Physics of the artificial and its beauty

Prof.dr. Ton de Kok
June 12, 2025

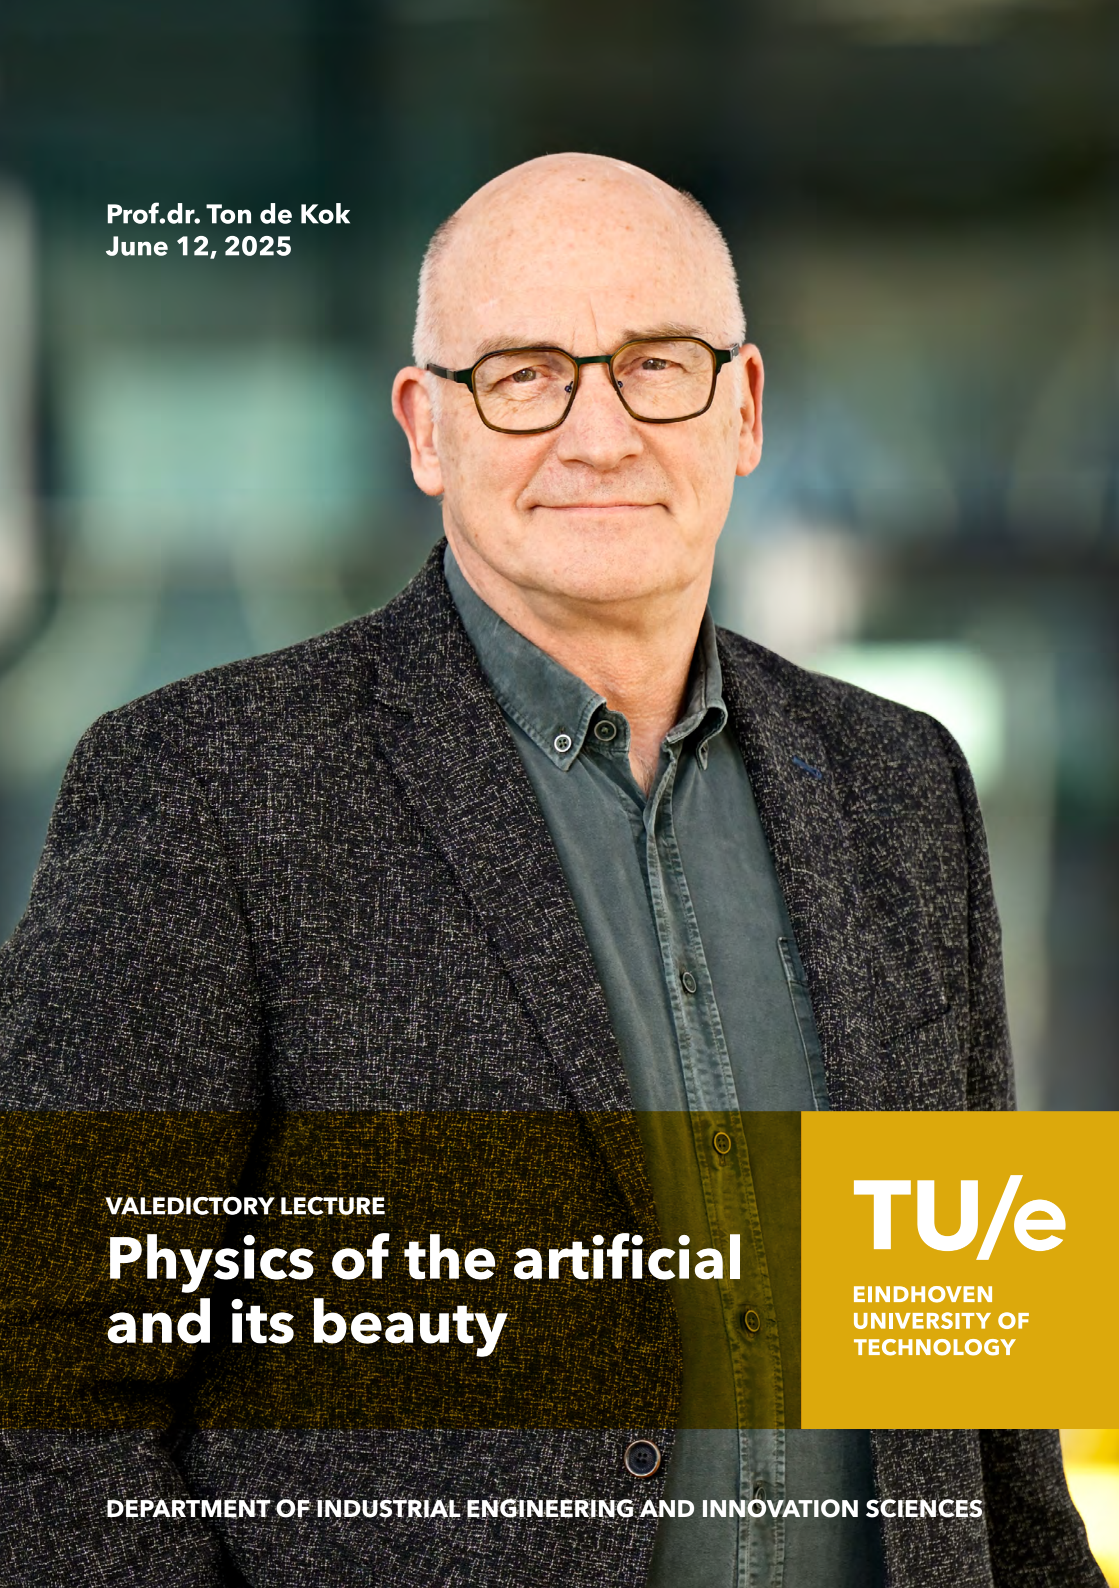

VALEDICTORY LECTURE

Physics of the artificial and its beauty

EINDHOVEN UNIVERSITY OF TECHNOLOGY

DEPARTMENT OF INDUSTRIAL ENGINEERING AND INNOVATION SCIENCES

VALEDICTORY LECTURE PROF.DR. TON DE KOK

Physics of the artificial and its beauty

June 12, 2025
Eindhoven University of Technology

Introduction

In my inaugural lecture when I joined Eindhoven University of Technology in 1993, I compared decision-making in industry with the well-known circus act of spinning plates: the decision-maker must pay attention to many things simultaneously. Now, I realize that this metaphor did not do justice to the situation confronting decision-makers close to the shopfloor of global supply chains. The complexity of their problem is far greater than that of the circus act. You may say that they are not only dealing with plates, but also with cups and cutlery, and that plates, cups and cutlery appear and disappear at random moments in time. In this lecture, I hope to be less metaphorical about what operational decision-making in global supply chains entails and how mathematical modeling and analysis can help both outside observers and the decision-makers themselves to understand their struggles and to continually improve their decisions. I also hope to convey the beauty of the results obtained from the rigorous analysis of mathematical models of human-designed processes. The beauty comes in the form of simplicity, generality and sometimes counterintuitivity.

A system's perspective

To understand what decision-making in global supply chains is about, you first must understand what supply chains are composed of. Supply chains are composed of items and resources. Items can be natural resources and materials, components, sub-assemblies or final products. An item always has the attribute location. A carton of detergent at P&G is different from the same carton at a retailer's warehouse, which is different from the same carton at the retailer's store. Resources can be people, machines, vehicles or tools. Resources transform sets of items into another item, e.g. flour, salt, butter and sugar into a cake using a human, an oven and some tools. Items not used can be stored for future use, but unused resource time cannot be stored. This is a fundamental difference between items and resources.

While I spent quite a lot of time modeling production systems from a resource perspective during my Philips career, after returning to academia I devoted most of my time to modeling supply chains and production systems from an item control perspective. It is appropriate to state here that the work on modeling production systems using queueing theory was quite successful. Much of this success was due to the extensive collaboration my colleague Frans Nijenhuis at Philips' Centre for Quantitative Methods and I had with the 3Bs of the Philips Industrial Engineering organization. Bas Brinkman, Rinus van Breukelen and Jaap Boorsma collectively represented about 100 years of extensive manufacturing systems design knowledge at a company that, I can safely say now, was at the time superior in manufacturing to its Japanese competition. The 3Bs provided the input to create a software tool capable of modeling real-life production systems comprising a variety of resource types. A course on designing manufacturing systems, which we offered to process engineering managers with a mechanical engineering background, showed us that the results produced made sense: it revealed the bottleneck resources, the impact of production batches and transportation batches, and the estimated throughput time with sufficient accuracy to link utilization to flexibility. The IDEAL software was a clear example of the fruitful blend of tacit knowledge, scientific knowledge and engineering solutions using real-life data.

Modeling supply chains will be discussed later in more detail, but a picture is worth a thousand words (and a formula a thousand pictures). Figure 1 depicts a real-life supply chain as a network of relations between items, with stocking locations depicted as triangles and the edges between them showing what is known as the parent-child relationship. The item on the left of an edge is called the child and the item on its right the parent. A child may have multiple parents and a parent may have multiple children. The number attached to the edge indicates the nominal time (in weeks) between releasing an order for the parent item and completing it, i.e. the lead time of the parent item. An order can only be released if all child items are available in sufficient quantity. In today's enterprise resource planning (ERP) systems, information about item orders released, items in stock, and the customer order backlog of items are available in real time, anywhere. At any point in time, we can extract detailed information about the objects of our study.

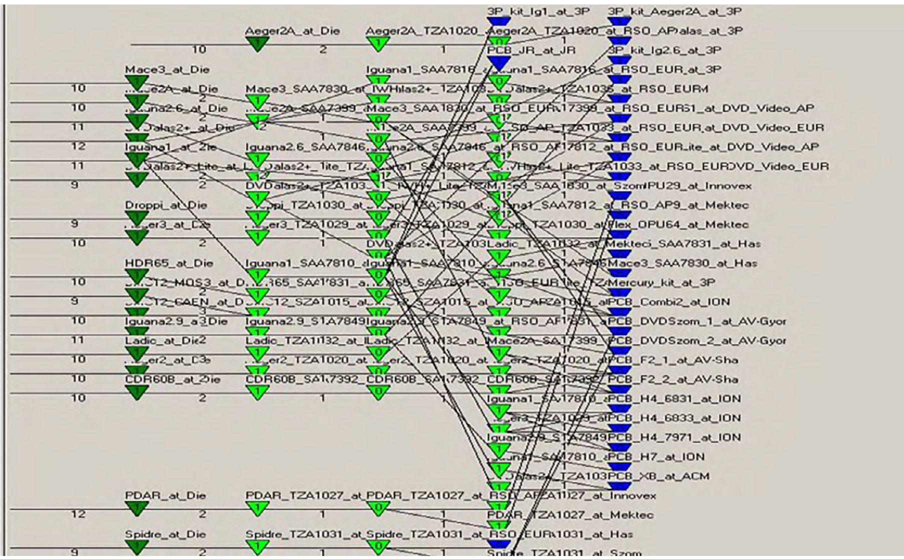

Figure 1. A typical supply chain, i.e. a network

The objectives of our study of global supply chains are several. The goal of a firm is to consistently generate a return on investment in line with shareholders' expectations. In trying to realize this goal, the firm is confronted with conflicting operational objectives (see Figure 2). A high return on investment can be achieved by low investments and high revenues. Low inventories help to reduce investments, but do not help to achieve high revenues as customers may often find inventory items out of stock. Similarly, high utilization reduces investment in resources but increases investments in work-in-progress and increases lead times to customers, which may deter them. Supply chain optimization is a balancing act between investment in resources and material and customer service.

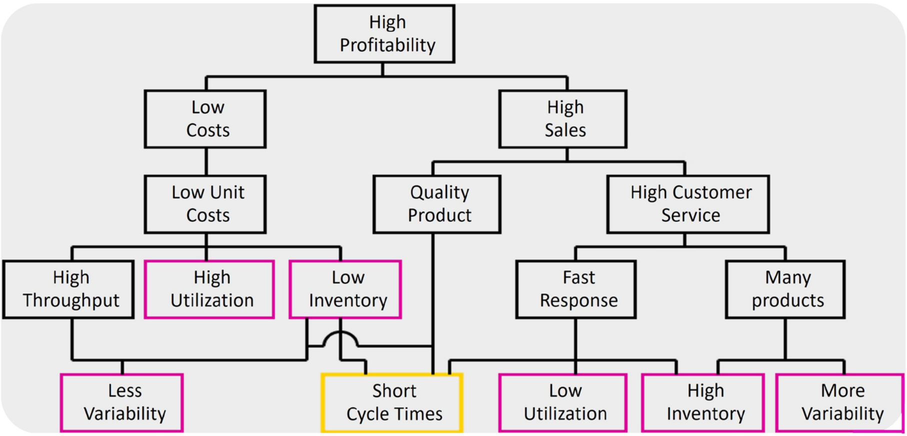

Figure 2. Conflicting objectives when managing a firm (cf. Figure 2 from Hopp and Spearman (2001))

Ideally, management would like to know the firm's efficient frontier, as depicted in Figure 3, which describes the relationship between a target (aggregate) customer service level and the minimal investment in resources and items needed. Such an efficient frontier should be based on a model that explains the relationship between the current investment and (aggregate) customer service.

Over the past 30 years, we have made progress in calculating this efficient frontier for real-life supply chains using mathematical models, finding that substantial reductions in investments were indeed possible while maintaining the required service level. Several MSc and PDEng projects at companies revealed that while the factory under consideration was considered best-in-class, considerable performance improvement was possible. In most situations, we found a religious attitude towards reducing inventory and production batch sizes that often resulted in the firm moving beyond a critical point where processes become inefficient. The most striking example was where we emptied a warehouse that had been filled with scrap over time, returning this scrap to the production site for reuse and

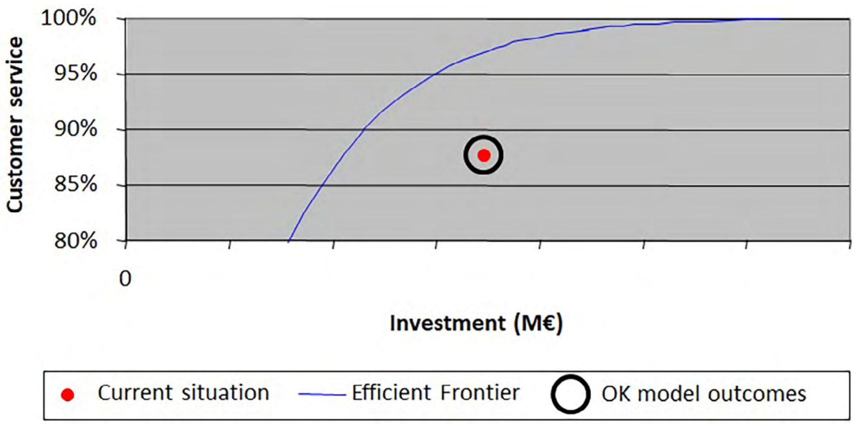

Figure 3. Efficient frontier relating optimal investments to customer service

shutting down two production lines while still maintaining the required output levels with shorter lead times.

In my inaugural speech, I presented my research plans, which could be summarized in two research goals:

1. Determine safety stocks in MRP-I systems to ensure target customer service
2. Create a mock-up of the supply chain that supports concurrent engineering of products and processes

Both research goals followed from my experience at Philips Electronics. Material requirements planning (MRP-I) systems had been implemented on a large-scale in industry, substantially improving material management. However, for each item managed by an MRP-I system, at least three parameters were needed: safety stock, lot size, and lead time. As stated earlier, tools were available at the time to support setting lot sizes and lead times, focusing on the trade-off between efficiency and flexibility. Tools for the analysis of multi-item multi-echelon inventory systems, as described by MRP-I systems, were still not available despite about 30 years of scientific research on these systems, primarily in the US.

The mock-up of the supply chain was motivated by the common belief that 90% of the life cycle cost was determined in the initial product and process engineering phase that itself only accounted for 10% of the life cycle cost. The mock-up was supposed to be a software tool for supply chain design, like the software tool IDEAL for production system design. Nowadays, I would describe such a mock-up as a *strategic twin*.

Research methodology

Before presenting several results from our research, it is appropriate to discuss the research methodology used. We use the research model by Mitroff et al. (1974) depicted in Figure 4 that distinguishes between reality or a real problem situation, a conceptual model, a scientific model and an implemented solution, allowing us to characterize scientific research and problem-solving.

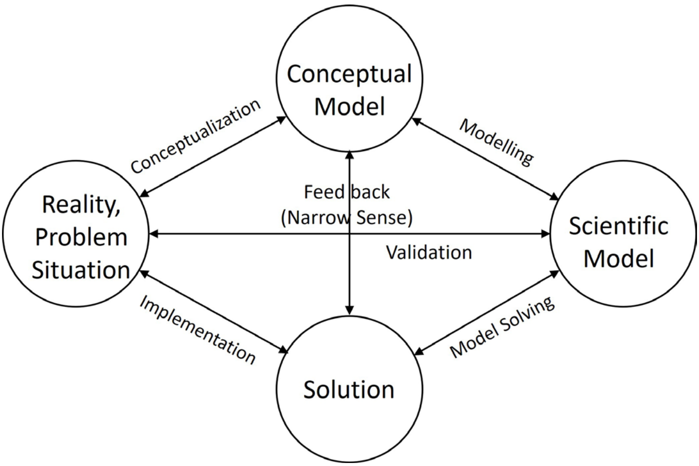

Figure 4: Research model by Mitroff et al. (1974)

A typical problem-solving approach is used by consultants. Based on whatever concepts happen to be *en vogue* at the time, problems are classified and conceptualized towards a solution, which is implemented across a large set of customers. These steps constitute the *consultancy cycle*. Typical examples are lean manufacturing and, more recently, demand-driven MRP. As early as the days when lean was called JIT, we came across implementation of zero-stock concepts that were detrimental to factory output. The above-mentioned example of the factory producing scrap was another, more recent, instance. Let me emphasize that lean concepts are suited for reducing work content and its volatility. But lean concepts such as supermarkets and the associated 'best practice formulas' are simply of no use when trying to manage the trade-off between customer service and investments. In 2010, we did a supply chain optimization study at a company producing a commodity product for a wide range of different markets. We predicted, using mathematically rigorous algorithms, that the inventory targets based on these 'best practice formulas' would yield 70% customer service instead of the 95% customer service targets. In 2011, it was publicly reported that the company's market share had dropped considerably due to bad supply chain performance.

The *conceptual model* typically involves a set of constructs and the relationships between them. Conceptual models help us to understand reality qualitatively. Mathematical models allow for quantification of the input-output behavior of the constructs and their relationships, creating a *scientific model* from which we can derive formal mathematical proofs of the consequences of behavior and relationships. Typical examples are micro-economic models, where game theory is a tool to provide a deeper understanding of the competitive or cooperative behavior of agents, using quantitative models for qualitative insights. In hindsight, the Cobb-Douglas production function is another result of this *axiomatic cycle*. The supply chain management literature is dominated by game-theoretic contributions, which helped to explain why silo-ed supply chains are suboptimal but do not help the operational management of real-life global supply chains (cf. Cachon (2003)).

Physics has been the role model for many sciences, like economics (cf. Von Neumann and Morgenstern (1953), Chapter I.1) and operations research, in its interplay between mathematical modeling and laboratory experiments, which we denote as the *basic science cycle*. Physics is not about optimization but observation. Mathematical models are a tool to describe the relations between measurable variables. The Bohr model of the atom is a prime example of a model that explains what we measure but is not a picture of the real atom, as we learned from electron microscopy.

Around the turn of the century, my colleague Will Bertrand proposed a term for the research methodology employed at the LBS group in Eindhoven, nowadays known as the OPAC group: *model-based empirical research*. For sure, we used the axiomatic cycle, creating stylized realities to explain the observations made during hundreds, if not thousands, of MSc projects in industry. But we also pursued the development of planning and control concepts and mechanisms that we implemented in practice. My research did not focus on qualitative insights from quantitative models. The mathematical models from my research should support operational decision-making and, ideally, propose decisions that could be accepted in 95% of situations, allowing the human planner to focus on the 5% of critical decisions where mathematical models would inevitably be too stylized to produce the right results. This perspective resulted in the development of mathematical models of multi-item multi-echelon (MIME) inventory systems that allowed the modeling of real-life situations and hence could be validated with historical data. This perspective also resulted in a conceptual model of the interactions between a human planner or scheduler and a planning tool with e.g. the MIME system model under the hood, illustrated in Figure 5 below.

It follows from Figure 5 that users can intervene in many ways. Most often, users change solutions obtained by decision support systems (DSS) to create solutions that are infeasible according to the DSS mathematical model, e.g. by speeding up orders and scheduling out customer orders after some negotiation. These human interventions improve customer service, even to 100%, but create a performance gap between customer service according to the model and actual customer service. In De Kok (2018b), this is addressed by the introduction of intervention-independent performance (IIP) indicators, typically measured before humans can intervene in processes. The empirical validation of models discussed below is based on identifying such IIPs. Fortunately, supply chain professionals are aware of the need to measure IIPs.

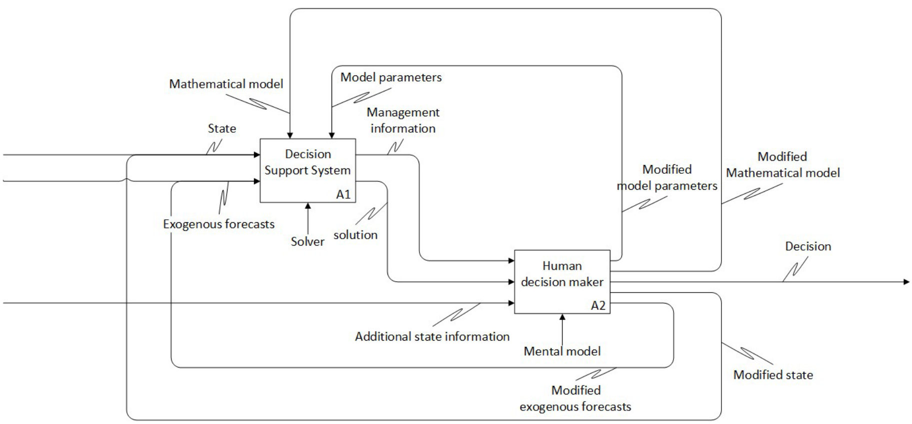

Figure 5. Generic decision function: the interplay between the human and the decision support system, *Wiers and De Kok (2018)*

A major difference between the basic science cycle and model-based empirical research in industrial engineering is the absence of a real laboratory in which you can manipulate reality. We had to deal with data gathered from ERP systems, typically covering a period of at most a few years. The data concerned forecasts of item demand, realized item demand, evolution of item inventory over the period, actual orders released, actual orders completed, cost data, and key performance indicators. We empowered MSc students with Excel spreadsheets and prototype software as the mathematical complexity of the models developed was beyond their capabilities.

For a better understanding of the conceptual models used, we describe a stylized supply chain using a conceptual model. This conceptual model has been used for over four decades and was developed by Will Bertrand, Jacob Wijngaard and Hans Wortmann (cf. Bertrand et al. (1990, 1998)).

Conceptual supply chain model

We first define *production units* that transform input items into output items (see Figure 6a). The inventory levels of input and output items are controlled by the *goodsflow control* decision function that releases output item orders into the production units with the input items needed. The item inventory is depicted by a downward-pointing triangle, which implies that the inventory is controlled (by the goodsflow control decision function). Within the production unit, the output item order is processed by resources, following a specified routing. At each resource, the output item order may have to wait for processing, which creates work-in-process inventory at this resource. This uncontrolled stockpoint is depicted by an upward-pointing triangle.

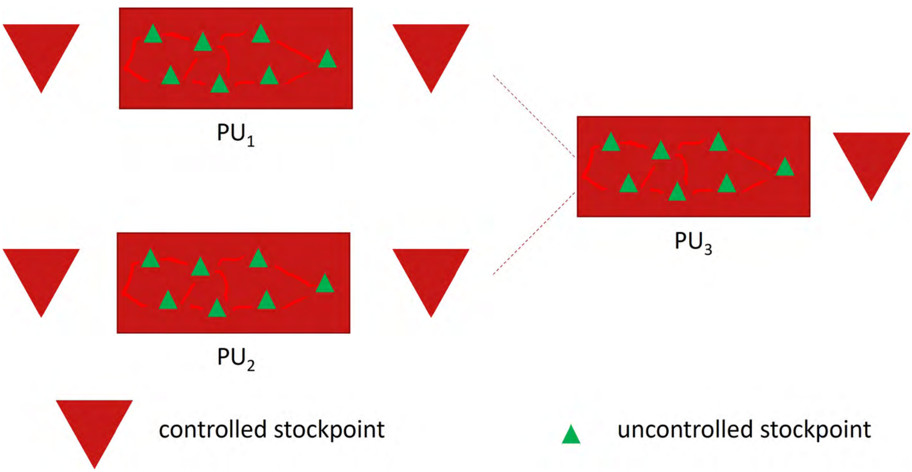

Figure 6a. Conceptual model of the supply chain

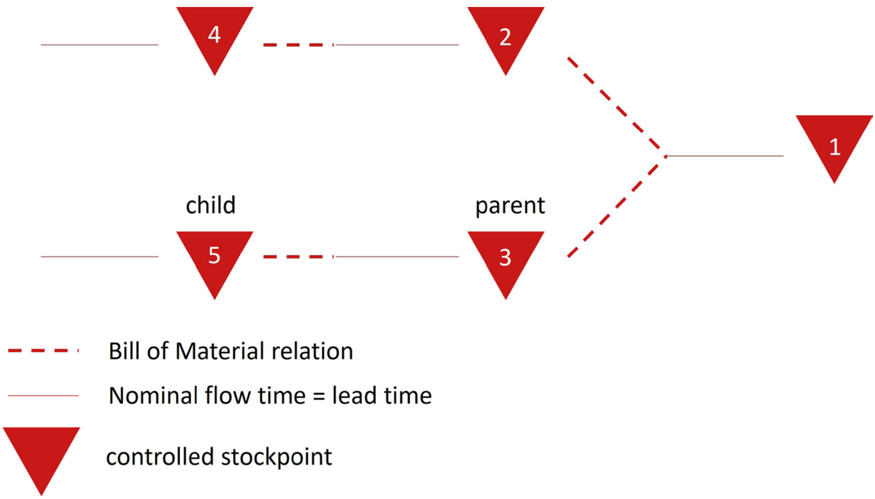

Figure 6b. Simplified conceptual model of the supply chain

Further simplifying the model in Figure 6b, we abstract from the production unit further by assuming a flow time and a lead time, which is a norm for the flow time. Thus, the production unit is only modeled as a delay. These modeling steps create a hierarchical model of supply chains (see Figure 6c), distinguishing between the conceptual (global) supply chain model and the production units with their (shopfloor) control. In most situations, we assume that the delays caused by the production units are equal to the nominal lead time.

In the production unit, the primary trade-off is between utilization, flexibility, and due date reliability, as we mentioned above. At the goodsflow control level, the primary trade-off is between capital investment and customer service. This creates a hierarchy, which we characterize as loosely coupled: the production units are treated as black boxes for the goodsflow control function.
Having defined the conceptual model, which consists of the bill-of-material structure and lead times, we add the demand for an end item to the model in a formal way. In this context, it is typically assumed that decisions are made daily or weekly, that demand in different periods is independent and identically distributed. From historical data, we derive the mean and standard deviation of end-item demand per period. Clearly, one can argue against these assumptions, like one can argue against assuming constant flow times. We will return to these

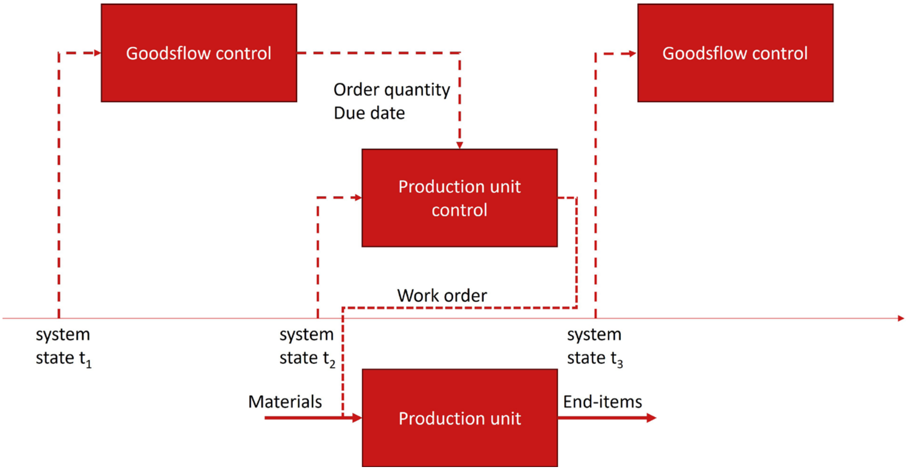

Figure 6c. Loosely coupled hierarchical control of the supply chain

issues when we discuss the main findings of model-based empirical research of global supply chains.

Uncertainty

We have assumed that demand is independent and identically distributed. This implies that we assume that demand is stochastic: we do not know demand before it is actually observed. We assume that we have empirical data that can be used to infer a forecast. When we assume a probability distribution of demand in a period, we take this to the extreme: we assume that we have demand data over a period of infinite length. In order to grasp the significance of this observation, consider the basic experiment of throwing a die. Most people know that the probability of throwing six equals 1/6. But most people also know that it sometimes takes quite some time before you throw your first six. This is the difference between throwing a die infinitely often, after which we find that indeed one out of six throws is a six, but if you throw a die 50 or 100 times you may end up with the following situation.
In Figure 7, the grey bars indicate the normative number of times a number between one and six should be thrown and the orange bar the actual outcomes. As neither 50 nor 100 are divisible by six, we can never realize the norm, but we see that the actual outcome can be quite different from the norm. The same applies to supply chains facing uncertain end-item demand (not to mention uncertain processing and transportation times). It is even worse: there is only a limited amount of relevant data available to infer something about the future demand and, as demand realizes itself, it may well be deviating from the behavior we inferred from our historical data. Interestingly, this leads to practices where uncertainty is either completely ignored or is dealt with by oversimplified methods, such as classifying end items in three or nine categories and planning for some number of periods of inventory for each of the categories. We should be aware that the latter approach is considered a best practice (remember the consultancy cycle).

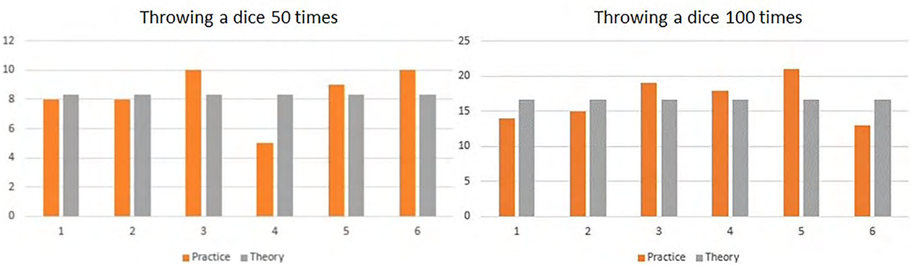

Figure 7. Throwing a dice

Sequential decision-making under uncertainty

Another phenomenon emerging when operationally managing a supply chain is revealed by the infamous Monty Hall problem, although I realize this will not be immediately clear. Consider three doors. Behind one door is a prize, e.g. a car. Behind the other two doors are nothing. You are asked to choose a door. After you make a choice, the quiz master opens a door behind which there is no prize and asks you: do you stay or do you swap to the other closed door? The optimal decision is to swap to the other closed door. By doing so you double the probability of winning the prize. No guarantee, but double the probability! This outcome was debated in the US in the 1970s, even by mathematicians who were convinced that it did not matter. It does!

The Monty Hall problem is an archetypal example of sequential decision-making under uncertainty, where uncertain information becomes known to the decision-maker over time. In this case, the quiz master reveals that there is no prize behind the open door. This new information has an impact on the decision to be taken. In operations management, uncertain information about demand in some future period changes as time passes until this period is the past and we know the actual demand. Typically, period demand uncertainty decreases as the period gets nearer.

Order release decisions must be feasible. This implies that, upon the release of an order quantity, the child items needed must be available in the right quantity to start processing immediately. Most goodsflow control policies calculate a parent item order release quantity under the assumption that child item material availability is sufficient. However, as the child items have been ordered one or more periods earlier, this may not be the case under demand uncertainty – and there are often multiple parent items using the same child items. In this case, the goodsflow control policy must have an *allocation mechanism* that reduces the size of the parent item order to ensure that a feasible parent order can be released.

In today's ERP systems, information is available about all outstanding orders and available stock for all items. This implies that we have information about the future availability of long-lead-time items in the current order pipeline, which may constrain future parent item order releases and thereby future order releases of short-lead-time child items of this parent item. Ignoring this information implies that short-lead-time child items are ordered too early, creating useless inventory, which I refer to as dead stock. In short, we must *synchronize* new orders of short-lead-time child items with outstanding orders of long-lead-time items.

In Figures 8a and 8b, we present one of the simplest possible examples of such a situation. We consider the goodsflow control of five items, where item 1 is assembled from items 3 and 4 within one week, item 2 is assembled from items 4 and 5 within one week, and items 3, 4 and 5 have an order lead time of three, five and one weeks, respectively.

Let us first consider the allocation of items 3-5. The situation at hand is that item 1 wants to order 650 units of 3 and 4 to restore its target availability, while item 2 wants to order 100. Item 3 has only 100 available, so only an item 1 order of 100 can be released. This leaves 150 items 4 and 150 items 5, which allows for an item 2 order release of 100.

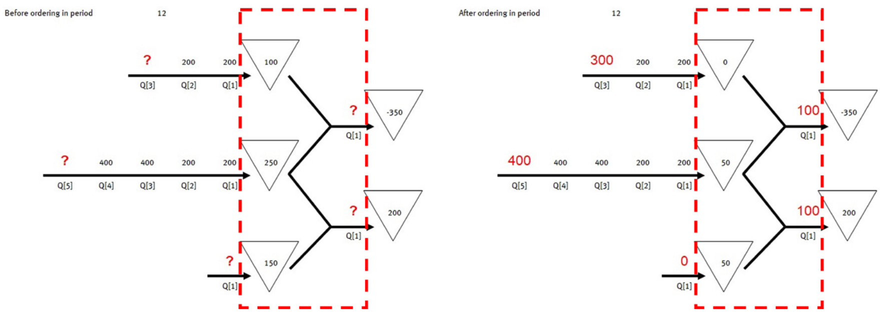

Figure 8a. Allocation in a simple supply chain

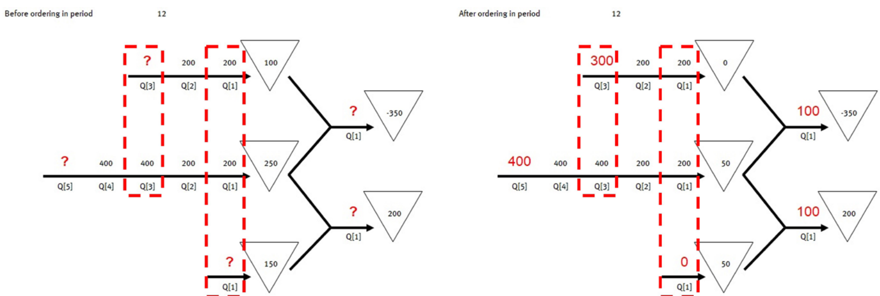

Figure 8b. Synchronization in a simple supply chain

Synchronization concerns the order releases for items 3 and 5. Here, we must realize that the pipeline stocks and on-hand stocks jointly constrain future sales of items 1 and 2. Item 1 has a high backlog due to high sales in the immediate past. It can be shown that an item 3 order release of 500 would yield the fastest recovery of that situation. However, item 4 covers future demand of both items 1 and 2. As this has been taken into account in past order releases, we can focus on Q[3] of item 4, which is 400 units. These 400 units will be assembled into items 1 and 2 in three weeks. Thus, ordering 500 of item 3 is creating at least a stock of 100 in three weeks' time that cannot be used. Using a policy discussed below in more detail, we find that an item 3 order of 300 can prevent future excess stock. Similarly, we can conclude from the quantities Q[1] for items 3 and 4 that these have been synchronized, leaving no items 4 for allocation to item 2. Thus, no order is released for item 5.

As noted above, ERP systems make this information available to planning systems for goodsflow control. If the planning system has a solver such as CPLEX or Gurobi under the hood, then both *allocation* and *synchronization constraints* are taken into account. As goodsflow control concerns thousands of items to be planned over hundreds of days at high frequency, the computational burden for such solvers is very high.

Most ERP systems use the computationally efficient MRP-I logic for goodsflow control. With the way it is commonly explained in textbooks and online videos, MRP-I logic is easy to understand. But these explanations never discuss allocation and synchronization constraints. Allocation constraints are only implicitly addressed as their violation can be identified as 'past-due' (i.e. too late) orders or orders to be 're-scheduled in'. Such orders are brought to the attention of material planners as error messages, which can number in hundreds and even thousands and have to be dealt with manually. This explains why companies making complex products have tens or hundreds of planners. These planners deal with a problem of mind-blowing complexity that involves thousands of decisions to be taken daily on a problem that seems to change every day, if not every hour. Before addressing a computationally efficient alternative for MRP-I logic that respects allocation and synchronization constraints, we discuss the impact of uncertainty on tactical and strategic decisions.

Inventory capital allocation in the supply chain

Operational planning of global supply chains is supported by planning systems. As mentioned earlier, such planning systems need *planning parameters* for each item: lead time, lot size and safety stock. The safety stock parameter addresses the impact of uncertainty in demand and supply. By increasing (decreasing) the safety stock of an item, we increase (decrease) the inventory capital invested for this item and increase (decrease) the customer service for its parents. Setting safety stocks is a tactical decision, but the resulting investment in inventory capital is of strategic importance (cf. discussion on the efficient frontier).

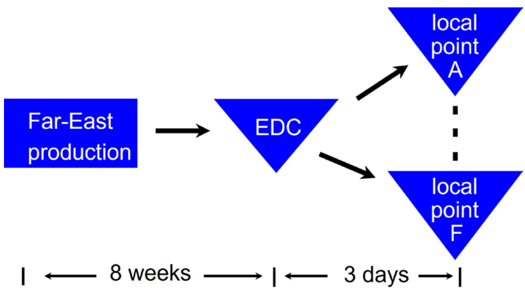

Figure 9. A European supply chain

As an example of this, we consider the situation at a consumer electronics company, the structure of whose supply chain in 1991 is shown in Figure 9. Thanks to the elimination of the EU borders, it was no longer necessary to store end products in each EU country, so it was decided to establish a European Distribution Center (EDC) in Eindhoven and six regional warehouses to satisfy the EU demand. The company now had to take an unprecedented decision: how much inventory to hold at the EDC and how much at the regional warehouses? Based upon this real-life case, I built a quiz that I played with thousands of students and professionals in the 1990s. The quiz is as follows.

An end product is produced in the far east. The lead time of an EDC order for this product is eight weeks, which includes waiting for order release, producing the

order and transporting the order by sea. Regional warehouses A to F order from the EDC and receive their orders from the EDC three days later by truck, unless there is not enough inventory at the EDC. Before implementing the EDC, the EU countries kept, on average, seven weeks of inventory. With a demand of 100 per week, seven weeks of average inventory equals 700. So, it was decided to allocate these seven weeks of inventory between the EDC and regional warehouses.
As this could be done in many different ways, it was decided to first consider two scenarios. In scenario 1, five weeks of inventory is held at the EDC and two weeks in each of the regional warehouses. In scenario 2, two weeks of inventory is held at the EDC and five weeks in each of the regional warehouses. Which scenario is best regarding customer service (as the amount of inventory is the same)?

In order to answer this question, we decided to use discrete event simulation as anyone could verify this, even in 1991. We used some real-world cases to determine a reasonable (compound Poisson) demand process and used continuous time (s,nQ) policies for all seven stockpoints. The policy parameters were chosen to ensure the allocation of average stock was as desired. The results of the simulation study are summarized in Table 1.

	Scenario 1 (5,2)	Scenario 2 (2,5)	Scenario 3 (1,3)	Scenario 4 (1,2)
Stock in weeks	7.1	7.1	4.1	3.1
Regional fill rate	94%	100%	97%	92%
EDC fill rate	100%	95%	80%	80%

Table 1. Inventory capital allocation in a distribution system

Clearly, scenario 2 is superior as it provides (close to) 100% customer service against only 94% for scenario 1. Interestingly, roughly 90% of the respondents to the quiz chose scenario 1. Again, the finding from the mathematical model of this distribution system is counterintuitive. Also interesting was the resistance to accepting this finding among professionals. It has taken more than 20 years for this finding to achieve widespread acceptance, despite it being observed earlier by others. For example, Graves (1985) noted that when optimizing a similar two-echelon system for spare parts, hardly any stock was held at the upstream stock location.

A policy for operational control of general supply chains under demand uncertainty

While distribution systems like those discussed above are widespread, real supply chains, as mentioned earlier, are networks, since parent items may have multiple children and child items multiple parents. It is commonly accepted in the field that there is no hope of finding the optimal policy for such a general supply chain structure. In fact, there is no hope to find the optimal policy for divergent systems due to the *curse of dimensionality*: the optimization problem is too big to store and solve (cf. Powell (2007); Dogru et al. (2009)). An optimal policy for convergent systems where each child item has at most one parent item has been derived by Rosling (1989). In De Kok and Visschers (1999), a policy is proposed that generalizes the known optimal policies for divergent and convergent systems to a policy for general systems with a bill-of-material matrix with only (0.1) elements. In De Kok (2018a), a description of this policy is provided for any bill-of-material structure.

The policy is based on an idea that cuts the Gordian knot of operational control of real-life supply chains: allocation before synchronization. This idea enables us to determine one or more *decision node structures* from the supply chain bill of material and the item lead times (see Figure 10). Each decision node structure constitutes a distribution system and, at the same time, extends the synchronization principle developed by Rosling (1989). The decision node structures give insight into the natural decision hierarchy imposed by the bill of material and the item lead times. For the decision node structures, we can determine near-optimal policies (cf. Diks and De Kok (1998, 1999)) to determine order release quantities for all items in the original supply chain that respect all material availability constraints. The policy proposed is called a synchronized base stock (SBS) policy.

In De Kok and Fransoo (2003) and Spitter (2005), simulation experiments show that the SBS policy outperforms a rolling scheduling policy using LP to solve the planning problem in each period. The SBS policy has *linear complexity* for determining the order releases for all items. As it builds on the recursive algorithm

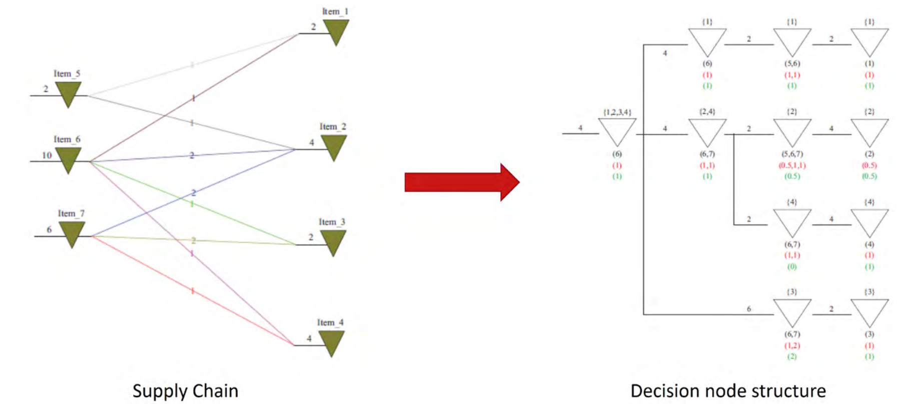

Figure 10. From supply chain to decision node structure

proposed in Diks and De Kok (1999), the computation of close-to-optimal control parameters (i.e. base stock levels) also has *linear complexity*.

The SBS policies are not optimal. But they have an important feature: they explain the relationship between inventory capital investment and operational customer service. De Kok (2018a) refers to a number of MSc theses using historical data to determine average stocks and average lot sizes for all supply chain items and the mean and variance of demand. With this information, the parameters of the SBS policy that would yield the same average stocks and lot sizes could be determined. Using this SBS policy, the end-item customer service levels were computed and compared against the actual service levels measured over the data collection period. In all cases reported, the computed aggregate service levels were close to the service levels measured, to the extent possible, as argued above with the impact of uncertainty in the context of throwing a die.

SBS policies can be extended with the concept of customer lead times. The customer lead time of an item is the (nominal) time between the moment a customer orders an item and the moment this order is shipped to (or received by) the customer. By introducing the customer lead time, the SBS policy concept unifies the well-known customer order decoupling points (Hoekstra and Romme (1992)) of make-to-stock, assemble-to-order and make-to-order and generalizes the concept from the perspective of an individual company to that of the supply chain.

With the SBS policy and efficient algorithms to compute its parameters, we had the ingredients to develop software for the analysis and optimization of real-life supply chains. An STW valorization grant provided funding for developing the ChainScope software that has been used since 2009 for several MSc thesis projects (see Figure 10).

Admittedly, the versatility of the ChainScope software, allowing both model validation and supply chain optimization, comes at a price. The underlying mathematical model incorporates aspects of real-life supply chains, such as item yield and item lot sizes, for which no exact methods yet exist. The extensive empirical testing of the software indicates that the heuristics developed generate results that can be used in practice. Clearly, further research is needed to formally assess the quality of the heuristics.

Two birds with one stone: operational planning and strategic twins for practical use
With the SBS policies, we can efficiently compute safety stock parameters for all items. This implies that we can build a model of the supply chain before making important design decisions. This *strategic twin* allows us to make trade-offs, such as flexibility versus costs, commonality versus diversity. We have a policy that can be used operationally for day-to-day decision-making, and also as a basis for tactical and strategic decision-making. There is not really a restriction on the number of items modeled, although for tactical and strategic decisions it is recommended to model only the most expensive items.

In addition to empirically validating SBS policies, we also implemented a modified SBS policy for operational decision-making in the context of collaborative planning across a high-volume electronics supply chain. At Philips Semiconductors, an SBS-based planning tool was used for weekly planning from 2001 to 2007 with a number of key customers and the electronic manufacturing services (EMS) companies in between (cf. De Kok et al. (2005)). This showed that these (straightforwardly) modified SBS policies can be used in forecast-driven, highly volatile, non-stationary demand and supply situations.

The empirical validation experiments reveal another striking phenomenon: the performance of supply chains is primarily driven by the average item stocks and average lead times. The specific control policy used and the manual modifications of the proposed orders that result in these observed quantities do not seem to be important. In practice, these modifications would be very hard to identify, especially when the control policy generates infeasible plans.

Figure 11. Companies for which ChainScope software was used for supply chain analysis

Our observations of empirical validity also imply that the assumption of constant item flow times is defensible. A policy for a multi-item multi-echelon inventory system takes into account the fact that, due to high demand in the past, material availability today may be low, requiring us to ration child item availability among parent items. The policy ensures the feasibility of orders released and the only reason for lateness is a lack of resources. We observed in practice that there are many ways to allocate additional resources to production orders that run the risk of being late: examples include re-routing, expediting, and overtime. It seems in practice that production management understands the detrimental impact of low due date reliability on the performance of the supply chain.

A critical remark could be made here: empirical validation uses information not known to the decision-maker when decisions are taken. Indeed, we use historical demand data and fit a (gamma) distribution to the data. But given the availability of these data in hindsight, the mathematical model should be able to explain the actual customer service. That is exactly what the mathematical model is doing!

We do not use the demand data themselves, as we often see in discrete event simulation studies. Our validation procedure is comparable to using data from the die-throwing experiment. We assume equal probability and fit a distribution using the mean and standard deviation of the data, which look like that presented in Table 2 for five experiments where we consider 50 and 100 throws.

		Experiment					Aggregate	Theory
N	Parameter	1	2	3	4	5		
50	μ	3.28	3.70	3.50	3.40	3.32	3.44	3.50
	σ	1.67	1.69	1.74	1.63	1.90	1.73	1.71
100	μ	3.45	3.60	3.43	3.38	3.49	3.47	3.50
	σ	1.70	1.70	1.74	1.66	1.71	1.70	1.71

Table 2. Mean μ and standard deviation σ of historical data from die-throwing experiments

We argue that it is more robust to use average and standard deviation as aggregate information from our experiments than empirical distribution that uses detailed data; similar observations have been reported again and again in the forecasting literature. Single (or simple) exponential smoothing (SES) is often one of the best performing forecasting rules, provided the smoothing coefficient is chosen to be low, say 0.1. SES implicitly weighs all past demand data and aggregates these into a single number, which is used as a forecast of demand for future periods. The success of SES is another indication that demand at the item level is indistinguishable from being stationary.

Algorithms for single-echelon models under the assumptions mentioned here have been implemented in practice, showing the validity of the models but also the validity of using historical demand data for extrapolation into forecasts for future demand. Implementation of these models implies that the parameters of control policies are calculated using data about demand, lead times and service level targets. The parameters are then used in the inventory management systems to generate replenishment orders. The actual service performance is closely monitored. In many cases, service performance was on target. A supply chain manager at an office supplies wholesaler stated: with the new inventory management tool, we have shown ourselves to be in control, which implies that we will get the inventory budget for next year as requested from our general management. Some case studies with software based on SBS policies showed similar results, but the large-scale implementation of this software is a challenge for the future.

Safety times in the order-driven supply chain

The SBS policy can be used for goodsflow control upstream of the customer order decoupling points (CODPs) of the supply chain. Downstream of the CODP, production and distribution activities are order-driven. The key uncertainty lies in the stochasticity of these activities. This implies that we must protect against time uncertainty, as opposed to quantity uncertainty upstream of the CODP. For this purpose, we use time buffers or safety times. In Atan et al. (2016), we describe a mathematical model for this situation under the assumption that an activity has at most one successor. This fits many real-life situations, as typically an order comprises a single complex product that is produced from a set of modules and components, which are assembled in several steps, and where subassemblies and the final product are tested regularly and rework is executed when necessary. The work of Atan et al. (2016) for two-echelon assembly systems is extended in Jansen (2019) to multi-echelon assembly systems.

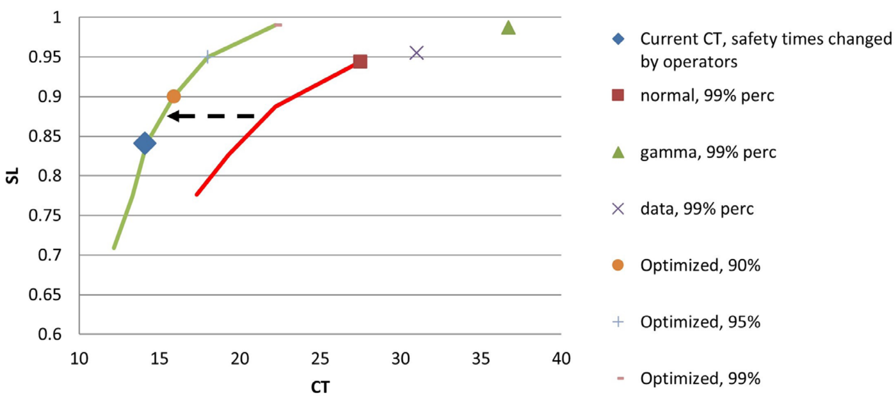

Figure 12. Efficient frontier of cycle time (CT) versus customer service (SL)

The model was validated at a manufacturer of capital goods, for which we also created a strategic twin by developing prototype software. This allowed us to compare the existing heuristic approach at the company based on setting a buffer time for each activity, based on the assumption of a normally distributed activity time and some probability of meeting the activity's lead time. The key trade-off is the length of cycle time (CT, in days), i.e. the nominal time between the start of assembly activities for a subassembly or final product and its completion versus the service level (SL), i.e. the reliability of delivery before or at this cycle time. The longer the cycle time, the higher the reliability of on-time completion. Figure 12 shows the efficient frontier computed for a complex module with our validated model, the non-efficient solution of the heuristic and the current situation, which resulted from manual modification of the safety times by operators, unknown to the business engineers who proposed buffer times based on the heuristic. It is clear that the 'stage-by-stage' heuristic is highly inefficient. It is remarkable, and beautiful, to see that the manual changes made by the operators put the current situation on the efficient frontier of the module.

Mathematically, the problem of setting safety times in the order-driven part of the supply chain is similar to that of setting safety stocks in the forecast-driven part of the supply chain. We learned from the EDC quiz that most of the inventory buffer must be held at the end of the supply chain. This result generalizes to any supply chain structure, and it also holds for safety time buffers in the order-driven part of the supply chain: most of the safety time should be at the end of the production and distribution process involved.

The beauty of this generic outcome is that the trade-off between customer service and inventory capital investment is concentrated at the most downstream part of the supply chain. The upstream part of the supply chain hardly plays a role in this trade-off. This implies that we can focus here on the trade-off between work-in-progress and batch stock investments on the one hand and utilization of fixed assets on the other. These decoupled trade-offs are in favor of sustainable production by efficient use of assets. Further research is needed to provide more extensive support to this observation that is similar to the finding that the cost performance of a heuristic that uses the economic order quantity for the lot sizing parameter of an inventory control policy under demand uncertainty is very good.

Aggregate customer service targets and inventory capital allocation

Another inventory capital allocation problem emerged at the office supply wholesaler mentioned above. This wholesaler supplies multiple items to a single customer, leading naturally to agreements on an aggregate service level. This implies that we take a weighted sum of the item service levels over all items delivered to this customer. It is clear that this weighing is no more than adding apples and oranges. However, it gives the wholesaler an opportunity to show a more reliable performance, like when throwing multiple dice and taking averages of the outcomes (cf. Table 2, aggregate column). But it also creates a new optimization problem: minimizing the inventory end-item capital invested subject to the aggregate service level constraint.

To solve the problem, we must decide on the weights of the individual service levels and on the service level criterion. Regarding the latter, we have already provided arguments to use the ready rate as the service level measure, i.e. the probability of no stockout at the end of an arbitrary time unit. This is easy to compute from the data in the inventory management system, and we can use holding and penalty costs to set cost-optimal end-item service levels. However, the choice of the weights and the aggregate ready rate constraint will impact these optimal end-item service levels. But how?

Again, we used the proven technology of mathematical models for inventory systems. We considered two basic and natural weighing policies:

a. The end-item service level is weighted with the number of items sold per time unit. This is denoted as the volume-based weighing policy.
b. The end-item service level is weighted with the financial turnover per time unit. This is denoted as the revenue-based weighing policy.

Our analysis of the optimization led to the following insights (cf. van Donselaar et al (2021)):

a. Under the volume-based weights, the optimal policy creates (relatively) high service levels for cheap items and (relatively) low service levels for expensive items.
b. Under the revenue-based weights, the optimal policy creates **equal service levels** for all items, i.e. equal to the aggregate service level required.

The striking and attractive result for the revenue-based weights holds not only for single echelon models but also for any multi-item multi-echelon system. It is a corollary of a generic finding that under linear holding cost h and linear penalty costs p for an end-item, the optimal control parameters of an inventory control policy must be chosen such that the ready rate of this end item equals $\frac{p}{p+h}$. This result holds for any multi-echelon inventory system where end items are controlled by well-known single-echelon control policies, and also under the non-stationary autocorrelated forecast errors that come with classical forecasting methods. The result has been obscured by the emphasis on the fill rate as the preferred service level for inventory management. Under the fill rate measure, the outcome is similar to the volume-based weight outcome under the ready rate for both weighing schemes.

The beauty of the result is that we do not need to differentiate between service levels for different items. This greatly simplifies management of large numbers of end items for different customers, as is the case for most wholesalers. The result requires rigorous mathematical analysis, but once this is done, it can be easily implemented in inventory management systems.

Epilogue

Looking back on my work as a researcher, it is clear that I was lucky to be trained as a pure and applied mathematician at Leiden University, to be introduced to various ideas, methods and techniques to tackle mathematically intractable stochastic problems by my PhD supervisor Henk Tijms, and to find at Philips Electronics that the mathematical rigor and the approximation ideas were well suited to solving real-world problems. Solving urgent problems for Philips factories, sales organizations and ICT support stimulated my modeling skills. I am well known as a sender, a person that always talks. But I can assure you, I have used my time well to listen to many people with tacit knowledge, people that knew where I had to model and analyze.

But I took nothing for granted. Over the years, I developed the belief that if I could not understand a line of thought, a proof, a data set, it was likely that there was something wrong. This may sound arrogant, but it turned out to be a good strategy. I found fundamental errors in conceptual models for production and supply chain management. I found errors in scientific papers. This is no fun when the papers are written by reputed scientists. As I primarily focused on stochastic problems, most of the time the errors were caused by the authors' ignorance of fundamentals of mathematics and, in particular, probability theory. In many cases, a rigorous analysis revealed the beauty of the problem, as I tried to show above. I also found errors in budget systems for educational work. Let me say this about it: every hour spent on teaching obligations should receive the same financial compensation.

Being a teacher for such a long time, I started appreciating the learning cycle that brings you from 'not knowing that you don't know' to 'knowing that you do not know' to 'knowing that you know' to the state of bliss of 'not knowing that you know'. This appreciation stems from the observation that there are many professionals who have developed tacit knowledge that awaits further study by scientists. This appreciation stems from another observation that many professionals and scientists that are confronted with the state of 'knowing that you do not know' after their education at high school and university seem to develop resistance to this fact, downplaying the matter they do not know, ignoring it or simply denying it. Above, I presented a number of research results that may be against your intuition and mine. I leave it to you to do some introspection on how you felt about being wrong. Like most of you, as a high school pupil and at university, I knew that I did not know. But over my career, I again and again realized: I know nothing about the matter I have to deal with in the near future. I found out that this is a realistic perspective, but it creates a permanent state of doubt.

Hopefully, you are not angry with me about my subjective view on the human condition. It has never been my intention to show off or to discredit other people. It was my intention to understand the physics of the artificial, much like when I wanted to understand the physics of nature. Once understood, all benefit from more effective and efficient processes, such that stewardship comes naturally. We can build a better world on accepting each other's knowledge and expertise. As Will Bertrand explained to me, and here I am paraphrasing: build a wall around your expertise, put a clear sign on it and bend over the wall to find the expertise of others, and collaborate to move the edge of science and to tackle the societal challenges to be tackled.

Word of thanks

Over all those years of studying and working, I met a lot of people from whom I have learned and with whom I have enjoyed working. It is impossible to name everyone, that should be clear. Above, looking back on my research, I have mentioned a few people by name. I would like to thank groups of people.

First of all, I would like to thank the ladies of the student administration because they always answered my questions directly and always did their work with great dedication. My PhD students helped me to answer the questions that came naturally from working with industry. Without you, the content of my farewell speech would have been paltry. I would like to thank my colleagues at LBS and OPAC, both support staff and scientists, for their collegiality and opposition. Indeed, I am not always right. We have all worked together to build a group that has gained an international reputation so that the OPAC group is now brimming with young talent.

I thank all my scientific colleagues from all over the world, some of whom are here. The transition from Philips to academia was tough, but I feel privileged to have met so many beautiful people from whom I have been able to learn so much. I realize that my managers have put their trust in me. I am very grateful for that. It made development possible at all times. I haven't always been easy to work with. I said what I thought of something and sometimes forgot that this was unpleasant for the recipient. I have apologized for it regularly, and rightly so. I hope it has dawned on me that I was always committed to the bigger picture: the group, the department, the university.

I thank my friends from the Christian community De Hooge Berkt for their support to my family over many years and creating a place where I found the foundation from which I tried to serve the various communities I have been part of.

Everything I have done would not have been possible without my home front. Irene, you ended up on a rollercoaster with me, partly because of me and partly because of what life seemed to have in store for us. We have done it together to this day. Thank you for who you are. Merel and Diede, what beautiful women you have become. I look at you with admiration and hope you don't blame me when I say that I see the same passion for what you are doing as I have. Always remain a little naïve. Finally, I would like to thank Casper, our son. His life force is unimaginable. There is no greater contrast between a person with many gifts like me and a person who is dependent on others in everything. But also a person with the gift of showing that in the small things there is everything that makes life valuable. You and all your group members over the years have taught me to have an eye for the needs of the people around me, to always end up with solidarity and to let go of all appearances. It is my faith in the Love that transcends us all and the Love that resides in all of us for each other that makes me believe that there will be a heaven here on earth. Let's work together, pooling our distinct talents.

Ik heb gezegd.

Acknowledgements

Herewith I thank Will Bertrand for his continual support and being a source of inspiration over so many years. I thank Reha Uzsoy for his friendship, stimulating creativity, and for his willingness to critically read the manuscript that preceded the above text.

References

Atan, Z., de Kok, T., Dellaert, N.P., van Boxel, R. and Janssen, F., 2016. Setting planned leadtimes in customer-order-driven assembly systems. *Manufacturing & Service Operations Management*, *18*(1), pp.122-140.

Bertrand, J.W., Wortmann, J.C. and Wijngaard, J., 1990. *Production control: a structural and design oriented approach*. Elsevier Science Inc..

Bertrand, J.W.M., Wortmann, J.C., Wijngaard, J (1998). Productiebeheersing en material management. Houten, Nederland. Educatieve Partners Nederland BV.

Cachon, G.P., 2003. Supply chain coordination with contracts. *Handbooks in operations research and management science*, *11*, pp.227-339.

De Kok, T., 2018a. Inventory management: Modeling real-life supply chains and empirical validity. *Foundations and Trends® in Technology, Information and Operations Management*, *11*(4), pp.343-437.

De Kok, T.G., 2018b. Modelling short-term manufacturing flexibility by human intervention and its impact on performance. *International Journal of Production Research*, *56*(1-2), pp.447-458.

De Kok, T., Janssen, F., Van Doremalen, J., Van Wachem, E., Clerkx, M. and Peeters, W., 2005. Philips electronics synchronizes its supply chain to end the bullwhip effect. *Interfaces*, *35*(1), pp.37-48.

De Kok, T.G. and Fransoo, J.C., 2003. Planning supply chain operations: definition and comparison of planning concepts. *Handbooks in operations research and management science*, *11*, pp.597-675.

De Kok, T.G. and Visschers, J.W., 1999. Analysis of assembly systems with service level constraints. *International Journal of Production Economics*, *59*(1-3), pp.313-326.

Diks, E.B. and De Kok, A.G., 1998. Optimal control of a divergent multi-echelon inventory system. *European journal of operational research*, *111*(1), pp.75-97.

Diks, E.B. and De Kok, A.G., 1999. Computational results for the control of a divergent N-echelon inventory system. *International Journal of Production Economics*, *59*(1-3), pp.327-336.

Doğru, M.K., De Kok, A.G. and van Houtum, G.J., 2009. A numerical study on the effect of the balance assumption in one-warehouse multi-retailer inventory systems. *Flexible services and manufacturing journal*, *21*, pp.114-147.

Graves, S.C., 1985. A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management science*, *31*(10), pp.1247-1256.

S. Hoekstra, J. Romme, Integrated Logistics Structures: Developing Customer Oriented Goods Flow, McGraw-Hill, London, 1992.

Jansen, S., 2019. Quantitative models for stochastic project planning. TUE PhD thesis.

Mitroff, I. I., Betz, F., Pondy, L. R., and Sagasti, F. (1974). On Managing Science in the Systems Age: Two Schemas for the Study of Science as a Whole Systems Phenomenon. Technical Report 3

Powell, W.B., 2007. *Approximate Dynamic Programming: Solving the curses of dimensionality* (Vol. 703). John Wiley & Sons.

Rosling, K., 1989. Optimal inventory policies for assembly systems under random demands. *Operations Research*, *37*(4), pp.565-579.

Spitter, J.M., 2005. Rolling schedule approaches for supply chain operations planning. TUE PhD thesis

Van Donselaar, K., Broekmeulen, R. and de Kok, T., 2021. Heuristics for setting reorder levels in periodic review inventory systems with an aggregate service constraint. *International Journal of Production Economics*, *237*, p.108137.

Von Neumann, J. and Morgenstern, O., 1953. *Theory of games and economic behavior, 3rd edition*. Princeton University Press.

Wiers, V.C. and A. (Ton) G. De Kok, 2018. *Designing, selecting, implementing and using APS systems*. Springer International Publishing.

Curriculum Vitae

Prof.dr. Ton de Kok was appointed full-time professor of Quantitative Analysis of Logistics Control Systems at the Department of Industrial Engineering and Innovation Sciences at Eindhoven University of Technology (TU/e) on August 1, 1992. He also served as a part-time professor of Industrial Mathematics from 1991 to 1998.

Ton de Kok is a full professor at the School of Industrial Engineering at Eindhoven University of Technology and the director of CWI in Amsterdam. His research concerns the optimization of operational business processes under uncertainty in the context of supply chain management, transportation management and production management. His work has been implemented in many different industries, ranging from transportation process industries to capital goods industries. The empirical validity of the models and their analysis has provided clear evidence of the importance of stationary stochastic models. Before returning to academia in 1992, Ton worked at Philips Electronics from 1985 as an operations research consultant and logistics innovation manager of consumer electronics.

Colophon

Production
Communication Expertise Center

Cover photography
Bart van Overbeeke Photography, Eindhoven

Design
Grefo Prepress, Eindhoven

Digital version:
research.tue.nl/en/

Visiting address
Building 1, Auditorium
Groene Loper, Eindhoven
The Netherlands

Navigation
De Zaale, Eindhoven

Postal address
PO Box 513
5600 MB Eindhoven
The Netherlands
Tel. +31 (0)40 247 9111
www.tue.nl/map

EINDHOVEN
UNIVERSITY OF
TECHNOLOGY