OPERATIONS RESEARCH
Vol. 59, No. 4, July–August 2011, pp. 866–881
ISSN 0030-364X | EISSN 1526-5463 | 11 | 5904 | 0866
http://dx.doi.org/10.1287/opre.1110.0927

Managing a Noncooperative Supply Chain with Limited Capacity

Rodney P. Parker
The University of Chicago, Chicago, Illinois 60637, rodney.parker@chicagobooth.edu

Roman Kapuściński
The University of Michigan, Ann Arbor, Michigan 48109, roman.kapuscinski@umich.edu

We consider a two-stage serial supply chain with capacity limits, where each installation is operated by managers attempting to minimize their own costs. A multiple-period model is necessitated by the multiple stages, capacity limits, stochastic demand, and the explicit consideration of inventories. With appropriate salvage value functions, a Markov equilibrium policy is found. Intuitive profit dominance allows for existence of a unique equilibrium solution, which is shown to be a modified echelon base-stock policy. This equilibrium policy structure is sustained in the infinite horizon. A numerical study compares the behavior of the decentralized system with the first-best integrated capacitated system. The performance of this decentralized system relative to the integrated system across other parameters can be very good over a broad range of values. This implies that an acceptable system performance may be attained without the imposition of a contract or other coordinating mechanism, which themselves may encounter difficulties in implementation in the form of negotiation, execution, or enforcement of these agreements. We find instances where tighter capacities may actually enhance channel efficiency. We also examine the effect of capacity utilization on the system suboptimality.

Subject classifications: inventory; capacity; supply chain; competition; Markov games.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received February 2006; revisions received October 2008, February 2010, June 2010; accepted August 2010. Published online in *Articles in Advance* July 12, 2011.

1. Introduction

The inclusion of noncooperative behavior into operational models adds an additional layer of realism, reflecting the relationships between independent firms in real supply chains acting in their own self-interest. We include such noncooperative behavior in a serial supply chain, together with several other elements of realism: production capacity limits, time dynamics, and stochastic consumer demand. All of this is done under the auspices of a model where the firms make production decisions and physical inventory and consumer backlogs are carried from period to period. While horizontal competition is studied in numerous papers (e.g., Parlar 1988, Lippman and McCardle 1997), vertical settings have attracted less attention. Furthermore, the common lexicon used in practice in these settings often de-emphasizes the agents' differing objectives by referring to them as "partners."

In this paper we are particularly interested to see how two elements observed in virtually all real situations, limited capacity and independent players, affect operations in vertical (serial) settings. We study the simplest setting where both elements are present, i.e., a two-stage competitive capacitated serial system. We assume that each of the two stages is owned separately, but the firms in the channel have the obvious commercial relationship of buyer and seller and a shared economic interest in satisfying the final customers. The model is evaluated over multiple time periods. All parameters and states are considered common knowledge so that all players have full information. We are interested in the following questions: Is there a unique equilibrium, and if not, is there a means of choosing an appropriate equilibrium? Does the structure of the equilibrium policy mirror that of the centralized case? Is there a significant performance gap due to not following the first-best solution? How does the split of benefits between retailer and supplier influence the efficiency of the whole chain? What is the effect of a more restrictive capacity upon the system efficiency?

The relevant literature can be divided into three categories: multiechelon research, capacitated research, and competitive research. For the literature in multiechelon systems, see Tayur et al. (1999). The first papers in the operations literature dealing with capacity decisions in simple one-stage systems include Arrow et al. (1958). Later, structural properties of operating under capacity constraints dominated a significant portion of this literature. Federgruen and Zipkin (1986a, b) demonstrate that a modified base-stock policy is optimal for a single-stage capacitated inventory system over an infinite horizon. This policy recommends ordering up to a particular inventory level, if possible, but to the capacity level otherwise. Kapuściński and Tayur (1998) extend this to a nonstationary (periodic) demand scenario. In Glasserman and Tayur (1994) several properties are shown for a serial multiechelon system with capacity limits when a base-stock policy is *assumed*, but this policy is not demonstrated to be optimal. Parker and Kapuściński (2004) examine a serial two-echelon system under centralized control, where each installation has a production capacity limit and the smaller capacity is at the retailer. The optimal policy is a modified echelon base-stock type, where the ongoing inventory operation can be reserved to a self-reinforcing "band" of inventory state space. The modified echelon base-stock policy is where the supplier will stock no more than the retailer can process in a single period, and both installations will attempt to reach desired echelon up-to levels (formally defined in Definition 2). Janakiraman and Muckstadt (2009) consider the same system and show a multilevel up-to policy is optimal for a system with both a single- and two-period leadtime upstream of the supplier, and derive a bound for the number of up-to levels in longer systems. The objective function in the papers cited in both groups above (multiechelon analysis and capacity) is to minimize the total system costs, i.e., multiple independent decision makers are not considered. Clearly, the advancement of capacity-limited inventory research has been surprisingly slow, indicating that the addition of capacity constraints dramatically increases the difficulty of these problems.

In the competitive setting, Cachon and Zipkin (1999) consider a decentralized two-stage (retailer and supplier) inventory game without capacity limits. The cost functions of the retailer and supplier in their paper are identical to those in our model. In their settings, the one-period problem is equivalent to the infinite-time problem, due to the assumed infinite capacities, as the supplier is always able to reach his desired target level, thus making it a regeneration point, and the retailer can achieve the same inventory distribution in every period. Due to this equivalence, local and echelon games can be compared with the supply chain optimal result from Clark and Scarf (1960). Such equivalence (between the one-period and infinite-horizon problems) is not readily available for systems with limited capacity, which we consider, because the initial inventory in every period may be arbitrarily low and we cannot guarantee any form of regeneration, so myopic policies cannot be considered.

There have been numerous recent studies investigating various mechanisms to "coordinate" the supply chain, beginning with Pasternack (1985) and including Chen (1999) and Porteus (2000). While this is a burgeoning line of inquiry (see Cachon 2003, for a summary), we do not design a coordinating mechanism or contract but seek the structure of the equilibrium inventory policy and examine the effects of capacity restrictions upon the system. Moreover, other than the below-cited papers (and references therein), single-period or infinite-horizon static settings are usually considered, whereas our model is inherently dynamic. Donohue (2000) illustrates the degree of complication of determining and executing contracts which span multiple periods (two in that case). The usage of Markov Games to inventory settings so far has been limited. Hall and Porteus (2000) consider a market-size dependent newsvendor under competition where the inventory perishes in one period and the only state variable that evolves over time is market size. Olsen and Parker (2008) consider a similar retailer competition setting but they allow for durable inventories and backlogging of unmet demand, so inventory levels in addition to market sizes are carried as state variables from period to period.

This paper provides contributions on three levels. First, the capacitated inventory theory literature is extended. It removes the central ownership assumption present in Parker and Kapuściński (2004). Second, this paper demonstrates an application of dynamic games to an inventory system. The dynamic game methodology has been used only in a limited sense previously in inventory models, with the exception of the above-mentioned papers; namely, for infinite-horizon games. It is noteworthy that such a representation is potentially susceptible to subgame deviations (see Başar and Olsder 1999). This weakness is not present in our dynamic game formulation. Last, the significance of our model is that it addresses three operational issues that are amongst the most important for a firm: (i) capacity constraints, (ii) inventory-related costs, and (iii) inter-firm coordination.

The paper is structured as follows. In §2 we describe the model and relate it to the fundamental requirements imposed in the economics literature. Section 3 contains key structural results, followed by related remarks in §4. Section 5, based on a numerical study, illustrates system dynamics and sensitivity to critical parameters, and it allows us to interpret the effects of capacity and competition. We conclude in §6. Proofs appear in Appendix A.

2. Model Description

Consider a multiple-period inventory model with a retailer ("her") and a supplier ("him") acting in their own interests. In each period the firms simultaneously decide the amounts to order, mindful that they cannot order any more than is available at the immediate upstream installation (the inventory availability constraint) or can be processed in that period (the capacity constraint). It is assumed that the outside supplier providing goods to the upper installation in our system has limitless availability of inventory. The demand is stochastic and stationary and realizations are unknown until after decisions are made. Unsatisfied consumer demand is fully backlogged. The distribution of demand is known to both players. The supplier's and retailer's objectives are to minimize their expected discounted backlogging and holding costs. Since capacity is limited, a large demand realization could result in a significant backlog that may take the system many periods to "catch up," resulting in each player being unable to attain their desired (presumed) stock levels. Thus, we explicitly consider a multiple-period (finite-horizon) game with a state (inventory position) which affects the equilibrium decisions in each period. Our objective is to characterize the behavior of both the retailer and supplier. Such behavior should be a subset of all equilibria.

2.1. Markov Equilibria and Strategies

We start with a justification of our choice of model in the broader landscape of models of competitive behavior. Game theory has many concepts of equilibrium. For our purposes, the Markov Equilibrium concept is more appropriate since we wish an explicit dependence upon some state variables and the stochasticity of our model as explained below. A Markov Equilibrium is associated with a Markov Game (see Başar and Olsder 1999), where the state can differ from period to period.[1] The current state can represent an aspect of the game that is manipulated by the players' actions, often a physical quantity or state.

Kirman and Sobel (1974) recognize the dynamics of physical inventory, noting the lack of intertemporal dependence in the extant literature: "whatever any player does in one period does not affect the game in the next period." Cachon and Zipkin (1999) avoid this dependence by considering the infinite-horizon version, endowing them with a static inventory policy. Due to the lack of capacity restrictions, the physical scenario is unchanged at the beginning of every time period resulting in the same equilibrium. Using the infinite-horizon version of Clark and Scarf's (1960) induced penalty cost functions (established by Federgruen and Zipkin 1984), they consider infinite-horizon versions of each installation's separated value function. For a finite horizon, the transferral of material between players and the delivery of goods to the consumer can actually be described by explicit functions, reminiscent of Clark and Scarf (1960) and Kirman and Sobel (1974). In the presence of capacity constraints and demand stochasticity, the assumption that all periods begin identically, as in a repeated game context, cannot be used. We apply, therefore, the Markov Equilibrium concept to a finite-horizon inventory model.

The appropriateness of ME is described in the economics literature. Maskin and Tirole (2001) provide a good discussion of the merits of MEs. First, Markov strategies prescribe the simplest form of behavior that is consistent with rationality. They depend upon a limited number of variables upon which decisions are based. Second, the notion of "bygones are bygones" is reflected in the use of the Markov concept. This suggests that the outcome of a game should only be affected by the strategic elements of that subgame. The third element is the principle that "minor causes should have minor effects," suggesting that the game should only be influenced greatly by factors that are significant. This selection of payoff-relevant history needs to be done carefully because the entire usage of the equilibrium concept depends upon it. In our modeling situation of a multiechelon inventory application, the selection of inventory levels at the beginning of each period as state variables seems natural.[2] We are conscious of the importance of this decision and are aware that any results are qualified by this selection. For further discussions of MEs, see also Fudenberg and Tirole (1991, chapter 13), Filar and Vrieze (1997), and Başar and Olsder (1999).

2.2. Formulation

In this section, the model is formulated, the notation is defined, and the assumptions are stated. We are modeling a multiperiod two-stage ($N = 2$) serial system. Each stage denotes a separate firm, attempting to minimize their own cost function. Each stage $j \in \{1, 2\}$ has a capacity limit, K_j, where we denote the retailer as stage 1, or installation 1, and the supplier as stage 2, or installation 2. Demand in period n, D, is stochastic and unsatisfied demand is backlogged. Costs include linear physical holding and backlogging penalty costs, charged after demand is realized.

The backlogging penalty cost is shared between the retailer and supplier. This setup is identical to that in Cachon and Zipkin (1999) and similar to Pasternack (1985). This is natural since a supplier (e.g., a manufacturer) would be concerned with the ultimate retail sale of his good; if the retailer is not selling the good, in turn the retailer will not be ordering goods from the supplier. Note that the supplier's share of the backlog cost does not imply a monetary exchange between the parties but an internal cost absorbed due to the insufficient channel service to the customers.[3]

It is appropriate to briefly discuss alternatives to the framework we are modeling for unsatisfied demand. Most economic models consider either one or two periods only and there is usually no modeling difference between lost sales and backlogging, and it is typical to interpret lost revenue as the cost of not satisfying the customer. While for short horizons, there may be a difference between sell-in (how much a supplier sells to a retailer) and sell-through (how much a retailer sells to the final customers), in the longer term these are identical and the entire demand not satisfied hurts both the retailer and supplier. The actual cost to each of them may be higher or lower than lost or delayed revenue, due to externalities such as the cost of lost goodwill (perhaps resulting in reduced future demand, as in Olsen and Parker 2008), the cost of transporting a product from an alternative retailer, or buying complementary products. Our approach of assigning the backlogging cost to both the supplier and retailer is very close to these interpretations. Obvious alternatives we do not model are to consider lost sales, partially lost sales, or self-imposed particular service levels. In these cases, however, the interaction between the firms' policies is modest and, most importantly, by exogenizing service levels, a corresponding level would not reflect one of the critical economic tradeoffs (between having too much and too little inventory) that is central to our model.

We assume the following:

- Demand, D, is a random variable with finite mean and second moment and $0 < \mathrm{E}[D] < K_1$. Its continuous probability distribution function is known to both firms. Demands are independent and identically distributed between periods.
- Any unsatisfied consumer demand is backlogged into the next period.
- Deliveries are made in the same periods as the orders placed for those goods, if availability and capacity limits permit.
- All costs and capacity levels are stationary, deterministic, and known by all firms.
- All payoff-relevant information is contained in the state variables, namely the inventory levels at the beginning of each period. Both firms know the inventory levels at both echelons at the beginning of each period.
- Both firms are rational and risk neutral. Both firms discount money at the same rate, $\beta \in [0, 1)$, although this is not necessary for any results.
- When equilibrium A results in lower costs than equilibrium B for both of the firms, B will be discarded (*Pareto refinement*).
- The retailer and supplier incur inventory carrying cost per unit per period of $h_1 + h_2$ and h_2, respectively ($h_j > 0$, $j = 1, 2$).
- p_j is the cost assessed to *installation* j for each unit of backlog at installation 1, $p_j > 0$.
- Capacities are ordered $K_2 \geqslant K_1 > 0$.

The timing of events within a period is as follows: (1) players observe inventory levels and place orders with their respective suppliers; (2a) the supplier delivers goods to the retailer; (2b) the outside supplier delivers goods to the supplier; (3) consumer demand is realized, and the retailer attempts to satisfy as much demand as possible; and (4) costs are assessed. We restrict our attention to pure strategies only. The orders by the supplier will be available for usage at the beginning of the following period. This translates into a single-period delivery leadtime for each firm. Leadtimes can be introduced between the supplier and retailer as described in §4.

The ordering of capacities ($K_1 \leqslant K_2$) limits the universality of the model, but still allows applicability to numerous industries. For example, in the steel industry the final stage of the supply chain is often a cold-rolling mill which is frequently a bottleneck. Clearly, the model's serial nature makes it inappropriate for assembly or distribution systems. The *stability condition* ($\mathrm{E}[D] < K_1$) is required to guarantee finite backlogs and stable operation of the system, similar to Parker and Kapuściński (2004).

Time is counted backwards from the end of the horizon with period 1 being the final period. We define the local inventory at installation j at the beginning of period n as x_n^j and the echelon inventory at echelon j in period n as $X_n^j := \sum_{i=1}^{j} x_n^i$. That is, the echelon inventory for a firm is the sum of all inventory at and downstream of that firm. Denote the inventory vector as $\tilde{x} = (x^1, x^2)$ and $\widetilde{X} = (X^1, X^2)$; the time subscript will be omitted generally. Let a_n^j be the actual amount ordered by installation j in period n. $\mathcal{A}^j(\widetilde{X})$ denotes the feasible action set for installation j, noting the possible dependence upon the current inventory position, and $\mathcal{A}(\widetilde{X}) = \bigtimes_{j=1}^{2} \mathcal{A}^j(\widetilde{X})$. The state space is $\{\widetilde{X} \subset \Re^2 \mid X^1 \leqslant X^2\}$. The inventory transition functions are $x_{n-1}^1 = x_n^1 + a_n^1 - D$ and $x_{n-1}^2 = x_n^2 + a_n^2 - a_n^1$. We can substitute with the standard echelon variables as follows: $X^1 = x^1$, $X^2 = x^1 + x^2$, $Y^1 = X^1 + a^1$, $Y^2 = X^2 + a^2 = x^1 + x^2 + a^2$, which implies, $X_{n-1}^1 = Y_n^1 - D$, $x_{n-1}^2 = x_n^2 - a_n^1 + a_n^2 = Y_n^2 - Y_n^1$, $X_{n-1}^2 = Y_n^2 - D$. Using these definitions, the periodic cost functions are:

$$L^1(\widetilde{Y}) = (h_1 + h_2)\mathrm{E}[(Y^1 - D)^+] + p_1 \mathrm{E}[(D - Y^1)^+], \quad (1)$$

$$L^2(\widetilde{Y}) = h_2(Y^2 - Y^1) + p_2 \mathrm{E}[(D - Y^1)^+], \quad (2)$$

where $(x)^+ = \max(0, x)$. The retailer's holding cost $(h_1 + h_2)$ reflects the value added by each of the installations. Denote the minimizing point of the one-period cost function for echelon 1, L^1, by y_{my}^*. Now we can define the value functions of the model for each player. As backorders are accepted at installation 1 but not at 2,[4] installation 1 may order up to the minimum of K_1 or x^2. Therefore, the action sets of feasible $\widetilde{Y}$'s are $\mathcal{A}^1(\widetilde{X}) = [X^1, \min(X^1 + K_1, X^2)]$ and $\mathcal{A}^2(\widetilde{X}) = [X^2, X^2 + K_2]$, and the game is defined as follows:

$$V_n^1(\widetilde{X}) = \mathrm{eqm}_{\widetilde{Y} \in \mathcal{A}(\widetilde{X})} J_n^1(\widetilde{Y}), \quad V_n^2(\widetilde{X}) = \mathrm{eqm}_{\widetilde{Y} \in \mathcal{A}(\widetilde{X})} J_n^2(\widetilde{Y}),$$

$$J_n^1(\widetilde{Y}) = L^1(\widetilde{Y}) + \beta \mathrm{E}[V_{n-1}^1(\widetilde{Y} - \widetilde{D})],$$

$$J_n^2(\widetilde{Y}) = L^2(\widetilde{Y}) + \beta \mathrm{E}[V_{n-1}^2(\widetilde{Y} - \widetilde{D})],$$

$$V_0^1(\cdot) = 0, \quad V_0^2(\cdot) = 0,$$

where the equilibrium operator (eqm) describes the value of a Nash equilibrium in each period, given the current state, $\widetilde{X}$, and constraints upon the players' actions, $\mathcal{A}(\widetilde{X})$. Vector $\widetilde{D}$ is defined as $\widetilde{D} = (D, D)$ (the same realization of demand for each vector element, corresponding to echelons). Our formulation assumes that the initial inventory levels are visible to both players. In the Remarks section (§4) we discuss the effect of various informational assumptions. Clearly, such a game is only well defined if the values associated with equilibria in all later periods are specified. As we show later, in the presence of multiple equilibria, each equilibrium has a different value for each player, thus requiring an equilibrium be uniquely determined, the focus of the next section. If uniqueness, or some other means of choosing a single equilibrium, were not possible, each equilibrium path would lead to different costs, thus eliminating the possibility of distinguishing a policy. Noticeably, although V_0^1 and V_0^2 are defined to be zero, separate salvage value functions will be applied in the final two periods. These will be dealt with in the following section.

3. Competing Through Echelon Levels

A natural question is whether the retailer and supplier compete by making decisions about their echelon or installation inventory levels. In the analysis of the centralized model, these two frameworks are equivalent. These two approaches are, however, not equivalent in decentralized settings.

To compare them, we first broaden the scope of potential alternatives and consider the sequential and simultaneous games. In a sequential game, with the retailer ordering first and the supplier second, similar to the centralized case, there is no difference between the echelon and the installation game—when the supplier chooses his order, he optimizes his value function. The sequential game might be a reasonable one to consider, if no informational or transportation delays took place. When some leadtimes are present, it is reasonable to expect that some of the supplier's decisions will be taken in anticipation of the retailer's orders, rather than knowing them for sure, which argues for a simultaneous game when any delays exist.

In the simultaneous game, two known formulations are the installation and echelon games. Both are for mathematical convenience. The actual decisions are order quantities a^1 by the retailer and a^2 by the supplier. Note that deciding a^1 is equivalent to deciding the ending echelon inventory Y^1 for a given X^1, as $Y^1 = X^1 + a^1$. Similarly, deciding a^2 is equivalent to deciding $Y^2 = X^2 + a^2$, given X^2. The intuition for the echelon game is straightforward: the supplier controls the total inventory in the two-echelon system but not his installation inventory, while the retailer controls how much of that two-echelon inventory should be moved to her warehouse and stored at a cost. Such a simple correspondence does not take place in installation variables. Consider the supplier who starts with inventory x^2 and intends to raise it to y^2. As $y^2 = x^2 + a^2 - a^1$, deciding y^2 would require knowing the retailer's order a^1 (the supplier does not control inventory withdrawn a^1). In case the retailer withdraws a different quantity than anticipated, the installation level will be different from the intended one as well. In this sense, to guarantee that the installation game results in the desired inventory at the supplier, the game would have to be sequential, with the retailer's decision being observed by the supplier before he makes his own decision.

Since it is difficult to provide any intuition that would justify a simultaneous game in installation levels, we consider a model when agents compete through their echelon levels. As explained below, we are able to limit our analysis to a particular band of the state space, $\mathscr{B}$, defined below. Even with this assumption, many properties that naturally hold in centralized models fail in competitive environments. Additional constructs will allow us to overcome these difficulties.

3.1. Salvage Value Functions

Salvage value functions are frequently used to (i) reflect economic reality, (ii) overcome undesirable and unrepresentative behavior at the end of the time horizon, or (iii) endow a model with analytical tractability. At the end of the horizon, there will frequently be nonstationary ordering behavior in inventory applications. This behavior is unrepresentative of the generally stationary ordering behavior for lengthy time horizons. Using salvage value functions, we can modify the end-of-horizon behavior (for our analytical convenience). As these one-time costs in the distant future (for reasonably long-time horizons) have little effect on the actions across most of the preceding periods, they are a device for generating sustainable behavior for the model.

We define salvage values which allow for sustainable (well-defined) equilibria. In general, the end-of-horizon effect of decreasing target levels exists, while, as we explain below, in our competitive setting it is critical that the target inventory levels for both firms do not decrease at the end of the horizon. To achieve this, we design two salvage value functions: a quadratic function for the retailer (Lemma A.1 in the Appendix illustrates such a quadratic function's effect upon a single-stage capacitated installation), and a kinked linear function for the supplier. Let us first consider the retailer's salvage value function:

$$S_0^1(\widetilde{X}) = \lambda_1 (X^1 - \gamma_1)^2,$$

where the appropriately chosen scaling factor $\lambda_1 > 0$ and translation factor $\gamma_1 \geqslant 0$ influence the retailer's order up-to level in period 1, the final period where the retailer makes a decision. Thus, $J_1^1(\widetilde{Y}) = L^1(\widetilde{Y}) + \beta \mathrm{E}[S_0^1(\widetilde{Y} - \widetilde{D})]$. Clearly, J_1^1 is a function of Y^1 only. Let $z_1^1 := \arg\min_{Y^1} J_1^1(\widetilde{Y})$. The advantage the quadratic salvage value function presents over a linear one is the facility to precisely locate a *finite* and *unique* minimizer in period 1, permitting the usage of a smaller portion of the state space at the end of the horizon. (This contributes to algorithmic efficiency when solving the model numerically, both through solving over a smaller state space in each period and a shorter time horizon until convergence.) Moreover, the quadratic function facilitates the derivative dominance analysis in Theorem 2, since the quadratic function's contribution to the derivative is a straight line. This enables the accurate placement of z_1^1, the retailer's period 1 minimizer. Since the retailer's salvage value function is applied in period 1, z_1^1 depends only on λ_1, γ_1, and her economic parameters.

The salvage value function for the supplier is defined as

$$S_1^2(\widetilde{X}) = \lambda_2 (\gamma_2 - X^2)^+,$$

and it is applied in period 2 according to the supplier's value function, $J_2^2(\widetilde{Y}) = L^2(\widetilde{Y}) + \beta \mathrm{E}[V_1^2(\widetilde{Y} - \widetilde{D})] + \beta \mathrm{E}[S_1^2(\widetilde{Y} - \widetilde{D})]$. We define $\gamma_2 := z_1^1$.[5] (Notice the time subscript indicates S_1^2 will affect the supplier's ordering decision in period 2, the final period in which he orders.) This choice of γ_2 forces the supplier's minimizing point to be located near the retailer's minimizing point. We define λ_2 as

$$\lambda_2 \begin{cases} = 0, & p_2 \leqslant \dfrac{1-\beta}{\beta} h_2, \\[2ex] \geqslant \dfrac{\beta p_2 \Pr(D + D > K_1)}{1-\beta}, & \text{otherwise.} \end{cases} \tag{3}$$

Note that $\Pr(D+D > K_1)$ in Equation (3) is referring to the sum of the random demands from two periods. Also note that if the penalty p_2 is small compared to the holding costs, it is not worthwhile for the supplier to hold stock. Formally, if $p_2 \leqslant (1-\beta)/\beta h_2$, in every period J_n^2 has a positive slope with respect to Y^2, and the effect of the accumulated penalties, starting from the following period (when the current supplier inventory could influence the retailer's inventory position), will *never* overcome the one-period holding cost and thus a salvage value is not needed; under this extreme circumstance, a salvage function is not necessary, and we assume $\lambda_2 = 0$. This result is shown in Theorem 3. When $p_2 > h_2(1-\beta)/\beta$, we need to define a value of λ_2 that guarantees both that the supplier's cost minimizer is positive and finite, as well as that the up-to levels, which describe an equilibrium, are nonincreasing in n. Within the proof of Theorem 2, we demonstrate how this value of λ_2 allows the dominance of cost function derivatives in neighboring periods.

3.2. Analysis

The following definition and lemma allow us to simplify notation and emphasize the structural elements of the main proofs.

Definition 1. • Let $x \mid [a, b] = \min(\max(a, x), b)$, i.e., the point within $[a, b]$ that is closest (or equal) to x.

• Let $\mathscr{D}$ denote the derivative, and let ∂_j denote the partial derivative with respect to the jth variable.

For convenience, let us define a subset of the state space, hereafter known as the "band" and the "modified echelon base-stock" policy.

Definition 2 (Parker and Kapuściński 2004). • The *band* is defined as $\mathscr{B} = \{\widetilde{X} \subset \Re^2 \mid X^1 \leqslant X^2 \leqslant X^1 + K_1\}$, the set of inventory states with the second-stage installation inventory not exceeding K_1.

• A policy is a *modified echelon base-stock* (MEBS) policy if there exist targets Z^{1*} and Z^{2*}, such that $Y^1 = Z^{1*} \mid [X^1, \min(X^1 + K_1, X^2)]$ and $Y^2 = Z^{2*} \mid [X^2, X^2 + K_1]$.

For the equilibrium policies, we only need to consider actions that result in inventory states in the band $\mathscr{B}$. The intuition behind this result is clear. Since the lower of the two capacities is at the retailer, the supplier can find no benefit in holding more inventory than the retailer's capacity since no more than this level can be drawn in any single period. And since the supplier's capacity exceeds the retailer's and the supplier can order and receive K_1 in one period, the retailer will never be starved of material if the supplier orders K_1 instead of a larger quantity. This logic is formalized in the following lemma.

Lemma 1. *Consider a system with beginning inventory $X_n^2 - X_n^1 \leqslant K_1$ for $K_1 \leqslant K_2$. When the set of equilibria is nonempty, a feasible policy such that $Y_n^2 - Y_n^1 > K_1$ for any n, cannot be an equilibrium.*

Based on the above lemma (the proof is in the appendix), we assume for the remainder of the paper that both echelons 1 and 2 are aware that it is not beneficial for echelon 2 to be outside of band $\mathscr{B}$.

We state without proof the pure-strategy existence theorem from Fudenberg and Tirole (1991, Theorem 1.2), rephrased for our purposes with a cost-minimization criterion.

Theorem 1 (Fudenberg and Tirole 1991). *Consider a strategic-form game, whose strategy spaces are nonempty compact convex subsets of a Euclidean space. If the payoff functions J^i are continuous in all players' strategies, $\widetilde{Y}$, and quasi-convex in Y^i, then there exists a pure-strategy Nash equilibrium.*

The following two theorems show the existence and properties of the equilibria. Theorem 2 applies when $p_2 > h_2(1-\beta)/\beta$ and Theorem 3 applies otherwise. Let us first define the following:

Definition 3. • $z_n^1 := \arg\min_{Y^1 \leqslant z_{n-1}^2 - K_1} J_n^1(Y^1, Y^1 + K_1)$.

• $z_n^2 := \arg\min_{Y^2} J_n^2(Y^1, Y^2)|_{Y^1 = Y^2}$.[6]

• $$\mathscr{B}_n := \begin{cases} \{\widetilde{X} \subset \Re^2 \mid X^1 \leqslant X^2 \leqslant X^1 + K_1, X^1 \leqslant z_{n-1}^2 - K_1\} \\ \quad \text{for } n > 2, \\ \mathscr{B} \quad \text{for } n \leqslant 2. \end{cases}$$

Definition 3 describes three attributes of the model in each period: z_n^1 is the retailer's minimizing point, assuming that the supplier will hold K_1 units of inventory (the upper edge of the band); z_n^2 describes the supplier's minimizing point; and $\mathscr{B}_n$ describes the subset of the band bounded above by the supplier's target stocking level. These points will be later shown to be uniquely defined.

It is convenient to define equilibria in two steps. In the first step we ignore a number of constraints defining feasible states, namely $X^1 \leqslant Y^1 \leqslant X^2 \leqslant Y^2 \leqslant X^2 + K_2$, and we incorporate only $Y^1 \leqslant Y^2$ and $Y^2 \leqslant Y^1 + K_1$ (based on Lemma 1). We will refer to such constructs as unconstrained response functions, $r_n^1(Y^2)$ and $r_n^2(Y^1)$, and unconstrained equilibria. In the second step, we incorporate all the omitted constraints. Lemma 2 demonstrates the shape of the unconstrained response functions, which is a critical element of the inductional step in the proof of Theorem 2. A specific period index has been suppressed in Lemma 2 since we are demonstrating the shape of the unconstrained best response functions based on the properties of the functions in any period. In Lemma 2, (a)–(d) are assumptions that allow the derivation of desired properties of each firm's best response functions. Its proof is in the appendix.

Lemma 2. *Consider constants $U > 0$ (upper bound) and $\phi > 0$. Define $\mathscr{B}^\phi := \{\widetilde{Y} \mid Y^1 + K_1 \leqslant Y^2 \leqslant Y^1 + K_1 + \phi\}$. Assume that $Y^1 \leqslant U$ and*

(a) *$J^i(Y^1, Y^2)$ are continuous and convex in Y^i, for $i = 1, 2$ and $\widetilde{Y} \in \mathscr{B} \cap \{Y^1 \leqslant U\}$; $J^1(Y^1, Y^2)$ is convex in Y^1 for $Y^1 + K_1 \leqslant Y^2 \leqslant U + K_1\}$;*

(b) *J^i is separable within the band, i.e., $J^i(Y^1, Y^2) = J^{i1}(Y^1) + J^{i2}(Y^2)$ for $\widetilde{Y} \in \mathscr{B} \cap \{Y^1 \leqslant U\}$*;

(c) *J^{11} and J^{22} are convex with minima at $\underline{z}^1$ and z^2, respectively*; *and*

(d) *$z^1 := \arg\min_{Y^1}(J^1(Y^1, Y^1 + K_1))$, $\underline{z}^1 \leqslant z^1 \leqslant U$, and there exists $\phi > 0$ such that $\arg\min_{Y^1} J^1(Y^1, Y^2)$ is constant in Y^2 for $z^1 + K_1 \leqslant Y^2 \leqslant z^1 + K_1 + \phi$ and $\widetilde{Y} \in \mathscr{B}^{\phi}$.*

Then, the best response functions are $r^2(Y^1) = z^2 \mid [Y^1, Y^1 + K_1]$, $r^1(Y^2) = \underline{z}^1 \mid [Y^2 - K_1, Y^2]$ for $Y^2 \leqslant z^1 + K_1$ and $r^1(Y^2) = z^1$ for $z^1 + K_1 \leqslant Y_2 \leqslant z^1 + K_1 + \phi$, and $r^1(Y^2) < Y^2 - K_1$ for $z^1 + K_1 + \phi < Y^2 \leqslant U - K_1$.

Potential cases considered in Lemma 2 are illustrated in Figure 1: in the top graph (A) the equilibrium up-to levels, z^1 and z^2, intersect outside the band; in the center graph (B) they intersect within the band; and in the bottom graph (C) the supplier up-to level is within K_1 of the retailer's myopic level, $\underline{z}^1 = y^*_{my}$. Consider Figure 1(A). The solid black line shows the unconstrained best-response function of the supplier. It is clearly limited to the band, based on Lemma 2. Since the supplier's cost functions are minimized at z^2, the response function remains as close to that level as possible while staying within the band. The retailer's best response function, shown as a solid grey line, is partly within the band. The retailer's one-period costs are minimized at $\underline{z}^1$, and her best-response attempts to stay within the band, as long as such behavior decreases the retailer's cost, while knowing that the supplier has no incentive to depart from the band. However, for higher levels of supplier inventory, the retailer's incentives change, and there exists a level z^1 above which she will never order. Notice the interval over which the two best-response functions overlap, indicating the existence of multiple equilibria. Multiple equilibria lead to ambiguity as to the appropriate value of a cost-to-go function. We later show (Theorem 2) that the supplier and retailer both prefer the same equilibrium thus removing this ambiguity; we formally use Pareto refinement to identify this equilibrium. When utilizing Lemma 2 within Theorem 2, z^1 will be become z^1_n, etc.

Without salvage value functions, the equilibrium up-to levels may be increasing in the horizon length, as in a conventional single-installation inventory model. This presents a possibility of the starting inventory in period 1 being above the period 1 equilibrium, which may lead to a loss of convexity for the following two reasons. First, as we illustrate in more detail below, a change in X^2 modifies the feasible region and requires that the operand needs to be convex in both Y^1 and Y^2. Second, above the equilibrium the cost functions are not convex. The cost function of the supplier echelon inventory is always flat as a function of the retailer's inventory, below the equilibrium base-stock levels. However, for the inventory above the up-to levels, the cost functions will decrease in the other firm's initial inventory. The technical implication of this is that the convexity of the cost functions is no longer guaranteed except below the equilibrium up-to levels. Thus, the motivation for the equilibrium up-to levels to decrease monotonically in the horizon length towards the steady-state levels, so that starting inventory levels in the next period are below the equilibrium levels. We show the monotonicity of equilibria in the following theorem through the usage of salvage value functions.

Figure 1. The "unconstrained" response functions for the retailer (thick grey line) and the supplier (thick black line).

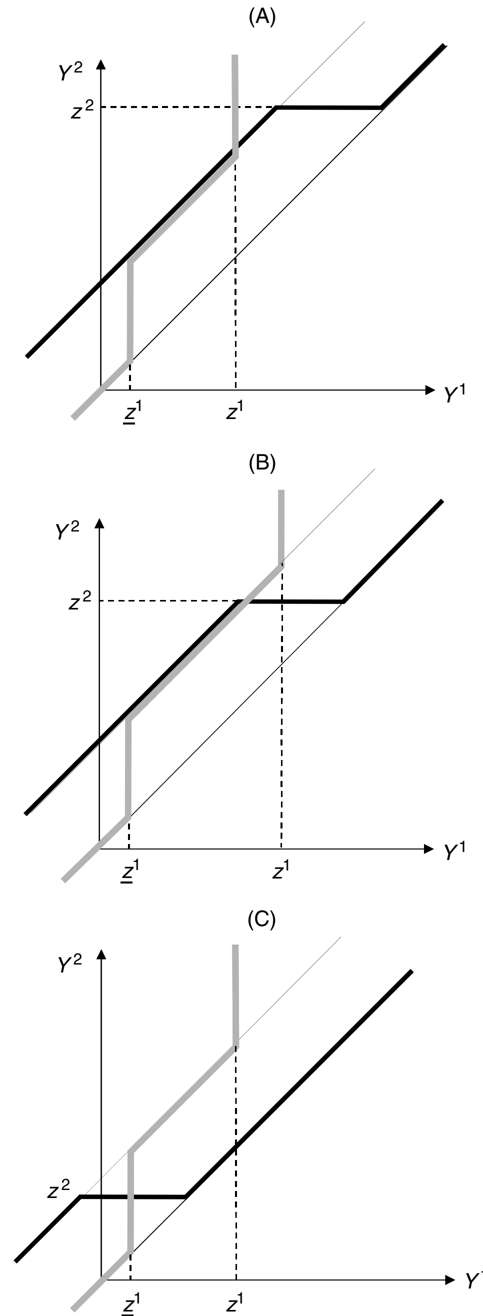

Notes. The top graph (A) applies when $z^1 \leqslant z^2 - K_1$, the center graph (B) when $z^1 > z^2 - K_1 > \underline{z}^1$, and the bottom graph (C) when $z^2 - K_1 \leqslant \underline{z}^1$. (Note that we always have $z^2 \geqslant \underline{z}^1$.)

Theorem 2 demonstrates the equilibrium structure and the monotonicity of the equilibrium up-to levels and the separability of the value functions within the band. When we discuss separability, we express it as a function of the initial or ending inventories and claim separability of those functions in the range up to the target levels. When referring to separability, we do not imply one-period cost minimization but instead minimization across multiple-period cost functions, which take advantage of the equilibrium policy structure in future periods and separability of costs to go. Separability in the context of our model (capacity limited two-echelon inventory game) is not trivial but is driven by the interaction of the following: the feasibility set (physical constraints), the equilibrium behavior of both players in the future, and the anticipation of economically justified responses. Furthermore, separation does not take place outside of what we label the band, and not above the equilibrium levels. Recall that z_n^1 and z_n^2 are defined in Definition 3.

THEOREM 2. *Assume* $p_2 > h_2(1-\beta)/\beta$. *Let salvage value functions* $S_0^1(\widetilde{X}) = \lambda_1(X^1-\gamma_1)^2$ *and* $S_1^2(\widetilde{X}) = \lambda_2(\gamma_2 - X^2)^+$. *There exist* $\lambda_1, \lambda_2, \gamma_1$, *and* γ_2 *such that for each starting inventory* $\widetilde{X} \in \mathcal{B}_k$:

(i) *Using Pareto refinement, there exists a unique pure-strategy Nash equilibrium, which is a modified echelon base-stock policy, in period* k;

(ii) $V_k^j(\widetilde{X}) = V_k^{j1}(X^1) + V_k^{j2}(X^2)$ *for* $j = 1, 2$; *and*

(iii) $z_{k+1}^j \leqslant z_k^j$ *for* $(j = 1, k \geqslant 1)$ *and* $(j = 2, k \geqslant 2)$.

To facilitate the understanding of the proof of Theorem 2, we provide the following "roadmap." The formal proof is in the electronic companion which is available as part of the online version at http://or.journal.informs.org/.

Roadmap to Theorem 2. The proof of Theorem 2 is by induction, similar to proofs in dynamic programming models. Since the induction step captures the most important logical elements, it is done first, while the basis step is the last element of the proof.

Induction step: For period $n-1$, the convexity and separability of each value function, for each decision variable, for portions of the feasible region is assumed. Similarly, we assume existence and ordering of the equilibrium order up-to levels for period $n-1$. These properties will then be demonstrated for period n. The existence of the equilibrium is established based on Theorem 1 (due to the convexity in each firm's own variable, continuity in both firms' variables, and compactness in the strategy spaces). When there are multiple equilibria, we demonstrate they are Pareto improving for increasing echelon inventories and thus both the supplier and retailer choose the same equilibrium. (Details are explained in the next bullet points.)

• *Unconstrained response functions and equilibria:* Recall that the unconstrained response functions ignore inventory constraints but do observe the supplier's economic incentive to operate within the band. Based on the convexity of the value function operand (induction assumption for period $n-1$), the results of Lemma 2 apply. Figure 1 demonstrates the three possible positions of response functions. When the firms' best-reply functions overlap at the upper edge of the band (Figures 1(A) and 1(B)), all these points represent legitimate equilibria. We establish that the cost functions of both the supplier and retailer decrease in their own and the other firm's inventory for overlapping portions of the response functions. Thus Pareto refinement implies that they both will choose the same unconstrained equilibrium.

• *Constrained equilibria:* The next step demonstrates the main result (the solution is a modified echelon base-stock policy) and also the separability of the value function and the convexity of the value function operands over certain portions of the state space. This is achieved by using the structure of the solution in period $n-1$ and is carried through to period n. The unconstrained response functions are now limited by the feasible regions (illustrated as rectangles in Figure 2), and thus the constrained equilibria will be contained within these feasible regions, which are defined by the period's initial echelon inventory position. The constrained equilibria are denoted by the solid circles in Figure 2.

• *Ordering of derivatives and up-to levels:* To show the monotonicity of the equilibrium up-to levels, we show that the derivatives of the value function operands are ordered. Specifically, if the derivative of the cost function in one period exceeds that in the following period (in our convex model), the up-to levels will increase as time moves forward. We separate the analysis into three parts corresponding to each of the cases in Figure 1. The state space is divided into multiple regions between the equilibrium up-to levels in periods $n-1$ and n, corresponding to seven possible reference inequalities in the online appendix, (A.1) to

Figure 2. Each rectangular area represents a feasible region for different initial echelon inventory positions.

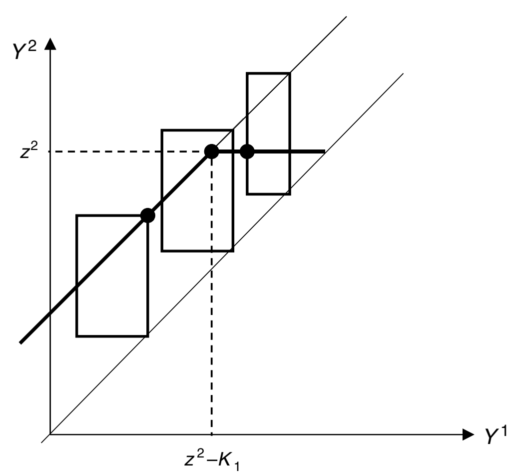

Notes. The left-lower corner of each rectangle represents the initial echelon inventory, $\widetilde{X}$. The feasible area is defined as $(Y^1, Y^2) \in [X^1, X^2] \times [X^2, X^2 + K_1]$. Consider the most left of the three rectangles: all the points $(Y^1, Y^1 + K_1)$ for $Y^1 \in [X^1, X^2]$ are equilibria, but the point $(X^2, X^2 + K_1)$ (represented by the solid circle) is a *dominating* equilibrium, as the costs of *both* echelons are lowest there among all equilibrium points.

(A.7). We demonstrate the same derivative dominance for the value function *operand* in period $n+1$.

• *Independence of the best response outside the band:* We show that the retailer's best-reply function rises vertically from the upper edge of the band at $Y^1 = z^1$. This technical property guarantees that the response function departs from the upper edge of the band and stays away from the band, thus uniquely identifying the equilibrium (since r^2 is contained in the band). If this were not the case, the response functions could potentially overlap on the upper edge of the band with conflicting directions (the retailer's costs decreasing and the supplier's costs increasing in Y^1 over (Y^1, Y^1+K_1)), thus destroying our use of Pareto improvement.

Induction basis: The final part of the proof establishes the basis for the induction, taking place over two periods, 1 and 2, instead of a single period. The reason for this is that in period 1 (the last period), the supplier will order nothing as he will receive no further orders. The actual decisions of the supplier that we capture are analyzed starting in period 2. The salvage value functions allow the initial ordering of the equilibria as well as satisfy all other conditions (ordering of derivatives). Applying the retailer's salvage function generates a finite period 1 cost minimizer at a point above the myopic minimizer, yielding the necessary results for the retailer. This minimizer establishes the translation factor (γ_2) for the supplier's salvage value function, applied in period 1. The necessary results are then determined for both the retailer and supplier in period 2.

The following theorem holds for $p_2 \leqslant h_2(1-\beta)/\beta$, that is, where the unit holding cost exceeds the discounted infinite sum of penalty costs for the supplier. The structure of the equilibrium policy, while similar to that of Theorem 2, has the supplier never ordering goods, resulting in a much simpler analysis. However, the retailer wishes to bring her inventory position to the newsvendor level. Thus, the firms' desired inventory positions are incompatible. The proof to Theorem 3 appears in the online appendix.

Theorem 3. *If $p_2 \leqslant h_2(1-\beta)/\beta$ for each starting inventory $\widetilde{X} \in \mathscr{B}$ $S_0^1(\widetilde{X}) = 0$ and $S_1^2(\widetilde{X}) = 0$, there exists a unique pure-strategy Nash equilibrium where the retailer orders up to a myopic base-stock level, y^*_{my}, if possible, and the supplier orders no goods at all. Also, $V_n^i(\widetilde{X}) = V_n^{i1}(X^1) + V_n^{i2}(X^2)$ for $i = 1, 2$, and $n > 0$.*

Now we present a result (the proof appears in the online appendix) demonstrating how the equilibrium up-to levels change with the economic parameters and the constraining capacity.

Theorem 4. *The equilibrium up-to levels are non-increasing when K_1, h_1, or h_2 increase or when p_1 or p_2 decrease.*

The final analytical result pertains to the applicability of the finite-horizon equilibrium policy to the infinite horizon. While extending finite-horizon results to the infinite horizon has become fairly standard in in traditional DP settings with one decision maker, it is not quite the case in game-theoretic settings. It is possible, however, to lean on some of the convergence results in the economics literature, as described below.

For the remainder of this section, we will count time forward (with an increasing index). To avoid confusion, we will use an index T rather than the n used previously for counting time backwards. We use the following standard terminology applied to our setting. The strategy space truncated at period T will be labeled $\mathscr{S}(T)$.

Definition 4. $g^* \in \mathscr{S}(T)$ is an ε-perfect Nash equilibrium if for each time $0 \leqslant s \leqslant T$, history x, strategy $g \in \mathscr{S}(T)$, and player i, $V^i(x^s(g^*)) - V^i(x^s(g^i, g^{*-i})) \leqslant \varepsilon$. 0-perfect Nash equilibrium will be labeled a perfect Nash equilibrium.

The following result comes from Fudenberg and Levine (1983) (hereafter FL83), and it will be used directly in the proof of Theorem 6. While Theorem 3.3B of FL83 assumes a deterministic game, footnote 2 on page 253 of §6 in FL83 states that the results hold for stochastic systems.[7]

Theorem 5 (Fudenberg and Levine 1983, Theorem 3.3B). *Suppose V is uniformly continuous. A necessary and sufficient condition that g^* be perfect in $\mathscr{S}(\infty)$ is that there be sequences ε_k, $T(k)$, and g_k such that g_k is ε_k-perfect in $\mathscr{S}(T(k))$ and as $k \to \infty$, $\varepsilon_k \to 0$, $T(k) \to \infty$, and $g_k \to g^*$.*

Theorem 6. *The MEBS equilibrium policy shown for the finite-horizon is also the equilibrium policy in the infinite horizon for $\beta \in [0, 1)$.*

Proof. We will use Theorem 5. Specifically, we will demonstrate that the following sufficient conditions are satisfied: (i) the value function is uniformly continuous; (ii) $\varepsilon_k \to 0$; (iii) $g_k \to g^*$, where g_k is ε_k-perfect in $\mathscr{S}(k)$. g_k denotes an ε_k-equilibrium for a k period problem, within $\mathscr{S}(k)$ the space of strategies, and g^* is the perfect equilibrium in the infinite horizon.

For our purposes, we will choose an arbitrary sequence ε_k monotonically decreasing to zero and (ii) trivially holds. We will show that (iii) holds for any arbitrary $\varepsilon > 0$. Convergence of policies (and of equilibria) is in the same topology as in FL83, where the distance between policies v and w is defined as $d(v, w) := \sup_T[(1/T)\min\{|\widetilde{Y}(v, T) - \widetilde{Y}(w, T)|, 1\}]$, v is policy, and $\widetilde{Y}(v, T)$ is the state in period T when following policy v. This distance notation may accommodate any of the standard metrics: sum of absolute differences ("Manhattan distance"), maximum of absolute differences (we use a modification of this, adopted directly from FL83), and square root distance.

For (iii), we need to first define g^*. g_k is the (perfect) equilibrium for the k-period game defined by targets (z_k^1, z_k^2) in the corresponding MEBS policy. Let $(z^{1*}, z^{2*}) := \lim_{k\to\infty}(z_k^1, z_k^2)$ for $k \geqslant 2$, which exists due to monotonicity and boundedness of (z_k^1, z_k^2), and let g^* be the corresponding MEBS policy. For an initial starting point $\widetilde{X}_1$ in period 1 and a demand sample path $d_1, d_2, \ldots, d_{i-1}$, denote $\widetilde{Y}_i = g_i(\widetilde{X}_1, d_1, d_2, \ldots, d_{i-1})$ and $\widetilde{Y}_i^* = g^*(\widetilde{X}_1, d_1, d_2, \ldots, d_{i-1})$.

We will show that $g_k \to g^*$. Consider any $\varepsilon > 0$. We will take advantage of the fact that the distance metric does not exceed 1. We will define T and m such that $|\widetilde{Y}_i - \widetilde{Y}_i^*| < \varepsilon$ for period $i = 1, \ldots, m$, while for periods $i = m+1, \ldots, T$, we have $1/(T-m) < \varepsilon$ implying that $\min\{|\widetilde{Y}_i - \widetilde{Y}_i^*|, 1\}/(T-m) < \varepsilon$. Specifically, we choose period m such that $z_m^j - z^{j*} < \varepsilon$, for $j = 1, 2$, and $1/m < \varepsilon$. Let $T \geqslant 2m$. For the first m periods, we have by induction ($j = 1, 2$): $Y_i^{j*} \leqslant Y_i^j \leqslant Y_i^{j*} + \varepsilon$, $z_i^{j*} \leqslant z_i^j \leqslant z_i^{*j} + \varepsilon$, and $X_{i+1}^{j*} \leqslant X_{i+1}^j \leqslant X_{i+1}^{j*} + \varepsilon$. This implies that the inventory differences are smaller than the difference in z's which are smaller than ε, i.e., $(1/i)\min\{|\widetilde{Y}_i - \widetilde{Y}_i^*|, 1\} \leqslant |\widetilde{Y}_i - \widetilde{Y}_i^*| < \varepsilon$ for $i = 1, 2, \ldots, m$. For the later periods $1/i \leqslant 1/(T-m) \leqslant 1/m < \varepsilon$. Thus, $d(\widetilde{Y}, \widetilde{Y}^*) \leqslant \varepsilon$ for $T \geqslant 2m$. This establishes the required convergence $g_k \to g^*$ (we have used strong sample-path bounds here that are *independent* of demand realizations). To show (i) that the value function is uniformly continuous, it is sufficient to show that $|v^k - w^k| \to 0$ implies $|V(v^k) - V(w^k)| \to 0$, where total cost $V(v)$ is expressed as a function of ordering policy v. This is straightforward due to the combination of the following three factors: the definition of distance between strategies, linearity of costs, and capacity constraints. Specifically, let $\varepsilon > 0$, and let $m_0 := \lfloor 1/\varepsilon \rfloor - 1$. $|v^k - w^k| < \varepsilon$ implies that for $i \leqslant m_0$, $|\widetilde{Y}(v, i) - \widetilde{Y}(w, i)| < i\varepsilon$. For $i > m_0$, due to the capacity constraint, we have $|\widetilde{Y}(v, i) - \widetilde{Y}(w, i)| < m_0\varepsilon + (i - m_0)K_1$. This implies that $|V(v^k) - V(w^k)| \leqslant (h_1 + h_2 + p)[\sum_{i=1}^{m_0} \beta^{i-1} i\varepsilon + \sum_{i=m_0+1}^{\infty} \beta^{i-1}(m_0\varepsilon + (i - m_0)K_1)] = (h_1 + h_2 + p)[\sum_{i=1}^{m_0} \beta^{i-1} i\varepsilon + \beta^{m_0} \cdot \sum_{i=1}^{\infty} \beta^{i-1}(m_0\varepsilon + iK_1)] \leqslant (h_1 + h_2 + p)[\varepsilon/(1-\beta)^2 + \beta^{m_0}(1/(1-\beta) + K_1/(1-\beta)^2)]$ (as $m_0\varepsilon \leqslant 1$). Clearly, as $\varepsilon \to 0$ we have $\beta^{m_0} \to 0$, and uniform continuity follows. □

4. Remarks

In this section we discuss some variants of our model and the corresponding results.

- *Longer supply chains*: A natural question is whether the current two-installation model can be extended to three or more installations. The enabling result, Lemma 1, fails for serial systems greater than two installations. The equilibrium dominance for the upstream installations which maintain stocking levels of K_1 or less, can no longer be guaranteed. Without this result, the primary result, Theorem 2, cannot follow.

- *Leadtimes*: We assume that delivery leadtimes are "natural" (or single period) leadtimes upstream of both the supplier and retailer. As aptly illustrated by Glasserman and Tayur (1994), leadtimes upstream of the supplier effectively mimic additional installations with single-period leadtimes, and as discussed in the preceding discussion point, supply chains of greater than two echelons cannot easily be analyzed. Leadtimes which are integer multiples of the period can be incorporated between the supplier and retailer, which might be representative of an outsourcing situation where, say, an Asian-based supplier is upstream of a U.S.-based retailer.

- *Capacity constraints vs. storage constraints*: In a one-period setting the production capacity constraints could be interpreted as storage limits. This equivalence does not hold for longer horizons, but storage limits actually would be far *easier* to characterize and analyze since they merely act as a limit upon the up-to decision variables as currently formulated in §2.

- *Common knowledge assumptions*: The current informational requirements of the system are that each firm knows the beginning inventory levels of themselves and the other firm. Given the decentralized nature of the problem, the obvious question is whether sufficient information can be contained in the retailer's orders. Cachon and Zipkin (1999) consider a decentralized model with infinite capacities. In their model, both the supplier and retailer have access to demand information (or equivalently retailer inventory) right away. Otherwise, the supplier would learn about demand with a one-period delay. We assume, as in Cachon and Zipkin (1999), that demand information is available immediately to both the retailer and supplier. Note that with capacity constraints, the actionable knowledge might remain the same as in Cachon and Zipkin (1999). The choice of ordering policy, however, influences the information passed. Assume temporarily that the target stocking levels are stationary and consider three different implementations. The retailer may (i) reorder up to her target level, (ii) simply order the last observed demand, or (iii) truncate the order to the capacity level. Clearly in case (i), all information about demand is passed to the supplier, as it is per definition in case (ii). In case (iii), the retailer's orders are sometimes truncated by the capacity. Despite the order truncation, the actions of the supplier are not always unduly influenced: when demand is larger than capacity, the supplier will not order more than capacity. However, if the difference between the supplier's and retailer's echelon stocking targets is less than capacity, there will be instances when the supplier's orders will not reflect all the demand information since he will stock to a local level less than capacity. Clearly, with target levels decreasing in the horizon length, the same logic holds, but the supplier needs to compensate for the difference between successive target levels. Thus, the specifics of the implemented ordering policy may influence the transfer of information.

- *Lower capacity at the supplier*: The sufficient conditions for the existence of an equilibrium (functional convexity) disappear when the smaller of the capacities is held by the supplier. Given the lack of easy-to-describe structure for the centralized version of the problem, it is not entirely surprising that convexity does not hold in the decentralized supply chain.

• *Coordinating mechanisms*: While we do not explicitly seek coordinating contracts but merely consider enhancement of channel efficiency, we are aware that this is an active research area. Cachon and Zipkin (1999) have shown that a three-parameter contract can coordinate in an uncapacitated two-echelon system. Clearly, three parameters are sufficient to coordinate a capacitated system. Consider, for example, sharing each firm's physical inventory cost and each firm sharing the retailer's backlog cost in the same proportion. Since such a contract is equivalent to total cost sharing, these contracts are considered complicated, are resisted, and are not observed in practice, to the best of our knowledge.

5. Numerical Examples

In this section we consider some numerical examples drawn from solving the model for sample data. In all the examples we use discretized probability distributions approximating a normal distribution with a mean of 9.8 with different values of variance and various levels for the discount factor, β. The models were iterated until *both* firms' value functions converged to within a margin of 0.005 (difference between subsequent periods). For all the following examples, we assign $h_1 + h_2 = 1$ and $\beta = 0.95$. In all the numerical examples, the echelon base-stock levels of the decentralized system were always lower than the echelon base-stock levels of the integrated system. This suggests competition reduces channel stocks, an observation compatible with the well-known *double marginalization* result. We present the results using relative additional cost which is defined as the percentage of the decentralized total cost over the integrated total cost.

First, consider Figure 3(a) which shows how the relative additional cost varies across the split α of the total unit penalty cost ($p = p_1 + p_2$, $p_1 = \alpha p$, $p_2 = (1-\alpha)p$) in the system for various values of the total penalty cost. We observe a high degree of consistency in percentage value above the integrated cost, for a wide range of total penalty values p from 1 to 50. Although not shown, this degree of consistency is maintained for far greater values of p also. This implies that the system's total cost is relatively independent of the values of the total penalty, for nonextreme values of α. We also observe for many parameter combinations that the efficiency can be shown to get to 5–10% above the integrated cost, establishing fairly efficient solutions of the system through the natural operation of the decentralized channel. Clearly, at extremely low values of α (the retailer's share of the total penalty), it is not in the retailer's interest to hold inventory; thus the customer is poorly served. At extremely high values of α, the retailer absorbs the bulk of the penalty costs but the inefficiency costs are not nearly as costly as for low values of α, implying that the retailer perhaps has more control over the supply chain.

A third observation from Figure 3(a) is the relative flatness of the curves across a broad range of values of α

Figure 3. The effect of penalty split across parameter values ($K_1 = 10, \beta = 0.95$).

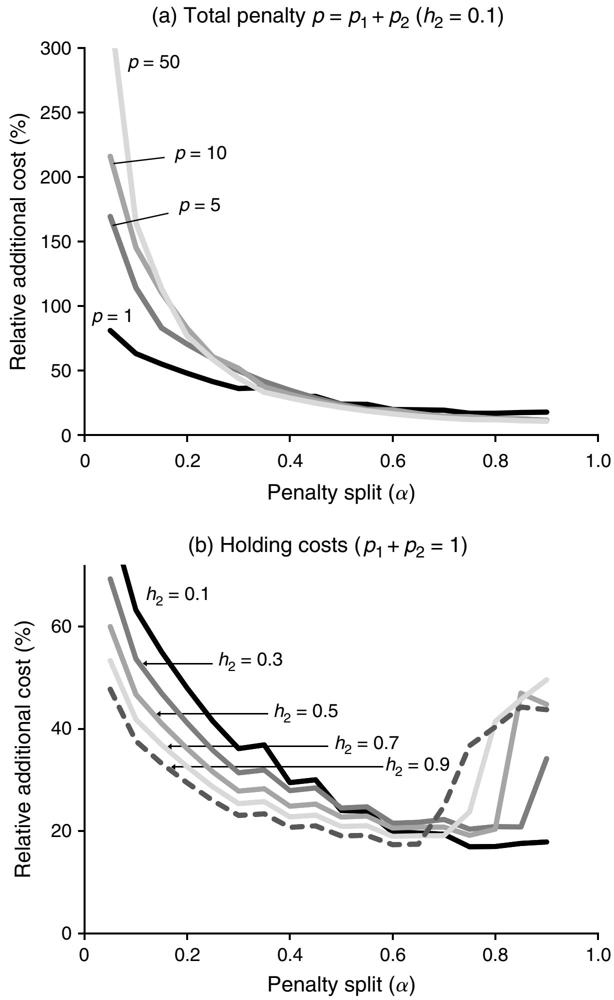

(the horizontal axis). The direct implication of this is that "good" solutions (closer to first-best) may be achieved for a variety of the allocations of the total penalty ($p_1 + p_2$) between the supplier and the retailer. So while a contract to coordinate the system may be difficult to find, let alone enforce, the system may be reasonably close to the coordinated one if the economic values of not satisfying customers for the supplier and retailer happen to fall in a fairly broad range of penalty allocation. Finding these economic values of not satisfying customers may be difficult to determine accurately. Fortunately, we observe a robustness in α indicating such a determination may not be necessary. While there is some mutual interest in avoiding the penalty cost, on the other hand, there is an incentive for shirking an effort related to holding inventory. The lower inventory levels for the decentralized system, compared to the centralized one, confirm this. Interestingly, despite the incentive for each firm to shirk effort and their potential lack of clarity about their respective responsibility for the backlogging penalty, the "decentralization" penalty upon the channel profits is not overwhelming. The attractiveness of this observation is that no third party ("principal") needs to impose their will upon the players in the system or attempt to force potentially unwilling players to sign a contract to coordinate the system; a reasonably good system solution is achieved merely from the normal operation of the supply chain.

In Figure 3(b) the unit holding cost of the supplier, h_2, is varied from 0.1 to 0.9 while the unit holding cost of the retailer is kept constant, $h_1 + h_2 = 1$. Note that a wide range of this ratio (between 0 and 100%) is driven by a possibility that a different portion of the value-added activities are performed in one of these two stages. As in Figure 3(a), the additional system cost relative to the integrated cost is plotted against the penalty split, α. What we first observe is that at lower values of α, the inefficiency is higher, as before. The curve with the highest cost here (at $\alpha = 0.05$) is $h_2 = 0.1$, and the relative costs decrease as h_2 increases. At the other end of the scale (at $\alpha = 0.95$), the sequence of the curves is reversed. A potential explanation for the ordering at each extreme value of α is that there is an imbalance of the values of the parameters. For example, at $\alpha = 0.05$, when $h_2 = 0.1$, the ratio of holding to penalty cost is low for the supplier while it is high for the retailer, so the system becomes unresponsive to customers, driving up costs; when $h_2 = 0.9$ and $\alpha = 0.95$, both of the ratios are very different and incentives are again misaligned—for example, the retailer anxious to carry extra stock faces an unmotivated supplier. A further observation is that generally the curves appear to be fairly closely aligned across the scale of α for nonextreme values and that there is a great deal of "flatness" of the curves (loosely speaking), implying a robustness with respect to α, a desirable property as indicated earlier.

In the next few charts, we consider the effect of varying the capacity constraint (in these examples, the mean of the distribution is 9.8 and the standard deviation is about 4, resulting in a coefficient of variation of over 0.4). Figure 4(a) shows that at the extreme values of α the more constrained systems (lower values of K_1) have higher relative additional system costs, but as the penalty costs are shared more equitably, these more constrained decentralized systems have lower relative additional system costs than the less constrained systems. There are two potential reasons for the players to attempt to coordinate: (1) they are both sharing the backlogging costs, and thus have an incentive to satisfy customers' demand; and (2) there are fewer alternatives, and this forces the retailer and supplier to coordinate even more closely.

This effect of having a tighter constraint which results in closer-to-coordinating behavior is not initially intuitive. In Figure 4(b), we find that the most constrained system ($K_1 = 10$) has the highest absolute costs. In most cases, while the "decentralization" penalty is small for a tight capacity, it does not imply low costs in absolute terms. In the centralized system, we observe the systemwide costs

Figure 4. The effect of penalty split for various capacity levels ($p_1 + p_2 = 1$, $h_2 = 0.9$, $\beta = 0.95$).

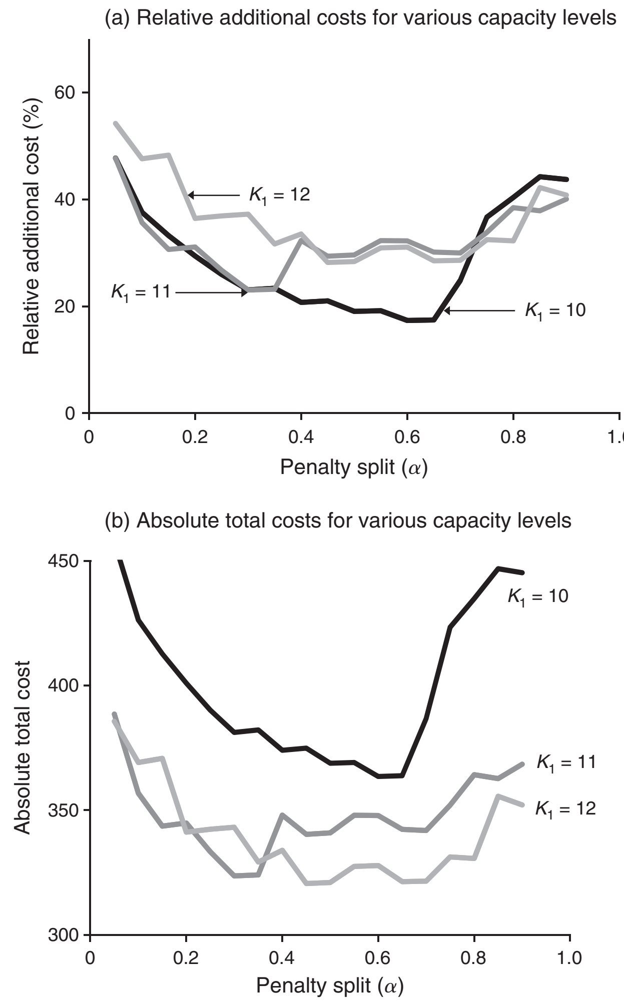

are monotonically increasing as the constraining capacity (K_1) decreases. In the decentralized system, this property is frequently true. However, we have observed circumstances when a tightening capacity may actually decrease absolute total system costs, as seen in Figure 4(b) when comparing the $K_1 = 11$ and $K_1 = 12$ curves.

In Figure 5(a) we have duplicated two of the curves from Figure 4(a), namely those for $K_1 = 10$ and $K_1 = 11$. We can observe the effect more clearly where the $K_1 = 10$ curve reaches closer to the integrated model cost for intermediate values of α while the $K_1 = 11$ curve has lower relative costs for extreme values of α. In Figure 5(b), these total cost curves from Figure 5(a) are separated for each player, and displayed as a percentage of the total system costs. First, the supplier's share of the total cost curves decreases and the retailer's share of the total costs increases as α increases, as we expect. However, a more interesting observation is that as there is an increase in capacity from $K_1 = 10$ to $K_1 = 11$, both players' absolute costs decrease. But the retailer's share of the total cost decreases and supplier's

Figure 5. The effect of penalty split for $K_1 = 10$ and $K_1 = 11$.

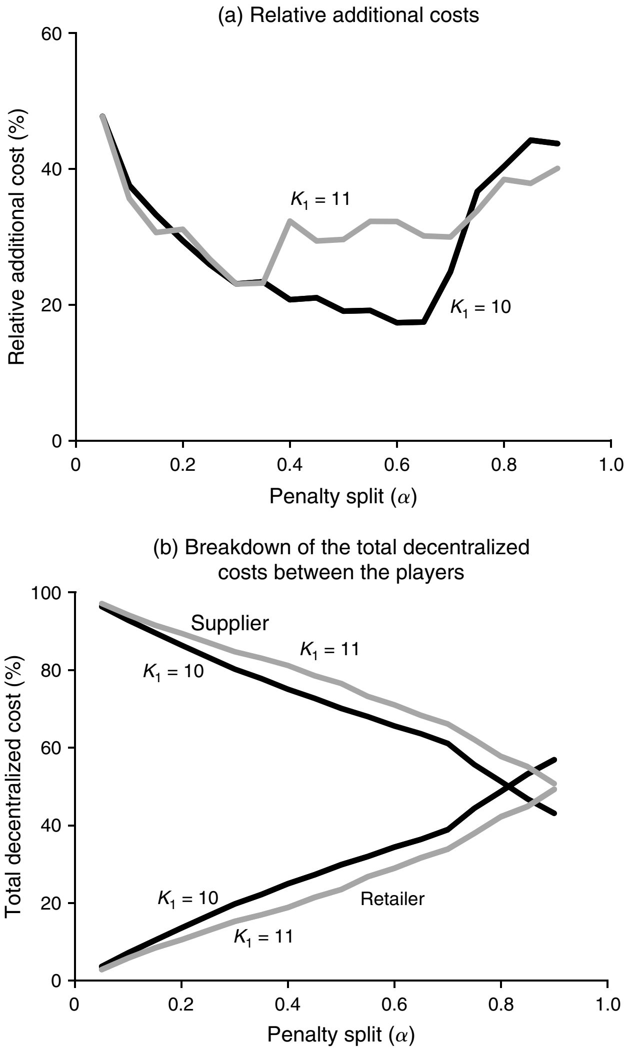

share of the total cost increases. An intuitive explanation for this could be that since the retailer is not constraining the system as much as before, she should not carry as much of the cost burden (similar to the induced penalty cost functions in Parker and Kapuściński 2004). We can consider the holding and penalty costs of each player. The total holding costs for the retailer will go down since she does not need to hold as much safety stock, and her total penalty costs will decrease since she will be able to satisfy the customers more easily since the capacity will not be limiting the flow of stock through the system. For much the same reason, the supplier's total penalty costs will be reduced also. However, there is ambiguity as to whether the supplier's holding costs will increase or decrease. They could increase as the larger capacity allows more material to be processed. They could also decrease similarly to the decrease in the retailer's base-stock level. Clearly, the sum of the supplier's holding and penalty costs (as a share of the total) increase as the capacity increases.

6. Conclusion

In this paper we model a decentralized serial two-stage multiechelon inventory system with capacity constraints and stochastic demand. Using the Markov Equilibrium solution concept, we demonstrate the existence and properties of the equilibrium policy of our model in a multiple-period context. The equilibrium policy is a modified echelon base-stock policy, identical in structure to the optimal policy in the centralized system. The modification is where the inventory levels are naturally restricted to a band of the state space. While the usual convexities are lost in this competitive framework, we utilize salvage value functions to generate base-stock levels monotonically decreasing in the length of the horizon, and we establish existence of either a unique equilibrium or Pareto dominating equilibria in every period. There are some subtle differences depending on the value of the supplier's unit penalty cost. Specifically, when the supplier's penalty cost is sufficiently small, the supplier will stock no inventory at all and the retailer will attempt to order up to the myopic level. We formally show that the equilibrium echelon base-stock levels are nonincreasing when K_1 increases, when h_1 or h_2 increase, or when p_1 or p_2 decrease.

In a numerical study, we evaluate the behavior of the decentralized system and compare it to the first best. It appears that the penalty costs $(p_1 + p_2)$ and split of the penalty costs (α), for nonextreme values, do not affect the relative total costs greatly, but the capacity limits do and in interesting ways. It appears that a tighter capacity can coordinate the system quite effectively by forcing the players to cooperate more closely, resulting in a total system cost closer to that of the integrated model. Thus, through the natural operation of the supply chain, *some* enhancement (not coordination) of the channel efficiency is observed through tighter capacity, rather than attempting to apply a coordinating contract.[8] We observe that, the decentralization penalty is not heavily influenced by the split of the penalty between the players, suggesting reasonably robust predictions of the efficiency gap. Altogether, the numerical study reinforces the importance of explicitly modeling the capacity constraints rather than merely taking an uncapacitated model as an approximation to a capacity constrained one.

7. Electronic Companion

Online appendices are available in the electronic companion to this paper, which is available as part of the online version at http://or.journal.informs.org/.

Appendix A. Additional Proofs

Lemma 1. *Consider a system with beginning inventory* $X_n^2 - X_n^1 \leqslant K_1$ *for* $K_1 \leqslant K_2$. *When the set of equilibria is nonempty, a feasible policy, such that* $Y_n^2 - Y_n^1 > K_1$ *for any* n, *cannot be an equilibrium.*

Proof. This proof is by sample path. Assume there exists a feasible dominant policy π, such that for a certain n, $Y_n^2 - Y_n^1 > K_1$. Choose the minimal n for which this condition takes place, i.e., $Y_m^2 - Y_m^1 \leqslant K_1$ for all $m < n$. We show that π cannot be a dominant equilibrium policy by constructing the following alternate policy, π', $Y_n^{2'} - Y_n^1 = K_1$ and $Y_n^{1'} = Y_n^1$ and $\pi' = \pi$ for all future periods (the variables under policy π' are denoted with the $'$ notation). Specifically, the supplier orders a smaller amount in period n under policy π' compared with policy π ($a_n^2 - a_n^{2'} = Y_n^2 - Y_n^{2'}$), but he compensates in period $n-1$ by ordering more, resulting in the same inventory levels at the end of period $n-1$ under both policies. Notice that the policies are identical from period $n-2$ onwards. The actions in period $n-1$ are not the same, but these result in the same ending inventories at the end of period $n-1$.

Recall from the formulation in §2 that a's denote the amount ordered. To check the feasibility of π', we need to check the following aspects of π' that differ from π: (i) $a_{n-1}^1 \leqslant x_{n-1}^{2'}$, (ii) $a_{n-1}^1 \leqslant K_1$, (iii) $a_n^{2'} \leqslant K_2$, and (iv) $a_{n-1}^{2'} \leqslant K_2$. (ii) is established from the feasibility of π. $Y_n^{2'} < Y_n^2$ implies that

$$a_n^{2'} < a_n^2 \leqslant K_2,$$

yielding (iii). Since $x_{n-1}^{2'} = K_1$ and from (ii), (i) is satisfied. Since $x_{n-2}^{2'} = x_{n-1}^{2'} - a_{n-1}^1 + a_{n-1}^{2'} \leqslant K_1$, and $x_{n-1}^{2'} - a_{n-1}^1 \geqslant 0$ from (i), $a_{n-1}^{2'} \leqslant K_1 \leqslant K_2$, satisfying (iv). Thus π' is feasible.

Because the total costs (for the current and future periods) for player 2 satisfy

$$\mathrm{Cost}^2(\pi) - \mathrm{Cost}^2(\pi') = h_2(x_{n-1}^2 - K_1) > 0$$

while player 1's inventory (or backlog) level are identical between policies π and π' at the beginning of period $n-1$, this results in identical costs thereon. Thus, the supplier can improve his costs under policy π by reducing his stocking level to K_1 without affecting the availability for the retailer, and therefore π cannot be an equilibrium. □

Lemma A.1. *Consider a single echelon with a capacity limit K, incurring unit holding and backlogging costs, h and p, stochastic demand D, and periodic cost, $L(y) = hE[(y-D)^+] + pE[(D-y)^+]$. Let $V_0(x) = \lambda(x-\gamma)^2$ be a salvage value function in period 0 and $\lambda > 0$ and $\gamma > 0$. Define the myopic minimum $y_{my}^* := \arg\min_y L(y)$, $y_n^* := \arg\min_y J_n(y)$, $J_n(y) := L(y) + \beta E[V_{n-1}(y-D)]$, and $V_n(x) = \min_{x \leqslant y \leqslant x+K} J_n(y)$, where $0 < \beta < 1$ for $n \geqslant 1$. For a chosen $\epsilon \in (0, h/4)$, let τ_0 be such that $L'(\tau_0) > h - \epsilon$, and τ_1 such that for all $\tau \geqslant \tau_1$, we have $L'(\tau) - \int_K^\infty L'(\tau + K - D) f(D) dD < \epsilon$. Setting* (a) $\gamma = \max(y_{my}^*, \tau_0, \tau_1) + E[D]$ *and* (b) $\lambda = h/(4\beta E[D])$, *then*

(i) *V_{n-1} is convex*;
(ii) *J_n is convex*;
(iii) *$J'_{n+1}(x) \geqslant J'_n(x)$ for $x \leqslant y_n^*$; and*
(iv) *$y_{n+1}^* \leqslant y_n^*$.*

Proof. (Intuition: A value of γ is chosen to be sufficiently high in the salvage value function to generate a final period minimizer that is large enough for the monotonicity of the up-to levels.)

V_0 clearly satisfies (i). If V_n is convex, due to the convexity of $\mathrm{E}[V_n(y-D)]$ and $L(y)$, J_{n+1} is also convex, which immediately implies convexity of V_{n+1}. Thus (i) and (ii) hold for all n. Also note that, for any n, (iii) implies (iv). Thus we only need to justify (iii).

Since J_1 and J_2 are convex, we need only show that $J'_2(y_1^*) \geqslant 0$ satisfies (iii). Since $L'(\tau_0) > h - \epsilon$ and (a) and (b), we have $L'(y_1^*) = -2\beta\lambda(y_1^* - \gamma - \mathrm{E}[D]) \leqslant h$ from the period 1 first-order condition, thus $y_1^* > \gamma - \mathrm{E}[D]$. Note, L' is a monotone nondecreasing function with limit h, thus establishing the existence of τ_1 (as $L'(\tau)$ and the weighted average of $L'(\tau - D)$, both increase and converge to h, the difference converges to 0 and is guaranteed not to exceed ϵ for sufficiently high τ's). We are interested in the territory $y \leqslant y_1^*$:

$$J_2(y) = L(y) + \beta\int_0^\infty \begin{Bmatrix} J_1(y+K-D) \\ \qquad y - D < y_1^* - K \\ J_1(y_1^*) \\ \qquad y_1^* - K \leqslant y - D \leqslant y_1^* \end{Bmatrix} f(D)\,dD$$
$$= L(y) + \beta\left\{\int_{y-y_1^*+K}^\infty J_1(y+K-D) f(D)\,dD + \int_0^{y-y_1^*+K} J_1(y_1^*) f(D)\,dD\right\},$$
$$J'_2(y) = L'(y) + \beta\int_{y-y_1^*+K}^\infty \{L'(y+K-D) + 2\beta\lambda(y+K-D-\gamma-\mathrm{E}[D])\} f(D)\,dD - \beta J_1(y_1^*) f(y-y_1^*+K) + \beta J_1(y_1^*) f(y-y_1^*+K).$$

Thus,

$$J'_2(y_1^*) = L'(y_1^*) + \beta\int_K^\infty \{L'(y_1^*+K-D) + 2\beta\lambda(y_1^* - \mathrm{E}[D] - \gamma + K - D)\} f(D)\,dD$$
$$= L'(y_1^*) + \beta\int_K^\infty \{L'(y_1^*+K-D) - L'(y_1^*) + L'(y_1^*) + 2\beta\lambda(y_1^* - \mathrm{E}[D] - \gamma) + 2\beta\lambda(K-D)\} f(D)\,dD$$
$$\geqslant (h-\epsilon) + \beta(-\epsilon) + 0 + \beta(-h/2),$$

which will be positive for $\epsilon < h/4$.

Assume (i) through (iv) hold in period $n-1$. Note that definition of V_{n-1} combined with convexity of J_{n-1} implies that $V'_{n-1}(x) = \max(\min(J'_{n-1}(x+K), 0), J'_{n-1}(x))$. Therefore, (iii) immediately implies that we have $V'_n(x) \geqslant V'_{n-1}(x)$ for $x \leqslant y_{n-1}^*$ and, therefore, also for $x \leqslant y_n^* \leqslant y_{n-1}^*$ and (iii) and (iv) follow. □

Lemma 2. *Consider constants $U > 0$ (upper bound) and $\phi > 0$. Define $\mathscr{B}^{\phi} := \{\tilde{Y} \mid Y^1 + K_1 \leqslant Y^2 \leqslant Y^1 + K_1 + \phi\}$. Assume that $Y^1 \leqslant U$ and*

(a) *$J^i(Y^1, Y^2)$ are continuous and convex in Y^i, for $i = 1, 2$ and $\tilde{Y} \in \mathscr{B} \cap \{Y^1 \leqslant U\}$, $J^1(Y^1, Y^2)$ is convex in Y^1 for $Y^1 + K_1 \leqslant Y^2 \leqslant U + K_1\}$;*

(b) *J^i is separable within the band, i.e., $J^i(Y^1, Y^2) = J^{i1}(Y^1) + J^{i2}(Y^2)$ for $\tilde{Y} \in \mathscr{B} \cap \{Y^1 \leqslant U\}$;*

(c) *J^{11} and J^{22} are convex with minima at z^1 and z^2, respectively; and*

(d) *$z^1 := \arg\min_{Y^1}(J^1(Y^1, Y^1 + K_1))$, $z^1 \leqslant z^1 \leqslant U$, and there exists $\phi > 0$ such that $\arg\min_{Y^1} J^1(Y^1, Y^2)$ is constant in Y^2 for $z^1 + K_1 \leqslant Y^2 \leqslant z^1 + K_1 + \phi$ and $\tilde{Y} \in \mathscr{B}^{\phi}$.*

Then, the best response functions are $r^2(Y^1) = z^2 \mid [Y^1, Y^1 + K_1]$, $r^1(Y^2) = z^1 \mid [Y^2 - K_1, Y^2]$ for $Y^2 \leqslant z^1 + K_1$, and $r^1(Y^2) = z^1$ for $z^1 + K_1 \leqslant Y_2 \leqslant z^1 + K_1 + \phi$, and $r^1(Y^2) < Y^2 - K_1$ for $z^1 + K_1 + \phi < Y^2 \leqslant U - K_1$.

Proof. Note that echelon 2 has no incentive to be outside the band $\mathscr{B}$. Since $J^i(Y^1, Y^2) = J^{i1}(Y^1) + J^{i2}(Y^2)$ for $\tilde{Y} \in \mathscr{B} \cap \{Y^1 \leqslant U\}$ and J^{22} is convex with minimizer z^2, we immediately have that, $r^2(Y^1) = z^2 \mid [Y^1, Y^1 + K_1]$ for $Y^1 \leqslant U$.

Now consider the response function for echelon 1, r^1. We separately describe r^1 (i) within $\mathscr{B}$ and (ii) outside the band $\mathscr{B}$, including $\mathscr{B}^{\phi}$.

(i) From (b) and (d) J^1 is separable and J^{11} is convex with minimizer $\underline{z}^1$, thus when r^1 is within the band $\mathscr{B}$, then $r^1(Y^2) = z^1 \mid [Y^2 - K_1, Y^2]$ for $Y^2 \leqslant U + K_1$.

(ii) From (d) $\arg\min_{Y^1} J^1(Y^1, Y^2)$ is constant in Y^2 for $\tilde{Y} \in \mathscr{B}^{\phi}$. Thus, for any Y^2, $J^1(Y^1, Y^2)$ decreases when $Y^1 < z^1$ and increases for $Y^1 > z^1$. Since (from (a)) $J^1(Y^1, Y^2)$ is convex in Y^1, we immediately have that $r^1(Y^2) = \underline{z}^1 \mid [Y^2 - K_1, Y^2]$ for $Y^2 \leqslant z^1 + K_1$ and $r^1(Y^2) = z^1$ for $z^1 + K_1 \leqslant Y^2 \leqslant z^1 + K_1 + \phi$. Due to (a) and (d), we also have $r^1(Y^2) \leqslant Y^2 - K_1 + \phi$. □

Endnotes

1. The term "Markov Perfect Equilibrium" (MPE) is used in some literature to emphasize that these equilibria are subgame perfect.

2. Clearly, the state variables might include the whole history of decisions, and actions may depend on inter-temporal properties of the firm's own or the competitor's behavior.

3. If a portion of the supplier's backlog cost *were* a monetary exchange, such a payment *could* be easily included in the model by redefining the coefficients of the penalty costs.

4. Allowing backorders at installation 2 without financial penalty would lead to an equivalent dynamics, while complicating the notation. We discuss this case in the Remarks section (§4).

5. This is correctly defined. z_1^1 is determined in the following period (period 1), thus in backward induction it is already known (while in a natural setting it is deducible).

6. This minimization occurs across the second argument for a specific value of the first argument. We subsequently show within the proof of Theorem 2 that z_n^2 minimizes J^2 for all feasible values of Y^1.

7. This assertion is corroborated by Börgers (1989). This is done by extending the result to games "in which Nature may make moves that are not completely observable; we will, however, continue to assume that the past actions of the *players* are common knowledge" (FL83). The stochastic game with no observability of actions requires the constraining assumption of a finite-action space. Such a restriction is not needed for full-information games, such as ours.

8. Most of the time, decreasing capacity increases the channel costs. There exist, however, cases where a tighter capacity will decrease the total costs.

Acknowledgments

The authors appreciate the support of Area Editor Jeannette Song, the associate editor, and two anonymous referees, whose persistence and efforts have considerably improved the paper. Financial support from the University of Chicago Booth School of Business is gratefully acknowledged.

References

Arrow, K. J., M. J. Beckmann, S. Karlin. 1958. The optimal expansion of the capacity of a firm. K. J. Arrow, S. Karlin, H. Scarf, eds. *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, Stanford, CA.

Başar, T., G. J. Olsder. 1999. *Dynamic Noncooperative Game Theory*, 2nd ed. Society for Industrial and Applied Mathematics, Philadelphia.

Börgers, T. 1989. Perfect equilibrium histories of finite and infinite horizon games. *J. Econom. Theory* **47** 218–227.

Cachon, G. P. 2003. Supply chain coordination with contracts. A. G. de Kok, S. C. Graves, eds. *Supply Chain Management: Design, Coordination, and Operation*. Handbooks in Operations Research and Management Science, Vol. 11. Elsevier, Amsterdam.

Cachon, G. P., P. H. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Sci.* **45**(7) 936–953.

Chen, F. 1999. Decentralized supply chains subject to information delays. *Management Sci.* **45**(8) 1076–1090.

Clark, A. J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6**(4) 475–490.

Donohue, K. L. 2000. Efficient supply contracts for fashion goods with forecast updating and two production modes. *Management Sci.* **46**(11) 1397–1411.

Federgruen, A., P. H. Zipkin. 1984. Computational issues in an infinite-horizon multiechelon inventory model. *Oper. Res.* **32**(4) 818–836.

Federgruen, A., P. H. Zipkin. 1986a. An inventory model with limited production capacity and uncertain demands I: The average-cost criterion. *Math. Oper. Res.* **11**(2) 193–207.

Federgruen, A., P. H. Zipkin. 1986b. An inventory model with limited production capacity and uncertain demands II: The discounted-cost criterion. *Math. Oper. Res.* **11**(2) 208–215.

Filar, J., K. Vrieze. 1997. *Competitive Markov Decision Processes*. Springer-Verlag, New York.

Fudenberg, D., D. Levine. 1983. Subgame-perfect equilibria of finite- and infinite-horizon games. *J. Econom. Theory* **31**(2) 251–268.

Fudenberg, D., J. Tirole. 1991. *Game Theory*. MIT Press, Cambridge, MA.

Glasserman, P., S. Tayur. 1994. The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy. *Oper. Res.* **42**(5) 913–925.

Hall, J., E. Porteus. 2000. Customer service competition in capacitated systems. *Manufacturing Service Oper. Management* **2**(2) 144–165.

Janakiraman, G., J. A. Muckstadt. 2009. A decomposition approach for a class of capacitated serial systems. *Oper. Res.* **57**(6) 1384–1393.

Kapuściński, R., S. Tayur. 1998. A capacitated production-inventory model with periodic demand. *Oper. Res.* **46**(6) 899–911.

Kirman, A. P., M. J. Sobel. 1974. Dynamic oligopoly with inventories. *Econometrica* **42**(2) 279–287.

Lippman, S. A., K. F. McCardle. 1997. The competitive newsboy. *Oper. Res.* **45**(1) 54–65.

Maskin, E., J. Tirole. 2001. Markov perfect equilibrium, I: Observable actions. *J. Economic Theory* **100** 191–219.

Olsen, T. L., R. P. Parker. 2008. Inventory management under market size dynamics. *Management Sci.* **54**(10) 1805–1821.

Parker, R. P., R. Kapuściński. 2004. Optimal policies for a capacitated two-echelon inventory system. *Oper. Res.* **52**(5) 739–755.

Parlar, M. 1988. Game theoretic analysis of the substitutable product inventory problem with random demands. *Naval Res. Logist.* **35** 397–409.

Pasternack, B. A. 1985. Optimal pricing and return policies for perishable commodities. *Marketing Sci.* **4**(2) 166–176.

Porteus, E. 2000. Responsibility tokens in supply chain management. *Manufacturing Service Oper. Management* **2**(2) 203–219.

Tayur, S., R. Ganeshan, M. Magazine. 1999. *Quantitative Models for Supply Chain Management*. Kluwer Academic Press, Boston.