MANAGEMENT SCIENCE
Vol. 53, No. 8, August 2007, pp. 1269–1288
ISSN 0025-1909 | EISSN 1526-5501 | 07 | 5308 | 1269
DOI 10.1287/mnsc.1070.0700

Risk Mitigation in Newsvendor Networks: Resource Diversification, Flexibility, Sharing, and Hedging

Jan A. Van Mieghem
Kellogg School of Management, Northwestern University, Evanston, Illinois 60208,
vanmieghem@kellogg.northwestern.edu

This paper studies how judicious resource allocation in networks mitigates risk. Theory is presented for general utility functions and mean-variance formulations and is illustrated with networks featuring resource diversification, flexibility (e.g., inventory substitution), and sharing (commonality). In contrast to single-resource settings, risk-averse newsvendors may invest more in networks than risk-neutral newsvendors: some resources and even total spending may exceed risk-neutral levels. With normally distributed demand, risk-averse newsvendors change resource levels roughly proportionally to demand variance, while risk-neutral agents adjust only proportionally to standard deviation.

Two effects explain this operational hedge and suggest rules of thumb for strategic placement of safety capacity and inventory in networks: (1) Risk pooling suggests rebalancing capacity toward inexpensive resources that serve lower-profit variance markets. This highlights the role of profit variance (instead of demand variance) in risk-averse network investment. (2) Ex post revenue maximization suggests rebalancing capacity toward substitutable flexible but away from shared capacity when markets differ in profitability. Capacity imbalance and allocation flexibility thus mitigate profit risk and truly are operational hedges.

Key words: risk aversion; mean-variance; capacity; inventory; investment; balance; real options; network design
History: Accepted by Paul H. Zipkin, operations and supply chain management; received March 25, 2004. This paper was with the author 13 months for 3 revisions. Published online in *Articles in Advance* July 20, 2007.

1. Introduction

This paper studies how resource allocation in networks can mitigate risk exposure. It presents theory and insight on how risk attitude and network configuration drive the strategic placement of safety capacity and inventory for operational hedging. The networks considered here are designed and managed by a single expected utility maximizer. Design involves the sizing of resources, which include inventories as well as capacities, and management means processing to best fill market demands. Timing follows a two-stage recourse model: Resources are sized ex ante when the demand vector is uncertain but its probability distribution is known, while processing occurs after observing demand.

Sections 2 and 3 present the model, theory, and general results in terms of statistical quantities that allow for computation by simulation. These propositions hold for any portfolio of real options with general network topology and utility functions. To bring that theory to life, however, the remaining sections focus on newsvendor networks, which are linear recourse models that feature parsimony, tractability, and effectiveness in yielding insights into planning under uncertainty. Sections 4 through 7 each analyze a specific network in two steps. First, observations are made from a numerical study and intuitive explanations are proffered. Second, the insights are generalized as properties, which are statements for a specific network under certain conditions that are specified and proved analytically in the appendix.

After reviewing the single-resource case in §4, special attention is devoted to the three canonical newsvendor networks shown in Figure 1. All three serve two markets and are building blocks for general networks. The dedicated network features a dedicated resource for each market and pure diversification benefits that provide a natural or passive hedge: While its two resources lack operational dependence, the network profit has less variability than the sum of the individual resource profits. Aversion to financial variability thus induces resource investment dependence. The other two networks add a third resource, either in series or in parallel, that features operational flexibility benefits: demand pooling and ex post revenue maximizing allocation provide an active hedge that increases value by exploiting upside variations. In serial networks, each market requires some upstream dedicated work that is followed by a shared resource requirement. Examples are disk drive and computer manufacturing, where a common set of computers

Figure 1 **Three Canonical Building Blocks for Multiresource Networks That Are Analyzed in Detail**

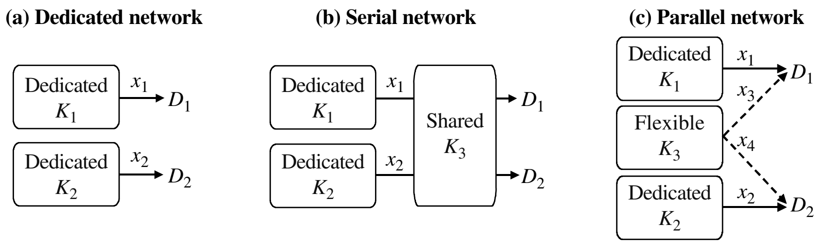

perform final burn-in and test routines. The serial network also is a core model for pure component commonality. The parallel network is a dedicated network augmented with auxiliary flexible capacity; it also models inventory substitution and resource redundancy. It allows a tailored response to uncertainty where dedicated capacity mainly fills base demand, while flexible capacity supplies variable demand. Economically, the shared resource is a complementary real asset, while the flexible resource is a substitute.

Risk exposure (or simply risk) refers to the undesirable consequence of a random prospect. In newsvendor models, the random prospect is typically called *demand* and is modeled by an exogenous probability distribution. The operational consequence of random demand is a likely, but undesirable mismatch between supply and demand manifested as overage or underage. The expected cost of overages and underages is called *mismatch cost*, following Cachon and Terwiesch (2006). The financial consequence of random demand is profit risk: profit variability risk as well as a decrease in expected profit.

Risk attitude describes how a decision maker perceives risk. Risk-averse agents prefer the expected value over the random variable. While risk-neutral newsvendors only care about the mismatch cost, risk-averse agents also care about profit variability risk. Traditionally, an increase in risk is equated with an increased mean-preserving spread (e.g., see Gollier 2001 for general definitions). For univariate normal random variables, this is equivalent to an increase in variance while keeping the mean constant. While demand variance impacts risk exposure (and will be used in our graphs), the remainder will illustrate that profit variances are the natural descriptors of risk exposure in networks. They summarize market and network interactions and capture the important impact of demand correlations, an increase of which typically increases profit risk.[1]

This paper establishes that risk attitude and network structure fundamentally change resource allocation. In contrast to single-resource settings, risk-averse newsvendors may invest *more* in networks than risk-neutral newsvendors: some resources and even total spending may exceed risk-neutral levels. With normally distributed demand, risk-averse newsvendors change resource levels roughly proportionally to demand variance (i.e., levels are quadratic in standard deviation), while risk-neutral agents adjust only proportionally to standard deviation.

These findings are explained in terms of hedging. Hedging is the action of a decision maker to mitigate a particular risk exposure. Operational hedging is risk mitigation using operational instruments. This definition is deliberately general to include risk-neutral agents as well as univariate settings. Holding excess assets such as stock or capacity reserves by a risk-neutral newsvendor is interpreted as operational hedging because it mitigates mismatch costs. This illustrates that operational hedging impacts expected profits; operational flexibility can even exploit risk and add value. Yet, the standard financial motivation for hedging is mitigation of profit variability risk,[2] which is the main topic of investigation in this paper. Hedging by "betting on two horses" or "not holding all eggs in one basket" presumes access to at least two risks whose counterbalancing effect is to reduce total risk. With multiple risks (demands) and multiple baskets (resources), newsvendor networks are a natural vehicle to study how operational instruments reduce total risk and may create value.

The analysis of the three canonical networks identifies three types of operational hedging that are summarized in Table 1. Risk mitigation through pure diversification or demand pooling steers the portfolio mix toward assets supplying lower-profit variance markets. These need not be the lower demand variance markets, which highlights the importance of profit variance to understanding risk-averse network

[1] Corbett and Rajaram (2006) show that intuition from multivariate normals extends beyond normal distributions. For example, for a broad class of distributions, aggregation of inventories is more valuable as demands are less-positively dependent.

[2] The *Oxford English Dictionary Online* (2004) defines hedging as "to surround with a hedge or fence as a boundary, or for purposes of defence, or to confine or restrict movement," while *Merriam-Webster's Collegiate Dictionary* (1998) states "to protect oneself from losing or failing by a counterbalancing action." "Most businesses insure or hedge to reduce risk, not to make money" by "taking on one risk to offset another," according to Brealey and Myers (2000, pp. 760 and 763).

Table 1 **Three Operational Instruments That Mitigate Profit Variability Risk and May Increase Expected Profit (Operational Flexibility Can Exploit Risk and Create Value) by Rebalancing Network Capacity**

Operational instrument	Portfolio of dedicated resources	Resource sharing	Resource flexibility
Driver			
Pure diversification	✓	✓	✓
Demand pooling	—	✓	✓
Ex post revenue maximization	—	✓	✓
Impact			
Mismatch cost	—	↓	↓
Profit variability risk	↓	↓	↓
Hedging			
Rebalance capacity mix	Toward assets serving lower profit variance markets	Away from (unique) shared asset	Toward (redundant) flexible asset

design. Capacity imbalance in the serial and parallel networks can remain even with perfectly positive correlations (i.e., in the absence of risk pooling). This isolates the contingent optimization option imbedded in shared and flexibile resources: they can steer and allocate production toward the higher-profit market. When markets differ in profitability, risk aversion rebalances capacity toward the (redundant) flexible resource, but away from the (unique) shared resource. (Given that both types of resources are "product-flexible," this means that the appropriate hedging action for product-flexible resources depends on their network position: increase when in parallel with dedicated resources, decrease when in series.) Capacity imbalance and allocation flexibility thus mitigate profit risk which confirms and refines their interpretation as operational hedges.

Section 8 concludes with managerial take-aways and discusses model limitations and extensions. The appendix contains one key proof; all other proofs are given in the online appendix (provided in the e-companion).[3]

Three research areas are most related to this article: risk-averse single-resource newsvendor models, newsvendor networks, and operational hedging. This article is a natural successor to the seminal work by Eeckhoudt et al. (1995), who prove that the optimal level of a single-resource newsvendor is always decreasing in risk aversion for general concave utility functions. This article extends their ingenious proof technique to a newsvendor network and shows that their unambiguous result does not hold for networks.

[3] An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Other studies of risk-averse single-resource newsvendor models include Atkinson (1979), Lau (1980), Spulber (1985), Anvari (1987), Lau and Lau (1999), Agrawal and Seshadri (2000a, b), Gan et al. (2005), Gaur and Seshadri (2005), Caldentey and Haugh (2006), and Chod et al. (2006), with multiperiod extensions in Bouakiz and Sobel (1992) and Chen et al. (2004).

This article is also a natural successor to Van Mieghem and Rudi (2002), who define and analyze newsvendor networks under expected profit maximization. Flexibility in risk-neutral parallel networks was first studied by Fine and Freund (1990) with a discrete math-programming model and by Van Mieghem (1998) with a newsvendor network model. Other related newsvendor network studies of the risk-neutral parallel network include Bassok et al. (1999), Hale et al. (2000), Rudi (2000), Netessine et al. (2002), Van Mieghem (2004), Bish and Wang (2004), and Goyal and Netessine (2007). The risk-neutral serial network was studied in Harrison and Van Mieghem (1999) and extended in Van Mieghem (2003). As far as we are aware, the only other paper on risk-averse newsvendor networks is Tomlin and Wang (2005), which complements this one in terms of research question and treatment of risk attitude. They investigate flexibility and dual sourcing in unreliable newsvendor networks and consider both loss aversion and conditional value-at-risk. With unreliable resources and risk aversion, inherent redundancy in a dedicated network can make it the preferred strategy to a flexible resource even if the latter is cheaper.

This article also relates to the literature on operational hedging, a term promulgated by Huchzermeier and Cohen (1996). They provided a valuation model and numerical evidence that embedded real options like contingent supply and production switching reduce downside risk in the presence of exchange-rate uncertainty. Ding et al. (2007) review and add recent analytical advances on joint operational and financial hedging of exchange-rate risk. Hedging in that setting of price uncertainty also may lead to an increase in capacity, similar to our finding under demand uncertainty. Operational hedging by means of flexibility and capacity imbalance in newsvendor networks under demand uncertainty was studied in Harrison and Van Mieghem (1999) and extended in Van Mieghem (2003). Boyabatlı and Toktay (2004) survey and critically discuss papers on operational hedging, most of which assume expected profit maximization. Hedging obviously requires the presence of uncertainty but its standard objective is to reduce risk, not to make money. This paper shows that risk aversion magnifies these operational constructions, establishing that they mitigate risk and strengthening their interpretation as operational hedges.

2. Model

2.1. Decision Problem

Consider a firm that has n different real assets or "means of processing," which we will call *resources*. We adopt the notation of Van Mieghem (2003) and denote its *resource portfolio* by the nonnegative resource vector $\mathbf{K} \in \mathbb{R}^n_+$, whose ith component represents the level of resource i available for processing during the period. The resulting operating profit gained at the end of the period is a random variable that is a function of the available resources. Let $\pi(\mathbf{K}, \omega)$ denote this operating profit function, where ω is a sample point in the sample space Ω. The operating profit function is concave in the resource vector $\mathbf{K}$, reflecting the natural assumption of decreasing marginal returns from investment. The financial investment cost to install resource levels $\mathbf{K}$ is denoted by $C_0(\mathbf{K})$. As usual, C_0 is assumed to be convex to guarantee a well-behaved concave optimization problem. A typical economic assumption, however, is that C_0 exhibits economies of scale and slightly concave functions or the addition of a fixed cost often does not pose a problem.

The research problem is to decide on the resource vector $\mathbf{K}$ given a probability distribution P on the sample space Ω. The risk-averse decision maker has von Neumann-Morgenstern preferences and maximizes expected utility of terminal wealth. (This can model the behavior not only of a small firm, but also of an individual decision maker in a large firm. After all, decisions ultimately are made by a few individuals who may be risk averse. Shortcomings of the utility approach are discussed in the conclusion.) The agent has an initial endowment or wealth W_0 and can borrow and lend without limitations at the risk-free interest rate r. Let $W = (1+r)W_0$ and $C(\mathbf{K}) = (1+r)C_0(\mathbf{K})$ denote the initial endowment and resource costs in end-of-period monetary units. Let V denote the net future value of the firm's resource portfolio: $V(\mathbf{K}, \omega) = \pi(\mathbf{K}, \omega) - C(\mathbf{K})$. Under these assumptions, a risk-averse investor will choose a resource vector $\mathbf{K}^u$ that maximizes expected future utility $U(\mathbf{K}; W)$, where

$$U(\mathbf{K}; W) = \mathbb{E}u(V(\mathbf{K}, \omega) + W) \tag{1}$$

and $u(\cdot)$ is strictly increasing (such that the investor prefers "more over less") and concave (such that the investor is "risk averse" and prefers the expected value over the risky outcome) on $\mathbb{R}$. Thus, this is an investment model of a risk-averse agent who can invest in a portfolio with one riskless asset and n risky assets.

2.2. Technical Assumptions

The only essential requirements are concavity: The function $V(\cdot, \omega)$ is concave on $\mathbb{R}^n_+$ for almost every ω, and $u(\cdot)$ is concave on $\mathbb{R}$. The analysis below presents expressions that allow computationally efficient, simulation-based optimization, also known as infinitesimal perturbation analysis (IPA). These expressions require the interchange of differentiation and integration, which is typically justified by a monotone or bounded convergence theorem. This requires certain technical conditions that, in one way or another, bound the derivatives as shown in the appendix.

2.3. Notation

Let $\mu(\mathbf{K})$ and $\sigma^2(\mathbf{K})$ denote the mean and variance of the value of the resource portfolio $\mathbf{K}$ expressed in end-of-period monetary units: $\mu(\mathbf{K}) = \mathbb{E}V(\mathbf{K}, \omega)$ and $\sigma^2(\mathbf{K}) = \mathbb{E}(V(\mathbf{K}, \omega) - \mu(\mathbf{K}))^2 = \mathbb{E}\pi^2(\mathbf{K}, \omega) - (\mathbb{E}\pi(\mathbf{K}, \omega))^2$. It will be useful to denote the marginal operating profit $\nabla_{\mathbf{K}}\pi(\mathbf{K}, \omega)$ by $\boldsymbol{\lambda}(\mathbf{K}, \omega)$, the marginal resource cost $\nabla C(\mathbf{K})$ by $\mathbf{c}(\mathbf{K})$, and the Hessian of the expected value $\nabla^2_{\mathbf{K}}\mathbb{E}V(\mathbf{K}, \omega)$ by $H(\mathbf{K})$.

2.4. Newsvendor Networks

To illustrate the general theory, we will use newsvendor networks, a tractable class of linear recourse models introduced in Van Mieghem and Rudi (2002) and defined via three data sets: (1) Demand data: when results apply only to newsvendor networks, the sample point ω will be replaced by the demand vector $\mathbf{D}$ which has continuous probability measure P. (2) Financial data: activity vector $\mathbf{x}$ yields gross margin $\mathbf{m}'\mathbf{x}$, where $\mathbf{m}$ equals price minus any marginal processing and transportation cost. The unit capacity investment costs are $\mathbf{c}_{\mathrm{K}}$ while inventory $\mathbf{S}$ incurs unit purchasing and holding costs $\mathbf{c}_{\mathrm{S}}$ and $\mathbf{c}_{\mathrm{H}}$; and unmet demand incurs shortage cost $\mathbf{c}_{\mathrm{P}}$. (3) Network data: the input-output matrices $\mathbf{R}_{\mathrm{S}}$ and $\mathbf{R}_{\mathrm{D}}$ and the capacity consumption matrix $\mathbf{A}$. Here $R_{\mathrm{S},ij}$ denotes the amount of input stock i consumed per unit of activity j, $R_{\mathrm{D},ij}$ is the amount of output i per unit of activity j, and A_{kj} is the amount of resource k capacity consumed per unit of activity j.

Let $\mathbf{v} = \mathbf{m} + \mathbf{R}'_{\mathrm{D}}\mathbf{c}_{\mathrm{P}} + \mathbf{R}'_{\mathrm{S}}\mathbf{c}_{\mathrm{H}}$ denote the *net value vector* associated with the various processing activities. Denote an ex post optimal activity vector by $\mathbf{x}^*(\mathbf{K}, \mathbf{S}, \mathbf{D})$. It solves the linear program $\max \mathbf{v}'\mathbf{x}$ subject to $\mathbf{x} \geq 0$, $\mathbf{R}_{\mathrm{S}}\mathbf{x} \leq \mathbf{S}$, $\mathbf{R}_{\mathrm{D}}\mathbf{x} \leq \mathbf{D}$, and $\mathbf{A}\mathbf{x} \leq \mathbf{K}$. The operating profit function then is $\pi(\mathbf{K}, \mathbf{S}, \mathbf{D}) = \mathbf{v}'\mathbf{x}^*(\mathbf{K}, \mathbf{S}, \mathbf{D}) - \mathbf{c}'_{\mathrm{P}}\mathbf{D} - \mathbf{c}'_{\mathrm{H}}\mathbf{S}$, which is concave for any newsvendor network. For notational simplicity and without loss of generality, we will treat all resources as capacities so that the relevant resource levels are $\mathbf{K}$. For the three networks in Figure 1, we will label market or product 1 as the more profitable one: $v_1 \geq v_2$. Also, μ_i and σ_i $(i = 1, 2)$ will denote the mean and standard deviation of market-i demand D_i and ρ the correlation coefficient between D_1 and D_2. (To avoid confusion with the mean and standard deviation of value, the demand mean and standard deviation will always have a subscript.) As usual, let Φ and ϕ denote the standard normal cumulative distribution and density function.

3. General Theory and Results

This section reviews relevant concepts and summarizes theory and results for general resource portfolio problems. These results are applied and extended to specific newsvendor networks in the remaining sections.

3.1. Optimality Conditions and Results for General Utility Functions

Proposition 1. *The expected utility function $U(\mathbf{K}; W)$ is concave in $\mathbf{K}$ for any W.*

Given that adding risk aversion does not destroy concavity, the optimization problem remains well behaved with sufficient first-order conditions:

Proposition 2. *An interior optimal investment $\mathbf{K}^u$ for a risk-averse utility function u solves the necessary and sufficient condition $\nabla U(\mathbf{K}^u) = 0$, where*

$$\nabla U(\mathbf{K}) = \mathbb{E}[(\boldsymbol{\lambda}(\mathbf{K}, \omega) - \mathbf{c}(\mathbf{K}))u'(V(\mathbf{K}, \omega) + W)]. \quad (2)$$

Let $\mathbf{K}^n$ denote an optimal resource vector for the risk-neutral case. (The superscript n is mnemonic for "risk-neutral" and also for "newsvendor.") In the risk-neutral setting, the utility function has no curvature so that u' is a constant and (2) simplifies to

$$\nabla U^n(\mathbf{K}^n) = \mathbb{E}[\boldsymbol{\lambda}(\mathbf{K}^n, \omega)] - \mathbf{c}(\mathbf{K}^n) = 0, \quad (3)$$

in agreement with Proposition 1 in Van Mieghem and Rudi (2002). Condition (3) is the multidimensional generalization of the familiar critical fractile condition. In a newsvendor network, the operating profit is the maximum of an underlying linear program, as discussed earlier. Then, there exists a partition of $\Omega = \bigcup_j \Omega_j(\mathbf{K})$ such that $\mathbb{E}\boldsymbol{\lambda}(\mathbf{K}, \omega) = \sum_j \boldsymbol{\lambda}_j P(\Omega_j(\mathbf{K}))$, where $\boldsymbol{\lambda}_j$ is the constant shadow vector of the resource constraint $\mathbf{Ax} \leq \mathbf{K}$ in event $\Omega_j(\mathbf{K})$. The optimality condition $\sum_j \boldsymbol{\lambda}_j P(\Omega_j(\mathbf{K}^n)) = \mathbf{c}$ generalizes the critical fractile condition of the single-resource newsvendor to higher dimensions. It sets the likelihood of a resource being a bottleneck proportional to its cost. In the risk-averse setting, Proposition 2 shows that this condition is perturbed in the sense that it is multiplied by the state-dependent marginal utility of money u'. Given that only changes from a constant u' matter, the perturbation impact increases proportionally to the curvature of u, exactly what risk aversion is expected to do.

Proposition 2 directly suggests a simulation-based gradient (steepest ascent) optimization method that is computationally efficient and unbiased, a characteristic of IPA. It is computationally efficient because one need only draw a single random sample $\{\omega_i\}$ that can be used during the entire optimization using the following algorithm. (This method with 10,000 sample demand vectors was used to produce the graphs in this paper.) Choose an initial value $\mathbf{k}_0$ and set $m = 0$. Compute the shadow vector $\boldsymbol{\lambda}(\mathbf{k}_m, \omega_i)$ and its weight $u'(V(\mathbf{k}_m, \omega_i) + W)$ for each point ω_i in the sample and compute their weighted sum to find an unbiased estimate for $\nabla U(\mathbf{k}_m)$. Find the maximizer $\mathbf{k}_{m+1}$ of U along the half-line $\mathbf{k}_m + t\nabla U(\mathbf{k}_m)$, where $t \geq 0$. Iterate until $\|\nabla U(\mathbf{k}_m)\|$ or $(\mathbf{k}_{m+1} - \mathbf{k}_m)$ are below a tolerance.

Much of the intuition behind our results will be explained in terms of the gradient of the value variance (or variance of operating profits) $\nabla\sigma^2(\mathbf{K})$, which also can be expressed in terms of the shadow vector $\boldsymbol{\lambda}(\mathbf{K}, \omega)$. Recall that the covariance between x and y is defined as $\operatorname{Cov}(x, y) = \mathbb{E}xy' - \mathbb{E}x\,\mathbb{E}y'$. (Some care is needed with dimensions: if y is scalar, then $\operatorname{Cov}(x, y)$ has the dimensions of x and contains the componentwise covariances.)

Proposition 3. *The gradient of the value variance is*

$$\nabla\sigma^2(\mathbf{K}) = 2\operatorname{Cov}(\boldsymbol{\lambda}(\mathbf{K}, \omega), \pi(\mathbf{K}, \omega)),$$

and evaluated at the risk-neutral solution $\mathbf{K}^n$, the following are equivalent:

$$\begin{aligned}\nabla\sigma^2(\mathbf{K}^n) &= 2\mathbb{E}[(\boldsymbol{\lambda}(\mathbf{K}^n, \omega) - \mathbf{c})V(\mathbf{K}^n, \omega)] \\ &= 2\operatorname{Cov}(\boldsymbol{\lambda}(\mathbf{K}^n, \omega), \pi(\mathbf{K}^n, \omega)).\end{aligned}$$

To discuss how risk-averse managers should adjust their resource portfolio relative to the risk-neutral one, we need a measure of risk aversion. Pratt (1964) showed that the *coefficient of absolute risk aversion*, $\gamma(x) = -u''(x)/u'(x)$, is a simple measure of local risk aversion, while there is no simple measure of risk aversion "in the large." (The coefficient of absolute risk aversion also meshes with our focus on additive risk, while relative risk aversion is better for multiplicative risk.) Nevertheless, comparisons of risk aversion between two utility functions u_1 and u_2 can be made simply: u_1 is (weakly) more risk averse than u_2 if and only if $\gamma_1(x) \geq \gamma_2(x)$ for all x, which is equivalent to the statement that $u_1(\cdot) = h(u_2(\cdot))$ for some strictly increasing, concave function h according to Pratt (1964, Theorem 1). We say that an optimal decision variable K_i is "increasing in risk aversion" if $K_i^{u_1} \leq K_i^{u_2}$ whenever u_2 is more risk averse than u_1. This means that, when comparing two managers, the more risk-averse manager will invest more than the other. We also will loosely say that a utility function u is moderately risk averse if there exists an $\varepsilon > 0$ such that if $\max_x |\gamma(x)| < \varepsilon$, then its optimal $\mathbf{K}^u$ is in an ε-neighborhood of $\mathbf{K}^n$.

3.2. Local Impact of Risk Aversion and CARA Utility

It is instructive to analyze the impact of a small, or local, increase in risk aversion, relative to the risk-neutral case. To characterize the associated optimal resource adjustment vector, we consider a parameterized class of utility functions and apply the implicit function theorem. While there is no simple measure to rank general utility functions in terms of degrees of risk aversion, the class of utility functions with constant absolute risk aversion (CARA) is a very useful exception. The condition that $\gamma(x) = -u''(x)/u'(x)$ be a constant yields a differential equation that is satisfied only by exponential functions:

$$u_{\text{CARA}}(x) = \alpha e^{-\gamma x} + \beta, \tag{4}$$

where $\alpha \le 0$ and $\gamma \ge 0$ to ensure the utility function is increasing concave. (It is convenient to choose $\alpha = -\beta = -\gamma^{-1}$ so that $u_{\text{CARA}}(x;\gamma) \to x$ when $\gamma \downarrow 0$, as will be assumed in later graphs.)

Exponential utility functions are theoretically and mathematically appealing. First, the scalar γ gives a simple cardinal measure of risk aversion. Second, its optimal actions are independent of the initial wealth W (under the earlier assumptions on borrowing and lending). Indeed, the optimality condition (2) simplifies to

$$\nabla U(\mathbf{K}) = 0 \Leftrightarrow \mathbb{E}[(\boldsymbol{\lambda}(\mathbf{K},\omega) - \mathbf{c}(\mathbf{K}))\exp(-\gamma V(\mathbf{K},\omega))] = 0. \tag{5}$$

Therefore, the CARA-optimal resource vector is a function of the scalar γ only, which we will denote by $\mathbf{K}^{\text{CARA}}(\gamma)$. Similarly, let $U^{\text{CARA}}(\gamma)$ denote the associated utility $\mathbb{E}u_{\text{CARA}}(V(\mathbf{K}^{\text{CARA}}(\gamma),\omega))$. Third, with CARA, the expected utility function is directly expressed in terms of the characteristic function of the value function, which is defined as $\psi_{V(\mathbf{K},\omega)}(t) = \mathbb{E}\exp(itV(\mathbf{K},\omega))$. It directly follows that $U_{\text{CARA}}(\mathbf{K}) = \alpha e^{-\gamma W}\psi_{V(\mathbf{K},\omega)}(i\gamma) + \beta$. (The characteristic function has simple closed-form expressions if π has a normal, exponential, gamma, or uniform distribution.) Finally, CARA allows us to characterize the optimal resource adjustment vector for small increases in risk around $\mathbf{K}^n$. (Recall that $H(\mathbf{K}^n)$ is the Hessian matrix of $\mathbb{E}V(\mathbf{K}^n,\omega)$.)

Proposition 4. *In the neighborhood of the risk-neutral case, small* CARA *risk aversion sets* $\mathbf{K}^{\text{CARA}}(\gamma) = \mathbf{K}^n + \gamma(d/d\gamma)\mathbf{K}^{\text{CARA}}(0) + o(\gamma)$, *where the optimal resource adjustment vector is*

$$\frac{d}{d\gamma}\mathbf{K}^{\text{CARA}}(0) = \tfrac{1}{2}H^{-1}(\mathbf{K}^n)\nabla\sigma^2(\mathbf{K}^n). \tag{6}$$

While it is known that the value variance should be decreasing with risk aversion around the risk-neutral solution (see the next subsection), the proposition explains how that is done. Recall that concavity implies that the Hessian is negative definite so that its diagonal elements are negative, but not necessarily its off-diagonal elements. Often, however, H^{-1} is diagonally dominant (a sufficient, but not necessary condition for concavity) so that the optimal adjustment vector is roughly in the negative or opposite direction of $\nabla\sigma^2(\mathbf{K}^n)$ (it really is a linear combination). In other words, an increase in risk aversion should adjust resource levels in the direction that reduces value variance.

So, should risk-averse managers set smaller resource levels than risk-neutral managers? While the answer is affirmative for single-resource problems, the remainder of this paper will show that such an unambiguous result does not extend to a bundle of resources. The gradient $\nabla\sigma^2(\mathbf{K}^n)$ captures the effect of increasing resource levels on the variance of operating profits. In a newsvendor setting, increasing the capacity or the feasible set of activity vectors $\mathbf{x}$ typically increases the variance of operating profits $\mathbf{v}'\mathbf{x}$. A positive gradient $\nabla\sigma^2(\mathbf{K}^n)$ together with a negative Hessian imply that the expected effect of risk aversion is to decrease resource levels. The remaining sections, however, will show that in a network the benefit from *increasing* some resource levels may outweigh its cost when correlations are sufficiently negative. In other words, *increasing* capacity in a network may decrease value variance.

3.3. The Mean-Variance Formulation and Hedging Effectiveness

Similar to financial portfolios, the effect of risk aversion on the configuration of a portfolio of real assets is often illustrated using a mean-variance (MV) formulation, which seeks to maximize

$$U_{\text{MV}}(\mathbf{K}) = \mu(\mathbf{K}) - \frac{\gamma}{2}\sigma^2(\mathbf{K}). \tag{7}$$

Let $\mathbf{K}^{\text{MV}}(\gamma)$ denote a maximizer of (7) and $U^{\text{MV}}(\gamma)$ the associated MV utility $U_{\text{MV}}(\mathbf{K}^{\text{MV}}(\gamma))$. Expressions (2) and (3) directly yield an expression for the gradient

$$\begin{aligned}\nabla U_{\text{MV}}(\mathbf{K}) &= \mathbb{E}[\boldsymbol{\lambda}(\mathbf{K},\omega)] - \mathbf{c}(\mathbf{K}) \\ &\quad - \gamma\,\text{Cov}(\boldsymbol{\lambda}(\mathbf{K},\omega), \pi(\mathbf{K},\omega)),\end{aligned} \tag{8}$$

which is useful for gradient-based numerical optimization via simulation. It is well known that the MV solution exhibits smaller value variance than the risk-neutral solution.[4] Proposition 4 shows that risk aversion adjusts resource levels in the direction that reduces variance, which also follows from (8).

MV formulations are hard to justify from a theoretical perspective and have several undesirable properties.[5] They are better viewed as approximations—the remainder suggests that the MV approximation is

[4] By definition, $U_{\text{MV}}(\mathbf{K}^{\text{MV}}) \ge U_{\text{MV}}(\mathbf{K}^n)$ and $\mu(\mathbf{K}^n) \ge \mu(\mathbf{K}^{\text{MV}})$ so that $\sigma^2(\mathbf{K}^{\text{MV}}) \le \sigma^2(\mathbf{K}^n)$.

[5] Typical justifications include normally distributed returns or quadratic preferences, both of which are problematic. In newsvendor

Figure 2 An Estimate of the Coefficient of Risk Aversion γ Suffices to Identify the Resource Investment on the Frontier That the Mean-Variance Formulation Recommends

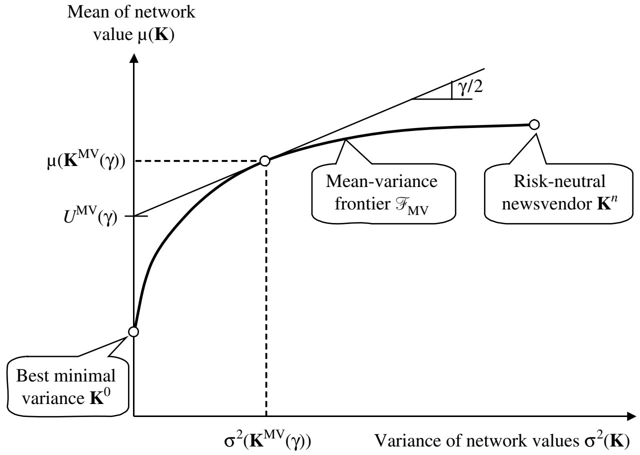

quite reasonable for newsvendor networks with moderate risk aversion—with the two practical benefits discussed in Van Mieghem (2003): (a) They are implementable because only two moments are required, which can be estimated. (b) They are useful in the sense that they provide "good recommendations" even when the decision maker does not know her utility function. "Good recommendations" are those investments that are on the MV frontier $\mathcal{F}_{\mathrm{MV}}$, which is the set of MV efficient investments that embody the trade-off between expected return and risk. Given a frontier, all that is needed is a reasonable estimate of the decision-maker's risk attitude γ, which directly identifies the recommended investment on the frontier as shown in Figure 2.

The standard objective of hedging suggests to measure hedging effectiveness by assessing how much risk can be reduced while giving up little expected return. In MV terms: How much can variance be reduced with little impact on mean value? The local sensitivity of expected return to optimal reduction of variance risk around the optimal risk-neutral portfolio $\mathbf{K}^n$ is measured by the curvature of the MV frontier, as shown in the appendix:

$$\mathcal{F}''_{\mathrm{MV}}(\sigma^2(\mathbf{K}^n)) = \frac{1}{\nabla'\sigma^2(\mathbf{K}^n)H^{-1}(\mathbf{K}^n)\nabla\sigma^2(\mathbf{K}^n)}. \quad (9)$$

This quantitative theoretical measure of hedging effectiveness reconfirms the importance of $\nabla\sigma^2(\mathbf{K}^n)$: effectiveness increases if the gradient of variance is large, in line with intuition.

networks, capacity constraints preclude normally distributed operating profits. In general, quadratic preferences $u(x) = \alpha(x-\beta)^2$ require the value function $V(\mathbf{K}^n, \omega)$ to be bounded a.s. by β for the expected utility to be concave increasing. MV utilities also have increasing absolute risk aversion and a bliss point after which utility is decreasing in wealth.

The remainder of this paper applies the general results of this section first to the single-resource problem, which will serve as our base case, and then to the three canonical networks of Figure 1.

4. The Base Case: The Single-Resource Newsvendor

Let us review how risk affects the single-resource newsvendor problem before delving into networks. To ensure a profitable investment, return should exceed cost: $v_1 \geq c_1$.

4.1. Risk Exposure

The optimal resource level for a risk-neutral newsvendor exposed to univariate demand risk solves (3) or $v_1 P(D_1 > K_1^n) = c_1$. The optimal service probability for any demand distribution is $1 - c_1/v_1$ and depends solely on the ratio of unit cost to unit return; it is also called the *critical fractile*. Thus, K_1^n is below, at, or above the demand median if the critical fractile is below, equal to, or above 1/2, respectively. For normal demand, the risk-neutral resource level is linear in demand standard deviation: $K_1^n = \mu_1 + z_1^n \sigma_1$, where $\Phi(z_1^n) = 1 - c_1/v_1$. Increased risk exposure (in the sense of an increasing mean-preserving spread, i.e., increase σ_1) then leads to a decrease, no change, or increase of the resource level if the critical fractile is below, equal to, or above 1/2, respectively. Figure 3 illustrates these facts with a numerical study with normal demand with $\mu_1 = 1$. (All graphs in this paper have γ going from 0 to 2.5 in 0.25 increments, unless indicated otherwise. The risk-neutral case is denoted by "o," which is mnemonic for $\gamma = 0$.)

4.2. Risk Aversion

The seminal work by Eeckhoudt et al. (1995) proves that the optimal resource level for a single-resource risk-averse newsvendor is always below the risk-neutral level for any concave utility function. Moreover, the optimal level always decreases in risk aversion so that the total monetary investment always decreases in risk aversion.[6] (This result is "global" in that it holds whenever a concave utility function $u(\cdot)$ is replaced by a "more concave" $h(u(\cdot))$ as discussed earlier.) They also show that the impact of an increased risk exposure on the risk-averse optimal

[6] The celebrated two-fund separation result of financial economics has established that the investment in the risky asset $C(K^n)$ decreases in risk aversion for the standard one-safe, one-risky financial asset portfolio. That result does not apply here because our real assets have nonlinear payoffs (nor does its key insight that the optimal relative configuration of risky assets is independent of an investor's coefficient of risk aversion hold here). Nevertheless, Eeckhoudt et al. (1995) showed that the investment in the risky real asset does always decrease in risk aversion for the single-resource newsvendor.

Figure 3 **With Normally Distributed Demand, Risk-Neutral Single-Resource Levels Are Linear (Labeled "o") in Demand Standard Deviation with Slope Depending on the Critical Fractile. Risk-Averse Capacities Are Roughly Quadratic and Below Risk-Neutral Levels**

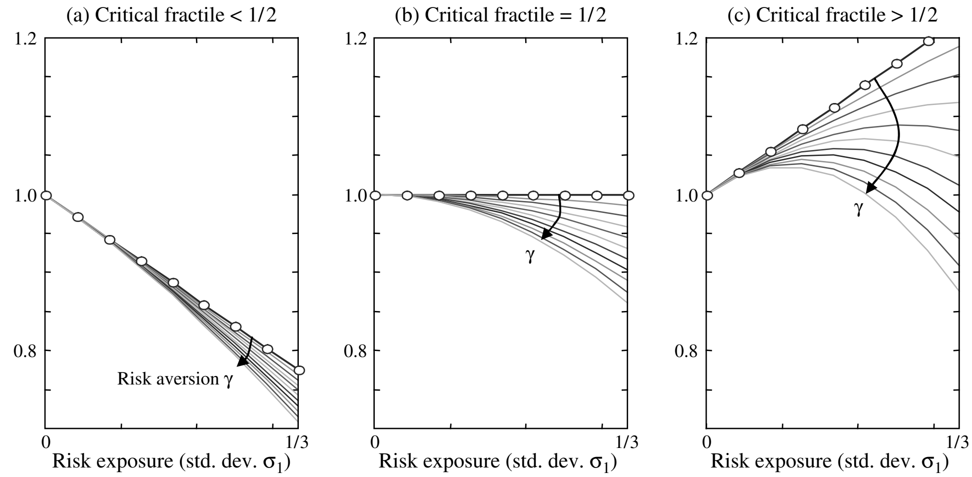

level is ambiguous, as Figure 3 illustrates for CARA utility.

This agrees with the optimal local risk-adjustment specified in Proposition 4, which can be solved analytically:

Property 1. *The optimal risk-adjustment* (6) *for the single-resource newsvendor problem with normally distributed demand and* CARA *utility is*

$$\frac{d}{d\gamma}K(0) = -\sigma_1^2 c\left(1+\frac{z^n(1-c/v)}{\phi(z^n)}\right) \leq 0,$$

$$\text{where } z^n = \Phi^{-1}\left(1-\frac{c}{v}\right).$$

The interesting insight from this property is that the local optimal adjustment due to risk aversion is proportional to the demand variance. This implies that, for normal demand, risk-averse resource levels (at least near the risk-neutral line) are quadratic in demand standard deviation. This is in stark contrast with the risk-neutral case, where investment levels are only proportional to demand standard deviation. This quadratic dependence explains the ambiguity of the impact of increased risk exposure on risk-averse resource levels for a single-resource newsvendor and extends to the three canonical networks (and is conjectured to be a general effect).

The property also illustrates that risk-averse investment is analytically complex, even in the single-resource case. It depends on all model parameters: demand uncertainty (mean μ_1 and std. dev. σ_1) and financials v and c. The independence of mean demand is unique to CARA; all other utility functions have wealth dependence and are thus also dependent on average demand.[7] Even in this simplest of networks, one cannot guarantee that the adjustment magnitude is monotone in c or v, except if $c/v > 1/2$, in which case the magnitude of the adjustment is increasing in v. It is even harder to characterize the conditions for effective hedging. According to (9), this amounts to identifying the parameters that yield a modest frontier curvature. While this curvature can be computed analytically for the single-resource case, the expression

$$\nabla'\sigma^2(K^n)H^{-1}(K^n)\nabla\sigma^2(K^n) = -4\sigma_1^3 c^2 v\phi(z^n)\left(1+\frac{(1-c/v)z^n}{\phi(z^n)}\right)^2$$

defies simple insights.

Numerical optimization with normal demand[8] allows some additional observations on the impact of risk aversion. Figure 4 compares $\mathbf{K}^{\mathrm{CARA}}(\gamma)$ with $\mathbf{K}^{\mathrm{MV}}(\gamma)$ and the associated utility levels $U^{\mathrm{CARA}}(\mathbf{K}^{\mathrm{CARA}}(\gamma))$ with $U^{\mathrm{CARA}}(\mathbf{K}^{\mathrm{MV}}(\gamma))$ for a representative situation (normal demand with $\mu_1 = 1$, $\sigma_1 = 1/3$, $c = 0.5$, and $v = 1$, so that $K^n = 1$). The left panel of Figure 4 shows an optimal risk adjustment that is near linear in risk aversion and larger than what MV suggests, a finding that extends to the three canonical networks. (All our numerical work suggests that $K^{\mathrm{CARA}}(\gamma)$ is convex decreasing if $c/v < 1/2$, linear if $c/v = 1/2$, and concave decreasing otherwise, but with very weak nonlinearity.) Near linearity implies that the risk-adjustment gradient is almost constant, which broadens the applicability of the local risk-adjustment result of Proposition 4. While the MV approximation underadjusts resources, the right panel of Figure 4

[7] For example, for DARA utility $u(x) = x^q$ with $x > 0$ and $0 < q < 1$, the optimal risk adjustment $(d/dq)\mathbf{K}(0)$ is $-H^{-1}\mathbb{E}(\boldsymbol{\lambda} - \mathbf{c})\ln(V + W)$ and depends on wealth and mean demand. No simple closed-form expressions are available.

[8] CARA and MV objective functions can be expressed analytically with normal demand so no simulation is needed.

Figure 4 Comparison of Optimal and Mean-Variance Capacity and Utilities for the Single-Resource Newsvendor Problem as a Function of Risk Aversion (Normal Demand $\mu_1 = 1$, $\sigma_1 = 1/3$, $c = 0.5$, $v = 1$)

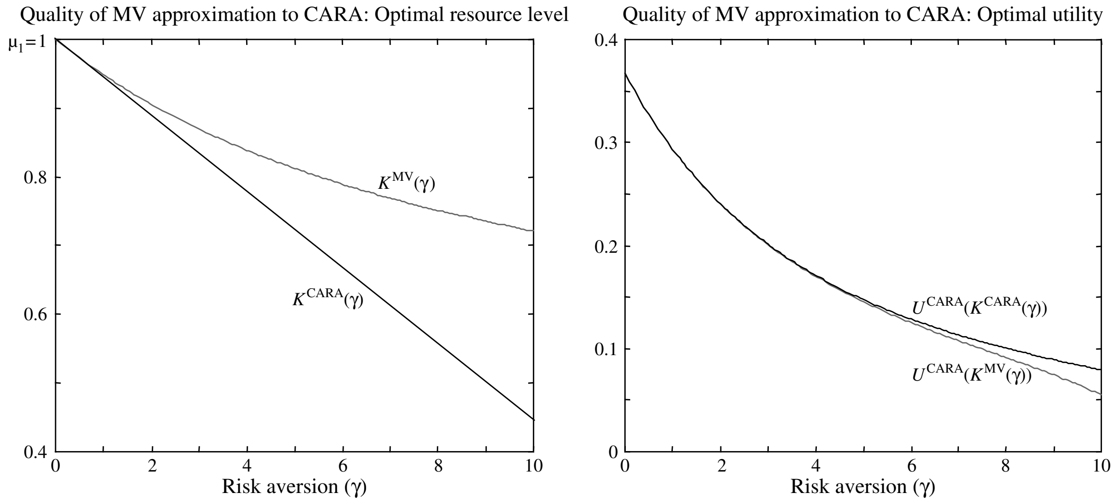

suggests that the impact on utility remains small for moderate levels of risk aversion. (Note that, as single-variable trajectories, the efficient frontier and the MV frontier coincide for single-resource problems.)

4.3. Operational Hedging

In the absence of risk, any newsvendor invests in a resource level equal to the deterministic demand μ_1. The change in the resource level due to risk exposure is called its *safety level*, e.g., safety capacity or safety inventory. For a risk-neutral newsvendor facing normal demand, the hedge is $K_1^n - \mu_1 = z_1^n \sigma_1$ and a change in risk σ_1 leads to a proportional decrease, no change, or increase depending on the critical fractile. A risk-averse newsvendor reacts stronger to a change in normal risk exposure: her safety-level adjustment is quadratic in σ_1. This is expressed more elegantly in terms of the standardized safety level $z_i(x)$ of a product i resource level x:

$$\text{standardized safety level} \quad z_i(x) = \frac{x - \mu_i}{\sigma_i}. \tag{10}$$

With normal demand, the standardized safety level is independent of risk exposure for a risk-neutral newsvendor (equal to the constant z_1^n), but decreases linearly in σ_1 for a risk-averse newsvendor.

The remaining sections will show that several of these single-resource insights extend to networks with the important exceptions that some risk-averse capacities as well as total spending may exceed risk-neutral levels. We start with the dedicated network before analyzing the serial and parallel networks.

5. Resource Diversification: The Dedicated Network

The dedicated network (Figure 1(a)) is the simplest network. While its two resources lack operational dependence, as a portfolio the network enjoys pure financial diversification and provides a "natural" or passive hedge: total profit variability risk is reduced by not putting all eggs into one basket. Risk aversion thus induces resource investment dependence and will favor the resource with lower-profit variance. Using a numerical study, we first discuss the insights from the single-resource newsvendor that carry over to the dedicated network and highlight the differences. Then, we generalize insights with analytically proven properties. To ensure that both investments are profitable, both market values should exceed their costs: $v_i \geq c_i$.

5.1. Risk Exposure

For a risk-neutral newsvendor, the dedicated network decomposes into two independent single-resource problems. With normal demand, capacities are linear in standard deviations with a slope determined by the critical fractile but independent of demand correlation. The "o"-connected lines in Figure 5 illustrate these facts with a numerical study assuming bivariate normal demand with unit means [1, 1] and $\sigma_1 = 0.75\sigma_2$, equal unit investment costs $c = [0.5, 0.5]$, and net values $v = [2, 1]$. (Market 1 thus is more profitable *and* has smaller demand risk.) The critical fractiles are 3/4 and 1/2 so that K_1^n is increasing in σ_1, while K_2^n is constant in σ_2.

5.2. Risk Aversion

As for a single-resource newsvendor, the numerical analysis of Figure 5 suggests that CARA utility-optimal capacities are decreasing in risk aversion and roughly quadratic in standard deviation for nonnegatively correlated normal demand. The striking difference is that K_2^u *increases* in risk aversion and standard deviation for negative correlation.

Figure 5 **With Normally Distributed Demand, the Risk-Neutral Dedicated Network Adjusts Capacities Linearly (Labeled "o") in Demand Standard Deviation Independent of Correlation. Risk-Averse Levels Are Quadratic and Can Be Above Risk-Neutral Levels with Strong Diversification Benefits at Low Correlation**

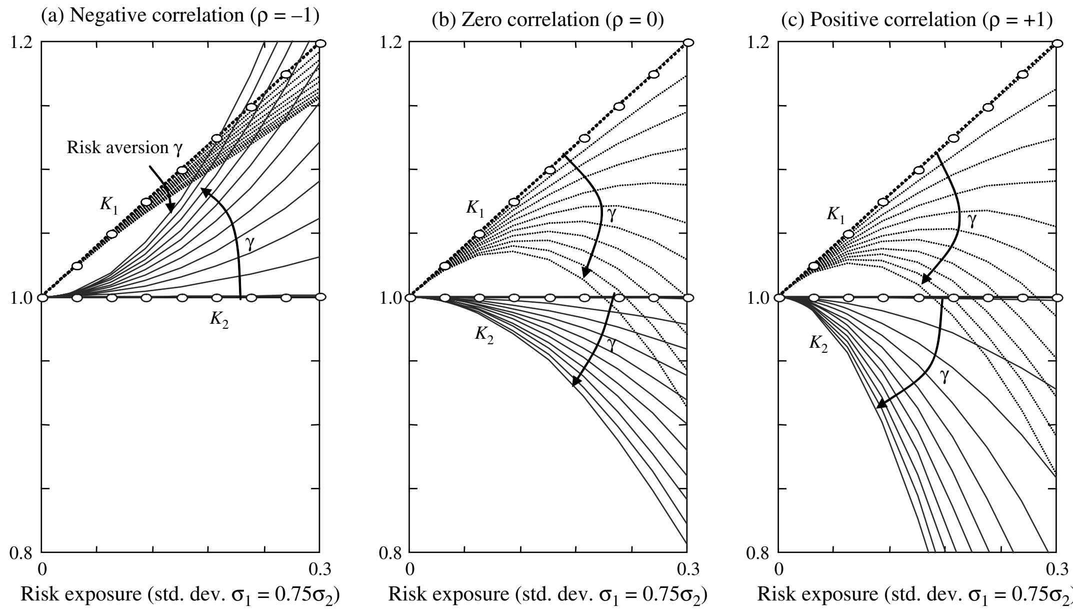

Consider Figure 6 to further investigate and explain the crucial role of correlation (for the same data as Figure 5). Figure 6(b) shows that K_2^u increases in risk aversion for any correlation coefficient ρ below about -0.45. Furthermore, for ρ below about -0.7, its increase outweighs the decrease of K_1^u so that the total spend of a risk-averse newsvendor *exceeds* the risk-neutral monetary investment (Figure 6(d)). Finally, Figure 6(c) shows that a risk-averse newsvendor strictly prefers more negatively correlated demand. Put another way: correlation is another measure of risk exposure. (Figure 6(c) also shows that the mean-variance frontiers dominate (by definition!) yet closely approximate the CARA efficient frontier in the relevant domain of hedging—up to 50% variance reduction with less than 10% value cost compared to the risk-neutral configuration.) These effects do not exist in the single-resource case and can be explained by risk diversification or hedging.

5.3. Operational Hedging

The dedicated network enjoys classic diversification benefits. Investing in two resources yields a diversified portfolio that pools risks and can be configured in the same manner that financial investors (re)balance portfolios to optimize their MV utilities. Increased risk aversion reduces profit variance at the expense of some mean value. Abstracting from capacity constraints, market i has profit variability risk (std. dev.) $v_i\sigma_i$. If we were to invest one dollar and allocate weight w_i to market i ($w_1 + w_2 = 1$), the profit risk of the portfolio would be $[(w_1v_1\sigma_1)^2 + 2\rho(w_1v_1\sigma_1)(w_2v_2\sigma_2) + (w_2v_2\sigma_2)^2]^{1/2}$, which increases in correlation ρ, but never exceeds the sum of market profit risks. In other words, risk-pooling benefits are zero at $\rho = +1$, but grow with smaller correlation and more equal profit variances. (A zero-variance portfolio or perfect hedge obtains with two perfectly negatively correlated assets with equal $v_1\sigma_1 = v_2\sigma_2$.) This risk-pooling benefit from diversification suggests that increased risk aversion will rebalance investment to favor the market with lower-profit standard deviation.

With capacity constraints, the diversification benefit must be balanced with its cost in terms of investment costs and changes in expected profits. The numerical analysis of Figure 6(d) shows that at sufficiently negative correlation, the diversification benefits are so large that a risk-averse newsvendor will invest more than a risk-neutral newsvendor! In our example, resource 2 has lower profit risk ($v_2\sigma_2 = (2/3)v_1\sigma_1$) and its fraction of total capacity (and investment) indeed increases in risk aversion except for high positive correlation, according to Figure 7. (This demonstrates the importance of profit over demand standard deviations; recall that market 1 is more profitable and has less demand risk: $\sigma_1 = 0.75\sigma_2$.) Note that rebalancing the capacity portfolio toward resource 2 is beneficial even with uncorrelated demand (although it then coincides with a reduction of both capacity levels). At $\rho = +1$, diversification benefits disappear and capacity mix choice is dominated by expected profit maximization, which favors the higher margins of market 1.

The key insights are that risk-averse newsvendors gain diversification benefits by rebalancing network

Figure 6 **Capacities, Frontiers, and Monetary Investment in the Dedicated Network as a Function of Correlation and Risk Aversion** ($\sigma_1 = 0.3$, $\sigma_2 = 0.4$)

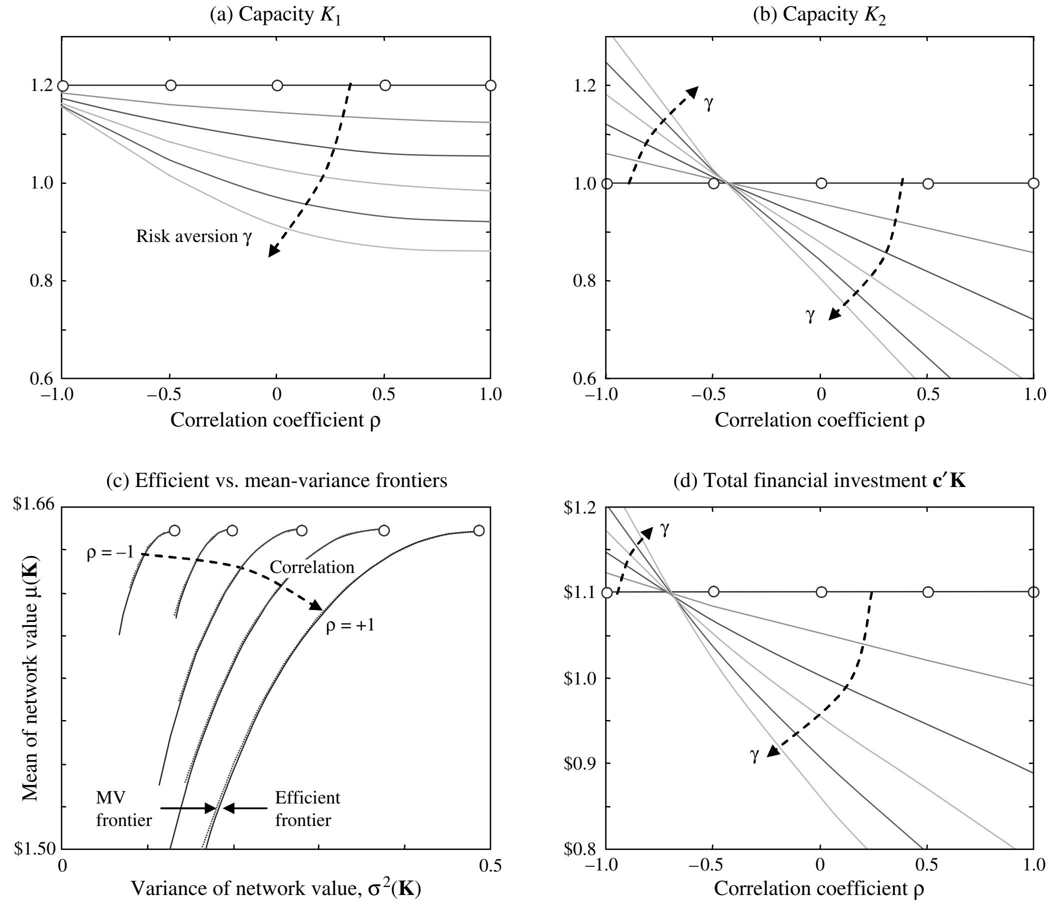

capacity toward resources that supply the lower-profit variance market. Total spending may even exceed the risk-neutral monetary investment.

5.4. Generalizations

These insights from the numerical analysis can be generalized to general demand distributions and utility functions under certain conditions. (Precise analytic conditions of all properties are relegated to their proofs in the appendix.)

Property 2 (Dedicated Network). $\mathbf{K}^u \leq \mathbf{K}^n$ *and resource levels always decrease in risk aversion if correlation* $\rho = 1$. *However,* $K_2^u \geq K_2^n$ *and increases in risk aversion with correlation* $\rho = -1$, $v_2\sigma_2 < v_1\sigma_1$, *moderate risk aversion, and standardized safety levels* $z_i(K_i^n) > 0$.

Remark. For symmetric demand distributions, it suffices that $c_i < 2v_i$ to yield positive standardized safety levels $z_i(K_i^n)$. The proof specifies (weaker) conditions analytically and extends the ingenious technique of Eeckhoudt et al. (1995) of bounding the marginal utility to a newsvendor network.

The risk-averse investor rebalances the capacity mix toward the lower-profit variance market, and the condition $v_2\sigma_2 < v_1\sigma_1$ suggests that this lowers the variance of the network profit. Indeed, profit variance is decreasing in K_2 so that risk aversion increases the level of K_2 beyond the risk-neutral level:

Property 3 (Dedicated Network). $\nabla_2\sigma^2(\mathbf{K}^n) \leq 0$ *and* $(d/d\gamma)K_2(0) \geq 0$ *with correlation* $\rho = -1$, $v_2\sigma_2 < v_1\sigma_1$, *moderate risk aversion, and standardized safety levels* $z_i(K_i^n) > 0$.

These theoretical results formalize and generalize the intuitive explanation in terms of the diversification benefits and costs. The result of Eeckhoudt et al. (1995) extends to dedicated networks if correlation $\rho = 1$, but fails with negative correlation and moderate costs and risk aversion; diversification benefits are the culprit. A continuity argument directly extends these results to a neighborhood of $\rho = \pm 1$: At high correlation, diversification benefits become negligible and both resource levels always decrease in risk aversion. It is easy to show that this also happens with CARA utility if

Figure 7 **Operational Hedging in the Dedicated Network Is Manifested by Capacity Rebalancing Toward the Minimal-Profit Variance Market if Diversification Benefits Outweigh Costs (Panels (a) and (b))**

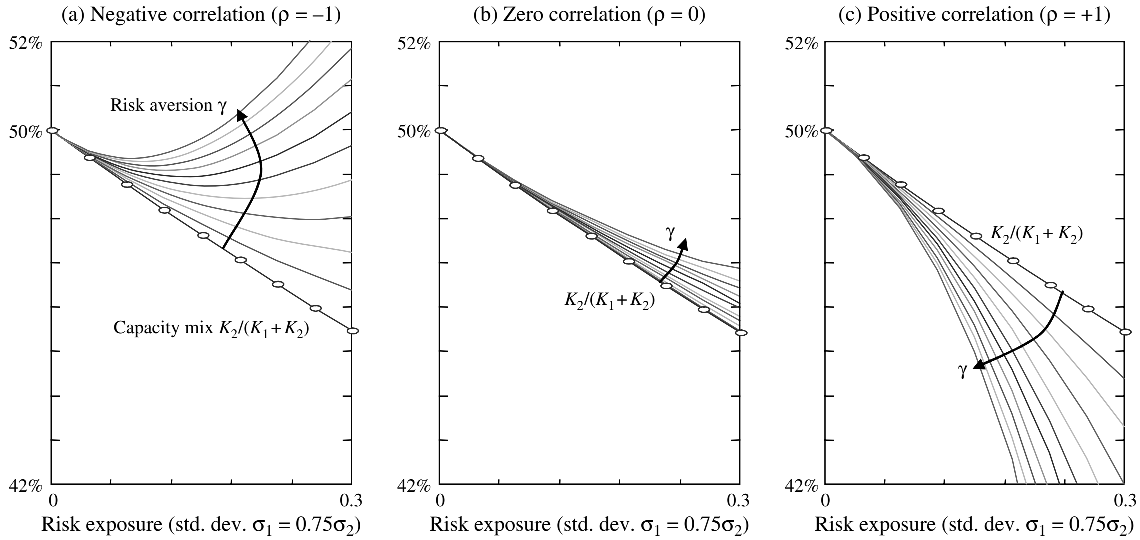

demands are independent (and thus $\rho = 0$).[9] At sufficiently low correlations, however, the diversification benefit is sufficient for a risk-averse newsvendor to increase the resource supplying the lower-profit variance market provided capacity is not too expensive. (This suggests that a negative correlation threshold—like $\rho \simeq -0.45$ in Figure 6(b)—may exist for CARA and normal demand below which $K_2^u \geq K_2^n$ and above which $K_2^u \leq K_2^n$.)

6. Resource Sharing and Complementarity: The Serial Network

To see how the insights from the dedicated network carry over to networks whose resources feature operational dependence, consider the serial network of Figure 1. Besides financial diversification, this network also benefits from operational flexibility through the shared resource. We will see that allocation flexibility adds two additional benefits: demand pooling and ex post revenue maximization, which provides an active hedge.

The question in the two-stage serial network is what the three resource levels should be, and the tension is about complementarity (between dedicated and shared resources) and bottlenecks. Expected profit optimization (3) sets the probability that a resource is a bottleneck proportional to its cost. Thus, any cost vector **c** is admissible in the serial network as long as both products or markets are economically viable: $c_i + c_3 < v_i$.

Again, we first discuss insights from a numerical study and then generalize. The numerical data is as earlier except that marginal investment costs are $c = (0.1, 0.1, 0.4)$. This choice allows meaningful comparisons because, abstracting from resource sharing, it yields the same critical fractile as for the dedicated network.

6.1. Risk Exposure

Risk-neutral capacities in newsvendor networks with normal demand are linear in (marginal) demand standard deviations. Indeed, the multivariate normal distribution scales linearly in standard deviations and the risk-neutral conditions (3) are expressed purely in probability fractiles. In contrast to the dedicated network, however, the slopes now depend on demand correlation as well as on the value-cost ratios c/v. The "o"-connected lines in Figure 8 illustrate these facts through a numerical study where both risk-neutral dedicated resource levels decrease in correlation, while the risk-neutral shared resource level K_3^n increases in correlation. The shared resource provides two benefits that were absent in the dedicated network and explains the impact of correlation on risk-neutral capacities:

(1) *Demand pooling* refers to serving multiple markets by one resource, e.g., through inventory centralization as first studied by Eppen (1979). *Diversification* refers to serving multiple markets from one portfolio or network. Demand pooling thus is a special form of diversification and risk pooling. It provides the risk mitigation benefits of §4 that are valued by risk-averse investors. In addition, demand pooling reduces

[9] The optimality equations then decouple and both resources behave as single-resource systems. CARA is necessary because in general, the wealth effect couples investments even with independent demand.

Figure 8 With Normally Distributed Demand, the Risk-Neutral Serial Network Adjusts Capacities Linearly (Labeled "o") in Demand Standard Deviation with Slope Dependent on Correlation. Risk-Averse Levels Are Quadratic and Can Be Above Risk-Neutral Levels at Low Correlation Reflecting Strong Risk-Pooling Benefits and Complementarity

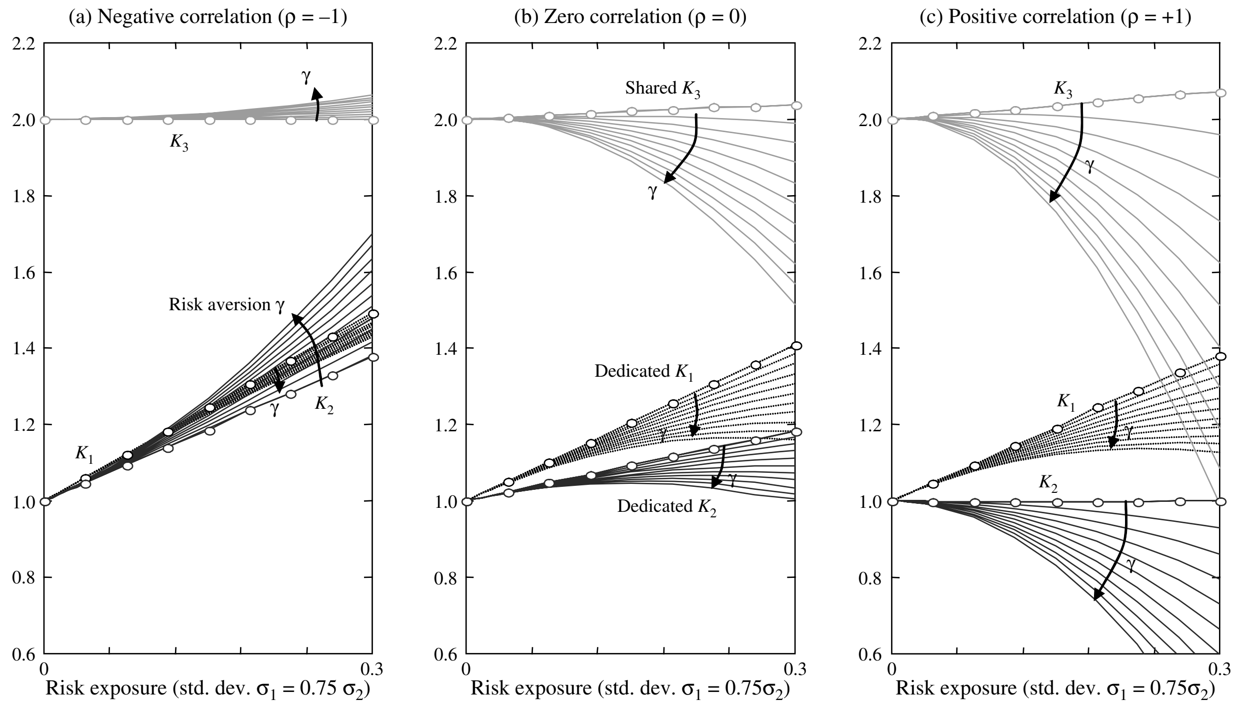

mismatch costs and affects the sizing of the shared resource[10] which also benefits a risk-neutral newsvendor (but less so as correlation increases).[11]

(2) *Ex post revenue* (*profit*) *maximization* refers to the ex post allocation option to steer or switch the output mix toward the more profitable market when the shared resource is capacity constrained. This benefit was first identified in Van Mieghem (1998, Proposition 3) and requires a market profit differential ($v_1 \neq v_2$). It is distinct from risk pooling as it survives even with perfect positive correlation (as the remainder will demonstrate). Indeed, by exploiting upside variations, it increases profits and thus also benefits risk-neutral newsvendors.

6.2. Risk Aversion

The numerical study suggests that the CARA utility dedicated capacities with normal demand mirror those in the dedicated network: In Figure 8, both appear roughly quadratic in standard deviation and K_2^u as well as the shared resource level K_3^u increase in risk aversion for strongly negatively correlated demand. In this example, the combined benefits from risk pooling and revenue maximization are so large with strongly negative correlations that a risk-averse newsvendor invests more money in the serial network than a risk-neutral investor (Figure 9(b)).

6.3. Operational Hedging

Serial network newsvendors purposely imbalance capacity to mitigate profit risk. A capacity portfolio $\mathbf{K}$ is balanced if all resources can be fully utilized simultaneously. In a newsvendor network, this requires the existence of an activity vector $\mathbf{x} \geq 0$ such that $\mathbf{Ax} = \mathbf{K}$ (Van Mieghem 2003, Definition 1). In the serial network, this condition simplifies to $K_1 + K_2 = K_3$, which is optimal with deterministic demand. With demand uncertainty, however, Harrison and Van Mieghem (1999) first showed that it is optimal for a risk-neutral newsvendor to under-invest in the shared resource ($K_3 < K_1 + K_2$). Three effects now explain such capacity imbalance. First, demand pooling is an obvious driver and suggests that relative capacity imbalance would decrease in correlation. Figure 10 confirms that suggestion in our numerical study, but also shows that capacity imbalance may remain optimal with perfect positive correlation. This highlights the revenue maximization option as the second driver (given that there is no risk-pooling benefit at $\rho = 1$). Finally, to

[10] The shared capacity is driven by a critical fractile of the sum of market profit distributions whose standard deviation is less than individual distributions with strong negative correlations. Eppen (1979) called this effect *statistical economies of scale*.

[11] This agrees with Proposition 3 in Van Mieghem and Rudi (2002), which states that the mean value of any newsvendor network whose operating profit is submodular in $\mathbf{D}$ (which is the case for the serial network) is decreasing in any correlation coefficient. Figure 9(a) also illustrates that even a risk-neutral newsvendor strictly prefers more negatively correlated demand.

Figure 9 **Serial Network Newsvendors Prefer Smaller Correlations (Frontiers Move Northwest). Risk Aversion May Increase Spending Above Risk-Neutral Levels with Low Correlation** ($\sigma_1 = 0.3, \sigma_2 = 0.4$)

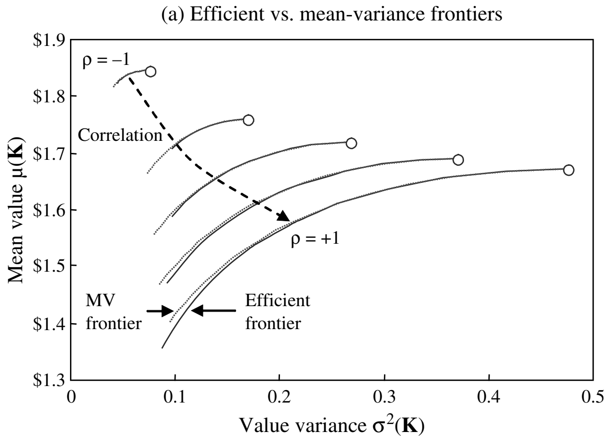

exercise its switching option, the shared resource must have upstream capacity leeway.

Relative capacity imbalance decreases in correlation, but is amplified by risk aversion, especially at correlation $\rho = 1$ according to the numerical study in Figure 10. In the absence of risk pooling, this reflects the revenue maximization option. While all three resource levels decrease in risk for $\rho = 1$, the reduction in the shared resource exceeds that in the dedicated resources, thereby increasing the relative potential of the switching option. (The smaller curvature of the frontiers in Figure 9(a) suggests that capacity imbalance may be more effective at higher correlations. Numerical analysis also shows that capacity imbalance increases in the profit differential $v_1 - v_2$, but disappears at $\rho = 1$ if $v_1 = v_2$.)

Given that the optimal risk-averse resource vector seeks to reduce profit variance, this means that more resource imbalance reduces more profit risk. Resource imbalance thus truly is an operational hedge.

6.4. Generalizations

The insights from the numerical study hold for general demand distributions and utility functions under certain conditions:

Property 4 (Serial Network). *$K_2^u \geq K_2^n$ and increases in risk aversion with correlation $\rho = -1$, $v_2\sigma_2 < v_1\sigma_1$, and moderate costs and risk aversion.*

(The proof details the cost conditions which can be expressed in terms of the safety resource levels and thus c/v fractions.) The driving force behind the increase of K_2 again is the reduction of profit variance:

Property 5 (Serial Network). *Assume correlation $\rho = -1$ and moderate costs and risk aversion. If*

Figure 10 **Operational Hedging in the Serial Network Rebalances Capacity Away from the Shared Resource. Imbalance Decreases in Correlation but Increases in Risk Aversion**

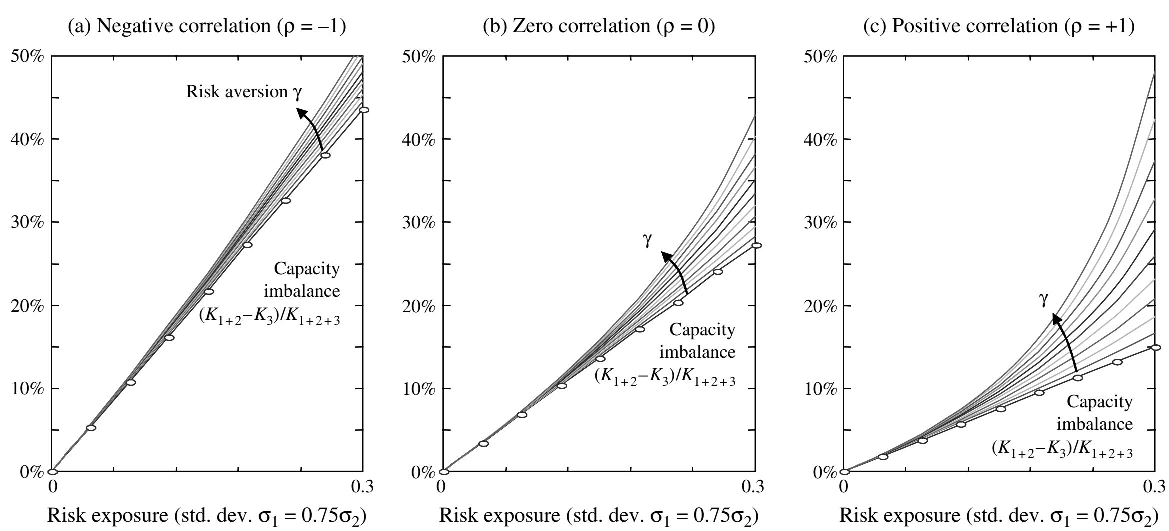

Figure 11 With Normally Distributed Demand, the Risk-Neutral Parallel Network also Adjusts Capacities Linearly (Labeled "o") in Demand Standard Deviation with Slope Dependent on Correlation. Risk-Averse Levels Are Roughly Quadratic. Flexibility Always Increases in Risk Aversion Reflecting Strong Risk Pooling and Revenue Maximization Benefits and Substitution

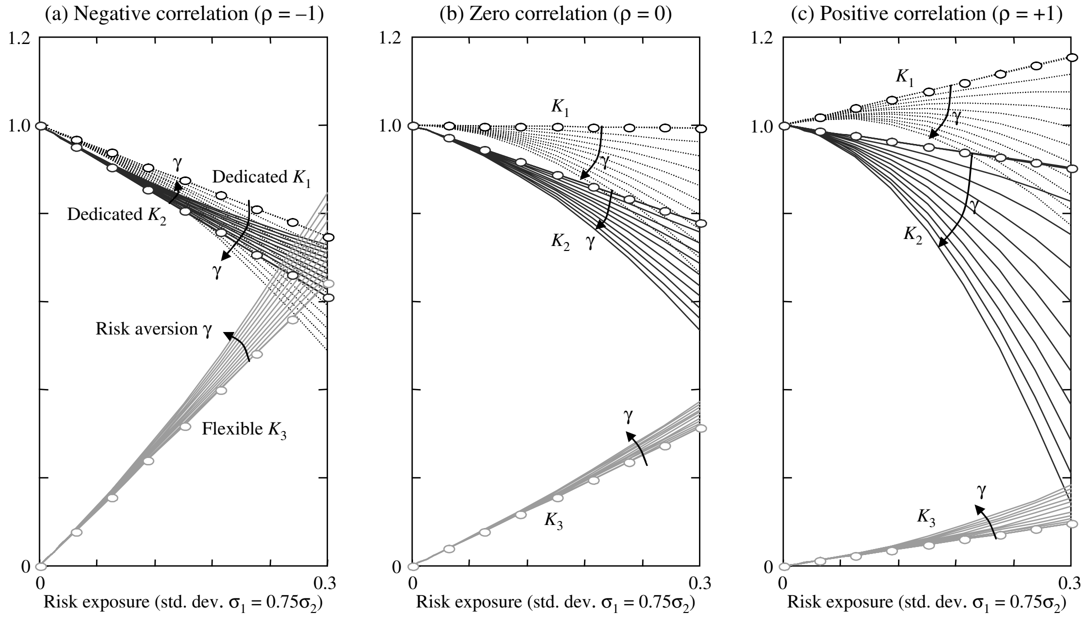

$v_2\sigma_2 < v_1\sigma_1$, *then* $\nabla_2\sigma^2(\mathbf{K}^n) \le 0$ *and* $(d/d\gamma)K_2(0) \ge 0$. *If* $v_1 = v_2$, *then* $\nabla_3\sigma^2(\mathbf{K}^n) \ge 0$.

The impact of risk aversion on the shared resource level is parameter dependent. Figure 8 shows that the shared resource can increase in risk aversion with strongly negative correlations, reflecting the complementarity in the network. The diversification benefit drives more risk-averse agents to increase K_2, and a sufficiently strong increase induces an increase in the shared resource to alleviate its potential of being a bottleneck. With higher correlations, the diversification benefit weakens as does its complementarity on K_3, which becomes decreasing in risk aversion. Aside from this complementarity, the shared resource level is also driven by the revenue maximization option: With $v_1 = 2v_2$, its value at $\rho = -1$ was sufficient for K_3 to increase in risk aversion in our numerical results, while K_3 decreases in risk aversion without revenue maximization (i.e., if $v_1 = v_2$) according to Property 5.

7. Resource Flexibility and Substitution: The Parallel Network

The parallel network can be viewed as a dedicated network augmented with the option to use a third flexible resource which can only improve upon the dedicated network's performance. Allocation flexibility in the parallel network again provides an active hedge, but its dependence differs from that in the serial network. In the latter, the shared resource was a potential bottleneck and a complement of the upstream dedicated resources, but its capacity was underweight to enable its switching option. Here, the flexible resource adds redundancy and is a substitute of the dedicated resources. The question in the single-stage parallel network is whether the substitutable resource 3 is a viable alternative and what its resource level should be. The substitution tension depends on the relative cost of resource 3, and the natural and simplest assumptions for the parallel network are $\max(c_1, c_2) < c_3 < c_1 + c_2$ and $c_i \le v_i$.

Again, we first make observations from a numerical study and then generalize. The numerical data is the same as before, but with more expensive flexible capacity costs $c_3 = 0.7 > c_1 = c_2 = 0.5$.

7.1. Risk Exposure

Again, risk-neutral capacities with normal demand remain linear in (marginal) demand standard deviations with slopes depending on demand correlation as well as on the value-cost ratios c/v. The "o"-connected lines in Figure 11 illustrate these facts through a numerical study where both risk-neutral dedicated resource levels increase in correlation, while the risk-neutral flexible resource level decreases. The latter stems from decreasing risk-pooling benefits, and the substitution effect explains the former. Interestingly, flexibility is used even at perfect positive correlation; while there is no risk pooling, the revenue maximization option of flexibility remains.

7.2. Risk Aversion

The numerical study suggests that the CARA utility dedicated capacities with normal demand again mirror those in the dedicated network. In Figure 11, both

Figure 12 **Parallel Network Newsvendors Prefer Smaller Correlations (Frontiers Move Northwest). Risk Aversion May Increase Spending Above Risk-Neutral Levels with Low Correlation** ($\sigma_1 = 0.3, \sigma_2 = 0.4$)

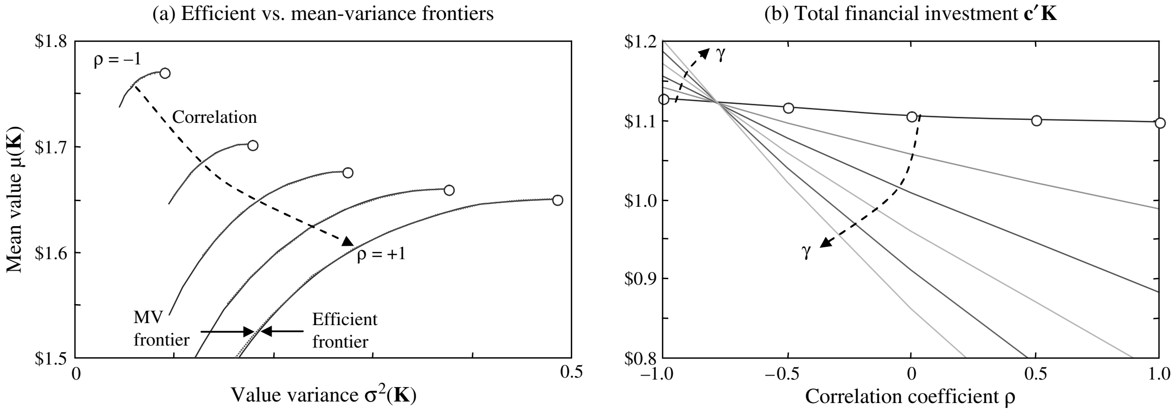

appear roughly quadratic in standard deviation and K_2^u increases in risk aversion for strongly negatively correlated demand. Interestingly, more risk-averse newsvendors use more flexibility for any correlation in this example, i.e., the result of Eeckhoudt et al. (1995) does not extend to serial networks (even not at $\rho = 1$). With strongly negative correlations, its benefit is so large that it is optimal to "oversubstitute" dedicated capacity with flexible at higher total spending than under risk neutrality (Figure 12(b)). Figure 12(a) shows that even a risk-neutral newsvendor again strictly prefers more negatively correlated demand and that the parallel network dominates the dedicated network (compare with Figure 6(c)).

7.3. Operational Hedging

First, compared to the dedicated network, the flexible resource provides additional risk mitigation and value creation: the frontiers of Figure 12 lie northwest relative to those of Figure 6(c). Second, parallel network newsvendors move the investment mix toward the flexible asset to mitigate profit risk. The flexible capacity share decreases in correlation (reflecting decreased risk pooling benefits), but does not disappear at $\rho = 1$ (reflecting revenue maximization benefits) in the numerical study of Figure 13. More importantly, the flexible share increases in risk aversion, especially at correlation $\rho = 1$. (The dedicated resources K_1 and K_2 are strongly decreasing in risk at $\rho = 1$, thereby freeing up funds that are partially invested to increase the flexible resource level. Numerical analysis shows that flexibility increases in the profit differential $v_1 - v_2$, but disappears at $\rho = 1$ if $v_1 = v_2$.)

Flexibility has been interpreted as an operational hedge in risk-neutral models. This paper shows that the capacity mix moves even more toward flexibility

Figure 13 **Operational Hedging in the Parallel Network Rebalances Capacity Toward the Flexible Resource. The Flexible Share Increases in Risk Aversion, but Decreases in Correlation**

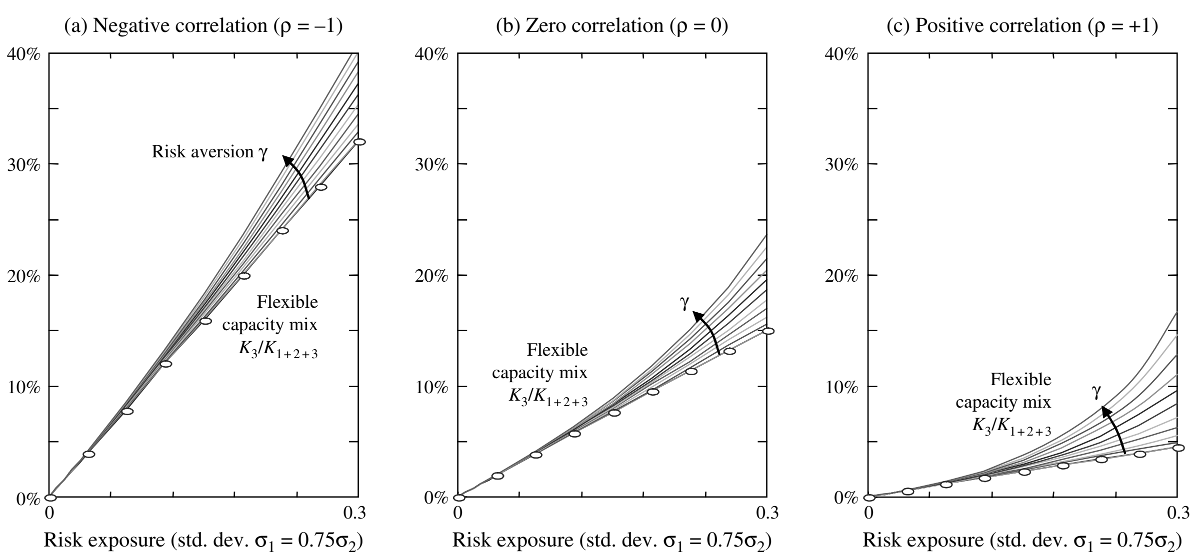

as risk aversion increases. Given that the optimal risk-averse resource vector seeks to reduce (variance) risk, this means that more flexibility reduces more profit variability, reinforcing its interpretation as a hedge.

7.4. Generalizations

The insights from the numerical study can again be generalized to other demand distributions and utility functions (but somewhat less so than earlier). Assuming that a rational risk-neutral agent invests in flexibility (see Van Mieghem 1998, Proposition 7 for conditions), risk aversion moves the capacity mix toward the lower-profit variance resource 2 and the flexible resource 3:

Property 6 (Parallel Network). *Assume correlation $\rho = -1$ and moderate risk aversion and flexible cost so that $K_3^n > 0$. If $\sigma_1 = \sigma_2$ and $v_1 > v_2$, then $K_2^u + K_3^u \geq K_2^n + K_3^n$ and increases in risk aversion, while $K_1^u \leq K_1^n$ and decreases in risk aversion.*

8. Conclusion, Limitations, and Extensions

This paper has shown how resource allocation in networks can mitigate risk exposure. Risk attitude and network configuration drive the strategic placement of operational resources like safety capacity and inventory to reduce financial risk such as mismatch costs and profit variability. This joint operational and financial perspective was adopted to develop theory and insight into capacity imbalance and allocation flexibility to mitigate risk and serve as operational hedges. They may even exploit risk and create value.

Risk-averse newsvendors may increase network capacity and total spending above risk-neutral levels because rebalancing capacity may decrease profit variance. (Increasing capacity captures the multivariate demand distribution more fully and increases risk-pooling benefits, especially with negative correlations. In contrast, with access to only a single asset and a single risk, decreasing capacity is the only way to decrease profit variance.)[12] With normally distributed demand, risk-averse newsvendors change resource levels roughly proportionally to demand variance, while risk-neutral agents adjust only proportionally to standard deviation. This was explained in terms of the benefits from pure diversification, demand pooling, and ex post optimization; Table 1 provides a summary.

The managerial take-away is that risk-averse newsvendors faced with increased risk exposure should overadjust their resource portfolio relative to their risk-neutral counterparts. Sometimes they should increase capacity, but the appropriate actions depend on market profit (co)variances and network structure. The theory suggests some rules of thumb for strategic placement of safety capacity and inventory in networks: Inexpensive resources supplying the lower-profit variance market may be increased for operational hedging, especially with strong negative correlations. When markets differ in profitability, capacity may be rebalanced toward substitutable flexible resources, but away from shared resources. Given that both types of resources are "product-flexible," this means that the appropriate hedging action for product-flexible resources depends on their network position: increase when in parallel with dedicated resources, decrease when in series. (We suspect that the redundancy inherent in parallelism plays a role in this difference in response, but leave that investigation to future work.)

Some testable hypotheses of our theory include: the capacity mix of firms with a natural hedge is negatively correlated with market profit variances. Firms with redundant flexible capacity have higher value and smaller profit variability than similar dedicated firms. Resource imbalance in firms with shared capacity is positively correlated with market demand standard deviation, but decreases in market correlation.

Our analysis and insights, however, have many limitations and much remains to be done. While newsvendor networks have several advantages (see the introduction), their main disadvantage is that they may be too stylized to capture details necessary for practical decision-support systems. The single-period model abstracts from real dynamics. The recent working paper by Zhu and Kapuscinski (2006) provides a first and promising extension of the analysis to dynamic risk-averse newsvendor networks just like Bouakiz and Sobel (1992) and Chen et al. (2004) did for the single-resource model. Few organizations are controlled by a single decision maker and only allow for input inventories (yet have general resource networks). Multiagent newsvendor networks are a natural extension (e.g., Van Mieghem 1999, Goyal and Netessine 2007) as is allowing inventory stocking at multiple stages, but adding risk aversion will further complicate analysis. von Neumann and Morgenstern's celebrated utility approach to decision making under uncertainty has well-known limitations: it is not the most general or basic way to describe human behavior (e.g., see Fishburn 1982, Heyman and Sobel 1984); the axioms postulated to guarantee the existence of a utility function are often violated in practice; human behavior is far more complex than that implied by increasing concave utility functions (see a review paper by Rabin 1998, Tomlin and Wang 2005 for considering loss aversion and value-at-risk measures);

[12] The access conditions are necessary: Chod et al. (2006) and Ding et al. (2007) show that single-resource risk-averse newsvendors with access to a financial (i.e., second) asset or two risks may also increase resource capacity beyond the risk-neutral level.

and constructing a utility function or soliciting preferences in practice is a daunting task. In light of this, it would be interesting to analyze whether certain networks are more robust to parameter estimation errors than others.

This paper provides some general mathematical expressions, but only applies them in the limited setting of three networks whose properties mostly focus on the local neighborhood of the risk-neutral solution. The future task is to expand the set of networks and the generality of the analysis to the extent possible. A first and natural approach is to use the model here as a numeric optimization tool to quickly identify and compare a number of promising network configurations, following the approach of Graves and Willems (2000) and Graves and Tomlin (2003) in response to industry practitioners' need. From a research perspective, the next task is to increase structural insight for general networks, like Jordan and Graves (1995) successfully did for risk-neutral network configuration for flexibility. One would like to characterize which network structures and parameters yield effective hedging and how each resource should be adjusted as risk exposure changes. This surely is a difficult assignment that probably is best addressed piecemeal wise. Eventually, one would like to formulate rules that specify which network modules, and even complete designs, are appropriate for given environments.

9. Electronic Companion and Proofs

An electronic companion to this paper is available as part of the online version that can be found at http://mansci.journal.informs.org/.

Acknowledgments

The author is grateful for stimulating discussions with Philipp Afeche (who suggested analyzing the dedicated network as the simplest network), Jiri Chod, Marty Lariviere, Lauren Lu, Panos Kouvelis, Nils Rudi, and Costis Skiadas. The author greatly benefited from the excellent comments and questions from three peer reviewers and the associate editor, which substantially improved this paper.

Appendix. Proof of Property 2

Optimal activities are simple: $x_i(\mathbf{K},\mathbf{D})=\min(K_i,D_i)$ so that $\lambda_i(\mathbf{K},\mathbf{D})=\nabla_{K_i}\pi(\mathbf{K},\mathbf{D})=v_i 1_{\{D_i\geq K_i\}}$. The demand space can be partitioned accordingly, as shown in Figure A1. Abbreviate $\Omega_i\cup\Omega_j$ by Ω_{ij} and $P(\Omega_{ij}(\mathbf{K}))$ by $P_{ij}(\mathbf{K})$. According to Proposition 2, the optimality conditions for the risk-averse resource vector $\mathbf{K}^u$ are

$$0=(v_1-c_1)\mathbb{E}_{34}u'(V(\mathbf{K}^u,\mathbf{D})+W)\\ -c_1\mathbb{E}_{01}u'(V(\mathbf{K}^u,\mathbf{D})+W), \tag{11}$$

$$0=(v_2-c_2)\mathbb{E}_{13}u'(V(\mathbf{K}^u,\mathbf{D})+W)\\ -c_2\mathbb{E}_{04}u'(V(\mathbf{K}^u,\mathbf{D})+W), \tag{12}$$

where $\mathbb{E}_i f=\int_{\Omega_i} f\,dP$ denotes partial expectation over Ω_i.

Part 1. $\rho=+1$. Any perfectly positively correlated distribution on $\mathbb{R}^2_+$ has an upward-sloping line as support:

$$\frac{D_1-\mu_1}{\sigma_1}=\frac{D_2-\mu_2}{\sigma_2}\Leftrightarrow z_1(D_1)=z_2(D_2)\Leftrightarrow D_1=aD_2+b_+ \quad \text{with probability } 1, \tag{13}$$

using the standardized safety level notation (10) and where $a=\sigma_1/\sigma_2>0$ and $b_+=\mu_1-a\mu_2$. There are two possible cases.

Case 1. $\mathbf{K}^u$ falls below the demand line so that $P_4(\mathbf{K}^u)=0$, as in Figure A1(b). The proof uses the point $\mathbf{x}_1(\mathbf{K})=(aK_2+b_+,K_2)$ to partition and bound marginal utilities as follows. Recall that $V(\mathbf{K},\mathbf{D})=\mathbf{v}'\mathbf{x}(\mathbf{K},\mathbf{D})-C(\mathbf{K})$, where $\mathbf{x}(\mathbf{K},\mathbf{D})$ falls on the the bold line in Figure A1, which is the smaller of capacity (boundary of Ω_0) and demand. Define $V_1(\mathbf{K})=\mathbf{v}'\mathbf{x}_1(\mathbf{K})-C(\mathbf{K})$. Clearly, $\mathbf{v}'\mathbf{x}(\mathbf{K},\mathbf{D})$, and thus V is increasing in the direction of $\mathbf{v}$ so that $V(\mathbf{K},\mathbf{D}_1)\leq V_1(\mathbf{K})\leq V(\mathbf{K},\mathbf{D}_2)$ for any $\mathbf{D}_1\in\Omega_0(\mathbf{K})$ and $\mathbf{D}_2\in\Omega_{13}(\mathbf{K})$. Because u' is strictly decreasing, $u'(V(\mathbf{K},\mathbf{D}_1))\geq u'(V_1(\mathbf{K}))\geq u'(V(\mathbf{K},\mathbf{D}_2))$, so that (12) yields

$$0\leq u'(V_1(\mathbf{K}^u)+W)[(v_2-c_2)P_{13}(\mathbf{K}^u)-c_2P_0(\mathbf{K}^u)]\\ \overset{P_0=1-P_{13}}{\Rightarrow} P_{13}(\mathbf{K}^u)=P(D_2>K_2^u)\geq\frac{c_2}{v_2}.$$

Similarly, use the point $\mathbf{K}^u$ and $V_0(\mathbf{K}^u)=\mathbf{v}'\mathbf{K}^u-C(\mathbf{K}^u)$ to establish that (11) yields

$$0\leq u'(V_0(\mathbf{K}^u)+W)[(v_1-c_1)P_3(\mathbf{K}^u)-c_1P_{01}(\mathbf{K}^u)]\\ \Rightarrow P_3(\mathbf{K}^u)=P(D_1>K_1^u)\geq\frac{c_1}{v_1}.$$

$\mathbf{K}^n$ satisfies $P(D_i>K_i^n)=c_i/v_i$ so that $K_i^u\leq K_i^n$. An increase in risk aversion is equivalent to a concave increasing transformation h of the utility function u, where $h'>0$ and $h''<0$. The gradient of $U_{h\circ u}(\mathbf{K};W)=\mathbb{E}h(u(V(\mathbf{K},\omega)+W))$ at $\mathbf{K}^u$ can also be signed:

$$\begin{aligned}\nabla_1 U_{h\circ u}&(\mathbf{K}^u;W)\\ &=\mathbb{E}(\lambda_1(\mathbf{K}^u,\omega)-c_1)h'(u(V(\mathbf{K}^u,\omega)+W))u'(V(\mathbf{K}^u,\omega)+W)\\ &=(v_1-c_1)\mathbb{E}_{34}h'(u(V(\mathbf{K}^u,\mathbf{D})+W))u'(V(\mathbf{K}^u,\mathbf{D})+W)\\ &\quad-c_1\mathbb{E}_{01}h'(u(V(\mathbf{K}^u,\mathbf{D})+W))u'(V(\mathbf{K}^u,\mathbf{D})+W)\\ &\leq h'(u(V_0(\mathbf{K}^u)+W))[(v_1-c_1)\mathbb{E}_3u'(V(\mathbf{K}^u,\mathbf{D})+W)\\ &\qquad -c_1\mathbb{E}_{01}u'(V(\mathbf{K}^u,\mathbf{D})+W)]=0,\end{aligned}$$

where we used (11) and the fact that h' is decreasing. A similar argument shows that also $\nabla_2 U_{h\circ u}(\mathbf{K}^u;W)\leq 0$. This holds for arbitrary h and u, which together with concavity of $U_{h\circ u}$ means that $\mathbf{K}^{h\circ u}\leq\mathbf{K}^u$.

Case 2. $\mathbf{K}^u$ falls on or above the demand line so that $P_1(\mathbf{K}^u)=0$. Partition using the points $\mathbf{x}_2(\mathbf{K})=(K_1,a^{-1}(K_1-b_+))$ and $\mathbf{K}^u$ to establish that $P_{34}(\mathbf{K}^u)=P(D_1>K_1^u)\geq c_1/v_1$ and $P_3(\mathbf{K}^u)=P(D_2>K_2^u)\geq c_2/v_2$, respectively.

Part 2. $\rho=-1$. The demand support is now downward sloping: $z_1(D_1)+z_2(D_2)=0\Leftrightarrow D_1+aD_2=b$, where $b=\mu_1+a\mu_2>0$. In the setting of Figure A1(a), use the point $\mathbf{x}_1(\mathbf{K})=(b-aK_2,K_2)$ and scalar $k_1(\mathbf{K})=\mathbf{v}'\mathbf{x}_1/v_1$ (the horizontal intercept of the normal to $\mathbf{v}$ through $\mathbf{x}_1$) to again partition and bound marginal utilities to establish that $P_1(\mathbf{K}^u)=$

Figure A1 **Dedicated Network Activities x Depend on Resource Levels K and Demand D**

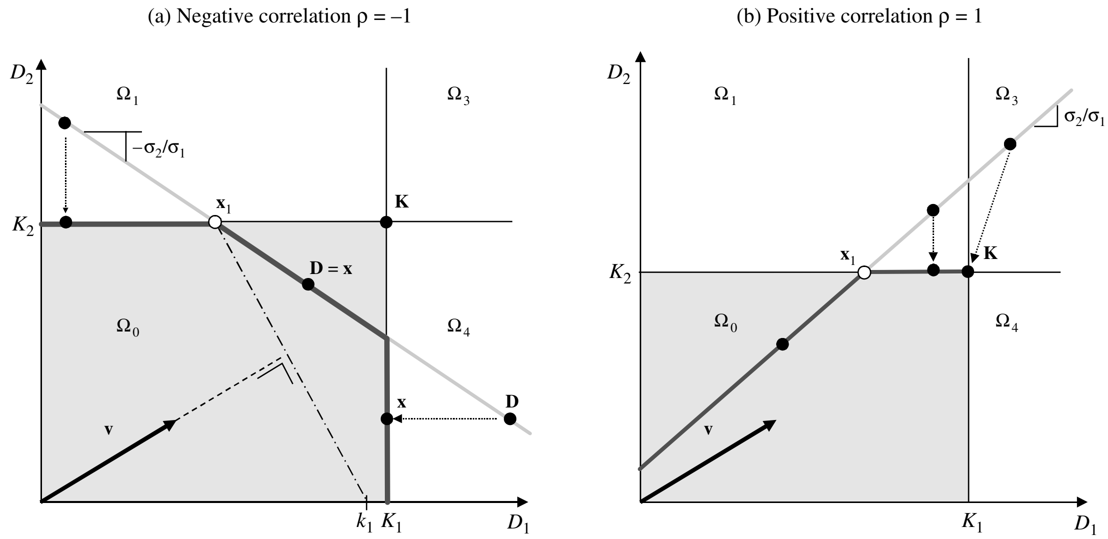

Note. Demand is on the downward-sloping line if $\rho = -1$ (a) and on the upward-sloping line if $\rho = 1$ (b).

$P(D_2 > K_2^u) \leq c_2/v_2$. In contrast to Part 1, Part 2 requires conditions: (1) v is below the demand normal or $v_2/v_1 < \sigma_1/\sigma_2$; (2) $k_1(\mathbf{K}^u) < K_1^u$ or $z_1(K_1^u) + (1 - v_2\sigma_2/v_1\sigma_1)z_2(K_2^u) > 0$; and (3) $\mathbf{K}^u$ falls above the demand line or $z_1(K_1^u) + z_2(K_2^u) > 0$. If risk aversion is moderate so that $\mathbf{K}^u$ falls in a $\mathbf{K}^n$ neighborhood, then conditions (2) and (3) evaluated at $\mathbf{K}^n$ suffice. Similarly, establish that $\nabla_2 U_{hou}(\mathbf{K}^u; W) \geq 0$. Given concavity, this means that an infinitesimal increase h in risk aversion relative to any u that satisfies the conditions leads to $K_2^{hou} > K_2^u$. □

References

Agrawal, V., S. Seshadri. 2000a. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing Service Oper. Management* **2**(4) 410–423.

Agrawal, V., S. Seshadri. 2000b. Risk intermediation in supply chains. *IIE Trans.* **32** 819–831.

Anvari, M. 1987. Optimality criteria and risk in inventory models: The case of the newsboy problem. *J. Oper. Res. Soc.* **38**(7) 625–632.

Atkinson, A. A. 1979. Incentives, uncertainty, and risk in the newsboy problem. *Decision Sci.* **10**(3) 341–355.

Bassok, Y., R. Anupindi, R. Akella. 1999. Single-period multiproduct inventory models with substitution. *Oper. Res.* **47**(4) 632–642.

Bish, E., Q. Wang. 2004. Optimal investment strategies for flexible resources, considering pricing and correlated demands. *Oper. Res.* **52**(6) 954–964.

Bouakiz, M., M. J. Sobel. 1992. Inventory control with an exponential utility criterion. *Oper. Res.* **40**(3) 603–608.

Boyabatlı, O., L. B. Toktay. 2004. Operational hedging: A review with discussion. Technical report, INSEAD, Fontainebleau, France.

Brealey, R. A., S. C. Myers. 2000. *Principles of Corporate Finance*, 6th ed. Irwin McGraw-Hill, Boston, MA.

Cachon, G. P., C. Terwiesch. 2006. *Matching Supply with Demand*. Irwin McGraw-Hill, Boston, MA.

Caldentey, R., M. Haugh. 2006. Optimal control and hedging of operations in the presence of financial markets. *Math. Oper. Res.* **31**(2) 285–304.

Chen, X., M. Sim, D. Simchi-Levi, P. Sun. 2004. Risk aversion in inventory management. *Oper. Res.* Forthcoming.

Chod, J., N. Rudi, J. A. Van Mieghem. 2006. Operational flexibility and financial hedging: Complementarity and substitution effects. Working paper, Northwestern University, Evanston, IL.

Corbett, C. J., K. Rajaram. 2006. A generalization of the inventory pooling effect to nonnormal dependent demand. *Manufacturing Service Oper. Management* **8**(4) 351–358.

Ding, Q., L. Dong, P. Kouvelis. 2007. On the integration of production and financial hedging decisions in global markets. *Oper. Res.* Forthcoming.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The risk-averse (and prudent) newsboy. *Management Sci.* **41**(5) 786–794.

Eppen, G. D. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.* **25**(5) 498–501.

Fine, C. H., R. M. Freund. 1990. Optimal investment in product-flexible manufacturing capacity. *Management Sci.* **36**(4) 449–466.

Fishburn, P. C. 1982. *The Foundations of Expected Utility*. D. Reidel Publishing Company, Dordrecht, The Netherlands.

Gan, X., S. Sethi, H. Yan. 2005. Coordination of a supply chain with risk-averse agents. *Prod. Oper. Management* **14**(1) 80–89.

Gaur, V., S. Seshadri. 2005. Hedging inventory risk through market instruments. *Manufacturing Service Oper. Management* **7**(2) 103–120.

Gollier, C. 2001. *The Economics of Risk and Time*. MIT Press, Cambridge, MA.

Goyal, M., S. Netessine. 2007. Strategic technology choice and capacity investment under demand uncertainty. *Management Sci.* **53**(2) 192–207.

Graves, S. C., B. Tomlin. 2003. Process flexibility in supply chains. *Management Sci.* **49**(7) 907–919.

Graves, S. C., S. P. Willems. 2000. Optimizing strategic safety stock placement in supply chains. *Manufacturing Service Oper. Management* **2**(1) 68–83.

Hale, W., D. F. P. Pyke, N. Rudi. 2000. An assemble-to-order system with component substitution. *Proc. 4th INFORMS M&SOM Conf., Ann Arbor, MI.*

Harrison, J. M., J. A. Van Mieghem. 1999. Multi-resource investment strategies: Operational hedging under demand uncertainty. *Eur. J. Oper. Res.* **113**(1) 17–29.

Heyman, D., M. Sobel. 1984. *Stochastic Models in Operations Research*, Vol. 2. McGraw-Hill, New York.

Huchzermeier, A., M. A. Cohen. 1996. Valuing operational flexibility under exchange rate risk. *Oper. Res.* **44**(1) 100–113.

Jordan, W. C., S. C. Graves. 1995. Principles on the benefits of manufacturing process flexibility. *Management Sci.* **41**(4) 577–594.

Lau, H. S. 1980. The newsboy problem under alternative optimization objectives. *J. Oper. Res. Soc.* **31** 525–535.

Lau, H. S., A. H. L. Lau. 1999. Manufacturer's pricing strategy and return policy for a single-period commodity. *Eur. J. Oper. Res.* **116** 291–304.

Merriam-Webster's Collegiate Dictionary. 1998. 10th ed. Merriam-Webster, Springfield, MA.

Netessine, S., G. Dobson, R. A. Shumsky. 2002. Flexible service capacity: Optimal investment and the impact of demand correlation. *Oper. Res.* **50**(2) 375–388.

Oxford English Dictionary Online. 2004. 2nd and following ed. Oxford University Press, Oxford, UK.

Pratt, J. W. 1964. Risk aversion in the small and in the large. *Econometrica* **32**(1–2) 122–136.

Rabin, M. 1998. Psychology and economics. *J. Econom. Literature* **36**(1) 11–46.

Rudi, N. 2000. Some models of risk pooling. Ph.D. dissertation, The Wharton School, University of Pennsylvania, Philadelphia, PA.

Spulber, D. F. 1985. Risk sharing and inventories. *J. Econom. Behav. Organ.* **6** 55–68.

Tomlin, B., Y. Wang. 2005. On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. *Manufacturing Service Oper. Management* **7**(1) 37–57.

Van Mieghem, J. A. 1998. Investment strategies for flexible resources. *Management Sci.* **44**(8) 1071–1078.

Van Mieghem, J. A. 1999. Coordinating investment, production and subcontracting. *Management Sci.* **45**(7) 954–971.

Van Mieghem, J. A. 2003. Capacity management, investment and hedging: Review and recent developments. *Manufacturing Service Oper. Management* **5**(4) 269–302.

Van Mieghem, J. A. 2004. Commonality strategies: Value drivers and equivalence with flexible capacity and inventory subsitution. *Management Sci.* **50**(3) 419–424.

Van Mieghem, J. A., N. Rudi. 2002. Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing Service Oper. Management* **4**(4) 313–335.

Zhu, W., R. Kapuscinski. 2006. Optimal operational versus financial hedging for a risk-averse firm. Working paper, School of Business, University of Michigan, Ann Arbor, MI.