Pricing in a Vertical Market with Upstream Demand Uncertainty*

Stefanos Leonardos† and Costis Melolidakis‡[1]

[1]*Department of Mathematics, National and Kapodistrian University of Athens, Panepistimioupolis GR-157 84, Athens, Greece*

June 10, 2018

Abstract

We study a classic Cournot market, which we extend to a two-stage game with endogenous cost formation: the retailers' marginal cost represents purchases from a price-setting, revenue-maximizing supplier. Any demand uncertainty falls to the supplier, who acts first and sets the wholesale price under incomplete information concerning the retailers' willingness-to-pay for ordering the product. We introduce the *generalized mean residual life* (GMRL) function of the supplier's belief distribution F and show that his revenue function is unimodal, if the GMRL function is decreasing – *(DGMRL)* property – and F has finite second moment. In this case, we show that the supplier's optimal price corresponds to the unique fixed point of the MRL function of his belief about the demand. Based on this characterization, we derive a tight bound on the probability of a stockout due to demand uncertainty and study supply chain efficiency in terms of the realized Price of Uncertainty. Furthermore, we employ the rich theory of stochastic orderings and perform a comprehensive comparative statics analysis that challenges previously established economic insights about the effects of market size and demand variability on wholesale prices. We illustrate our results with numerical examples.

Keywords: Cournot-Nash Equilibrium, Decreasing Generalized Mean Residual Life, Incomplete Information, Comparative Statics, Stochastic Orderings
Mathematics Subject Classification (2000): 91A10, 91A40

1 Introduction

Pricing decisions are often made under incomplete information, as sellers may not know what exact valuations buyers assign to their products. In these cases, revenue functions of decision makers who act under uncertainty are stochastic and their maximization is not straightforward. Such pricing problems often arise in supply chain management, when a wholesale supplier sells a product (or its main resource) to retailers without knowing the exact demand that the retailers are facing or, equivalently, their exact willingness-to-pay for their orders. A common problem in the relevant literature is to determine conditions on the distribution of the parameter of uncertainty that ensure a unimodal and hence "well behaved" revenue function. If the uncertainty does not directly affect the supplier, a widely applicable unimodality condition is the *increasing generalized failure rate* (IGFR) property, introduced by Lariviere [1999] and Lariviere and Porteus [2001].

In the present work, we develop a unimodality condition for stochastic revenue functions, by suitably generalizing the *mean residual life* (MRL) function of the parameter of uncertainty. The condition, termed the *decreasing generalized mean residual life* (DGMRL) property, arises naturally in the general context of a Cournot competition with endogenous cost formation under stochastic demand and provides a generalization of the IGFR property.

Specifically, to study the pricing decision problem of a supplier who is seeking the optimal wholesale price to maximize his revenue, we employ the classic Cournot model of a single homogeneous product and extend it to a two-stage game. In the first stage the supplier sets a wholesale price r under incomplete information about the market demand and, hence, about the retailers' exact valuation of the product. In the second stage, with no uncertainty about the market demand, the retailers engage in a classic Cournot competition. Their constant marginal cost is equal to the wholesale price r set by the supplier and their decision variable is the quantity q that they will release to the market. The market is cleared at a price p determined by the affine inverse demand function $p = \alpha - \sum q$.

*This research has been supported by the Research Funding Program ARISTEIA II: "Optimization of stochastic systems under partial information and applications" of NSRF.

†sleonardos@math.uoa.gr (corresponding author), cmelol@math.uoa.gr.

The implications of our study are twofold. From a modelling perspective, our approach aims to endogenize the formation of the cost parameter in the classic Cournot market in a way that is both realistic and mathematically tractable. As Marx and Schaffer [2016] point out, despite numerous applications of Cournot's model, there has been very little discussion about the origin of the competing firms' costs. The approach using exogenously given, constant marginal costs does not address the strategic considerations that emerge in prevailing economic practice, in which costs commonly represent purchases from a third party. Due to this negligence, Marx and Schaffer [2016] question the robustness of the results obtained thus far in Cournot oligopoly theory.

Viewed as a two-echelon supply chain, the present setting differs from the commonly studied newsvendor problem, seeMandal et al. [2018], Xu et al. [2011, 2010], Yao et al. [2006] and Lariviere and Porteus [2001], in that the second stage entails a game instead of a single decision maker (newsvendor). However, it remains amenable to similar analysis under the subgame-perfect equilibrium concept: for each wholesale price r, the second-stage competing retailers have unique equilibrium strategies (order quantities), which may be directly derived from the standard Cournot theory. Hence, under subgame perfection, the supplier knows precisely the total quantity that he will receive as an order for any possible price r and realized demand α, and the exact number of second stage retailers proves inconsequential. Cachon and Netessine [2004] survey the applications of game theoretical concepts in supply chain analysis and state that the mostly used concept is that of static non-cooperative games. They highlight the need for a game theoretical analysis of more dynamic settings and at the time of their survey, they report of only two papers that apply the solution concept of subgame perfect equilibrium, both in a setting quite different from ours.

Still, our main contribution is from a technical/methodological perspective, in which we identify and study a mild unimodality condition that ensures a well-behaved stochastic revenue function. If demand uncertainty affects only the retailers and the supplier's revenue function is deterministic, this problem has been conclusively studied in Lariviere and Porteus [2001]. They model the source of uncertainty in the demand as a nonnegative random variable X with distribution function F, survival function $\bar{F} := 1 - F$ and density $f = F'$, for all $r > 0$. X has an *increasing failure rate* (IFR) or simply, X is IFR if the failure rate function $\mathrm{h}(r) := \frac{f(r)}{\bar{F}(r)}$ is increasing for r such that $\bar{F}(r) > 0$. Lariviere [1999] and Lariviere and Porteus [2001] define the *generalized failure rate* (GFR) function of X as

$$g(r) \overset{\text{def}}{=} r\mathrm{h}(r)$$

for r in the support of X. They show that, if the GFR function is increasing and the first moment of F is finite, then the supplier's revenue function is unimodal. IFR distributions are trivially IGFR, and many DFR distributions are also IGFR, see Lariviere [2006], Banciu and Mirchandani [2013]. Hence, the IGFR condition strictly generalizes the previously used IFR unimodality condition.

Hall and Wellner [1981] and references discussed therein, studied the GFR function and its properties under different terminology and in the distinct context of reliability and statistical analysis. The first to realize its importance in economic applications were Singh and Maddala [1976], who used the IGFR property in the modelling of income distributions. Due to its intuitive interpretation, they used the term *proportional failure rate* function. Its properties were studied by Belzunce et al. [1995] and Belzunce et al. [1998]. Independently of these papers, the use of IGFR distributions in pricing/revenue management applications was introduced by Lariviere [1999] and Lariviere and Porteus [2001]. Following their influential work, the IGFR class of probability distributions has proven particularly useful in the analysis of pricing/revenue management applications, and its properties have been thoroughly studied, see Ziya et al. [2004], Paul [2005], Lariviere [2006] and Banciu and Mirchandani [2013].

The present setting differs from the supply chain studied by Lariviere and Porteus [2001], mainly in that the demand uncertainty affects the supplier instead of the retailers. The optimal wholesale price remains the supplier's lone decision variable. However, his revenue function is now stochastic. To ensure unimodality in this setting, we utilize the *mean residual life (MRL)* function $\mathrm{m}(r)$ of F, defined as $\mathrm{m}(r) := \mathbb{E}(X - r \mid X > r)$, for r such that $F(r) < 1$. X has decreasing mean residual life (DMRL) or simply, X is DMRL, if $\mathrm{m}(r)$ is decreasing. A sufficient condition for the existence and uniqueness of the supplier's optimal price, is stated in terms of the *generalized mean residual life* (GMRL) function $\ell(r)$, which we define as

$$\ell(r) \overset{\text{def}}{=} \frac{\mathrm{m}(r)}{r}$$

for $0 < r$ and $F(r) < 1$. Its interpretation is straightforward: while the MRL function $\mathrm{m}(r)$, yields the expected additional demand given the current demand, the GMRL function, yields the expected additional demand as a percentage of the current demand. Based on this interpretation, the function $\ell(r)$ could be aptly termed the *proportional mean residual life* function, as in Belzunce et al. [1998]. However, we prefer the term GMRL to highlight its close connection to the GFR function. In the current setting, we show that if the GMRL function is decreasing and the second moment of F is finite, then the supplier's revenue function is unimodal. Again, DMRL distributions are trivially DGMRL, but it is also true that many IMRL distributions are DGMRL. Hence, the DGMRL generalizes the DMRL class.

Hall and Wellner [1981] study the GMRL function and its connection with the GFR function. Among other properties, they show that both the FR and the MRL functions, and, hence, the GFR and GMRL functions, uniquely determine the underlying distribution. Guess and Proschan [1988] survey earlier advances concerning the MRL function, its properties, and applications. The IFR and DMRL properties are closely related to the more general concept of log-concave probability. In an inspiring survey, Bagnoli and Bergstrom [2005] examine a series of theorems relating to the log-concavity and/or log-convexity of probability density, distribution, and reliability functions, and their integrals.

In the context of pricing/revenue management, the MRL and GMRL functions arise naturally in problems that entail the maximization of stochastic functions. In an active stream of literature, Mandal et al. [2018], Luo et al. [2016], Song et al. [2009, 2008] and Petruzzi and Dada [1999], as well as references cited therein, study the tail of the distribution of the source of uncertainty, for which they use the term $\Theta(r) := \int_r^{+\infty} u f(u)\,\mathrm{d}u$, and provide simplified expressions, see e.g. Song et al. [2009], Lemma 1 and Song et al. [2008], equation (2). Using the definition of the MRL function, one may observe that $\Theta(r) = \mathrm{m}(r)\bar{F}(r)$. Therefore, the study of $\Theta(r)$ may be equivalently done in terms of the MRL function $\mathrm{m}(r)$, thus exploiting the rich literature on its properties. Similarly, unimodality conditions that could have been formulated in terms of the MRL function are developed in Bernstein and Federgruen [2005] and Lu and Simchi-Levi [2013]. Lagerlöf [2006] derives conditions in terms of the failure rate function that ensure equilibrium uniqueness in a Cournot market with demand uncertainty. In Section 6, we show that the DMRL and DGMRL conditions apply in his setting and generalize the framework under which his results apply.

1.1 Summary of results

The introduction of the MRL and GMRL function in stochastic revenue management problems, as in the aforementioned papers, is more than a formality and has non-trivial implications. As we demonstrate in the present paper, the advantage of this approach is that alternative representations and properties of the MRL function may be utilized to gain economic intuition. In Theorem 3.1, we characterize the supplier's optimal price as a fixed point of the MRL function of his belief distribution F. If F has the DGMRL property and finite second moment, then this fixed point is unique. Under the more restrictive DMRL sufficiency condition, we provide a bound on the market inefficiency, as measured by the probability of no transaction between the supplier and retailers when the supplier is uncertain about their willingness-to-pay. This bound holds for any DMRL distribution and is shown to be tight over the DMRL class, see Theorem 4.1. Our results are illustrated through several examples. As in the study of IGFR random variables, the Pareto distribution plays an important role as the unique distribution with a constant GMRL function.

To study market performance, we evaluate the supplier's share of the realized market profits. If the realized demand is too low, then a transaction between the supplier and the retailers does not occur and the market experiences an immediate stockout. For all distributions with the DMRL property – which form a subset of DGMRL distributions – we show that this probability is upper-bounded by $1 - e^{-1}$. Over the family of DMRL distributions, this bound is robust, i.e. independent of the particular distribution F, and tight, as it is attained by the exponential and, asymptotically, by a parametric Beta distribution, see Examples 4.2 and 4.3. For higher values of realized demand (for which a transaction between the supplier and the retailers occurs), the supplier captures a smaller share of the realized profit as the realized demand increases and a larger share of the system profits as the number of competing retailers increases.

1.2 Outline

The rest of the paper is structured as follows. In Section 2 we build the formal setting, then in Section 3, we study its equilibrium behavior and state the main result, Theorem 3.1. In Section 6 we return to the current model and demonstrate its fit in real economic applications and connections with other models in the literature. Our conclusions are summarized in Section 7.

2 The Model

We consider a classic Cournot duopoly with two identical (symmetric) retailers[1] who sell a single homogeneous good to consumers. The retailers face affine inverse demand and marginal cost r. Instead of the standard assump-

[1] Our analysis extends readily to any $n \geq 1$. To avoid unnecessary notation, we restrict our exposition to $n = 2$ and restate our results for arbitrary n in Subsection 3.3.

tion of exogenously given marginal cost, we consider r as the wholesale price of the product (or the main resource required to produce it), which is set by a single, revenue-maximizing supplier (manufacturer). The supplier has enough capacity to cover any possible demand from the retailers, and produces at a constant marginal cost, which we normalize to zero[2]. Consequently, his lone decision variable is his wholesale price r – equivalently, his profit margin – that he will use to sell his product to the retailers which is determined prior to and independently of the retailers' order-decision.

2.1 Two-stage game

We model this market as an extensive, two-stage game. The supplier acts first (Stackelberg leader) and chooses a wholesale price $r \geq 0$. The retailers observe r (a price-only contract[3]) as well as the demand realization α and choose simultaneously and independently their order-quantities $q_i(r \mid \alpha), i = 1, 2$. Formally, a strategy q_i for a retailer $i = 1, 2$, is a function $q_i : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, which specifies the quantity that the retailer will order for any price r set by the supplier. The retail price p is determined by an affine inverse demand function $p = \alpha - q(r \mid \alpha)$, in which α is the *demand parameter* and $q(r \mid \alpha) := q_1(r \mid \alpha) + q_2(r \mid \alpha)$ is the total quantity that the retailers release to the market. It is immediate, that $q_i(r \mid \alpha)$, or slightly abusing notation, $q_i, i = 1, 2$, may not exceed α. Hence, $q_i \in [0, \alpha]$ for $i = 1, 2$. The players receive their payoffs via the strategy profile $(r, \mathbf{q}(r))$, where $\mathbf{q}(r) = (q_1(r), q_2(r))$. Given the cost r, the profit function $u_i(\mathbf{q}(r) \mid r)$ or simply $u_i(\mathbf{q} \mid r)$, of retailer $i = 1, 2$, is

$$u_i(\mathbf{q} \mid r) = q_i(\alpha - q) - rq_i = q_i(\alpha - r - q) \ . \tag{1}$$

For a given value of α, the supplier's revenue function, is $u_s(r \mid \alpha) = rq(r \mid \alpha)$, for $0 \leq r < \alpha$, where $q(r \mid \alpha)$ depends on α via (1).

2.1.1 Demand uncertainty

The retailers face no uncertainty about the demand – parameter α in the demand function – and the quantity that they order from the supplier is equal to the quantity that they sell to the market (at equilibrium). Contrarily, the supplier may have incomplete information about the exact value of the demand parameter α when he sets the wholesale price. To model this situation, we assume that, after the pricing decision of the supplier, but prior to the order-decisions of the retailers, a value for α is realized from a continuous distribution F, with finite mean $\mathbb{E}\alpha < +\infty$ and nonnegative values, i.e. $F(0) = 0$. Equivalently, F represents the supplier's belief about the demand parameter and, hence, about the retailers' willingness-to-pay his price. For the tail of F, we write $\bar{F} := 1 - F$ and for the support of F, let $\alpha_L = \sup\{r \geq 0 : F(r) = 0\} \geq 0$ and $\alpha_H = \inf\{r \geq 0 : F(r) = 1\} \leq +\infty$. The case $\alpha_L = \alpha_H$ is not excluded[4] and corresponds to the situation where the supplier is also completely informed about the true demand. Under these assumptions, the supplier's payoff function becomes stochastic and hence, assuming expected utility maximization, he will optimize its expectation with respect to the distribution F of the unknown demand parameter α, i.e., $u_s(r) = \mathbb{E}u_s(r \mid \alpha)$.

All of the above are assumed to be common knowledge among the participants in the market (the supplier and the retailers).

3 Market equilibrium: existence and uniqueness

We restrict attention to subgame perfect equlibria[5] of the extensive form, two-stage game, in which for any value of the cost parameter, $r \in (0, \alpha_H)$, set by the supplier, the retailers use their Nash equilibrium strategies, $\mathbf{q}^*(r) = \left(q_1^*(r), q_2^*(r)\right)$, in the game induced by r. Accordingly, let r^* denote the optimal strategy of the supplier, and $q^*(r) := q_1^*(r) + q_2^*(r)$ the total quantity in the order that the supplier receives when the retailers respond optimally to price r.

3.1 Second stage: Cournot competition

Given the wholesale price r set by the supplier and the realized value of the demand parameter α, the retailers play a standard Cournot duopoly with affine inverse demand function and constant marginal cost r. Their equilibrium

[2]For a discussion of the model assumptions see Section 6.

[3]There is no return option and the salvage value of the product is zero.

[4]Formally, this case contradicts the assumption that F is continuous or non-atomic. However, it is allowed so that unnecessary notation is avoided. It should cause no confusion.

[5]Technically, these are perfect Bayes-Nash equilibria, since the supplier has a belief about the retailers' types, i.e. their willingness-to-pay his price, that depends on the value of the stochastic demand parameter α.

strategies are well known to be unique, symmetric, and equal to $q_i^*(r) = \frac{1}{3}(\alpha - r)^+$ for $i = 1, 2$. Under subgame perfection, uniqueness of the equlibrium strategies implies that the supplier correctly predicts the retailers' response, i.e. orders of quantities, for any wholesale price $r \in (0, \alpha_H)$ that he may set.

3.2 First stage: supplier's optimal pricing decision

Given the second-stage equilibrium strategies, we turn to the supplier's optimization problem in the first stage. First, we treat the deterministic case, i.e. the case with no supplier uncertainty about the demand parameter α, which, according to the notation introduced in Subsection 2.1.1, corresponds to the case $\alpha_L = \alpha_H$. Since $q^*(r \mid \alpha) = \frac{2}{3}(\alpha - r)^+$, the supplier's payoff on the equilibrium path is

$$u_s(r \mid \alpha) = rq^*(r \mid \alpha) = \frac{2}{3}r(\alpha - r)^+, \quad \text{for } 0 \leq r. \tag{2}$$

Maximization of equation (2) with respect to r yields that the complete information two-stage game has a unique subgame perfect Nash equilibrium, under which the supplier sells with optimal price $r^*(\alpha) = \frac{1}{2}\alpha$ and each of the retailers orders quantity $q^*(r) = \frac{1}{3}(\alpha - r)^+$.

The equilibrium behavior of the market in the stochastic case $\alpha_L < \alpha_H$, i.e., under the assumption that the supplier has incomplete information about the true value of the demand parameter, α, when he sets his price, is less straightforward. Assuming that the retailers respond to the supplier's choice r with their unique equilibrium strategies, the supplier's payoff at equilibrium is given by taking the expectation in (2) with respect to the distribution F of α

$$u_s(r) = \frac{2}{3}r\,\mathbb{E}(\alpha - r)^+, \quad \text{for } 0 \leq r \leq \alpha_H.$$

Hence, the supplier is interested in finding conditions for the existence and/or uniqueness of a finite optimal price r^* that maximizes $u_s(r)$. For an arbitrary distribution however, $u_s(r)$ may not be concave and, hence, not unimodal, in which case the solution to the supplier's optimization problem is not immediate. To obtain a general unimodality condition, we proceed by differentiating the supplier's revenue function $u_s(r)$. By the monotone convergence theorem, $\mathbb{E}(\alpha - r)^+$ is continuously differentiable for $r \in (0, \alpha_H)$ with $\frac{d}{dr}\mathbb{E}(\alpha - r)^+ = -\bar{F}(r)$. Since $(\alpha - r)^+$ is nonnegative, we may write (e.g. see Billingsley [1986]),

$$\mathbb{E}(\alpha - r)^+ = \int_0^\infty P\left((\alpha - r)^+ > u\right) \mathrm{d}u = \int_r^\infty \bar{F}(u)\,\mathrm{d}u, \quad \text{for } 0 \leq r < \alpha_H$$

Using this formulation, both the revenue function of the supplier and its first derivative can be expressed in terms of the *mean residual life* (MRL) function of α. The MRL function $\mathrm{m}(\cdot)$ of a nonnegative random variable α with cumulative distribution function (cdf) F and finite expectation, $\mathbb{E}\alpha < +\infty$, is defined as

$$\mathrm{m}(r) := \mathbb{E}(\alpha - r \mid \alpha > r) = \frac{1}{\bar{F}(r)}\int_r^\infty \bar{F}(u)\,\mathrm{d}u, \quad \text{for } r < \alpha_H$$

and $\mathrm{m}(r) := 0$, otherwise, see, e.g., Shaked and Shanthikumar [2007], Lai and Xie [2006] or Belzunce et al. [2016]. Using this notation, we obtain that

$$u_s(r) = \frac{2}{3}r\mathrm{m}(r)\bar{F}(r) \tag{3a}$$

$$\frac{\mathrm{d}u_s}{\mathrm{d}r}(r) = \frac{2}{3}(\mathrm{m}(r) - r)\bar{F}(r) = \frac{2}{3}r\left(\frac{\mathrm{m}(r)}{r} - 1\right)\bar{F}(r) \tag{3b}$$

for $0 < r < \alpha_H$. Based on (3b), the first order condition (FOC) for the supplier's revenue function is that $\mathrm{m}(r) = r$ or equivalently that $\mathrm{m}(r)/r = 1$. We call the expression

$$\ell(r) := \frac{\mathrm{m}(r)}{r}, \quad 0 < r < \alpha_H \tag{4}$$

the *generalized mean residual life* (GMRL) function, due to its connection to the generalized failure rate $\mathrm{g}(r) := \frac{rf(r)}{\bar{F}(r)}$, defined and studied by Lariviere [1999] and Lariviere and Porteus [2001]. Its meaning is straightforward: while the MRL function $\mathrm{m}(r)$ at point $r > 0$ measures the expected additional demand, given the current demand r, the GMRL function measures the expected additional demand as a percentage of the given current demand. From an economic perspective, $\ell(r)$ is related to the *price elasticity* $e_{<r>}$ *of expected demand*, since

$$\ell(r) = \frac{\mathrm{m}(r)}{r} = \left(\frac{-\bar{F}(r)}{\mathrm{m}(r)\bar{F}(r)} \cdot r\right)^{-1} = \left(-r \cdot \frac{\frac{d}{dp}\mathbb{E}q^*(r \mid \alpha)}{\mathbb{E}q^*(r \mid \alpha)}\right)^{-1} = e_{<r>}^{-1} \tag{5}$$

Hence, the FOC shows that the supplier's payoff is maximized at the point of unit elasticity. Realistic problems must have a price elasticity that eventually becomes greater than 1, see Lariviere [2006]. Hence, for an economically meaningful analysis, we focus on distributions for which $\mathrm{m}(r)/r$ eventually becomes less than 1, i.e., distributions for which $\bar{r} := \sup\{r \geq 0 : \mathrm{m}(r) \geq r\}$ is finite. Moreover, for a nonnegative random demand α with continuous distribution F and finite expectation $\mathbb{E}\alpha$, $\mathrm{m}(0) = \mathbb{E}\alpha > 0$ and hence $\bar{r} > 0$.

Based on these considerations, it remains to derive conditions that guarantee the existence and uniqueness of an r^* that satisfies the FOC and to show that this r^* indeed corresponds to a maximum of the supplier's revenue function. This is established in the following Theorem which is the main result of the present Section.

Theorem 3.1. *Assume that the supplier's belief about the unknown, nonnegative demand parameter, α, is represented by a continuous (non-atomic) distribution F, with support inbetween α_L and α_H with $0 \leq \alpha_L < \alpha_H \leq \infty$.*
(a) Necessary condition: If an optimal price r^* for the supplier exists, then r^* satisfies the fixed point equation

$$r^* = \mathrm{m}(r^*) \tag{6}$$

(b) Sufficient conditions: If the generalized mean residual life (GMRL) function $\ell(r), r > 0$, of F is strictly decreasing and $\mathbb{E}\alpha^2$ is finite, then under equilibrium, the optimal price r^* of the supplier exists and is the unique solution of (6)*. In this case, if $\frac{1}{2}\mathbb{E}\alpha < \alpha_L$, then $r^* = \frac{1}{2}\mathbb{E}\alpha$. Otherwise, $r^* \in [\alpha_L, \alpha_H)$.*

The last claim of part (b) establishes the link between the equilibrium wholesale price in the stochastic and the wholesale price in the deterministic market. To see this, consider a given demand level α and a distribution F with support $[\alpha_L, \alpha_H]$ and $\alpha \in [\alpha_L, \alpha_H]$. If the supplier has less uncertainty about α, then we mau

Proof. Since $\bar{F}(r) > 0$ for $0 < \alpha < \alpha_H$, the sign of the derivative $\frac{\mathrm{d}u_s}{\mathrm{d}r}(r)$ is determined by the term $\mathrm{m}(r) - r$ and any critical point r^* satisfies $\mathrm{m}(r^*) = r^*$. Hence, the necessary part of the theorem is obvious from (3b) and the continuity of $\frac{\mathrm{d}u_s}{\mathrm{d}r}(r)$. For the sufficiency part, it remains to check that such a critical point exists and corresponds to a maximum under the assumptions that $\ell(r)$ is strictly decreasing and $\mathbb{E}\alpha^2 < +\infty$. Clearly, $\mathrm{m}(r) - r$ is continuous and $\lim_{r\to 0^+} \mathrm{m}(r) - r = \mathbb{E}\alpha > 0$. Hence, $u_s(r)$ starts increasing on $(0, \alpha_H)$. However, the limiting behavior of $\mathrm{m}(r) - r$ and hence of $\frac{\mathrm{d}u_s}{\mathrm{d}r}(r)$ as r approaches α_H from the left, may vary depending on whether α_H is finite or not. If α_H is finite, i.e. if the support of α is bounded, then $\lim_{r\to\alpha_H^-}(\mathrm{m}(r) - r) = -\alpha_H$. In particular, for bounded support, it is trivial that $\ell(r)$ eventually becomes less than 1, since it goes to 0 as r approaches α_H. Hence, in this case, a critical point r^* that corresponds to a maximum exists without any further assumptions. Uniqueness of r^* is also established by the strict monotonicity of $\ell(r)$. If $\alpha_H = +\infty$, then an optimal solution r^* may not exist, see Pareto distribution in Example 3.3, as the limiting behavior of $\mathrm{m}(r)$ as $r \to +\infty$ may vary, Bradley and Gupta [2003]. In this case, the condition of finite second moment ensures that $\bar{r} < +\infty$. In particular, as we show in Leonardos and Melolidakis [2018], if the GMRL function $\ell(r)$ of a random variable α with unbounded support is decreasing, then $\lim_{r\to+\infty} \ell(r) < 1$ if and only if $\mathbb{E}\alpha^2$ is finite. Strict monotonicity of $\ell(r)$ is required to eliminate intervals of the form $\mathrm{m}(r) = r$ that give rise to multiple optimal solutions. These conditions combined are clearly sufficient to guarantee the existence and uniqueness of a critical point r^*, i.e. of an optimal solution for the supplier's decision problem.

To prove the second claim of the sufficiency part, note that $\mathbb{E}\alpha < 2\alpha_L$ is equivalent to $\mathrm{m}(\alpha_L) < \alpha_L$. Then, the DGMRL property implies that $\mathrm{m}(r) < r$ for all $r > \alpha_L$, hence $r^* < \alpha_L$. In this case, $\mathrm{m}(r^*) = \mathbb{E}\alpha - r^*$ and hence r^* will be given explicitly by $r^* = \frac{1}{2}\mathbb{E}\alpha$, which may be compared with the optimal r^* of the complete information case. On the other hand, if $\mathbb{E}\alpha \geq 2\alpha_L$, then for all $r < \alpha_L$, $\mathrm{m}(r) = \mathbb{E}\alpha - r \geq 2\alpha_L - r > r$ which implies that r^* must be in $[\alpha_L, \alpha_H)$. □

Remark 3.2. Strict monotonocity may be relaxed to weak monotonicity in the statement of Theorem 3.1 without significant loss of generality. This relies on the fact that distributions that contain MRL functions with linear segments have been charachterized in Proposition 10 of Hall and Wellner [1981]. Namely, $\mathrm{m}(r) = r$ on some interval $J \subseteq [\alpha_L, \alpha_H]$ if and only if the underlying random variable α has the Pareto distribution with shape parameter 2 on J. If J is unbounded, then $\mathbb{E}\alpha^2 = +\infty$, see Example 3.3, and hence, this case may not occur under the requirement that $\mathbb{E}\alpha^2 < +\infty$. However, to eliminate bounded intervals J, on which the supplier's payoff function admits consecutive optimal values, strict monotonicity of the GMRL function is necessary. In view of the above, this is equivalent to excluding the Pareto distribution with parameter 2, on every subinterval $J \subseteq [\alpha_L, \alpha_H]$.

Example 3.3 (Pareto distribution)**.** The Pareto distribution is the unique distribution with constant GMRL and GFR functions. Let α be Pareto distributed with pdf $f(r) = k\alpha_L^k r^{-(k+1)}\mathbf{1}_{\{\alpha_L \leq r\}}$, and parameters $0 < \alpha_L$ and $k > 1$ (for $0 < k \leq 1$ we get $\mathbb{E}\alpha = +\infty$, which contradicts the basic assumptions of our model). To simplify, let $\alpha_L = 1$, so that $f(r) = kr^{-k-1}\mathbf{1}_{\{1\leq r<\infty\}}$, $F(r) = \left(1 - r^{-k}\right)1_{\{1\leq r<\infty\}}$, and $\mathbb{E}\alpha = \frac{k}{k-1}$. The mean residual life of α is given by $\mathrm{m}(r) = \frac{r}{k-1} + \frac{k}{k-1}(1-r)\mathbf{1}_{\{0\leq r<1\}}$ and, hence, is decreasing on $[0, 1)$ and increasing on $[1, \infty)$. However, the GMRL function $\ell(r) = \frac{1}{k-1}$ is constant over the support of α and, hence, α is DGMRL. Similarly, for $1 \leq r$ the failure (hazard) rate $\mathrm{h}(r) = kr^{-1}$ is decreasing, but the generalized failure rate $\mathrm{g}(r) = k$ is constant and, hence, α is IGFR. The payoff function of the supplier is

$$u_s(r) = \frac{2}{3} r\, \mathrm{m}(r)\, \bar{F}(r) = \frac{2}{3(k-1)} \begin{cases} r(r - rk + k) & \text{if } 0 \leq r < 1 \\ r^{2-k} & \text{if } r \geq 1, \end{cases}$$

which diverges as $r \to +\infty$, for $k < 2$ and remains constant for $k = 2$. In particular, for $k \leq 2$, the second moment of α is infinite, i.e. $\mathbb{E}\alpha^2 = +\infty$, and also $\lim_{r\to+\infty} \ell(r) = \frac{1}{k-1} \geq 1$ and $\lim_{r\to+\infty} \mathrm{g}(r) = k \leq 2$. This shows that for DGMRL distributions, we may not drop the assumption that the second moment of F is finite, cf. Theorem 3.1. Contrary, for $k > 2$, we get a unique solution as expected, namely $r^* = \frac{k}{2(k-1)}$, which is indeed the unique fixed point of $\mathrm{m}(r)$.

3.3 General case with n identical retailers

To ease the exposition, we restricted our presentation to $n = 2$ identical retailers, but the present analysis applies to arbitrary $n \geq 1$. Formally, let $N = \{1, 2, \ldots, n\}$, with $n \geq 1$ denote the set of identical retailers. As in Section 2, a strategy profile is denoted with $\mathbf{q} = (q_1, q_2, \ldots, q_n)$. The payoff function of retailer i, for $i \in N$, depends on the total quantity of the remaining $n - 1$ retailers and is given by (1), where now q denotes the total quantity sold by all n retailers, i.e. $q = \sum_{j=1}^{n} q_j$. Following common notation, let $\mathbf{q} = (\mathbf{q}_{-i}, q_i)$ and $q_{-i} = q - q_i$, for $i \in N$. In the second stage, the n-identical retailers play a Cournot oligopoly with linear inverse demand function and cost r, and hence their equilibrium strategies $q_i^*(r)$ are given by $q_i^*(r) = \frac{1}{n+1}(\alpha - r)^+$, for $r \geq 0$. In the first stage, the payoff function of the supplier under complete information is given by $u_s(r \mid \alpha) = rq^*(r) = \frac{n}{n+1} r(\alpha - r)^+$ and hence it is maximized again at $r^*(\alpha) = \frac{1}{2}\alpha$. If the supplier knows only the distribution F and not the true value of α, his payoff function becomes

$$u_s(r) = \frac{n}{n+1} r\mathbb{E}(\alpha - r)^+ = \frac{n}{n+1} r\mathrm{m}(r)\, \bar{F}(r), \quad \text{for } 0 \leq r < \alpha_H.$$

Hence, the number of the second-stage retailers affects the supplier's revenue function only up to a scaling constant and Theorem 3.1 is stated unaltered for any $n \geq 1$. Intuitively, for any $n \geq 1$, there exists a unique second-stage equilibrium, and hence, under the subgame perfect equilibrium concept and given α, the supplier knows precisely what order he will receive for any price $r > 0$ he may set. Thus, the approach to the supplier's expected revenue maximization in the first-stage remains the same independently of the number of second-stage retailers.

4 Market Efficiency

Markets with incomplete information are usually inefficient in the sense that trades that are profitable for all market participants may actually not take place. In the current model, such inefficiencies appear as values of α for which a transaction does not occur in equilibrium under incomplete information, although such a transaction would have been beneficial for all parties involved, i.e., supplier, retailers and consumers.

If $\alpha < r^*$, then the retailers buy 0 units and there is an immediate stockout. Hence, for a particular continuous distribution F of α, the probabilitiy that a transaction does not occur in equilibrium under incomplete information is equal to $P(\alpha \leq r^*) = F(r^*)$. To study this probability as a measure of market inefficiency, we restrict attention to the family of DMRL distributions, i.e., distributions for which $\mathrm{m}(r)$ is nonincreasing. In this case, we have

Theorem 4.1. *For any distribution F of α with the DMRL property, the probability $F(r^*)$ that a transaction does not occur in the equilibrium of the incomplete information case cannot exceed the bound $1 - e^{-1}$. This bound is tight over all DMRL distributions.*

Proof. By expressing the distribution function F in terms of the MRL function, e.g. see Guess and Proschan [1988], we get

$$F(r^*) = 1 - \frac{\mathrm{m}(0)}{\mathrm{m}(r^*)} \exp\left\{-\int_0^{r^*} \frac{1}{\mathrm{m}(u)}\, \mathrm{d}u\right\}.$$

Hence, by the DMRL property and the monotonicity of the exponential function, it follows that $F(r^*) \leq 1 - \frac{\mathrm{m}(0)}{\mathrm{m}(r^*)} \exp\left\{-\frac{1}{\mathrm{m}(r^*)} \cdot r^*\right\}$. Since $r^* = \mathrm{m}(r^*) \leq \mathrm{m}(0)$, we conclude that $F(r^*) \leq 1 - e^{-1}$. If the MRL function is constant, as is the case for the exponential distribution, see Example 4.2, then all inequalities above hold as equalities, which establishes the second claim of the theorem. □

Example 4.2 (Exponential distribution)**.** Let $\alpha \sim \exp(\lambda)$, with $\lambda > 0$, and pdf $f(r) = \lambda e^{-\lambda r}\mathbf{1}_{\{0\le r<\infty\}}$. Since $\mathrm{m}(r) = \frac{1}{\lambda}$, for $r > 0$, the MRL function is constant over its support and, hence, F is both DMRL and IMRL but strictly DGMRL, as $\ell(r) = \frac{1}{\lambda r}$, for $r > 0$. By Theorem 3.1, the optimal strategy r^* of the supplier is $r^* = \frac{1}{\lambda}$. The probability of no transaction $F(r^*)$ is equal to $F(r^*) = F(1/\lambda) = 1 - e^{-1}$, confirming that the bound derived in Theorem 4.1 is tight. Thus, the exponential distribution is the least favorable, over the class of DMRL distributions, in terms of efficiency at equilibrium.

Example 4.3 (Beta distribution)**.** This example refers to a special case of the Beta distribution, also known as the Kumaraswamy distribution, see Jones [2009]. Let $\alpha \sim Beta(1,\lambda)$ with $\lambda > 1$, and pdf $f(r) = \lambda(1-r)^{\lambda-1}\mathbf{1}_{\{0<r<1\}}$. Then, $\bar{F}(r) = (1-r)^{\lambda}$ and $\mathrm{m}(r) = \frac{1-r}{1+\lambda}$ for $0 < r < 1$. Since the MRL function is decreasing, Theorem 3.1 applies and the optimal price of the supplier is $r^* = \frac{1}{\lambda+2}$. Hence, $F(r^*) = 1 - \left(1 - \frac{1}{\lambda+2}\right)^{\lambda} \to 1 - e^{-1}$ as $\lambda \to +\infty$. This shows that the upper bound of $F(r^*)$ in Theorem 4.1 is still tight over distributions with strictly decreasing MRL, i.e., it is not the flatness of the exponential MRL that generated the large inefficiency.

Example 4.4 (Lomax, Generalized Pareto or Pareto II distribution)**.** This example shows that the bound of Theorem 4.1 does not extend to the class of DGMRL distributions. Let $\alpha \sim \mathrm{Lomax}(A, B = 1, k = 2+\epsilon)$, with $A < 1, \epsilon > 0$ and cdf $F(r) = 1 - [1 - A + r]^{-k}$. By a standard calculation, $\mathrm{m}(r) = \frac{1-A+r}{1+\epsilon}$, hence F is DGMRL but not DMRL. In this case, $r^* = \frac{1-A}{\epsilon}$ and $\bar{F}(r^*) = \left(\frac{\epsilon}{(1+\epsilon)(1-A)}\right)^{(2+\epsilon)}$, which shows that the probability of a stockout may become arbitrarily large as ϵ approaches 0 from above. The "pathology" of this example relies on the fact that $\mathbb{E}\alpha^2 \to \infty$ as $\epsilon \searrow 0$.

In the case that a transaction takes place, i.e., for values of $\alpha > r^*$, we measure market efficiency in terms of the *realized* market profits. To this end, we fix a demand distribution F with support $S \subseteq [\alpha_L, \alpha_H]$ and a realized demand level $\alpha \in S$ and compare the individual realized profits of the supplier and each retailer between the two scenarios: the scenario in which the supplier prices before demand realization and the scenario in which the supplier prices after demand realization. For clarity, the necessary quantities are summarized in Table 1.

	Upstream Demand for Supplier	
	Uncertain $\alpha \sim F$	Deterministic α
Equilibrium Wholesale Price	$r^* = \mathrm{m}_F(r^*)$	$r^* = \alpha/2$
	Realized Profits in Equilibrium	
Supplier	$\Pi_s^U = \frac{n}{n+1} r^*(\alpha - r^*)^+$	$\Pi_s^D = \frac{n}{n+1}(\alpha/2)^2$
Retailer i	$\Pi_i^U = \frac{1}{(n+1)^2}\left((\alpha - r^*)^+\right)^2$	$\Pi_i^D = \frac{n}{(n+1)^2}(\alpha/2)^2$
Aggregate	$\Pi_{\mathrm{Agg}}^U = \Pi_s^U + \sum_{i=1}^n \Pi_i^U$	$\Pi_{\mathrm{Agg}}^D = \Pi_s^U + \sum_{i=1}^n \Pi_i^U$

Table 1. Wholesale price in equilibrium and realized profits when the supplier prices under stochastic demand (left column) and under deterministic demand (right column).

First, we focus on the stochastic market and evaluate the division of the *realized* system profit between supplier and retailers which is given by the ratio $\Pi_s^U / \left(\Pi_s^U + \sum_{i=1}^n \Pi_i^U\right)$. Recall from Theorem 3.1 that As mentioned above, if $\alpha \le r^*$, then there is no transaction and the profits of all participants are equal to zero. For $\alpha > r^*$, we have that

$$\frac{\Pi_s^U}{\Pi_s^U + \sum_{i=1}^n \Pi_i^U} = \frac{\frac{n}{n+1} r^*(\alpha - r^*)}{\frac{n}{n+1} r^*(\alpha - r^*) + n\left(\frac{1}{n+1}(\alpha - r^*)\right)^2} = \frac{nr^* + r^*}{nr^* + \alpha}$$

Hence, the division of realized profit between supplier and retailers depends on the realized demand α and the number n of retailers. The supplier captures a smaller share of the realized profit as the realized demand α increases and a larger share of the realized market profits as the number n of competing retailers increases (for fixed α).

If there is no demand uncertainty, then, given the demand parameter α, the supplier sets a wholesale price equal to $r^* = \alpha/2$. I this case,

$$\frac{\Pi_s^D}{\Pi_s^D + \sum_{i=1}^n \Pi_i^D} = \frac{\frac{n}{n+1}\left(\frac{\alpha}{2}\right)^2}{\frac{n}{n+1}\left(\frac{\alpha}{2}\right)^2 + \frac{n}{(n+1)^2}\left(\frac{\alpha}{2}\right)^2} = \frac{n+1}{n+2}$$

which is independent of the realized demand α and increasing in the number n of competing second-stage retailers. Comparing the two ratios – with and without demand uncertainty – we see that the supplier captures a larger share of the market profits if he prices under demand uncertainty for realized values of α in $[r^*, 2r^*]$.

4.1 Price of Uncertainty

We next turn to the comparison of the individual and aggregate market profits in the two scenarios: with and without demand uncertainty for the supplier. For each retailer in the stochastic market, we have $\frac{1}{(n+1)^2}\left((\alpha - r^*)^+\right)^2 \geq \frac{1}{(n+1)^2}\left(\frac{\alpha}{2}\right)^2$ for all values of $\alpha \geq 2r^*$. This implies that for larger values of the realized demand, the retailers are better off if the supplier prices under demand uncertainty. In contrast, the supplier is never better off when he prices under demand uncertainty, as is intuitively expected. Indeed $\frac{n}{n+1} r^* (\alpha - r^*)^+ \leq \frac{n}{n+1} (\alpha/2)^2$ for all values of α, with equality if and only if $\alpha = 2r^*$. The ratio of the supplier's realized profit in the scenario with demand uncertainty to the scenario without demand uncertainty is equal to $4 \cdot \frac{r^*}{\alpha}\left(1 - \frac{r^*}{\alpha}\right)$ and hence it has the shape shown in Figure 1, independently of the underlying demand distribution.

Similar findings are obtained when we compare the market's *aggregate* realized profits (supplier and retailers) between these two scenarios. Specifically, for a fixed demand distribution F with support $S \subseteq [\alpha_L, \alpha_H]$, we evaluate the ratio of the aggregate realized market profits under deterministic demand Π^D_{Agg} to the aggregate realized market profits under stochastic demand Π^U_{Agg}. Again, to retain equilibrium uniqueness, we restrict attention to the class of continuous, nonnegative DGMRL distributions. For a realized demand $\alpha \leq r^*$, there is a stockout and the realized aggregated profits Π^U_{Agg} are equal to 0. In this case, the stochastic market performs arbitrarily worse than the deterministic market. Hence, to obtain a non-trivial analysis, we study the opposite question, i.e., of *how good* can the market perform under stochastic demand compared to deterministic demand. Motivated by a similar notion that is studied in Blum and Mansour [2013], we fix a demand distribution F with support $S \subseteq [\alpha_L, \alpha_H]$ and study the *realized Price of Uncertainty* (PoU_r)

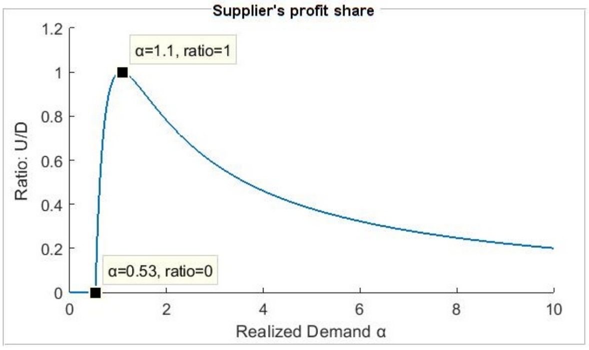

Figure 1. Ratio Π^U_s/Π^D_s of the supplier's realized profits with and without demand uncertainty for $\alpha \sim$ Weibull$(1, 2)$ and arbitrary n.

$$\text{PoU}_\text{r} := \inf_{\alpha \in S}\left\{\frac{\Pi^D_{\text{Agg}}}{\Pi^U_{\text{Agg}}}\right\} = \inf_{\alpha \in S}\left\{\frac{\Pi^D_s + \sum_{i=1}^n \Pi^D_i}{\Pi^U_s + \sum_{i=1}^n \Pi^U_i}\right\} \tag{7}$$

Whenever $\Pi^U_{\text{Agg}} = 0$ we set the inner ratio equal to $+\infty$. Lower values of PoU_r correspond to a lower *price of uncertainty*, i.e., to better performance of the market under stochastic demand relative to the underlying deterministic case. This implies that as defined here, the PoU_r is a measurement of the inefficiencies in the *best case scenario*. Intuitively, one expects the market to perform worse under demand uncertainty which translates to PoU being bounded below by 1. Whereas for realized demand values $\alpha \leq r^*$ the stochastic indeed performs arbitrarily worse than the deterministic market, for values $\alpha > r^*$, the aggregate market profits of the supplier and the retailers may be larger if the supplier prices under demand uncertainty, and hence, the PoU_r may be less than 1, as stated in Theorem 4.5.

Theorem 4.5. *The realized Price of Uncertainty, PoU_r, of the stochastic market is* $1 - O\left(n^{-2}\right)$*. For any demand distribution F, the* inf *is attained for realized demand* $\alpha^* := \min\left\{\alpha_H, \frac{2n}{n-1} r^*\right\}$.

Proof. By Table 1, a direct substitution yields that the inner ratio is equal to

$$\frac{\Pi^D_s + \sum_{i=1}^n \Pi^D_i}{\Pi^U_s + \sum_{i=1}^n \Pi^U_i} = \frac{\frac{n}{n+1}\left(\frac{\alpha}{2}\right)^2 + \frac{n}{(n+1)^2}\left(\frac{\alpha}{2}\right)^2}{\frac{n}{n+1} r^* (\alpha - r^*)^+ + n\left(\frac{1}{n+1}(\alpha - r^*)^+\right)^2}$$

Hence, $\text{PoU}_\text{r} = \inf_{\alpha > r^*}\left\{\frac{(n+2)\alpha^2}{4(\alpha - r^*)(\alpha + nr^*)}\right\}$. Taking the partial derivative of the inner ratio with respect to $\alpha > r^*$ yields

$$\frac{\partial}{\partial \alpha}\left(\frac{(n+2)\,\alpha^2}{4(\alpha - r^*)(\alpha + nr^*)}\right) = \frac{(n+1)\,\alpha r^*}{4(\alpha - r^*)^2(\alpha + nr^*)^2}\left(\alpha(n-1) - 2nr^*\right)$$

which shows that the ratio is decreasing on $[r^*, \frac{n}{n-1}2r^*)$, and increasing thereafter. If $\alpha_H \geq \frac{n}{n-1}2r^*$, then the ratio is minimized for $\alpha^* := \frac{2n}{n-1}r^*$, yielding a value of $1 + \frac{1}{n^2+2n} = 1 + O\left(n^{-2}\right)$. This lower bound does depend on the underlying distribution F and is less than 1 for any number n of competing second-stage retailers. If $\alpha_H < \frac{2n}{n-1}r^*$, the lowest ratio is obtained for α_H and hence, the previously obtained bound PoU_r still applies (however, in this case, it is not attained). □

The values for which the ratio is less than 1, depend on n. For $n \geq 3$, we have that $\frac{(n+2)\alpha^2}{4(\alpha-r^*)^+(\alpha+nr^*)} \leq 1$ for values of α in $S \cap [2r^*, \frac{2n}{n-2}r^*]$. In this case, the upper bound decreases to $2r^*$ as $n \to \infty$. For $n = 2$, the upper bound is equal to infinity, i.e., the range of α for which the ratio is below 1 is equal to $S \cap [2r^*, +\infty)$. In all cases, the lower bound is independent of n. Finally, by taking the partial derivative with respect to n, we find that the ratio $\Pi^D_{\text{Agg}}/\Pi^U_{\text{Agg}}$ is nonincreasing in n for realized values of α in $[r^*, 2r^*]$ and increasing in n thereafter, again independently of the underlying demand distribution. This is illustrated in Figure 2.

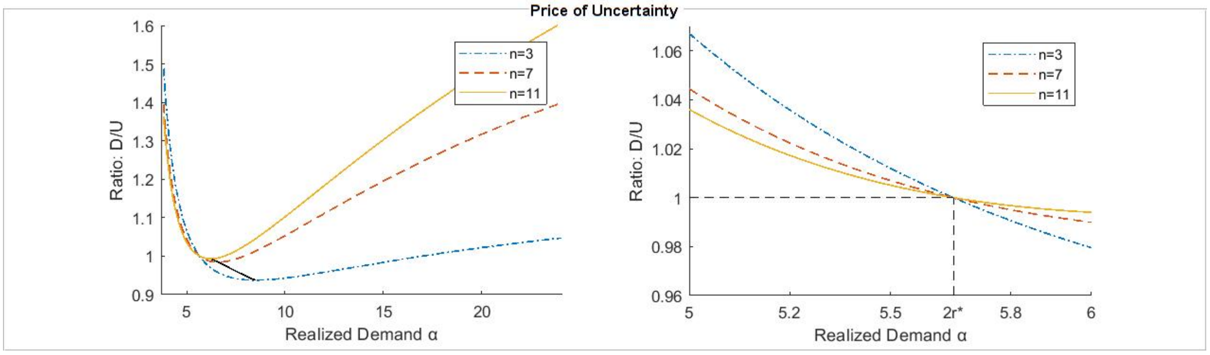

Figure 2. Ratio $\Pi^D_{\text{Agg}}/\Pi^U_{\text{Agg}}$ of the aggregate realized market profits with and without demand uncertainty for $n = 3$ (solid line) and $n = 10$ (dashed line) with $\alpha \sim \text{Gamma}(2, 2)$. The right figure amplifies the interval $[4, 7]$ in which the two curves intersect. Prior to the intersection point, $2r^* \approx 5.657$, the ratio is higher for $n = 10$, whereas after the intersection the ratio is higher for $n = 3$.

5 Comparative Statics

The closed form expression of (6) is the basis for the comparative statics and sensitivity analysis on the demand distribution's paremeters via the rich theory of *stochastic orders*, see Shaked and Shanthikumar [2007], Lai and Xie [2006] and Belzunce et al. [2016]. In equilibrium, the total quantity $q^* = \frac{n}{n+1}(\alpha - r^*)^+$ that will be sold to the market and the retail price $p^* = \alpha - q^*$ are both monotone in the wholesale price r^*. Accordingly, we restrict attention to the wholesale price r^* for varying distribution parameters.

For a meaningful comparison between different markets, we assume throughout equilibrium uniqueness. Hence, unless stated otherwise, we consider only continuous, strictly DGMRL distributions with finite second moment[6].To proceed, we introduce some notation. Let $X \sim F, Y \sim G$ be two nonnegative random variables with MRL functions $\text{m}_F(r)$ and $\text{m}_G(r)$, respectively, such that $\text{m}_F(r) \leq \text{m}_G(r)$ for every $r \geq 0$. Then, X is less than Y in the *mean residual life* order, denoted by $X \leq_{\text{mrl}} Y$. Using this notation, the following Lemma captures the importance of the characterization in (6).

Lemma 5.1. *Let $X \sim F, Y \sim G$ be two nonnegative and continuous, strictly DGMRL demand distributions with finite second moment. If $X \leq_{mrl} Y$, then $r^*_F \leq r^*_G$.*

Proof. By definition of $X \leq_{\text{mrl}} Y$, it holds that $\text{m}_F(r) \leq \text{m}_G(r)$ for all $r > 0$ which in turn is equivalent to $\ell_F(r) \leq \ell_G(r)$ for all $r > 0$. Hence, by (6), $1 = \ell_F\left(r^*_F\right) \leq \ell_G\left(r^*_F\right) < \ell_G(r)$ for all $r < r^*_F$, where the second inequality follows from the assumption that Y is strictly DGMRL. Since r^*_2 is the unique solution of $\ell\left(r^*_G\right) = 1$, it follows that $r^*_F \leq r^*_G$. □

[6]Since the DGMRL class is particularly inclusive, see Banciu and Mirchandani [2013] vale edw ton allo apo ton Koba/Popescu. and Leonardos and Melolidakis [2018] and finiteness of the second moment of the demand captures the assumption of increasing price elasticity, these conditions do not pose significant restrictions. Still, since they are sufficient but not necessary, the analysis applies to any other setting that guarantees equilibrium existence and uniqueness.

Although trivial to prove once Theorem 3.1 has been established, Lemma 5.1 is the key for the subsequent comparative statics analysis. Utilizing (5), we have that $X \leq_{\text{mrl}} Y$ if and only if for any $r \geq 0$, the price elasticity of expected demand $e^F_{<r>}$ in market X is less than the price elasticity of expected demand $e^G_{<r>}$ in market Y, i.e., if $e^F_{<r>} \leq e^G_{<r>}$ for every $r \geq 0$. This motivates the following definition: we will say that market X is *less elastic* than market Y, denoted by $X \leq_{\text{el}} Y$, if $e^F_{<r>} \leq e^G_{<r>}$ for every $r \geq 0$. Combininig the above, we have established that $X \leq_{\text{mrl}} Y$ if and only if $Y \leq_{\text{el}} X$ and that the supplier will charge a higher price in a less elastic market. Hence, the task of studying the response of the optimal wholesale price r^* to varying demand distribution parameters – such as market size or demand variability – largely reduces to ordering markets accordings to their elasticities or equivalently to their mrl functions. Such conditions can be found in Shaked and Shanthikumar [2007] and provide the framework for the subsequent analysis.

5.1 Market Size

We start with the response of the equilibrium wholesale price r^* to transformations that intuitively correspond to a larger market.

5.1.1 Reestimating Demand

Let X denote the demand distribution in an instance of the market under consideration. Let $c \geq 1$ denote a positive constant and Z an additional random source of demand that is independent of X. Moreover, let r^*_X denote the equilibrium wholesale price in the initial market and r^*_{X+Z} the equilibrium wholesale price in the market with random demand $X + Z$. How does r^*_X compare to r^*_{cX} and r^*_{X+Z}?

While the intuition that the larger markets cX and $X+Z$ give rise to higher wholesale prices is largely confirmed, see Theorem 5.2, the results do not hold in full generality and hinge on additional conditions on Z. For instance, since DGMRL random variables are not closed under convolution, see Leonardos and Melolidakis [2018], the random variable $X+Z$ may not be DGMRL. This may lead to multiple equilibriam in the $X+Z$ market, irrespectively of whether Z is DGMRL or not. To deal with the possilble multiplicity of equilibrium wholesale prices after some transformation of the original demand distribution X, we will write $\mathbf{r}^*_{\mathbf{W}} = \{r : r = \mathrm{m}_W(r)\}$ for the set of all such possible wholesale equilibrium prices, where m_W denotes the mrl function of a $W \sim F_W$ demand distribution. In Theorem 5.2, we will use $W := X + Z$. To ease the notation, we will also write $\mathbf{r}^*_{\mathbf{W}} \leq \mathbf{r}^*_{\mathbf{V}}$, when all elements of the set $\mathbf{r}^*_{\mathbf{W}}$ are less or equal than all elements of the set $\mathbf{r}^*_{\mathbf{V}}$.

Theorem 5.2. *Let $X \sim F$ be a nonnegative and continuous demand distribution with finite second moment.*

(i) If X is DGMRL and $c \geq 1$ is a positive constant, then $r^*_X \leq r^*_{cX}$.
(ii) If X is DMRL and Z is a nonnegative, continuous random variable with finite second moment and independent of X, then $r^*_X \leq r^*_{X+Z}$ for any equilibrium wholesale price r^*_{X+Z} of the $X + Z$ market.

Proof. The proof of part (i) follows directly from the preservation property of the $\leq_{\text{mrl}}$-order that is stated in Theorem 2.A.11 of Shaked and Shanthikumar [2007]. Specifically, since $\mathrm{m}_{cX}(r) = c\mathrm{m}_X(r/c)$ is the mrl function of cX, we have that $\mathrm{m}_{cX}(r) = r \cdot \frac{\mathrm{m}_X(r/c)}{r/c} = r \cdot \ell(r/c) \geq r \cdot \ell(r) = \mathrm{m}_X(r)$, for all > 0, with the inequality following from the assumption that X is DGMRL. Hence, $X \leq_{\text{mrl}} cX$ which by Lemma 5.1 implies that $r^*_X \leq r^*_{cX}$.

Part (ii) follows from Theorem 2.A.11 of Shaked and Shanthikumar [2007]. The proof necessitates that X is DMRL and hence requiring that X is merely DGMRL is not enough. Since, X is DMRL, we know that $r < \mathrm{m}(r)$ for all $r < r^*_X = \mathrm{m}_X(r^*_X)$. Together with $X \leq_{\text{mrl}} X + Z$, this implies that $r < \mathrm{m}_X(r) \leq \mathrm{m}_{X+Z}(r)$, for all $r < r^*_X$. Hence, $\mathbf{r}^*_{\mathbf{X+Z}} \subseteq [r^*_X, +\infty)$, which implies that in this case, r^*_X is a lower bound to the set of all possible wholesale equilibrium prices in the $X + Z$ market. □

5.1.2 Preservation of Market Size & Wholesale Price

Next, we turn our attention to operations that preserve the $\leq_{\text{mrl}}$-order. Let $X_1 \sim F_1, X_2 \sim F_2$ denote two different demand distributions, such that $X_1 \leq_{\text{mrl}} X_2$. In this case, we know by Lemma 5.1 that $r^*_1 \leq r^*_2$. We are interested in determining transformations of X_1, X_2 that preserve the $\leq_{\text{mrl}}$-order and hence the ordering $r^*_1 \leq r^*_2$. Again, to avoid technicalities, we assume that X_1, X_2 are such that Theorem 3.1 applies.

Theorem 5.3. *Let $X_1 \sim F_1, X_2 \sim F_2$ denote two nonnegative, continuous and strictly DGMRL demand distributions, with finite second moments, such that $X_1 \leq_{mrl} X_2$.*

(i) If Z is a nonnegative, IFR distribution, independent of X_1 and X_2, then $\mathbf{r}^*_{\mathbf{X_1+Z}} \leq \mathbf{r}^*_{\mathbf{X_2+Z}}$.
(ii) If ϕ is an increasing, convex function, then $r^*_{\phi(X_1)} \leq r^*_{\phi(X_2)}$.

(iii) If $X_p \sim F_1 + (1-p) F_2$ for some $p \in (0,1)$, then $r^*_{X_1} \leq \mathbf{r}^*_{\mathbf{X_p}} \leq r^*_{X_2}$.

Proof. For the proof, we utilitze Lemma 2.A.8 and Theorems 2.A.18 and 2.A.19 from Shaked and Shanthikumar [2007] and the closure properties of DGMRL random variables from Leonardos and Melolidakis [2018]. Specifically, part (i) follows from Lemma 2.A.8. Since the resulting distributions $X_i + Z, i = 1, 2$ may not be DGMRL nor DMRL, the setwise notation is necessary. Part (ii) follows from Theorem 2.A.19. Since the DGMRL class of distributions is closed under increasing, convex transformations, both $\phi(X_i), i = 1, 2$ remain DGMRL and hence equilibrium uniqueness is retained. Finally, part (iii) follows from Theorem 2.A.19. However, the DGMRL class is not closed under mixtures and hence, in this case, the X_p market may have multiple equilibria, which necessitates the setwise statement for the wholesale equilibrium prices of the X_p market. □

If instead of $X_1 \leq_{\text{mrl}} X_2$, X_1 and X_2 are ordered in the stronger $\leq_{\text{hr}}$-order, i.e., if $X_1 \leq_{\text{hr}} X_2$, then part (ii) of Theorem 5.3 remains true by Lemma 2.A.10 of Shaked and Shanthikumar [2007], even if Z is merely DMRL (instead of IFR). Although Theorems 5.2 and 5.3 are immediate *once* Theorem 3.1 and Lemma 5.1 have been established, they provide non-trivial economic intuitions. In general, both Theorems imply that if the supplier reestimates upwards his expectations about the stochastic demandl then he will charge a higher wholesale price. Although intuitive, this conclusion depends on the conditions that imply the $\leq_{\text{mrl}}$-order, e.g. that Z is IFR in part (i) of Theorem 5.3, and does not hold in general, as shown in Subsection 5.1.3. Following the exposition of Shaked and Shanthikumar [2007], the above collection of statements can be extended to incorporate more case-specific results.

5.1.3 Stochastically Larger Market

The usual $\leq_{\text{st}}$-order does not imply nor is implied by the $\leq_{\text{mrl}}$-order, see Shaked and Shanthikumar [2007]. This implies that there exist instances of stochastically larger markets, in which the supplier may still charge a lower price. This in line with the intuition of Lariviere and Porteus [2001] that "size is not everything" and that prices are driven by different forces. To see this, consider the following example from Shaked and Shanthikumar [1991]. Let $X \sim F$ be uniformly distributed on $[0, 1]$ and let $Y \sim G$ have a piecewise linear distribution with $G(0) = 0, G(1/3) = 7/9, G(2/3) = 7/9$ and $G(1) = 1$. Then, as the shown in Figure 3, $Y \leq_{\text{st}} X$ (right panel) but $r^*_F \leq r^*_G$ (left panel). This example considers distributions on a bounded interval. However, the same situation may arise

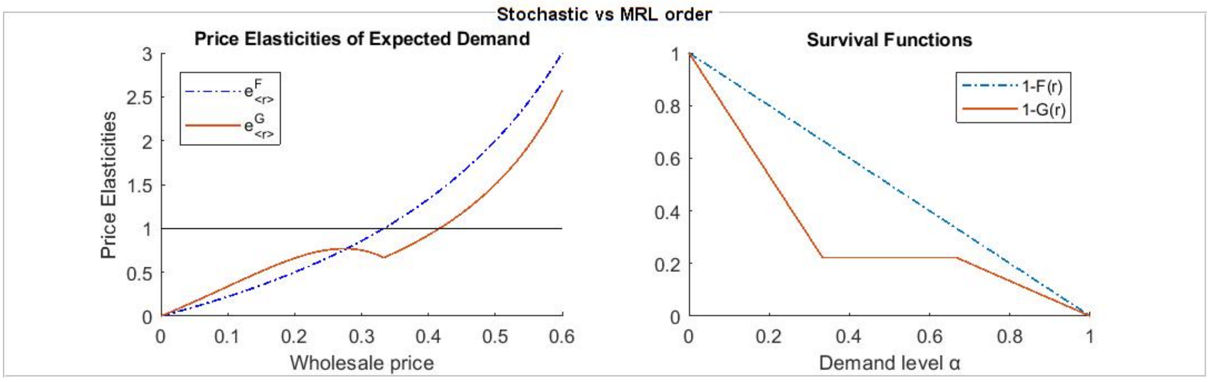

Figure 3. F stochastically dominates G, however $r^*_G > r^*_F$.

eveon over undbounded intervals. Such an example, albeit technical, is provided in Figure 4. Let

$$f(r; \omega, \kappa, \phi) := \frac{\kappa\left(\kappa^2 + \omega^2\right)}{\kappa^2 \cos(\omega\phi) + \kappa^2 + \kappa\omega \sin(\omega\phi) + \omega^2} \cdot e^{-\kappa r} \left(\cos(\omega(r - \phi)) + 1\right)$$

for $r \geq 0$, denote the densities of a parametric family of exponentially decaying sinusoids. For $(\omega, \kappa, \phi) = (0, \kappa, 0)$, f corresponds to the exponential distribution with parameter κ. Figure 4 depicts the optimal wholesaleprices r^*_F and r^*_G (left panel) and the logarithm of the ratio of the survival functions, $\log\left(\bar{F}/\bar{G}\right)$ (right panel), for $X \sim F$ corresponding to $(\omega, \kappa, \phi) = (\pi, 0.8, 1.2)$ and $Y \sim G$ to $(\omega, \kappa, \phi) = (0, 0.9, 0)$. Since the log-survival ratio remains throughout positive, it follows that $X \leq_{\text{st}} Y$. However, as shown graphically in the left panel,$r^*_F = 1.03 < r^*_G = 1.11$. This relies on the fact that X and Y are not ordered in the $\leq_{\text{mrl}}$-order. In both examples, the anomalies arise from distributions that are not *DGMRL*. We have not a conclusive answer on whether stochastically larger DGMRL (or DMRL) demand distributions also imply higher wholesale prices.

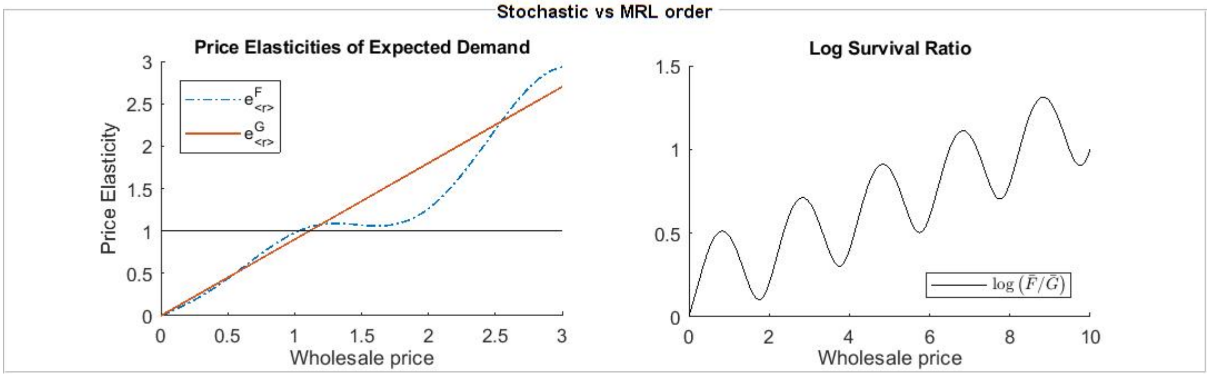

Figure 4. F stochastically dominates G, however $r^*_G > r^*_F$.

5.2 Market Demand Variability

The response of the equilibrium wholesale price to increasing (decreasing) demand variability is even less straightforward. There exist several notions of stochastic orders that compare random variables in terms of their variability and depending on which we employ, we derive conflicting results. First, we introduce some notation.

5.2.1 Variability or Dispersive Orders

Let $X_1 \sim F_1$ and $X_2 \sim F_2$ be two nonnegative distributions with equal means, $\mathbb{E}X_1 = \mathbb{E}X_2$, and finite second moments. If $\int_r^{+\infty} \bar{F}_1(u)\,\mathrm{d}u \le \int_r^{+\infty} \bar{F}_2(u)\,\mathrm{d}u$ for all $r \ge 0$, then X_1 is said to be smaller than X_2 in the *convex order*, denoted by $X_1 \le_{\text{cx}} X_2$. If F_1^{-1} and F_2^{-1} denote the right continuous inverses of F_1, F_2 and $F_1^{-1}(r) - F_1^{-1}(s) \le F_2^{-1}(r) - F_2^{-1}(s)$ for all $0 < r \le s < 1$, then X_1 is said to be smaller than X_2 in the *dispersive order*, denoted by $X_1 \le_{\text{disp}} X_2$. Finally, if $\int_{F_1^{-1}(p)}^{\infty} \bar{F}_1(u)\,\mathrm{d}u \le \int_{F_2^{-1}(p)}^{\infty} \bar{F}_2(u)\,\mathrm{d}u$ for all $p \in (0,1)$, then X_1 is said to be smaller than X_2 in the *excess wealth order*, denoted by $X_1 \le_{\text{ew}} Y$. Shaked and Shanthikumar [2007] show that $X \le_{\text{disp}} Y \implies X \le_{\text{ew}} Y \implies X \le_{\text{cx}} Y$ which in turn implies that $\text{Var}(X) \le \text{Var}(Y)$. Further insights and motivation about these orders are provided in Chapter 3 of Shaked and Shanthikumar [2007].

Does less variability imply a lower (higher) wholesale price? The answer to this question, largely depends on the notion of variability that will be employed. First, we consider the more general $\le_{\text{cx}}$-variability order. In general, the $\le_{\text{mrl}}$-order is not implied by the $\le_{\text{cx}}$-order. This implies that it does not suffice to order two demand distributions $X \sim F$ and $Y \sim G$ in the $\le_{\text{cx}}$-order, to conclude that wholesale prices in the X, Y markets will be ordered respectively. This is illustrated in Figures 5 and 6. In Figure 5, we consider two demand distributions,

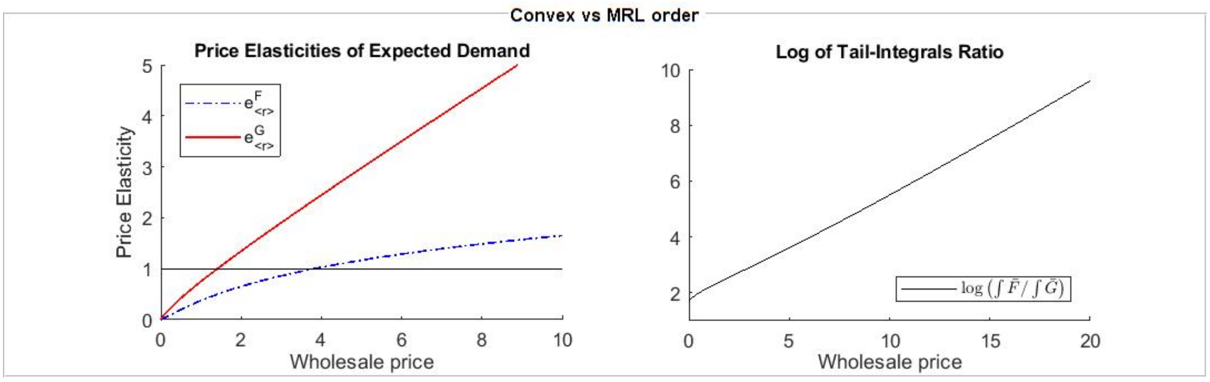

Figure 5. Comparison of $F \sim \text{Lognormal}(\mu = 0.5, \sigma = 1)$ and $G \sim \text{Gamma}(\alpha = 2, \beta = 0.25)$.

$X \sim F$ with F Lognormal$(\mu = 0.5, \sigma = 1)$ and $Y \sim G$ with G Gamma$(\alpha = 2, \beta = 0.25)$. For this choice of parameters, $\mathbb{E}X = \mathbb{E}Y = 0.5$ and hence X, Y are ordered in the $\le_{\text{cx}}$-order if and only if the tail-integrals of F and G can be ordered, see Shaked and Shanthikumar [2007] Theorem 3.A.1. The right panel depicts the log of the ratio of these integrals, i.e., $\log\left(\int_r^\infty \bar{F}\,\mathrm{d}u \,/\, \int_r^\infty \bar{G}\,\mathrm{d}u\right)$ which remains throughout positive (and increasing). Hence, $Y \le_{\text{cx}} X$. The left panel depicts the MRL functions of X and Y and their fixed points, i.e. the optimal wholesales prices. As can be seen, in the X market (distribution F) the supplier charges a higher pricer than in the less variable (according to the $\preceq_{cx}$-order) Y market (distribution G).

The above conlcusion is reversed in the case of Figure 6. In this example, we consider two demand distributions, $X \sim F$ with F, as above, Lognormal $(\mu = 0.5, \sigma = 1)$ and $Y \sim G$ with G Gamma $(\alpha = 8, \beta = 0.25/4)$. The choice of parameters ensures that the equality $\mathbb{E}X = \mathbb{E}Y = 0.5$ is retained and hence that X, Y can be ordered in the $\preceq_{cx}$-order if and only if the tail-integrals of F and G can be ordered. Again, the right panel depicts the log of the ratio of these integrals which remains throughout positive (and increasing). Hence, $Y \preceq_{cx} X$. However, the picture in the left panel is now reversed. As can be seen, in the X market (distribution F) the supplier now charges a lower price than in the less variable (according to the $\preceq_{cx}$-order) Y market. If we restrict attention to the $\preceq_{ew}$- and

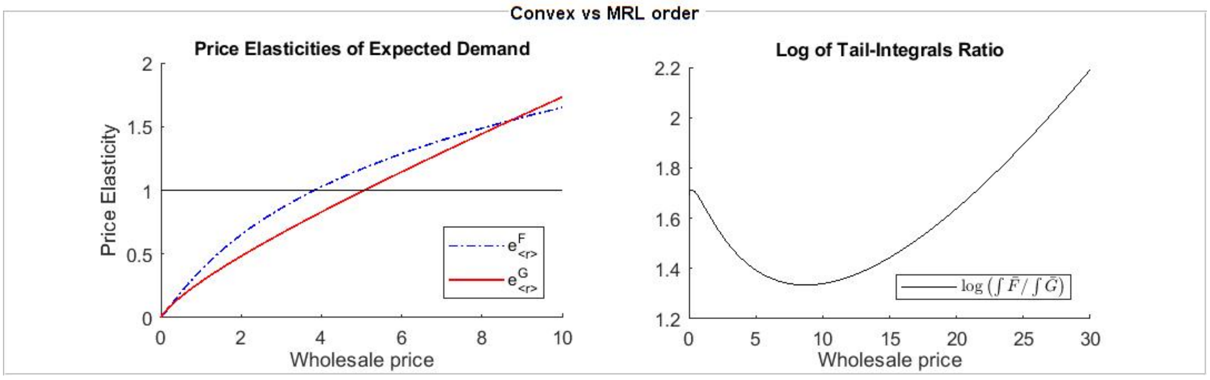

Figure 6. Comparison of $F \sim$ Lognormal $(\mu = 0.5, \sigma = 1)$ and $G \sim$ Gamma $(\alpha = 8, \beta = 0.25/4)$.

$\preceq_{disp}$-orders, then more can be said. Recall that α_L denotes the left end of the support of a variable X. Accordingly, we will write α_{Li} to denote the left end of the support of variable X_i for $i = 1, 2$. Again, to avoid technicalities, we restrict attention to random variables that satisfy the assumptions of Theorem 3.1.

Theorem 5.4. *Let $X_1 \sim F_1, X_2 \sim F_2$ be two nonnegative, DGMRL demand distributions with finite second moment and $\alpha_{L1} \le \alpha_{L2}$. If either X_1, X_2 or both are DMRL and $X_1 \preceq_{ew} X_2$, then $r_1^* \le r_2^*$.*

Theorem 5.4 follows directly from Theorem 3.C.5 of Shaked and Shanthikumar [2007]. Based on its proof, the assumption that at least one of the two random variables is DMRL (and not merely DGMRL) cannot be relaxed. Belzunce et al. [2016] provide equivalent characterizations of the $\preceq_{ew}$-order to address the difficulty in the evaluation of incomplete integrals of quantile functions that is involved in its definition. A result of similar flavor can be obtained if we use the $\preceq_{disp}$ order. Again, the condition that both X_1 and X_2 are DGMRL does not suffice and we need to assume that at least one is IFR. Recall, that IFR $\subset$ DMRL $\subset$ DGMRL with all inclusions being strict, see e.g. Leonardos and Melolidakis [2018].

Theorem 5.5. *Let $X_1 \sim F_1, X_2 \sim F_2$ be two nonnegative, DGMRL demand distributions with finite second moment. If either X_1, X_2 or both are IFR and $X_1 \preceq_{disp} X_2$, then $r_1^* \le r_2^*$.*

Theorem 5.5 follows directly from Theorem 3.B.20 (b) of Shaked and Shanthikumar [2007] and the fact that the $\preceq_{hr}$-order implies the $\preceq_{mrl}$-order. Again, more case specific results can be drawn from the analysis of Shaked and Shanthikumar [2007]. The first lesson from the preceding discussion is that for a range of distributions that can be ordered in terms of their variability, less (more) variability implies lower (higher) wholesale prices. This is in sharp contrast with the results of Lariviere and Porteus [2001] and sheds light on the effects of demand uncertainty. If uncertainty affects the retailer, then the supplier charges a higher price and captures an increasing share of all supply chain profits as variability reduces. Contrarily, if uncertainty falls to the supplier, then the supplier charges a lower price as variability increases.

The second is that these results, albeit general, do not apply to all distributions that are comparable according to *some* variability order. As illustrated in Figures 5 and 6, there exist notions of variability and demand distributions that can be ordered according to these notions, for which less variability implies higher wholesale prices. This demonstrates the usefulness of the characterization in (6). In contrast to most existing studies, it shows that conclusions regarding the effect of demand variability on prices, may well differ depending on the exact notion of variability that will be employed and may be contradictory even in the standard setting of linear demand that is employed in the present study.

5.2.2 Parametric families of distributions

To further elaborate on the fact that different variability notions may lead to different conclusions, we reconsider the approach of Lariviere and Porteus [2001]. Given a random variable X with distribution F, Lariviere and Porteus [2001] consider the random variables $X_i := \delta_i + \lambda_i X$ with $\delta_i \geq 0$ and $\lambda_i > 0$ for $i = 1, 2$. They show that in this case, the wholesale price is dictated by the coefficient of variation, $CV_i = \sqrt{\mathrm{Var}(X_i)}/\mathbb{E}X_i$. Specifically, if $CV_2 < CV_1$, then $r_1^* < r_2^*$, i.e., in their model, a lower CV, or equivalently a lower relative variability, implies a higher price.

To establish a similar comparison in the present setting, we consider two normal demand distributions $X_1 \sim N\left(\mu_1, \sigma_1^2\right)$ and $X_2 \sim N\left(\mu_2, \sigma_2^2\right)$. By Table 2.2 of Belzunce et al. [2016], if $\sigma_1 < \sigma_2$ and $\mu_1 \leq \mu_2$, then $X_1 \leq_{\mathrm{mrl}} X_2$ and hence, by Lemma 5.1, $r_1^* \leq r_2^*$. However, by choosing σ_i and μ_i appropriately, we can trivially achieve an arbitrary ordering of their relative variability in terms of their CV's. The reason is that the conclusions from this approach are obscured by the fact that changing μ_i for $i = 1, 2$, does not only affect CV_i but also the central location of the respective demand distribution. In contrast, under the assumption that $\mathbb{E}X_1 = \mathbb{E}X_2$, the stochastic-orderings approach that is employed above, isolates the effect of demand variability and provides more robust insights.

6 Model assumptions & related literature

In this section, we return to the present economic model and discuss its assumptions and their fit to real economic applications, as well as its relation to other models that have been studied in the relevant literature.

a) Supplier's production capacity: The supplier has enough capacity to cover any possible demand. Hence, despite having incomplete information about the retailers' exact orders, the supplier decides only on an optimal wholesale price and not on a combined capacity-pricing policy. Technically, this assumption is implemented through the normalization of the supplier's production cost to 0. In this way, one part of the supplier's uncertainty is absorbed by the fact that his production decision does not depend on it. This is a common assumption in models that focus on price-only decisions as in Lariviere and Porteus [2001], Lu and Simchi-Levi [2013] a.o.

b) Timing of the demand realization: The supplier is uncertain about the value of the demand parameter, which is realized only after his pricing decision. As pricing may be considered flexible, one may argue that the supplier can postpone his decision to the point that the uncertainty is resolved. Although *rigidity* or *inflexibility* of prices is hardly uncommon in contemporary economic practice, see e.g. Fabiani et al. [2007] and references contained therein, our model aims to capture the idea that a supplier may have to make decisions with incomplete information about the retailers' willingness-to-pay for ordering the product or some other aspect of the retail market (like the demand). Hence, the timing of the demand realization after the first stage is only the means to facilitate the implementation of incomplete information. As a consequence, the revenue function of the supplier in the current setting is stochastic, which is an important change from the model studied in Lariviere and Porteus [2001].

In terms of technical conditions, a closely related model is considered in Lagerlöf [2006]. He studies profit maximization of Cournot firms and establishes equilibrium uniqueness under the assumptions that the hazard rate of the demand uncertainty function is monotone (or changes sign at most once) and the standard assumption that the expected value of the demand is higher than the firms' cost. Under mild additional conditions, equilibrium existence and/or uniqueness can be achieved in his model under the less restrictive DMRL or DGMRL conditions. To see this, we rewrite the equilibrium condition in equation (1) of Lagerlöf [2006] as

$$\int_{bX^*}^{\alpha^*} u f(u)\,\mathrm{d}u - c = \frac{(n+1)\,bX^*}{n}\bar{F}(bX^*) \iff \mathrm{m}(bX^*) - c\bar{F}(bX^*)^{-1} - \frac{bX^*}{n} = 0$$

Letting $q(X) := \mathrm{m}(bX) - c\bar{F}(bX)^{-1} - \frac{bX}{n}$, we have that $q(0) = \mathbb{E}(a) - c > 0$ by assumption, and, hence, the DMRL condition is sufficient for equilibrium uniqueness. Further, if we rewrite $q(X)$ as $q(X) = bX\left(\ell(bX) - \frac{1}{n}\right) - c\bar{F}(bX)^{-1}$, we see that the more restrictive condition that $\ell(bX)$ is decreasing and becomes less than $\frac{1}{n}$, or, equivalently, that the $(n+1)$-th moment of F is finite, see Leonardos and Melolidakis [2018], is sufficient to yield existence, but, however, not necessarily uniqueness. If c is set equal to 0, then the DGMRL condition and finiteness of the $(n+1)$-th moment, yield both existence and uniqueness.

7 Conclusions

We studied a classic Cournot oligopoly with affine demand function. Contrary to the classic assumption of exogenously given, constant marginal cost, we endogenized the formation of the cost parameter by considering a two-stage model. In the first stage, a single, revenue-maximizing supplier sets the wholesale price under demand uncertainty or equivalently under partial ignorance of the retailers' exact valuations of the product. His uncertainty is modelled through a continuous distribution F for the demand parameter α with nonnegative support and finite expectation. To determine his optimal pricing policy, we identified two suffiency conditions. The first is the DMRL property of the supplier's belief. To formulate a less restrictive condition, we introduced the generalized mean residual life (GMRL) function $\ell(r) := \frac{\mathrm{m}(r)}{r}$ and showed in Theorem 3.1 that the supplier's revenue function is unimodal if the GMRL function is decreasing and the second moment of F is finite. Based on the characterization of all critical points of the supplier's revenue function as fixed points of the MRL function, we derived in Theorem 4.1 a robust and tight over the DMRL class, upper bound on market inefficiency due to demand uncertainty.

Acknowledgements

Stefanos Leonardos gratefully acknowledges support by a scholarship of the Alexander S. Onassis Public Benefit Foundation.

References

M. Bagnoli and T. Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, 2005. doi:10.1007/s00199-004-0514-4.

M. Banciu and P. Mirchandani. Technical note – new results concerning probability distributions with increasing generalized failure rates. *Operations Research*, 61(4):925–931, 2013. doi:10.1287/opre.2013.1198.

F. Belzunce, J. Candel, and J. M. Ruiz. Ordering of Truncated Distributions through Concentration Curves. *Sankhy: The Indian Journal of Statistics, Series A (1961-2002)*, 57(3):375–383, 1995. URL http://www.jstor.org/stable/25051063.

F. Belzunce, J. Candel, and J. M. Ruiz. Ordering and Asymptotic Properties of Residual Income Distributions. *Sankhyā: The Indian Journal of Statistics, Series B (1960-2002)*, 60(2):331–348, 1998. URL http://www.jstor.org/stable/25053042.

F. Belzunce, C. Martinez-Riquelme, and J. Mulero. Chapter 2: Univariate stochastic orders. In F. Belzunce, C. Martnez-Riquelme, and J. Mulero, editors, *An Introduction to Stochastic Orders*, pages 27 – 113. Academic Press, 2016. doi:10.1016/B978-0-12-803768-3.00002-8.

F. Bernstein and A. Federgruen. Decentralized Supply Chains with Competing Retailers Under Demand Uncertainty. *Management Science*, 51(1):18–29, 2005. doi:10.1287/mnsc.1040.0218.

P. Billingsley. *Probability and measure, 2nd edition*. Wiley, NewYork, 1986.

M.-F. Balcan A. Blum and Y. Mansour. The Price of Uncertainty. *ACM Trans. Econ. Comput.*, 1(3):15:1–15:29, 2013. doi:10.1145/2509413.2509415.

D. Bradley and R. Gupta. Limiting behaviour of the mean residual life. *Annals of the Institute of Statistical Mathematics*, 55(1):217–226, 2003. doi:10.1007/BF02530495.

G. Cachon and S. Netessine. Game theory in supply chain analysis. In David Simchi-Levi, David Wu, and Zuo-Jun Shen, editors, *Handbook of Quantitative Supply Chain Analysis*, volume 74 of *International Series in Operations Research & Management Science*, pages 13–65. Springer US, 2004. doi:10.1007/978-1-4020-7953-5_2.

S. Fabiani, C. Loupias, F. Martins, and R. Sabbatini. *Pricing decisions in the Euro area: How firms set prices and why*. Oxford University Press, Inc., 2007. ISBN 978-0-19-530928-7.

F. Guess and F. Proschan. Mean residual life: theory and applications. In P. R. Krishnaiah and C. R. Rao, editors, *Quality Control and Reliability*, volume 7 of *Handbook of Statistics*, pages 215–224. Elsevier Amsterdam, 1988. doi:10.1016/S0169-7161(88)07014-2.

W. Hall and J. Wellner. Mean residual life. In M. Csørgø, D. A. Dawson, J. N. K. Rao, and A. K. Md. E. Saleh, editors, *Proceedings of the International Symposium on Statistics and Related Topics*, pages 169–184. North Holland Amsterdam, 1981.

C. Jones. Kumaraswamy's distribution: A beta-type distribution with some tractability advantages. *Statistical Methodology*, 6:70–81, 2009. doi:10.1016/j.stamet.2008.04.001.

J. Lagerlöf. Equilibrium Uniqueness in a Cournot Model with Demand Uncertainty. *The B.E. Journal in Theoretical Economics*, Vol. 6: Iss 1. (Topics), Article 19:1–6, 2006. doi:10.2202/1534-598X.1333.

C.-D. Lai and M. Xie. *Stochastic Ageing and Dependence for Reliability*. Springer Science + Business Media, Inc., 2006.

M.A. Lariviere. Supply Chain Contracting and Coordination with Stochastic Demand. In S. Tayur, R. Ganeshan, and M. Magazine, editors, *Quantitative Models for Supply Chain Management*, pages 233–268. Springer US, Boston, MA, 1999. doi:10.1007/978-1-4615-4949-9_8.

M.A. Lariviere. A Note on Probability Distributions with Increasing Generalized Failure Rates. *Operations Research*, 54(3):602–604, 2006. doi:10.1287/opre.1060.0282.

M.A. Lariviere and E. Porteus. Selling to the Newsvendor: An Analysis of Price-Only Contracts. *Manufacturing & Service Operations Management*, 3(4):293–305, 2001. doi:10.1287/msom.3.4.293.9971.

S. Leonardos and C. Melolidakis. Fixed Points of the Mean Residual Life Function. *ArXiv e-prints*, 2018. URL http://adsabs.harvard.edu/abs/2018arXiv180506327L.

Y. Lu and D. Simchi-Levi. On the Unimodality of the Profit Function of the Pricing Newsvendor. *Production and Operations Management*, 22(3):615–625, 2013. doi:10.1111/j.1937-5956.2012.01419.x.

S. Luo, S.P. Sethi, and R. Shi. On the optimality conditions of a price setting newsvendor problem. *Operations Research Letters*, 44(6):697–701, 2016. doi:https://doi.org/10.1016/j.orl.2016.08.005.

P. Mandal, R. Kaul, and T. Jain. Stocking and pricing decisions under endogenous demand and reference point effects. *European Journal of Operational Research*, 264(1):181–199, 2018. doi:10.1016/j.ejor.2017.05.053.

L.M. Marx and G. Schaffer. Cournot Competition with a Common Input Supplier. Working paper, Duke University, 2016. URL https://faculty.fuqua.duke.edu/~marx/bio/papers/MarxShaffer_Cournot_2015.pdf.

A. Paul. A Note on Closure Properties of Failure Rate Distributions. *Operations Research*, 53(4):733–734, 2005. doi:10.1287/opre.1040.0206.

G. Perakis and G. Roels. The Price of Anarchy in Supply Chains: Quantifying the Efficiency of Price-Only Contracts. *Management Science*, 53(8):1249–1268, 2007. doi:10.1287/mnsc.1060.0656.

N.C. Petruzzi and M. Dada. Pricing and the Newsvendor Problem: A Review with Extensions. *Operations Research*, 47(2):183–194, 1999. doi:10.1287/opre.47.2.183.

M. Shaked and G. Shanthikumar. *Stochastic Orders*. Springer Series in Statistics, New York, 2007. ISBN 0-387-32915-3.

M. Shaked and J. G. Shanthikumar. Dynamic Multivariate Mean Residual Life Functions. *Journal of Applied Probability*, 28(3):613–629, 1991.

S. K. Singh and G. S. Maddala. A Function for Size Distribution of Incomes. *Econometrica*, 44(5):963–70, 1976.

Y. Song, S. Ray, and S. Li. Structural Properties of Buyback Contracts for Price-Setting Newsvendors. *Manufacturing & Service Operations Management*, 10(1):1–18, 2008. doi:10.1287/msom.1060.0140.

Y. Song, S. Ray, and T. Boyaci. Technical Note – Optimal Dynamic Joint Inventory-Pricing Control for Multiplicative Demand with Fixed Order Costs and Lost Sales. *Operations Research*, 57(1):245–250, 2009. doi:10.1287/opre.1080.0530.

M. Xu, Y.(F.) Chen, and X. Xu. The effect of demand uncertainty in a price-setting newsvendor model. *European Journal of Operational Research*, 207(2):946–957, 2010. doi:10.1016/j.ejor.2010.06.024.

X. Xu, X. Cai, and Y.(F.) Chen. Unimodality of price-setting newsvendor's objective function with multiplicative demand and its applications. *International Journal of Production Economics*, 133(2):653–661, 2011. doi:10.1016/j.ijpe.2011.05.013.

L. Yao, Y.(F.) Chen, and H. Yan. The newsvendor problem with pricing: Extensions. *International Journal of Management Science and Engineering Management*, 1(1):3–16, 2006. doi:10.1080/17509653.2006.10670993.

S. Ziya, H. Ayhan, and R.D. Foley. Relationships Among Three Assumptions in Revenue Management. *Operations Research*, 52(5):804–809, 2004. doi:10.1287/opre.1040.0134.