Coordination and Risk-Sharing in E-Business

Victor F. Araman[1], Jochen Kleinknecht[2], and Ram Akella[3]

Technical Report

June 2002

[1]Stanford University
[2]Stanford University
[3]State University of New York at Buffalo

Abstract

Internet-based marketplaces represent an additional channel for the procurement of material or components. In this paper, we study the effects of such a spot market on a supplier-retailer problem, where the retailer operates in a make-to-order environment and where the supplier offers a long-term contract to the retailer. We determine a class of contracts that allow both means of supply (spot market and long-term contract) to coexist, and we give a necessary and sufficient condition for those contracts to achieve channel coordination. We introduce a new type of contract, a linear risk sharing contract, that achieves system efficiency and enables a range of profit split between the supplier and the retailer. We also study closely the retailer's problem under the linear risk sharing contract, discuss numerical examples, and provide managerial insights.

Keywords: Supply Contracts; Channel Coordination; Internet-based Exchanges.

1   Introduction

Recently, internet-based marketplaces and on-line trading communities have been praised for enabling the new revolution in procurement. As a consequence, Business-to-Business companies have been mushrooming, with many of them, e.g., Ariba, Commerce One, Chemdex, or FreeMarkets, attaining breath-taking stock market capitalizations before the downturn in the stock market. Even traditional bricks-and-mortar companies such as Ford and General Motors or Boeing have been setting up their own Web markets. However, despite the hype around these new channels of procurement, only very few manufacturers use them actively, as is manifested by the low transaction volumes of most of the internet-based marketplaces (see [5, 2001]). Part of the reason for this slow acceptance is that manufacturers struggle with managing the risk that the varying price and delivery schedules of the on-line trading communities introduce into their procurement process. Likewise, suppliers are hesitant to utilize these marketplaces as a sales channel.

Moreover, although it is fairly evident that the Web can lead to tremendous cost savings by automating transactions and that the buyers should be at least as well off in a world with internet-based exchanges where everything else remains the same, experts are unclear about the underlying economics. In particular, there is an apparent need for analytical models that identify the circumstances under which B2B marketplaces will succeed, and to assess whether they will replace traditional long-term supply contracts, or whether they will coexist with long-term contracts in the long run. Although the current industry trend is towards collaborative commerce between manufacturers and suppliers, and customers (see [18, 2001]), signifying that companies now again try to establish long-term contractual relationships with their suppliers, this may not be optimal in the long run and the critical issue becomes: Which type of contractual relationship between a retailer and her supplier will be efficient, i.e., lead to maximum system profit in a decentralized setting? - And what is the optimal procurement

strategy for the retailer under these efficient contracts? Since spot markets offer more flexibility to the buyer, they are expected to be more expensive than long-term contracts (at least in expectation), and one might thus argue that spot markets will only be used relatively infrequently, as needed. In the event of a spot market being less expensive, there still remain issues of the impact of the actual (possibly uncertain) unit price on the spot market and total available capacity. Yet another issue is whether the spot markets' current low levels of transaction volumes prevent manufacturers from using spot markets or whether this is more temporary phenomenon?

In this paper, our objective is to provide a generic framework for the study of supply chain coordination in general, including in the presence of internet-based marketplaces. In particular, we consider a retailer who can procure components or raw materials either via a traditional long-term contract, an internet-enabled spot market, or by using a combination of both means of supply. The supplier can sell only via the long-term contract and defines the pricing schedule associated with the contract. The retailer then chooses the cost minimizing level of contract quantity and uses the spot market only if the demand exceeds the reserved quantity level.

We first introduce a general set of supply contracts and derive necessary and sufficient conditions for them to be efficient in such a situation, i.e., to achieve maximum system profits in a decentralized setting. Notice that the class of supply contracts under consideration is general enough to include prominent examples such as quantity flexibility, buy-back, and fixed quantity commitment contracts. Another contribution of this paper is that we introduce a new class of supply contracts, which we refer to as linear risk sharing scheme. This type of contracts, considered as very realistic by companies such as General Motors, permits to reach efficiency in a decentralized setting, and is very simple to implement.

The questions we address for this particular contract are the following: 1. Under what conditions is it optimal to use both supply channels? 2. What determines the

optimal production quantity level? - Using a single period setting (e.g., corresponding to a quarter or year) with the objective of determining the profit-maximizing procurement mix, we demonstrate that both procurement channels provide economic value for buyers and that they will thus coexist with long-term supply channels. We prove and demonstrate that the optimal value of the production quantity level committed by the retailer to the long-term supplier is in most cases close to the mean value of demand. To our knowledge, for the first time, we are demonstrating that spot markets have an economic value and will thus be used on a sustained basis.

This paper is organized as follows: In Section 2 we give a brief summary of the related financial and supply chain management literature. We also state how our approach differs from those used by the various authors. Section 3 presents the general framework and discuss how more general models are embedded in the one we are analyzing here. In Section 4 we introduce the class of contracts that we will be focusing on, and determine a necessary and sufficient condition for channel coordination. Finally, in Section 5 we introduce a new type of contract, which we call "linear risk sharing contract", that belongs to the class of contracts we are analyzing in this paper, and we show that we can reach system efficiency with this type of contract. Using the linear risk sharing contract, we study the retailer's problem more closely and we discuss numerical examples, which provide managerial insights. We conclude with a summary of our results and discuss plans for future research that we deem important. Most proofs are collected in the Appendix.

2   Related Literature

Two major lines of literature are relevant to this paper: the operations and supply chain management literature on supply contracts and dual sourcing and the finance and economics literature on options and dual supply.

3

The literature on supply chain contracts and its impact on supply chain coordination has exploded over the past few years. Recent review papers such as the ones by Lariviere ([10]), by Tsay et al. ([17, 1998]), and by Anupindi and Bassok ([1, 1998]) summarize many of the results. Li and Kouvelis ([11, 1999]) develop valuation methodologies for different types of supply contracts under deterministic demand but uncertain prices. They also consider a linear "risk-sharing" contract similar to the one we adopt. Prominent examples of research regarding channel coordination are the following: Tsay ([16]) shows that a quantity flexibility contract, which specifies a minimum and a maximum purchasing quantity, can achieve maximum system profits in a decentralized setting. Donohue ([7]) finds that buy-back contracts can coordinate retailer and supplier, too. More recently, Barnes et al. ([4]) investigate the role of options in a buyer-supplier system and show that channel coordination can be achieved only if the options exercise price is piecewise linear.

Channel coordination has also been extensively studied in the marketing literature. Pasternack studies buy-back contracts with wholesale pricing and return policies and shows that they allow to achieve channel coordination while ensuring individual rationality.

A number of recent papers study the effect of dual sourcing on a production system. Fong et al. ([9, 2000]) as well as Rudi ([14, 1999]) analyze the effect of dual sourcing on the optimal inventory. Fong et al. study an inventory system with a choice between two supply options having different lead-times. They allow for normally distributed demand and Erlang-distributed supplier lead-times. In contrast to Fong et al., Rudi does not assume that orders for both supply options are placed concurrently and considers sequential decision-making. He uses a two-stage stochastic linear programming formulation to analyze the split between make-to-stock and assemble-to-order. Both papers assume that all cost parameters are deterministic. Van Mieghem ([19, 1999]) studies the option of subcontracting to improve financial performance by analyzing a

competitive stochastic investment game with recourse between a retailer and a subcontractor. He considers a two-stage stochastic model of the investment decision process of the two firms for three different contract types. Notice that the setting of his model can be re-interpreted in our framework such that the subcontractor can play the role of a spot market and the retailer owns a component production facility.

The recent paper by Cohen and Agrawal ([6, 1999]) considers a stochastic dynamic program formulation to analyze the trade-off between long- and short-term contracts. Their model, however, does only allow for the usage of either contract at a time, i.e., a mixed strategy is not considered as is the case in our paper. Seifert and Thonemann ([15, 2000]) consider a situation similar to the one we consider here, in that they determine the optimal mix between long-term contract and spot market under the assumption that demand and spot market price follow a Bivariate Normal distribution. Contrary to our paper, they do not model the supply side.

A recent example of finance research dealing with the role of inventories and the economic theory of storage is the work of Routledge et al. ([13, 2000]), who devise a one-factor equilibrium model of forward prices for commodities with Markovian demand shocks as stochastic input factor. This formulation allows the authors to provide economic insights into the effect of inventories on spot prices, but does not consider the effect of spot markets on supply chain operations.

In this paper, we contribute to the existing literature by considering the optimal procurement strategy obtained by choosing between long-term contracts, spot market, and a mix between the two. Our paper defines a general set of contracts that allow channel coordination in a simple setting in the presence of a spot market. The model is broad enough to cover a dual sourcing problem, or even a one source of supply with lost sales. We are careful in making only very broad assumptions about the spot market price-quantity uncertainty and the demand distribution when formulating our models. We also consider a risk-sharing agreement between the retailer and her

long-term supplier with a linear unit price penalty for under-using supplier reserved production quantity. Using analytical results and numerical computation, we provide managerial insights into the chosen reservation level and the cost savings that the retailer can achieve by following the optimal procurement strategy, which always is a mixed strategy between long-term and spot market channel.

3 The Model

We consider a retailer operating in an assemble-to-order mode who needs to satisfy a non-negative random end-demand d in every period. We assume that d has a continuous density function, f_d, and we call μ_d the mean demand and σ_d its standard deviation. The retailer has the choice of purchasing components (potentially a critical component) from the spot market and/or negotiating a long-term deal with a particular supplier. The long-term contract is defined by two quantities, an amount of units K that the supplier will have available in every period and a pricing scheme denoted by $P(d, K)$, that determines the unit price that the retailer will pay the supplier when d units are bought out of the K reserved initially.

If both players agree to enter into such a contractual relationship, i.e., decide on a unit pricing scheme $P(d, K)$ and a reservation quantity K, then as the end-customer demand d is realized, two cases are possible. First, if the demand d is smaller than the reserved amount K, then the retailer will meet her entire demand from the long-term supplier and pay the associated unit price $P(d, K)$. Second, if the demand d is bigger than K, then the retailer procures K units from the long-term supplier at a unit price $P(K, K) = p$ and buys the remaining $d - K$ units on the spot market at the current price π. Notice that we assume here that the retailer will always meet all her demand d. This is possible as the spot market is not capacity constrained. The retailer will sell the goods to her end-customer at a fixed unit price b. The unit production cost for the

supplier is m. All units remaining in the system are salvaged at unit price s.

Contrary to the long-term contract, the unit price on the spot market, π, is random. Evidently, the spot market price, π, will be affected by the overall demand for the end product, for which the particular component is being used. Since the retailer's demand, d, will most probably depend on the industry-wide demand, we assume that the spot market unit price is correlated with the overall (end) demand the retailer receives.

As the demand is realized before the spot market is cleared and the price π is revealed, the conditional expectation of the unit spot market price given a certain realization of total demand, d, is relevant to our analysis. We denote this quantity by $\pi(d) = \mathrm{E}[\pi|d]$ and we let $\bar{\pi}$ be the expected spot market unit price, $\bar{\pi} = \mathrm{E}[\pi] = \int_0^\infty \pi(u) f_d(u) du$. Furthermore, we assume $\pi(d)$ to be a continuous function on $[0, \infty)$. The cost of ordering ζ units from the spot market when the total end-demand is d, is thus equal to $\Pi(\zeta, d) = \zeta \cdot \pi(d)$. The following assumption related to the different costs involved will be assumed to hold throughout the paper.

A1 $s < m < p < \bar{\pi}$ and for all u, $s <= \pi(u)$.

The inequality $p < \bar{\pi}$ suggests that the spot market is not the cheapest way of buying the goods, in expectation. The unit price charged by the long-term supplier when the retailer buys all the K units reserved is denoted by p and it is strictly smaller than the mean spot market price. The other inequalities are self-explanatory.

3.1 Shortage Cost and Spot Market

It is quite evident that one could interpret the spot market price as a general shortage cost. Indeed, consider the reserved production quantity, K, as the safety stock in a newsvendor-type model. In this interpretation, the spot price is equivalent to the

(underage) cost per unit of demand in excess of K, that cannot be met from the safety stock. (Notice that since we focus on a make-to-order environment, the holding cost does not enter into the equations.) Let H_{SM} be the expected procurement cost to paid on the spot market, we have that $H_{SM}(K) = \mathrm{E}\left\{[d-K]^+\pi(d)\right\}$. Having said that, there still is a difference between charging a penalty cost and buying the units from the spot market. In the former case, no revenues are generated, whereas in the latter, the retailer is selling the units bought on the spot market and is thus generating revenues of b per unit. However, by defining a slightly modified spot market curve one can easily show that both versions of the problem are equivalent. To this end, note that by charging a penalty cost z for every unit of unfilled demand the retailer will only choose to not fulfill a unit of demand if its effective cost $z+b$ is smaller than the spot price, $\pi(d)$. Therefore, the new cost function $\tilde{H}_{SM}$ is given by,

$$
\begin{aligned}
\tilde{H}_{SM}(K) &= (b+z)\,\mathrm{E}\left\{[d-K]^+ 1_{\pi(d)>z+b}\right\} + \mathrm{E}\left\{\pi(d)[d-K]^+ 1_{\pi(d)\leq z+b}\right\} \\
&= \mathrm{E}\left\{\min(\pi(d),(b+z))[d-K]^+\right\}
\end{aligned}
$$

By defining for all d, $\tilde{\pi}(d) = \min(\pi(d),(b+z))$, the version with shortage cost becomes similar to the one we are considering here, namely assuming that the entire end-demand d has to be satisfied using both means of supply as we explained earlier. In a similar fashion, our formulation can also account for a situation without a spot market, i.e., where the retailer can procure only via the long-term contract and faces a (possibly non-constant) shortage penalty per unit of unsatisfied demand. In summary, as this part is concerned, our paper generalizes other formulations in the literature.

3.2   Symmetric Information

The information available to both players as we described it earlier will be an important factor in determining whether they will reach an equilibrium that is efficient for the total system. More precisely, symmetric information between the two agents happens

8

to be a necessary condition to create an over-all efficient system. However, in this paper inefficiency due to asymmetric information will not be our primary concern. As we will argue later that even under symmetric information regarding the demand distribution and the spot market, inefficiency of the system is inherently present unless an additional condition is enforced which relates to the nature of the contract between the retailer and the supplier. We refer the reader to Tsay ([16]) for a similar but more complete discussion of this issue.

4 Channel Coordinating Contracts

In what follows we will first analyze the centralized case and compare it to the decentralized case under symmetric information. As we do that, we will define the pricing scheme specified in the long-term contract. Our definition determines a class of contracts that is quite general as it comprises many well-known contract types. We then study necessary and sufficient conditions for the system to be efficient based on our general class of contracts.

4.1 Centralized Case

To establish a benchmark, we first consider the centralized problem, where for instance the supplier and the retailer are owned by one firm. Quite evidently, there exists no contract between the supplier and the retailer in the centralized case and the incentives of both players are automatically aligned. Therefore, the profit of the centralized case is the maximum one can achieve for the total system.

The centralized firm maximizes its expected profit, $\Pi_{cc}(K)$, given by,

$$\Pi_{cc}(K) = b\mu_d + s\,\mathrm{E}[K - d]^+ - H_{SM}(K) - Km$$

where H_{SM} is the expected cost of goods bought from the spot market:

$$H_{SM}(K) = \int_K^\infty (u - K)\pi(u)f_d(u)du. \tag{1}$$

Thus the optimal production quantity, K_{cc}^*, is defined by

$$K_{cc}^* = \arg\max_K\{\Pi_{cc}(K)\} \tag{2}$$

The assumption A1 stated earlier will guarantee the existence and uniqueness of K_{cc}^*, which is generally referred to as the first best solution.

Proposition 1 *Under A1, K_{cc}^* exists and is the unique solution of*

$$H'_{SM}(K_{cc}^*) = -m + sF_d(K_{cc}^*) \tag{3}$$

K_{cc}^* *is the first best solution.*

Proof. $\Pi'_{cc}(0) = H'_{SM}(0) - m = \bar\pi - m > 0$, and $\lim_{K\to\infty} \Pi'_{cc}(K) = -(m - s) < 0$. Therefore, there exists at least one finite value, K_{cc}^* where Π_{cc} reaches its global maximum. In order to show that this value is unique we prove that Π_{cc} is concave in K under *A1*. As a matter of fact,

$$\Pi''_{cc}(K) = -\pi(K)f_d(K) + sf_d(K) > 0$$

By the same token we just showed that H_{SM} is convex in K. ∎

4.2 Decentralized Case

We come back to the decentralized case, where the supplier and the retailer are two independent entities. Each of them is maximizing her own profit but still has the common interest of achieving a high total system profit.

4.2.1 The Supplier

Our primary focus in this paper is on the retailer. Therefore, from the supplier's side we chose to only model the supply contract with the retailer. The supplier's revenues are then generated only from the long-term contract, and other sources of revenues such as, possibly, the spot market are omitted. If we denote by $H_{LT}(K)$ these revenues given a contract reservation quantity, K, we have that

$$H_{LT}(K) = \int_0^K uP(u,K)f_d(u)du + Kp(1 - F_d(K))$$

and the associated profit, $\Pi_s(K)$, is then given by

$$\Pi_s(K) = H_{LT}(K) + s\,\mathrm{E}\,[K - d]^+ - mK$$

We assume that the following conditions hold true regarding the long-term pricing scheme:

A2 $P(u,K)$ is increasing and concave in K, and

$$\left.\frac{\partial P(u,K)}{\partial K}\right|_{u=K} \leq \frac{p-s}{K} \tag{4}$$

The fact that the unit price $P(u,K)$ increases with K, means that the retailer will be charged more for buying the same amount of units when the initial reservation amount is higher, or equivalently when the amount of under-used production quantity is higher. This assumption will make sure for instance that the retailer does not reserve a higher amount than needed without being penalized. The concavity assumption reflects a certain economy of scale, e.g., stemming from reduced set-up costs or improved production efficiency due to the higher contract volume. The last part of assumption A2 given by equation (4) is a technical condition. This way, A2 is enough to ensure the concavity of the supplier's total expected revenues as we will see later.

Before we continue our analysis of the supplier's profit function, Π_s, we briefly indicate the generality of the contract pricing scheme under consideration in providing examples of well-studied supply contracts that are but special cases.

Examples of Supply Contracts Satisfying A2 Many of the supply contracts in the literature follow A2. Consider, for instance, *a fixed quantity commitment contract* where the retailer pays a fixed unit price p for ordering a constant amount K in every period. In this case, there exists no flexibility in terms of the units bought and the retailer salvages the leftover units, $[K - d]^+$, at a unit value, s. Since in our formulation the supplier keeps the salvage value, the contract price has to be adjusted accordingly in order to have an equivalent setting to ours. Thus, the effective unit price, $P_e(d, K)$ (to be compared to our $P(d, K)$ defined above), that the retailer has to pay in order to satisfy a demand level d equals

$$P_e(d, K) = s + \frac{p - s}{d} K \quad \text{for } K \geq d$$

which is linear increasing in K for all d and hence follows A2. Another well-known contract is the *quantity flexibility contract* defined for instance in Tsay ([16]), where the contract specifies a fixed unit price, p, and the retailer needs to buy a minimum amount $Q_{\min} = \alpha K$, with $\alpha \in (0, 1)$ while the supplier ensures the availability of $Q_{\max} = K$ units. Again, assuming that the retailer can salvage any leftover units, $[\alpha K - d]^+$, at a unit value s, the effective price

$$P_e(d, K) = \begin{cases} s + \frac{p-s}{d} \alpha K & \text{for } K \geq d/\alpha \\ p & \text{for } d <= K <= d/\alpha \end{cases}$$

is increasing and concave in K for all d and again follows A2.

Finally, consider the *buy-back contracts* defined for instance by Pasternack ([12]) and Donohue ([7]). With this type of contracts, there exists a flexibility related to the number of units bought from the supplier, in the following sense. The retailer buys K units at a price p, however, after the demand d is realized, $[K - d]^+$ units are returned to the supplier at a unit price b, with $s <= b < p$ (s is again the salvage value). The effective unit price is thus

$$P_e(d, K) = b + \frac{p - b}{d} K \text{ for } K \geq d$$

that is, $P_e(d, K)$ is linearly increasing in K and again follows A2.

It should then be clear that the class of contracts defined by A2, is quite broad. Now, we use A2 to easily verify that H_{LT} is well-behaved:

Lemma 1 *Under A2, H_{LT} is increasing and concave in K.*

Proof. The proof follows directly from A2 and using the second derivative of $H_{LT}(K)$

$$H_{LT}''(K) = \int_0^K u \frac{\partial^2 P(u, K)}{\partial K^2} f_d(u) du + \left. \frac{\partial P(u, K)}{\partial K} \right|_{u=K} K f_d(K) - p f_d(K))$$

∎

Furthermore, A1 and A2 together ensure that Π_s is concave and that it has a unique maximum.

Proposition 2 *Under A1 and A2, Π_s is concave and admits a unique maximum at K_s^* such that*

$$H_{LT}'(K_s^*) = m - s F_d(K_s^*) \tag{5}$$

K_s^* *is non-zero but could be infinite in some cases.*

The above proposition provides us with the supplier's maximum profit for any contract pricing scheme that follows assumptions A1 and A2.

4.2.2 The Retailer

Similarly, the retailer maximizes her expected profit, $\Pi_R(K)$, given by,

$$\Pi_R(K) = b\mu_d - H_{LT}(K) - H_{SM}(K) \tag{6}$$

Figure 1: The effective unit price $\omega_e(d, K)$

Notice that $\Pi_R(K)$ is governed by the retailer's total expected procurement cost, $G(K)$, with $G(K) = H_{LT}(K) + H_{SM}(K)$ or, combining the two terms,

$$G(K) = \int_0^\infty \omega_e(u, K) u f_d(u) du$$

where $\omega_e(d, K)$ is the effective unit price for any amount of demand d non-negative, when the reserved amount is K.

$$\omega_e(d, K) = \begin{cases} [\pi(d)d - K(\pi(d) - p)]/d & \text{for } K < d \\ P(d, K) & \text{for } K \geq d \end{cases}$$

By plotting the graph of this effective unit price as a function of K when d is fixed and assuming $P(d, K)$ to follow assumption A2, one easily notices that $\omega_e(d, K)$ is unimodal admitting a minimum at d if $\pi(d)$ is greater than or equal to p. (see Figure 1). Therefore, as $\bar{\pi} > p$ (by A1) one would expect this characteristic to be preserved by taking the expected value of the total cost, $G(K) = \mathrm{E}[d \cdot \omega_e(d, K)]$, at least for well behaving demand distributions or for many pricing schedules, $P_e(d, K)$ and $\pi(d)$ (e.g., with minor conditions for the three examples satisfying A2 mentioned above). We will assume the following to hold:

A3 G is unimodal.

We note here that we can prove in many cases that Π_R admits a maximum (or equivalently that G admits a minimum), for instance if H_{LT} is *not* bounded (which is the case in the different type of contracts we discussed earlier). However, the uniqueness of this maximum, although it appears to hold numerically for many common demand distributions (such as normal and exponential), is a difficult question to prove analytically.

Proposition 3 *Under A2 and A3, Π_R is unimodal and admits a unique, finite maximum at $K_R^* > 0$, such that*

$$H'_{LT}(K_R^*) + H'_{SM}(K_R^*) = 0 \tag{7}$$

Proposition 2 together with Proposition 3 show that both the retailer and the supplier will try to maximize their own profit by pushing for a contract based on their optimal reservation quantity (K_S^* or K_R^*). As these two desired contract quantities are not necessarily equal, a negotiation between the retailer and her supplier should take place in order to determine the contract quantity. As both players' profit functions are unimodal, it is easy to see that a Nash equilibrium will always be reached, i.e., a contractual agreement can be reached and a long-term contract will be established between the two players. Proposition 4 also asserts that for all type of contracts that verify the previous conditions, both means of supply, the long-term contract and the spot market, will be used as procurement channels by the retailer ($K^* > 0$).

Proposition 4 *For any pricing scheme and demand distribution, defined by the assumptions A1-A3, there exists a Nash equilibrium between the retailer and the supplier, defined by*

$$K^* = \min(K_S^*, K_R^*)$$

and K^* is strictly positive.

We are now ready to "compare" the system profits resulting from the Nash equilibrium to the ones achieved in the centralized system, i.e., the maximum system profits:

Theorem 1 *Consider a long-term contract that follows conditions A1-A3. The contract is system-efficient if and only if*

$$H'_{LT}(K^*_{cc}) = m \tag{8}$$

Where K^*_{cc} is defined by (3). In this case,

$$K^* = K^*_S = K^*_R = K^*_{cc}.$$

Otherwise, the system is strictly inefficient and the retailer will always under-consume, i.e.,

$$K^* < K^*_{cc} \tag{9}$$

Proof. By the uniqueness of the solutions of (7) and (5) in addition to the condition given by (8) we conclude that $K^*_S = K^*_R = K^*_{cc} = K^*$. On the other hand, since the functions $-H'_{SM}$ and H'_{LT} intersect at a unique point K^*_R, we conclude, as both functions are non-increasing, that either $H'_{LT}(K) > -H'_{SM}(K)$ for K such that $K^*_R \geq K$ and $H'_{LT}(K) < -H'_{SM}(K)$ for $K > K^*_R$, or the other way around. Therefore, as K^*_S and K^*_{cc} are defined by the intersection of H'_{LT} (resp. H'_{SM}) with $m - sF_d(K)$, which is decreasing in K, we necessarily have either $K^*_S \geq K^*_{cc} \geq K^*_R$ or $K^*_R \geq K^*_{cc} \geq K^*_S$. This leads in both cases to K^* being smaller than K^*_{cc}. ∎

Based on Theorem 1, the inequality (9) holds, and hence the system is inefficient for any supply contract and demand distribution of the type we are considering in this paper (including the contracts we discussed earlier such as fixed quantity commitment, quantity flexibility, and buy back) unless equation (8) is fulfilled. Thus, the double marginalization phenomenon is present resulting interestingly always in under-consumption by the retailer with respect to the centralized system. Only if one can

design a contract such that (8) holds will the system be efficient and the total system profit is maximized. The optimal production/contract-reservation quantity will in this case necessarily be equal to K_{cc}^*. We will observe here that all the contracts that we introduced earlier could potentially be used to reach system efficiency and again one will only need to solve for (8). We should observe that what makes some contracts more interesting than others is the level of flexibility they allow. For instance, the fixed quantity commitment contract, i.e., which is based on a fixed quantity and fixed price, will enable system efficiency if and only if the price charged by the supplier is equal to the marginal production cost, $m = p$, which would lead to the supplier earning zero profit. Other types of contracts are more realistic as it is possible for both the supplier and the retailer to achieve positive profits in the efficient case by using these contracts.

5 Linear Risk-Sharing Contract

With a fixed quantity commitment contract, commonly used in the literature, the retailer decides on the order quantity K via the long-term contract and she commits to purchasing it before the actual demand is known. Any leftover remaining inventory is salvaged at the unit salvage value, s. Obviously, this approach leaves the retailer with a high risk due to the demand uncertainty.

The contract pricing scheme that we propose in the following represents a mechanism of sharing this risk between both parties (while still allowing for channel coordination). More specifically, we assume that the contract's pricing scheme, $P(u, K)$, is linear in the retailer's order quantity. Thus, it is completely determined by two parameters, to which we will refer as p and c. The first one, $p = P(K, K)$, determines the unit price when the reserved quantity is fully used by the retailer, and the second one, c, with $p < c <= 2p - s$, defines the maximum amount per unit that the retailer

has to pay when "under-buying". We thus have

$$P(u, K) = -\frac{c-p}{K} \cdot u + c \quad \text{for all } u <= K \tag{10}$$

and we obtain that H_{LT}, the expected cost to the retailer for purchases from the long-term supplier, equals

$$H_{LT}(K) = \int_0^K u \left(-\frac{c-p}{K}u + c \right) f_d(u)du + Kp(1 - F_d(K)) \tag{11}$$

Remark: Note that if $c = p$ then the pricing scheme is constant and similar to the fixed quantity commitment contract defined in the first part of the paper. The only difference is that now the retailer is not forced to purchase K units regardless of her end demand. This implies that the entire risk due to the end demand resides with the supplier. Hence, in order to ensure a "sharing" of the risk associated with the demand uncertainty we assume that $c > p$.

Before going further, we should verify that this linear unit pricing scheme belongs the class of contracts we discussed earlier mainly that A2 holds.

Lemma 2 *The linear pricing scheme verifies A2.*

Proof. We have

$$\frac{\partial P(u, K)}{\partial K} = \frac{c-p}{K^2} \cdot u \quad \text{for all } u <= K$$

and

$$\frac{\partial^2 P(u, K)}{\partial K^2} = -\frac{c-p}{K^3} \cdot u \quad \text{for all } u <= K$$

Furthermore,

$$\frac{\partial P(u, K)}{\partial K}\Big|_{u=K} = \frac{c-p}{K} <= \frac{p-s}{K}$$

by assumption on c. ∎

Concerning assumption A3, we note that it is equivalent to say that Π_R admits a finite unique maximum. Now that we are dealing with a specific linear pricing scheme, A3 is not anymore a condition on the pricing scheme but should be viewed as an assumption on the demand distribution. What we were able to show (next Proposition) is that Π_R admits a maximum under this pricing scheme (this is not trivial as H_{LT} in this case is bounded). However, the uniqueness of the maximum seems very difficult to show analytically, although it seems to hold numerically for most known distributions. For the sake of generality we keep A3 as an assumption.

Proposition 5 *Consider a linear pricing scheme and assume that A1 holds. If*

$$\int_0^\infty u^2 \pi(u) f_{dm}(u) du < \infty \tag{12}$$

then Π_R admits a finite maximum.

Proof. We consider the first derivative of Π_R, $\Pi'_R(K) = H'_{LT}(K) + H'_{SM}(K)$. Note first that from A1, $\Pi'_R(0) = \bar{\pi} - p > 0$. It is enough to show that Π'_R is non-positive for K big enough. From the defining equations for $H_{LT}(K)$ and $H_{SM}(K)$ together with the fact that $\sigma_d < \infty$, we can easily see that,

$$K^2 H'_{LT}(K) = (c - p) \int_0^K u^2 f_{dm}(u) du + p K^2 [1 - F_{dm}(K)]$$

hence

$$0 <= K^2 H'_{LT}(K) - (c-p) \int_0^{+\infty} u^2 f_{dm}(u) du <= (c-p) \int_K^{+\infty} u^2 f_{dm}(u) du + p \int_K^{+\infty} u^2 f_{dm}(u) du$$

So we conclude that $K^2 H'_{LT}(K) \to (c-p) \int_0^{+\infty} u^2 f_{dm}(u) du < \infty$ as K goes to infinity. On the other hand, from (12) we get

$$0 <= -K^2 H'_{SM}(K) = K^2 \int_K^\infty \pi(u) f_{dm}(u) du$$

$$<= \int_K^\infty u^2 \pi(u) f_{dm}(u) du \to 0 \quad \text{as } K \to \infty$$

19

Thus, $K^2\Pi'_R(K) = -K^2(H'_{SM}(K) + H'_{LT}(K)) \rightarrow -(c-p)\int_0^{+\infty} u^2 f_{dm}(u)du < 0$. Hence there exists $\hat{K}$, such that $\forall K \geq \hat{K}$

$$\Pi'_R(K) < 0$$

which complete our proof. ■

The condition given by equation (12) is assumed to hold for technical reasons that becomes clear in the proof of the previous Proposition. But one should note that this condition has a little effect on the type of demand distributions that one chooses to use, especially when assuming that the standard deviation of the demand is already finite.

5.1 Channel Coordinating Contract with Risk Sharing

Now that we made sure that such type of contracts belong to the class we studied in the previous section, we are interested in studying whether it could ensure system-efficiency. Based on Theorem 1, we only need to verify that the linear pricing scheme as defined above verifies equation (8) to conclude that the linear pricing scheme leads to channel coordination.

Proposition 6 *Under a linear pricing scheme as described above, system-efficiency is reached when one chooses p and c such that $p \in (p_1, \bar{p})$ with $p_1 < \bar{p}$ and*

$$c(p) = p + \frac{p_1 - s}{\bar{p} - p_1}(\bar{p} - p)$$

where,

$$p_1 = \left[s + K_{cc}^{*2} \frac{m - sF_d(K_{cc}^*)}{\int_0^{K_{cc}^*} u^2 f_d(u)du} \right] * \left[1 + K_{cc}^{*2} \frac{(1 - F_d(K_{cc}^*))}{\int_0^{K_{cc}^*} u^2 f_d(u)du} \right]^{-1}$$

and

$$\bar{p} = \min(\bar{\pi}, p_2), \qquad with \ p_2 = \frac{m - sF_d(K_{cc}^*)}{(1 - F_d(K_{cc}^*))}$$

Again K_{cc}^* is given by equation (3).

Proof. We know that in order to achieve the same production (and the same profit) level as in the centralized case, we need the following to hold

$$H'_{LT}(K^*_{cc}) = (c - p) \int_0^{K^*_{cc}} \frac{u^2}{K^{*2}_{cc}} f_d(u) du + p\left(1 - F_d(K^*_{cc})\right) = m - sF_d\left(K^*_{cc}\right)$$

Thus, for any fixed p, $c = c(p)$, which achieves the centralized profits in a decentralized system, will be given by

$$c(p) = \frac{mK^{*2}_{cc} - sK^{*2}_{cc}F_d\left(K^*_{cc}\right) + p\left[\int_0^{K^*_{cc}} u^2 f_d(u)du - K^{*2}_{cc}\left(1 - F_d(K^*_{cc})\right)\right]}{\int_0^{K^*_{cc}} u^2 f_d(u)du}$$

For there to be a risk sharing in the sense that we defined above, we need to enforce that $p < c <= 2p - s$. It is clear that

$$c(p) = p + K^{*2}_{cc} \frac{m - sF_d\left(K^*_{cc}\right) - p\left(1 - F_d(K^*_{cc})\right)}{\int_0^{K^*_{cc}} u^2 f_d(u)du}$$

$$= K^{*2}_{cc} \frac{m - sF_d\left(K^*_{cc}\right)}{\int_0^{K^*_{cc}} u^2 f_d(u)du} + p\left[1 - K^{*2}_{cc} \frac{\left(1 - F_d(K^*_{cc})\right)}{\int_0^{K^*_{cc}} u^2 f_d(u)du}\right]$$

Hence, as $c(.)$ is linear in p, there exists two constants p_1 and p_2 such that for all $p \in (p_1, p_2)$, $p < c < 2p - s$. p_1 (resp. p_2) is the solution of $c(p) = p$ and $c(p) = 2p - s$. Since $c(0) > 0$ we necessarily have $p_1 < p_2$. In the case where $\pi(u) = \bar{\pi}$ for all u, Equation (3) gives us that $p_2 = \bar{\pi}$. ∎

The previous result is interesting as it shows that the linear risk sharing contracts can lead to system efficiency, when the constants p and c are well chosen. However, another important characteristic of such contracts, is the fact that they enable for a wide range of profit split between the retailer and the supplier. Every value of p in the interval $(p_1, \bar{p})$ defined above, corresponds to a certain profit split. We observe that both profits Π_S and Π_R are linear in p and by definition they sum to $\Pi_{cc}(K^*_{cc})$, as we can see in Figure 2. This could also allow for a fair total profit sharing based on the negotiation power of each player. By conducting a sensitivity analysis on the previous results (derived numerically), and studying the effect of an increase in the volatility of the demand, we noticed the following. As expected all the profit decrease with an

21

Figure 2: Profit split between the retailer and the supplier for an efficient system

Figure 3: Expected profit vs. changes in demand variability

increase in demand variability. However two interesting points are worth mentioning. The first one is that the optimal production quantity K takes values around the mean demand and is not much affected by this change of variability. The second concerns the profit of the retailer and the supplier. We observe that the supplier gets highly affected by the change in the standard deviation of the demand distribution, (see Figure 3) where the retailer profit decreases in a very moderate way. The retailer is hedging, through the dual sourcing, against the demand variability. The spot market is bringing here an important advantage to the retailer. Due to the way the problem is structured and especially since the long-term contract is the preferred means of supply, a natural question that arises is to ask how often the spot market will be used. Indeed, the higher the amount of reserved quantity, K, the less frequent the supplier purchases on the spot market. We will address this point in our numerical analysis later in the paper.

5.2   Retailer's Problem Revisited

Since our model does not allow the supplier to participate in the spot market and the supplier's decision problem is thus not our primary concern, we are going to focus in the remainder of this paper primarily on the retailer's problem, given that a contract with linear pricing scheme has been imposed. By this token, the main objective of this section is therefore to study the optimal procurement strategy from a retailer's point of view when a supply contract with linear pricing scheme is being used. The retailer will choose the expected minimum cost procurement channel mix between the spot market channel and the reservation quantity level of the long-term reservation quantity channel (including the possibility of one channel alone). Note again that for the retailer, it is equivalent to consider in the analysis either the cost or the profit function. We chose in this section to work with the former. Finally, the only assumption that we are keeping here is A1 related to the different fixed unit costs.

We now state the main result (based mainly on Proposition 5) that indicates that the spot market is valuable under the fairly general conditions stated earlier. We will also characterize the values of the (optimal) contracted quantity K.

Theorem 2 *Consider a linear pricing scheme, and let K_R^* be the value of K that maximizes the retailer's expected profit.*

- *Suppose that equation (12) holds together with A1. If the demand is random on the real line, then the retailer's profit function $\Pi_R(K)$ admits a maximum at K_R^* such that $0 < K_R^* < \infty$.*

- *If the total demand d has finite support $[0, X]$, with $0 < X < \infty$, and A1 holds then the retailer's profit function $\Pi_R(K)$ admits a maximum at K_R^* such that $0 < K_R^* < X$.*

It is important to recognize that $K_R^* = X$ ($X = \infty$ in the first case and X finite in the second) corresponds to the case when the retailer uses the long-term supplier exclusively and, on the other extreme, $K_R^* = 0$ is equivalent to the retailer purchasing only from the spot market. A value of K_R^* equal to an interior point of the support of d means that a mixed strategy is optimal. Note that if the demand is deterministic, there will clearly be no mix, i.e., the retailer will select the supplier with the cheapest unit price and hence, in our case, will go with the long-term contract as Assumption A1 holds. Note the following remarks:

- By stating that K_R^* never takes on one of the extreme points of the domain, Theorem 2 indicates that for a wide range of demand distributions a mixed procurement strategy is always optimal for the retailer, i.e., she will always choose to meet part of her demand with the long-term supplier and the remainder via the spot market.

- The intuition behind Theorem 2 is that the realizations of the higher demand values occur only with a small probability, and thus the retailer is better off paying a high price on the spot market for these "rare" ("low probability") events than to almost always pay a penalty for under-using the supplier's reserved quantity. Nevertheless, it is surprising that a mixed strategy is optimal, regardless of the spot market price (e.g., even for a very high one in comparison to the long-term contract pricing schedule), and for a fairly general class of demand distributions.

5.2.1 Example: Exponential Demand Distribution

To further illustrate the findings of Theorem 2, we provide a specific example assuming that the demand for the retailer's products follows an exponential distribution ($f_d(u) = \lambda \exp\left(-\lambda u\right)$) and that the spot market price is independent of the retailer's order, i.e., $\pi(u) = \pi$ and thus $\bar{\Pi} = \bar{\pi}\mu_d$. Furthermore, in order to make the calculations simpler

to follow, we put $s = 0$ and $c = 2p$ (which may not result in system efficiency). Under these assumptions, the retailer's objective becomes:

$$\max_{K}\Pi_R\left(K\right) = \frac{b}{\lambda} - 2p\frac{e^{-\lambda K} - 1 + \lambda K}{\lambda^2 K} - \pi\frac{e^{-K\lambda}}{\lambda}$$

It is then easy to verify that if Assumption A1 applies, then $K_R^* \in (0, \infty)$, that is, a mixed procurement strategy is optimal. We can furthermore establish that the optimal level of reservation quantity is less than the mean value of demand when the spot market is not too expensive in expectation. Or, more precisely,

Proposition 7 *If the retailer's demand follows an exponential distribution and the spot market price is independent of the retailer's order then,*

$$K_R^* < \mu_d \quad \text{if and only if} \quad 1 < \frac{\pi}{p} < 2\left(e - 2\right) \tag{13}$$

Notice that $2\left(e - 2\right) \approx 1.4$, that is the optimal level of quantity reservation, K_R^*, is less than the mean value of demand, μ_d, even if the expected price premium on the spot market price with respect to the minimum price via the long-term contract is as high as 40%. Furthermore, since for an efficient contract $K^* = \min\left(K_s^*, K_R^*\right)$ Proposition 7 indicates that K^* may well lie below the mean value of the demand, μ_d, implying that both means of supply may be used by the retailer on a frequent basis in the efficient contract.

5.2.2 Numerical Results & Managerial Insights

Consider the retailer's expected procurement cost, $G\left(K\right)$ only. Moreover, we assume for the sake of simplicity that the contract with the long-term supplier specifies a linear penalty for under-ordering, salvage value $s = 0$, and we model demand d and spot market unit price π as a Bivariate Normal (BN) distribution with correlation ρ: $(d, \pi) \sim BN\left[\mu_d, \sigma_d^2, \mu_\pi, \sigma_\rho^2, \rho\right]$. For the computations, we use that $\pi\left(d\right) = E\left[\pi|d\right] = \mu_\pi + \rho\frac{\sigma_\pi}{\sigma_d}\left(d - \mu_d\right)$.

Figure 4: Example of $G(K)$, the expected total procurement cost

Optimal Level of Reservation quantity, K^*   Figure 4 illustrates the impact of varying the long-term quantity K on the total cost $G(K)$. Note that this example reflects the remarkable result that the optimal value of long-term quantity K^* is in most cases close to the mean of the demand faced by the retailer (rather than greatly exceeding the mean demand level, as we might intuitively have expected without our analysis leading to Theorem 2). This implies that both procurement channels will be used on a regular, almost equal basis. According to Proposition 7 and our numerical studies, this result holds true, even when the spot market is more expensive than the long-term supplier. Obviously, as the spot market becomes more expensive with respect to the long-term supplier, the optimal value of reserved long-term quantity K^* increases.

Optimal Total Cost, $G(K^*)$ Notice that the cost reduction due to the existence of the spot market can be very significant. In the example of Figure 4, the retailer saves around 40% in expected total cost with respect to exclusive procurement via long-term contract, which is still the current practice in many corporations. In general, the "value" of the mixed strategy can be computed as the difference between the expected cost associated with an exclusive long-term supplier relationship, $G(\infty)$, and the minimum expected cost under the optimal (mixed) procurement strategy, $G(K^*)$. (Recall

27

Figure 5: Robustness of $G\left(K^{*}\right)$ for $BN\left(40, \sigma_{d}, 4, 1\right)$ and $p = 2$

that $G\left(\infty\right) = H_{LT}\left(\infty\right) = c\mu_{d}$, which is for many demand distributions independent of the standard deviation). The optimal procurement mix also yields significant savings when compared to exclusive sourcing on the spot market.

We can also see that the unimodality assumption A3 is verified numerically for the example of Figure 4.

Robustness of Results   Most importantly, from a practitioner's point of view, the total expected cost, $G\left(K\right)$, is relatively stable in the neighborhood of the optimal quantity reservation level, K^{*}. This means that a mixed strategy captures great savings even when the quantity reservation level is not at its optimal level, e.g., due to errors in demand forecast or cost parameter estimates. For instance, Figure 5 shows that even for a relatively expensive spot market, where the average spot price, μ_{π}, equals the most expensive price via the long-term contract, $c = 2p = 4$, the difference in expected total cost when the optimal long-term reservation quantity is chosen assuming $\mu_{d} = 40$ and $\sigma_{d} = 10$, $G\left(K_{-}\right)$ and the expected total cost when the reservation quantity level is chosen optimally for each value of standard deviation of demand, $G\left(K^{*}\right)$, is minimal. In addition, we plotted the absolute and relative difference between the two

costs (denoted by Delta and % Delta in Figure 5) and the optimal value of long-term reservation quantity, K^*, itself.

Impact of Demand Uncertainty   We can also study the effect of demand uncertainty, modeled as standard deviation, σ_d, on the optimal procurement strategy. For instance, the optimal level of reservation quantity, K^*, increases with σ_d in most cases, in particular if $\mu_\pi \approx p$. However, K^* decreases with demand uncertainty (standard deviation) for $\mu_\pi \approx p$, i.e., the spot market price is close to the lowest price available from the long-term supplier.

Not surprisingly, the corresponding total cost, $G(K^*)$, increases with demand uncertainty for most cases. Only if the spot price π is very close to p, the lowest price via the long-term contract, and the demand and the spot price are strongly negatively correlated (e.g., $\rho < -0.5$) may $G(K^*)$ decrease with σ_d. In general though, $G(K^*)$ increases with σ_d in a fairly robust way (see Figures 5 and 7). Hence, the value of the mixed procurement strategy, modeled as the difference between the optimal solution, $G(K^*)$, and $H_{LT}(\infty)$, corresponding to the cost associated with the current practice for most corporations, which only use a single long-term supplier, decreases with demand uncertainty, since $H_{LT}(\infty) = c\mu_d$, which is for many distributions, in particular for Binormal Distributions, independent of the value of the standard deviation of demand.

Impact of Spot Market Price and Uncertainty   As the spot market becomes more expensive with respect to the long-term contract price schedule, the retailer prefers to purchase more from the long-term supplier. Our numerical study confirms that both K^* and $G(K^*)$ are increasing with (relative increases in) the mean spot market price, μ_π. We do not show any graphs regarding the impact of μ_π, since the results confirm common intuition.

Figure 6: Impact of change in μ_d on K^* and $G\left(K^*\right)$ for $BN\left(\mu_d, \mu_d/2, 4, 1, \rho\right)$ and $p = 2$

Impact of Coefficient of Variation of Demand A surprising finding was, however, that the optimal level of reservation quantity, K^*, and the associated expected procurement cost, $G\left(K^*\right)$, increase almost linearly with the mean value of demand, μ_d, for a constant coefficient of variation cv_d (where $cv_d = \mu_d/\sigma_d$). Figure 6 depicts an illustrative example of how K^* and $G\left(K^*\right)$ change with the mean demand, μ_d, for constant coefficient of variation, $cv_d = 2$.

Impact of Correlation between Spot Market Price and Demand For positive correlation ρ, we would expect that the spot market becomes less attractive, because the spot price is more likely to be high when the demand is high, whereas the opposite should hold true for negative correlation. Indeed, our numerical study confirms that K^* and $G\left(K^*\right)$ increase with the value of correlation ρ. Moreover, the difference in expected total cost for different values of ρ becomes more pronounced for high values of standard deviation of demand (while mean demand remains constant), i.e., $G(K^*)$ grows faster with the standard deviation of demand for higher values of ρ. See Figure 7 for an example of how the correlation ρ impacts the total expected cost, $G\left(K^*\right)$.

Notice that both positive and negative correlation between spot price π and demand d could occur: Negative correlation would be typical for a mature industry where players merely fight for market share and the overall industry sales are flat. In such a situation, an increase in demand for one retailer's products is accompanied by a

Figure 7: Impact of σ_d and ρ for $BN\left(40, \sigma_d, 4, 1, \rho\right)$ and $p = 2$

decrease in demand for competing products. Hence, the amount of orders that the retailer's competitors will place on the spot market is low, resulting in a low spot price (while the retailer's demand is high). The opposite scenario (positive correlation between spot market price and demand) could realize in a growing industry, where increasing demand for one player's products is an indication of overall industry growth. The spot price would tend to be higher in such a scenario, reflecting the overall strong industry demand. Hence, a high value of d means that higher values of π are likely, too.

Spot Price Volatility Surprisingly, the impact of the spot market price variability depends on the correlation between spot market price and demand. When the standard deviation of the spot market price, σ_π, is low, the spot market price is relatively stable and predictable. Accordingly, the optimal level of reservation quantity, K^*, and the associated expected total cost, $G\left(K^*\right)$, increase with σ_π for positive values of correlation, ρ. However, our computations revealed (see Figure 8) that for negative correlation ρ, the spot market becomes more attractive with increasing volatility, σ_π, i.e., both K^* and $G\left(K^*\right)$ decrease with σ_π, because when demand is high, the spot

Figure 8: Effect of σ_π on K^* and $G(K^*)$ for $BN(20, 5, 3, \sigma_\pi, \rho)$ and $p = 2$

price is likely to be low due to the negative correlation. In such a situation, the high volatility of the spot price, as modeled by σ_π, increases the chances that a very low unit price is present on the spot market when the retailer's end demand, d, is high even more than already by the negative correlation between π and d.

6   Conclusions

In this paper we offered a generic framework for a simple setting to construct efficient supply contracts. We showed that the possible types of contracts include many of the well known ones from the academic literature. We introduced a linear pricing scheme with risk-sharing that allows channel coordination for which we study numerically and analytically the retailer's problem. An interesting conclusion is that although on average the spot market is considered a more expensive means of supply, in many cases the efficient contract quantity, K^*, will be around the mean value of the retailer's end demand, μ_d.

Despite the mushrooming of internet-based marketplaces, manufacturers (and suppliers) are reluctant to embrace them wholeheartedly and integrate them in their procurement (and sales) processes. To the best of our knowledge, our paper is one of the

first to address the issue of how to make best use of these spot markets by suggesting the optimal procurement strategy between long-term contracts, spot markets, or any mix between the two channels. More specifically, we show that such spot markets are beneficial from the retailer's perspective and that both spot markets and long-term contracts coexist under fairly general conditions. Since little is now known from existing spot markets because of their short history, we were careful to avoid any restrictive assumptions. Knowing that such spot markets can significantly reduce supplier costs, even if the unit price on the spot market is much higher than via the long-term contract or when the availability of products on the spot markets is limited, procurement managers should increase their companies' efforts towards internet-based procurement and prepare supply chains for such a change. Capitalizing on our models, these managers can effectively manage the risk of such a move and determine the optimal procurement channel mix.

We also studied the sensitivities of the results with respect to the various input parameters. Our computational study provides significant managerial insights into how the different parameters affect the optimal choice of the long-term reservation quantity level, K^*, and the associated expected total cost to the retailer, $G(K^*)$. In particular, the computational result that in most cases the chosen value of K^* is close to the mean demand value, is not at all evident and hence of very significant value. It is important to recognize that our results are even applicable for customized components, which are specifically made and designed for one customer. In such a situation, the spot market could be designed to trade production quantity instead of finished components, as long as the same machines can produce several customized products for potentially different companies (possibly requiring a set-up cost). The semiconductor industry represents an example, where it is common practice to negotiate supplier contracts based on wafer starts per week, i.e., in production capacity. Quality and delivery performance become more important in such a setting, since the retailer is highly dependent on the supplier's ability to deliver high quality products on time. However, as more and

more internet-based marketplaces offer quality and credibility verification on both the potential buyers and sellers, the main decision-making criteria will be the price of the good offered.

There are many avenues that we deem important to pursue in our future work. An essential one is a more complete analysis of the class of contracts that we introduced in this paper. More precisely, study more thoroughly the advantages and disadvantages of the linear pricing scheme over the other type of contracts such as the Buy Back or the Quantity Flexibility contract. We are also currently working on a more exhaustive model for the supplier, in a multi-period setting, that studies its optimal capacity management in the presence of a spot market. One issue that we have not fully addressed is the supply lead-time. We are aware that the impact of different lead-times of the long-term supplier and the spot market may tend to increase or decrease the importance of a spot market. To be able to quantify this, we plan to extend our model accordingly. To incorporate additional realism, as most retailers deal with a great number of components and suppliers, we would like to derive solutions for a multiple supplier/multiple components setting and see how robust the solution we obtained for the 1-supplier/1-component case is. A further extension of the last point would be to investigate multiple tiers of suppliers and see how the interactions and interconnections between different layers of the supply chain affect the overall supply chain performance and costs.

6.1   Acknowledgments

tive, and of AIM, the Alliance for Innovative Manufacturing, at Stanford University.

References

[1] Anupindi, R. and Y. Bassok, "Supply Contracts with Quantity Commitments and Stochastic Demand", in *Quantitative Models for Supply Chain Management*, S. Tayur, M. Magazine, and R. Ganshan (eds.), Kluwer Academic Publishers, Boston, MA, 1998

[2] Araman, V., J. Kleinknecht, R. Akella, E. Tse, T. Costy, and J. Tew, "Supplier and Procurement Risk Management in E-Business: Optimal Long-term and Spot Market Mix", GM Contract Report, Warren, MI, 2000

[3] Araman, V., J. Kleinknecht, and R. Akella, "Supplier and Procurement Risk Management in E-Business: Optimal Long-term and Spot Market Mix", Working Paper, Stanford University, Stanford, CA, 2001

[4] Barnes-Schuster, D., Y. Bassok, and R. Anupindi, "Coordination and Flexibility in Supply Contracts with Options", Working Paper, University of Chicago, 2000

[5] Berryman, K. and S. Heck, "Is the third time the charm for B2B?", *The McKinsey Quarterly*, 2001, Number 2, On-line tactics

[6] Cohen, M. and N. Agrarwal, "An Analytical Comparison of Long and Short Term Contracts", IIE Transactions, 31, 8 (1999), 783-796

[7] Donohue, K. L., "Efficient Supply Contracts for Fashion Goods with Forecast Updating and Two Production Modes", Management Science, 46, 11 (2000), 1397-1411

[8] Durlacher Research, "Business to Business E-Commerce, Investment Perspective", 2000

[9] Fong, D. K. H., V. M. Gempesaw, and J. K. Ord, "Analysis of a Dual Sourcing Inventory Model with Normal Unit Demand and Erlang Mixture Lead Times", European Journal of Operations Research, 120, 1 (2000), 97-107

[10] Lariviere, M. A., "Supply Chain Contracting and Coordination with Stochastic Demand", S. Tayur, R. Ganeshan, M. Magazine (eds.), *Quantitative Models for Supply Chain Management*, Kluwer Academic Publishers, Boston, MA, 1998

[11] Li, C. and P. Kouvelis, "Flexible and Risk-Sharing Supply Contracts Under Price Uncertainty", Management Science, 45, 10 (1999), 1378-1398

[12] Pasternack, B. A., "Optimal Pricing and Return Policies for Perishable Commodities", Marketing Science, 4, 2 (1985), 166-176

[13] Routledge, B., S., Duane J., and C. S. Spatt, "Equilibrium Forward Curves for Commodities", Journal of Finance, 55, 3 (2000), 1297-1338

[14] Rudi, N., "Dual Sourcing: Combining Make-to-Stock and Assemble-to-Order", Working Paper, The Simon School, University of Rochester, Rochester, NY, 1999

[15] Seifert, R. and U. Thonemann, "Optimal Procurement Strategies for Online Spot Markets", Working Paper, Stanford University, June 2000

[16] Tsay, A., "The Quantity Flexibility Contract and Supplier-Customer Incentives", Management Science, 45, 10 (1999), 1339-1358

[17] Tsay, A., S. Nahmias, and N. Agrawal, "Modeling Supply Chain Contracts: A Review", S. Tayur, R. Ganeshan, M. Magazine, (eds.), *Quantitative Models for Supply Chain Management*, Kluwer Academic Publishers, Boston, MA, 1998

[18] The Economist, "A great leap, preferably forward", Jan 18th, 2001

[19] Van Mieghem, J., "Coordinating Investment, Production, and Subcontracting", Management Science, Vol. 45, No. 7, July 1999, pp. 954-971