

Online Appendix for

Strategic Inventory Placement in Supply Chains: Non-Stationary Demand

Stephen C. Graves
Sean P. Willems

Leaders for Manufacturing Program and
A. P. Sloan School of Management
Massachusetts Institute of Technology
Cambridge MA 02139-4307

Boston University
School of Management
Boston, MA 02215

February 2006

This online appendix accompanies Graves and Willems (2006). We first present a detailed numerical analysis and a real-world example that demonstrate how a constant service time policy performs relative to a dynamic service time policy. A constant-service-time policy is equivalent to assuming stationary locations for safety stocks; nevertheless, the actual base and safety stock levels will change as the demand process changes. A dynamic policy dynamically modifies the service times in every time period. We then describe and document the planning problem that is implied by the adoption of the periodic-review base-stock replenishment policy with time-varying base-stock levels.

1. Introduction

As discussed in Graves and Willems (2006), the primary intent of this research is to examine the problem of locating safety stocks in a supply chain, in a way that accounts for both uncertain and non-stationary demand processes. We use the modeling framework from Graves and Willems (2000), and propose a demand model that permits consideration of non-stationary demand. A key result is to show that the optimization from Graves and Willems (2000) for finding the safety-stock placement in a supply chain applies to this case of non-stationary demand.

In this online appendix, Section 2 reports on a computational study that explores the near-optimality of a constant-service-time policy versus a dynamic-service-time-policy. In section 3, we present an example to illustrate how the proposed policy might apply in practice. In section 4, we present the planning problem for adjusting the base stocks that would be prescribed by a CST policy.

2. Serial-Line Numerical Experiments

In this section, we develop a numerical experiment to examine how well the constant-service-time (CST) policy performs relative to a dynamic-service-time (DST) policy that optimizes the service times in each period. We consider an N -stage serial-line supply chain where stage 1 is the raw material stage and N is the finished goods stage. The total per unit cost of the product is \$100 and the total production time is 100 days. The horizon length is 300 days split into three 100-day phases. The holding cost rate is 35%, and the maximum service time at the finished goods stage is zero.

We consider three profiles for cost and time accrual in the supply chain. We define the cumulative position x of the stage i to be $x = i/N$ for $i=1,2, \dots,N$. The profile $f(x)$ denotes the

cumulative cost, or time, up to and including that stage. For both cost and time, we consider three profiles: $f(x) = x^{.25}$, x , and x^2 . For example, if the cost (time) profile were $f(x) = x^{.25}$, then the cost added (production time) at stage i is $(\frac{i}{N})^{.25} - (\frac{i-1}{N})^{.25}$; for the numerical results that follow, the profile calculations are rescaled to \$100 or 100 days and rounded to two significant digits.

These profiles allow us to capture the breadth of supply chain structures that exist in practice. For example, traditional consumer-electronics-manufacturing supply chains consist of high raw material cost and long component lead-times (where both cost and time accrual profiles follow $f(x) = x^{.25}$) whereas an OEM relying on outsourced manufacturing may see a similar cost profile but a time profile much more like $f(x) = x^2$ due to the fact that its supply chain responsibilities are primarily distribution-based.

We characterize demand by (6) – (8) from Graves and Willems (2006) with safety factor $k = 2.05$, and where each stage's demand in period t has mean per period demand μ_p and standard deviation σ_p for all periods in phase $p = 1, 2, 3$; we assume time prior to phase 1 has the same demand characterization as in phase 1. We set the demand parameters in each phase to be high, medium or low. We define high and low demand as a fixed percentage ε greater, or less, than the medium-demand, where we assume the mean daily demand in the medium phase equals 1000 units/day. Thus, the high demand has a mean per period demand $(1+\varepsilon)1000$ and low demand equals $(1-\varepsilon)1000$. We assume the per-period standard deviation of demand equals the mean per period demand in all instances. A daily coefficient-of-variation of 1 is consistent with the level of variability faced by short life-cycle products.

For the numerical experiments, we define a demand profile as a specification of the demand over all three phases. If each phase can experience high, medium or low demand, then there are 27 possible profiles. However, we restrict our attention to the twelve profiles that do not

have the same demand characterization in two adjacent phases; low-medium-low is an example of one profile.

We define a scenario in terms of the number of stages in the serial line and ϵ . Thus, scenario (5, 0.3) corresponds to a five-stage serial supply chain where the high demand phase has mean per period demand 1300, the medium phase has a mean 1000 and the low phase has mean demand 700. For each scenario, we evaluate and solve 108 safety-stock optimization problems that correspond to the permutations of three cost-accrual profiles, three time-accrual profiles, and twelve demand profiles.

We consider a five-stage serial supply chain and let ϵ vary from 0.01 to 0.99. For each of the 99 scenarios, Figure 1 compares the optimal CST and DST policies.




Figure 1: Comparison of constant and dynamic service time policies for (5, ϵ)

Even at $\epsilon=0.99$, the safety stock cost of CST only exceeds DST by an average of 13%, yet DST employs on average 7.1 different inventory policies. Indeed, for the test problems in the (5, 0.99) scenario, the DST policy differs from the CST policy on 57.2% of the days in the transient time window.

For the 108 test problems in the (5, 0.3) scenario, on average the safety stock cost for CST exceeds DST by 1.6%; the median difference is 1.4% and the maximum is 6.0%. On average, DST adopts 4.56 of the 16 possible decoupling inventory stocking policies during the horizon; the median is four and the maximum is nine.

In order to understand the structure of DST in more detail, we will focus on one permutation in the (5, 0.3) scenario. We consider the test problem where demand is low-high-low, cost is $f(x) = x$ and time is $f(x) = x^{.25}$. In this instance, the five stages each have a direct cost of \$20 and the production times, starting from stage 1 are 36, 28, 20, 12 and 4 days. For this test problem, the cost difference is 5.1%, which is the highest among the six permutations in the (5, 0.3) scenario where DST produce nine safety stock policies.

For this permutation, CST places decoupling safety stocks at stages 2 and 5. Table 1 shows the optimal safety stock policy for each time period. There is a decoupling safety stock at a stage whenever its outbound service time is zero. Thus, from the table, we can see how the structure of the DST policy changes over the course of the transient time window. Indeed, we see that there are ten changes in policy over the 300 day horizon.

Transition Interval	Starting Period	Ending Period	Outbound service time by stage				
			1	2	3	4	5
A	1	100	36	0	20	32	0
B	101	113	0	28	48	60	0
C	114	178	36	64	84	96	0
D	179	201	36	0	20	32	0
E	202	208	0	28	48	0	0
F	209	212	0	28	0	0	0
G	213	213	0	0	0	0	0
H	214	231	0	0	0	12	0
I	232	233	0	0	20	0	0
J	234	262	0	0	20	32	0
K	263	300	36	0	20	32	0

Table 1: Using DST, the inventory policy adopted for (5, 0.3) permutation where demand profile is low-high-low, cost profile is x and time profile is $x^{.25}$

The primary insight drawn from Table 1 is that DST changes the inventory policy in order to cover as much demand variability as possible from a low-demand phase, subject to the cost trade-off of increasing net replenishment times at more expensive downstream stages. In general, when transitioning from a low-demand-variability phase to a high-demand-variability phase, DST reduces the number of stocking locations in order to increase the net replenishment time at the remaining inventory locations so that demand requirements remain partially in the low-variability phase. When transitioning from a high-variability phase to a low-variability phase, DST increases the number of stocking locations in order to have as much variability as possible covered from the low-variability phase.

As with the example from section 4 in Graves and Willems (2006), we conclude from the numerical experiments that the CST policy is near optimal for a broad set of supply chain specifications. Furthermore, we provide some insight into the behavior of the optimal DST policy. Nevertheless, the DST policy remains problematic due to the difficulty in implementing such a policy.

3. Example – Consumer Packaged Goods Manufacturer

In this section, we illustrate the impact of non-stationary demand and validate the modeling approach presented in this paper by presenting a real-world example from a consumer packaged goods company¹. Figure 2 graphically represents the supply chain for a product family:

¹ The data in this section has been disguised to protect company-confidential data. The insights drawn from the disguised data are identical to the conclusions based on the actual data.

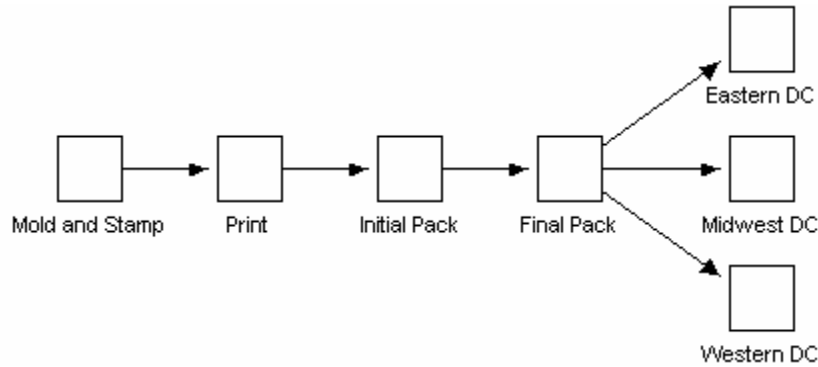


Figure 2: Supply chain map for consumer packaged goods company

The product is created through an injection molding and stamping process. Logos are then printed onto the plastic part. After the product goes through two packing operations, it is sent to one of three regional distribution centers where it serves the region’s demand. The (disguised) costs and production lead-times, in days, are:

Stage Name	Production Lead-Time	Cost Added
Mold and Stamp	15	\$0.85
Print	3	\$0.60
Initial Pack	3	\$0.10
Final Pack	3	\$0.25
Eastern DC	25	\$0.05
Midwest DC	20	\$0.05
Western DC	15	\$0.05

Table 2: Time and cost parameters for supply chain

The actual demand data for the three DCs are:

	Phase One		Phase Two		Phase Three	
	μ	σ	μ	σ	μ	σ
Eastern DC	1068.5	756.0	1402.0	918.8	740.6	498.5
Midwest DC	670.5	411.3	1035.0	494.3	663.2	292.2
Western DC	322.0	257.0	577.5	379.4	272.0	208.2

Table 3: Daily demand parameters by phase

Each phase covers four months (120 days) of the product’s life cycle. The data in Table 3 represents the daily averages and standard deviations for each phase. We model the demand bound over the twelve-month horizon in the same way as was done for the two-stage example;

that is, we use equations (6) – (8) from Graves and Willems (2006) to model the demand bounds, with safety factor $k = 1.645$.

The CST policy is given in Table 4:

Stage Name	Optimal Service Times (days)	Phase One		Phase Two		Phase Three	
		Safety stock (units)	Safety stock (cost)	Safety stock (units)	Safety stock (cost)	Safety stock (units)	Safety stock (cost)
Mold and Stamp	0	5722	\$567	7072	\$701	3913	\$388
Print	3	0	\$0	0	\$0	0	\$0
Initial Pack	6	0	\$0	0	\$0	0	\$0
Final Pack	9	0	\$0	0	\$0	0	\$0
Eastern DC	0	7251	\$1,565	8812	\$1,902	4781	\$1,032
Midwest DC	0	3643	\$786	4378	\$945	2588	\$559
Western DC	0	2071	\$447	3057	\$660	1678	\$362

Table 4: Optimal inventory results for constant-service-time policy

The CST policy holds a decoupling inventory at the Mold and Stamp stage and at the DCs. Using the company’s 35% annual holding cost rate, the total annual safety stock cost is \$9,915. Table 5 shows the optimal DST policy.

Starting Period	Ending Period	Mold and Stamp	Print	Initial Pack	Final Pack	Eastern DC	Midwest DC	Western DC
1	46	15	18	21	24	0	0	0
47	123	0	3	6	9	0	0	0
124	160	15	3	6	9	0	0	0
161	248	0	3	6	9	0	0	0
249	265	0	3	6	0	0	0	0
266	360	0	3	6	9	0	0	0

Table 5: Optimal outbound service times by period for dynamic-service-time policy

At the start of both the first and second phases, the DST policy only holds inventory at the DCs. This keeps the net replenishment time as long as possible, keeping the demand bound in the adjacent lower-demand phase. In phase three, near the start of the phase it is optimal to add a decoupling stock at Final Pack in order to lower the net replenishment times at the DCs and force as much demand as possible into phase three. Nevertheless, for all phases, the optimal policy reverts back to the optimal CST policy, as the influence from the prior phase diminishes.

The safety stock cost of the DST policy is \$9,629. Looking at only safety stock cost, the CST policy is 2.9% more expensive than the DST policy. If we factor in the pipeline stock cost

of \$46,162 that is common to both policies, then the inventory cost differs by 0.5%. For another example, we cite Coughlin (1998), who describes an application to a high-tech supply chain. He also partitioned the demand for the high-tech products into several stationary phases for the purposes of determining the placement of strategic inventory across a supply chain.

4. Deterministic Planning Problem for Adjusting the Base Stock Levels

The purpose of this section is to document the planning problem that is implied by the adoption of the periodic-review base-stock replenishment policy with time-varying base-stock levels.

From Graves and Willems (2006), we found in equation (2) that at the end of each period t , the inventory at stage j should be exactly:

$$I_j(t) = B_j(t) - d_j(t - SI_j - T_j, t - S_j) \quad (A1)$$

where the base stock level is given by $B_j(t) = D_j(t - SI_j - T_j, t - S_j)$.

We wish to determine what the production needs to be at each stage so as to accomplish (A1). We define $P_j(t)$ to be the production completed at stage j in period t , and $R_j(t)$ to be the production starts (releases) into stage j at time t . Thus, by the definition of the lead time, we have $R_j(t) = P_j(t + T_j)$.

Demand Node j : Suppose stage j is a demand stage that serves end-item demand. The inventory balance equation for stage j is:

$$I_j(t) = I_j(t-1) + P_j(t) - d_j(t - S_j)$$

We can use (A1) to substitute for $I_j(t)$ to get a recipe for finding $P_j(t)$:

$$\begin{aligned} P_j(t) &= I_j(t) - I_j(t-1) + d_j(t - S_j) \\ &= B_j(t) - d_j(t - SI_j - T_j, t - S_j) - B_j(t-1) + d_j(t-1 - SI_j - T_j, t-1 - S_j) + d_j(t - S_j) \\ &= \Delta B_j(t) + d_j(t - SI_j - T_j) \end{aligned} \quad (A2)$$

where $\Delta B_j(t) = B_j(t) - B_j(t-1)$ is the first difference. Then from (A2) we have that:

$$R_j(t) = P_j(t+T_j) = \Delta B_j(t+T_j) + d_j(t - SI_j)$$

Thus we release into production in time period t the demand observed at time $t-SI_j$, plus the anticipated change to base-stock levels in period $t+T_j$. But these base-stock levels are known a priori as they only depend on the known demand bounds.

Internal Stage j : For internal stage j , the inventory balance equation for stage j is:

$$I_j(t) = I_j(t-1) + P_j(t) - \sum_{i:(j,i) \in A} \phi_{ji} P_i(t+T_i)$$

As explanation, the demand experienced at an internal stage j is the production starts for its customers. The production starts at stage i in period t equal the production completions at time $t+T_i$, by the definition of T_i .

We can use (A1) to substitute for $I_j(t)$ to get a recipe for finding $P_j(t)$:

$$\begin{aligned} P_j(t) &= I_j(t) - I_j(t-1) + \sum_{i:(j,i) \in A} \phi_{ji} P_i(t+T_i) \\ &= B_j(t) - d_j(t - SI_j - T_j, t - S_j) - B_j(t-1) + d_j(t-1 - SI_j - T_j, t-1 - S_j) + \sum_{i:(j,i) \in A} \phi_{ji} P_i(t+T_i) \quad (A3) \\ &= \Delta B_j(t) - d_j(t - S_j) + d_j(t - SI_j - T_j) + \sum_{i:(j,i) \in A} \phi_{ji} P_i(t+T_i) \end{aligned}$$

We can simplify this expression, given that there are no cycles in the supply chain. In particular, we will show that for all stages:

$$P_j(t) = \Delta E_j(t) + d_j(t - SI_j - T_j) \quad (A4)$$

where

$$\Delta E_j(t) = \Delta B_j(t) + \sum_{i:(j,i) \in A} \phi_{ji} \Delta E_i(t+T_i). \quad (A5)$$

Proof of (A4): We will show this by induction over the stages. Assume the stages are ordered from 1 to N such that if $(i,j) \in A$ then $i > j$. Thus, stage 1 is an end item (a demand node).

Consider stage 1. Since end items have no descendants, from (A5) we have that $\Delta E_j(t) = \Delta B_j(t)$. Thus we see that (A4) is true for stage 1, as it is equivalent to (A2).

Indeed, we see that (A4) is true for all end items.

Suppose (A4) is true for each stage 1,2, ... j-1. We now show it to be true for an internal stage j. From (A3), we have:

$$P_j(t) = \Delta B_j(t) - d_j(t - S_j) + d_j(t - SI_j - T_j) + \sum_{i:(j,i) \in A} \phi_{ji} P_i(t + T_i).$$

We can now apply the induction hypothesis to re-write (A3) as:

$$\begin{aligned} P_j(t) &= \Delta B_j(t) - d_j(t - S_j) + d_j(t - SI_j - T_j) + \sum_{i:(j,i) \in A} \phi_{ji} (\Delta E_i(t + T_i) + d_i(t - SI_i)) \\ &= \left(\Delta B_j(t) + \sum_{i:(j,i) \in A} \phi_{ji} \Delta E_i(t + T_i) \right) - \left(d_j(t - S_j) - \sum_{i:(j,i) \in A} \phi_{ji} d_i(t - SI_i) \right) + d_j(t - SI_j - T_j) \end{aligned} \quad (A6)$$

From (A5), we note that the first bracketed term in (A6) is $\Delta E_j(t)$. The second bracketed term in (A6) is zero, since for an internal stage j, we have that $d_j(t) = \sum_{i:(j,i) \in A} \phi_{ji} d_i(t)$ and we assume that for all $(j,i) \in A$, we have $S_j = SI_i$. If instead we have $S_j < SI_i$, then we need to modify the definition of stage j demand to be: $d_j(t) = \sum_{i:(j,i) \in A} \phi_{ji} d_i(t + S_j - SI_i)$. This modification reflects the fact that demand from node i placed on stage j will be delayed by $SI_i - S_j$ periods.

Thus, we can simplify (A6) to get the desired result, namely (A4):

$$P_j(t) = \Delta E_j(t) + d_j(t - SI_j - T_j)$$

This completes the proof.

We find the production starts: $R_j(t) = P_j(t + T_j) = \Delta E_j(t + T_j) + d_j(t - SI_j)$. Thus, the productions starts correspond to observed demand, plus an adjustment term. The interpretation of

the adjustment term is that it is the change in the echelon base stock for stage j , accounting for lead times to flow the changes through the supply chain. We can calculate this adjustment term from the non-stationary base-stock levels, using (A5).

In this development of the planning process to achieve the non-stationary base-stock levels, we raise a couple of caveats.

We have ignored non-negativity constraints on production. The above derivation finds the desired production without consideration of non-negativity constraints on P . When the prescribed P is negative, one would presumably set $P = 0$ and the inventory at the stage will be more than required. Subsequent production starts would be reduced to net out the excess inventory.

We have ignored any capacity constraints on P . One might formulate this planning problem as a linear optimization, where the prescribed production levels (A4) could be used to determine lower bounds on production in order to provide 100% service within the service times. One could solve a linear optimization to find the best non-negative production plan that satisfies the service requirements and does so within the capacity constraints with the least amount of inventory. To guarantee feasibility, one might need to set up the linear optimization to permit violation of the capacity constraints and/or of the service requirements.

Finally, we note that to determine what to start in period t , we need to know the base stocks (i.e., the demand bounds) from period t to $t+M$ where M is the lead time of the longest path through the supply chain. We see this from (A4) and (A5), in which the production at a stage depends on the time-phased future adjustments to all of its downstream stages.

5. Conclusion

In this online appendix, we show that the constant-service-time policy need not be optimal; rather, an optimal policy might entail changing the service times, and thus the locations of safety stocks, as demand evolves over time. Nevertheless, we argue and provide limited evidence that applying the constant-service-time policy is near optimal and has obvious implementation advantages. We illustrate the model with an application to a consumer-packaged goods company. We have also developed and solved the planning problem for adjusting the base stock levels, as would be required when demand is non-stationary.

References

Coughlin, R. L., "Optimization and Measurement of a World-Wide Supply Chain," S.M. Thesis, Leaders for Manufacturing Program, MIT, Cambridge, MA., June 1998.

Graves, S. C., and S. P. Willems, "Optimizing Strategic Safety Stock Placement in Supply Chains," *Manufacturing & Service Operations Management*, 2 (2000), 68-83.

Graves, S. C., and S. P. Willems, "Strategic Inventory Placement in Supply Chains: Non-Stationary Demand," working paper, March 2000, revised February 2006, available from <http://web.mit.edu/sgraves/www/papers/>