

Chapter 1

THE STATE OF PRACTICE IN SUPPLY-CHAIN MANAGEMENT: A RESEARCH PERSPECTIVE*

Leroy B. Schwarz

Krannert Graduate School of Management

Purdue University

West Lafayette, Indiana 47907

lschwarz@mgmt.purdue.edu

Abstract In this paper, I will describe examples of state-of-the-art practice in supply-chain management; e.g., vendor-managed inventory, quick response, and other contemporary systems, such as Wal-Mart's *RetailLink*. The perspective will be that of what I call the *IDIB* Portfolio; i.e., what *Information (I)*, *Decision-Making (D)*, *Implementation (I)*, and *Buffer (B)* systems are employed in managing real-world supply chains. Most operations-research models consider only two components of this portfolio: the decision-making and the buffer systems. More specifically, most operations-research models involve selecting a decision-rule to minimize expected buffer (i.e., inventory-holding) and backorder cost given a fixed level of information. Implementation as a decision variable is typically ignored. However, in the real world, with changing information, communication, and implementation technologies, supply-chain management can – and *should* – be viewed as changing the nature of the *entire* IDIB Portfolio. After interpreting current practice from the perspective of the IDIB Portfolio, I will forecast *future* practice using Collaborative Planning, Forecasting, and Replenishment (CPFR) as an example. I will describe the elements of CPFR, identify companies that are using it, and the challenges they face in realizing its potential. Finally, I will identify research opportunities in CPFR, and, more generally, research opportunities involving the IDIB Portfolio.

*Forthcoming in *Applications of Supply Chain Management and E-Commerce Research in Industry*, E. Akçali, J. Geunes, P.M. Pardalos, H.E. Romeijn, and Z.J. Shen (editors), Kluwer Academic Publishers, Dordrecht, The Netherlands, 2004.

1. Introduction

We've all seen one version of it or another: The consulting company's Power-Point slide representing the evolution of supply-chain management (SCM). A staircase ascends left to right: on the bottom step is "Basic Supply-Chain Management", typically intra-company MRP or ERP systems; and on the top step, "Advanced Supply-Chain Management", typically described as "wireless, broad-band, web-based, truly collaborative, etc.". The next slide asks: "*Where is your company?*"

Sridhar Tayur (Carnegie Mellon University) displayed one of these staircase slides at a recent meeting of supply-chain thought leaders at Harvard University. Marshall Fisher (University of Pennsylvania) asked: "Does anyone know of any companies that are at or near the top step?" No one raised a hand. That slide, Marshall's question, and *no one's* answer provided the motivation for this paper.

In what follows I will do three things: First, I will describe the state-of-the-art in supply-chain management practice. A caveat: my assessment will be general; no details of specific buyer-supplier practice will be described. Second, I will introduce a paradigm – called the "IDIB Portfolio" – for understanding the evolution of supply-chain management to date, and for predicting its future. Third, using this framework, I will suggest topics for research.

My focus will be on managing the link(s) between buyers and suppliers that are independently owned and managed. Although centrally-owned and managed links provide a valuable benchmark, the real challenges in supply-chain management involve two or more independently-owned and managed companies. Two fundamental challenges are posed in improving the management of such supply chains: (1) the development of techniques to improve overall supply-chain performance (e.g., increasing total supply-chain profit); and (2) the development of contracting mechanisms that will motivate *all* the partners to implement these techniques. In other words: (1) how to "enlarge the pie"; and (2) how to "provide larger slices" to all the partners. I will focus on the first challenge – the development of techniques to improve overall performance. See Cachon, 2004 for a review of the supply-chain literature on the management of incentive conflicts with contracts.

This paper is organized as follows: Section 2 provides an overview of contemporary supply-chain management systems. In Sections 3-4 I will introduce the IDIB Portfolio paradigm, describe what "managing the IDIB Portfolio" means, and contrast the IDIB Portfolio paradigm with the operations-research paradigm. Section 5 uses the IDIB Portfolio to provide a perspective on contemporary supply-chain management prac-

tice. Section 6 describes two “Axioms of the IDIB Portfolio”. In section 7, I use these axioms to forecast the future of supply-chain management. Section 8 provides an overview of Collaborative Planning, Forecasting, and Replenishment (CPFR), which, I believe, is one “contender” for the future of supply-chain management. In Section 9, I identify several research topics in CPFR, in Section 10, several research topics in supply-chain collaboration, and, in Section 11, several research topics related to the IDIB Portfolio paradigm. Section 12 cites closely-related references. Section 13 provides a summary.

2. An Overview of Contemporary Supply-Chain Management Systems

Vendor-Managed Inventory (VMI), introduced by Kurt Solomon Associates in 1992 (<http://www.kurtsalmon.com>), is perhaps the most widely-known system for managing supply chains. Under VMI, the buyer authorizes the supplier (i.e., vendor) to manage the inventory of a set of stock-keeping units (SKUs) at the buyer’s site(s) under agreed-upon parameters (e.g., minimum and maximum inventory targets). The buyer provides the supplier with sales and/or inventory-status information; and the supplier makes and implements decisions about replenishment quantities and timings.

VMI reduces information distortion, which is one cause of the “bull-whip” effect (Lee et al., 1997). In addition, VMI provides the supplier with the opportunity to better manage its own production, inventory, and transportation costs. (See, for example, Çetinkaya and Lee, 2000). In exchange, the buyer typically receives price discounts or improved terms of payment from the supplier.

Quick Response (QR) was innovated by Milliken & Company (<http://www.milliken.com>) in the early 1990’s and subsequently codified by the Voluntary Interindustry Commerce Standards (VICS) Association. QR has four levels of application and technology. Levels 1 and 2, for example, involve retailer inventory-status information-sharing and automatic order-processing between retailer and supplier. Levels 3 and 4 include VMI and cross-docking warehouses. See Fiorito et al., 1994 for more information.

Although VMI and QR might be the “best-known” management systems among both practitioners and academics, perhaps the most *highly-regarded* systems are proprietary systems developed by large retailers, such as Wal-Mart’s *RetailLink*, Kmart’s *Workbench*, and Target’s *Partners Online*. Although the detailed inner workings of these systems are closely-guarded secrets, they all have two common characteristics: (1)

the sharing of transactions-level data among partners; and (2) the use of agreed-upon metrics (e.g., in-stock, inventory-turnover, and on-time delivery measures) and targets to assess partner performance. *RetailLink*, for example, captures sales, inventory, and delivery-related data for every SKU at every Wal-Mart facility (i.e., store and distribution center) and uploads it to a central database at least every 24 hours. These data, and metrics based upon them, are made available to *every* manager and *every* company up or down the supply chain whose performance is related to this SKU.

How this shared information is used – in particular, whether the decisions based on it are made centrally or decentrally – depends on the specific partnership agreement, and the type of product(s) being managed. Similarly, who does the implementation and how it is done is specific to the partnership and product(s). To illustrate: Wal-Mart generally centralizes decision-making but decentralizes implementation for Wal-Mart and Sam’s Club facilities, while delegating decision-making and implementation to its suppliers. However, regardless of who makes or implements the decisions, the quality of the decisions and their implementation are continually monitored by all parties and compared to the agreed-upon targets.

Before taking a closer look at current supply-chain practice, I will introduce a simple paradigm that will be useful in interpreting it.

3. Information, Decision, Implementation, and Buffer (IDIB) Systems

Managing anything, including managing a supply chain, can be viewed as four related activities: (1) *getting information* for decision-making; (2) *decision-making*; (3) *implementing decisions*; and (4) *buffering* against imperfections in (1)-(3). Correspondingly, every organization has systems for performing these four activities.

The Information System

The role of the “information system” is to provide past, present, and future-oriented information for decision-making. This information might be about demand, costs, materials, capacities, etc. An *ideal* Enterprise Resource Planning (ERP) system should be capable of capturing and providing *all* this information. In practice, however, information “in the information system” is captured and stored in dozens of different ways, among them paper records and computer files. In many organizations, much of this information is, literally, in the heads of management personnel.

Hence, a firm's information system is the less "a thing" and more a collection of "things": the collection of all business processes – formal and informal – plus all the information technologies and systems that provide information for decision-making. In this sense, information systems are "decision-support systems". However, I prefer the label "information system" since most decision-support systems are limited in their domain, and, hence, are a subset of the information system.

The overall quality of an information system's information about anything can be measured as some combination of: (1) its *accuracy*; that is, the correspondence between past, present, and future reality and what the information system reports (or reported) about it; (2) its *leadtime*; i.e., the time between an event and when the information system reports it; (3) its *level of aggregation*; i.e., the categories and units in which information is provided; and (4) its *horizon*; i.e., how far into the future (or past) the information system looks.

At their worst, information systems provide grossly inaccurate, even irrelevant, information. "Better" information systems typically provide some historical information and some current-status information. Still better information systems provide future-oriented information, such as demand forecasts and cost projections. A "perfect" information system would be the proverbial crystal ball, capable of seeing perfectly into the past, present, or future.

Another important characteristic of a company's information system is its cost; that is, the cost of the people, equipment, facilities, and processes that, together, comprise the information system. Typically, for a given technology, the cost of an information system is an increasing and marginally-increasing function of its overall quality. In other words, improving a given information system costs more, and each additional increment in quality costs more than the last.

From a management perspective, of course, the *value of an information system* doesn't depend on the *quality of the information* it provides, but on the *quality of the decisions made* based on this information.

The Decision-Making System

The role of a decision-making system, of course, is to make decisions using the information provided by the information system. Decision-making takes many different forms and is performed by many different individuals or groups. Decision-making occurs throughout every organization, from the shop-floor to the executive suite. Strategic decisions (e.g., the organization of the supply-chain, product offerings) are typically made at the executive level. Managerial decisions (e.g., the master schedule, order-promising) are typically made by middle managers.

Tactical decisions (e.g., processor assignment, workload sequencing) are typically delegated to the shop floor. The decision-making processes themselves might be informal, even intuitive (e.g., “Ned uses his 25 years of experience to assign workloads to processors”), or they might be codified, even regulated. For example, in US pharmaceutical manufacturing, most of the production decisions (e.g., the process steps and lot sizes) have been specified and approved by the FDA and, hence, *must* be adhered to. The corresponding decision rules might be simple – for example, in master-scheduling “run-out rules” are often used – or mathematically sophisticated (e.g., master-schedule optimization using mathematical programming). The decision-rules might be formal and automated; or, they might be totally within the head of some individual or group of individuals. Typically, important decisions are the result of a complex set of activities, some logical or based on management judgment, some simply guess-work.

The overall quality of a decision-making system’s decisions can be measured as some combination of: (1) the *optimality* of the decisions made; and (2) the decision-making *leadtime*. “Optimality” means how desirable the decision is – given the quality of information provided by the information system – with respect to cost, profit, or some other measure of utility. The “decision-making leadtime” is the amount of time it takes to make the decision once the appropriate information has been provided. This leadtime might be short or long; it might be a fixed amount of time, or variable; it might be known or unknown in advance.

Like information systems, the cost-drivers for decision-making systems are the people, equipment, facilities, and processes that, together, comprise it. And, like information systems, for a given technology, the cost of a decision-making system is typically an increasing, and marginally-increasing function of its overall quality. In other words, improving a given decision-making system costs more, and each additional increment in quality – for example, each increment in the desirability of its decisions or each decrease in leadtime – costs more than the last.

Finally, like information systems, the *value* of a decision-making system doesn’t depend only on its quality characteristics. In particular, a “perfect” decision that doesn’t get implemented is of *no* value.

The Implementation System

Implementation usually involves some “paperwork” to authorize or initiate activity. For example, a decision to ship 200 units from Dock 4 at 12 o’clock on April 5th will typically involve inventory-withdrawal authorizations, transportation requisitions, etc. Often, some preliminary actions must also be taken. For example, if 200 units aren’t in inven-

tory, then production decisions and their corresponding implementations must take place.

The overall quality of an implementation system can be measured as some combination of: (1) the *implementation leadtime*; and (2) *implementation accuracy*. The “implementation leadtime” is the amount of time required to make the decision happen; in other words, the time between making the decision and having the corresponding actions completed (i.e., implemented). For example, the amount of time it actually takes to ship the 200 units from Dock 4 once the decision has been made to do it. Like decision-making leadtimes, implementation leadtimes might be short or long, fixed or variable, known or unknown in advance. “Implementation accuracy” measures how closely the implementation matches the decision. Perfect accuracy means that the implementation perfectly matched the decision; for example, that exactly 200 units were shipped from Dock 4 at 12 o’clock on April 5th. In practice, implementation is seldom perfect: differences may be small, as in a tightly-controlled JIT system, or large, as in high-density chip fabrication, where yield losses are unpredictable and difficult to control.

Often trade-offs occur between the implementation leadtime and the accuracy-of-implementation. The phrase “quick-and-dirty”, for example, means that decision-making and implementation are “quick” (i.e., that their combined leadtime is short) but that the decision and/or its implementation are “dirty” (i.e., that the decision isn’t very desirable and/or the implementation isn’t accurate).¹

Like information and decision-making systems, implementation systems cost money; that is, the cost of the people, equipment, facilities, and processes that involve making the decision happen. And, like information and decision-making systems, for a given level of technology, the cost of an implementation system is typically an increasing and marginally-increasing function of its overall quality. In other words, improving a given implementation system – that is, making it faster or more accurate – costs more, and each additional increment in quality costs more than the last.

The Buffer System

In a perfect world, information systems would provide perfect information and decision-making systems would make perfect decisions. Implementation would be perfect, too. However, in the real world, *none* of

¹In fact, decision-making and implementation are often iterative. For example, a decision might be made; then, during the process of implementation, new information is revealed that might lead to modifying the decision, etc.

these three systems is *ever* perfect. Management systems compensate for these imperfections using buffers and buffer systems.

What are buffers? Unlike information, decision-making, and implementation systems, which can be realized in a virtual infinity of different forms, buffers come in *only* three basic forms: inventory, leadtime, and capacity. A “buffer system” is a combination of inventory, leadtime, and capacity buffers, in various amounts, located within what might be called the “management-system supply chain”. Often, for example, inflated leadtimes or extra capacity are imbedded in the information, decision-making, and implementation systems. Work-in-process inventories are typically found throughout the implementation system, with raw-material inventories at the beginning and finished goods at the end of the implementation chain. Leadtime buffers (e.g., inflated promise dates) are typically found at the interface between the points of delivery and the customers.

Re-thinking Buffers

The best way to understand – indeed, appreciate – the roles that buffers play in a management system is this: *forget everything you’ve thought about them until now*. In particular, put aside any notion that buffers are inherently “bad”.

Yes, buffers are often thought to be “bad”. Why? Perhaps it’s because buffers cost money. Yet, the other three components of a management system cost money, too. It is also worth noting – and peculiar, I think – that despite the fact that information, decision-making, and implementation systems *also* cost money, these three elements of a management system are generally thought to be “good”. Further, it is generally accepted that “improving” – that is, increasing the capabilities of – an information, decision-making, or implementation system is a “good thing”. On the other hand, it is generally thought that “improving a buffer” *must* mean reducing its capabilities or eliminating it. *Put this notion aside, too*.

Another reason why buffer systems might be thought to be “bad” is that whatever amount of buffering management chooses to provide, it is typically the *wrong* amount. In other words, the amount of buffering is either too much – for example, leftover inventory at the end of a selling season – or not enough – for example, that sales were lost despite the fact that “lots” of safety stock was provided.

I believe it is more useful – and accurate – to think about buffers in terms of the role they play in a management system: to compensate for imperfections in the information, decision-making, and implementation systems. From this perspective, I believe that buffers can only rightfully

be called “too big” if they use *more* resources than necessary to do what they are intended to do: to compensate for imperfections *elsewhere* in the management system. Similarly, buffers should only be thought of as “bad” if they *imperfectly* compensate for those imperfections.

What would a “perfect” amount of buffering be?

Consider a business scenario like that of the newsvendor model, and a management system whose information system, I , perfectly forecast customer demand. Assume that this perfect forecast is provided to decision-making, D , and implementation, I , systems that are capable of providing exactly this number of units at the instant they are demanded in precisely the right quantity. How much buffering, B , is required to perfectly satisfy customer demand? None. In other words, the “perfect” amount of buffering in *this* system *is* zero.

Now change the management system slightly so that the information system can only provide a probability distribution of future demand. Given this level of imperfection in the information system, is *any* management system capable of *always* satisfying customer demand *without* lost sales or unused inventory, capacity, or leadtime?

The answer is “yes”, *provided* that the decision-making and implementation systems are perfect; that is, provided that, once demand is known, management is capable of deciding to provide this amount, and provided that the implementation system is capable of producing precisely this amount *instantaneously*. Here, too, the perfect amount of buffering would be zero. On the other hand, *unless* the decision-making and implementation systems are perfect, *some* buffering will be required. How much and what kind of buffering?

Suppose that the “IDI_ systems” – that is, the combined information, decision-making, and implementation systems – are perfect with respect to leadtime, but imperfect with respect to quantity. In other words, *some* amount will be provided at exactly the right time to satisfy demand, but that the quantity won’t necessarily equal the quantity demanded. In order to avoid lost sales, *inventory buffering* will be required, the amount depending on the overall imperfection, or variance, in the quantity the IDI_ provides. As the variance of this imperfection increases, in order to compensate, the corresponding amount of buffer inventory must also increase. In the extreme, as this variance increases to infinity, the “perfect” amount of buffering also increases to infinity.

Next, suppose that IDI_ systems above are *imperfect* in quantity *and* leadtime. In order to satisfy demand at the time that it occurs in the quantity demanded, a *leadtime buffer* must be added to the inventory buffer. And, as the uncertainty in “supply leadtime” increases, the corresponding leadtime buffer must also increase.

Finally, since in many situations, the uncertainty arising from the IDI_ systems is a consequence of scheduling conflicts on some constrained resource(s), capacity buffers can usually be substituted for some of the inventory and leadtime buffers. Hence, the appropriate *forms* of buffering – inventory, leadtime, or capacity – depend on the *nature* of the imperfection(s) in the IDI_ systems.

Finally, then, given some form of buffering, what is the “perfect” amount?

The “Perfect” Amount of Buffering?

If the role of the *management system* – the *IDIB* – is to provide what is required or demanded at the time and in the quantity that’s required or demanded; and if the role of a buffer is to compensate for imperfections in the IDI_ systems, then the “perfect” amount of leadtime, capacity, or inventory is the amount that does just that and no more: the amount that provides precisely what’s required or demanded despite the imperfections in the IDI_ systems. Hence, the “perfect” amount of buffering is *never* zero unless the IDI_ systems are capable of perfection. Furthermore, as the combined imperfections in the IDI_ increases, the corresponding perfect amount of buffering must also increase.

An aside on the “Zero Inventory” concept often associated with Just-in-Time systems: The concept of zero inventory, or, more generally, zero buffering can be extremely important in identifying the nature and magnitude of imperfections in an IDI_ system. In other words, reduce the amount of buffering and “see” what imperfections are revealed. Often such experiments uncover imperfections in the IDI_ systems that are inexpensive to reduce or eliminate. If so, then this should be done and the corresponding amount of buffering should be permanently reduced. However, to the extent that imperfections in the IDI_ systems remain, the perfect amount of buffering isn’t zero.

Hence, the *perfect* amount of buffering depends on the *imperfections* in the combined IDI_ systems. If the amount of imperfections – that is, the uncertainties associated with the IDI_ systems – increases, then the overall level of buffering also must be “improved” – increased – in order to compensate.

The Cost of Buffering

Like all the other components of a management systems, the cost of a buffer system is typically an increasing, and marginally-increasing function of its overall quality. In other words, improving – that is, increasing – the capability of any given buffer costs more, and, typically, each additional increment in quality costs more than the last.

An Alternative Perspective: the “Optimal” Amount of Buffering

The operations research (OR) paradigm provides an alternative perspective on buffers, and, based on this perspective, defines the “optimal” amount of buffering. The perspective of the OR paradigm is as follows: The realized amount of buffering provided by a given management system is the difference between what reality requires or demands and what the management system provides. For example, if demand was 100 units and the management system provided 120 units, then the buffer was 20 units. Or, if management provided capacity of 40 hours and capacity of 50 hours was required, then the buffer was –10 units. A negative buffer, whether it’s inventory, capacity, or leadtime, means that the management system didn’t provide enough of what was required or demanded; a positive buffer means that too much was provided.

Correspondingly, the OR paradigm takes the view that the cost of the buffer system is the cost of all the positive buffering *plus* all the costs associated with negative buffering. In other words, the cost of the buffer system is the cost of all the *unused* inventory, capacity, and leadtime that the management system provided – measured *after* demand for that resource occurs – *plus* the cost of all the corresponding shortages of inventory, capacity, and leadtime. (e.g., lost sales, backorder, goodwill cost).

The newsvendor model is, perhaps, the best-known example of the OR paradigm. The newsvendor model chooses the optimal “target inventory” based on three “drivers”: the per-unit cost of “not enough” buffering (i.e., opportunity cost of lost sales), the per-unit cost of “too-much buffering” (i.e., the out-of-pocket leftover cost), and the probability distribution of demand. The expected-cost minimizing target inventory is provided by the well-known “newsvendor fractile” of the cumulative demand distribution. The corresponding “safety-stock” (inventory) buffer is measured, *a priori*, as the difference between the chosen target inventory and the expected customer demand.

What’s Wrong with the OR Paradigm’s View of Buffers?

There is nothing “wrong” with the OR paradigm or with its view of buffers. “Near-sighted” would be a better description.

In particular, the OR-paradigm’s view of the too-much and not-enough costs associated with the imperfections of a management system is consistent with the IDIB-paradigm perspective. In other words, once, say, demand occurs, the relevant costs incurred because of an imperfect IDIB Portfolio are those associated with providing too much or not enough of whatever was demanded.

But why should these too-much and not-enough costs *only* be associated with – if you will, “blamed” on – the buffer system? Yes, buffers are typically imperfect, but what about the imperfections in the IDI_ systems? More on this below.

One way of looking at the difference between the OR and the IDIB paradigms is that the IDIB paradigm generalizes the OR paradigm. Specifically, the IDIB paradigm takes the view that the quality of *all* four components of a management system – not just the buffer system quality (i.e., size) – are decision variables; second, that the underlying quality-cost function for each of these components – that is, the cost to move them in the direction of perfection – is increasing and marginally increasing. Finally, that there is a cost associated with the entire IDIB Portfolio – again, not just the buffer system – for failing to provide whatever is required or demanded. Hence, the “optimal” IDIB Portfolio is the IDIB Portfolio that minimizes the total costs of *all* of its components *plus* the cost of failing to provide precisely what is/was required or demanded.

From this point of view, the OR paradigm is “near-sighted” to the extent that it takes the IDI_ as fixed, and focuses only on picking the amount of buffering that minimizes the corresponding too-much and not-enough buffer cost.

We will address the question of the optimal IDIB in Section 11. We turn now to the concept of the “IDIB Portfolio”.

4. The IDIB Portfolio

I label the *combination* of these four components of a management system – the information, decision-making, implementation, and buffer systems – a “portfolio”, because, like a financial portfolio, each of these systems involves an investment of dollars. And, like the performance of an investment portfolio, the performance of a management system depends on how well its components perform in *combination*, not as separate components. Finally, in assembling an IDIB Portfolio, as in assembling a financial portfolio, an almost unlimited number of combinations can be chosen. As an illustration, suppose the goal is to manage a supply chain to provide a 95% customer fill-rate at the retail level. This might be provided by managing every link of the chain with low-quality information, decision-making, and implementation systems, but large inventory buffers everywhere. Or, without changing the management of the other links in this supply chain, the manufacturer might substitute buffer capacity for some of its finished-goods inventory and still offer the same service to the distributor. Similarly, the distributor

might choose higher cost of express delivery (from the manufacturer) in order to reduce the expediting and inventory-holding cost on its safety-stock (buffer) inventory. Three different IDIB Portfolios all providing the same level of customer service.

Which is best?

If the goal of a management system is to maximize profit, then the “best” IDIB Portfolio is the portfolio with the least total cost: that is, the total cost of all the people, facilities, equipment, and technology associated with providing information, making decisions, implementing them, and buffering to compensate for imperfections in the IDI_ systems, *plus* the cost of lost sales, expediting, and goodwill loss resulting from failing to do all this perfectly.

Managing the IDIB Portfolio

“Managing” the IDIB Portfolio means making decisions about the nature and quality of its four components: the information, decision-making, implementation, and buffer systems. The *ideal* is to select the quality of each component, plus the cost of lost sales, etc., so that total cost is minimized. Is “managing the IDIB Portfolio” amenable to the tools of operations research? I think it *could* be, but most OR models are much too myopic to be called “IDIB optimizers”, much less, “IDIB improvers”.

Consider the well-known newsvendor inventory model. In this model, the information system provides the probability distribution of customer demand and estimates of the costs associated with buying and selling newspapers. Attention is focused on the decision about the number of newspapers to have on hand at the beginning of the day in order to maximize the newsvendor’s expected profit. The optimal decision-rule is well known: set this inventory equal to the “critical fractile” of the probability distribution of customer demand. The newsvendor’s implementation leadtime is not explicitly considered. Instead, it is usually assumed that whatever this leadtime is, it is short enough so that, once the newsvendor has decided how much to order, the chosen quantity will be delivered on time. The basic newsvendor model also ignores implementation accuracy; that is, it is implicitly assumed that whatever quantity is ordered will be delivered. Extensions of the newsvendor model consider accuracy; that is, the correspondence between what is ordered and what is delivered (See Karlin, 1958 and Ehrhardt and Taube, 1987, for example; see Yano and Lee, 1995 for other references).

I believe the newsvendor model is representative of virtually every operations-research model of supply-chain management. That is, the quality of the information provided by the information system is *assumed* to be fixed. The costs associated with the information, decision-

making are typically *ignored*. The costs associated with implementation are sometimes represented, but in a highly-stylized manner. For example, the only implementation cost associated with the newsvendor model is the marginal purchasing cost. Similarly, the only implementation cost associated with the EOQ model is the assumed-to-be-fixed order cost. The goal, as described above, is to determine the decision-rule that minimizes buffer-system cost. In IDIB Portfolio terms, the newsvendor model – like most supply-chain models – selects the decision-rule for the decision-making system (“D”) that minimizes expected total buffer-system (“B”) cost for a fixed quality of information, “I”. The implementation “I” is represented either as a cost or as a constraint; not as a decision-variable.

“*Managing the Newsvendor’s IDIB Portfolio*” is much more complex. It *does* involve selecting decision rules, but the objective is to minimize *total* portfolio cost, not just the cost of the buffer system. In particular, managing the newsvendor’s IDIB Portfolio *also* involves assessing the cost and value of an information system that would provide more (or, possibly, even less!) precise information about customer demand. It *also* involves assessing the cost and value of implementation systems with different leadtimes and accuracies. More broadly, managing the newsvendor’s IDIB portfolio might involve fundamental changes in the newsvendor’s operations.

Suppose, for example, that the newsvendor was able to make and implement her/his ordering decisions any number of times during the day and receive those newspapers instantaneously (i.e., decision and implementation leadtimes are zero). In such a scenario, the newsvendor wouldn’t inventory *any* newspapers, nor would he/she need a probability distribution of demand. Instead, the newsvendor would wait until a customer requested a newspaper and then provide it upon demand. Sound far fetched? Consider this:

For decades, the copier division of Xerox struggled with managing the inventory of owner’s manuals for its copiers: how much inventory of which manuals to have on hand, and when to replenish this inventory. Xerox eliminated this problem by developing a system for printing and binding manuals upon demand; i.e., whenever assembly of a copier is scheduled, the printing of its manuals is also scheduled.

How was this done? By developing a system for implementing the decision to produce manuals whose leadtime is less than or equal to the time required to implement the decision to assemble the copier.

“Managing” the IDIB Portfolio is no easy task. First, the different components of the portfolio are often difficult to identify. For example, managers, who are nominally decision-makers, also often play a role in

the information system; line personnel, who are nominally implementers, have roles to play in information, decision-making, and buffering. So, it is often difficult to separate the components of a firm's IDIB Portfolio. Second, many of the costs associated with a firm's IDIB Portfolio are "overhead" or "indirect" costs, which makes them difficult to estimate. These and other difficulties make it virtually impossible to find the truly optimal IDIB Portfolio; that is, the portfolio of information, decision-making, implementation, and buffering whose combined cost *plus* the cost of failing to provide enough of whatever was demanded or required (e.g., backordering or expediting cost) is the minimum possible cost.

Nonetheless, it is often relatively easy to verify that one given IDIB Portfolio has lower total cost than another. In the Xerox example above, it was relatively easy to verify, using back-of-the-envelope estimates, that Xerox's "new" portfolio was an order of magnitude less expensive than its old one.

Xerox's old information system ignored the fact that *Xerox management* decided which copiers to produce and when to produce them. Instead, it assumed that demand for a given manual was provided by a probability distribution. Xerox's new IDIB Portfolio uses information that management has scheduled a given copier to be assembled, and then implements the decision to print its manuals in a short enough leadtime so that the manual can be packaged with the copier at the end of the assembly line.

I believe that experienced operations researchers, practitioners, and consultants already recognize the tradeoffs that the IDIB Portfolio makes explicit. For example, consultants will often prescribe a less-than-perfect, heuristic decision-rule because it is less demanding of the information system and/or easier to implement, particularly if these imperfections are relatively inexpensive to buffer against. On a broader level, sensitivity analysis, which is a well-founded tool of the operations research theorist, can be viewed as assessing the impact of imperfections in the quality of the information and/or decision-making systems on system performance, and, hence, the level of buffering that might be required. For example, sensitivity analyses on the basic EOQ model can be viewed as assessing the sensitivity of lot-sizing decisions to inaccuracies (i.e., imperfections) in the information required to support it (e.g., estimates of company inventory-holding cost or set-ups). See Lowe and Schwarz, 1983, for example.

Similarly, but at a metaphysical level, operations-research theorists often prefer a less realistic (i.e., less perfect) model to a more realistic model because of the insight its analysis provides.

From this point of view, the fundamental difference between the IDIB paradigm and the OR paradigm is that experienced operations researchers, practitioners, and consultants make these tradeoff *a priori*; that is, without explicitly identifying all the alternative levels of quality in information, decision-making, and implementation systems that might be chosen. Although this is understandable, such *a priori* choices necessarily lead to locally, not globally, optimal choices for the corresponding IDIB Portfolio.

The Role of Information Technology and Economics

What “happened” at Xerox that led to its development of a new IDIB Portfolio for managing manuals? Was it management’s “discovery” that *their own decisions* created the demand for manuals? Not likely.

Obviously, the availability of technology played a role: printing-and-binding technology that could “quickly” produce an owner’s manual. In IDIB Portfolio terms, technology whose implementation leadtime was shorter than the implementation leadtime to assemble a copier. The other element, of course, is economics: Given that technology *can* facilitate a “new” IDIB Portfolio, its adoption only makes sense if the total cost of the “new” portfolio is less than the total cost of the “old” portfolio. So, obviously, technology and economics play a substantial role in the development of “new” IDIB Portfolios.

Given the proven success of information technology (e.g., microcomputers and the internet) as a significant facilitator of improved information, decision-making, and implementation, I believe it is inevitable that information technology will continue to create the opportunities for “new” IDIB Portfolios. Further, to the extent that the cost of information technology continues to fall, following Moore’s Law², these IDI_ systems will continue to become less and less expensive. And, buffers, “B”s – inventory, capacity, and leadtime – are, if anything, becoming more expensive over time.

Hence, I believe information technology and economics will continue to offer IDIB Portfolios to supply-chain managers whose total cost is

²The observation made in 1965 by Gordon Moore, co-founder of Intel, that the number of transistors per square inch on integrated circuits had doubled every year since the integrated circuit was invented. Moore predicted that this trend would continue for the foreseeable future. In subsequent years, the pace slowed down a bit, but data density has doubled approximately every 18 months, and this is the current definition of Moore’s Law, which Moore himself has blessed. Most experts, including Moore himself, expect Moore’s Law to hold for at least another two decades. *Source: webopedia.com* (definition last modified in March, 1998).

less than today's³. The challenge to supply-chain managers is: Which new IDIB Portfolio to adopt, and when to adopt it? The challenge to supply-chain modelers is: What techniques and models will aid managers in choosing new IDIB Portfolios?

5. An IDIB Portfolio Perspective on Supply-Chain Management

Until fairly recently, virtually every link of every real-world supply chain was managed using a *very* crude IDIB Portfolio. Typically, the only information a buyer shared with its supplier was its current order; and the *most* information the supplier shared with the buyer was this order's planned or promised shipping date. Status information (e.g., order status, inventory status) and future-oriented information (e.g., planned orders and production) were seldom, if ever, shared, often because neither partner had easy access to its *own* information about them. Further, in those instances when one partner *did* have access to this information, the technology for sharing it either wasn't available or was very expensive. As a consequence, supply-chain "partners" – if we can call them that – were figuratively blind to one another. In the absence of useful supplier/customer information for decision-making, each "partner" made decisions that were focused on what little information was available, typically, internally-focused information, such as processor utilization, hot lists, etc. Consequently, the decision-making and implementation based on this low-quality information required *huge* buffers: large buffer inventories of raw materials, work-in-process, and finished goods *plus* buffer leadtimes and capacity.

Given the high cost of these buffers, it's no surprise that when low-cost technology for information-sharing between buyers and suppliers became available, innovators seized the opportunity to substitute low-cost information-sharing for these high-cost buffers, thereby achieved significant improvements in performance and/or reduced total cost.

Wal-Mart, of course, must be credited with introducing many of the technological innovations (e.g., bar-coding, satellite communication of point-of-sale information) associated with contemporary supply-chain

³Although new technological capabilities are likely to be the primary drivers for changing an existing IDIB Portfolio, at least in the short term, it is important to note that even if technology were held constant, any significant change in the cost of one or more of the components of *any* given IDIB Portfolio is an equally-important driver. For example, holding everything else equal, a significant decrease in the cost of capital makes inventory and equipment buffering less expensive. Under these circumstances, increasing the quality of the buffer system and decreasing the quality of the corresponding information, decision-making or implementation system should reduce total cost.

management, and equally important, in *demonstrating* that substituting improved information for inventory and leadtime buffers reduced total cost.

And, so the “revolution” in supply-chain management began and continues today, the state-of-the-art represented by Wal-Mart’s *RetailLink* system.

Yet, what is really different about these systems? From an IDIB perspective, these innovations are actually fairly modest, at least in terms of what *could* be innovated. Under Vendor-Managed Inventory (VMI), for example, the buyer delegates the making and implementation of inventory-replenishment decisions to the vendor. In order to do so, the buyer’s information system provides the vendor with information about customer demand and inventory status. Levels 1 and 2 of Quick Response (QR) are the same, with the addition of automatic order-processing (i.e., implementation) between retailer and supplier. Levels 3 and 4 of QR include cross-docking warehouses (i.e., faster implementation of warehousing decisions). Even Wal-Mart’s *RetailLink*, although awesome in its scale (e.g., in the level of detail provided, in the type and number of partners with access to it), is fundamentally: (1) a system for rapidly sharing transactions-level data and metrics about the past and the present with its suppliers; and (2) a centralized system for making and implementing decisions for its own facilities.

So, why the “significant improvements” from what I describe as “modest innovations”? There is an old saying that: *“In the land of the blind, the one-eyed man is king”*. Historically, supply-chain “partners” were figuratively blind to one another. Add just a little vision – for example, under QR, customer-demand and retailer-inventory information – and, suddenly, there is no need for much of the buffering that had been required in the “land of the blind”. Reducing inventory produces cash. Reducing capacity increases productivity, and, hence, profitability. Reducing leadtimes attracts more customers.

Given the huge payoffs that supply-chain partners have derived from sharing a modest amount of information or from delegating *some* decision-making and implementation from buyer to supplier, will more of the same yield even larger payoffs? Where will it end? Will it end?

The “Axioms of the IDIB Portfolio” suggest that long before information-sharing and/or delegation of decision-making and implementation becomes “total” between supply-chain partners – if it ever does – something even more revolutionary will happen.

6. The Axioms of the IDIB Portfolio

The 1st Axiom of the IDIB Portfolio is this: *Given an existing IDIB Portfolio, increasing the quality of one of its components facilitates decreasing the quality of one or more of its other three components while maintaining the same level of customer service (e.g., fill-rate, leadtime).*

In an inventory-replenishment system, for example, reducing the leadtime to implement a replenishment decision, facilitates decreasing the safety-stock inventory or leadtime buffer – that is, *decreasing* its quality – without affecting customer fill-rates, expected backorders, etc.. Or, in the same setting, decreasing the variance of the leadtime to *make* a replenishment decision by, say, one unit, facilitates increasing the corresponding *implementation* leadtime variance by one unit.⁴ Many such tradeoffs are possible.

Schneider National Trucking Company’s innovation of satellite tracking to locate its trucks allowed Schneider to reduce two of its costly buffers while *improving* customer service. Here are some details:

Before introducing its satellite-based information system for locating its thousands of trucks, Schneider dispatchers relied on periodic telephone calls from its drivers in order to learn where its trucks and drivers were. The corresponding uncertainty about where and when its own trucks would be available led to inflated promised pick-up times to customers and to a significant amount of idle capacity (i.e., “deadheading”, which is a truck moving without a load). In other words, in the absence of accurate information about the location of its own trucks, Schneider buffered itself using leadtime and capacity. By adopting its satellite-tracking system, Schneider was able to reduce both these buffers *and* offer improved delivery performance.

The 1st Axiom of the IDIB Portfolio might be called the “trade-off axiom”; that is, by incurring increased cost for higher quality in one component of the IDIB Portfolio it should be possible to reduce the quality and cost of another component. *If the cost reduction is larger than the cost increase, then total portfolio cost has been reduced.* Further, as is often the case, some of the net savings can be invested in improving competitiveness (e.g., increasing fill-rates, reducing delivery leadtimes, offering higher levels of customization, reducing prices).

Other trade-offs are possible, too. For example, a make-to-stock manufacturer who has made the transition from a “push” management sys-

⁴Assuming these processes are independent, the same safety stock will provide the same customer fill-rate, etc.

tem (e.g., MRP) to a “pull” management system (e.g., JIT) has, everything else being equal, chosen to substitute buffer capacity for buffer inventory (and, often, to shift the inventory-buffering responsibilities to its suppliers and/or customers).

Sometimes, trade-offs are made to *increase* buffers. One example of this, experienced by every supply-chain manager who has replaced a domestic supplier with a lower-cost international supplier, is the increase in safety-stock inventory necessitated by the less-reliable transportation leadtimes from the off-shore supplier. In such cases, of course, the cost saving in implementation (i.e., delivery of items ordered from the supplier) isn’t typically in transportation – indeed, transportation cost may increase – but in the cost of acquiring the materials themselves.

I believe that virtually all of the dramatic improvements reported by companies in managing their own internal supply chains or by companies and their partners in managing their shared supply chains can be interpreted – indeed, could have been forecast – using the 1st Axiom of the IDIB Portfolio. That is, given the reduced cost of a “better” information system (as provided by innovations in information technology) and the already high cost of buffering, that information could be substituted for buffers – that is, information improved in quality and buffering reduced in quality⁵ – *without* reducing customer service *and* at reduced total cost.

What the 1st Axiom *doesn’t* suggest is the magnitude of improvement in cash position, productivity, and competitiveness that so many partners have reported. This magnitude, of course, depends on the cost of the buffering required by being “blind” and how much of a buffer reduction (i.e., reduced inventory, capacity, and/or leadtime) a little “vision” provides.

So, will the future of supply-chain management involve even more extensive information-sharing or delegation of decision-making and implementation between supply-chain partners? The 2nd Axiom suggests *some* more, but, depending on the partnership, perhaps not a great deal more.

The “2nd Axiom of the IDIB Portfolio” is this: ***Investment to improve the quality of any single component of an IDIB Portfolio will, over some range, decrease the total cost of the portfolio; but, beyond some quality level, increase the total cost of the portfolio.***

⁵Recall that, everything else being equal, reducing the amount of buffering reduces its capability to compensate for imperfections. Hence, reducing the amount of buffering reduces its “quality”.

Although “axioms” are supposed to be self-evident, a little discussion of the 2nd Axiom is appropriate. The 2nd Axiom considers what happens to total portfolio cost if one varies the quality level of any single component of an IDIB Portfolio and adjusts the quality levels of the other three components to minimize the corresponding total portfolio cost. For example, consider varying the quality of implementation in some existing IDIB Portfolio. Applied to this example, the first half of the proposition is that if implementation is very low in quality, for example, that implementation leadtime has a large mean and variance, or that the accuracy of implementation is poor, then the total cost of this IDIB Portfolio can be reduced by: (1) *increasing* the quality of implementation; and (2) *reducing* the quality of the corresponding information and/or decision-making systems, and/or reducing the amount of buffering (e.g., leadtime, capacity, or inventory).

The support for this proposition is that the cost of any of the four components of an IDIB Portfolio is an increasing and marginally-increasing function of its quality. In this instance, the proposition is that the cost of a low-quality implementation system is small, and that the cost to improve its quality is low compared to the high cost of the information, decision-making, and buffer systems *required to work in combination* with it. More specifically, given a low-quality implementation system, management would probably have been forced a very high-cost buffer system, possibly one with large inventories, or possibly one with large leadtime and/or capacity buffers. Hence, improving the quality of implementation – which should cost relatively little – and reducing the quality of its buffer system – which should save relatively more – will reduce total portfolio cost.

The second half of the proposition is that continuing to improve the quality of any single component and, correspondingly, decreasing the quality of one or more of the other three components of an IDIB Portfolio will, beyond some point, *increase* total portfolio cost. The support for this proposition, again, is that the cost of any of the four components of an IDIB Portfolio is an increasing and marginally-increasing function of its quality. In terms of the example, the proposition is that if the implementation system is already operating at a very high level of quality – e.g., nearly immediate, nearly perfect implementation – then the cost of any incremental improvement will be large relative to the savings generated by the corresponding decreases in the quality levels of buffering, information, or decision-making that this improvement in implementation facilitates.

The 2nd Axiom might be called the “golden mean” axiom; that is, it doesn’t make sense to invest too much in improving a single component

of the IDIB Portfolio without making corresponding improvements in its other components.

7. The Future of Supply-Chain Management

Given that the current state-of-the-art in supply-chain management involves *some*, perhaps even a great deal, of information-sharing and *some* delegation of decision-making and implementation, the axioms suggest that even if the net savings from more information-sharing and delegation is positive, some *other* change in the IDIB Portfolio may yield *larger* net savings. In other words, the question is *not* whether more information-sharing or more delegated decision-making or more delegated implementation will reduce total portfolio cost. Instead, the question is: *Which changes in which components of the supply chain's IDIB Portfolio will facilitate the largest reduction in total portfolio cost?*

I believe that the most likely candidates for large cost savings are in *collaborative* decision-making and/or *collaborative* implementation. What's "collaborative"?

Broadly speaking, "collaborative" and "shared" mean the same thing, but I use the word "collaborative" in order to make an important distinction between visibility and participation. "*Shared* information" is about visibility: that is, within some given domain, all the partners "see" the same thing. Hence, "shared" decision-making or "shared" implementation might be interpreted to mean decision-making or implementation that is visible to all the partners. Although such visibility is important, *participating* in decision-making or *participating* in implementation means something much more significant. Participation means that both partners' objective functions, constraints, and relative capabilities are considered. From an operations-research perspective, collaborative decision-making and implementation involve *joint* optimization, not independent optimization.

The 2nd Axiom suggests that if partners are doing little or no collaborative decision-making or implementation, then it is possible that the greatest potential for improvement is in precisely these areas. Moreover, the higher the quality of the components of the supply chain's IDIB Portfolio is in *other* respects, the more likely shared decision-making and implementation are to provide the most significant total cost reductions.⁶

⁶In addition, independent decision-making based on the same information, or delegated decision-making and implementation, *at best*, yield locally-optimal decisions and actions.

The IDIB Portfolio, *The Goal*, and the Theory of Constraints

In his ground-breaking book, *The Goal* (Goldratt and Cox, 1985), Eli Goldratt introduced the concept of bottlenecks in a firm's production capacity that limit its "throughput"; i.e., the rate at which a production system generates money through sales. From the perspective of the IDIB paradigm, throughput is a dollar-oriented quality characteristic of a firm's implementation (i.e., production) system. Correspondingly, Goldratt's bottleneck concept is that this quality characteristic – the capability of the implementation system to generate dollars through sales – is limited by the bottleneck process(es) internal to the implementation system.

Of course, a firm's implementation system doesn't generate throughput all by itself. Decisions must be made about what to produce, when to produce it, etc. This is the role of the decision-making system. And, in order to make decisions, the decision-making system requires information that's provided by the information system. Finally, buffer systems are there to provide throughput by compensating for whatever imperfections might exist in the IDI_ systems. In other words, it is a firm's *entire* IDIB Portfolio that generates money through sales, not just its implementation system.

Although Goldratt doesn't recognize the IDIB Portfolio explicitly, his prescriptions most certainly apply to it. For example, Goldratt prescribes that decision-making in the management system should be focused on the "drum" of the implementation-system's bottleneck(s) and the "ropes" that feed it. Goldratt also has prescriptions about the form, location, and amount of buffering that should be provided⁷, and about the nature of the information systems⁸ management should use in decision-making. In brief, Goldratt recommends that all four components of a firm's IDIB Portfolio should be focused on the supply (and demand) bottlenecks. Hence, the IDIB paradigm and the bottleneck paradigm are consistent, indeed, complementary, to one another.

Specifically, I believe that the IDIB paradigm enriches the paradigm of *The Goal* in several ways. For example, the IDIB paradigm suggests that the "goal" of making money can be achieved by virtually unlimited number of different IDIB Portfolios, each component contributing in its characteristic way (i.e., gathering information, making decisions, etc.), each imperfect in different ways, and each compensating for imperfec-

⁷For example, buffer inventory in front of bottlenecks and buffer leadtimes to protect delivery dates.

⁸See Goldratt, 1991.

tions in the others. Second, of course, that the quality of each of these components is a decision variable.

Next, recall that Goldratt points out that the only guaranteed way for a firm to make money is to simultaneously increase throughput and reduce “inventory” and “operating expenses”.⁹ However, Goldratt offers relatively little guidance about *how* to reduce inventory and operating expense, much less minimize them. The IDIB paradigm suggests how: select the least total-cost IDIB Portfolio. In other words, since most of a firm’s operating expenses are driven by management’s chosen IDI_ systems and most of its inventory is in its safety-stock and capacity buffers, by carefully choosing its IDIB Portfolio, management will minimize its combined operating expenses and inventory

Finally, note that the “bottleneck” concept is *imbedded* in the 2nd Axiom. Here’s how: Consider a company with medium-to-high quality decision-making and implementation systems, but a low-quality information system. Under these circumstances, the overall quality of this company’s IDI_ systems is limited by its low-quality information-system. According to the 2nd Axiom, if this company invests in a better information system, the overall quality of its IDI_ systems will improve. This improvement facilitates the reduction in its associated buffers. Given some (low) range of information-system quality, the effect should be to reduce the total cost of this company’s IDIB Portfolio.

Next, consider a company with an excellent implementation system, say a state-of-the-art cellular system, but with mediocre-quality information and decision-making systems. Additional improvement in this company’s implementation system will yield an improvement in its overall IDI_ systems, and facilitate a reduction of the corresponding buffers. However, the money saved on the buffer system may be less than the additional cost of the improved implementation system. The result is an increase in total portfolio cost.

From Goldratt’s viewpoint, the “bottleneck” in the quality of first company’s IDI_ systems is its information system: a dollar invested there yielded more than a dollar saved on the buffer system, thereby reducing the firm’s total inventory and operating expenses. On the other hand, the “bottleneck” in the quality of second company’s IDI_ systems *wasn’t* its implementation system. Hence, a dollar invested there is a dollar wasted.¹⁰

⁹According to Goldratt, “inventory” is the dollars that a firm has invested in things that it intends to sell, while operating expense is the cost of things that the firm does to turn inventory into throughput.

¹⁰Of course, neither of these improvements, whether they would reduce its total IDIB Portfolio cost or not, necessarily increases the company’s throughput. Throughput would only be

The Future is Wow for Some Supply-Chain Partners

Some supply-chain partners are already sharing decision-making. Since 1995, Heineken USA, Inc. and its independent distributors have been sharing information and decision-making about the replenishment of Heineken's beer products under a system called HOPS: the Heineken Operational Planning System (<http://64.158.250.111/news/heineken.html> and <http://64.158.250.111/news/archive99/06091999.html>). Intel and its customer computer-assemblers (e.g., IBM, Dell, Compaq) have been using a collaborative information-sharing and decision-making system to manage the assembler's inventories of computer chips under Intel's Supply-Line Management (SLM) program.

One well-known and widely-implemented system for information-sharing and collaborative decision-making in supply chains is Collaborative Planning, Forecasting, and Replenishment (CPFR).

8. Collaborative Planning, Forecasting, and Replenishment (CPFR)

CPFR is a process model, shared by a buyer and supplier, through which inventory-status, forecast-, and promotion-oriented information are shared and replenishment decisions are made. In IDIB Portfolio terms, a process model for sharing information and decision-making between a buyer and supplier.

CPFR began with a pilot program between Wal-Mart and Warner-Lambert, called CFAR: "Collaborative Forecasting and Replenishment". In 1997, the Voluntary Interindustry Commerce Standards (VICS) Association (<http://vics.org>) developed the "CPFR Initiative" (<http://www.cpfr.org>). In 1998, VICS published the first "CPFR Guidelines" (<http://www.cpfr.org/Guidelines.html>). Since then, a large number of partners have developed CPFR pilots. Appendix 1 provides a partial list. Several partnerships have subsequently adopted CPFR as a standard way of doing business with one another.

The CPFR Process

CPFR consists of 9 process steps, as follows:

- Step 1.** Develop Front-End Agreement: Roles, Measurement, Readiness

increased if the changes to the entire IDIB Portfolio increase availability or otherwise make the company's products more competitive. Nonetheless, even if throughput isn't increased, reducing the total IDIB Portfolio cost reduces a company's operating expense, which increases profits; i.e., makes more money.

- Step 2.** Create Joint Business Plan: Strategies and Tactics
- Step 3.** Create Sales Forecasts: Buyer and supplier both create customer-demand forecasts
- Step 4.** Identify Exceptions in Sales Forecasts
- Step 5.** Resolve Exceptions: Agree on single forecast or agree to disagree
- Step 6.** Create Order Forecasts: Buyer and supplier both create plans for buyer orders
- Step 7.** Identify Exceptions in Order Forecasts
- Step 8.** Resolve Exceptions: Agree on single plan for buyer orders
- Step 9.** Order Generation

More details about these steps and the roles of the buyer and supplier in each step are provided in Figure 1.1 and at the VICS CPFR website (<http://www.cpfr.org>). However, the basics of CPFR are straightforward: First, the partners share information about demand. If the buyer is a retailer – and so far, most buyers using CPFR *are* retailers, then demand is retail customer demand. If the buyer is a manufacturer or assembler then demand is generated by the manufacturer or assembler’s trial master-production schedule. Then, significant differences between the buyer’s and seller’s demand forecast, labeled “exceptions”, are discussed and resolved. These are Steps 3-5 above. Then, buyer and supplier share plans for orders that the buyer will place with the supplier, based on the shared demand forecasts. Again, exceptions are identified and resolved (Steps 6-8). Subsequently, using the shared order plan, actual orders are generated (Step 9). The foundation for Steps 3-9 is the so-called “front-end agreement”, under which the roles of the buyer and supplier, and their capabilities to perform these roles are assessed. In this step, targeted performance and measures are also adopted. In Step 2, strategies and tactics are specified in detail.

The benefits reported by CPFR partners, as might be predicted by the axioms, are increased inventory turns (i.e., lower buffer inventory) and increased fill-rates for the SKU’s involved; that is, higher levels of customer service.

Several consulting firms offer software systems and support for CPFR, among them Logility, Inc. (<http://www.logility.com>) and Syncra Systems, Inc. (<http://www.syncrasystems.com>). CPFR is also being implemented on B2B exchanges such as Worldwide Retail Net, Transora, and NetXchange.

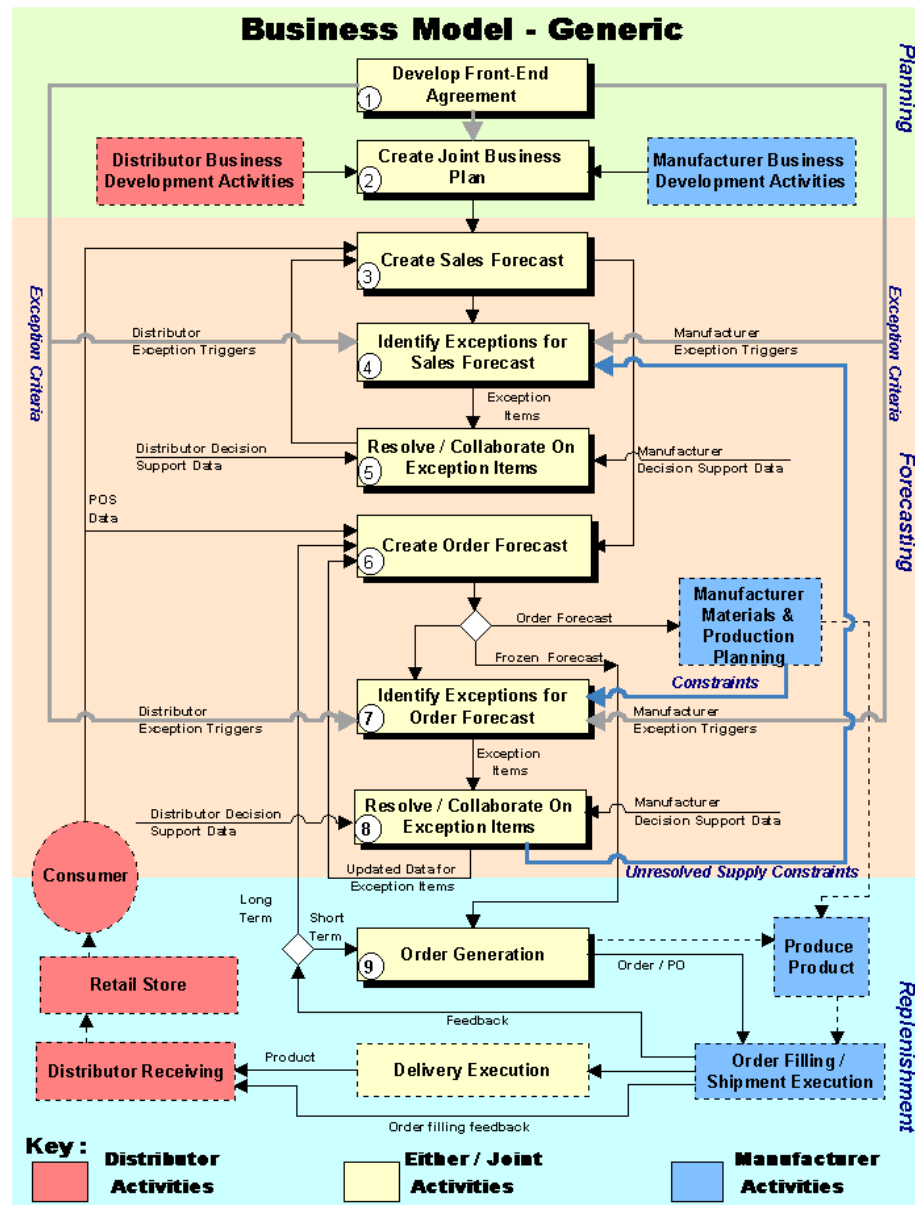


Figure 1.1. The 9-Step CPFR Process.

Based on the success of CPFR between single buyer-supplier pairs, thought leaders in CPFR have suggested its extension to include collaboration with the carrier that transports goods between the buyer and supplier. This is called “CTM”: Collaborative Transportation Management (<http://www.cpfr.org/WhitePapers/CTMwhitepaper.pdf>). It has also been suggested that in order to be truly successful, collaboration should involve all of the links of a supply chain, under a scheme labeled “n-Tier Collaboration” (<http://www.cpfr.org/WhitePapers/nTierProposal.doc>). Under such a scheme, for example, not only would Kimberly-Clark collaborate with Wal-Mart to plan Wal-Mart’s orders of Kimberly-Clark products, but Kimberly-Clark would, in turn, collaborate with its suppliers to determine Kimberly-Clark’s orders for materials for its paper products. So, in effect, Wal-Mart and Kimberly-Clark’s suppliers would be collaborating with Kimberly-Clark to determine plans to supply and order Kimberly-Clark’s paper products.

The Challenges of CPFR

Although CPFR has enormous potential for reducing the total cost of any supply chain’s IDIB Portfolio, there are also enormous challenges. At the most fundamental level, buyers and suppliers must develop trust that each will treat the other fairly and honestly. Prerequisite for this are incentives to do so. Again, see Cachon, 2004 for a review of the supply-chain literature on the management of incentive conflicts with contracts. On a technical level, buyers and suppliers must develop a common language for identifying products and making decisions about them. (See <http://www.cpfr.org/WhitePapers/CollaborationDataModelingA.pdf>). Similarly, systems must be developed for linking the buyer’s and supplier’s business processes. This will involve a great deal of system change and training. Third, security protocols must be implemented that will safeguard both partners from leaks of proprietary information.

Nonetheless, I believe that buyers and suppliers who find ways to overcome these challenges will achieve a competitive advantage, particularly for products that compete primarily on price and availability (e.g., consumer products). This advantage will either force the competitors of CPFR partners to adopt similar techniques or force them out of business.

9. Research Topics in CPFR

CPFR poses many interesting questions for supply-chain researchers, questions whose answers involve models that have yet to be developed.

At the broadest level, there are questions involving the “drivers” for collaboration. For example, what are characteristics of buyers, suppliers,

and the environments in which they operate that promotes a desire on either of their parts to collaborate using CPFR? It is well known that agency losses occur in decentralized supply chains that involve hidden information and/or hidden actions. What is less well known are the circumstances under which *both* the buyer and supplier will be better off by collaborating. See, for example, Monahan, 1984; Weng, 1995 and Taylor, 2001. In the absence of these circumstances or other incentives for it, collaboration, in general, and CPFR, in particular, seem doomed.

With respect to CPFR in particular, consider a buyer-supplier pair. Among the most interesting research questions are: First, how should the “front-end” agreement be structured in order to maximize – or merely, increase – the profits for both partners? What role will each partner play? How should performance improvement be measured? Perhaps most important: how should the benefits of improved performance be shared between the buyer and the supplier so that both will be better off?

Another set of interesting questions involves defining the elements of the data to be shared. For example, given the cost of data processing and security considerations, should SKU-level data be shared or should only aggregate information be shared? What aggregation/disaggregation procedures are best?

Finally, given that there are costs and benefits associated with exception-processing, regardless of *how* they are processed, how should exceptions be defined? More fundamentally, what does it mean to process an exception; that is, given a significant difference between and the buyer’s and seller’s forecast (planned orders), how should the difference be resolved?

10. Research Topics Related to Supply-Chain Collaboration

Another set of interesting research questions involves the examination of supply-chain collaboration in general, whether or not the tools used are those of CPFR. For example, given a supply chain with some given level of collaboration in information-sharing, decision-making, implementation, and buffering – including no collaboration whatsoever – then, *assuming* that collaboration will be increased, *which links* should be involved, and *in which components* of the IDIB Portfolio will collaboration yield the largest payoffs to the partners involved and to the entire supply chain? Further, how will the benefits of collaboration be measured and shared among partners and along the chain?

An equally interesting set of questions involves collaboration itself. For example, consider a supply chain with some given level of collaboration in information-sharing, decision-making, implementation, and buffering – including no collaboration whatsoever. Then, given the costs and challenges of collaboration, is collaboration the most cost-effective way to improve supply-chain performance? In particular, improving the quality of one or more of the components of one of the partner’s IDIB Portfolios may yield larger improvements at lower cost. If so, then, again, how will benefits be measured and shared?

11. Research on the IDIB Portfolio

In generalizing the OR paradigm, the IDIB Portfolio paradigm provides rich and challenging opportunities for operations researchers.

I believe that the OR paradigm has focused almost entirely on modeling only a single component of supply-chain IDIB Portfolios, typically on the decision-making component. To the extent that the associated information and implementation systems are represented at all, these models treat them as fixed, as costs, or as constraints, but *not* as decision-variables. For example, such models typically take the quality of the information system as given (e.g., demand is known or given by some particular probability distribution with fixed parameters) and implementation systems are represented as parameterized leadtimes or by their cost drivers (e.g., the parameterized cost of performing a set up or the unit cost of holding inventory). The *form* of buffering (i.e., inventory, leadtime, or capacity) is usually also fixed. In general, the objective function is to select the decision (or decision-rule) that minimizes the associated implementation costs plus the too-much and not-enough costs associated with the buffer system. Other models, in particular, those that treat either the information system or the implementation system as a decision variable, typically ignore the *other* three components of the IDIB Portfolio as decision variables.

That’s the “bad news”.

The “good news” is that the OR paradigm’s focus has facilitated the development of fairly sophisticated decision rules and provided limited insight into information and implementation systems. It is also “good news” that many of the techniques employed in the OR paradigm, for example, well-proven estimation and optimization techniques, can be applied to improving, if not optimizing, the IDIB Portfolio model of a management system. For example, cost-estimating techniques that are already being applied to estimating the cost of a given implementation

system can be applied to determining the cost of *alternative* information, implementation, and buffer systems.

Notwithstanding the availability of OR tools, the challenges of IDIB Portfolio optimization, even if done heuristically, are daunting. Recall, the “optimal IDIB Portfolio” is the portfolio of information, decision-making, implementation, and buffering whose combined cost *plus* the cost of failing to provide precisely what was demanded or required (e.g., back-ordering or expediting cost) is the minimum possible cost. Hence, IDIB Portfolio optimization involves *four* decision-variables, not simply one. Further, in order to be useful in selecting an IDIB Portfolio, IDIB models *must* be able to capture complex cost-quality interactions *among* these components, not just the cost-quality characteristic of each component.

Practice-oriented research also presents interesting research opportunities. For example, consider a firm’s current IDIB Portfolio and assume management has the desire to “improve” it. The question is: Where (i.e., which component) is the “bottleneck” of the current IDI_ systems and what improvement should be made to it? *Before* this question can be answered, operations researcher must first develop techniques and measures to assess the capabilities of each of its components. Next, in order to make the most cost-effective improvement, managers will require *a priori* estimates of the marginal costs and benefits of the *next* increment in quality in their current components. Further, there are questions relating to the timing of IDIB Portfolio changes. For example, although adopting current technology might reduce total portfolio cost, should management forego adopting it, and, instead, wait for new technology that might provide even larger total cost reductions? Technology forecasting must play a role in answering this question.

Despite the daunting nature of the challenges to improving a firm’s current IDIB Portfolio or optimizing a proposed IDIB portfolio, operations researchers will be foregoing a tremendous opportunity if they continue to focus on optimizing one component of the IDIB portfolio at a time. Worse, they commit the Cardinal Sin of Operations Research: *Sub-Optimization*.

12. Some Related Literature

Most of the literature in supply-chain management can be viewed from the perspective of the IDIB Portfolio. As explained above, I believe that most of this literature examines the choice of the decision (e.g., order quantity, target inventory) or decision-rule (e.g., EOQ) to minimize total too-much and not-enough costs. See the discussion of the newsvendor model above, for example. Some of this literature also examines the

impact of implementation issues on decision-making. For example, in a highly-stylized way the EOQ formula represents the impact of set-up (i.e., implementation) cost on the optimal order quantity. However, there are some signs of interesting developments.

A growing body of literature examines the value of information-sharing in managing a supply chain. Lee et al., 2000, for example, have shown that information-sharing can dampen the so-called “bullwhip effect” so often observed in supply chains. Chen et al., 2000, have shown that the bullwhip effect can be reduced by centralizing information. Cachon and Fisher, 2000, study the value of sharing demand and inventory-level information in a supply chain. Aviv, 2001, examines the effect of collaborative forecasting on supply-chain performance.

A small, but growing body of research involves collaborative replenishment decision-making to take advantage of shared information. For example, Song and Zipkin, 1996, develop an inventory-replenishment policy to take advantage of information about supply conditions. Aviv, 2002, examines joint forecasting and replenishment processes. Iyer and Ye, 2000, develop a model to assess the value of information-sharing in a retail environment in which retailers share promotional information with their suppliers.

However, I am not aware of *any* research that addresses itself to “managing the IDIB Portfolio”; that is, making decisions about the nature and quality of *all* its four components: the information, decision-making, implementation, and buffer systems.

Nonetheless, there are two papers of particular relevance to “managing the IDIB Portfolio” that I would like to draw attention to:

Milgrom and Roberts, 1988, develop a model of a manufacturer and the market for its products, and, for this particular model, establish that the optimal management system will either produce to inventory or produce to customer order. In IDIB Portfolio terms, Milgrom and Roberts examine a model in which there is a cost to acquire information about market demand and a cost to produce. Their analysis concludes that a profit-maximizing firm will either: (1) acquire no additional information about customer demand and produce entirely to inventory; or (2) acquire complete information about customer demand and produce entirely to customer order.

It should be noted that Milgrom and Robert’s result is specific to the assumptions of the model they propose. Other assumptions would yield different results. For example, Milgrom and Robert’s model ignores the production (implementation) leadtime, customer preferences, and competition. Given significant production leadtimes and customer preference for availability over variety, a manufacturer that might otherwise

choose make-to-order might be forced to make to stock. Alternatively, given customer preference for variety over availability, a manufacturer that might otherwise choose make-to-stock might be forced to make to order. Nonetheless, Milgrom and Roberts are the first to suggest tradeoffs between information and inventory (i.e., one form of buffer).

Hariharan and Zipkin, 1995 implicitly suggest tradeoffs between the characteristics of a company's information system and implementation system. In particular, within a supply-chain setting, they establish the equivalence between the information system's ability to "see" one more (less) period of demand into the future and the ability of the implementation system to produce one period faster (slower).

The use of the "ICB Portfolio" paradigm, which is closely related to the IDIB paradigm, in teaching operations management can be found in Schwarz, 1998.

13. Summary

In this chapter, I have introduced a new paradigm for management in general and for managing supply chains, in particular, called the IDIB Portfolio. The IDIB Portfolio, which can be viewed as a generalization of the OR paradigm, takes the view that the quality of *all* four components of a management system – not just the buffer system quality (i.e., size) – are decision variables; and that the underlying quality-cost function for each of these components – that is, the cost to move them in the direction of perfection – is increasing and marginally increasing. Finally, that there is a cost associated with the entire IDIB Portfolio – again, not just the buffer system – for failing to provide whatever is required or demanded. Hence, the "optimal" IDIB Portfolio is the portfolio that minimizes the total costs of all of its components *plus* the cost of failing to provide whatever is required or demanded.

The IDIB Portfolio and its axioms provide insight into the evolution of supply-chain practice to date, and, I believe, suggest that the future of supply-chain management practice will involve a significant level of collaborative decision-making and implementation: a level of collaboration in decision-making and implementation that is comparable to the level of information-sharing in contemporary supply-chain management. I have described the VICS CPFR initiative as one example of this future. Finally, I have suggested several research topics involving CPFR and, more broadly, challenging new research into the IDIB Portfolio paradigm.

Acknowledgement

I am indebted to two anonymous reviewers for their helpful comments on an earlier draft. I am particularly indebted to Nick Petruzzi for many insightful suggestions, including the notion that the IDIB paradigm generalizes the OR paradigm and the links between the IDIB Portfolio paradigm and Goldratt's *Goal*-oriented paradigm.

Appendix: List of Buyers and Suppliers participating in CPFR partnerships

Source: The VICS CPFR[®] Matrix, January 2002. This list is maintained monthly by MoonWatch Media, Inc. See <http://www.retailsystems.com/communitycenters/cccc/cpfirmatrix.pdf>.

Table 1.A.1. *Buyers*

10 Internal Affiliates	McDonald's France
4 Retailers	Meijer
850 n-Tier Partners	Mervyn's
Ace Hardware	Radio Shack
Albertson's	RiteAid
Best Buy	Royal Ahold
Canadian Tire	RONA
CVS	Safeway
Dansk	Safeway (UK)
Dealers	Sainsbury
Delhaize le Lion	SAKS
Distributors	Sears Roebuck
Do It Best	Somerfield
Eckerd	Sports Authority
Federated Department Stores	Staples
H.E. Butt	Superdrug
Home Depot	Target
J.C. Penney	Tesco
Jusco	TruValue
Londis	Walgreens
Marshall Field's	Wal-Mart
Match Supermarket	Wickes Furniture
McDonald's	Woolworth UK

Table 1.A.2. Suppliers

12 Suppliers	Levi
20+ Suppliers	Levi Strauss
Ashley Furniture	Liquid Nails
Ball Sports	Liz Claiborne
Black & Decker	Manco
Broyhill	Mars
Chanel	Master Lock
Chapin	Meriat
Colgate-Palmolive	Mitsubishi Motors
Compaq	Nelson
Eastman Chemicals	Nestle UK
ECPG3	New Balance
Eli Lilly	Pacific Coast
Feather Fruit Growers' Cooperative	Panasonic
FujiFilm	Philips Consumer
GE Appliances	Pillowtex
General Mills	PlumbPak
Genovs	Polo Ralph Lauren
Georgia Pacific	Proctor & Gamble
Harley-Davidson	Reynolds Metals
Hasbro	Rowe Companies
Heineken	Sara Lee
Henkel	Schering-Plough
Herlitz	Solo Cup
Hewlett-Packard	Spectrum
HYKo	Thomson Electronics
Inland Paperboard & Packaging	Timberland
International Paper	Truya
John Deere	Unilever Argentina
Johnson & Johnson	Vandemoortele of Belgium
Kao	Warner-Lambert
Kimberly Clark	Whitehall Robbins
Kraft	Woodstream
Lever-fabrege	YKK

References

- Aviv, Y. (2001). The effect of collaborative forecasting on supply-chain performance. *Management Science*, 47(10):1326–1343.
- Aviv, Y. (2002). Gaining benefits from joint forecasting and replenishment processes: The case of auto-correlated demand. *Manufacturing & Service Operations Management*, 4(1):1–18.
- Cachon, G.P. (2004). Supply chain coordination with contracts. In Kok, A.G. de, and Graves, S.C., editors, *Handbooks in Operations Research*

- and Management Science: Supply Chain Management*. North-Holland, Amsterdam, The Netherlands. Forthcoming.
- Cachon, G.P. and Fisher, M. (2000). Supply-chain inventory management and the value of shared information. *Management Science*, 46(8): 1032–1050.
- Çetinkaya, S. and Lee, C. (2000). Stock replenishment and shipment scheduling for vendor-managed inventory systems. *Management Science*, 46:217–232.
- Chen, F., Drezner, Z., Ryan, J.K., and Simchi-Levi, D. (2000). Quantifying the bullwhip effect in a simple supply chain: The impact of forecasting, leadtimes, and information. *Management Science*, 46(3):436–443.
- Ehrhardt, R. and Taube, L. (1987). An inventory model with random replenishment quantities. *International Journal of Production Research*, 25:1795–1804.
- Fiorito, S.S., May, E.G., and Straughn, K. (1994). Quick response in retailing: Components and implementation. *International Journal of Retail and Distribution Management*, 23:12–21.
- Goldratt, E. (1991). *The Haystack Syndrome: Sifting Information Out of the Data Ocean*. North-River Press, Groton-on-Hudson, New York.
- Goldratt, E. and Cox, J. (1985). *The Goal*. North-River Press, Groton-on-Hudson, New York.
- Hariharan, R. and Zipkin, P. (1995). Customer-order information, lead-times, and inventories. *Management Science*, 41:1599–1607.
- Iyer, A.V. and Ye, J. (2000). Assessing the value of information-sharing in a promotional retail environment. *Manufacturing & Service Operations Management*, 2(1):128–143.
- Karlin, S. (1958). One-stage models with uncertainty. In Arrow, K.J., Karlin, S., and Scarf, H., editors, *Studies in the Mathematical Theory of Inventory and Production*, chapter 8. Stanford University Press, Stanford, California.
- Lee, H., Padmanabhan, V., and Whang, S. (1997). Information distortion in a supply chain: The bullwhip effect. *Management Science*, 43:546–558.
- Lee, H., So, K.C., and Tang, C.S. (2000). The value of information-sharing in a two-level supply chain. *Management Science*, 46(5):626–643.
- Lowe, T. and Schwarz, L.B. (1983). Parameter estimation for the EOQ lot-size model: Minimax and expected value choices. *Naval Research Logistics*, 30(2):367–376.

- Milgrom, P. and Roberts, J. (1988). Communication and inventory as substitutes in organizing production. *Scandinavian Journal of Economics*, pages 275–289.
- Monahan, J.P. (1984). A quantity-discount pricing model to increase vendor profits. *Management Science*, 30:720–726.
- Schwarz, L.B. (1998). A new teaching paradigm: The information/control /buffer portfolio. *Production and Operations Management*, 7:125–131.
- Song, J.S. and Zipkin, P.H. (1996). Inventory control with information about supply conditions. *Management Science*, 42(10):1409–1419.
- Taylor, T.A. (2001). Channel coordination under price promotion, midlife returns, and end-of-life returns in dynamic markets. *Management Science*, 47:1220–1234.
- Weng, Z.K. (1995). Channel coordination and quantity discounts. *Management Science*, 41:1509–1522.
- Yano, C. and Lee, H. (1995). Lot sizing with random yields: A review. *Operations Research*, 43:311–334.