set at $\underline{l}$. In contrast, if $\sigma > 0$, then $G(l)$ is greater than 0 and less than 1 for any quality level l, although $G(l)$ approaches to 1 for sufficiently high-quality level l. This corresponds to vague specification, where there is a 50-50 chance that the minimum quality level is $\underline{l}$. The vague specification described here can be viewed as a symmetric adaptation of the "anchoring-and-adjustment" approach introduced by Einhorn and Hogarth [13].

The supplier could form a crude estimation of $G(\cdot)$ based on the buyer's implicit (albeit vague) quality requirements, as well as its past knowledge with other buyers and general industry/market benchmarks. For example, the supplier could estimate that a tolerance level of x_i may be accepted with probability y_i for $i = 1, \ldots, n$, given that industry benchmark is $\bar{x}$ and the particular buyer "seems tougher" than other buyers. These crude estimations can subsequently be combined to statistically find the most likely distribution for $G(\cdot)$.

2.2. Production Cost and Revenue

To satisfy the buyer's delivery requirement (quantity and quality), the supplier sets a production lot size $Q \geq q$, that is, the supplier anticipates its quality risk and therefore builds quantity provisions by launching a larger than required production run. We assume that the supplier makes one production run and it has to commit to the production quantity Q before observing actual quality realizations. This can be the case, for example, when the buyer's initial test order is small, the setup cost is prohibitively high, the lead time is long, or the product life cycle is short.

[3] This conceptual framework requires that the supplier and the buyer share a similar perception about what is "ideal." Such similar perception may be established, for example, through their prior working relationships. Otherwise, without a similar perception the notion of exact specification (by definition) is not meaningful.

The supplier's unit production cost is c, and the buyer pays a unit wholesale price $w > c$ for the (accepted) supplier's delivery. In addition, the supplier incurs a linear penalty cost of p for the fraction of the output that fails to meet the buyer's quality expectations. Leftovers of the supplier's output is salvaged at a unit value of $s(x) < c$, where $s'(x) \geq 0$ and $s''(x) \leq 0$, that is, the salvage value is concave increasing in quality level x. To avoid the trivial case where salvaging a product is more profitable than selling to the buyer, we need the following assumption.

ASSUMPTION 1: $(w+p)G(x) - s(x)$ is nonnegative and increasing in quality level x.

For a given unit with quality x, the probability of acceptance is $G(x)$ and so $(w + p)G(x)$ is the expected payoff. Hence, Assumption 1 says that the expected payoff of selling a unit is greater than salvaging it, and that as quality level x increases, it is increasingly better to sell than to salvage.

Assumption 1 is not required when $x < \underline{l}$ under exact specification, or, when salvage value is independent of quality level (i.e., $s'(x) = 0$). The latter may be the case, for example, when the product is customized or when salvage value mainly derives from material cost.

2.3. Allocation Policy

After production is finished, the supplier observes the realized quality distribution $F_\epsilon(\cdot)$ and then allocates a fraction of its output Q to satisfy the delivery requirement q. Allocating the best q units to the buyer is a reasonable policy, but depending on the buyer's quality expectations such a policy may or may not be optimal for the supplier. Hence, in what follows we describe a generic allocation policy first and characterize the optimal policy later.

Let $[0, Q]$ be the set of output ranked in ascending quality order. An allocation rule π_ϵ over $[0, Q]$ defines a subset $A^{\pi_\epsilon} \subseteq [0, Q]$ such that $\pi_\epsilon(y) = 1$ for any $y \in A^{\pi_\epsilon}$ and $\pi_\epsilon(y) = 0$ for any $y \notin A^{\pi_\epsilon}$, that is, π_ϵ indicates the subset A^{π_ϵ} for delivery to the buyer. For an arbitrary allocation rule π_ϵ, A^{π_ϵ} can be a collection of disjoint subsets, that is, $A^{\pi_\epsilon} = \cup_i A_i^{\pi_\epsilon}$, where $A_i^{\pi_\epsilon} = [a_i^0, a_i^1]$ and $0 \leq a_i^0 \leq a_i^1 \leq Q$. In addition, A^{π_ϵ} needs to satisfy the delivery requirement q, that is, $|A^{\pi_\epsilon}| = \sum_i |A_i^{\pi_\epsilon}| = \sum_i (a_i^1 - a_i^0) = q$. See Fig. 1 for an illustration. Note that any feasible allocation is completely determined by $\vec{a} = (\ldots, a_i^0, a_i^1, \ldots)$, and we will characterize $\vec{a}$ after describing the supplier's problem.

2.4. Summary and Problem Formulation

2.4.1. *Summary of Notation*

- q: buyer's delivery quantity requirement (order quantity).
- Q: Supplier's production quantity.

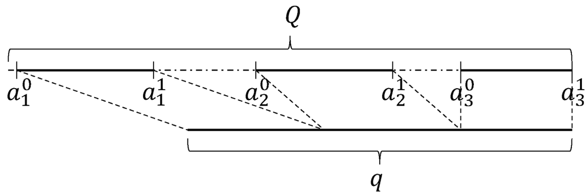

Figure 1. An illustration of an arbitrary allocation policy.

- ϵ: Random production shocks.
- $F_\epsilon(x)$: Fraction of the supplier's output with quality level less than x.
- π_ϵ: Allocation policy for a realized production shock ϵ.
- A^{π_ϵ}: Subsets of output delivered to the buyer according to allocation policy π_ϵ.
- l: Minimum quality level imposed by the buyer.
- $G(\cdot)$: Distribution function of the minimum quality level l.
- c: Supplier's unit production cost.
- w: Unit wholesale price paid (for accepted units) by the buyer to the supplier.
- p: Unit penalty cost incurred by the supplier for unmet delivery quantity requirement.
- $s(x)$: Unit salvage value for a given quality level x.

2.4.2. *Sequence of Events*

1. Supplier learns procurement requirement q and the quality requirement $G(\cdot)$.
2. Supplier sets production lot size Q before observing production shocks ϵ.
3. After production is finished and ϵ is realized, supplier observes realized quality distribution and determines an allocation policy π_ϵ.
4. Supplier delivers q units (if available) to the buyer according to the allocation scheme A^{π_ϵ}, receives unit revenue w for the fraction accepted, incurs unit penalty p for the fraction (if any) that is rejected, and salvages the leftovers at a unit value of $s(\cdot)$.

2.4.3. *Problem Formulation*

The supplier's expected profit for any given allocation policy π_ϵ is

$$\begin{aligned} v(Q|\pi_\epsilon) = \mathsf{E}_\epsilon \int_x \Bigg(w \int_{y \in A^{\pi_\epsilon}} \mathbf{1}\left(F_\epsilon^{-1}(y/Q) \geq x\right) dy \\ - p \int_{y \in A^{\pi_\epsilon}} \mathbf{1}\left(F_\epsilon^{-1}(y/Q) < x\right) dy \Bigg) dG(x) \\ + \mathsf{E}_\epsilon \int_{y \notin A^{\pi_\epsilon}} s\left(F_\epsilon^{-1}(y/Q)\right) dy - cQ, \qquad (2) \end{aligned}$$

where $\mathbf{1}(\cdot)$ is an indicator function such that $\mathbf{1}(\chi) = 1$ if χ is true and 0 otherwise, and $F_\epsilon^{-1}(y/Q)$ describes the quality level associated with the yth unit (ranked by quality) among Q. In (2), the first two terms capture the expected revenue for accepted units and expected penalty cost for rejected units, respectively, and the third term captures the expected salvage value, which depends on the supplier's allocation scheme A^{π_ϵ} and associated quality levels.[4]

3. PRELIMINARIES

3.1. Optimal Allocation Policy

The following proposition proves that, as one might expect, allocating the best q units is optimal under vague, but not necessarily exact, specification.

PROPOSITION 1: Suppose Assumption 1 holds. Let $[0, Q]$ be the set of output ranked in ascending quality order. (a) The optimal allocation policy π_ϵ defines a single subset $A^{\pi_\epsilon} \subseteq [0, Q]$. (b) The optimal subset A^{π_ϵ} is defined by a single vector $\vec{a} = (a^0, a^1)$ such that $a^1 = a^0 + q$. (c) Under exact specification, the optimal allocation sets $a^0 = \min\{F_\epsilon(\underline{l})Q, Q - q\}$. (d) Under vague specification, the optimal allocation sets $a^0 = Q - q$, that is, it is optimal to allocate the best q units.

Part (a) of Proposition 1 proves that there is no gap in the optimal allocation, and the explicit form of which is given by part (b). Part (c) tells us that, with exact specification, the supplier does not necessarily allocate the best q units to the buyer; instead, the supplier allocates the output in ascending quality order starting with the unit that just meets the buyer's quality specifications. In contrast, part (d) says that the supplier always allocates the best q units under vague specification. Hence, for any given production lot size Q, the expected quality in vague specification weakly dominates that in exact specification.

3.2. Structural Property of the Supplier Objective Function

Leveraging part (b) of Proposition 1, we can rewrite (2) as

[4] In (2), the supplier incurs a penalty cost (but not salvage value) on the fraction of q units delivered to the buyer but failed to meet the quality expectations. This is the case when penalty cost subsumes salvage value. If both salvage values and penalty costs are simultaneously incurred on rejected units, then (i) if $s'(x) = 0$ all results continue to hold, and (ii) if $s'(x) > 0$ all results in Section 4.1 continue to hold, but subsequent analysis in Sections 4.2 and 4.3 becomes intractable.

$$v(Q|\pi_\epsilon) = \mathsf{E}_\epsilon \int_x \left(w \int_{a_0}^{a_0+q} \mathbf{1}\left(F_\epsilon^{-1}(y/Q) \geq x\right) dy \right.$$
$$\left. - p \int_{a_0}^{a_0+q} \mathbf{1}\left(F_\epsilon^{-1}(y/Q) < x\right) dy \right) dG(x)$$
$$+ \mathsf{E}_\epsilon \int_0^{a_0} s\left(F_\epsilon^{-1}(y/Q)\right) dy$$
$$+ \mathsf{E}_\epsilon \int_{a_0+q}^{Q} s\left(F_\epsilon^{-1}(y/Q)\right) dy - cQ. \tag{3}$$

In what follows, we characterize the supplier's optimal production decision under vague specification. The analysis of exact specification is somewhat similar, and, for succinctness, we do not present the details here but they can be found in the Appendix. Leveraging part (d) of Proposition 1, we prove that the supplier's objective function (3) is well-behaved.

PROPOSITION 2: Suppose Assumption 1 holds. With vague specification, the supplier's objective $v(Q|\pi_\epsilon)$ is concave. Furthermore, the optimal (interior) production lot size satisfies

$$\mathsf{E}_\epsilon \left[(w+p) \int_{F_\epsilon^{-1}(1-q/Q)}^{F_\epsilon^{-1}(1)} \overline{F}_\epsilon(x) dG(x) \right.$$
$$\left. + \int_0^{1-q/Q} s\left(F_\epsilon^{-1}(z)\right) dz + (q/Q) s\left(F_\epsilon^{-1}(1-q/Q)\right) \right] = c. \tag{4}$$

Proposition 2, and particularly expression (4), helps to investigate the implications of specification vagueness on the supplier's optimal production decisions and the expected delivery quality.

4. IMPLICATIONS OF VAGUENESS IN QUALITY SPECIFICATION

In this section, we study the implications of exact and vague specifications on the supplier's expected delivery quality. For any given production quantity Q, define the expected delivery quality under vague specification as

$$\tau(Q) = \mathsf{E}_\epsilon \int_{Q-q}^{Q} F_\epsilon^{-1}(y/Q) dy. \tag{5}$$

Similarly, for any given Q, the expected delivery quality under exact specification is $\tau^E(Q) = \mathsf{E}_\epsilon \int_{\min\{F_\epsilon(\underline{l})Q, Q-q\}}^{\min\{F_\epsilon(\underline{l})Q+q, Q\}} F_\epsilon^{-1}(y/Q) dy$. All else being equal, then, $\tau(Q) - \tau^E(Q) \geq 0$, which is consistent with our earlier observation that the expected delivery quality under vague specification weakly dominates that under exact specification. For notational ease, we simply use τ (τ^E) to denote the expected delivery quality under optimal production quantity. Note that τ depends on quality specification $G(\cdot)$ through (4).

4.1. General Properties

We first study whether the supplier will choose a sufficiently large production lot size Q to guarantee that at least q units will satisfy the minimum quality level $\underline{l}$. The following proposition suggests that an exact specification cannot guarantee (a sufficiently large Q such that) at least q units exceed quality level $\underline{l}$, but a vague specification may lead to at least q units exceeding quality level $\underline{l}$.

PROPOSITION 3: Suppose the random production shocks ϵ is nondegenerate, that is, ϵ is not deterministic. (a) With exact specification, there exist some realizations of ϵ such that $\overline{F}_\epsilon(\underline{l}) < q/Q$. (b) With vague specification, the supplier may choose a production lot size Q such that $\overline{F}_\epsilon(\underline{l}) \geq q/Q$ for any realization of ϵ.

Part (a) of Proposition 3 says that the supplier's quantity provision will be insufficient to fully satisfy the delivery requirement under exact specification, that is, a fraction of delivery (for some realized ϵ) will not satisfy the buyer's quality specifications. Although it may seem counter-intuitive at first, this is not surprising: the supplier needs to balance the cost of a larger production run (to increase delivery quality) with the benefit of a smaller run (but with some quality penalty cost). In contrast, part (b) says that with vague specification the supplier may choose a sufficiently high-production lot size Q such that the delivery quality (of q units) always exceeds $\underline{l}$.

REMARK 3: Part (b) of Proposition 3 is an existence result, and additional conditions are required for part (b) to be nonvacuous and comparable to part (a). In particular, let ϵ_l and ϵ_h denote the production shocks that result in least and most favorable quality outputs, respectively. Define $\widehat{Q}$ as the unique solution to $F_{\epsilon_l}^{-1}(1 - q/\widehat{Q}) = \underline{l}$. Then, following the proof of Proposition 3 (see Appendix), a sufficient condition for part (b) to be nonvacuous is $F_{\epsilon_h}^{-1}(1 - q/\widehat{Q}) < \infty$ and $G(F_{\epsilon_h}^{-1}(1 - q/\widehat{Q})) < 1$, that is, the quality level of the qth unit (ranked in descending quality order) is finite, and there is a positive probability that the minimum quality level imposed by the buyer may be greater than the quality of the qth unit. If in addition $G(\cdot)$ is symmetric with mean $\underline{l}$, then part (a) and (b) of Proposition 3 are comparable as they have identical expected quality level $\underline{l}$. These conditions are fairly reasonable: Eq. (1), for example, readily satisfies the above conditions.

With some vague specification, then, the supplier may always have more than q units with quality level exceeding $\underline{l}$. Note that this does not guarantee that all q units will be accepted by the buyer, because the minimum quality level imposed at delivery can be greater than $\underline{l}$. The buyer may

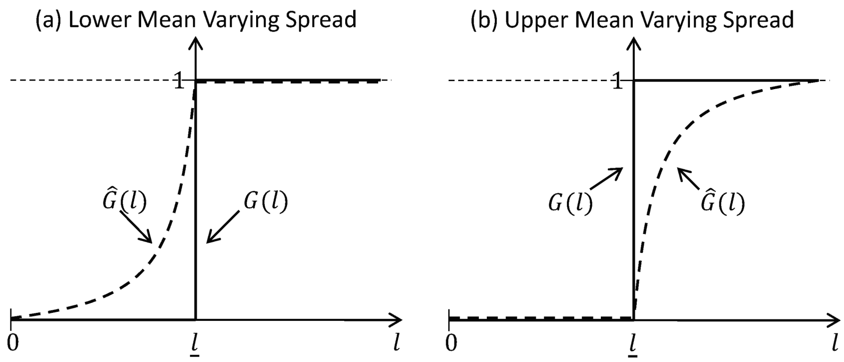

Figure 2. Mean varying spread of vague specification.

therefore still incur a quantity shortage despite the fact that all q units exceed quality level $\underline{l}$. Nevertheless, having more than q units exceed quality level $\underline{l}$ is more desirable to the buyer, who could subsequently use such higher (and achievable) quality standard for its regular orders with the supplier.

We note that Proposition 3 does not imply that vague specification always leads to a higher expected delivery quality than the exact specification does. It is possible that in vague specification a fraction of delivery will never satisfy $\underline{l}$ even under the most favorably realized ϵ. Hence, from the buyer's perspective, vague specification may or may not outperform exact specification, depending on system parameters.

The key to a more precise understanding of specification vagueness lies in characterizing how specification vagueness affects the optimal production quantity Q, as the expected delivery quality $\tau(Q)$ is a monotone increasing function in Q. It is worth pointing out, however, that with exact specification the expected delivery quality may not increase in Q, because the supplier may not always allocate the best q units (Proposition 1(c)). This can have a magnitude, but not directional, impact on Proposition 4 in next section under certain conditions, see proof of Proposition 4 for more details.

In what follows, we consider two distinct classes of vague specification: $\widehat{G}(\cdot)$ being a mean-varying spread of $G(\cdot)$ and $\widehat{G}(\cdot)$ being a mean-preserving spread of $G(\cdot)$.

4.2. Vague Specification through Mean Varying Spread

Let $\widehat{G}(\cdot)$ be a mean-varying spread of $G(\cdot)$, such that the quality requirement prescribed by $\widehat{G}(\cdot)$ is more uncertain than that by $G(\cdot)$. We consider two special types of mean varying spread: lower mean varying spread such that $G(\cdot)$ first-order stochastically dominates (FSD) $\widehat{G}(\cdot)$, and upper mean varying spread such that $\widehat{G}(\cdot)$ FSD $G(\cdot)$. See Fig. 2 (dashed line denote $\widehat{G}(\cdot)$).

With lower mean varying spread, $\widehat{G}(l) \geq G(l)$ for $l < \underline{l}$ and $\widehat{G}(l) = G(l)$ for $l \geq \underline{l}$. Hence, the quality expectation from the buyer is lower under $\widehat{G}(\cdot)$ than that under $G(\cdot)$. Conversely, with upper mean varying spread, $\widehat{G}(l) = G(l)$ for $l \leq \underline{l}$ and $\widehat{G}(l) \leq G(l)$ for $l > \underline{l}$. Hence, the quality expectation is higher. One therefore expects that all else being equal the expected delivery quality is lower in the former case and is higher in the latter case.

PROPOSITION 4: All else being equal, (a) a lower mean varying spread in $G(\cdot)$ decreases the expected delivery quality. (b) An upper mean varying spread in $G(\cdot)$ increases the expected delivery quality.

The above proposition tells us that an upper mean varying spread increases in the expected delivery quality. This is not surprising because in this case the quality expectation from the buyer is higher under $\widehat{G}(\cdot)$ than that under $G(\cdot)$. This result is consistent with the fact that in this case $\widehat{G}(\cdot)$ first-order stochastically dominates $G(\cdot)$. Proposition 4 suggests that all else being equal an upper mean varying spread in specification vagueness is more effective in motivating the supplier to increase production quantity provision and hence the expected delivery quality.

In what follows, we study the more interesting (and more challenging) question of whether specification vagueness is still attractive when the mean quality expectation is identical between the exact and the vague specification.

4.3. Vague Specification through Mean Preserving Spread

Let $\widehat{G}(\cdot)$ be a mean-preserving spread of $G(\cdot)$, such that $G(\cdot)$ second-order stochastically dominates (SSD) $\widehat{G}(\cdot)$, that is, the quality requirement prescribed by $\widehat{G}(\cdot)$ is more uncertain than that by $G(\cdot)$. Figure 3 depicts two classes of mean preserving spread: the elementary increase in vagueness (Fig. 3a) and the more general SSD case (Fig. 3b).

The elementary increase in vagueness is a fairly special type of mean-preserving spread, so we omit its analysis here and focus instead on the general mean preserving spread case. A detailed treatment of the elementary increase in vagueness can be found in an unabridged version of the article. Note

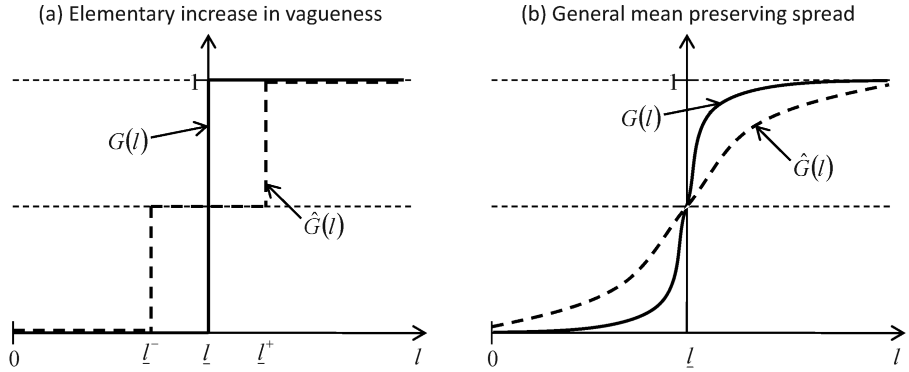

Figure 3. Mean preserving spread of vague specification.

that the analysis of elementary increase in vagueness does offer a useful hint: the curvature (e.g., concave or convex) of the quality distribution $F_\epsilon(\cdot)$ has a significant impact on the expected delivery quality. For general mean-preserving spread, it turns out, we need a more nuanced measure (other than concavity or convexity) to describe the curvature of $F_\epsilon(\cdot)$.

DEFINITION 1: For any given quality distribution $F_\epsilon(x)$, define $k_F(x) = -\frac{d}{dx}\log f_\epsilon(x)$ as proneness to low-quality outcome (PLQ), where $f_\epsilon(x)$ is the derivative of $F_\epsilon(x)$ with respect to x.

The measure $k_F(x)$ describes the curvature of $F_\epsilon(x)$, and it is also commonly known as absolute risk aversion (Pratt [35]) in the utility function context. A higher $k_F(x)$ is associated with a higher fractile of low quality output, that is, is more prone to low-quality outcome. A particularly useful family of distribution functions is the one that exhibits constant PLQ, that is, $k_F(x) = \lambda$, where λ can be positive or negative. This family of quality distribution serves as a useful benchmark such that we can partially characterize effect of specification vagueness even if $F_\epsilon(\cdot)$ does not exhibit constant PLQ.

In this section, we focus on the class of mean preserving spreads such that $G(\cdot)$ and $\widehat{G}(\cdot)$ cross only once, where $\widehat{G}(\cdot)$ is a mean-preserving spread of $G(\cdot)$. We refer to $G(\cdot)$ and $\widehat{G}(\cdot)$ as simply related (Hammond [19]), which is also known as second-order stochastic dominance with tail-dominance (Eeckhoudt and Hansen [12]), or simply referred to as single-crossing mean preserving spread. The above definition is easily satisfied, for example, when $G(\cdot)$ satisfies (1) and $\widehat{G}(\cdot)$ is generated by setting $\widehat{\sigma} > \sigma$ in (1). When the two distributions are simply related, let $\widehat{G}(\cdot) \succsim_L G(\cdot)$ denote the fact that all else being equal a low-quality delivery is more likely to be accepted by $\widehat{G}(\cdot)$ than by $G(\cdot)$.

The following proposition tells us that there exists a threshold level of PLQ, $\overline{\lambda}$, such that increasing specification vagueness leads to higher quantity provision and hence higher expected delivery quality for any production process with a PLQ greater than $\overline{\lambda}$ (less capable). Let $\tau_{\widehat{G}}$ and τ_G denote the expected delivery quality (see (5)) under quality requirements $\widehat{G}(\cdot)$ and $G(\cdot)$, respectively. In addition, define l_A as the first point where $d\widehat{G}(l) - dG(l)$ changes sign from positive to negative at $l = l_A$. Similarly, define l_B as the last point where $d\widehat{G}(l) - dG(l)$ changes sign from negative to positive at $l = l_B$.

PROPOSITION 5: Suppose $k_F(x) = \lambda$ and $\widehat{G}(\cdot) \succsim_L G(\cdot)$. (a) If $F_\epsilon^{-1}(1 - q/Q) \le l_A$, there exists a $\overline{\lambda} > 0$ such that $\tau_{\widehat{G}} > \tau_G$ for any $\lambda > \overline{\lambda}$. (b) If $F_\epsilon^{-1}(1) > l_B$, there exists a $\underline{\lambda} < 0$ such that $\tau_{\widehat{G}} < \tau_G$ for any $\lambda < \underline{\lambda}$.

Proposition 5 tells us that if the supplier's quality distribution exhibits significant PLQ (less capable), then a vague specification may induce the supplier to choose a larger quantity provision with increased delivery quality. Conversely, if the supplier's quality distribution does not exhibit much PLQ (more capable), then a more precise (exact) specification is more effective.

The following proposition further strengthens Proposition 5 by proving that there exists a unique switching point in PLQ, such that specification vagueness leads to higher (lower) quantity provision (and hence higher expected delivery quality) for any production process exhibiting higher (lower) PLQ.

PROPOSITION 6: Suppose $k_F(x) = \lambda$ and $\widehat{G}(\cdot) \succsim_L G(\cdot)$. If (a) $F_\epsilon^{-1}(1 - q/Q) < l_A$, (b) $F_\epsilon^{-1}(1) > l_B$, and (c) $\widehat{G}(F_\epsilon^{-1}(1 - q/Q)) - G(F_\epsilon^{-1}(1 - q/Q)) < \iota$ where ι is an arbitrarily small nonnegative constant, then there exists a λ^s such that $\tau_{\widehat{G}} > \tau_G$ for $\lambda > \lambda^s$ and $\tau_{\widehat{G}} < \tau_G$ for $\lambda < \lambda^s$.

Proposition 6 tells us that under the regulatory conditions (a)–(c), there exists a unique threshold PLQ (λ^s) such that a more exact specification is appropriate for any production process that is less prone to low quality ($\lambda < \lambda^s$), and a more vague specification is effective for any production process that is more prone to low quality ($\lambda > \lambda^s$).

Up to now, we have focused on the case where the quality distribution has a constant PLQ, that is, $k_F(x)$ being a constant. Next, we prove that the above characterization still holds even if the quality distribution does not have a constant PLQ.

PROPOSITION 7: Suppose $\widehat{G}(\cdot)$ and $G(\cdot)$ are simply related and $\widehat{G}(\cdot) \succsim_L G(\cdot)$. Suppose (a) $F_\epsilon^{-1}(1-q/Q) < l_A$, (b) $F_\epsilon^{-1}(1) > l_B$, (c) $\widehat{G}(F_\epsilon^{-1}(1-q/Q)) - G(F_\epsilon^{-1}(1-q/Q)) < \iota$ where ι is an arbitrarily small nonnegative constant, and (d) $\widehat{G}(F_\epsilon^{-1}(1)) \approx 1$. (i) If $\tau_{\widehat{G}} > \tau_G$ with $k_{F_1}(x)$, then $\tau_{\widehat{G}} > \tau_G$ for any $k_{F_2}(x) > k_{F_1}(x)$. (ii) If $\tau_{\widehat{G}} < \tau_G$ with $k_{F_1}(x)$, then $\tau_{\widehat{G}} < \tau_G$ for any $k_{F_2}(x) < k_{F_1}(x)$.

Proposition 7 tells us that even if the quality distribution function does not exhibit constant PLQ, one can still use the constant case as a benchmark. For example, if the buyer prefers a vague specification when the supplier's quality distribution has a constant PLQ, then the buyer will continue to prefer vague specification for any quality distribution that exhibit higher PLQ, even if the buyer does not know the exact form of the quality distribution. This result is very useful in situations where it is difficult to precisely assess the supplier's quality distribution function.

It is worth pointing out that the above analysis only partially characterizes the attractiveness of the exact versus vague specification for a realized production shock ϵ. Because ϵ influences quality distribution's PLQ outcome ($k_F(x)$), the buyer's expected preference toward exact or vague specification is therefore weighted by the realized production shocks ϵ.

In summary, the preceding analysis (Sections 4.1–4.3) partially establishes that all else being equal a vague specification can be preferred to the exact specification when the supplier's quality distribution exhibits more PLQ. Also, it is possible that a vague specification can achieve 100% quality conformance (i.e., all delivered q units exceed $\underline{l}$ in quality), which cannot be achieved by an exact specification. A buyer could therefore use vague specification to "discover" achievable quality level in the initial stage of sourcing a new product, and subsequently use the discovered quality level to set more exact quality requirements for its regular orders.

Part of the intuition of the above result is that vague specification keeps the supplier "on its toes," as the downside penalty can be significant if the quality requirement turns out to be high. To mitigate the uncertain aspect of the quality requirement, the supplier may find it beneficial to produce somewhat more (and salvage the leftovers) than to produce less (and incur stiff penalties). This is particularly true when the supplier's production cost is low, where the supplier can afford to ramp up production quantity to meet the buyer's quality specifications. Thus, the fact that vague specification can be attractive is driven by two factors: (1) the supplier's production process is prone to low-quality outcomes (because otherwise it is not necessary for the supplier to build up production quantity), and (2) the supplier's production cost is not too high (because otherwise it will not be able to afford quantity provision).

REMARK 4: Although our earlier analysis suggests that introducing some vagueness to quality specification may improve the expected delivery quality from the supplier, too much vagueness (extreme mean preserving spread) will always result in lower expected quality. This can be seen, for example, by substituting $G(l) = 0.5$ (for any l) into (4) such that the left hand side of (4) is always less than the right hand side, suggesting that with extreme vague specification the production quantity Q exactly equals to q (or $Q = 0$ if we allow the supplier not to participate).[5] In other words, the supplier does not build in any buffer quantity to reduce the expected quality shortfall. Part of the intuition is that, with too much vagueness, the supplier could not reasonably assure that its delivery will be accepted even if it builds a very large quantity provision. The supplier in this case may simply "give up" as opposed to "stand on its toes."

5. DISCUSSION AND CONCLUSION

We analytically establish in Section 4 that specification vagueness could motivate the supplier to increase its production quantity provision and hence a higher expected delivery quality. In what follows, we discuss how our results could be influenced by a number of practical considerations. Note that for reasons of space, we only summarize our key observations below. Detailed analytical treatments and numerical experiments for this section can be found in an unabridged version of the article.

5.1. Implementation and Scope of Application

As alluded to in the introduction, vague specification can be effective when the buyer is in the early stage of sourcing a new product, where precise information about suppliers' cost and capability is not readily available. In such cases, specification vagueness could help the buyer to discover achievable delivery quality level, which can subsequently be used to set more concrete quality requirements for regular production orders. Our discussion with industry practitioners also confirms that vague specifications are more likely to be observed at early stages of sourcing relationships.

A practical issue with vague specification is how specification vagueness $G(\cdot)$ and production capability $F(\cdot)$ could

[5] The limit of the mean preserving spread is a uniform distribution. Using (1), for example, we have $\lim_{\sigma\to\infty} G(l) = 0.5$ for any l.

be inferred by the supplier and the buyer, respectively. As we discussed in Section 2.1, the supplier could form a crude estimation of $G(\cdot)$ by synthesizing several sources of information: the buyer's implicit (albeit vague) quality requirements, the perceived "toughness" of the buyer, the average quality requirements from other buyers, and industry benchmark reports. If the average quality requirements from other buyers are highly variable and the focal buyer provides little guidance, for example, the supplier's estimation of $\widehat{G}(\cdot)$ will have a higher variance (but with the same mean) than what would have been estimated had the focal buyer stipulated more precise requirements. This corresponds to the scenario where $\widehat{G}(\cdot)$ forms a mean-preserving spread of $G(\cdot)$.

On the other hand, the buyer can easily calculate the PLQ index if the supplier discloses its production capability information $F(\cdot)$. This may or may not happen in practice, and if not, the buyer could resort to market intelligence reports to estimate a baseline level $F(\cdot)$ and the corresponding PLQ index for the general supplier pool. The buyer could then use such benchmark to gauge whether the particular supplier is more or less capable than the benchmark one. As our analysis does not depend on precise forms of $F(\cdot)$, such crude estimation should be suffice for the buyer to decide whether it is advantageous to adopt vague specification.

Although in theory vague specification can also be used with regular production orders, one may encounter practical and legal issues implementing vague specification, especially in a multitier supply chain where quality often exhibits interdependencies among supply chain members. A buyer, for example, may have difficulty in rejecting a large order from the supplier (who may initiate legal challenges), or accepting a large order with vague specification and pass on to its downstream partners for further processing (downstream partners may insist on clear specifications based on industry standards). As such, one should view our results with caution when large, regular production orders are involved and the product is subject to further processing down the supply chain.

The supplier may also strategically cope with vague specification with multiple production runs. If the buyer places a small test order, however, the supplier may find it uneconomical (due to fixed setups) to further split the order for multiple production runs. This is unlikely to be a concern, therefore, if vague specification is used within the appropriate guidelines discussed above. In contrast, if the buyer use vague specification for large, regular production orders, the supplier could thwart the buyer's vague specification through multiple production runs, such that the supplier could deliver an initial batch and observe the realized acceptance quality level. The supplier can then produce all subsequent batches using the observed quality level, just as if the buyer has stipulated an exact quality requirement. This is somewhat similar to the scenario described in Terwiesch and Loch [40], in which the supplier builds "prototypes" for the buyer.

5.2. Supplier Improvement Effort

Besides quantity provision, the supplier can also invest in equipment, training, and research and development to improve process reliability or quality distributions to meet the buyer's quality requirement. Using the notion of the first-order stochastic dominance, we find that the supplier's optimal improvement effort is in general not monotonic in specification vagueness. Nevertheless, we observed that, when production cost is high and production is prone to low-quality outcomes, the supplier may exert quite significant efforts to improve its production process under vague specification. Hence, vague specification can be quite effective to motivate the supplier to improve production process when it has higher cost and low reliability.

Although it is intuitive that any improvement effort by the supplier can lead to improved delivery quality, it is unclear whether a higher improvement effort always yields higher return on expected quality, because the supplier may subsequently build less quantity provision due to increased reliability. We find that, with exact specification, even if the supplier exerts significant improvement effort, such improvement effort does not translate into higher expected delivery quality. The reason is that the supplier trades off improvement effort with production quantity provision: a higher improvement effort allows the supplier to reduce its quantity provision and yet achieve similar quality level. In other words, with exact specification the benefit of improvement effort mainly accrues to the supplier.

In contrast, with moderate specification vagueness (or even significant specification vagueness when production cost is low), supplier's improvement effort yields more pronounced quality benefits. We find that a moderate amount of specification vagueness can be quite beneficial. Therefore, even if specification vagueness may not lead to a large improvement effort, it extracts more "value" from the supplier, resulting in higher delivery quality. To summarize, we find that increasing specification vagueness typically do not always spur the supplier's improvement effort, but a moderate amount of specification vagueness seems to be a quite robust choice to improve delivery quality.

5.3. Pay a Premium or Subsidy for Quality

One might conjecture that the buyer can always pay a premium for higher delivery quality. Although this is largely true for a capable supplier, we find that adjusting payment terms may have limited effect on the expected delivery quality under exact specification, especially when the supplier's production process is less capable. From a managerial point of view, then, using vague specification can be a Pareto improvement over increasing unit payment w, because the former approach may achieve similar quality level without costly increase in unit payment to the supplier. Similarly, using vague specification can also be superior to raising unit penalty p, because the latter approach increases the supplier's default risk (or the supplier may choose not to participate).

In contrast, a subsidy on leftovers can be effective in raising the expected delivery quality. A key reason is that, besides making the supplier's production cheaper, with subsidy the buyer can always obtain the best q units, just as with vague specification. The increased delivery quality, however, should be balanced with the increased subsidy cost, which may reduce the attractiveness of full subsidy from the buyer's perspective. We note that the buyer does not have to fully subsidize all leftovers, and that a partial subsidy for limited number of leftovers can achieve an expected delivery quality level close to that under full subsidy. From the buyer's point of view, then, a partial subsidy policy can be more attractive because it is less costly and yet achieves similar quality levels.

In summary, subsidies on excess units can be more effective than raising unit payment w or penalty p, and a sufficiently high subsidy (under exact specification) can outperform vague specification in terms of the expected delivery quality. Subsidy, however, does require the buyer to share more supply risk (e.g., coping with excess units), and therefore the buyer may or may not prefer subsidy to vague specification, depending on the tradeoff between increased quality level and the added cost.

5.4. Limitations and Future Extensions

In this work, we largely ignore the role of demand uncertainty on the effectiveness of vague specification. Part of the reason is that demand uncertainty may only have limited impact on the supplier's incentives brought on by vague specification, because the buyer needs to determine its order quantity ex ante. From the supplier's perspective, then, demand uncertainty is insulated by the buyer. Demand uncertainty also has limited impact if the buyer uses vague specification in an initial small order to test the supplier's capability.

On the other hand, demand uncertainty may influence the attractiveness of vague specification in two important ways. First, if the supplier anticipates future demand uncertainty when completing the buyer's initial test order, the supplier's incentive to build quantity provision could be influenced by how large the future demand is going to be. Significant demand uncertainty will cloud the supplier's estimation of future orders, and thus skew the supplier's quantity provision incentives. Such effect could be particularly acute if the supplier is risk averse. Second, demand uncertainty may also influence the buyer's bargaining power with the supplier. With significant demand uncertainty, the buyer may not be able to convince the supplier of any significant future business, potentially depriving the buyer of its decision right. A significant loss of decision right may render vague specification un-implementable by the buyer. In such cases, one may have to resort to ex-post bilateral Nash bargaining framework to investigate the impact of vague specification on the supplier's quantity provision and expected delivery quality.

Given that quality requirements play a central role in any business transactions, we hope that future work on the above extensions, for example, demand uncertainty, risk aversion, and bilateral bargaining, will further our understanding of how quality requirements influence the supplier's quality provisions.

APPENDIX: PROOFS

PROOF OF PROPOSITION 1: Part (a). We prove the proposition statement by contradiction. If π_ϵ is an optimal allocation policy and it indicates A^{π_ϵ} as a sum of at least two disjoint subsets $A_1^{\pi_\epsilon}$ and $A_2^{\pi_\epsilon}$, then there exists at least one element $x \notin A_1^{\pi_\epsilon} \cup A_2^{\pi_\epsilon}$, such that $\arg\min_{y\in A_1^{\pi_\epsilon}} F_\epsilon^{-1}(y/Q) < F_\epsilon^{-1}(x/Q) < \arg\min_{y\in A_2^{\pi_\epsilon}} F_\epsilon^{-1}(y/Q)$. Now construct an allocation policy π'_ϵ by replacing the smallest element in $A_1^{\pi_\epsilon}$ with x. By (2), $v(Q|\pi'_\epsilon) \ge v(Q|\pi_\epsilon)$ for any given Q. If x is a set instead of an element, then continue the construction of π'_ϵ until the gap between $A_1^{\pi_\epsilon}$ and $A_2^{\pi_\epsilon}$ is filled. Thus, we have obtained an alternative allocation policy π'_ϵ which indicates a single set $A^{\pi'_\epsilon}$ that is no worse than π_ϵ. Part (b). Follows directly from part (a). Part (c). By part (b), A^{π_ϵ} is defined by a single vector $\vec{a} = (a^0, a^1)$ such that $a^1 = a^0 + q$. By (3) and after some algebra, we have

$$\begin{aligned}\frac{\partial v(Q|\pi_\epsilon)}{\partial a^0} &= \mathsf{E}_\epsilon\left[(w+p)\left(G\left(F_\epsilon^{-1}((a^0+q)/Q)\right) - G\left(F_\epsilon^{-1}(a^0/Q)\right)\right)\right] \\ &\quad + \mathsf{E}_\epsilon\left[s\left(F_\epsilon^{-1}(a^0/Q)\right) - s\left(F_\epsilon^{-1}((a^0+q)/Q)\right)\right]. \qquad \text{(A1)}\end{aligned}$$

The proposition statement is true if $\partial v(Q|\pi_\epsilon)/\partial a^0 \le 0$ for any $a^0 \ge F_\epsilon(\underline{l})Q$ under exact specification. Note $a^0 \ge F_\epsilon(\underline{l})Q \Rightarrow F_\epsilon^{-1}(a^0/Q) \ge \underline{l} \Rightarrow G(F_\epsilon^{-1}(a^0/Q)) = 1 \Rightarrow G(F_\epsilon^{-1}((a^0+q)/Q)) = 1$. Substituting this into (A1), we have $\partial v(Q|\pi_\epsilon)/\partial a^0 = \mathsf{E}_\epsilon[s(F_\epsilon^{-1}(a^0/Q)) - s(F_\epsilon^{-1}((a^0+q)/Q))] \le 0$, where the inequality follows from the fact that $s'(x) \ge 0$. Hence, the supplier's expected profit increases in a^0 up to the point at $a^0 = F_\epsilon(\underline{l})Q$ and then decreases beyond that point. It follows that it is optimal to set $a^0 = \min\{F_\epsilon(\underline{l})Q, Q-q\}$ under the exact specification. Part (d). The proposition statement is true if $v(Q|\pi_\epsilon)$ increases in a^0. Let $x_1 = F_\epsilon^{-1}(a^0/Q)$ and $x_2 = F_\epsilon^{-1}((a^0+q)/Q)$. Then, (A1) can be simplified to $\partial v(Q|\pi_\epsilon)/\partial a^0 = \mathsf{E}_\epsilon[\{(w+p)G(x_2) - s(x_2)\} - \{(w+p)G(x_1) - s(x_1)\}] \ge 0$, where the inequality follows from Assumption 1 and the fact that $x_2 > x_1$. □

PROOF OF PROPOSITION 2: Substituting $a^0 = Q - q$ into (3) (and after some algebra), we have

$$\begin{aligned}v(Q|\pi_\epsilon) = -cQ + \mathsf{E}_\epsilon\Bigg[&(w+p)qG\left(F_\epsilon^{-1}(1-q/Q)\right) - pq \\ &+ (w+p)Q\int_{x>F_\epsilon^{-1}(1-q/Q)}^{x\le F_\epsilon^{-1}(1)} \overline{F}_\epsilon(x)dG(x) \\ &+ Q\int_0^{1-q/Q} s\left(F_\epsilon^{-1}(z)\right)dz\Bigg]. \qquad \text{(A2)}\end{aligned}$$

It follows that

$$\begin{aligned}v'(Q|\pi_\epsilon) = -c + \mathsf{E}_\epsilon\Bigg[&(w+p)\int_{x>F_\epsilon^{-1}(1-q/Q)}^{x\le F_\epsilon^{-1}(1)} \overline{F}_\epsilon(x)dG(x) \\ &+ \int_0^{1-q/Q} s\left(F_\epsilon^{-1}(z)\right)dz + (q/Q)s\left(F_\epsilon^{-1}(1-q/Q)\right)\Bigg], \qquad \text{(A3)}\end{aligned}$$

and $v''(Q) = \mathsf{E}_\epsilon\left[-(w+p)\frac{q^2}{Q^3}\frac{G'(F_\epsilon^{-1}(1-q/Q))}{F'_{l|\epsilon}[F_\epsilon^{-1}(1-q/Q)]} + \frac{q^2}{Q^3}\frac{s'(F_\epsilon^{-1}(1-q/Q))}{F'_{l|\epsilon}[F_\epsilon^{-1}(1-q/Q)]}\right] < 0$, where the inequality follows from Assumption 1. Note that (4) is obtained by setting (A3) equal to zero. □

PROOF OF PROPOSITION 3: Part (a). First we adapt (3) to the exact specification. By Proposition 1, the optimal allocation sets $a^0 = \min\{F_\epsilon(\underline{l})Q, Q-q\}$. Let $\overline{\epsilon}$ be the unique solution to $F_{\overline{\epsilon}}(\underline{l}) = 1 - q/Q$. We have

$$\begin{aligned}
v(Q|\pi_\epsilon) &= \mathsf{E}_{\epsilon|\epsilon\le\overline{\epsilon}}\int_x\left(w\int_{Q-q}^{Q}\mathbf{1}\left(F_\epsilon^{-1}(y/Q)\ge x\right)dy\right.\\
&\quad\left.-p\int_{Q-q}^{Q}\mathbf{1}\left(F_\epsilon^{-1}(y/Q)<x\right)dy\right)dG(x)\\
&\quad+\mathsf{E}_{\epsilon|\epsilon\le\overline{\epsilon}}\int_0^{Q}s\left(F_\epsilon^{-1}(y/Q)\right)\mathbf{1}\left(F_\epsilon^{-1}(y/Q)<x\right)dy\\
&\quad+\mathsf{E}_{\epsilon|\epsilon>\overline{\epsilon}}\int_x\left(w\int_{F_\epsilon(\underline{l})Q}^{F_\epsilon(\underline{l})Q+q}\mathbf{1}\left(F_\epsilon^{-1}(y/Q)\ge x\right)dy\right.\\
&\quad\left.-p\int_{F_\epsilon(\underline{l})Q}^{F_\epsilon(\underline{l})Q+q}\mathbf{1}\left(F_\epsilon^{-1}(y/Q)<x\right)dy\right)dG(x)\\
&\quad+\mathsf{E}_{\epsilon|\epsilon>\overline{\epsilon}}\int_0^{F_\epsilon(\underline{l})Q}s\left(F_\epsilon^{-1}(y/Q)\right)dy\\
&\quad+\mathsf{E}_{\epsilon|\epsilon>\overline{\epsilon}}\int_{F_\epsilon(\underline{l})Q+q}^{Q}s\left(F_\epsilon^{-1}(y/Q)\right)dy-cQ\\
&=\mathsf{E}_{\epsilon|\epsilon\le\overline{\epsilon}}\left[(w+p)\overline{F}_\epsilon(\underline{l})Q-pq+Q\int_0^{F_\epsilon(\underline{l})}s\left(F_\epsilon^{-1}(z)\right)dz\right]\\
&\quad+\mathsf{E}_{\epsilon|\epsilon>\overline{\epsilon}}\left[wq+Q\int_0^{F_\epsilon(\underline{l})}s\left(F_\epsilon^{-1}(z)\right)dz\right.\\
&\quad\left.+Q\int_{F_\epsilon(\underline{l})+q/Q}^{1}s\left(F_\epsilon^{-1}(z)\right)dz\right]-cQ. \qquad \text{(A4)}
\end{aligned}$$

We now prove the proposition statement by contradiction. Suppose the supplier chooses an optimal production lot size Q^* such that $\overline{F}_\epsilon(\underline{l}) \ge q/Q$ for any realized ϵ. Then, (A4) simplifies to $v(Q|\pi_\epsilon) = \mathsf{E}_\epsilon[wq + Q\int_0^{F_\epsilon(\underline{l})} s\left(F_\epsilon^{-1}(z)\right)dz + Q\int_{F_\epsilon(\underline{l})+q/Q}^{1} s\left(F_\epsilon^{-1}(z)\right)dz] - cQ$. It follows that $v'(Q|\pi_\epsilon) = \mathsf{E}_\epsilon[\int_0^{F_\epsilon(\underline{l})} s(F_\epsilon^{-1}(z))dz + \int_{F_\epsilon(\underline{l})+q/Q}^{1} s(F_\epsilon^{-1}(z))dz + (q/Q)s(F_\epsilon^{-1}(F_\epsilon(\underline{l}) + q/Q))] - c < \mathsf{E}_\epsilon[\int_0^{F_\epsilon(\underline{l})} c\,dz + \int_{F_\epsilon(\underline{l})+q/Q}^{1} c\,dz + (q/Q)c] - c = c - c = 0$. This suggests $v(Q|\pi_\epsilon)$ is strictly decreasing in Q if $\overline{F}_\epsilon(\underline{l}) \ge q/Q$ for any realized ϵ, thus contradicting the assumption that Q^* is optimal. Part (b). First, consider the special case where salvage value is constant, that is, $s(x) = s$. Then, (4) can be simplified to $\mathsf{E}_\epsilon[(w+p)\int_{F_\epsilon^{-1}(1-q/Q)}^{F_\epsilon^{-1}(1)} \overline{F}_\epsilon(x)dG(x) + s] = c$. It follows that

$$\mathsf{E}_\epsilon\left[\int_{F_\epsilon^{-1}(1-q/Q)}^{F_\epsilon^{-1}(1)}\overline{F}_\epsilon(x)dG(x)\right]=\frac{c-s}{w+p}. \qquad \text{(A5)}$$

Consider the LHS of (A5) inside the expectation operator, we have

$$\begin{aligned}
\int_{F_\epsilon^{-1}(1-q/Q)}^{F_\epsilon^{-1}(1)}\overline{F}_\epsilon(x)dG(x) &= \int_{F_\epsilon^{-1}(1-q/Q)}^{F_\epsilon^{-1}(1)}G(x)dF_\epsilon(x)\\
&\quad-(q/Q)G(F_\epsilon^{-1}(1-q/Q))\\
&\ge\int_{F_\epsilon^{-1}(1-q/Q)}^{F_\epsilon^{-1}(1)}G(F_\epsilon^{-1}(1-q/Q))dF_\epsilon(x)\\
&\quad-(q/Q)G(F_\epsilon^{-1}(1-q/Q))=0,
\end{aligned}$$

where the inequality is strict if $G(F_\epsilon^{-1}(1-q/Q)) < 1$. Hence, even if $\overline{F}_\epsilon(\underline{l}) \ge q/Q \Rightarrow F_\epsilon^{-1}(1-q/Q) \ge \underline{l}$ for any realized ϵ, as long as $G(\underline{l}+\delta) < 1$ (where $\delta = F_\epsilon^{-1}(1-q/Q) - \underline{l}$), then (A5) can hold as an interior solution for sufficiently small $(c-s)/(w+p)$, that is, the optimal lot size Q can guarantee that $\overline{F}_\epsilon(\underline{l}) \ge q/Q$ for any realized ϵ. Now, if unit salvage value is increasing in quality level x, then it only serves to increase the optimal lot size, which makes the above analysis more likely to hold. □

PROOF OF PROPOSITION 4: We prove the proposition statement for upper mean varying spread, that is, when $\widehat{G}(\cdot)$ first-order stochastically dominates $G(\cdot)$. The case with lower mean-varying spread can be analogously proved. First note that expected delivery quality $\tau(Q)$ (see (5)) is a monotone increasing function in Q, and hence we only need to characterize the effect of $\widehat{G}(\cdot)$ on the optimal production quantity Q. By (4), specification vagueness in $G(\cdot)$ affects the optimal production quantity through the first term in LHS of (4) only. For expositional ease, define

$$\psi_G(\epsilon)=\int_{F_\epsilon^{-1}(1-q/Q)}^{F_\epsilon^{-1}(1)}\overline{F}_\epsilon(u)dG(u), \qquad \text{(A6)}$$

Because $\psi_G(\epsilon)$ is a decreasing function in Q, if specification vagueness in $G(\cdot)$ results in an increase in $\psi_G(\epsilon)$ (for any given Q), then condition (4) implies that the optimal production quantity Q must increase, which leads to increased expected delivery quality $\tau(Q)$. Now, $\psi_{\widehat{G}}(\epsilon) - \psi_G(\epsilon) = \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} \overline{F}_\epsilon(u)d[\widehat{G}(u) - G(u)] = [G(F_\epsilon^{-1}(1 - q/Q)) - \widehat{G}(F_\epsilon^{-1}(1-q/Q))]q/Q + \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} f_\epsilon(u)[\widehat{G}(u) - G(u)]du$ $= \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} f_\epsilon(u)[\widehat{G}(F_\epsilon^{-1}(1-q/Q)) - G(F_\epsilon^{-1}(1-q/Q))]du - \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} f_\epsilon(u)[G(u) - \widehat{G}(u)]du$. Now, if $F_\epsilon^{-1}(1-\frac{q}{Q}) <= \underline{l}$, $v_{\widehat{G}}(\epsilon) - v_G(\epsilon) = -\int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} f_\epsilon(u)[G(u) - \widehat{G}(u)]du > 0$. Consider $F_\epsilon^{-1}(1 - \frac{q}{Q}) > \underline{l}$. Note that $G(x) - \widehat{G}(x)$ decreases in x for $x > \underline{l}$. Hence, $\int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} f_\epsilon(u)[\widehat{G}(F_\epsilon^{-1}(1-q/Q)) - G(F_\epsilon^{-1}(1-q/Q))]du \ge \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} f_\epsilon(u)[G(u) - \widehat{G}(u)]du \Rightarrow \psi_{\widehat{G}}(\epsilon) - \psi_G(\epsilon) \ge 0$. The proposition statement then follows from the fact that $\psi_{\widehat{G}}(\epsilon) \ge \psi_G(\epsilon)$ for any realized production shock ϵ and therefore the expectation does not change the sign of $\psi_{\widehat{G}}(\epsilon) - \psi_G(\epsilon)$.

Note that when the salvage value strictly increases in quality level, that is, $s'(x) > 0$, and when $\widehat{G}(\cdot) \downarrow G(\cdot)$, the expected delivery quality is in general not continuous at the limit. To see this, note that when $s'(x) > 0$, the expected delivery quality under exact specification is $\tau^E(Q) = \mathsf{E}_\epsilon\int_{\min\{F_\epsilon(\underline{l})Q,Q-q\}}^{\min\{F_\epsilon(\underline{l})Q+q,Q\}} F_\epsilon^{-1}(y/Q)dy$. Partition ϵ into $\Omega_1 = \{\epsilon : F_\epsilon(\underline{l})Q > Q - q\}$ and $\Omega_2 = \{\epsilon : F_\epsilon(\underline{l})Q \le Q - q\}$. We have $\tau^E(Q) = \mathsf{E}_{\epsilon|\Omega_1}\int_{Q-q}^{Q} F_\epsilon^{-1}(y/Q)dy + \mathsf{E}_{\epsilon|\Omega_2}\int_{F_\epsilon(\underline{l})Q}^{F_\epsilon(\underline{l})Q+q} F_\epsilon^{-1}(y/Q)dy = Q[\mathsf{E}_{\epsilon|\Omega_1}\int_{1-q/Q}^{1} F_\epsilon^{-1}(z)dz + \mathsf{E}_{\epsilon|\Omega_2}\int_{F_\epsilon(\underline{l})}^{F_\epsilon(\underline{l})+q/Q} F_\epsilon^{-1}(z)dz]$. It follows that $\partial\tau^E(Q)/\partial Q = \mathsf{E}_{\epsilon|\Omega_1}[\int_{1-q/Q}^{1} F_\epsilon^{-1}(z)dz - (q/Q)F_\epsilon^{-1}(1-q/Q)] + \mathsf{E}_{\epsilon|\Omega_2}[\int_{F_\epsilon(\underline{l})}^{F_\epsilon(\underline{l})+q/Q} F_\epsilon^{-1}(z)dz - (q/Q)F_\epsilon^{-1}(F_\epsilon(\underline{l}) + q/Q)]$. Notice that the first term in the square bracket is positive, whereas the second term is negative, and hence an increase in Q does not necessarily improves the expected delivery quality. Now, suppose $\widehat{G}(\cdot)$ is an epsilon change from $G(\cdot)$, that is, $\widehat{G}(\cdot)$ is a limiting case of the vague specification, then the supplier allocates the best q units under $\widehat{G}(\cdot)$. Adapting (4) to the exact specification case [following similar approach as that in (A4)], one can prove that the optimal Q under $\widehat{G}(\cdot)$ is less than that under $G(\cdot)$, even as $\widehat{G}(\cdot) \downarrow G(\cdot)$. Hence, at limit the expected delivery quality between the vague and the exact specification can have a persistent gap, which is caused by the fact that the allocation policy changes from allocating the best q units to allocating from the unit that just meets the specification. Therefore, when $s'(x) > 0$ and $\widehat{G}(\cdot)$ is arbitrarily close to $G(\cdot)$, the magnitude effect of vague specification can be ambiguous. □

PROOF OF PROPOSITION 5: Given $k_F(x) = \lambda$, let $F_\epsilon(x) = 1 - \exp(-\lambda x)$. There is no loss of generality in assuming this functional form, because alternative forms of exponential family such as $e^c F_\epsilon(x) + d$, where c and d are constants, do not change the rankings of $\psi_{\widehat{G}}(\epsilon)$ and $\psi_G(\epsilon)$ - see Ref. [35] (p. 126). Part (a). We adopt the similar proof approach in Ref. [19] (p. 1054). By (A6),

$$\begin{aligned}\psi_{\widehat{G}}(\epsilon) - \psi_G(\epsilon) &= \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} \overline{F}_\epsilon(u)d\widehat{G}(u) - \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} \overline{F}_\epsilon(u)dG(u)\\ &= \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} e^{-\lambda u}[d\widehat{G}(u) - dG(u)]\\ &= \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{l_A} e^{-\lambda u}[d\widehat{G}(u) - dG(u)]\\ &\quad + \int_{l_A}^{F_\epsilon^{-1}(1)} e^{-\lambda u}[d\widehat{G}(u) - dG(u)],\end{aligned}$$

where the first term is positive because $d\widehat{G}(u) - dG(u) > 0$ for $u < l_A$. Now, if the magnitude of the first term is greater than that of the second term, then $\psi_{\widehat{G}}(\epsilon) - \psi_G(\epsilon)$ is positive regardless of whether the second term is positive or negative. This indeed is the case, because, the first term declines slower than the second term as the exponents in the first term are smaller (less than l_A) than that in the second term. Hence, for sufficiently large $\lambda > \overline{\lambda}$, the magnitude of the first term will surpass that of the second term, with a result that $\psi_{\widehat{G}}(\epsilon) - \psi_G(\epsilon) > 0$. Part (b) follows analogously from part (a). □

PROOF OF PROPOSITION 6: To capture the case where λ can be either positive or negative, let $F_\epsilon(x) = (1 - e^{-\lambda x})/\lambda$. We have

$$\begin{aligned}\psi_{\widehat{G}}(\epsilon) - \psi_G(\epsilon) &= \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} \left(1 - (1 - e^{-\lambda u})/\lambda\right)[d\widehat{G}(u) - dG(u)]\\ &= [\widehat{G}(u) - G(u)]\left(1 - (1 - e^{-\lambda})/\lambda\right)\Big|_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)}\\ &\quad + \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} e^{-\lambda u}[\widehat{G}(u) - G(u)]du\\ &= [\widehat{G}(F_\epsilon^{-1}(1 - q/Q)) - G(F_\epsilon^{-1}(1 - q/Q))]q/Q\\ &\quad + \int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} e^{-\lambda u}[\widehat{G}(u) - G(u)]du.\end{aligned}$$

By assumption (c), $0 \le [\widehat{G}(F_\epsilon^{-1}(1 - q/Q)) - G(F_\epsilon^{-1}(1 - q/Q))]q/Q < \iota \cdot q/Q < \iota$, hence we need only focus on the second term in the above expression. By Ref. [23] (p. 282), $e^{-\lambda u}$ belongs to the class of Pólya Type ∞ distribution. Also, because $\widehat{G}(\cdot)$ and $G(\cdot)$ are simply related, $\widehat{G}(u) - G(u)$ changes sign once. Hence, by Theorem 3 in Ref. [23] (p. 291), $\int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} e^{-\lambda u}[\widehat{G}(u) - G(u)]du$ changes sign at most once. By Proposition 5, $\int_{F_\epsilon^{-1}(1-\frac{q}{Q})}^{F_\epsilon^{-1}(1)} e^{-\lambda u}[\widehat{G}(u) - G(u)]du$ changes sign at least once as λ changes from $-\infty$ to ∞. The proposition statement then follows. □

PROOF OF PROPOSITION 7: We follow the similar approach in the proof of Theorem 3 in Ref. [19] (p. 1056). Part (i). Note that the proposition statement is true if $\psi_{\widehat{G}|F_2}(\epsilon) - \psi_{G|F_2}(\epsilon) \ge e^{-M}(\psi_{\widehat{G}|F_1}(\epsilon) - \psi_{G|F_1}(\epsilon))$ for any arbitrary $M > 0$. Recall that

$$\begin{aligned}\psi_{\widehat{G}|F_i}(\epsilon) - \psi_{G|F_i}(\epsilon) &= \int_{F_{i|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{i|\epsilon}^{-1}(1)} \overline{F}_{i|\epsilon}(u)d\widehat{G}(u)\\ &\quad - \int_{F_{i|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{i|\epsilon}^{-1}(1)} \overline{F}_{i|\epsilon}(u)dG(u)\\ &= \overline{F}_{i|\epsilon}(u)[\widehat{G}(u) - G(u)]\Big|_{F_{i|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{i|\epsilon}^{-1}(1)}\\ &\quad + \int_{F_{i|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{i|\epsilon}^{-1}(1)} f_{i|\epsilon}(u)[\widehat{G}(u) - G(u)]du\\ &= -o(\iota) + \int_{F_{i|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{i|\epsilon}^{-1}(1)} f_{i|\epsilon}(u)[\widehat{G}(u) - G(u)]du,\end{aligned} \tag{A7}$$

where $o(\iota) = [G(F_{i|\epsilon}^{-1}(1 - q/Q)) - \widehat{G}(F_{i|\epsilon}^{-1}(1 - q/Q))]q/Q = \iota \cdot q/Q < \iota$. Consider the last term in (A7). For any given $k_{F_i}(x)$, its corresponding quality distribution function can be expressed as $F_{i|\epsilon}(x) = \int \exp(-\int_0^x k_{F_i}(y)dy)dx$ [35, p. 126]. Hence, $f_{i|\epsilon}(x) = \exp(-\int_0^x k_{F_i}(y)dy)$. Substituting $f_{i|\epsilon}(x)$ into the last term in (A7), we have

$$\psi_{\widehat{G}|F_i}(\epsilon) - \psi_{G|F_i}(\epsilon) = -o(\iota) + \int_{F_{i|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{i|\epsilon}^{-1}(1)} e^{-\int_0^u k_{F_i}(y)dy}[\widehat{G}(u) - G(u)]du.$$

Hence, the proposition statement is proved if

$$\begin{aligned}&\int_{F_{2|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{2|\epsilon}^{-1}(1)} e^{-\int_0^u k_{F_2}(y)dy}[\widehat{G}(u) - G(u)]du\\ &\qquad \ge e^{-M}\int_{F_{1|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{1|\epsilon}^{-1}(1)} e^{-\int_0^u k_{F_1}(y)dy}[\widehat{G}(u) - G(u)]du.\end{aligned} \tag{A8}$$

Because $k_{F_2}(x) > k_{F_1}(x)$, we have $F_{2|\epsilon}^{-1}(1 - q/Q) \le F_{1|\epsilon}^{-1}(1 - q/Q)$ and $F_{2|\epsilon}^{-1}(1) \le F_{1|\epsilon}^{-1}(1)$. Also, by assumption (d), $\widehat{G}(u) - G(u) \approx 0$ for $u \ge F_{2|\epsilon}^{-1}(1)$. Therefore, (A8) is true if

$$\int_{F_{1|\epsilon}^{-1}(1-\frac{q}{Q})}^{F_{2|\epsilon}^{-1}(1)} \left(e^{-\int_0^u k_{F_2}(y)dy} - e^{-M-\int_0^u k_{F_1}(y)dy}\right)[\widehat{G}(u) - G(u)]du > 0. \tag{A9}$$

Recall that $\widehat{G}(u)$ and $G(u)$ are simply related and $\widehat{G}(u) > G(u)$ for $u < \underline{l}$ and $\widehat{G}(u) < G(u)$ for $u > \underline{l}$. Hence, if $e^{-\int_0^u k_{F_2}(y)dy} - e^{-M-\int_0^u k_{F_1}(y)dy}$ has the same sign pattern as $\widehat{G}(u) - G(u)$ then (A9) is true. Select M such that $e^{-\int_0^{\underline{l}} k_{F_2}(y)dy} = e^{-M-\int_0^{\underline{l}} k_{F_1}(y)dy}$. Because $k_{F_2}(x) > k_{F_1}(x)$, we have $\int_0^u k_{F_2}(y)dy < M + \int_0^u k_{F_1}(y)dy$ for $u < \underline{l}$ and $\int_0^u k_{F_2}(y)dy > M + \int_0^u k_{F_1}(y)dy$ for $u > \underline{l}$. Hence, $e^{-\int_0^u k_{F_2}(y)dy} - e^{-M-\int_0^u k_{F_1}(y)dy}$ indeed has the same sign pattern as $\widehat{G}(u) - G(u)$ in (A9). Part (ii) of the proposition statement can be analogously proved. □

ACKNOWLEDGMENTS

The author thanks Brian Tomlin, Jing Li, seminar participants at Cornell University, two anonymous referees, and the editor for helpful comments.

REFERENCES

[1] U.T. AG, Intelligent sourcing at detail. Available at: http://www.uster.com/en/service/intelligent-sourcing/. Last retrieved: September 18, 2012.
[2] D. Ariely, In praise of the handshake, Harv Bus Rev 89 (2011), 40–40.
[3] S. Baiman, P.E. Fischer, and M.V. Rajan, Information, contracting, and quality costs, Manag Sci 46 (2000), 776–789.
[4] BCG, Special report: New challenges for 'Made in China'. The Boston Consulting Group, Boston, MA, June, 2009.
[5] S. Beall, C. Carter, P. Carter, T. Germer, T. Hendrick, S. Jap, L. Kaufmann, D. Maciejewski, R. Monczka, and K. Petersen, The role of reverse auctions in strategic sourcing, Technical report, CAPS Research, 2003.
[6] R.E. Bellmand and L.A. Zadeh, Decision-making in a fuzzy environment, Manag Sci 17 (1970), B141–B164.
[7] R.R. Britney, Optimal screening plans for nonserial production systems, Manag Sci 18 (1972), 550–559.
[8] J. Chen, D.D. Yao, and S. Zheng, Optimal replenishment and rework with multiple unreliable supply sources, Oper Res 49 (2001), 430–443.
[9] T.Y. Choi and Y. Hong, Unveiling the structure of supply networks: case studies in Honda, Acura, and DaimlerChrysler, J Oper Manag 20 (2002), 469–493.
[10] M. Dada, N.C. Petruzzi, and L.B. Schwarz, A newsvendor's procurement problem when suppliers are unreliable, Manuf Serv Oper Manag 9 (2007), 9–32.
[11] L. Dong and B. Tomlin, Managing disruption risk: The interplay between operations and insurance, Manag Sci 58 (2012), 1898–1915.
[12] L. Eeckhoudt and P. Hansen, Mean-preserving changes in risk with tail-dominance, Theor Deci 33 (1992), 23–39.
[13] H.J. Einhorn and R.M. Hogarth, Decision making under ambiguity, J Bus 59 (1986), s225–s250.
[14] J. ERG, How to determine target yields. Available at: http://www.eng-tips.com/viewthread.cfm?qid=247080. Last retrieved: August 16, 2012.
[15] J.R. Evans and W.M. Lindsay, Managing for quality and performance excellence, 7Ed., Thomson South-Western, Mason, OH, 2008.
[16] E. Gal-Or, M. Gal-Or, and A. Dukes, Optimal information revelation in procurement schemes, Rand J Econ 38 (2007), 400–418.
[17] A. Grosfeld-Nir and Y. Gerchak, Production to order with random yields: Single-stage multiple lot-sizing, IIE Trans 28 (1996), 669–676.
[18] S.J. Grossman and O.D. Hart, The costs and benefits of ownership: A theory of vertical and lateral integration, J Polit Econ 94 (1986), 691–719.
[19] J.S. Hammond, Simplifying the choice between uncertain prospects where preference is nonlinear, Manag Sci 20 (1974), 1047–1072.
[20] O. Hart and J. Moore, Incomplete contracts and renegotiation, Econometrica 56 (1988), 755–785.
[21] O. Hart and J. Moore, A theory of debt based on the inalienability of human capital, Q J Econ 109 (1994), 841–879.
[22] I. Hwang, S. Radhakrishnan, and L. Su, Vendor certification and appraisal: Implications for supplier quality, Manag Sci 52 (2006), 1472–1482.
[23] S. Karlin, Pólya type distributions, II, Ann Math Stat 28 (1957), 281–308.
[24] M. Kaya and Ö. Özer, Quality risk in outsourcing: Noncontractible product quality and private quality cost information, Naval Res Logis 56 (2009), 669–685.
[25] B. Kazaz and S. Webster, The impact of yield-dependent trading costs on pricing and production planning under supply uncertainty, Manuf Serv Oper Manag 13 (2011), 404–417.
[26] H.L. Lee and M.J. Rosenblatt, Optimal inspection and ordering policies for products with imperfect quality, IIE Trans 17 (1985), 284–289.
[27] R.E. Levitan, The optimum reject allowance problem, Manag Sci 6 (1960), 172–186.
[28] W.S. Lim, Producer-supplier contracts with incomplete information, Manag Sci 47 (2001), 709–715.
[29] S. Liu, K.C. So, and F. Zhang, Effect of supply reliability in a retail setting with joint marketing and inventory decisions, Manuf Serv Oper Manag 12 (2010), 19–32.
[30] R. Monczka, R. Trent, and R. Handfield, Purchasing and supply chain management, South-Western, Thomson, Mason, OH, 2005.
[31] J. Mula, R. Poler, and F. Lario, Models for production planning under uncertainty: A review, Int J Prod Econ 103 (2006), 271–285.
[32] D. Murthy and L. Ma, MRP with uncertainty: A review and some extensions, Int J Prod Econ 25 (1991), 51–64.
[33] J. Ou and L.M. Wein, Sequential screening in semiconductor manufacturing, II: Exploiting lot-to-lot variability, Oper Res 44 (1996), 196–205.
[34] M.H. Peters, H. Schneider, and K. Tang, Joint determination of optimal inventory and quality control policy, Manag Sci 34 (1988), 991–1004.
[35] J.W. Pratt, Risk aversion in the small and in the large, Econometrica 32 (1964), 122–136.
[36] D.J. Reyniers and C.S. Tapiero, Contract design and the control of quality in a conflictual environment, Eur J Oper Res 82 (1995a), 373–382.
[37] D.J. Reyniers and C.S. Tapiero, The delivery and control of quality in supplier-producer contracts, Manag Sci 41 (1995b), 373–382.
[38] A. Sheopuri and E. Zemel, Costly enforcement of quality standards in decentralized supply chains. Working Paper, IBM Research Center, Yorktown Heights, NY, 2008.
[39] R. Swinney and S. Netessine, Long-term contracts under the threat of supplier default, Manuf Serv Oper Manag 11 (2009), 109–127.
[40] C. Terwiesch and C.H. Loch, Collaborative prototyping and the pricing of custom-designed products, Manag Sci 50 (2004), 145–158.
[41] J. Tirole, Incomplete contracts: Where do we stand? Econometrica 67 (1999), 741–781.
[42] O.E. Williamson, The economic institutions of capitalism, The Free Press, New York, NY, 1985.
[43] A. Wolinsky, Procurement via sequential search, J Polit Econ 113 (2005), 785–810.
[44] C.A. Yano and H.L. Lee, Lot sizing with random yields: A review, Oper Res 43 (1995), 311–334.
[45] D.D. Yao and S. Zheng, Dynamic control of quality in production-inventory systems, Springer, New York, NY, 2002.