Risk-Averse Newsvendor Networks: Resource Flexibility, Sourcing, and Hedging

Jan A. Van Mieghem

March 2004 (revised January 2006)
COSM-04-001

Working Paper Series

Center for Operations and Supply Chain Management

Kellogg School of Management

Northwestern University

Risk-averse Newsvendor Networks: Resource Flexibility, Sharing, and Hedging

Jan A. Van Mieghem

Kellogg School of Management, Northwestern University

March 25, 2004; Revised January 11, 2006

Abstract

This paper studies how differences in risk exposure and risk attitude affect network configuration. Theory is presented for general utility functions and their mean-variance approximations and illustrated with networks featuring inventory substitution and commonality, or resource sharing and flexibility.

When faced with increased risk exposure, risk-neutral newsvendors change resource levels proportionally to standard deviations, but risk-averse agents adjust proportionally to variance. While single-resource risk-averse newsvendors always invest less than risk-neutral newsvendors, they may invest more in networks. Not only may some resource levels increase in risk aversion, even the entire monetary network investment may exceed its risk-neutral counterpart. Two drivers explain this operational hedge, suggesting rules of thumb for strategic placement of safety capacity and inventory in networks. With strongly negative correlations, diversification (risk pooling) suggests to increase inexpensive resources that supply lower profit variance products. This highlights the role of profit over demand variance in risk-averse network design. With product profit differences, revenue maximization suggests to increase (decrease) parallel (serial) flexibility. These insights confirm and refine the intuitive notion of safety capacity/inventory and flexibility as an operational hedge to mitigate risk.

(Key Words: risk aversion, mean-variance, capacity, inventory, investment, real options, network design.)

1 Introduction

Designing the organization so that it will handle future uncertainty well is a universal strategic mantra. From a financial perspective, this often involves limiting the variability in payoffs with risk management instruments that provide financial hedging. From an operational perspective, handling uncertainty well invariably involves adding safety buffers. Basic risk-neutral, single-resource newsvendor and queuing models have shown the merits of some combination of safety inventory, safety capacity, or safety time. Yet in real life, many organizations use multiple assets to produce multiple products and most managers care about risk. While finance has appropriately focused on portfolios and risk, there is surprisingly little research on network operations and financial risk. This paper is a first attempt to start filling that void.

This paper studies how differences in risk exposure and risk attitude affect network configuration. It presents theory and insight on the strategic placement of safety capacity and inventory and illustrates how processing

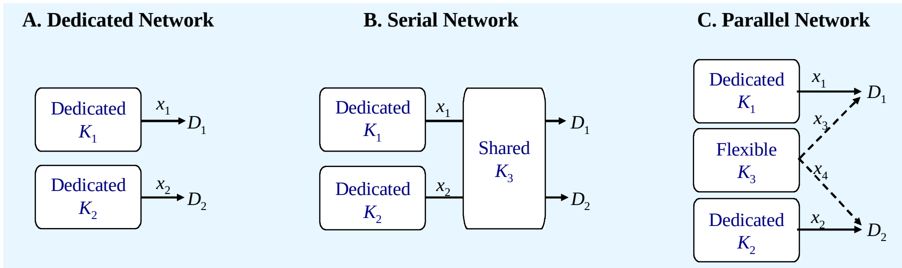

Figure 1 Three canonical building blocks for multi-resource networks.

networks can be designed as an operational hedge to mitigate risk. The networks considered here are designed and managed by a single expected utility maximizer. Design involves sizing of resources, which include inventories as well as capacities, and management refers to processing to best fill product 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 results hold for any portfolio of real options with general network topology. 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 stochastic planning. Special attention is devoted to the three canonical newsvendor networks shown in Fig. 1 because they are the building blocks for many multi-resource networks (and remain reasonably tractable). All three process two products with the dedicated "network" being the simplest possible: Each product requires a product-dedicated resource. The other two networks add a third product-flexible resource, and with it, processing interdependence. The parallel network is a dedicated network augmented with auxiliary flexible capacity. It allows a tailored response to uncertainty where dedicated capacity mainly fills base demand while flexible capacity supplies variable demand. In serial networks each product 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 perform final burn-in and test routines, or car manufacturing with a flexible assembly line. From a network design standpoint, these two networks differ in the position of the flexible resource relative to the dedicated resources. Economically, the flexible resource is either a substitutable or a complementary real asset. The parallel network is a core model for flexible technology and inventory substitution while the serial network models resource sharing and pure component commonality.

This paper establishes that risk attitude and network structure fundamentally change how organizations should be designed to handle uncertainty. First, while risk-neutral newsvendors change resource levels proportionally to standard deviations when faced with increased risk exposure (demand uncertainty), risk-averse agents adjust proportionally to variance (i.e., levels are quadratic in standard deviation). Second, while single-resource risk-averse newsvendors always invest less than risk-neutral newsvendors, they may invest more in networks. Not only may some resource levels increase in risk aversion, even the entire monetary network investment may exceed its risk-neutral counterpart.

Two drivers explain this operational hedge. The benefit from diversification (risk pooling) steers the portfolio mix towards assets supplying lower profit variance products. These need not be the lower demand variance products, which highlights the importance of profit variance to understanding risk-averse network design. Operational hedging is also manifested by capacity imbalance in the canonical three-resource networks which can remain even with perfectly positive correlations (i.e., in the absence of risk pooling). This isolates the revenue (profit) maximization option imbedded in flexible assets as the second driver of operational hedging. The appropriate hedging action depends on the network structure: It is shown that in the parallel (serial) network risk aversion increases (decreases) flexible investment relative to dedicated investments. These insights confirm and refine the intuitive notion of safety capacity/inventory and flexibility as a hedge or mitigation of risk.

The managerial take-away is that risk-averse newsvendors faced with increased risk exposure should over-adjust their resource portfolio relative to their risk-neutral counterparts. Sometimes they should increase capacity but the appropriate actions depend on product 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 that supply the lower profit variance product are candidates for hedging by increasing their levels, especially with strong negative correlations. In contrast, parallel (serial) flexibility should always be increased (decreased) with product profit differences.

This introduction finishes with a brief literature review of three research areas that 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, Gollier & Schlesinger (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. Other studies of the risk-averse single-resource newsvendor include Atkinson (1979), Lau (1980), Spulber (1985), Anvari (1987), Lau & Lau (1999), Agrawal & Seshadri (2000), Caldentey & Haugh (2003), Gan, Sethi & Yan (2005) and Gaur & Seshadri (2005) with multi-period extensions in Bouakiz & Sobel (1992) and Chen, Sim, Simchi-Levi & Sun (2004).

This article is also a natural successor to Van Mieghem & Rudi (2002), who defined and analyzed newsvendor networks under expected profit maximization. Flexibility in risk-neutral parallel network was first studied by Fine & 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, Anupindi & Akella (1999), Hale, Pyke & Rudi (2000), Rudi (2000), Netessine, Dobson & Shumsky (2002), Van Mieghem (2004), Bish & Wang (2004) and Goyal & Netessine (2005). The risk-neutral serial network was studied in Harrison & Van Mieghem (1999) and extended in Van Mieghem (2003). As far as we are aware, the only other paper on risk-averse newsvendor networks paper is Tomlin & 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.

This article also relates to the literature on operational hedging, a term first introduced by Huchzermeier & Cohen (1996). Operational hedging by means of flexibility and capacity imbalance in newsvendor networks was studied in Harrison & Van Mieghem (1999) and extended in Van Mieghem (2003). Boyabatlı & Toktay (2004) survey and critically discuss papers on operational hedging. Most of those works study exchange-rate uncertainty typically assuming expected profit maximization. Hedging obviously requires the presence of uncertainty but its standard objective is "to reduce risk, not to make money," according to Brealy & Myers (2000, p. 760). This paper shows that risk aversion magnifies these operational constructions, which establishes that they mitigate risk and strengthens their interpretation as operational hedges.

2 Model

▶ 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 non-negative resource vector $\mathbf{K} \in \mathbb{R}^n_+$ whose ith component represents the level of resource i that is 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 do not pose a problem. The research problem is to decide on the resource vector $\mathbf{K}$ given a probability distribution P on the underlying probability space $(\Omega, \mathcal{F}, P)$. 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.) 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.

► 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.

► 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 $\sigma^2(\mathbf{K}) = \mathbb{E}\left(V(\mathbf{K},\omega) - \mu(\mathbf{K})\right)^2 = \mathbb{E}\pi^2(\mathbf{K},\omega) - \left(\mathbb{E}\pi(\mathbf{K},\omega)\right)^2$. 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 $H(\mathbf{K})$.

► NEWSVENDOR NETWORKS To illustrate the general theory, we will use newsvendor networks, a tractable class of linear recourse models introduced in Van Mieghem & 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 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 $\mathbf{x}^*(\mathbf{K},\mathbf{S},\mathbf{D}) = \arg\max\left\{\mathbf{v}'\mathbf{x} : \mathbf{x} \geq 0, \mathbf{R}_{\mathrm{S}}\mathbf{x} \leq \mathbf{S}, \mathbf{R}_{\mathrm{D}}\mathbf{x} \leq \mathbf{D}, \mathbf{A}\mathbf{x} \leq \mathbf{K}\right\}$. 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 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 product-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 product 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.*

(Proofs are relegated to the Appendix.) 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*

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

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:

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

in agreement with Proposition 1 in Van Mieghem & Rudi (2002). Condition (3) is the multi-dimensional 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 $\Omega = \cup_j \Omega_j(\mathbf{K})$ such that $\mathbb{E}\boldsymbol{\lambda}(\mathbf{K}, \omega) = \sum_j \boldsymbol{\lambda}_j P\left(\Omega_j(\mathbf{K})\right)$, where $\boldsymbol{\lambda}_j$ is the constant shadow vector of the resource constraint $\mathbf{A}\mathbf{x} \leq \mathbf{K}$ in event $\Omega_j(\mathbf{K})$. The optimality condition $\sum_j \boldsymbol{\lambda}_j P\left(\Omega_j(\mathbf{K}^n)\right) = \mathbf{c}$ 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 a function 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 infinitesimal perturbation analysis (IPA). (This method was used to produce the graphs in this paper.) 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. 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 the definition of the covariance between x and y:

$$Cov(x,y) = \mathbb{E}xy' - \mathbb{E}x\mathbb{E}y'. \tag{4}$$

(Some care is needed with dimensions: If y is scalar, then $Cov(x,y)$ has the dimensions of x and is the scalar, vector, or matrix of component-wise covariances.)

Proposition 3 *The gradient of the value variance is*

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

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

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

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." 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 less 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:

$$u_{CARA}(x) = \alpha e^{-\gamma x} + \beta, \tag{5}$$

where $\alpha \leq 0$ and $\gamma \geq 0$ to ensure the utility function is increasing concave. (It is convenient to choose $\alpha = -\beta = -\gamma^{-1}$ so that $u_{CARA}(x;\gamma) \to x$ 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

$$\nabla U(\mathbf{K}) = 0 \Leftrightarrow \mathbb{E}\left[(\boldsymbol{\lambda}(\mathbf{K},\omega) - \mathbf{c}(\mathbf{K}))\exp\left(-\gamma V(\mathbf{K},\omega)\right)\right] = 0. \tag{6}$$

Therefore, with CARA utility, the optimal resource vector is a function of the scalar γ only, which we will denote by $\mathbf{K}^{CARA}(\gamma)$. Similarly, let $U^{CARA}(\gamma)$ denote the associated utility $\mathbb{E}u_{CARA}(\mathbf{K}^{CARA}(\gamma))$. Third, exponential utility functions can be interpreted as local second-order approximations of general risk-averse utility functions, an interpretation that will be useful for the mean-variance approximation in the next subsection. Also, with CARA, the expected utility function is directly expressed in terms of the characteristic function of the value function, which is defined as $\phi_{V(\mathbf{K},\omega)}(t) = \mathbb{E}\exp(itV(\mathbf{K},\omega))$. There are simple closed-form expressions if π has a normal (see below), exponential, Gamma, or uniform distribution. It directly follows that:

$$U_{CARA}(\mathbf{K}) = \alpha e^{-\gamma W}\phi_{V(\mathbf{K},\omega)}(i\gamma) + \beta. \tag{7}$$

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* $\mathbf{K}^{CARA}(\gamma) = \mathbf{K}^n + \gamma\frac{d}{d\gamma}\mathbf{K}(0) + o(\gamma)$*, where the optimal resource adjustment vector is:*

$$\frac{d}{d\gamma}\mathbf{K}(0) = \tfrac{1}{2}H^{-1}(\mathbf{K}^n)\nabla\sigma^2(\mathbf{K}^n). \tag{8}$$

While it is known that the value variance should be decreasing with risk aversion around the risk-neutral solution (see 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 (it really is a linear combination) in the negative direction of $\nabla\sigma^2(\mathbf{K}^n)$. 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 unambiguous result does not extend to a bundle of resources. Proposition 4 shows that *the sign of the resource adjustment is typically opposite to the sign of* $\nabla\sigma^2(\mathbf{K}^n)$ in the neighborhood of the risk-neutral solution. 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 Efficient Frontier and the Mean-Variance Formulation

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

$$U_{MV}(\mathbf{K}) = \mu(\mathbf{K}) - \frac{\gamma}{2}\sigma^2(\mathbf{K}). \tag{9}$$

Let $\mathbf{K}^{MV}(\gamma)$ denote a maximizer of (9) and $U^{MV}(\gamma)$ the associated MV-utility $U_{MV}(\mathbf{K}^{MV}(\gamma))$. Expressions (2) and (3) directly yield an expression for the gradient

$$\nabla U_{MV}(\mathbf{K}) = \mathbb{E}\left[\boldsymbol{\lambda}(\mathbf{K},\omega)\right] - \mathbf{c} - \gamma Cov(\boldsymbol{\lambda}(\mathbf{K},\omega), \pi(\mathbf{K},\omega)), \tag{10}$$

which is useful for gradient-based numerical optimization via simulation. It is well-known that the mean-variance solution exhibits smaller variance than the risk-neutral solution.[1] Proposition 4 shows that risk aversion adjusts resource levels in the direction that reduces variance, which also follows from (10).

Under the assumptions of perfect capital markets, constant absolute risk aversion and normally-distributed operating profits, the von Neumann-Morgenstern framework of maximizing utility $U(\mathbf{K})$ is equivalent to maximizing $U_{MV}(\mathbf{K})$. Indeed, if V is normally distributed, its characteristic function $\phi_\pi(t) = \exp(\mu(\mathbf{K})it - \frac{1}{2}\sigma^2(\mathbf{K})t^2)$, so that (7) yields:

$$U(\mathbf{K}) = \alpha e^{-\gamma W}\phi_{V(\mathbf{K},\omega)}(i\gamma) + \beta = \alpha \exp\left(-\gamma W - \gamma\mu(\mathbf{K}) + \frac{\gamma^2}{2}\sigma^2(\mathbf{K})\right) + \beta.$$

For financial portfolios, the profit function π (and thus also $V = \pi - cK$) is linear in the resource vector and in the random vector that models uncertainty: $\pi = \widetilde{\mathbf{r}}'\mathbf{K}$, where $\widetilde{\mathbf{r}}$ is the vector of random returns. Normally-distributed returns $\widetilde{\mathbf{r}}$ and CARA preferences thus yield the celebrated mean-variance formulation (9) of financial-portfolio theory and elegant closed-form solutions, as reviewed by its developer in Markowitz (1991).[2] Unfortunately, for portfolios of real options the profit function is typically non-linear. Indeed, for newsvendor networks, the profit function is piece-wise linear (similar to call options): $\pi(\mathbf{K}, D) = \mathbf{v}'\mathbf{x}(K, D) \leq \mathbf{v}'\mathbf{D}$, so that the distribution of π is a truncated demand distribution which cannot be normally distributed (except if capacity is abundant which is economically suboptimal).

While a mean-variance objective for newsvendor networks is generally not justifiable from an economic theory perspective, it has several merits from an applied perspective. It is perhaps best viewed as an approximating objective for a general utility function u. It can be interpreted as a second order expansion around the risk-neutral solution. The following sections will confirm that the MV approximation is often

[1] By definition, $U_{MV}(\mathbf{K}^{MV}) \geq U_{MV}(\mathbf{K}^n)$ and $\mu(\mathbf{K}^n) \geq \mu(\mathbf{K}^{MV})$ so that $\sigma^2(\mathbf{K}^{MV}) \leq \sigma^2(\mathbf{K}^n)$.

[2] A mean-variance formulation is also obtained if preferences are quadratic, $u(x) = \alpha(x-\beta)^2$, and the value function $V(\mathbf{K}^n, \omega)$ is bounded a.s. by β so that the expected utility is concave increasing.

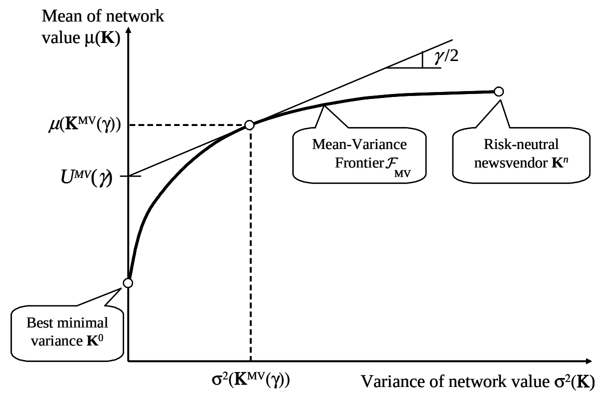

Figure 2 The mean-variance frontier $\mathcal{F}_{MV}$ and its correspondence to the MV-optimal investment vector and utility $\mathbf{K}^{MV}(\gamma)$ and $U^{MV}(\gamma)$ for a fixed coefficient of risk aversion γ.

good, if not excellent, for newsvendor networks with moderate risk aversion. (Other sufficient conditions are that $u''' \ll u''$ or the third moment $\mathbb{E}V^3 \ll \mathbb{E}V^2$.) Thus, at least in an approximate sense, the benefits of financial mean-variance formulations extend to newsvendor network MV approximations as discussed in Van Mieghem (2003):

1. they are *implementable* (i.e., only two moments are required, which can be estimated) and
2. *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 $\mathcal{F}_{MV}$, which is the set of MV-efficient investments. (An investment $\mathbf{K}$ is MV-efficient if it is not dominated, i.e., if there does not exist a resource investment $\mathbf{K}'$ such that either $\mu(\mathbf{K}') > \mu(\mathbf{K})$ while $\sigma^2(\mathbf{K}') = \sigma^2(\mathbf{K})$, or $\sigma^2(\mathbf{K}') < \sigma^2(\mathbf{K})$ while $\mu(\mathbf{K}') = \mu(\mathbf{K})$.) The MV-efficient investments and frontier thus embody the trade-off between expected return and risk. To conclude the investment problem, then all that is needed is a reasonable estimate of the decision-maker's risk attitude γ, which directly defines the optimal point on the frontier and the associated MV-optimal investment as shown in Figure 2. In practice, the decision maker can be presented with (a subset of) the MV investment path $\mathbf{K}^{MV}(\gamma)$, starting from the risk-neutral solution $\mathbf{K}^n$ and moving to the lowest risk efficient plan by tightening the coefficient of risk aversion γ.

3.4 Risk Exposure, Risk Aversion, and Hedging

Before analyzing the impact of uncertainty on specific network designs it is useful to distinguish between risk exposure, risk aversion, and hedging. Defining these concepts in full generality is difficult and beyond our scope; the following explanations are adequate for our purposes. Risk exposure refers to the consequence of a random prospect. Risk aversion is a measure that evaluates how the decision-maker perceives a risk exposure. Hedging is the action of a decision maker to mitigate a particular risk exposure.

In newsvendor models, the random prospect is typically called demand and is modeled by a probability distribution that is exogenous and a characteristic of the environment. The operational consequence of random demand is a likely but undesirable mismatch between supply and demand (overage or underage); the financial consequence is random profits or value. Comparing risk exposure thus ultimately boils down to comparing profit probability distributions. Most of the operations management literature considers a random variable "less risky" than another if its variance is smaller while both have equal means. We will follow this tradition of equating an increase in risk with an increased mean-preserving spread. (Restricting attention to variance comparisons, however, is not sufficient or even appropriate in general; see Eeckhoudt et al. (1995) or Gollier (2001) for the definition of an increase in risk.)

Risk exposure in the single-resource newsvendor is fully explained by univariate demand, which explains why most assessment of risk in the operations management literature focuses on demand uncertainty. This focus on demand uncertainty could extend to multivariate demand; a mean-preserving increase in risk for normally-distributed demand vectors is achieved by scaling up (one or several) variances. This paper, however, will demonstrate that one must focus on the consequence of the random prospect, profit uncertainty, to understand the impact of uncertainty in a network. Profit uncertainty summarizes product and network interactions and is the bottom line that matters to the decision maker. The remainder will show that any risk exposure conditions are specified in terms of profit (as opposed to purely demand) variances. Profit uncertainty also directly captures the important impact of demand correlations, an increase of which typically increases profit risk exposure.[3]

Our definition of hedging is deliberately general to include univariate settings and the actions of risk-neutral agents. It interprets as hedging the keeping of excess assets (stock or capacity) by a risk-neutral decision maker to mitigate her risk exposure in terms of expected profits lost to shortages or overages. Typically, however, hedging presumes access to at least two risks whose counterbalancing effect is to reduce total risk. Such "betting on two horses" or the interpretation of capacity imbalance in the serial network (see later) can increase expected value and thus appeals to a risk-neutral decision maker. Yet, the typical motivation for hedging is risk mitigation.[4] This paper will demonstrate that risk aversion magnifies these actions.

A natural objective of hedging is to reduce risk without giving up significant mean value. The mean-variance

[3] Using copulas, Corbett & Rajaram (2005) present evidence that most intuition from multivariate normals extend to non-normal distributions. For example, for broad classes of distributions, aggregation or pooling of inventories is more valuable as demands are less positively dependent.

[4] 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." In financial terms, hedging aims to secure or limit a risk. The standard financial textbook by Brealy & Myers (2000, pp. 763 and 760) describes hedging as "taking on one risk to offset another" and states that "most business insure or hedge to reduce risk, not to make money."

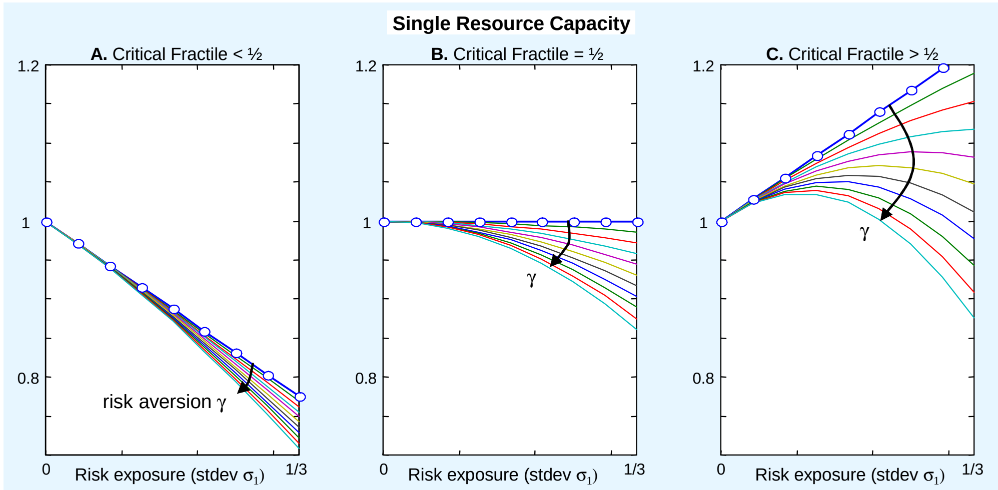

Figure 3 Single-resource newsvendor capacity as a function of risk exposure and risk aversion.

formulation approximation suggests an intuitive measure of the effectiveness of operational hedging. It is natural to consider how much risk can be reduced by adjusting the resource vector from the optimal risk-neutral portfolio $\mathbf{K}^n$, while giving up only a small amount of expected return. This amounts to asking how sensitive utility is to an increase in risk aversion. In mean-variance terms: How much can variance be reduced without greatly impacting mean value? The sensitivity of expected return to optimal reduction of variance risk around $\mathbf{K}^n$ is measured by the curvature of the MV-frontier, as shown in the Appendix:

$$\mathcal{F}''_{MV}(\sigma^2(\mathbf{K}^n)) = \frac{1}{\nabla'\sigma^2(\mathbf{K}^n)H^{-1}(\mathbf{K}^n)\nabla\sigma^2(\mathbf{K}^n)}. \tag{11}$$

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

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: $v_1 \geq c_1$.

Risk exposure. It is well-known that 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 $1 - c_1/v_1$, which depends solely on the ratio of unit cost to unit return and is also known as the critical fractile. It directly follows that K_1^n is below, at, or above the demand median if the critical fractile is respectively below, equal to, or above $1/2$. The impact of an increased risk exposure in the sense of an increasing mean-preserving spread also depends solely on the same critical fraction conditions: It leads to a decrease, no change, or increase of the resource level if the critical fractile is, respectively, below, equal to, or above 1/2. 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 $\Phi(z_1^n) = 1 - c_1/v_1$ as Fig. 3 illustrates for $\mu_1 = 1$. (The risk-neutral case is denoted by 'o' which is mnemonic for $\gamma = 0$.)

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.[5] (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 level is ambiguous, as Fig. 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 (8) for the single-resource newsvendor problem with normally distributed demand and CARA utility is*

$$\frac{d}{d\gamma}K(0) = -\sigma_1^2 c \left(1 + \frac{z^n \left(1 - \frac{c}{v}\right)}{\phi(z^n)}\right) \leq 0 \qquad \textit{where } z^n = \Phi^{-1}(1 - \frac{c}{v}).$$

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 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 proportional to demand standard deviation. This quadratic dependence not only explains the ambiguity of the impact of increased risk exposure on risk-averse resource levels for a single-resource newsvendor, but also 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. It depends on all model parameters: demand uncertainty (mean μ_1 and std. deviation σ_1) and financials v and c. The independence of mean demand is unique to CARA, all other utility functions have wealth dependence and thus also dependence on average demand[6]. 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 powerful hedging. According to (11), this

[5] The celebrated two-fund separation result of financial economics has established that the financial investment in the risky asset $C(K^u)$ decreases in risk aversion for the standard one-safe, one-risky financial asset portfolio model. But two-fund separation does not apply here because our assets have non-linear payoffs. (Indeed, its key insight that the optimal relative configuration of risky assets is independent of an investor's coefficient of risk aversion does not hold here.) Nevertheless, Eeckhoudt et al. (1995) showed that the investment in the risky asset does always decrease in risk aversion for the single-resource newsvendor case.

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

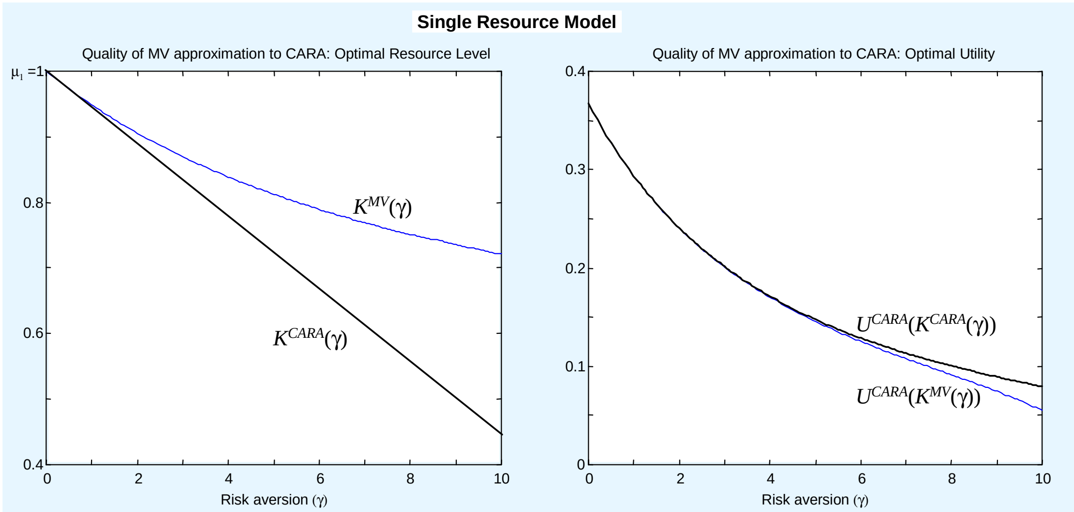

Figure 4 Optimal and mean-variance capacity and utilities for the single-resource newsvendor as a function of risk aversion. (Normal demand $\mu_1 = 1, \sigma_1 = 1/3, c = .5, v = 1.$)

amounts to identifying the parameters that yield a modest 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.

Figure 4 shows that optimal risk adjustment is near-linear in risk aversion, an effect that also extends to the three canonical networks. (Numerical work suggest that $K^{CARA}(\gamma)$ is convex decreasing if $c/v < 1/2$, linear if $c/v = 1/2$, and concave decreasing otherwise, but with very weak non-linearity.) The near-linearity suggests that the risk-adjustment gradient is almost constant, which broadens the applicability of the local risk-adjustment result of Proposition 4.

Figure 4 also allows some observations regarding the accuracy of the MV-approximation to CARA utility for the familiar single-resource newsvendor. With normal demand, both the CARA and MV objective functions can be expressed analytically so that only numerical optimization is needed, but no simulation. Figure 4 compares $\mathbf{K}^{CARA}(\gamma)$ with $\mathbf{K}^{MV}(\gamma)$ and the associated utility levels $U^{CARA}(\mathbf{K}^{CARA}(\gamma))$ with $U^{CARA}(\mathbf{K}^{MV}(\gamma))$ for normal demand with unit mean and $\sigma_1 = 1/3$. Let $v = 1$ and set $c = 1/2$ so that $K^n = 1$. Clearly, $K^{MV}(\gamma)$ is a continuous function of γ that is decreasing at small γ (because there $K^{MV}(\gamma) \simeq K^{CARA}(\gamma)$). Typically, $K^{MV}(\gamma)$ is also monotone so that the (utility-optimal) efficient frontier and the MV-frontier always coincide for single-resource problems. Figure 4 is representative of the typical situation: The MV resource levels are convex decreasing in γ, whereas the optimal levels are approximately linearly decreasing. Typically, the optimal adjustment is larger than what MV suggests.

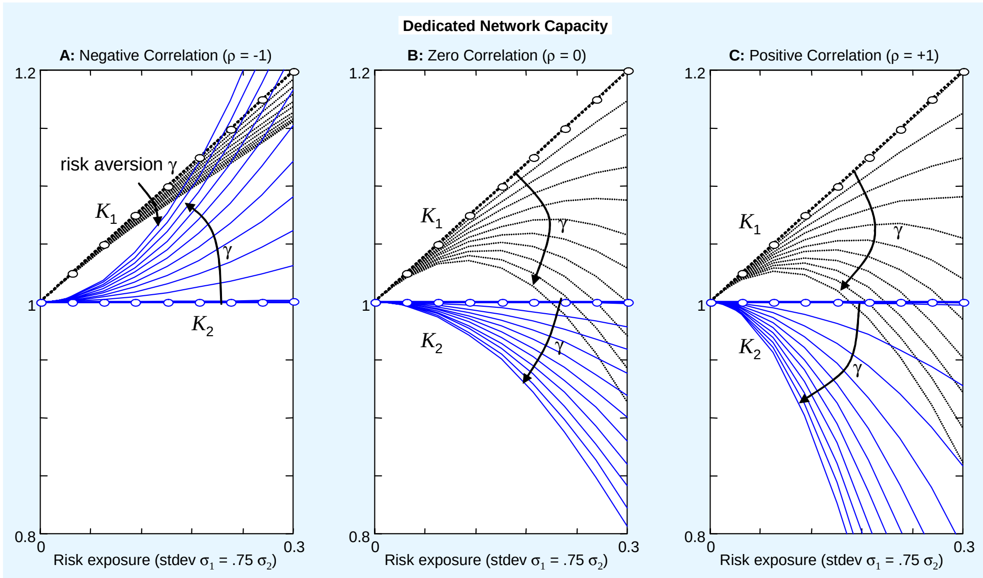

Figure 5 Dedicated network capacity as a function of risk exposure, risk aversion, and correlation.

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 (also called safety capacity or safety inventory depending on the resource type). For a risk-neutral newsvendor 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. For a risk-averse newsvendor a change in risk exposure has a magnified impact and leads to a safety level adjustment that 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:

$$\text{standardized safety level } z_i(x) = \frac{x - \mu_i}{\sigma_i}. \tag{12}$$

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 the total investment may exceed their risk-neutral counterparts. We first analyze the simplest network (the dedicated network) before addressing the serial and parallel network.

5 The Dedicated Network

Before highlighting the differences, we first discuss the insights from the single-resource newsvendor that carry over to the dedicated network, which really is just two operationally uncoupled resources. To ensure both investments are profitable, both product values should exceed their costs: $v_i \geq c_i$.

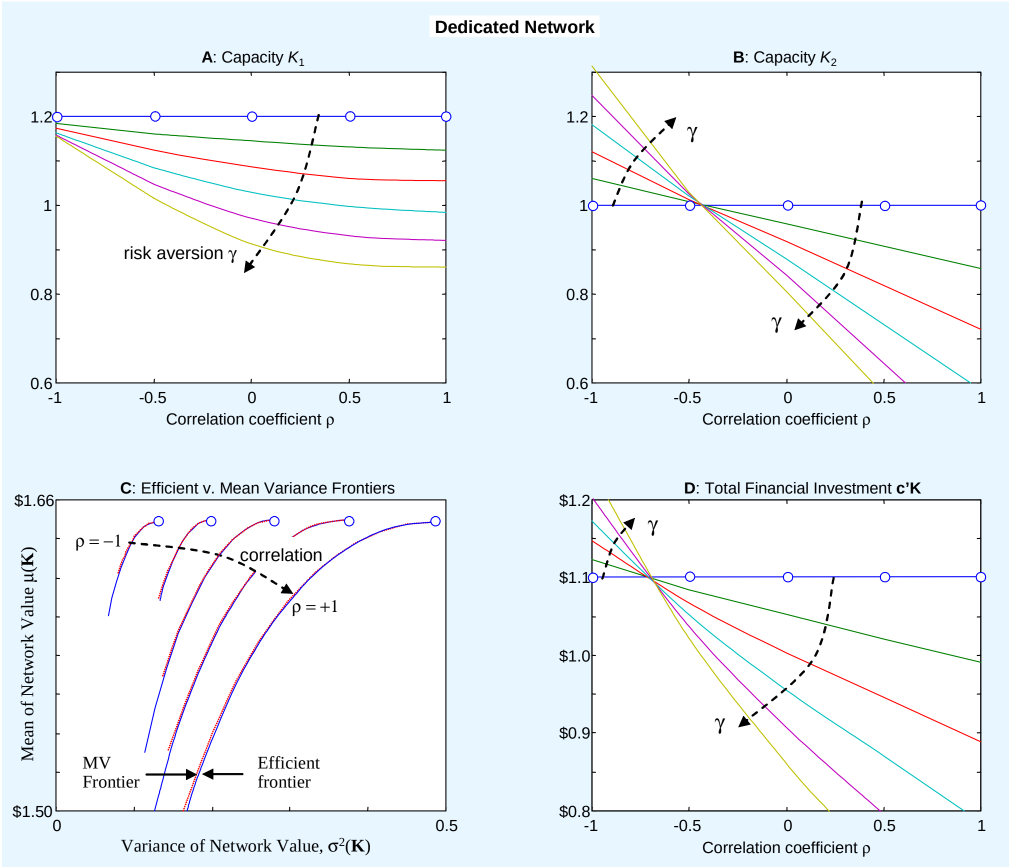

Figure 6 Capacities, frontiers and monetary investment in the dedicated network as a function of correlation and risk aversion. $(\sigma_1 = .3, \sigma_2 = .4)$.

Risk exposure. For a risk-neutral newsvendor, the dedicated network decomposes into two independent single-resource problems and capacities are linear in standard deviations with a slope determined by the critical fractile but independent of demand correlation. This is illustrated by the 'o'-connected lines in Fig. 5 for bivariate normal demand with unit means $[1, 1]$ and $\sigma_1 = .75\sigma_2$, equal unit investment costs $c = [.5, .5]$, 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.

Risk aversion. As for a single-resource newsvendor, Fig. 5 suggests that CARA-utility-optimal capacities are decreasing in risk aversion and quadratic in standard deviation *for non-negatively correlated demand*. The striking difference is that K_2^u *increases* in risk aversion and standard deviation for negative correlation. Consider Fig. 6 to further investigate and explain the crucial role of correlation (for the same data as Fig. 5). Panel B shows that K_2^u increases in risk aversion for any correlation coefficient ρ below about $-.45$. Furthermore, for ρ below about $-.7$, its increase outweighs the decrease of K_1^u so that the total financial investment for a risk-averse newsvendor *exceeds* the risk-neutral investment (Panel D). Finally, Panel C

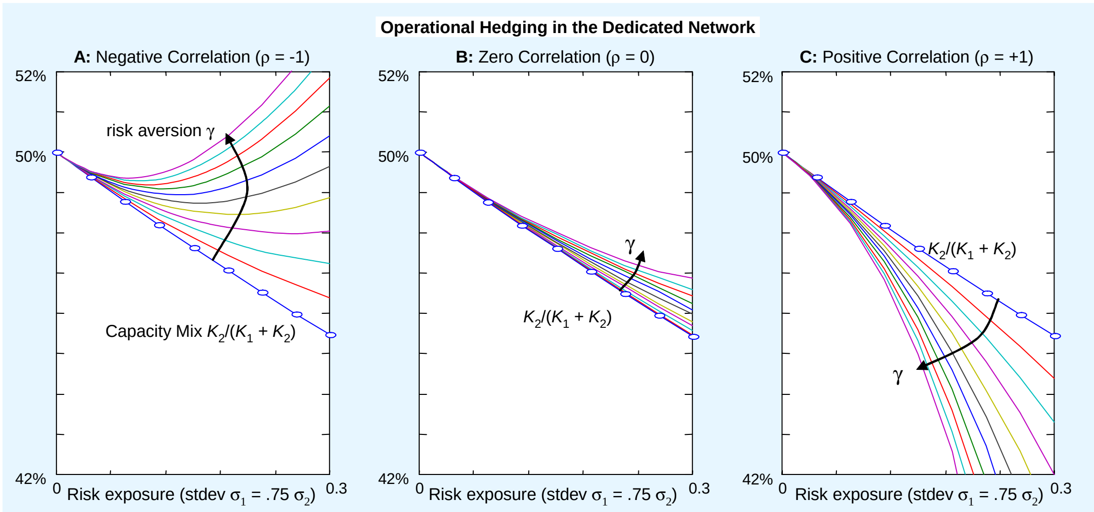

Figure 7 Operational hedging in the dedicated network is manifested by capacity rebalancing towards the minimal profit variance market if diversification benefits outweigh costs (panels A and B).

shows that a risk-averse newsvendor strictly prefers more negatively correlated demand. Put another way: Correlation is another measure of risk exposure. (Panel C also shows that the mean variance frontiers dominate (by definition!) yet closely approximate the (CARA)-efficient frontier.) These effects do not exist in the single resource case and can be explained in terms of risk mitigation or hedging.

Operational hedging. Financial risk aversion couples the resources into a diversified portfolio which can be configured in the same manner employed by financial investors (re)balancing portfolios to optimize their mean-variance utilities. Increased risk aversion increases *profit* variance reduction at the expense of some mean value. Abstracting from capacity constraints, market i has profit risk (st. dev.) $v_i\sigma_i$. If we were to invest one dollar with weights w_i to market i, 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. This risk reduction from diversification is well known and also called risk pooling. It has no impact at $\rho = 1$ but its benefit increases with smaller correlation and more equal $v_i\sigma_i$. (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 suggests that increased risk aversion will favor the market will lower profit standard deviation. This indeed is market 2 in our example ($v_1\sigma_1 = 1.5v_2\sigma_2$), which demonstrates the importance of profit over demand standard deviations (recall that market 1 is more profitable and has less demand risk: $\sigma_1 = .75\sigma_2$). With capacity constraints, the diversification benefit must be balanced with its cost in terms of investment costs and changes in expected profits. With sufficiently negative correlation the benefit can outweigh the cost so much that it even warrants a total investment above the risk-neutral investment! In the example, the benefit of favoring the lower profit variance market 2 continues to outweigh its cost even with uncorrelated demand as shown in Fig. 7. However, the variance reduction from only changing capacity mix then insuffi-

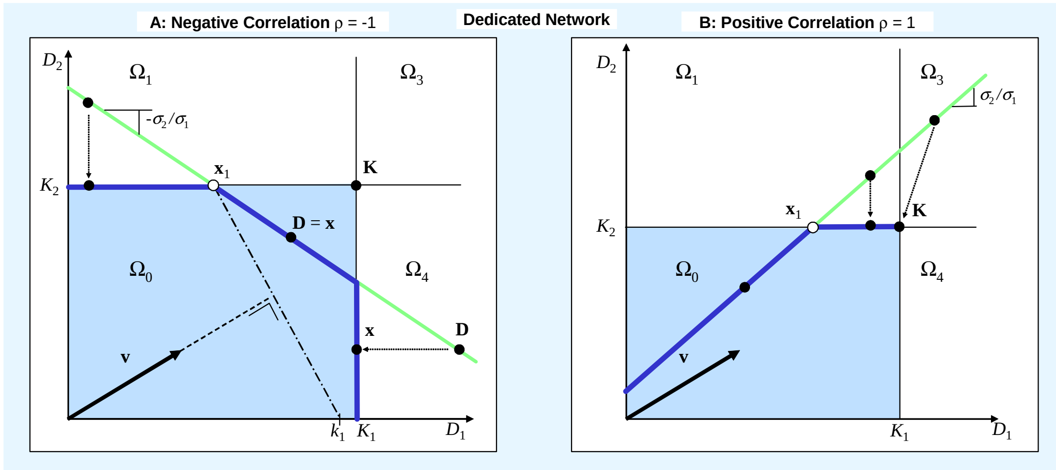

Figure 8 Activities $\mathbf{x}$ in the dedicated network depend on resource levels $\mathbf{K}$ and demand $\mathbf{D}$. Demand is on the downward-sloping line if $\rho = -1$ (left) and on the upward sloping line if $\rho = 1$ (right).

ciently reduces variance so that both capacity levels decrease in risk aversion (while the mix favors market 2). If $\rho = +1$, the diversification benefit disappears completely and capacity mix choice is dominated by expected profit maximization, which favors the higher margins of market 1.

The key insights are that, in a newsvendor network with strongly negatively correlated demand, risk-averse investors should increase investment of those resources supplying the lower profit variance market. The total investment may even exceed the risk neutral investment.

Generalizations. The insights from the numerical analysis assuming normal demand and CARA utility can be generalized. The seminal work by Eeckhoudt et al. (1995) proves that the optimal single-resource level is always decreasing in risk aversion for general concave utility functions. The remainder of this section shows that their strong result does not hold for dedicated networks with moderately concave utility functions, moderate c_i/v_i ratios, and negative correlation:

Property 2 *Consider the dedicated network. With strongly positive correlation, $\mathbf{K}^u \leq \mathbf{K}^n$ and decreases in risk aversion. However, with strongly negatively correlated demands, moderate risk aversion, $v_2\sigma_2 < v_1\sigma_1$ and standardized safety levels $z_i(K_i^n) > 0$, $K_2^u \geq K_2^n$ and increases in risk aversion.*

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 identifies weaker conditions and extends Eeckhoudt et al. (1995)'s ingenious technique of bounding the marginal utility to a newsvendor network so it merits placement in the main text:

Proof: 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 Fig. 8. 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:

$$0 = (v_1 - c_1)\,\mathbb{E}_{34}u'(V(\mathbf{K}^u, \mathbf{D}) + W) - c_1\mathbb{E}_{01}u'(V(\mathbf{K}^u, \mathbf{D}) + W), \tag{13}$$

$$0 = (v_2 - c_2)\,\mathbb{E}_{13}u'(V(\mathbf{K}^u, \mathbf{D}) + W) - c_2\mathbb{E}_{04}u'(V(\mathbf{K}^u, \mathbf{D}) + W), \tag{14}$$

where $\mathbb{E}_i f = \int_{\Omega_i} f dP$ denotes partial expectation over Ω_i. The proof assumes perfectly correlated demand; a continuity argument directly extends to a neighborhood of $\rho = -1$ or $+1$.

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

$$\frac{D_1 - \mu_1}{\sigma_1} = \frac{D_2 - \mu_2}{\sigma_2} \Leftrightarrow z_1(D_1) = z_2(D_2) \Leftrightarrow D_1 = aD_2 + b_+ \qquad \text{w.p. } 1, \tag{15}$$

using the standardized safety level notation (12) 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 panel B of Fig. 8. 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 Fig. 8, 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 (14) 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 (13) 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:

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

where we used (13) 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 $P_{34}(\mathbf{K}^u) = P(D_1 > K_1^u) \geq \frac{c_1}{v_1}$ and $P_3(\mathbf{K}^u) = P(D_2 > K_2^u) \geq \frac{c_2}{v_2}$, 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 Panel A of Fig. 8, 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 $P_1(\mathbf{K}^u) = P(D_2 > K_2^u) \leq \frac{c_2}{v_2}$. 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 $\nabla_2 U_{h \circ u}(\mathbf{K}^u; W) \geq 0$. Given concavity, this means that an infinitesimal increase h in risk aversion relative to any u that satisfies the conditions leads to $K_2^{h \circ u} > K_2^u$. ■

The risk-averse investor rebalances the capacity mix towards 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 *A dedicated network with the negative correlation conditions of Property 2 has* $\nabla_2\sigma^2(\mathbf{K}^n) \leq 0$ *and* $\frac{d}{d\gamma}K_2(0) \geq 0$.

These theoretical results formalize and generalize the intuitive explanation in terms of the diversification benefits and costs. At high correlation, diversification benefits become negligible and both resource levels *always* decrease in risk aversion. This also happens with CARA utility if demands are independent (and thus $\rho = 0$)[7]. At sufficiently low correlations, however, the diversification benefit is sufficient for a risk-averse newsvendor to increase the resource supplying the lower profit variance product provided capacity is not too expensive. (This implies that for CARA and normal demand, there exists a negative correlation threshold—like $\rho \simeq -.45$ in Panel B of Fig. 6—below which $K_2^u \geq K_2^n$ and above which $K_2^u \leq K_2^n$.)

6 Resource Sharing and Complimentarity: the Serial Network

To see how the insights from the dedicated network carry over to "true operational networks," consider the serial network of Fig. 1. Besides the financial portfolio aspect, the third resource now adds operational coupling because it is shared by both products. 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 $\mathbf{c}$ is admissible in the serial network as long as both products are economically viable: $c_i + c_3 < v_i$. Its optimal capacity portfolio is shown in Fig. 9 for the same numerical data 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.

Risk exposure. While risk and resource sharing complicates investment, risk-neutral capacities remain linear in (marginal) demand standard deviations as shown in Fig. 9. 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

[7] 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.

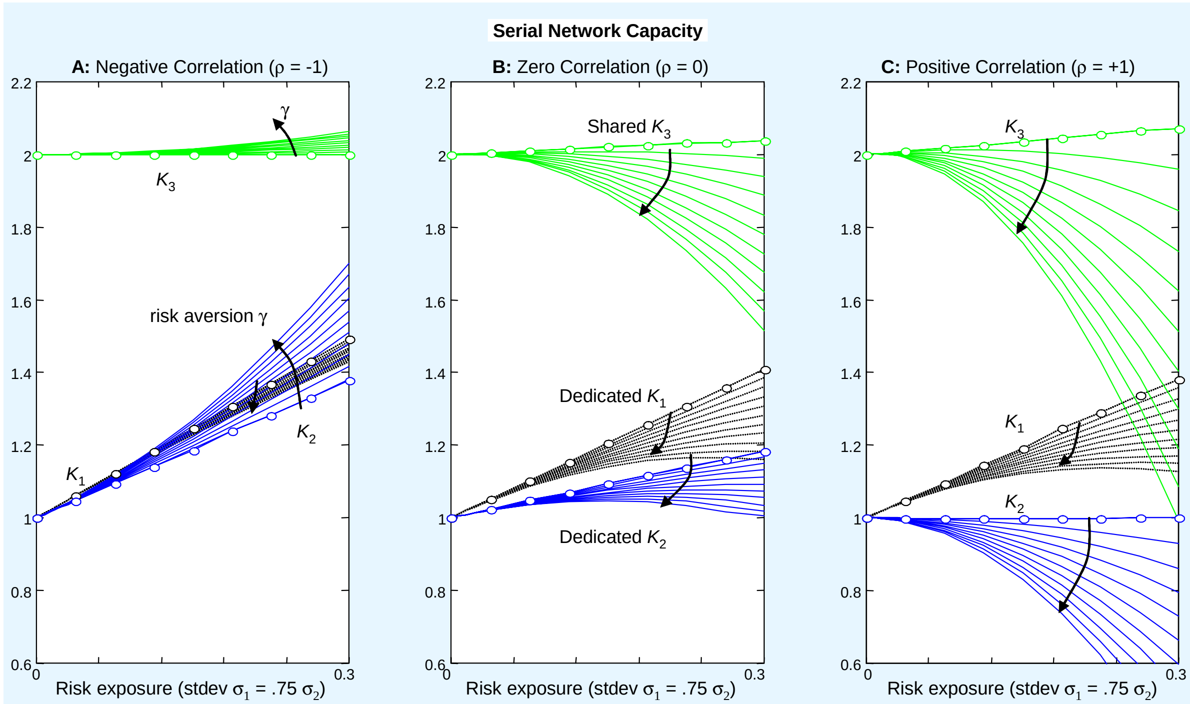

Figure 9 Serial network capacity as a function of risk exposure, risk aversion, and correlation.

correlation as well as on the value-cost ratios c/v. As before, both dedicated resource levels decrease in correlation, but the risk-neutral shared resource level K_3^n increases in correlation.

Risk aversion. The CARA-utility-optimal capacities for dedicated resource 1 and 2 have similar dynamics as in the dedicated network: Both appear quadratic in standard deviation and K_2^u *as well as* the added shared resource level K_3^u increase in risk aversion for strongly negatively correlated demand. Panel A in Fig. 10 shows that even a risk-neutral newsvendor strictly prefers more negatively correlated demand. Lower correlations strictly dominate higher correlations: The mean value of a newsvendor network is decreasing in any correlation coefficient because the operating profit is submodular in $\mathbf{D}$ for the serial network and risk pooling reduces variance more with more negative correlation. The shared resource has an additional benefit first identified in Van Mieghem (1998, Prop. 3) and absent in the dedicated network: Product-flexible resources have an ex-post *revenue (profit) maximization option* to steer or switch the output mix towards the more profitable product when the shared resource is capacity constrained (a bottleneck). This switching option is valuable only with a product profit differential ($v_1 \neq v_2$) but is independent from risk pooling as it survives even with perfect positive correlation (as the parallel network below will demonstrate). The combined benefits from risk pooling and revenue maximization are so large with strongly negative correlations that they warrant a risk-averse newsvendor to invest more money in the serial network than a risk-neutral investor would (Panel B in Fig. 10).

Operational hedging. Serial network newsvendors purposely imbalance capacity to mitigate profit risk.

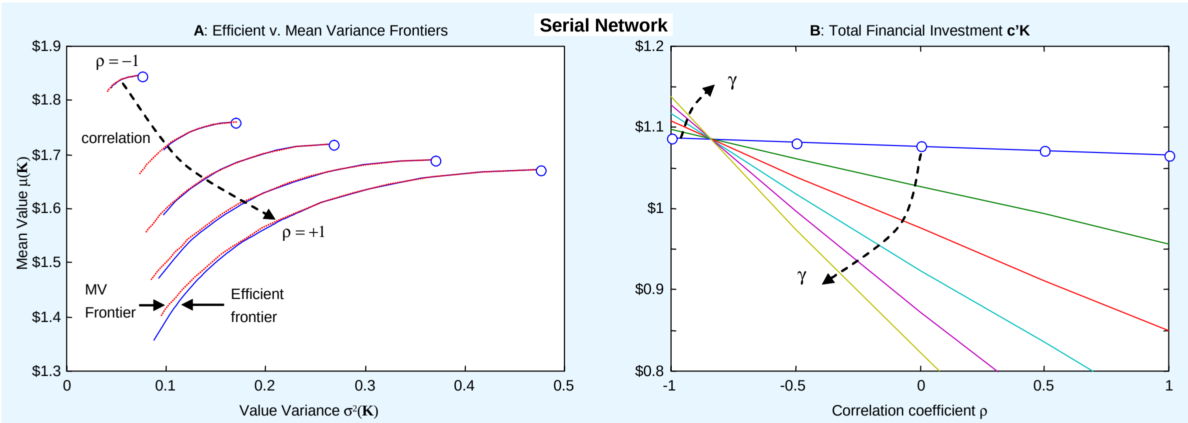

Figure 10 Frontiers and monetary investment in the serial network as a function of correlation and risk aversion. ($\sigma_1 = .3, \sigma_2 = .4$).

Harrison & Van Mieghem (1999) first identified capacity imbalance as a natural form of operational hedging. Van Mieghem (2003, Definition 1) defines a capacity portfolio $\mathbf{K}$ in a newsvendor network to be balanced if there exists an activity vector $\mathbf{x} \geq 0$ such that $\mathbf{Ax} = \mathbf{K}$ and all resources are simultaneously fully utilized. In the serial network, the condition for balance simplifies to $K_1 + K_2 = K_3$, which clearly is optimal in a deterministic setting. With demand uncertainty, however, those two articles showed that it is optimal for a risk-neutral newsvendor to under-invest in the shared resource because such resource imbalance ($K_1 + K_2 > K_3$) maximizes mean value. An obvious driver behind resource imbalance for a risk-neutral investor is risk pooling[8] which suggests that relative capacity imbalance would decrease in correlation. Fig. 11 confirms that suggestion but also shows that capacity imbalance remains optimal even with perfect positive correlations which highlights the revenue maximization option as the second driver. To exercise its switching option, the shared resource must have upstream capacity leeway which further explains the need for capacity imbalance. This paper shows that risk-averse newsvendors prefer even more capacity imbalance: Facing increased risk, they increasingly rebalance capacity away from the shared resource and more so the larger their risk aversion or correlation, as shown in Fig. 11. Given that the optimal risk-averse resource vector seeks to reduce (variance) risk, this means that more imbalance reduces more risk, in agreement with the intuitive notion of a hedge. The fact that the effect is magnified with positive correlations again highlights 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. (Recall, there is no risk pooling benefit at $\rho = 1$ and numerical analysis shows that capacity balance increases in the profit differential $v_1 - v_2$ but disappears at $\rho = 1$ if $v_1 = v_2$.)

Resource imbalance can thus be interpreted as an operational hedge to minimize the variability of network

[8] The shared resource level is driven by a critical fractile of the sum of product profit distributions whose standard deviation is less than individual distributions with strong negative correlations.

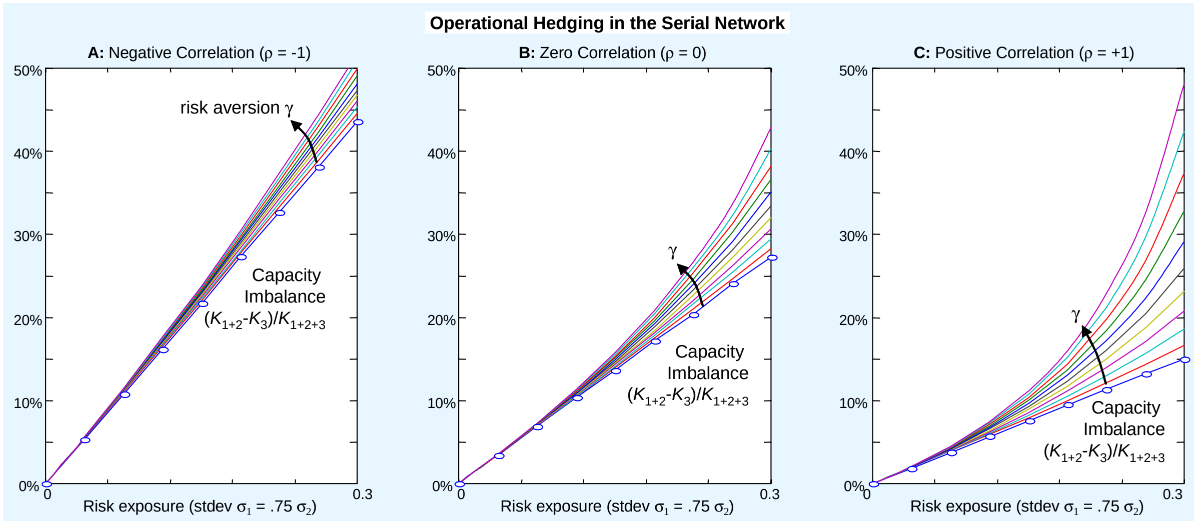

Figure 11 Operational hedging in the serial network rebalances capacity away from the shared resource.

value. Given that it also decreases mean network value, however, the hedge is not perfect. The smaller curvature of the frontiers (Panel A in Fig. 10) suggests that capacity imbalance may be more effective at higher correlations.

Generalizations. The insights from the numerical study with normal demand and CARA utility generalize:

Property 4 *Consider the serial network with strongly negatively correlated demands and moderate costs and risk aversion. If* $v_2\sigma_2 < v_1\sigma_1$*, then* $K_2^u \geq K_2^n$ *and increases in 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 *Consider the serial network with strongly negatively correlated demands and moderate costs and risk aversion. If* $v_2\sigma_2 < v_1\sigma_1$*, then* $\nabla_2\sigma^2(\mathbf{K}^n) \leq 0$ *and* $\frac{d}{d\gamma}K_2(0) \geq 0$*. If* $v_1 = v_2$*, then* $\nabla_3\sigma^2(\mathbf{K}^n) \geq 0$*.*

The earlier Fig. 9 showed that the shared resource can also 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 equal values $v_1 = v_2$, there is no revenue maximization option and the profit variance increases in the shared resource. The proposition thus highlights the role of the revenue maximization option imbedded in the shared resource: with $v_1 = 2v_2$

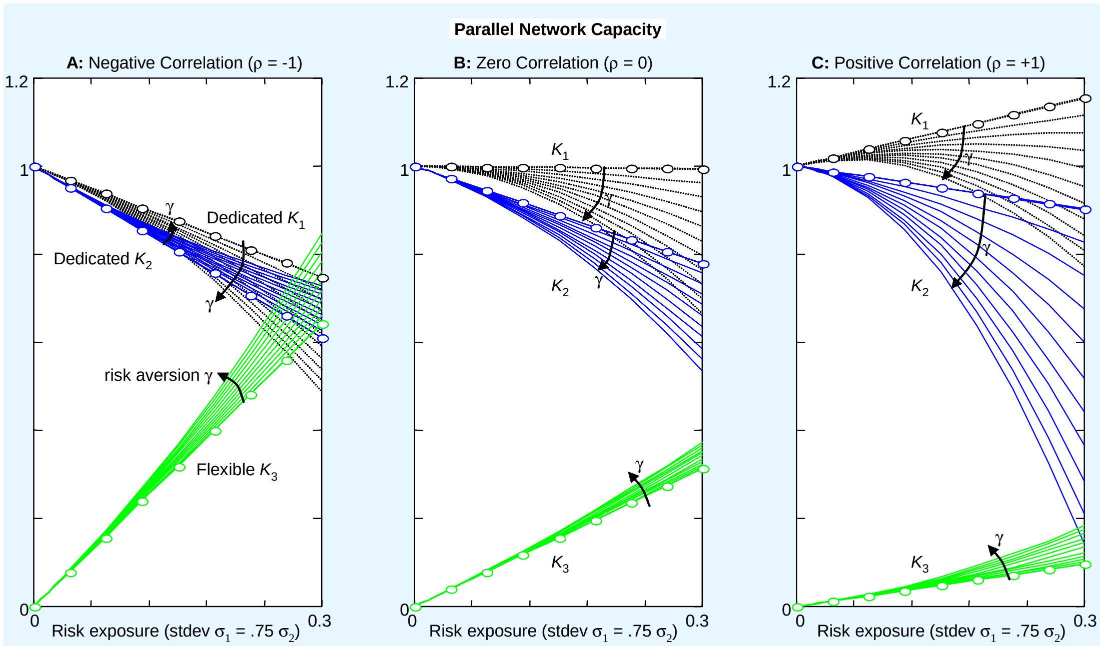

Figure 12 Parallel network capacity as a function of risk exposure, risk aversion, and correlation.

its value was sufficient to justify K_3 to increase in risk aversion in our numerical results while K_3 decreases in risk aversion if $v_1 = v_2$.

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. The effect of the product-flexible resource in the parallel network differs from that in the serial network. In the latter, the serial shared resource was a potential bottleneck. Its resource level was a complement of the upstream dedicated levels but was underweight to enable its switching option. Here, the flexible resource is a substitute of the dedicated resources as is evident from the optimal capacity portfolio shown in Fig. 12. (The numerical data is the same as before but with more expensive flexible capacity costs $c_3 = 0.7 > c_1 = c_2 = 0.5$.) 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 is driven strongly by 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 $c_i \leq v_i$.
Risk exposure. Again, risk-neutral capacities remain linear in (marginal) demand standard deviations with slopes depending on demand correlation as well as on the value-cost ratios c/v. Now, both dedicated resource levels increase in correlation while the the risk-neutral flexible resource level K_3^n decreases in correlation. The latter stems from decreasing 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

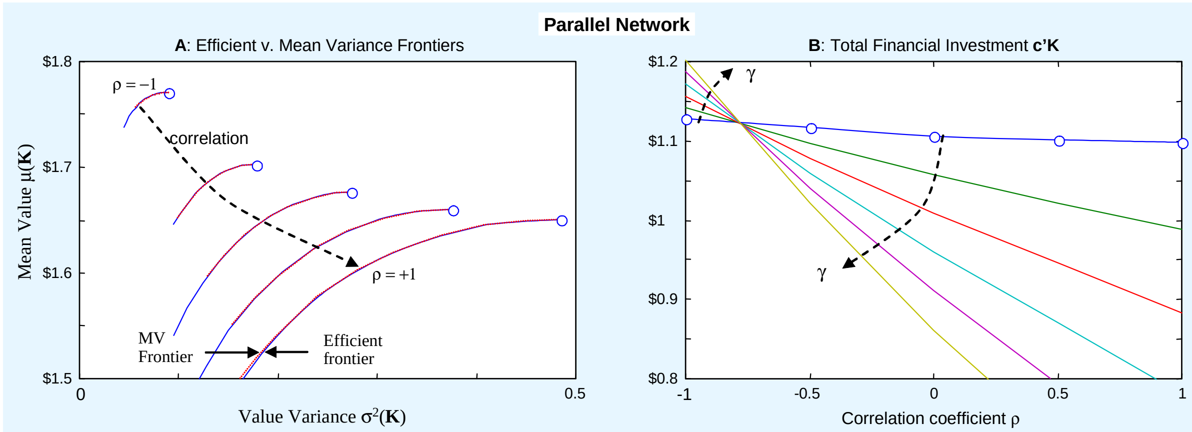

Figure 13 Frontiers and monetary investment in the parallel network as a function of correlation and risk aversion. ($\sigma_1 = .3, \sigma_2 = .4$).

option of flexibility remains.

Risk aversion. The CARA-utility-optimal capacities for dedicated resource 1 and 2 have similar dynamics as in the dedicated network: Both appear 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. With strongly negative correlations, its benefit is so large that it is optimal to "over-substitute" dedicated capacity with flexible at a higher total network investment than that under risk-neutrality (Panel B in Fig. 13). Panel 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 Panel C of Fig. 6).

Operational hedging. To mitigate profit risk, parallel network newsvendors move the investment mix towards the flexible asset, as shown in Fig. 14. While this flexibility preference decreases in correlation (reflecting decreased risk pooling benefits) it does not disappear at $\rho = 1$ (reflecting revenue maximization benefits). Flexibility has been interpreted as an operational hedge in risk-neutral models. This paper shows that the capacity mix moves even more towards the flexible asset 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 risk, reinforcing the the notion of flexibility as a hedge. Again the move towards a more flexible capacity mix is magnified with positive correlations and highlights the revenue maximization option. (Numerical analysis shows that flexibility increases in the profit differential $v_1 - v_2$ but disappears at $\rho = 1$ if $v_1 = v_2$.) Put another way, it remains optimal to reallocate funds towards the flexible asset if $v_1 > v_2$. Recall that at $\rho = 1$, the dedicated resources K_1 and K_2 are strongly decreasing in risk, thereby freeing up funds that are partially invested to increase the flexible resource level. By taking on more of the flexible asset 3 and less of the dedicated assets 1 and 2, increasing resource flexibility thus can be interpreted as

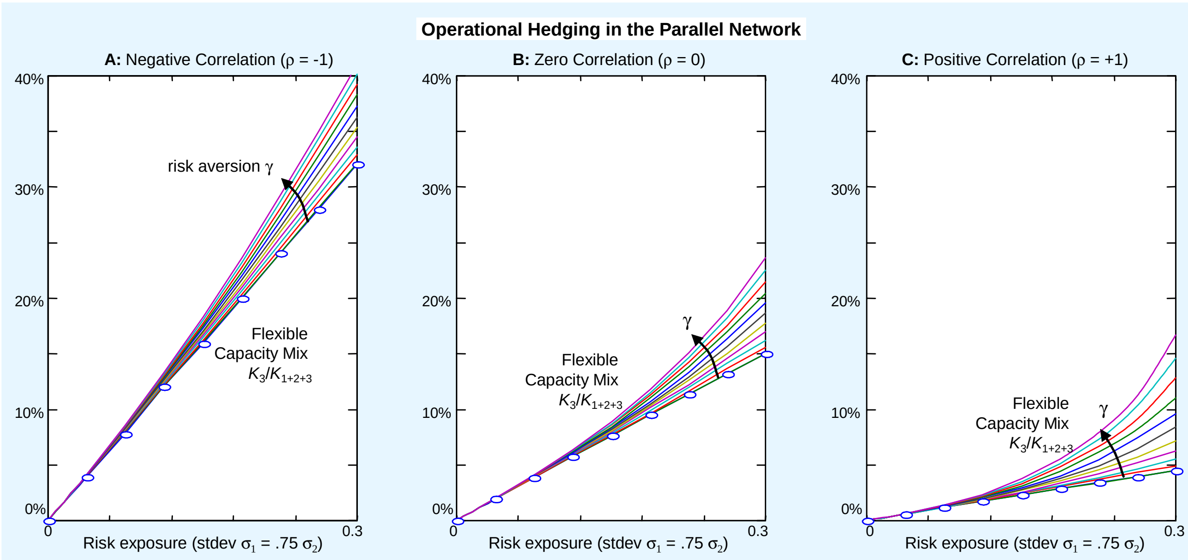

Figure 14 Operational hedging in the parallel network rebalances capacity towards the flexible resource.

an imperfect operational hedge, mitigating the impact of variability in network value at the expense of a decrease in mean value.

Generalizations. These insights from the numerical analysis assuming normal demand and CARA utility can again be generalized (but somewhat less so than earlier). Assuming that a rational risk-neutral agent invests in flexibility–Van Mieghem (1998, Proposition 7) gives conditions such that $K_3^n > 0$–risk aversion moves the capacity mix towards the lower profit variance resource 2 and the flexible resource 3:

Property 6 *Consider the parallel network with strongly negatively correlated demands 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 that risk exposure and risk aversion fundamentally change how organizational networks should be designed to handle uncertainty well. A joint operational and financial perspective was adopted to develop theory and insight on the strategic placement of safety capacity and inventory. This suggested some rules of thumb, summarized in the abstract, for designing processing networks to mitigate risk and serve as an operational hedge.

Our analysis and insights, however, have many limitations and much remains to be done. While newsvendor networks have several advantages (see 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. It may be possible to extend the analysis to dynamic newsvendor networks just like Bouakiz & 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 (but general resource networks). Multi-agent newsvendor networks are a natural extension (e.g., Van Mieghem (1999) and Goyal & Netessine (2005)) as is allowing inventory stocking at multi-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) and Heyman & 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) and Tomlin & Wang (2005) for considering loss aversion and value-at-risk measures); 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 article provided some general mathematical expressions, but only applied them in the limited setting of three networks. The future task is to expand the set of networks. 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 & Willems (2000) and Graves & 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 & Graves (1995) successfully did for risk-neutral network configuration for flexibility. One would like to characterize which network structures and parameters yield powerful 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.

Acknowledgments: I am grateful to Philipp Afeche who suggested analyzing the dedicated network as the simplest network and for feedback by Lauren Lu, Nils Rudi, and Costis Skiadas. I greatly benefited from the excellent comments and questions by three peer-reviewers and the associate editor which substantially improved this paper.

Appendix

Technical Assumptions Clearly, $V(K,\omega)$ must have an expected value; that is, it is a measurable and integrable function of ω for every K. Interchanging differentiation and integration requires conditions that bound the derivatives such as requiring that all functions are Lipschitz, as in the general conditions in Appendix A of Broadie & Glasserman (1996). In our setting all functions are concave and thus absolutely continuous on any compact subset of $\mathbb{R}^n_+$, where its right- and left-hand partial derivatives exist (and thus are finite) and are monotone increasing. Let $\nabla f(K,\omega)$ denote this vector of right-hand partial derivatives of a concave function f with respect to K. An absolutely continuous function f satisfies a Lipschitz condition if and only if $||\nabla f||$ is bounded. Clearly, for a concave function, $\nabla f(K,\omega)$is bounded, and thus Lipschitz, on any compact subset in $\mathbb{R}^n_+$. The only technical condition is to require finite derivatives also at 0 and ∞ and require a Lipschitz condition on the open set $\mathbb{R}^n_+$:

Assumption 1 *The value function $V(K,\omega)$ satisfies a Lipschitz condition on $\mathbb{R}^n_+$ almost surely: there is a $M_V(\omega)$ such that $|V(K_1,\omega) - V(K_2,\omega)| \le M_V(\omega)||K_1 - K_2||$ for all $K_1, K_2 \in \mathbb{R}^n_+$, where $\mathbb{E}M_V(\omega) < \infty$.*

Assumption 2 *The utility function $u(x)$ satisfies a Lipschitz condition on $\mathbb{R}$: there is a M_u such that $|u(x_1) - u(x_2)| \le M_u|x_1 - x_2|$ for all $x_1, x_2 \in \mathbb{R}$.*

For newsvendor networks, π and thus also V are the solution of a linear program and have finite partial derivatives so that Assumption 1 is always satisfied. It is not unrealistic to assume that demands are bounded which then obviates Assumption 2.

Proof of Proposition 1: The function $V(\mathbf{K},\omega)$ is concave in $\mathbf{K}$ for any ω as a sum of two concave functions. Because $u(\cdot)$ is concave increasing, the scalar composition $u(V(\mathbf{K},\omega) + W)$ is also concave in $\mathbf{K}$ for any ω and W. (The latter is directly shown for twice differentiable functions, but also holds without assuming differentiability, see Boyd & Vandeberghe (2004, p. 84).) Finally, the expected utility function is concave as a linear combination of concave functions.■

Proof of Proposition 2: Let $f(\mathbf{K},\omega) = u(V(\mathbf{K},\omega)+W)$, which is concave in K for any ω and wealth W according to the proof of Proposition 1 so that $\nabla U(\mathbf{K}^*) = 0$ is necessary and sufficient for an interior maximum.

Because f is concave, its right-hand partial gradient $\nabla f(\mathbf{K},\omega)$ exists for every ω. Thus, for all $\mathbf{K} \in \mathbb{R}^n_+$ and $m > 0$, $g_m = m\left(f\left(\mathbf{K} + m^{-1}\mathbf{e}_i, \omega\right) - f(x,\omega)\right) \to_m \nabla_i f(\mathbf{K},\omega)$, where $\mathbf{e}_i$ is the i-th unit vector. Given that the Lipschitz property is preserved by composition, the technical assumptions guarantee the existence of Lipschitz modulus $M(\omega)$ for f w.p. 1 that is integrable. Because $|g_m(\mathbf{K},\omega)| < M(\omega)$ with $\mathbb{E}M(\omega) < \infty$, the dominated convergence theorem shows that $\lim_{m\to\infty} \mathbb{E}g_m = \mathbb{E}\lim_{m\to\infty} g_m$. Thus, differentiation and integration interchange so that $\nabla U(\mathbf{K}) = \mathbb{E}\nabla f = \mathbb{E}u'(V + W)\nabla V$. ■

Proof of Proposition 3: Applying the definition of covariance and given that integration and differentiation can be interchanged, we have that:

$$\begin{aligned}\nabla\sigma^2(\mathbf{K}) &= \nabla_{\mathbf{K}}\left(\mathbb{E}\pi^2(\mathbf{K},\mathbf{D}) - (\mathbb{E}\pi(\mathbf{K},\mathbf{D}))^2\right) = \mathbb{E}\nabla_{\mathbf{K}}\pi^2(\mathbf{K},\mathbf{D}) - \nabla_{\mathbf{K}}(\mathbb{E}\pi(\mathbf{K},\mathbf{D}))^2 \\ &= 2\mathbb{E}\left[\pi(\mathbf{K},\mathbf{D})\lambda(\mathbf{K},\mathbf{D})\right] - 2(\mathbb{E}\pi(\mathbf{K},\mathbf{D}))\mathbb{E}\lambda(\mathbf{K},\mathbf{D}) = 2Cov(\lambda,\pi).\end{aligned}$$

For the second part, again applying the definition of covariance, we have that $\mathbb{E}\left[\left(\boldsymbol{\lambda}(\mathbf{K}^n,\omega) - \mathbf{c}\right)V(\mathbf{K}^n,\omega)\right] =$

$$\begin{aligned} & Cov(\boldsymbol{\lambda}(\mathbf{K}^n,\omega) - \mathbf{c}, V(\mathbf{K}^n,\omega)) + \mathbb{E}\left[\boldsymbol{\lambda}(\mathbf{K}^n,\omega) - \mathbf{c}\right]\mathbb{E}V(\mathbf{K}^n,\omega) \\ = \; & Cov(\boldsymbol{\lambda}(\mathbf{K}^n,\omega) - \mathbf{c}, \pi(\mathbf{K}^n,\omega) - C(\mathbf{K}^n)) \qquad \text{(second term = 0 by (3))} \\ = \; & Cov(\boldsymbol{\lambda}(\mathbf{K}^n,\omega), \pi(\mathbf{K}^n,\omega)) \qquad \text{(constants fall out)}. \qquad \blacksquare\end{aligned}$$

Proof of Proposition 4: According to the implicit function theorem, $\mathbf{K}(\gamma)$ is a continuous function of γ where $\frac{d}{d\gamma}\mathbf{K}(\gamma)$ is found by differentiating the first-order condition: $\frac{d}{d\gamma}\mathbb{E}\left[\left(\boldsymbol{\lambda}(\mathbf{K},\omega) - \mathbf{c}\right)\exp\left(-\gamma V(\mathbf{K},\omega)\right)\right] = 0$ or

$$\begin{aligned}\mathbb{E}\left[\left(\nabla_{\mathbf{K}}\boldsymbol{\lambda}(\mathbf{K},\omega)\right)\exp\left(-\gamma V(\mathbf{K},\omega)\right) + \left(\boldsymbol{\lambda}(\mathbf{K},\omega) - \mathbf{c}\right)\exp\left(-\gamma V(\mathbf{K},\omega)\right)(-\gamma)\nabla_{\mathbf{K}}V\right]\frac{d}{d\gamma}\mathbf{K}(\gamma) & \\ +\mathbb{E}\left[\left(\boldsymbol{\lambda}(\mathbf{K},\omega) - \mathbf{c}\right)\exp\left(-\gamma V(\mathbf{K},\omega)\right)(-V(\mathbf{K},\omega))\right] &= 0\end{aligned}$$

Recall that the Hessian $H(\mathbf{K}^n) = \mathbb{E}\left[(\nabla_{\mathbf{K}}\boldsymbol{\lambda}(\mathbf{K}^n,\omega))\right]$ and evaluate at the risk-neutral case $\gamma = 0$ to get:

$$H(\mathbf{K}^n)\frac{d}{d\gamma}\mathbf{K}(0) = \mathbb{E}\left[(\boldsymbol{\lambda}(\mathbf{K}^n,\omega) - \mathbf{c})\,V(\mathbf{K}^n,\omega)\right].$$

Given that Π is concave, its Hessian $H(\mathbf{K}^n)$ is negative-definite, and invertible.■

Proof of (11): As illustrated in Figure 2, the MV-frontier function $\mathcal{F}$ and the maximal utility $U^{MV}(\gamma)$, denoted as $\mathcal{U}(\gamma)$, are almost inverse functions in that they satisfy, except at possible inflection points: $\mathcal{U}'(\mathcal{F}'(x)) = -x$ and thus $\mathcal{U}''(\mathcal{F}'(x)) = -1/\mathcal{F}''(x)$. Evaluating at $x = \sigma^2(\mathbf{K}^n)$, where $z = \mathcal{F}'(x) = 0$, directly yields $\mathcal{U}'(0) = -\sigma^2(\mathbf{K}^n)$ and $\mathcal{F}''(\sigma^2(\mathbf{K}^n)) = -1/\mathcal{U}''(0)$. It only remains to find $\mathcal{U}''(0)$. Twice differentiate the defining condition of $\mathbf{K}(z)$:

$$\begin{aligned}
\frac{d}{d\gamma}\mathcal{U}(\gamma) &= \nabla'\mu(\mathbf{K}(\gamma))\frac{d}{d\gamma}\mathbf{K} - \frac{\gamma}{2}\nabla'\sigma^2(\mathbf{K}(\gamma))\frac{d}{d\gamma}\mathbf{K} - \frac{1}{2}\sigma^2(\mathbf{K}(\gamma)),\\
\frac{d^2}{d\gamma^2}\mathcal{U}(\gamma) &= \left(H(\mathbf{K}(\gamma))\frac{d}{d\gamma}\mathbf{K}\right)'\frac{d}{d\gamma}\mathbf{K} + \nabla'\mu(\mathbf{K}(\gamma))\frac{d^2}{d\gamma^2}\mathbf{K} - \frac{1}{2}\nabla'\sigma^2(\mathbf{K}(\gamma))\frac{d}{d\gamma}\mathbf{K}\\
&\quad -\frac{\gamma}{2}\frac{d}{d\gamma}\left(\nabla'\sigma^2(\mathbf{K}(\gamma))\frac{d}{d\gamma}\mathbf{K}\right) - \frac{1}{2}\nabla'\sigma^2(\mathbf{K}(\gamma))\frac{d}{d\gamma}\mathbf{K}.
\end{aligned}$$

Evaluate at $\gamma = 0$ and recall that $\nabla\mu(\mathbf{K}(0)) = 0$ and $H(\mathbf{K}(0))\frac{d}{d\gamma}\mathbf{K}(0) = \nabla\sigma^2(\mathbf{K}(0))$:

$$\mathcal{U}''(0) = -\nabla'\sigma^2(\mathbf{K}(0))'\frac{d}{d\gamma}\mathbf{K}(0) = -\nabla'\sigma^2(\mathbf{K}^n)H^{-1}(\mathbf{K}^n)\nabla\sigma^2(\mathbf{K}^n). \qquad ■$$

Proof of Property 1: For the single-resource newsvendor, we have that $V(K,D) = vx(K,D) - cK$, where $x = \min(K,D)$. Thus, $\lambda = v1_{\{D\geq K\}}$ with the familiar risk-neutral optimality condition $\mathbb{E}(\lambda - c) = v(1-F(K)) - c = 0$. Using the standard normal pdf ϕ and cdf Φ, $v(1-\Phi(z^n)) = c$ where $z^n = (K^n - \mu_1)/\sigma_1$. Hence, $H(K^n) = \frac{d}{dK}\mathbb{E}\lambda = -\frac{v}{\sigma_1}\phi(z^n) < 0$ and

$$\begin{aligned}
\mathbb{E}\left[(\boldsymbol{\lambda}(\mathbf{K}^n,\omega) - \mathbf{c})\,V(\mathbf{K}^n,\omega)\right] &= \int_{-\infty}^{K^n}(-c)(vx - cK^n)dF + \int_{K^n}^{\infty}(v-c)(vK^n - cK^n)dF\\
&= -cv\underbrace{\int_{-\infty}^{K^n} x dF}_{A} + \underbrace{c^2K^nF(K^n) + (v-c)^2K^n(1-F(K^n))}_{B}
\end{aligned}$$

For the normal distribution, integration by parts yields $A = \mu_1\Phi(z^n) - \sigma_1\phi(z^n) = \mu_1(1-c/v) - \sigma_1\phi(z^n)$. Moreover $B = \left[c^2(1-c/v) + (v-c)^2c/v\right]K^n = c\,(v-c)\,K^n$. Putting this together yields

$$\begin{aligned}
\mathbb{E}\left[(\boldsymbol{\lambda}(\mathbf{K}^n,\omega) - \mathbf{c})\,V(\mathbf{K}^n,\omega)\right] &= -cv\left(\mu_1(1-c/v) - \sigma_1\phi\right) + c\,(v-c)\left(\mu_1 + z\sigma_1\right)\\
&= \sigma_1 cv\left(\phi(z^n) + (1-c/v)z^n\right),
\end{aligned}$$

which is non-negative because $f(z) = \phi(z) + z\Phi(z)$ has $f(-\infty) = 0$ and is non-decreasing ($f'(z) = \Phi(z) \geq 0$).■

Proof of Property 3: Using the notation of the proof of Property 2:

$$\begin{aligned}
\frac{1}{2}\nabla_2\sigma^2(\mathbf{K}^n) &= \mathbb{E}\left(\lambda_2(\mathbf{K}^n,\omega) - c_2\right)V(\mathbf{K}^n,\omega)\\
&= (v_2 - c_2)\mathbb{E}_{13}V(\mathbf{K}^n,\mathbf{D}) - c_2\mathbb{E}_{04}V(\mathbf{K}^n,\mathbf{D})\\
&\leq V_1(\mathbf{K}^n)\left[(v_2-c_2)P_1(\mathbf{K}^n) - c_2P_{04}(\mathbf{K}^n)\right] = 0.
\end{aligned}$$

Also $H_{12} = \frac{\partial}{\partial K_2}\mathbb{E}\lambda_1(\mathbf{K}^n,\omega) = \frac{\partial}{\partial K_2}v_1P(D_1 > K_1^n) = 0$, so that

$$H^{-1}(\mathbf{K}^n) = \begin{bmatrix} \frac{-1}{v_1f_1(K_1^n)} & 0 \\ 0 & \frac{-1}{v_2f_2(K_2^n)} \end{bmatrix},$$

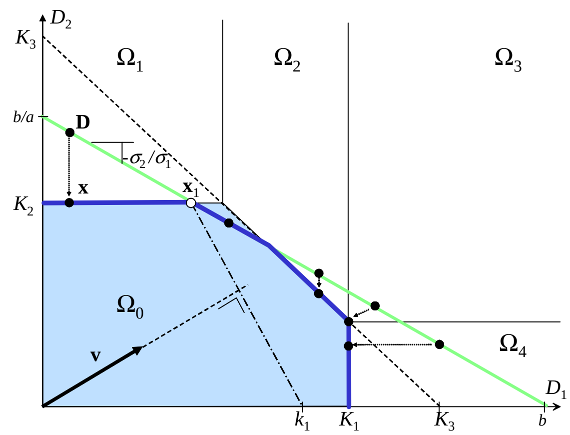

Figure 15 The activity vector $\mathbf{x}$ for the serial network when $\sigma_1 \geq \sigma_2$ and $\rho = -1$.

where $f_i(\cdot)$ is the pdf of D_i.■

Proof of Property 4 is similar to that of Property 2: The activity vector $\mathbf{x}(\mathbf{K}, \mathbf{D})$ again is a simple greedy solution: $x_1 = \min(D_1, K_1, K_3)$ and $x_2 = \min(D_2, K_2, K_3 - x_1)$. The optimal risk-neutral $\mathbf{K}^n$ satisfies the optimality conditions $\mathbb{E}\boldsymbol{\lambda}(\mathbf{K}^u, \omega) = \mathbf{c}$:

$$(v_1 - v_2)P_3(\mathbf{K}^n) + v_1 P_4(\mathbf{K}^n) = c_1, v_2 P_1(\mathbf{K}^n) = c_2, v_2 P_{2+3}(\mathbf{K}^n) = c_3. \quad (16)$$

According to Proposition 2, the (second) optimality condition for the risk-averse resource vector $\mathbf{K}^u$ is::

$$0 = (v_2 - c_2)\mathbb{E}_1 u'(V(\mathbf{K}^u, \mathbf{D}) + W) - c_2 \mathbb{E}_{0234} u'(V(\mathbf{K}^u, \mathbf{D}) + W).$$

Case 1: $\sigma_1 \geq \sigma_2$ as shown in Fig. 15. Notice that it is suboptimal for the point $(K_1^u, K_3^u - K_1^u)$ to be above the demand line because reducing K_3^u by ϵ would not change operating profits but reduce investment costs. Thus there are two possible cases: moderate c_3 so that $(K_3^n - K_2^n, K_2^n)$ falls above/at the demand line; with high c_3 it falls below. The property applies in the former case. The proof is similar to that of Property 2: define the vector $\mathbf{x}_1(\mathbf{K})$ and scalar $k_1(\mathbf{K})$ and establish that, if $k_1(\mathbf{K}^u) < K_1^u$, then $P_1(\mathbf{K}^u) \leq \frac{c_2}{v_2} = P_1(\mathbf{K}^n)$. Note that $P_1(\mathbf{K}) = P(D_2 > K_2)$ so that $K_2^u \geq K_2^n$. The required conditions are (1) and (2) of the proof of Property 2 and (3) the point $(K_3^n - K_2^n, K_2^n)$ falls above the demand line or $z_1(K_3^n - K_2^n) + z_2(K_2^n) > 0$. The increasing-in-risk aversion is proved similarly to the proof of Property 2.

Case 2: $\sigma_1 < \sigma_2$ proceeds similarly but uses the point $\mathbf{x}_1(\mathbf{K}) = (K_3 - K_2, K_2)$. The required conditions are (1) of the proof of Property 2; (2) $k_1(\mathbf{K}^n) < K_1^n$ or $K_3^n < K_1^n + (1 - v_2/v_1)K_2^n$; and (3) the point $(K_1^n, K_3^n - K_1^n)$ falls above the demand line or $z_1(K_1^n) + z_2(K_3^n - K_1^n) > 0$.■

Proof of Property 5: Use the notation of the proof of Property 4.

$$\begin{aligned} \frac{1}{2}\nabla_2 \sigma^2(\mathbf{K}^n) &= \mathbb{E}\left(\lambda_2(\mathbf{K}^n, \omega) - c_2\right) V(\mathbf{K}^n, \omega) \\ &= (v_2 - c_2)\mathbb{E}_1 V(\mathbf{K}^n, \mathbf{D}) - c_2 \mathbb{E}_{0234} V(\mathbf{K}^n, \mathbf{D}) \\ &\leq V_1(\mathbf{K}^n)\left[(v_2 - c_2)P_1(\mathbf{K}^n) - c_2 P_{0234}(\mathbf{K}^n)\right] = 0. \end{aligned}$$

It is also easily verified that

$$H^{-1}(\mathbf{K}^n) = \frac{1}{|H|}\begin{bmatrix} \cdots & 0 & \cdots \\ 0 & \cdots & 0 \\ \cdots & 0 & \cdots \end{bmatrix},$$

where all non-zero elements (denoted by $\cdots$) are positive and $|H| < 0$. Similarly, if $v_1 = v_2$:

$$\begin{aligned} \frac{1}{2}\nabla_3\sigma^2(\mathbf{K}^n) &= \mathbb{E}\left(\lambda_3(\mathbf{K}^n,\omega) - c\right)V(\mathbf{K}^n,\omega) \\ &= (v_2 - c_3)\mathbb{E}_{23}V(\mathbf{K}^n,\mathbf{D}) - c_3\mathbb{E}_{014}V(\mathbf{K}^n,\mathbf{D}) \\ &\geq V_3(\mathbf{K}^n)\left[(v_2 - c_3)P_{23}(\mathbf{K}^n) - c_3P_{014}(\mathbf{K}^n)\right] = 0, \end{aligned}$$

where $V_3(\mathbf{K}^n) = V(\mathbf{K}^n, \mathbf{x}_3(\mathbf{K}^n))$ where $\mathbf{x}_3(\mathbf{K}) = (K_1, K_3 - K_1)$.■

Proof of Property 6 is similar to that of Property 2: Notice that $K_1^u + K_2^u + K_3^u > b$ is suboptimal because reducing K_3^u by ϵ and increasing K_1^u by ϵ would not change operating profits but would decrease investment cost by $(c_3 - c_1)\epsilon > 0$. Thus there are two possible cases: $K_1^u + K_2^u + K_3^u < b$ if c_3 is high and $K_1^u + K_2^u + K_3^u = b$ otherwise. The property applies in the latter boundary case which is shown in Figure 16.

Let $\mathbf{K}_{1:2} = (K_1, K_2)$ be the independent variable for this boundary case where $K_3 = b - K_1 - K_2$. The associated two-dimensional shadow vector on this boundary has components $\lambda_1^b(\mathbf{K}_{1:2}^u, \mathbf{D}) = -v_2 1_{\{\mathbf{D}\in\Omega_1(\mathbf{K}_{1:2}^u)\}}$ and $\lambda_2^b(\mathbf{K}_{1:2}^u, \mathbf{D}) = -v_1 1_{\{\mathbf{D}\in\Omega_4(\mathbf{K}_{1:2}^u)\}}$ with effective marginal cost $c^b = (c_1 - c_3, c_2 - c_3) < 0$. The risk-neutral boundary solution satisfies $\mathbb{E}\boldsymbol{\lambda}^b(\mathbf{K}^n, \mathbf{D}) = \mathbf{c}^b$ so that $P_1(\mathbf{K}^n) = (c_3 - c_1)/v_2$ and $P_4(\mathbf{K}^n) = (c_3 - c_1)/v_2$. Define the vector $\mathbf{x}_1(\mathbf{K}) = (K_1, K_2 + K_3)$ to partition and bound marginal utilities similar to the proof of Property 2: The optimality conditions for $\mathbf{K}^u$ include

$$\begin{aligned} 0 &= \mathbb{E}\left(\lambda_1^b(\mathbf{K}^u,\mathbf{D}) - c_1^b\right)u'(V(\mathbf{K}^u,\mathbf{D}) + W) \\ &= (-v_2 - c_1 + c_3)\mathbb{E}_1 u'(V(\mathbf{K}^u,\mathbf{D}) + W) + (-c_1 + c_3)\,\mathbb{E}_{24}u'(V(\mathbf{K}^u,\mathbf{D}) + W) \\ &\leq u'(V_1(\mathbf{K}^u))\left[-v_2P_1(\mathbf{K}^u) - c_1 + c_3\right] \Rightarrow P_1(\mathbf{K}^u) \leq (c_3 - c_1)/v_2. \end{aligned}$$

Thus, $P_1(\mathbf{K}^u) \leq P_1(\mathbf{K}^n)$ so that $K_1^u \leq K_1^n$ and $K_2^u + K_3^u \geq K_2^n + K_3^n$. The increase in risk aversion is proved similarly to the proof of Property 2. The required conditions are (1) $v_1 > v_2$ and $\sigma_1 = \sigma_2$; (2) $k_1(\mathbf{K}^n) < K_1^n + K_3^n$ or $v_2K_2^n < (v_1 - v_2)K_3^n$; (3) $K_1^n + K_2^n + K_3^n = b$ or conditions (c) of Prop. 7 of Van Mieghem (1998, Proposition 7).■

References

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

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

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

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

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

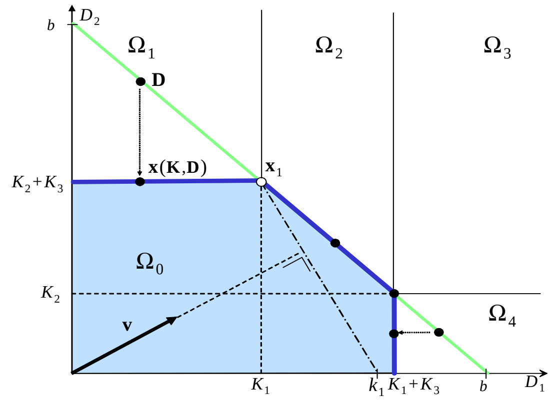

Figure 16 The activity vector $\mathbf{x}$ for the parallel network when $\sigma_1 = \sigma_2$ and $\rho = -1$.

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

Boyabatlı, O. & Toktay, L. B. (2004), Operational hedging: A review with discussion, Technical report, INSEAD 2004/12/TM.

Boyd, S. & Vandeberghe, L. (2004), *Convex Optimization*, Cambridge University Press, Cambridge, UK.

Brealy, R. A. & Myers, S. C. (2000), *Principles of corporate finance*, 6 edn, Irwin McGraw-Hill, Boston, MA.

Broadie, M. & Glasserman, P. (1996), 'Estimating security price derivatives using simulation', *Mgt. Sci.* **42**(2), 269–285.

Caldentey, R. & Haugh, M. (2003), 'Optimal control and hedging of operations in the presence of financial markets', *Math. of Oper. Res.* . Forthcoming.

Chen, X., Sim, M., Simchi-Levi, D. & Sun, P. (2004), 'Risk aversion in inventory management', *MIT working paper* .

Corbett, C. J. & Rajaram, K. (2005), A generalization of the inventory pooling effect to non-normal dependent demand, Technical report, UCLA. available at personal.anderson.ucla.edu/charles.corbett.

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

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

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

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

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

Gollier, C. (2001), *The Economics of Risk and Time*, The MIT press, Cambridge.

Goyal, M. & Netessine, S. (2005), Strategic technology choice and capacity investment under demand uncertainty, Technical report, Wharton. available from www.netessine.com.

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

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

Hale, W., Pyke, D. F. P. & Rudi, N. (2000), An assemble-to-order system with component substitution, *in* 'Proceedings of the 4th MSOM Conference', INFORMS.

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

Heyman, D. & Sobel, M. (1984), *Stochastic Models in Operations Research*, Vol. 2, McGraw-Hill.

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

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

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

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

Markowitz, H. M. (1991), 'Foundations of portfolio theory', *Journal of Finance* **46**(2), 469–477.

Merriam-Webster's Collegiate Dictionary (1998), tenth edn, Merriam-Webster, Springfield, MA.

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

Oxford English Dictionary Online (2004), second and following edn, Oxford University Press, Oxford.

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', *Journal of Economic Literature* **36**(1), 11–46.

Rudi, N. (2000), Some Models of Risk Pooling, PhD thesis, Wharton School at the University of Pennsylvania.

Spulber, D. F. (1985), 'Risk sharing and inventories', *Journal of Economic Behavior and Organization* **6**, 55–68.

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

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

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

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

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

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