Planned spontaneity for better product availability

Sridhar R. Tayur
Tepper School of Business, Carnegie Mellon University
Pittsburgh, PA 15213
stayur@andrew.cmu.edu

September 24, 2012

This invited article reflects the particular preoccupations of the author, which concern the parallel and connected development of academic models for managing inventories and their actual applications in industry.

Some *centralized* production-inventory models that have been of service to contemporary business practices in *decentralized decision making settings* are discussed here. I have limited my examples to those in which I have had direct experience. I have attempted to highlight four key points. (a) The strategic role of inventories in capturing revenue and market share, in addition to their traditional role as buffers, in our contemporary "customer-scarce Schumpeter's market." (b) Facilitating the acceptance of model outputs by decision makers across organizational silos and even across firm boundaries using "management mechanics." (c) Development and use of a general-purpose algorithm utilizing Infinitesimal Perturbation Analysis (IPA) derivatives. (d) The emergence of Enterprise Inventory Optimization (EIO) as mainstream software.

Introduction

Edgar Allan Poe, *The Murders in the Rue Morgue:*

> "I am not now writing a treatise, but simply prefacing a somewhat peculiar narrative by observations very much at random."

Several Operations Management (OM) issues in Fortune 500/Global 2000 (or smaller) firms involve the production of physical goods in a repetitive manner to be sold over a horizon of time. One prominent issue companies face is how best to invest in working capital, in the form of inventories, to achieve a business purpose. Some inventory models that have served useful purposes in such firms are the subjects of this article.

Practical problems in Operations Management (PPOM), like other problems in business faced by senior executives in their line of business (LOB), frequently have three features: (1) decisions are made under uncertainty; (2) more than one decision has to be made, either sequentially or jointly; and (3) several actors across different organizations, sometimes in different firms, need to "weigh in" (or be convinced, thus requiring "buy in") prior to implementing a chosen set of decisions.

Before we get to the inventory models and the examples of their successful application, it may help to place PPOM in a capitalist's worldview -- that of Joseph Schumpeter -- eloquently described by David Reisman in *Schumpeter's Market: Enterprise and Evolution (2004)*:

> "The market is supply and demand. Seen as statics, it is a gravitational field that produces equilibrium price as if guided by an omniscient auctioneer. Seen as dynamics, it is a voyage of discovery that, powered by profit or driven on by challenge, need never arrive at its final point of rest…. Enterprise is relentless transformation. Newness is the leitmotiv. Capitalism is newness."

In this paper I will not discuss equilibrium models (although they are useful), but rather highlight decision support models that have helped OM executives compete in this dynamic world of relentless change. This is in line with Koopmans (1975): "…*the production programs of the individual plant or the enterprise for a short period ahead*." A major difference, of course, is that we operate in a capitalist economy where firms compete to create and capture demand in an industry vertical, not in the planned socialist setting (or the central planner situation) for which Kantorovich and Koopmans developed their models. Thus, the models that are described here are not linear and deterministic, but non-linear and stochastic.

Also, over the years I have found Drucker's (1955) view of the role of Management Science has been very helpful:

> "Management Science should not aim at showing the manager 'one right solution.' It must define the 'right question,' and must bring out a full range of alternative solutions…Management Science should aim at showing the Manager what to expect from a given course of action and at warning him when events fail to live up to expectations. It should supply him with the vision needed to make rational decisions in respect to the business enterprise…The aim of Management Science should be to arm the manager's imagination."

This capitalist viewpoint differs from that which guided some of the first set of models after the Great Depression, around and after World War II. To see this, consider the last sentence of the Foreword (by A. R. Marchenko) prefacing the English translation of Kantorovich (1939) [*Management Science* (1960)]:

> "We hope this monograph will play a very useful role in the development of our socialist industry."

Coincident with the difference in world-views, the "customer-friendly" organizations that occupy the business world today reflect a fundamental shift from the "production-scarce" economy of the past to the "customer-scarce" economy of the present, where providing superior consumer convenience, personalization and immediate gratification is a central component of competitive strategy in firms. I coined the term "planned spontaneity" (in 1997, presented at Supply Chain Thought Leaders Conference, Tayur, 1999) to describe what companies are doing on their *supply side* to *create* (and capture) customer demand by making "product availability" an important part of their business strategy.

Product availability has three components: (a) providing *choices* (through a broad portfolio of products, called product variety), (b) providing *access* to the products either as a "pick-up" or a "delivery" (measured by service time and service level) and (c) fine tuning the *price* at which the product and access are offered. Therefore, through choices, access and price, firms determine the strategic role of inventories in creating and capturing demand. For an example, see the responsive supply chain for Caterpillar in Rao, Scheller-Wolf and Tayur (2000): In order to meet uncertain demand for a wide assortment of attachments, a dealer's own inventory is augmented with strategic inventory, pre-positioned (either at an upstream Caterpillar location or at another dealer) to be sourced after demand occurs. This allowed Caterpillar, a new entrant to this market segment, to grab market share from the entrenched leader through product availability while also creating a larger overall market.

Another example is the design of a postponement strategy (using *vanilla* boxes) to effectively provide mass customization at IBM. In Swaminathan and Tayur (1998), the inventories and structure of vanilla boxes are planned in advance so as to be able to react to customer demands quickly, providing a dazzling array of options, without having a large risk of unused inventory. A third example of planned strategic inventory (at a GE plant) is achieved by designing stochastic cyclic schedules that guarantee the availability of a set of components with correlated demands. This allows GE to manufacture the products (after seeing the demands) – make to order -- within a short lead time in a consistent manner, as described in Anupindi and Tayur (1998) and Tayur (2000). The two preceding examples show how companies manage their "push-pull" boundary: you "push" and stock components and vanilla boxes based on forecasts (and their errors); you "pull" these components and vanilla boxes to make the final products after demands are firm. A fourth example (again at Caterpillar) describes how stable product availability to customers, segmented by machine type, is achieved for various customer types. Caterpillar accomplishes this by optimally positioning raw materials, semi-finished products and ready to configure machines across a 4-stage supply chain spanning their suppliers, their own factories producing "base-models" that are ready to be configured, and independent dealers that carry fully configured machines. This "extended multi-enterprise" approach to inventory management increased their revenue by 2% while lowering the overall inventory investment by 15% (Keene et al, 2006).

The above examples (Table 1) fit the definition of "planned spontaneity" – by effectively planning inventory firms can "spontaneously" react to consumer demands.

Table 1 Some Practical Problems in Operations Management (PPOMs)

Business Problem	Company	PPOM	Inventory Model	References
Enter a new market	Caterpillar	PPOM-1: Rapid response supply chain	Discrete-time, stochastic, dual lead time, partial lost sales, two customer classes	Rao, Scheller-Wolf and Tayur (2000)
Reduce \$1b of inventory without reducing service level	Deere	PPOM-2b: Operate a multi-stage SC with lowest inventory investment while meeting service levels	Discrete time, 2-stage distribution, non-stationary; Postponement at upper stage, independent dealers at lower stage	Troyer et al (2005)
Stabilize availability	Caterpillar	PPOM-2b	Discrete-time, assembly and distribution, 2,3 and 4 stages, non-stationary, multiple customer classes	Keene et al (2006)
Manage broad product variety	IBM	PPOM-3: Vanilla Boxes	Discrete time, multi-product, common components, capacitated, stochastic 2-stage stochastic program with recourse	Swaminathan and Tayur (1998)
Quote accurate lead times	GE	PPOM-4: Stochastic Cyclic Schedules	Continuous time, chance constraints, multi-product, changeover times and costs, stochastic	Anupindi and Tayur (1998), Tayur (2000)

This strategic role of inventories is in stark contrast to the earlier literature in OM, as observed by Arrow, Karlin and Scarf (1958):

> "As we have already noted, virtually all work in inventory theory assumes that the demand is independent of the firm's control."

Furthermore, the ubiquity of high performance computing today in the enterprise information technology infrastructure that businesses have invested in (called the "IT stack") facilitates a complete *re-imagining* of what is possible with respect to inventory models and algorithms to solve complex models. We must not be limited by the conditions prior to the advent of computers. There is no need to develop simple formulas ("closed-form" or otherwise) that can be quickly done on paper for special situations like "stationary 3-stage un-capacitated, single item, serial systems facing Poisson demands" (Whitin, 1953), or to "hand over the task of finding a solution to a special computational assistant; then he can check the solution in 10 or 15 minutes with no difficulty at all" (Kantorovich, 1939). Instead, we can have a general-purpose algorithm that can handle significant complexity and scale (and run unattended by humans), providing solutions in seconds to several global locations simultaneously across firm boundaries. This is called "extended" (or multi-enterprise) inventory planning and optimization.

This article is the result of an opportunity to spend part of a summer in reading and reflecting upon centralized production-inventory models as they pertain to aiding decision-making in a decentralized setting. What is offered here is one person's perspective, based on experiences and observations in a journey spanning 20 years.

Two Practical Problems in Operations Management

Let me start with where Schumpeter may have begun, if he were to be interested in PPOM, with an example of "competitive capitalism," that is, *entry*. As in Reisman (2004):

> "The interloper wins because he mounts a credible challenge. The incumbent loses because he no longer has a capacity to compete. In general it is not the owner of stagecoaches who builds railways. Rationality and bureaucracy stamp out initiative. Creative destruction would disappear into the locked box of self-perpetuation."

Caterpillar enters a new market

In 1997, Dr. George Cusack from Caterpillar's Technical Services Division (TSD), at the request of a senior executive in a line of business at Caterpillar, contacted me with a marvelous request: could we, professors at Carnegie Mellon University, design a responsive supply chain for their new line of products – smaller machines as opposed to their large equipment -- so that they could penetrate a new market segment (for them) that was growing at a very healthy rate but was currently dominated by a competitor? Caterpillar executives had decided on a strategy to grab market share from an entrenched leader: they were not going to compete on price (consistent with their high quality brand image) but rather on *availability*, both in terms of the product *variety* offered and off-the shelf (or within a very short time window) *service level*.

How could we design a supply network with a *make-to-stock* strategy, possibly very different from their existing physical footprint, material flows and *operating policies* for their *make-to-order* large machines, but still take advantage of the supply base and facilities, for sourcing, production, assembly and distribution? How could we "piggy-back" exceptional skills in *service parts distribution* that they already possessed, and leverage those skills and information technology? How would we provide exceptional availability -- without charging a premium -- but still have good profit margins? How to account for the variety of products (machines and attachments) that differ in their complexity and margins, as well as various customer segments with different characteristics? As they were going to be selling these new products through their existing network of independent dealers (who exclusively sold Caterpillar products), how could we get the dealers to co-operate in this strategy, one that would require some to share their inventories and provide lateral *transshipment*?

How could we *operationalize a business strategy* such as this?

Beyond the obvious importance of a project like this (to Caterpillar), what was especially attractive to us was the mental delight that this could (and it did) bring us. As Veblen accurately pointed out in *The Instinct of Workmanship and the Irksomeness of Labor (1898)*, we have "a taste for effective work, and distaste for futile effort."

We had to determine how to "design in" responsiveness in a physical supply chain. *Network design* had always been done using (largely) deterministic models that used mixed-integer-linear programming (MILP) to propose efficient (lowest cost) solutions. But this was very different: responsiveness inherently supposes a certain amount of *ex-ante uncertainty in planning*, followed by an *ex-post reaction to uncertainty resolution*. What stochastic inventory (and other OR) models should we use? What data was needed? If the data were to be available, how could we acquire it? How would we know if the data were accurate? What if the data change over time due to the dynamic nature of the business and the industry? What if the data were not available? How could we get useful estimates? And then could we make rational decisions with incomplete and imperfect data?

This reminded me of Sherlock Holmes in Arthur Conan Doyle's *The Adventure of the Copper Beeches*:

> "Data! data! data!", he cried impatiently. "I can't make bricks without clay."

How would we know that the "answer," however found, is worth implementing? How could we convince the executives to spend tens of millions of dollars, allocate people and other resources towards our answer? How could we convince the independent dealers to go along? Let us call this "PPOM-1: Designing a responsive and efficient supply chain to which executives agree and their business partners co-operate." Let us get back to this later.

Fortunately, this was not the first time when a senior executive in the line of business of a Fortune 500/Global 2000 company had requested a "solution" to a pressing problem of importance (called "the burning platform").

IBM Executive requests a "Inventory-Service level" trade-off curve

Earlier, in 1991, it was at IBM Worldwide Logistics Headquarters (Somers) when Dr. Pasumarti Kamesam, a Member of Technical Staff, had brought me in for a summer to develop an inventory-service level trade-off curve for their disk-drive business, something that a very senior executive (a "Global EVP") really wanted to see. Let us call this "PPOM 2: Generating an Inventory-Service level tradeoff curve for an existing supply chain and actually making changes to current parameters and policies". (A single-stage, stationary version had been studied as "stock-exchange" curves by R. G. Brown, 1967.) Here too, implementation might require co-operation from business partners, as well as other members in the firm not reporting to this executive's line of business. This practical problem was intended to solve two sets of decisions: (a) what service level should IBM provide to its customers? (b) What inventory strategy in their global network would achieve this service level with a minimum investment in inventory? The service level metric can be *fill rate* or *non-stock out probability* at an item level or at a product family level ("joint service level"), this is measured based on the *service time* that is quoted to the customer.

Moreover, unlike a "*one-off*" (or relatively infrequent) decision like designing a responsive supply chain, operating an existing supply chain that meets a certain service level target with low inventory investment (PPOM-2b) is an *"on-going" activity*, performed every day/ week/ month, forever. The decision on what service level (PPOM-2a) to provide can be an annual exercise reflecting business conditions and competitive positioning. We will return to this story a bit later as well.

A leisurely stroll in the thick forest of inventory models

It is time to introduce some (more) vocabulary. A crucial modeling choice is how time is represented: if it is important to represent events as they happen continuously, and the solutions of such models can be expected to be reasonably implemented in practice using real systems, then working with a "continuous time" model may make sense. Otherwise, an alternative is to chop up time into "buckets," or periods, and use a "discrete-time" (or a "periodic review") model. For solving many practical problems in OM outside the factory shop floor, discrete-time models have been useful. I have used continuous time models for problems inside a factory.

The earliest inventory models – mentioned in Fairfield E. Raymond, *Quantity and Economy in Manufacturing (1931)* -- are those of F. W. Harris (1915) (what we now call the economic order quantity EOQ model), with finite production rate analyzed by Benjamin Cooper (1926), and a model with randomness studied in 1928 by Thornton C. Fry. These researchers were inspired by F. W. Taylor, *Shop Management (1903)* and had executive interest from folks at Westinghouse (in Pittsburgh) and Henry Ford. (Hence you can understand why they were concentrating on single item models with the objective of cost minimization.) It appears that they were not aware of related models on managing money, such as Edgeworth (1888) (we now call this the News-vendor model). In 1949, Arrow and Marschak were aware of both these sets of papers, and they naturally did what was needed (see Arrow, 2002): join the dynamic nature of the Harris model (1915) (deterministic and in continuous time) with the stochastic (single period) model of Edgeworth (1888). At that time, stochastic dynamic programming was not fully developed, so they recruited a newly minted Math PhD student from Princeton named T.E. Harris (a student of William Feller, different from the EOQ Harris), and we get Arrow, Harris and Marschak

(1951). This paper was considered as the "beginning of what may be called modern analysis of inventory systems," see Naddor (1966). The seductive mathematical elegance of these stochastic, dynamic, discrete time inventory models attracted more newly minted math PhDs (such as Karlin and Scarf) from where (the not-EOQ) Harris came from. This was not without some undesired consequences. Hadley and Whitin, in 1963, write in *Analysis of Inventory Systems:*

> "At one extreme a considerable amount of work is concerned strictly with practical applications, while, at the other extreme, work is being done on the abstract mathematical properties of inventory models without regard to possible practical applications."

They were airing their frustration about the discrete-time, single item, stochastic models studied in Arrow, Karlin and Scarf (1958)! Also, see the plea from Conway, Maxwell and Miller, *Theory of Scheduling* (1967):

> "The theory of scheduling has not attracted the attention of scholars comparable to Arrow, Karlin and Scarf; perhaps because the mathematical model which underlies the theory of scheduling cannot compare in elegance with that which underlies inventory theory. However, there is at least as much practical incentive to solve the problems of sequence…"

Models may consider the building of a product ("assembly network") or the servicing of customers with goods already produced ("distribution network"). Models that consider interactions between inventory levels at multiple locations at once are called "multi-echelon" (or "multi-stage") if there is a precedence relationship in the flow of materials between the locations or simply "multi-location" if the locations are "parallel" to each other and can transship items between each other if desired. Those that look at each inventory location in isolation are called "single stage." The lead time to obtain the material after placing an order can be deterministic or stochastic, the number of sources from which supply can be obtained can be one ("single sourcing"), two ("dual sourcing") or more, the lead time (and cost) can depend on the mode of transportation used ("truck load" or "less than truck load;" rail or van or air or ship), multiple modes can be simultaneously available ("regular" shipment versus "emergency" or "expedited" shipment), a discount on cost may be available if the quantity purchased at any time (or over a horizon) from a supplier is sufficiently large ("all units discount" or "incremental discount above a threshold" with several "break points"). The inventory can belong either to the customer or the supplier ("on consignment" or "vendor managed inventory"). If discounts are available only for a short interval of time -- a promotion of sorts – or if the price is expected to change sometime in the future, then customers may *time and batch* their purchases accordingly. Items may be perishable and so have a "shelf life." If an item is unavailable, a customer may substitute, wait, or forgo the purchase entirely.

Models typically aim to minimize the expected total cost: the holding cost of inventory is usually modeled as a linear (or convex) function of the inventory level; when a demand is not satisfied within a prescribed time window, either a "backlog" cost is incurred or there is a cost for lost margin if the customer leaves. An alternative formulation replaces the cost of backlog or lost margins with a constraint on (expected or probabilistically guaranteed) service level. Both formulations are useful.

If a discrete time model considers a one-shot decision, it is called a "single period" model; otherwise, it is "multi-period." A multi-period model is called non-stationary if the parameters may change with time; otherwise, "stationary." Models have inputs, constraints and objective functions; a solution to a model provides outputs. As time evolves – we are in a dynamic operating environment -- new information becomes available, consequences of past decisions become known and new decisions have to be taken (constrained by past decisions).

The outputs of a discrete-time, stochastic dynamic inventory model are a form of *policy* to follow with appropriate *parameters*. The simplest policy to study (also the easiest to implement in practice) is a base-stock (or order up-to) policy with a base stock (or order up-to) level -- by period, by item -- based on the inventory position -- inventory on hand plus on order minus any backlog. These base stock parameters are typically computed probabilistically and are recommended as target inventory positions to maintain before the next demand occurrence. The actual production (or order quantity) takes place after the next demand occurs with the aim to restore the (now lowered) inventory position back to the target inventory position (or as close to it as possible when constrained by available upstream inventory, capacity constraints or batch-size restrictions).

A model can be "centralized" (meaning there is only one decision maker in charge of all decisions) or "de-centralized" (there are at least two, not completely aligned, decision makers). The latter class of models is used in the equilibrium analysis of static models (using tools from game theory).

But: *"All models are wrong. Some are useful,"* said George Box (although this quote has been frequently improperly attributed to Albert Einstein).

How to create useful models?

I will describe how to *make the outputs of centralized models implementable in a decentralized setting, thus making the models useful*. This summarizes my approach over the past 20 years: construct centralized production-inventory models in such a way that demonstrates benefits of co-operation to the various parties. Thus, they agree to make appropriate changes in their policies and so *implement the model outputs in the real world*. See Erhun and Tayur (2003), where the merchandizing group of a grocery retailer, incentivized to get the lowest unit cost, buys in large quantities (from Heinz, Kellogg's, ConAgra Foods and so on) to obtain discounts. But the warehousing group of the same retailer, whose performance is measured by inventory "turns" prefers lower inventories. And the logistics group that is compensated based on transportation costs prefers full truckload shipments. All three groups were brought together to achieve the lowest net landed cost (NLC) for their firm. For a good discussion on organizational silos in the context of inventories, see Killeen (1969). In Troyer et al (2005), you can see how the results of a centralized 2-stage inventory model of distribution was used to drive inventory targets across the five warehouse of Deere and 2500+ independent dealers. This reduced the inventories by $1+Billion while increasing the service levels from 63% to 92%. This resulted in an annual increase of $120 million in the shareholder value added (SVA). In Keene et al (2006), you see an example where Tier-1 suppliers of Caterpillar adopted the solutions of a centralized model for the "greater good" of the entire chain, and the Caterpillar factories adjusted their component inventories (upward) in discord with their "lean philosophy" mantra of inventory reduction, for the same reason.

In the above three companies – and in several others – the senior executives (including the CEO) were keen to see what was happening and the initiatives required the approval of the Board of Directors. This was even more so during the financial crisis in 2008 where the importance of working capital skyrocketed to a top CEO initiative, triggering DuPont and Estee Lauder (among others) to accelerate inventory oriented projects. During the implementation process, the CEOs/CFOs (of publicly traded companies) have mentioned the use of inventory control models and software in their quarterly earnings call to Wall Street analysts (an example is Q3 2010 call by Celestica). Privately held companies, such as Kohler, are usually quite discreet about their initiatives and outcomes; however, they have presented the value of these models in public forum, and I have been told that the Kohler family is very pleased. This is in contrast to the frustration expressed by Herb Simon in Simon and Newell (1958):

> "Operations Research has more to do with the factory manager and the production-scheduling clerk than it has to with the vice president and the Board of Directors."

Why did Herb fail to engage senior executives using Operations Research? I think what may have happened is that the executives had shifted from a "production oriented" world into the "consumer scarce" world. He was prescient in other ways:

> "We are now poised for a great advance that will bring the digital computer and the tools of mathematics and the behavioral sciences to bear on the very core of managerial activity – on the exercise of judgment and intuition; on the process of making complex decisions."

Herb was looking for "executive centrifuges" so that we could have a science of "judgment mechanics" to match quantum mechanics. Let us use the phrase "management mechanics," because it sounds similar to "management engineering" (of the 1920s) and "management consulting" (coined in the 1930s and still in use).

Management Mechanics: The various functions of a model

Here is some advice that has helped me over the years:

(1) Even if the data exist, they are typically not easily accessible. One reason is that data can be in several different locations, in different information systems, or under the control of different organizational groups. It is important to have the support of a senior executive who can "make things happen" starting from, the "zero-th" task, namely, obtaining available data in a timely fashion.
(2) Analyzing a (mathematical) model and creating outputs is not the same as solving a (practical) problem in operations management. To be clear: if the recommendation of a model is not implemented then no value has been created (unless inaction was in fact the recommendation). To quote an ancient philosopher: "Analysis without action is no different from day dreaming." (He also cautioned: "Action without analysis can create nightmares.") Recently, and bluntly, Lou Gerstner, a former CEO of IBM writes in *Who Says Elephants Can't Dance? Inside IBM's Historic Turnaround*: "Ideas that are not implemented are no different from hallucinations."
(3) Work only on OM problems where senior executives are willing to spend a lot of money and are willing to engage by allocating appropriate time on their schedule. This is most likely the case when they have a "burning platform" and a non-trivial part of their executive compensation depends on the outcome of the implementation.

How to make sure that our modeling and analysis work is not futile, that we are not day dreaming or hallucinating?

I believe that before you model something and analyze the model, you should know how the results of the analysis are going to be used, who are involved in making decisions, how to overcome any impediments to their implementation and how to measure the value created due to this implementation. That is, avoid "pre-mature modeling." Sherlock Holmes in *Valley of Fear*:

> "The temptation to form pre-mature theories upon insufficient data is the bane of our profession...I should like a few more facts before I get so far as a theory."

Next, how does one put the recommendations of a model – a proposal of action -- into actual practice? In 1954, two Carnegie Mellon University professors, Charles Holt and Herbert Simon described a framework that is still being used today in their paper *Optimal Decision Rules for Production and Inventory Control*. Planning organizations in manufacturing firms follow a three-step approach: (1) forecasting; then (2) planning; and finally (3) reacting. Today, this is known as Sales Inventory & Operations Planning (SIOP), and pretty much every planning group of every manufacturing company in the world uses this process for ongoing operations. So, if one wants to use the outputs of a mathematical model in practice, in a real PPOM-2b situation, here is the process in which it can get imbedded.

Let us look at a realistic version of PPOM-2b that many senior executives at Global 2000/Fortune 500 manufacturing companies face. There are multiple stages of production and assembly to make the products. For example, at GlaxoSmithKline, several supply chains (such as the one that produces Paxil) have 11 stages of production and assembly; Eastman Chemical has 9 or more stages, and Kohler has 6 stages. There is a capacity constraint on how much one can produce in a given amount of time, all items are not produced in every period, and when produced, are done so in specific lot sizes. There are multiple stages in the distribution network that the (near final) product flows through from the final factory to the end customer. Several end products can share common components or sub-assemblies, likely in different quantities. As engineering changes occur, the bill-of-materials (BOM) can change (and BOMs therefore have "effective" dates). It is common to have at least two distribution stages, and many times there are three echelons. Along the distribution chain, a nearly finished item (due to a "postponement" strategy in place for effectively managing product variety) is configured before final sale; thus even in the distribution network there is a bill-of-material involved.

How to operate such a supply chain? What is the policy that the planning group should use? What are the parameters of the policy? How much inventory is needed across the chain? *Why?* (Magee, 1956, is an example of a well written executive communication on the role of inventories for single-stage models.) How does the investment change as the type of service, the service level, forecast errors, batch sizes, capacity constraints, lead times, schedule adherence, supply reliability and review frequency change?

The "answer" to a practical problem such as the above, in my experience, requires the following: (1) the ability to model several complexities at once; (2) the ability to perform several "what-if" analyses from different points of view corresponding to the various parties involved in decision making; and (3) the proposal of solutions that are understandable (by the executive as well as the planner/user), trust worthy (as their professional careers and compensation depends on it), sufficiently robust (as the actual real world setting cannot be fully captured by the models anyway) and implementable within (or with small modifications) of existing processes, organizational structures and information technology infrastructure. Additionally, depending on the risk averseness of the executive or inherent in the organizational culture, a "roll out" plan, in waves, is also required so that the organization can "walk" before they choose to "run." (Sometimes folks want to "crawl" even before they choose to "walk".)

Thus a modeling framework and solution proposal should allow for partial changes in the decisions in a sub-network holding the rest somewhat constant, and then, increase the range and scope of decisions being changed. What is needed is a comprehensive model that allows for what I call "staged optimization" deliberately restricting some variables to be within a certain range for the time being. That is, a "controlled release" in concert with the organization's capacity to absorb change, in rhythm with their existing processes and compatible with their IT systems.

What is described here is an essential aspect of "management mechanics," facilitating the implementation of (centralized) model outputs in the (decentralized, multi-agent) real world.

Now let us move to the analysis of models.

"Optimal" policies versus "implementable" policies

A policy is considered "optimal" for a chosen model if no other policy dominates it according to a specified objective. The point is that "optimality" is a feature of a solution to a mathematical model and may not have anything to do with the executive's or the firm's problem. It is believed that if one chooses a "realistic" model and finds good solutions to that model, these translate into useful solutions to the practical problem. Thus, we must first agree on a mathematical model to represent reality that is considered "satisfactory" by the eventual decision makers and users. Next, we have to find the optimal (or simply a good) policy for the model, crank through the data, and see if the outputs that are created are worth implementing. Outputs that are worth implementing are those that executives believe will create significant improvement in the key operating metrics of their firm (compared to what they are doing now) at this time when used by their people (or those that can be reasonably hired) using computers that they have (or can purchase at reasonable cost).

Einstein, *"Models should be as simple as possible, but no simpler."* Also as Joseph E. Stiglitz has warned:

> "Models help guide our thinking, but we should never let the analysis of simple models replace our thinking, or let us lose touch with reality."

Although a stylized mathematical model has an elegant solution, it is of no value to the decision maker if it is unrealistic. At the same time, "too realistic" models may be very difficult to solve. Thus what is needed is a "sufficiently accurate" representation of the real world, balancing decision relevant criteria with tractability, and providing understandable results that can be implemented.

"Optimality" of a policy in a mathematical inventory model is typically proved using dynamic programming. For several single-stage, discrete time production-inventory models, with complete backlogging of unsatisfied demand, a time varying "base stock" policy (perhaps modified due to capacity constraints or lot-size constraints) can be shown to optimal. If a fixed cost (for any production or procurement event) exists in addition to the usual variable per-item costs, but there is no capacity constraint, a (min,max) policy (with time varying parameters as needed) is optimal. That is good news as most planning systems in practice that are part of the information technology (IT) infrastructure can handle time varying versions of these two policies, at an (item, location) level of granularity. This means that the outputs of models can be inputs into IT systems that drive execution.

The bad news is any additional complexity (such as lost sales or multiple distribution stages or a capacitated assembly network) makes the "optimal" policy too complex to be implemented in any realistic setting. Thus most practical implementations require that the policy structure to be either base-stock or min-max. In some respect, you can think of adding a cost term to the typical academic models that ignore implementation: significant costs are incurred if the policy structure deviates from what the mental model is, or if changes are needed in the IT systems. This "total cost" optimization will get you close to the two implementable policies, as it is highly unlikely, in practice, that the "optimal policy without these costs" will perform so much better in inventory savings than the best policy within this class. (In case it is, an executive may be persuaded to consider a major upheaval of the IT systems.) Thus academic inventory models are not *complete* (to use the term from Little, 1970) if they do not account for these implementation costs. Thus, one can sympathize with Hadley and Whitin's complaint -- of (some papers in) Arrow, Karlin and Scarf (1958) -- that they assume unrealistic and simplistic model settings and/or propose complex un-implementable policies, *exactly the opposite of what is needed*. (These models can have other uses.)

So, where is such a comprehensive modeling framework that we require (as discussed earlier) and a computational procedure that allows for staged optimization accommodating the needs of various decision makers? In 1991, I could not find one. That was bad news for the IBM executive looking for answers. There was, however, some consolation for me, as his parting comments were something like:

> "If your model were more realistic in capturing our capacity constraints, changing patterns of demand and supply, and could account for the fact that we have common components that are shared by many products whose demands are correlated, and that these are produced and distributed using a physical network with shared resources and multiple production choices, it would have been so useful. We could have done so many interesting what-if's and decided how IBM can compete better. Call me when you get there."

This was great news for a new assistant professor about to join a university. One could not have asked for a better research plan!

My garden of inventory models and solution methods

Paul Samuelson:

> "Each new theorem, each new insight, is like money in the bank, waiting to be drawn in some unexpected connection."

In 1991, there was no known method to compute the optimal base stock level even for the simplest setting: a capacitated, stationary, discrete time, single stage, single item stochastic model being operated by a (modified) base stock policy. But I made an unexpected connection: a capacitated, production-inventory system, operated under a base stock policy is equivalent to a model of a dam, previously studied by P. A. P. Moran in 1954. In fact, the stochastic process – called "shortfall" to distinguish it from "backlog" – of the amount that one is below the base stock level due to the capacity constraint follows Lindley's equation identical to a D/G/1 queue. (See N. U. Prabhu's 1965 *Queues and Inventories.*) This allows for a recursive method to efficiently compute the optimal base stock policy; see Tayur (1993a).

And you get more! Beyond the dam model connection, I realized that complex capacitated, multi-stage stochastic systems could be studied via sample paths, not necessarily through distributions. Glasserman and Tayur (1995) provides optimal base stock levels for a capacitated multi-echelon production-inventory network (assembly, distribution, combined or any network that is an acyclic directed graph). Here we use *Infinitesimal Perturbation Analysis* (IPA), a simulation-based optimization method that is valid under a wide range of complexities that one encounters in practice. This works because Lindley type recursions can be written and analyzed in these situations as well (Glasserman and Tayur, 1994, that uses a method from Loynes, 1962). Simulation is used not only to compute the expected cost at any base stock level, but the same sample path can be used to compute a valid estimate of all the gradients (of expected cost with respect to each base stock level, one for each stage regardless of the number of stages in the supply chain). Thus, this is a very effective computational procedure: it can handle many complexities of the model and the computational time does not grow very much with scale. Although the expected cost function is not convex, it is uni-modal; a gradient search procedure converges to the minimum cost quite quickly. IPA can be used to compute base stock levels for both service level constrained models and expected total cost models, in stationary and time varying conditions, in single period, finite horizon and infinite horizon settings. By varying service levels, one can compute the inventory investment needed at each service level, and construct the inventory-service level curve for PPOM-2.

Time varying characteristics are analyzed in Kapuscinski and Tayur (1998) in the single stage setting. The proof of optimal policy – it is time varying modified base stock policy -- is through stochastic dynamic programming. The optimal base stock parameters are computed by the use of IPA. Multiple products sharing common capacity, re-entrant flow structure through multiple stages and loops, with several different capacity allocation rules are discussed in Bispo and Tayur (2001).

Thus, there are three different things going on in the analysis of models: (1) searching for an optimal policy of the model; (2) devising a computational procedure to obtain parameters for the policy chosen, which may or may not be optimal, for that model; and (3) showing that the computational procedure is valid. We can use different mathematical methods for each of the three things: Dynamic programming for (1), sample path derivatives via IPA for (2) and methods from real analysis (such as the dominated convergence theorem, Lipschitz continuity, Harris ergodicity, positive recurrence and so on) for (3).

Although IPA was originally developed for PPOM-2, it is also used to compute inventory parameters for the responsive supply chain (PPOM-1), vanilla boxes for postponement (let us call it PPOM-3) and stochastic cyclic schedules (call it PPOM-4), thus proving it to be a very versatile technique to solve complex models at industrial scale in a variety of situations. PPOM-3 is *formulated* as a two-stage stochastic program with recourse in Swaminathan and Tayur (1998), a framework developed by Dantzig (1955). PPOM-4 is modeled using chance-constrained programs in Anupindi and Tayur (1998), a framework developed by (Carnegie Mellon professors) Charnes and Cooper (1959).

I do want to emphasize one point here. Hilbert:

> "A perfect formulation of a problem is already half its solution."

For an entirely different method to solve chance-constrained programs where derivatives are not used, see Kannan, Mount and Tayur (1995), where a carefully selected random walk procedure leads us close to the optimal solution (in polynomial time) with high probability. The proof involves showing that the Markov Chain created by this walk *mixes* rapidly. A lower bound on the conductance of the Markov Chain provides the complexity result. If integer variables are involved, we can use the lattice walk procedure generated by "test sets" of an integer program, computed as Grobner Basis of certain polynomial ideals, an idea from Algebraic Geometry; see Tayur, Thomas and Natraj (1995) and Scarf (1997). An entirely different method of solving integer programs via Grobner Basis is in Bertsimas, Perakis and Tayur (1999), which looks at a different construct of a polynomial ideal, and obtains a weak duality result by an application of Hilbert's *Nullstellensatz*. This can be viewed as *Farkas lemma* for Integer Programs.

I will readily admit that these procedures that utilize random walks and Grobner Basis have not been of much use thus far in solving PPOMs. However, they have satisfied my "inner mathematician" needs for pure intellectual elegance uncorrupted by any desire to be of immediate practical use. I am reminded of a comment made by Hermann Weyl:

> "My work always tried to unite the true with the beautiful; but when I had to chose one or the other, I usually chose the beautiful."

These then are some of the contributions (beyond IPA derivatives for inventory models that have been immediately useful) to the "bank," hoping that someone, somewhere, sometime in the future, for some purpose, will find some unexpected connection and take off in an entirely different direction. Indeed, as G.H. Hardy concludes in *A Mathematician's Apology:*

"…that I have added something to knowledge, and helped others to add more.."

Enterprise Inventory Optimization

Let us return to the real world of today. How to implement multi-stage inventory models not just in a handful of companies, but in virtually every company where there is sufficient benefit in doing so?

Quoting again from Reisman (2004):

"The entrepreneur is a person who implements innovations to make his undertakings a success. Entrepreneurship is the propensity to pioneer new innovations…."

And, quoting Schumpeter directly on the entrepreneurial type (Thomas K. McCraw, *Prophet of Innovation*):

"…the joy of creating, of getting things done, or simply of exercising one's energy and ingenuity; our type seeks out difficulties, changes in order to change, delights in ventures. The entrepreneur is a driven man because that is what he is. Neither consideration of the effort nor satiation of his hedonic needs tames his lust for action."

In 2000, I founded SmartOps Corporation (and served as its CEO for 12 years) to bring new intellectual property to the market. We created a new segment called – we coined the phrase -- Enterprise Inventory Optimization (EIO). We designed our EIO software to work in concert with Enterprise Resource Planning (ERP) and Advanced Planning and Optimization (APS) systems. It was important that we had one product, not a custom product for each customer, both from a profit margin perspective and also for scaling the company and having a replicative machinery of pre-sales consulting, sales, implementation, and post-sale support for on-going use by the customer. Thus a key proprietary innovation was the creation of a discrete-time, stochastic, multi-stage inventory model that could handle the possible and conceivable complexity of any situation (within reason). A fast algorithm computes good outputs quickly. This model also allows for staged optimization and several what-if analyses and scenario planning. To make it convenient for executives, we created iPad compatible analytics and "apps." To increase IT convenience and reduce operating costs, we have architected EIO to work on a "cloud," and to provide a rapid response to what-if questions, we have architected it to run on an "in-memory database."

In the beginning, it seemed that prudent businessmen followed Alexander Pope's precept, also encountered by Herbert Simon as he discussed in his 1991 autobiography, *Models of my Life*:

"Be not the first by whom the new are tried; nor yet the last to lay the old aside."

This resistance to new technology is not specific to production-inventory planning. Indeed, as W. Brian Arthur notes in *The Nature of Technology:*

> "Still another reason is psychological. The old principle lives on because practitioners are not comfortable with the vision – and promise – of the new. Origination is not just a new way of doing things, but a new way of *seeing* things. And the new threatens. It threatens to make the old expertise obsolete."

Fortunately, there are always "innovators" and "early adopters" in Industry (these terms were used in Rogers, 1962, and formed the basis for Bass, 1969), we did benefit from those executives who were looking for something new. Recall from before that "Capitalism is newness."

But it is important to "cross the chasm" (Geoffrey Moore) between the "early adopters" and the "leading majority" to really become mainstream. There was another key innovation that was needed. A major hurdle that a "technically correct solution that can create huge practical value" needs to overcome to be of mainstream use by non-technical people is the *abstraction step*. (This issue is also mentioned in Bixby, 2002, with respect to the adoption of mixed integer linear programming, MILP.) Recall that we first translate the real world into a mathematical model ("the abstraction step"), then solve the mathematical model, and then output numbers. It is not possible to eliminate the abstraction step as it *creates* the "model!" But then, if you don't, you need experts around and this limits the companies in which the models will be implemented. So what we did was to make the abstraction step *invisible* to the typical users (but make it available to a few "super-users" in any firm to review and modify if needed) by creating an automated, transformation module that reads transactional data (in the form that the user knows and is comfortable with) from the ERP system, creates the mathematical model automatically (*no human being involved*!), solves the model, and outputs the answers in a form (absorbable and viewable in the APS system) that the user already is comfortable navigating.

This worked like magic in line with Arthur C. Clarke's observation:

> "Every sufficiently advanced technology is indistinguishable from magic."

But how can we make this "invisible secret sauce" available to hundreds of companies? What if the folks who sell ERP and APS systems, like SAP AG, also were allowed to sell EIO from SmartOps? In 2009, we signed a worldwide reseller agreement with SAP allowing their sales force (several hundred of them) to sell our EIO software along with theirs. A case study (for MBA students and for executives) on SmartOps Corporation discusses our "go-to-market" and channel strategy (with SAP AG); see Wilcox (2011).

Over the past twelve years, I have had the good fortune of working with nearly 100 companies, across the world, in several industry verticals who have now incorporated PPOM2-b into their on-going operations by imbedding our enterprise software (with proprietary algorithms). The software periodically (and automatically) calculates operational inventory targets, for each item, at each location, for each period, in their complex multi-stage global network spanning several continents.

Designing and implementing a responsive supply chain at Caterpillar

Let us return to PPOM-1 faced by Caterpillar that I mentioned earlier. The key inventory model to "build in" responsiveness had the following features: (1) discrete time, (2) stochastic demand, (3) two customer classes, (4) partial backlogging, (5) dual modes of supply with different lead times, (6) supply uncertainty, and (7) operated by a two-parameter base stock policy, with a level (say 10) for emergency shipments (using faster mode of transportation) and a level (say 35) for regular shipments. Thus, at the end of each period, after satisfying (partial or complete) demand, accounting for customer backlogging or abandonment, we look at the inventory position. If it is lower than 10 units, we order enough using the emergency mode to reach 10 and order (35-10=25) through the regular mode. If the inventory position was above 10, we only order using the regular mode to restore the inventory position to 35. This is not an optimal policy to the model; but it is implementable, and using IPA, we computed the pair of levels for all items and locations in the network, across several different network choices, and across a variety of parameter settings for robustness. The details are in Rao, Scheller-Wolf and Tayur (2000). Our proposals were implemented. A business article in *Fortune* (Seikman, 2000) *New Victories in the Supply Chain Revolution*:

> "Among the techniques the Carnegie-Mellon group used to attack this complex problem was so-called infinitesimal perturbation analysis, for which no complete explanation is possible for the faint-hearted or mathematically disadvantaged."

Continuing, the article quotes a senior executive from Caterpillar:

> "…the Carnegie-Mellon solutions are not what Cat would have come up with on its own. A couple of special tool-distribution centers, which the company had planned to build, were found unnecessary. Just as important, the response time in the system was sufficiently fast that the inventories that the dealers would have to carry were not high enough to require a subsidy from Caterpillar. …[Carnegie Mellon] gave us the highest response, lowest cost, lowest inventory [solution]..."

This *Fortune* article attracted the attention of several executives at Deere in 2001. I received a call on August 6th, 2001: "*Can you help us reduce $1 billion of inventory, over the next 5 years, without sacrificing service levels*?" A SmartOps team did a "proof of value" consulting project – "How low can you go?" -- between October and

December 2001. We showed that through EIO and operational discipline in PPOM2-b they can "fix their mix" of inventories, reduce the total investment while not reducing service levels. We met with CEO Bob Lane (and his senior staff of CFO and the Presidents of each division) in January 2002. We did a controlled roll out – a pilot limited to certain product families and key dealers from their dealer council -- in the spring of 2002 before launching the full implementation in June 2002. Four years into it, more than $1.1 billion of inventory was either eliminated or avoided, $100 million more than planned and one full year ahead of schedule while increasing service levels (Troyer et al 2005).

Dantzig (1963):

> "The final test of any theory is its capacity to solve the problems that originated it."

Concluding remarks

I have highlighted some academic papers and books – Holt and Simon (1954), P.A.P. Moran (1954), Dantzig (1955), Drucker (1955), Magee (1956), Charnes and Cooper (1959), Loynes (1962), Prabhu (1965), R. G Brown (1967) and Killeen (1969) – as an attempt to rescue them from undeserved neglect by our inventory research community today. Also, if the industrial examples have concentrated on situations described in my papers, it is not because I consider them the only (or even the most important) set of problems and models, but rather because, on the one hand, they are more within my own special competence, and on the other, it seems to me, based on my 20 years of interactions with nearly 100 firms in several continents and industry verticals (Table 2 provides some examples), that they likely constitute situations that many thousands of companies face.

Table 2 Industry Verticals and Representative Companies

Industry Vertical	Representative Companies
Chemicals	Bayer, PPG, Dow, Dupont, Eastman, Cabot, Monsanto, Lubrizol, Rohm & Haas
Consumer packaged goods (CPG)	Kellogg's, ConAgra, Unilever, Smuckers, Clorox, J&J, Scherring-Plough, Kohler, Estee Lauder, Campbell's
Hi-tech	Cisco, Celestica, Jabil, HP, LSI, Micron
Industrial machinery/components	Caterpillar, Deere, Danfoss, Honeywell
Retail/Distribution	Cardinal Health, Shaw's Supermarkets
Pharmaceuticals, Medical devices	GSK, Wyeth, Pfizer, Medtronic

I have cherry picked the quotations to advance my narrative. Similarly, I have intentionally constructed contrasts: (1) capitalism versus socialism; (2) stochastic, non-linear models versus linear, deterministic ones; (3) decision support in a dynamic environment versus static equilibrium models; (4) customer versus production scarcity; (5) general purpose computational methods rather than specifically constructed formulas; (6) real world implementation across silos and firm boundaries versus analysis of simplistic models; (7) staged optimization with exploration of several alternatives versus providing one solution; (8) enterprise software that runs on a cloud and renders analytics on an iPad for automated global planning versus manual or spreadsheet calculations done at a local level; (9) executive engagement and Board of Directors approval versus production-clerk support; and (10) a lust for action versus pure intellectual pursuits without concern for practical implications.

I want to emphasize that "management mechanics" through staged optimization is not dependent on a particular model (nor restricted to the use of IPA). It is not limited to global multi-enterprise inventory planning nor is to "planned spontaneity." This approach has facilitated the implementation of centralized model outputs in the multi-silo, multi-firm settings, *outside the "direct control" of the executive.* It is simply a way of developing models and "getting the outputs implemented" in the real world. Benjamin Franklin:

> "Wisdom is knowing what to do. Virtue is actually doing it."

Acknowledgements. There are so many people who have helped make inventory models useful to practice -- ranging from my PhD students, faculty colleagues, co-authors and collaborators, SmartOps' employees and channel partners, executives, IT staff and planners from Fortune 500/Global 2000 firms -- that to list them individually by name will exceed any reasonable page limit. I would like to thank Alan Scheller-Wolf, Mustafa Akan, Tinglong Dai, Steve Graves, Jack Muckstadt, Soo-Haeng Cho, Mike Trick, Nicola Secomandi, Jay Swaminathan, Paul Glasserman and David Simchi-Levi for their help and suggestions in the preparation of this manuscript.

End Notes

1. It is said that imitation is the sincerest form of flattery. The style of the article – especially the extensive use of quotes – pays homage to the thoughtful articles of S. Chandrasekhar collected in *Truth and Beauty: Aesthetics and Motivations in Science (1987).* The use of some language in the article is borrowed from the preface of Tjalling C. Koopmans (1957), *Three essays on the state of Economic Science,* while some language in the concluding section is from Milton Freidman, from his conclusion section of a 1955 report found in *Two Lucky People*, by Milton and Rose Friedman. I also benefited from reading the personal narratives in the collection of articles in 50th Anniversary Issue of *Operations Research* (Jan-Feb, 2002) and of *Management Science* (December 2004).

2. The desire for academics to "help senior management solve their most pressing problems" is not new. James O' McKinsey, a professor at U. Chicago GSB (now called Booth) founded McKinsey & Company in 1926. He wanted to

do more than what “efficiency experts” were doing under the banner of “management engineering.” Marvin Bower later coined the term “management consulting” which has remained in use since the 1930s. You should now understand why I prefer “management mechanics.”

3. The earliest reference I could find that seems to have put F. W. Taylor on the path towards *Scientific Management*, the precursor to “management engineering” – and looking at management from a scientific perspective, leading to the creation of the fields of Industrial Engineering, Industrial Administration and Operations Management -- is a 1886 paper by Henry R. Towne “The Engineer as an Economist.” Towne was the President of the Yale and Towne Manufacturing Company and a President of American Society of Mechanical Engineers (ASME). The young Taylor (who became a ASME member in 1885), impatient with what he called “just management of initiative and incentive”, wanted to think about how to “increase productivity and lighten labor’s efforts.”

4. The phrases “production scarce” and “customer scarce” are borrowed from Yuji Ijiri’s chapter in *The Innovative University (2004)*.

5. My PhD advisor’s PhD advisor – my “academic grandfather” so to say – is Arthur F. Veinott Jr, whose academic grandfather is none other than T. E. Harris! I discovered this through the *Mathematical Genealogy Project*.

6. One of my “minors” in graduate study was queuing theory and N. U. Prabhu was on my thesis committee. The other “minor” was Operations Management, and Dick Conway was also a committee member. My PhD advisor was Robin Roundy. My PhD thesis was on stochastic models of serial production lines operated by kanban systems; see Tayur (1992,1993b) and Muckstadt and Tayur (1995ab). The stochastic cyclic schedules in Tayur (2000) at GE are operated using kanban cards. My first introduction to inventory models was in a class taught by Jack Muckstadt at Cornell using *Hadley and Whitin* as a reference text. Bill Maxwell taught a course on scheduling. My first OR course was taught by T.T. Narendran at Indian Institute of Technology, Madras. Moving in the other direction, Nihat Altintas, Carlos Bispo, Feryal Erhun, Srinagesh Gavirneni, Roman Kapuscinski, Pinar Keskinocak and Jay Swaminathan are some of my PhD students.

7. I met Paul Glasserman in 1988 at Bell Labs in Holmdel, NJ. Beyond IPA, Paul and I have developed a large deviation approximation (based on extreme value theory) for capacitated multi-echelon inventory models (Glasserman and Tayur, 1996, Glasserman, 1998). A single stage formula connects inventory levels, service levels, excess capacity and variability in closed form. This “Glasserman-Tayur” formula (for discrete time, capacitated inventory model) can be considered a “cousin” of *Pollaczek-Khinchin* (M/G/1 queues).

8. Rekha Thomas was a PhD student at Cornell and her PhD thesis is on Grobner Basis and Integer Programs. Dimitris Bertsimas and Georgia Perakis were hosts during my 1997 sabbatical at MIT.

9. Holt and Simon worked closely with PPG. Magee (and Arthur D. Little) worked closely with Johnson & Johnson. Both PPG and Johnson & Johnson have integrated SmartOps EIO software into their IT stack, connected it to Oracle and SAP ERP and APS, and have made multi-stage inventory planning part of their monthly global SIOP process.

10. I have limited my discussion to four or so PPOMs in this article. Just as grocery retailers are optimizing NLC, their suppliers (like Heinz) are optimizing quantity discount schedules (see Altintas, Erhun and Tayur, 2008) and, like ConAgra Foods, are making their own production planning more flexible while reducing total costs (see Mehrotra et al, 2011).

11. I have not discussed the role of information (Gavirneni, Kapuscinski and Tayur, 1999), limited history of data (Akcay, Biller and Tayur, 2011), operations reversal as a strategy to manage variety (Kapuscinski and Tayur, 1999), guaranteed lead times (Kapuscinski and Tayur, 2007, Keskinocak, Ravi and Tayur, 2001) international operations affected by exchange rates (Scheller-Wolf and Tayur, 2009), real (capacity) options (Erhun, Keskinocak and Tayur, 2008a), supply chain co-ordination (Erhun, Keskinocak and Tayur, 2008b) or the role of the internet and e-business (Keskinocak and Tayur, 2003, Swaminathan and Tayur, 2004).

12. I have also not discussed the choice of the product portfolio itself. See Yunes et al (2007) for what Deere has. Deere has also tailored its distribution logistics by season to further improve operations performance with respect to responsiveness and cost; see Tardif et al (2010).

13. In 2005, we tested a massively parallel version of our algorithm on IBM Blue Gene (using their on-demand Deep Computing offering). An industrial instance with about a million inventory targets – one for each item, location and week -- solved in 0.04 seconds on a “half-rack” system of 512 parallel processors.

14. I have discussed monetizing the value of operations research in Camm and Tayur (2010). I presented EIO (and entrepreneurship) at UCLA’s Marschak Colloquium, along with Private Equity and Lean Operations (January 2012).

15. Entrepreneurship and significant contributions to practice (in diverse areas) by OR/OM professors have a long history: Harry Markowitz, Egon Balas, Art Geoffrion, Jack Muckstadt, Don Ratliff, Marshall Fisher, Bob Bixby, Morris Cohen, Steve Graves, Sunder Kekre, Hau Lee, David Simchi-Levi, Larry Wein and Dimitris Bertsimas, to name a few.

References

AKCAY, A., BILLER, B. and TAYUR, S., 2011, Improved inventory targets in the presence of limited historical demand data. *Manufacturing & Service Operations Management,* **13**(3), 297–309.

ALTINTAS, N., ERHUN, F. and TAYUR, S. 2008. Quantity discounts under demand uncertainty. *Management Science*, **54**(4), 777–792.

ANUPINDI, R., TAYUR, S., 1998, Managing stochastic multiproduct systems: model, measures and analysis, *Operations Research,* **46**, S98-111 (Supplement 3).

ARROW, K. J., HARRIS, T. E. and MARSCHAK, J., 1951. Optimal inventory policy. *Econometrica*, **19** (3), 250–72.

ARROW, K. J., 2002, The genesis of "optimal inventory policy." *Operations Research,* **50**(1), 1–2.

ARROW, K., KARLIN, J. S. and SCARF. H. E., 1958, *Studies in the Mathematical Theory of Inventory and Production* (Stanford, California: Stanford University Press).

ARTHUR, W. B., 2007, *The Nature of Technology: What it is and How it Evolves* (New York: The Free Press).

BASS, F. M., 1969, A new product growth model for consumer durables. *Management Science*, **15** (5), 215–27.

BERTSIMAS, D., PERAKIS, G., TAYUR, S., 2000, A new algebraic geometry algorithm for integer programming. *Management Science*, **46** (7), 999–1008.

BISPO, C., TAYUR, S., 2001, Managing re-entrant flow lines: theoretical framework and experimental results. *IIE Transactions*, **33**(8), 609–623.

BIXBY, R. E., 2002, Solving real-world linear programs: A decade and more of progress. *Operations Research*, **50** (1), 3–15.

BOX, G. E. P., DRAPER, N. R., 1987, *Empirical Model-Building and Response Surfaces* (New York: Wiley).

BROWN, R. G., 1967, *Decision Rules for Inventory Management* (New York: Holt, Rinehart, and Winston).

CAMM, J., TAYUR, S, 2010, Editorial: how to monetize the value of OR, *Interfaces*, **40** (6), 446–450.

CHANDRASEKHAR, S., 1987, *Truth and Beauty: Aesthetics and Motivations in Scienc*e (Chicago: Chicago University Press).

CHARNES, A., COOPER, W. W.,1959, Chance-constrained programming. *Management Science*, **6**(1), 73–79.

CONWAY, R. W., MAXWELL, W. L. and MILLER, L. W., 1967, *Theory of Scheduling* (Reading, Mass.: Addison-Wesley).

COOPER, B., 1926, How to determine economical manufacturing quantities. *Industrial Management*, **72** (4), 228–233.

DANTZIG, G. B., 1955, Linear programming under uncertainty. *Management Science,* **1** (3-4), 157-206.

DANTZIG, G. B., 1963, *Linear Programming and Extensions* (Princeton: Princeton University Press).

DRUCKER, P. F., 1955, "Management science" and the manager. *Management Science,* **1**(2), 115–126.

EDGEWORTH, F., 1888, The mathematical theory of banking. *Journal of Royal Statistical Society*, **51** (1), 113–127.

ERHUN, F., TAYUR, S., 2003. Enterprise-wide optimization of total landed cost at a grocery retailer. *Operations Research*, **51**(3), 343–353.

ERHUN, F., KESKINOCAK, P. and TAYUR, S., 2008a. Dynamic procurement in a capacitated supply chain facing uncertain demand. *IIE Transactions*, **40**(8), 733–748.

ERHUN, F., KESKINOCAK, P. and TAYUR, S., 2008b. Dynamic procurement, Quantity discounts and supply chain efficiency. *POMS*, **17**(4), 1-8.

FRIEDMAN, M., FRIEDMAN, R. D., 1998, *Two Lucky People: Memoirs* (Chicago: University of Chicago Press).

FRY, I. C., 1928, *Probability and its Engineering Uses* (New York: D. Van Nostrand Co., Inc.).

GAVIRNENI, S., KAPUSCINSKI, R. and TAYUR, S., 1999, Value of information in capacitated supply chains. *Management Science*, **45**(1), 16–24.

GERSTNER, L. V., 2002, *Who Says Elephants Can't Dance? Inside IBM's Historic Turnaround* (New York: Collins).

GLASSERMAN, P., 1998, Service levels and tail probabilities in multistage capacitated production-inventory systems. In *Quantitative Models for Supply Chain Management*, TAYUR, S., GANESHAN, R. and MAGAZINE, M.J., (Eds.). Springer.

GLASSERMAN, P, TAYUR, S., 1994, The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy. *Operations Research*, **42**(5), 913–925.

GLASSERMAN, P., TAYUR, S., 1995, Sensitivity analysis for base stock levels in multi-echelon production-inventory system. *Management Science*, **41**(2), 263–281.

GLASSERMAN, P., TAYUR, S., 1996, A simple approximation for a multistage capacitated production-inventory system. *Naval Research Logistics*, **43**(1), 41–58.

HADLEY, G., WHITIN, T.M. 1963. *Analysis of Inventory Systems* (Englewood Cliffs, N.J.: Prentice-Hall).

HARDY, G. H., 1940, *A Mathematician's Apology* (New York: Cambridge University Press).

HARRIS, F. W., 1915, *Operations Cost, Factory Management Series* (Chicago: Shaw).

HOLT, C. C., SIMON, H. A., 1954. Optimal decision rules for production and inventory control. In *Proceedings of the Conference on Operations Research in Production and Inventory Control* (Case Institute of Technology), 73–89.

IJIRI, Y., 2004, Science, technology, and the innovative business school, In *The Innovative University*, Resnick, D., Scott, D. (Eds). Pittsburgh, PA: Carnegie Mellon University Press.

KANNAN, R., MOUNT, J. and TAYUR, S. 1995. A randomized algorithm for optimizing over certain convex sets. *Mathematics of Operations Research*, **20**(3), 529–549.

KANTOROVICH, L. V., 1939, *Mathematical Methods of Organizing and Planning Production* (Leningrad State University, in Russian).

KANTOROVICH, L. V., 1960, Mathematical methods of organizing and planning production (Translated by MARCHENKO, A. R.). *Management Science*, **6**(4), 366–422.

KAPUSCINSKI, R., TAYUR, S., 1998, A capacitated production-inventory model with periodic demand. *Operations Research* **46**(6), 899–911.

KAPUSCINSKI, R., TAYUR, S., 1999, Variance vs. standard deviation: note on variance reduction through operations reversal in supply chain reengineering. *Management Science,* **45**(5), 765–767.

KAPUSCINSKI, R., TAYUR, S., 2007, Due date setting in capaitated MTO setting with two customer classes. *Operations Research,* **55**(1), 56-74.

KEENE, S., ALBERTI, D., HENBY, G., BROHINSKY, A. J. and TAYUR, S., 2006, Caterpillar's building construction products division improves and stabilizes product availability. *Interfaces*, **36**(4), 283–295.

KESKINOCAK, P., RAVI, R. and TAYUR, S., 2001, Scheduling and Reliable Lead Time Quotation for Orders with Availability Intervals and Lead Time Sensitive Revenues. *Management Science*, **47**(2), 264-279.

KESKINOCAK, P., TAYUR, S., 2003, Quantitative analysis of internet enabled supply chains. *Interfaces*, **31**(2), 70-89.

KILLEEN, L. M., 1969, *Techniques of Inventory Management* (N.Y., American Management Association).

KOOPMANS, T. C., 1975, Concepts of optimality and their uses. Nobel Memorial Lecture, December 11.

KOOPMANS, T. C., 1957, *Three Essays on the State of Economic Science* (New York: McGraw-Hill Book Co.).

LITTLE, J. D. C., 1970, Models and managers: the concert of a decision calculus. *Management Science*, **16**(8), B466-B485.

LOYNES, R. M., 1962, The stability of a queue with non-independent interarrival and service times. *Mathematical Proceedings of the Cambridge Philosophical Society,* **58** (3), 497–520.

MAGEE, J. F., 1956, Guides to inventory policy, I, II, III. *Harvard Business Review* **Jan.-Feb.**, 49–60, **Mar.-April**, 103–116, **May-June** 57–70.

MCCRAW, T. K., 2007, *Prophet of Innovation* (Cambridge: Harvard University Press).

MEHROTRA, M., DAWANDE, M., GAVIRNENI, S., DEMIRCI, M. and TAYUR, S., 2011, Production planning with patterns: a problem from processed food manufacturing. *Operations Research*, **59**(2), 267–282.

MOORE, G. A., 1991, *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers* (New York: Harper-Collins).

MORAN, P. A. P., 1954, A probability theory of dams and storage systems, *Australian Journal of Applied Science*, **5** (3), 116–124.

MUCKSTADT, J., TAYUR, S., 1995a, A comparison of alternate kanban control mechanisms: Part 1. Background and Structural Properties, Experimental Results, *IIE Transactions*, **27** (2), 140-150.

MUCKSTADT, J., TAYUR, S., 1995b, A comparison of alternate kanban control mechanisms: Part 2. Experimental Results, *IIE Transactions*, **27** (2), 151-161.

NADDOR, E., 1966, *Inventory Systems* (New York: John Wiley).

PRABHU, N. U. 1965, *Queues and Inventories* (New York: John Wiley & Sons).

RAO, U., SCHELLER-WOLF, A. and TAYUR, S.,2000, Development of a rapid-response supply chain at Caterpillar. *Operations Research,* **48**(2), 189–204.

RAYMOND, F. E., 1931, *Quantity and Economy in Manufacturing* (New York: McGraw-Hill).

REID, C., 1996, *Hilbert* (New York: Springer-Verlag).

REISMAN, D., 2004, *Schumpeter's Market: Enterprise and Evolution* (Northampton, Mass.: Edward Elgar).

ROGERS, E. M, 1962, *Diffusion of Innovation* (New York: Free Press).

SCARF, H. E.,1997, Test sets for integer programs. *Mathematical Programming*, **79**(1), 355–368.

SCHELLER-WOLF, A., TAYUR, S., 2009, Risk sharing in supply chains using order bands – analytical results and managerial insights. *International Journal of Production Economics*, **121**(2), 715–727.

SEIKMAN, P., 2000, New victories in the supply-chain revolution. *Fortune* (Oct 30, 2000).

SIMON, H. A., NEWELL, A., 1958, Heuristic problem solving: the next advance in operations research. *Operations Research*, **6**(1), 1–10.

SIMON, H. A., 1991, *Models of My Life* (New York: Basic Books).

SWAMINATHAN, J., TAYUR, S., 1998, Managing broader product variety through delayed differentiation using vanilla boxes. *Management Science,* **44** (12), 161–72.

SWAMINATHAN, J., TAYUR, S., 2004, Tactical supply chain models for E-business. *Management Science,* **49** (10), 1387-1406.

SZENBERG, M., GOTTESMAN, A. A. and RAMRATTAN, L., (Eds.), 2005, *Paul A. Samuelson: On Being an Economist* (New York: Jorge Pinto Books Inc).

TARDIF, V., TAYUR, S. REARDON, J., STINES, R. and ZIMMERMAN, P., 2010, OR Practice--Implementing seasonal logistics tactics for finished goods distribution at Deere and Company's C&CE division. *Operations Research*, **58**(1), 1–15.

TAYLOR, F.W., 1903, *Shop management* (New York, Harper & Row).

TAYLOR, F.W., 1911, *Principles of Scientific Management* (New York: Harper).

TAYUR, S., 1992, Properties of serial kanban systems. *Queueing Systems*, **12** (3-4), 297–318.

TAYUR, S., 1993a, Computing the optimal policy for capacitated inventory models. *Stochastic Models*, **9** (1), 585–598.

TAYUR, S., 1993b, Structural properties and a heuristic for a kanban controlled serial manufacturing system, *Management Science*, **39**(11), 1347–1368.

TAYUR, S., 1999, *Planned Spontaneity*, presented at Supply Chain Thought Leaders Conference, Brussels.

TAYUR, S., 2000, Improving operations and quoting accurate lead times in a laminate plant. *Interfaces*, **30**(5), 1–15.

TAYUR, S., THOMAS, R. and NATRAJ, N. R., 1995, An algebraic geometry algorithm for scheduling in presence of setups and correlated demands. *Mathematical Programming*, **69**(3), 369–402.

TOWNE, H. R., 1886, The engineer as an economist. *American Society of Mechanical Engineers Transactions*, **7** (2), 428–432.

TROYER, L., SMITH, J., MARSHALL, S., YANIV, E., TAYUR, S., BARKMAN, M., KAYA, A. and LIU, Y., 2005, Improving asset management and order fulfillment at Deere & Company's C&CE division, *Interfaces*, **35**(1), 76–87.

VEBLEN, T. B., 1898, The instinct of workmanship and the irksomeness of labor. *American Journal of Sociology*, **4** (2), 187–201.

WHITIN, T. M., 1953, *The Theory of Inventory Management* (Princeton University Press, Princeton, NJ).

WILCOX, R., 2011, How close can you stand to a software giant? *The Washington Post* (June 25).

YUNES, T., NAPOLITANO, D., SCHELLER-WOLF, A. and TAYUR, S., 2007, Building efficient product portfolios at John Deere and Company. *Operations Research*, **55**(4), 615–629.