Avoiding the capacity cost trap: Three means of smoothing under cyclical production planning

Carl Philip T. Hedenstierna*

Department of Industrial Economics and Technology Management, NTNU – Norwegian University of Science and Technology, 7491 Trondheim, Norway.

Stephen M. Disney

Logistics System Dynamics Group, Cardiff Business School, Cardiff University, Aberconway Building, Colum Drive, CF10 3EU, United Kingdom.

Abstract

Companies tend to set their master production schedule weekly, even when producing and shipping on a daily basis—the term for this is *staggered deliveries*. This practice is common even when there is no marginal cost of setting a new schedule. This paper argues that the practice is sound for companies that use the ubiquitous order-up-to (OUT) policy to control production of products with a significant capacity cost. Under these conditions, the length of the order cycle (time between schedule updates) has a damping effect on production, while a unit (daily) order cycle can cause significant capacity costs. We call this the *capacity cost trap*.

Developing an analytical model based on industrial evidence, we derive capacity and inventory costs under the staggered OUT policy, showing for this policy there is an optimal order cycle possibly greater than unity. To improve on this solution, we consider three approaches to smoothing: either levelling within the cycle, deferring excess production or idling to future cycles via a proportional OUT policy, or increasing the length of the cycle. By deriving exact cost expressions we compare these approaches, finding that smoothing by employing the proportional OUT policy is sufficient to avoid the capacity cost trap.

Keywords: Inventory, Order-up-to policy, Reorder period, Overtime cost

*Corresponding author
Email addresses: `carl.p.hedenstierna@ntnu.no` (Carl Philip T. Hedenstierna), `DisneySM@cardiff.ac.uk` (Stephen M. Disney)

1. Introduction

In many global supply chains production planning and shipping cycles are synchronised. That is, production planning cycles are generated once a week and containers of product are shipped to the customer once a week via a scheduled liner service. Most of our stylized production and inventory control studies make this assumption. However, in many local supply chains, especially those operating in lean or just-in-time mode, production plans are generated weekly, but trucks leave for the customer at the end of each day, loaded with the days production[1]. It then follows that an inventory deviation occurring *just after* a plan has been set will not be corrected as quickly as one occurring *just before* the plan is set. When each plan contains a sequence of deliveries we are said to use *staggered deliveries* [14], a setting rarely studied in the literature. We show how staggered deliveries complicate production smoothing, as capacity and inventory costs are related to the length of the order cycle, overtime work or idling can be deferred to future cycles, and the overtime or idling within the present cycle must be determined.

Present-day order cycle lengths tend to be one week, occasionally one month (see Table 1), but the conundrum of selecting an order cycle length can be traced back to 1924, when Alfred P. Sloan decided that General Motors (GM) Corporation's production plans should be reviewed every ten days instead of once every three months as was done before. During the 1920's GM managed to increase its total inventory turnover from 2 to nearly $7^1/_2$ times per annum by streamlining its production and distribution network [29]. Toyota Motor Corporation also considered fast reordering as important [23, p. 51] making a detailed forecast for a month at a time, but using an order cycle of ten days, while considering moving to weekly or even daily cycles [27, p. 129]. In the same vein, Burbidge [3] declared short order cycles as one of his "five golden rules to avoid bankruptcy".

1.1. Literature review

The theory of staggered deliveries and the selection of an appropriate order cycle length is only partially resolved by the literature. Flynn and Garstka [14] identified that a *staggered, traditional order-up-to* (STOUT) policy is optimal under piecewise-linear inventory costs and a once-per-cycle audit cost. Here demand was assumed to be independent and identically distributed (i.i.d.). Building on these results, Flynn and Garstka [15] proved the existence of an optimal order cycle length, and outlined a procedure for identifying it. The model was extended to a multi-product scenario in [12], while [11] showed a staggered (S,s) policy to be optimal when a fixed cost for non-zero orders was

---

[1]In a company setting that motivated this study, the production planner also occasionally received urgent requests from customers during the week, causing an in-week reschedule. As this effect is rare, and near impossible to model analytically, we defer this aspect to future work.

Company	Country	Industry	Order cycle	When
Tesco[a]	UK	Grocery	8h or 24h	2005
Lexmark[b]	USA	Printer	Weekly	2013
Harman Kardon	UK	Audio equipment	Weekly	2001
P&G	Global	Household goods	Weekly	2014
Princes	UK	Fruit Juice	Weekly	2003
TRW	Global	Automotive	Weekly	1999
BAT[c]	Global	Consumer goods	Monthly	2012
Hotai Motor Co[d]	Taiwan	Automotive	Monthly	2001
Renishaw	UK, India	Measuring equipment	Monthly	2014

[a] Reported in Potter and Disney [25] ;   [b] Reported in Disney et al. [10] ;
[c] Reported in Hedenstierna [17] ;   [d] Reported in Chiang [4].

Table 1: Examples of industrial order cycles (Source: Authors)

added. Another extension investigated a heuristic for finding the optimal order cycle length [13].

The inventory-optimal policy for autocorrelated demand was derived in Hedenstierna and Disney [16], who also identified the optimal order cycle length, and demonstrated that optimal ordering causes the fill rate to fluctuate over the cycle, while the availability remains constant. A related problem was investigated by Prak et al. [26] where inventory inspections and deliveries occurred continuously.

The staggered delivery problem was approached from a different angle by Chiang [4], who demonstrated that staggering deliveries via lot splitting reduces cycle stock in comparison to non-staggered models. Bradley and Conway [2] also considered the effect of cyclic scheduling on inventories, finding that lot-splitting can destroy value when set-up times are significant.

The link between staggering and stable production rates was studied by Chiang [5], where overtime work was allocated to the beginning of the cycle. Here orders were allowed to vary in the first few periods of each cycle, but were bounded by some upper value in the remaining periods of the cycle. Modigliani and Hohn [21] investigated how economic production plans can be made when demand is known in advance (i.e. in a make-to-order setting). They found that an order cycle should be no longer than one seasonal cycle, and shorter in cases with high inventory costs.

The problem of total inventory and overtime costs in a deterministic setting was studied by Kaku and Krajewski [20], who explored this relationship for Period Batch Control systems. They study how set-up times affect overtime costs, finding that short order cycles could lead to unnecessary costs. Nathan

and Venkataraman [22] investigated a specific industrial case with overtime costs. Using mixed integer goal programming, they found an optimum order cycle length of two months.

To recap: A handful of industrialists have called for short order cycles, albeit rarely of less than one week. Indeed, most of the companies in Table 1 have an order cycle of one week, but Tesco, a retailer with low capacity costs, manages to reschedule up to three times per day. The literature supports order cycles greater than unity when a fixed cost per order cycle is involved, but there is no knowledge of the impact of the order cycle length on capacity costs.

Systems with staggered deliveries are reasonably well understood from the perspective of inventory costs, under which the STOUT policy is optimal (Flynn and Garstka [14]; Prak et al. [26]; Hedenstierna and Disney [16]). Here we expand the analysis to include capacity costs: First we investigate how the inventory-optimal STOUT policy performs when both piecewise-linear capacity costs and inventory costs are present. Then we investigate the performance of a more general policy, termed the *staggered proportional order-up-to* (SPOUT), which does not correct the entire inventory deficit or surplus in the current order cycle, but only a fraction thereof. The SPOUT policy is an extension of the proportional order-up-to-policy implemented at Lexmark, [10]. It is also viable to distribute inventory deviations evenly over all orders in a cycle, as was done by Chiang [4]. We call these variations STOUT-E and SPOUT-E, the additional *E* denoting equal over-time. Figure 1 characterizes these four strategies. Table 2 summarises the staggered delivery literature, highlighting its relation to our four staggered strategies.
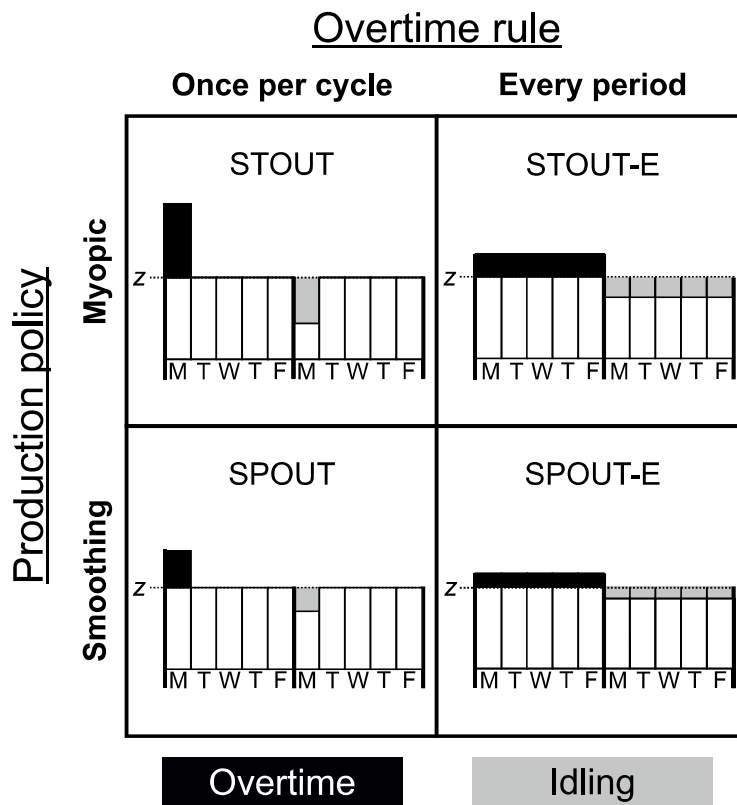
Figure 1: Categorization of the policies in this paper. Note: z is the nominal capacity available without idling or overtime.

Ref.	Demand	Products	Costs	Policy †	Unique contribution or findings
[2]	Constant	Multiple	Time-based	STOUT-E	More changeovers reduces inventory buildup under cyclical plans
[4]	Normal, i.i.d.	Single	Inventory, ordering	STOUT-E	It may be beneficial to split an order quantity into multiple lots
[5]	Poisson, i.i.d.	Single	Inventory	STOUT	A simplified STOUT policy under capacity constraints
[11]	Normal i.i.d.	Single	Inventory, audit	STOUT	A global search heuristic for optimal cycle length
[12]	Normal i.i.d.	Multiple	Inventory, audit	STOUT	Proves optimal policy exists, develops heuristic policies
[13]	Normal, i.i.d.	Multiple	Inventory, audit	STOUT	A heuristic for finding the optimal cycle length
[14]	Stochastic, i.i.d.	Single	Inventory, audit	STOUT	Shows STOUT to be optimal under pure inventory and audit costs
[15]	Normal i.i.d.	Single	Inventory, audit	STOUT	Proves optimal cycle length exists
[16]	Normal, Autocorrelated	Single	Inventory, audit	STOUT	Integrates STOUT with a forecast for autocorrelated demand
[20]	Normal, i.i.d.	Multiple	Inventory, overtime	STOUT-E	Shows workcell design to affect the optimal cycle length
[21]	Deterministic	Single	Inventory, capacity	Math. Prog.	Derives the optimal production quantity sequentially
[22]	Deterministic	Multiple	Inventory, overtime, setup	STOUT	Forecast window intervals
[26]	Normal, Gamma, i.i.d.	Single	Inventory	STOUT	Shows STOUT to be optimal when production is continuous
This paper	Normal, i.i.d.	Single	Inventory, capacity	STOUT[-E], SPOUT[-E]	Shows superiority of SPOUT over STOUT when facing capacity costs

Table 2: Summary of the staggered delivery literature (†: Actual, or closely related, policy)

1.2. Contribution

The main contribution of this paper is the consideration of capacity costs in staggered delivery systems, the identification of the order cycle length as a smoothing mechanism, and the evaluation of three distinct types of production smoothing: deferring overtime work to a future order cycle, levelling overtime work or idling within the cycle, and extending the length of the order cycle. The last two means of smoothing are exclusively available to systems with staggered deliveries and are not studied in the literature. A major result is the realization that companies using the STOUT policy will see increased capacity costs as the order cycle is shortened, but that it is sufficient to implement the SPOUT policy to avoid this cost increase.

Four propositions highlight our findings: Proposition 1 finds that under i.i.d. demand (of arbitrary distribution), it is optimal to absorb all the demand fluctuation by over-time or idling in the first period of the order cycle. Propositions 2 and 3 use an inverse function approach to identify the optimal order cycle length under the STOUT and STOUT-E policy. Proposition 4 shows that the optimal cycle length under the SPOUT policy is unity and this is always the lowest cost solution (arbitrary distribution). Propositions 2–4 hold for normally distributed demand.

1.3. Paper structure

We proceed with a description of the basic mechanism of our staggered delivery inventory setting, including a discussion on safety stocks and inventory costs, as well as on capacity levels and capacity costs. Section 3 defines the inventory optimal STOUT policy and its optimal order cycle length when capacity costs are incorporated into the objective function. We also consider the allocation of overtime and idling within the order cycle. Section 4 considers a more general replenishment policy, the SPOUT policy, capable of smoothing production across cycles. Section 5 provides numerical examples and illustrates all of our theoretical contributions which fully characterize our considered scenarios. Section 6 provides managerial insights, section 7 concludes. The proofs to all Lemmas, Propositions, and Corollaries are housed in the Appendix A. While we have been careful to define notation on first use, Appendix B provides a list of notation for convenience.

2. Model development

Consider a single-product inventory system operating in discrete time, where the integer t indexes time, counted in days. In each day there is a demand of d_t units. We assume $d_t = \mu + \varepsilon_t$, where μ is the mean demand and ε_t is a zero mean i.i.d. random variable with a variance of σ_d^2, drawn from an arbitrary distribution unless otherwise stated. Negative demand represents returns from customers.

At the integer moments of time t, the sequence of events is as follows. To the last period's ending inventory i_{t-1}, receive r_t from production, then satisfy demand d_t, thereafter tally the inventory i_t.

$$i_t = i_{t-1} + r_t - d_t. \tag{1}$$

Here, r_t and d_t reflect the completed production and consumption, respectively, in the time-span between the successive observations i_{t-1} and i_t. Note that r_t are the production orders released to the shop floor $L + 1$ periods ago where L is the physical production lead-time. For a simpler analytical treatment, we use the double subscript notation, $i_{t,k}$, to denote a set of P inventory levels, indexed by $k \in \{1, 2, 3, ..., P\}$, determined periodically when $t/P \in \mathbb{Z}$, e.g. $i_{t,k} = \{i_{t+k+L}|t/P \in \mathbb{Z}\}$. To avoid clutter we suppress the conditional statement $\{\cdot|t/P \in \mathbb{Z}\}$, whenever we use the double subscript on a state variable.

We assume that a sequence of P orders are placed once every P periods. Specifically, order quantities are decided when t/P is an integer, i.e. in the periods $t = \{0, P, 2P, ...\}$. Immediately after observing i_t, the P order quantities, $o_{t,k}$, are decided, with k giving the release sequence of the orders. The first order planned at time t is immediately released to production, the second order is released in the next period, $t + 1$, the k^{th} order is released at time $t + k - 1$, and the final order at time $t + P - 1$. The orders are registered as received in inventory $L + 1$ periods after they were released, i.e. in period $t + k + L$, i.e. the order quantities and their corresponding receipts $o_{t,k} = r_{t+k+L}$ are first recorded as received in i_{t+k+L}. The smallest possible physical lead time, $L = 0$, results in the first order of a cycle being determined at t, released for production immediately, and first included in the inventory tally i_{t+1}. While L represents the *physical* lead time, $k + L$ is the *effective* lead time for each order. Figure 2 portrays the order generation, production releases, and inventory receipts in our staggered delivery system.

Returning to the inventory balance equation (1), we use induction to express the inventory level in an arbitrary period as

$$i_{t+k+L} = i_t + \sum_{n=1}^{k+L} (r_{t+n} - d_{t+n}), \qquad \text{when } t/P \in \mathbb{Z}. \tag{2}$$

This formulation is helpful when planning $o_{t,k} = r_{t+k+L}$, as the inventory i_t and all preceding orders $\{..., r_{t+k+L-1}\}$ are known. Define

$$x_{t,k} = \{i_t + \sum_{n=1}^{k+L} r_{t+n}|t/P \in \mathbb{Z}\}, \tag{3}$$

as the inventory position *after* placing $o_{t,k}$. As $i_{t,k} = i_{t+k+L}$, it is possible to express the inventory level as

$$i_{t,k} = x_{t,k} - D_{t,k}, \tag{4}$$

where $D_{t,k} = \{\sum_{n=1}^{k+L} d_{t+n}|t/P \in \mathbb{Z}\}$ is the effective-lead-time demand, and $x_{t,k}$ is a decision variable as it includes $o_{t,k}$. We have left the precise mechanism
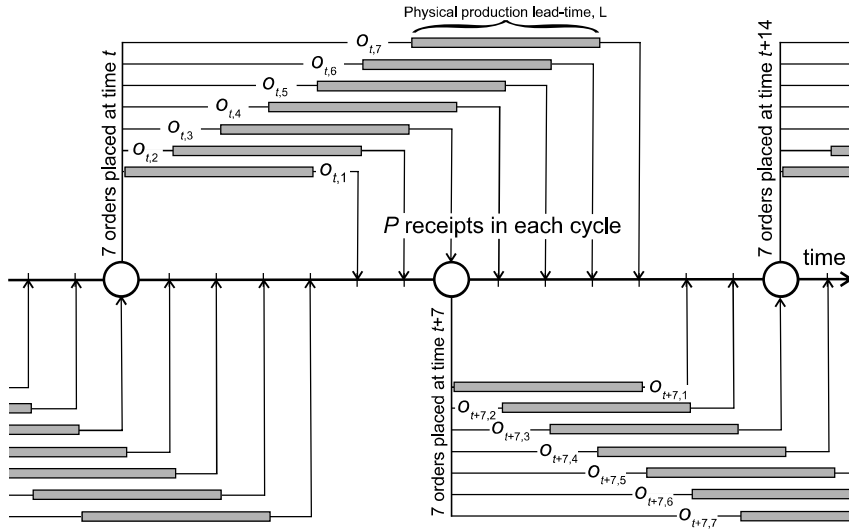
8

Figure 2: Sequence of events in a staggered delivery system ($P = 7, L = 4$).

for generating the orders, $o_{t,k}$, unspecified, to be treated in subsequent sections. Next, we shall introduce a cost model that acts on the production–inventory system, along with optimal inventory and capacity levels.

2.1. Inventory costs

Consider linear inventory holding costs, h, and backlog costs, b, following

$$j(i_{t,k}) = h(i_{t,k})^+ + b(-i_{t,k})^+. \tag{5}$$

where $(x)^+ = \max(x, 0)$, and the inventory level $i_{t,k}$ is a random variable defined on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$. We shall see that for arbitrary distributions of the inventory level, resulting from an optimal or a suboptimal policy, there exists an inventory position, and a corresponding safety stock setting, that minimizes the inventory costs as far as the policy in question permits. We refer to this as the optimal inventory position, $x_k^* = \mathbb{E}[x_{t,k}]$, and the optimal safety stock as $i_k^* = \mathbb{E}[i_{t,k}]$.

Lemma 1 (Patterned on Churchman et al., 1957, p. 212). *For the arbitrarily distributed random variable $i_{t,k}$, the inventory costs are minimized when*

$$i_k^* = F_{i,k}^{-1}\left(\frac{h}{b+h}\right), \tag{6}$$

where $F_{i,k}^{-1}$ is the inverse cumulative density function (CDF) of the inventory level in period $t + k + L$.

9

Remark Lemma 1 restates the standard newsvendor model in a manner that explicitly captures the safety stock and inventory costs in each period of the cycle. This is required as we later show that the inventory levels are heteroskedastic in staggered delivery systems.

The optimal safety stock, i_k^*, implies an optimal setting for $x_{t,k}$, namely

$$x_k^* = i_k^* + \mathbb{E}\left[D_{t,k}\right] = i_k^* + \mu(k + L). \tag{7}$$

This criterion applies to arbitrary policies. Although x_k^* is known as the base stock, or order-up-to level, it is applicable to any policy where the costs follow (5). A practical interpretation of the optimality criterion is that the long-run average of $x_{t,k}$, $x_{t+P,k}$, $x_{t+2P,k}$, $\cdots$ should equal x_k^*, regardless of the policy selected. When $x_{t,k} = x_k^*$ and the demand is normally distributed, (A.1) simplifies to

$$\mathbb{E}\left[j(i_{t,k})\right] = \sigma_{i,k}\left(b + h\right)\varphi\left[\Phi^{-1}\left(\frac{b}{b+h}\right)\right], \tag{8}$$

where $\varphi(\cdot)$ is the probability density function (PDF) of the standard normal distribution and $\Phi^{-1}(\cdot)$ is the inverse CDF of the standard normal distribution. Averaged over a cycle, the expected inventory cost per period is

$$J_P = \frac{1}{P}\sum_{k=1}^{P}\mathbb{E}\left[j(i_{t,k})\right] = (b+h)\bar{\sigma}_{i,P}\varphi\left[\Phi^{-1}\left(\frac{b}{b+h}\right)\right], \tag{9}$$

where $\bar{\sigma}_{i,P} = P^{-1}\sum_{k=1}^{P}\sigma_{i,k}$ is the average standard deviation of the inventory level. Observe that the average standard deviation does not equal the square root of the average inventory variance due to Jensen's inequality. For the same reason, the standard deviation of the inventory level (sampled over all k) differs from $\bar{\sigma}_{i,P}$.

2.2. Capacity costs

Workers are guaranteed compensation for a daily output of up to z_k products at the normal rate of u dollars per product; when the production quantity is greater than z_k, the excess is paid for at an overtime rate of v dollars per product. Assuming that it is incurred in the same period as goods are received $(t + k + L)$, the capacity cost can be expressed as

$$a(o_{t,k}) = uz_k + v(o_{t,k} - z_k)^{+}, \tag{10}$$

where the order quantity $o_{t,k}$ is an arbitrarily distributed random variable defined on the probability space $(\Omega, \mathcal{A}, \mathbb{P})$. If $o_{t,k} < z_k$ the workers are still paid for their guaranteed hours (uz_k). Thus z_k is the available labour capacity in the nominal working day when $o_{t,k}$ enters production. While we use the term *capacity limit* to denote z_k, this is only the limit of capacity before overtime is used to produce the remaining quantity $(o_{t,k} - z_k)^{+}$. In effect, we assume that there is no limit on the amount of over time available; either there is sufficient overtime to meet peak demands, or peak demands can be directed to a subcontractor who can process items with same lead time and quality at a unit cost of w.

Lemma 2 (Patterned on Hosoda and Disney, 2012). *For the arbitrarily distributed random variable $o_{t,k}$, the expected capacity costs, (10), are minimized when*

$$z_k^* = F_{o,k}^{-1}\left(\frac{v-u}{v}\right),\tag{11}$$

where $F_{o,k}^{-1}(x)$ is the CDF of $o_{t,k}$, i.e. the CDF of the k'th order placed in a cycle.

Remark We later show that the orders in staggered delivery systems are heteroskedastic, hence we have specified in Lemma 2 the capacity requirements and capacity costs in each period of the cycle.

While Lemma 2 remains true regardless of the distribution the demand, under normally distributed demand the optimal capacity level is $z_k^* = \sigma_{o,k}\Phi^{-1}\left[(v-u)/v\right] + x_k^* - x_{k-1}^*$, where $\sigma_{o,k}$ is the variance of $o_{t,k}$, i.e. the variance of the orders on the k^{th} period of the cycle [18]. In addition, Hosoda and Disney [18] provide an expression for expected capacity cost that can be readily adapted to our staggered setting when z_k^* is used under normal demand,

$$\mathbb{E}\left[a(o_{t,k})\right] = v\sigma_{o,k}\varphi\left[\Phi^{-1}\left(\frac{v-u}{v}\right)\right] + u\left(x_k^* - x_{k-1}^*\right).\tag{12}$$

The average capacity cost per period is

$$A_P = \frac{1}{P}\sum_{k=1}^{P}\mathbb{E}\left[a(o_{t,k})\right] = v\bar{\sigma}_{o,P}\varphi\left[\Phi^{-1}\left(\frac{v-u}{v}\right)\right] + u\mu,\tag{13}$$

where $\bar{\sigma}_{o,P} = P^{-1}\sum_{k=1}^{P}\sigma_{o,k}$ is the average standard deviation of the orders. The total average cost per period is then $C_P = J_P + A_P$, including both inventory costs and capacity costs.

3. The staggered traditional order-up-to policy (STOUT)

The order-up-to (OUT) policy is a popular replenishment policy for production planning and inventory control in practice as it is available native in many ERP/MRP systems. The OUT policy is the optimal policy for controlling inventory related costs in non-staggered settings. The staggered equivalent of the OUT policy is also optimal when only inventory costs are present [16]. Under i.i.d. demand, it places all stochastic corrections (overtime or idling) in the first period of the cycle, with the remaining periods having deterministic order quantities. The STOUT order quantity is

$$o_{t,k} = x_k^* - x_{t,k-1}.\tag{14}$$

Note, $x_{t,0} = i_t + \sum_{n=1}^{L}r_{t+n}$. When orders follow (14), the inventory position is immediately raised to the optimal OUT level, x_k^*, following $x_{t,k} = x_{t,k-1} + o_{t,k} = x_k^*$. Inserting $x_{t,k}$ into (4) provides

$$i_{t+k+L} = x_k^* - D_{t,k},\tag{15}$$

11

revealing that the inventory level is the difference between the constant term x_k^* and the random (not necessarily normal) variable $D_{t,k}$. With this, we are able to prove:

Lemma 3. *For the STOUT policy,*

(a) the inventory variance is

$$\sigma_{i,k}^2 = \sigma_d^2 \left(k + L\right);\tag{16}$$

(b) the order variance is

$$\sigma_{o,k}^2 = \begin{cases} \sigma_d^2 P & when\ k = 1, \\ 0 & otherwise. \end{cases}\tag{17}$$

Remark The variance expressions in (16) and (17) (and all the other variance expressions in this paper) hold when the i.i.d. demand is drawn from *any* distribution; we only require normally distributed demand when we are considering economic performance.

We notice that the inventory cost optimal policy, STOUT, accounts for all of the stochastic components of the cycle demand in the first order of the cycle. A natural question now arises; what are the consequences of absorbing this stochastic component in other periods of the cycle? This is explored in the following proposition.

Proposition 1. *When capacity and safety stock levels are optimal, and the planned overtime or idling brings the inventory position closer to its target, i.e. $x_0^* - x_{t,0} \geq \sum_{n=1}^{P} o_{t,n} - \mathbb{E}\left[o_{t,n}\right] \geq 0$, the total cost is minimized when all overtime is allocated to the first period of each cycle.*

Although the intuition behind this proposition is simply that an hour of overtime costs the same on Mondays as on any other weekday, it also captures the changes in the optimal safety stock and capacity levels as overtime is reallocated via q_k.

Having fully specified the STOUT policy, we observe that the inventory variance increases with k. We also note the following Corollary,

Corollary 1. *(a) The average standard deviation of the orders, $\bar{\sigma}_{o,P} = \sigma_d \sqrt{P^{-1}}$, is decreasing in P. (b) The average standard deviation of the inventory, $\bar{\sigma}_{i,P} = \sigma_d P^{-1} \sum_{k=1}^{P} \sqrt{k + L}$, is increasing in P.*

Corollary 1 shows that the order cycle contains a crude mechanism for production smoothing, and implies that the optimal reorder period may be greater than unity. This brings into question the belief in ever-shortening order cycles, and motivates the search for the P that minimizes the costs generated by the STOUT policy.

The P^* optimization problem is non-convex, but can be solved with an inverse-function approach. First, we write the total cost as

$$C_P\left(\lambda\right) = \psi\left[\bar{\sigma}_{i,P} + \lambda\left(\bar{\sigma}_{o,P} - \bar{\sigma}_{i,P}\right)\right] + \mu u, \tag{18}$$

where ψ is a scaling factor,

$$\psi = v\varphi\left[\Phi^{-1}\left(\frac{v-u}{v}\right)\right] + (b+h)\,\varphi\left[\Phi^{-1}\left(\frac{b}{b+h}\right)\right], \tag{19}$$

and

$$\lambda = \frac{v\varphi\left\{\Phi^{-1}\left[(v-u)/v\right]\right\}}{\psi}, \tag{20}$$

provides the balance between inventory costs and overtime costs. The setting $\lambda = 1$ represents capacity costs only ($b = 0$, or $h = 0$, or $b + h = 0$) and $\lambda = 0$ indicates an absence of overtime costs ($v = 0$, or $u = v$). The important result from (18) is that the total cost is a linear function of λ for any fixed P.

Proposition 2. *The order cycle length P minimizes the total cost $C_P\left(\lambda^*\right)$ for $\lambda^* \in [\lambda_{P-1}, \lambda_P]$, where $\lambda_0 = 0$, and*

$$\lambda_P = \frac{\Delta_i}{\Delta_i + \Delta_o}, \tag{21}$$

where $\Delta_i = \bar{\sigma}_{i,P+1} - \bar{\sigma}_{i,P}$ and $\Delta_o = \bar{\sigma}_{o,P} - \bar{\sigma}_{o,P+1}$.

We can exploit (21) to find the cost balances λ for which a given P is optimal. As λ_P is increasing in P, there is never any doubt if the optimum is greater than, equal to, or less than some candidate value of P. The values of λ for which $P^* < 30$ are plotted in Figure 3, illustrating two properties that follow from (A.6): P^* is increasing in L and in λ. $\lambda = 1$ implies zero cost of inventory deviations, only capacity costs are present. Since the capacity cost is decreasing in P, $P^* \to \infty$ as $\lambda \to 1$. Intuitively, as the cycle length increases, more temporal aggregation (pooling) and smoothing occurs. In cases when $\lambda \to 1$, the optimal cycle length, P^* increases.

3.1. The staggered order-up-to policy with equal overtime (STOUT-E)

In the STOUT policy, overtime work is concentrated to the first period of each order cycle. We may also be interested to see the effects of distributing the overtime work evenly over every period in the cycle compares to all-at-the-beginning-of-the-cycle allocation; not because it is more economical (Proposition 1 revealed that the all-at-the-beginning allocation is most economic), but to see how much money could be taken off the table from moving away from the practice of evenly distributed overtime. Starting with the production requirements of the STOUT policy, we divide the overtime work into P equal parts. This provides the STOUT-E policy:

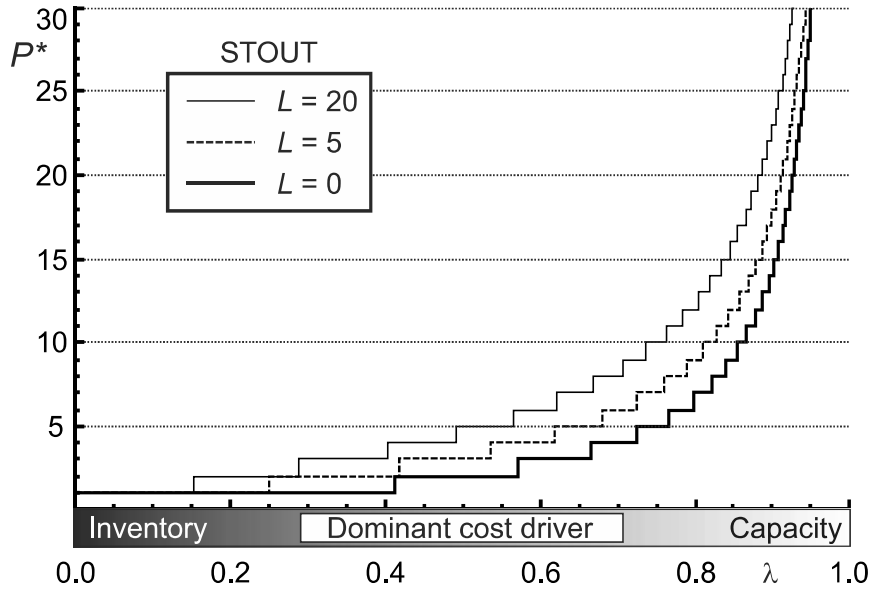$$o_{t,k} = x_k^* - x_{k-1}^* + P^{-1}\left(x_0^* - x_{t,0}\right), \tag{22}$$

13

Figure 3: Optimal order cycle lengths (P^*) under the STOUT policy.

where $x_0^* = \mathbb{E}\left[x_{t,0}\right] = x_P^* - \mu P$ is the observed inventory position at the start of the cycle before any orders have been placed. New values of x_k^* must be computed as the variances of the inventory and the orders differ from the variances generated by the STOUT policy.

Lemma 4. *For the STOUT-E policy,*

(a) the inventory variance is

$$\sigma_{i,k}^2 = \sigma_d^2 \left[k + L + \frac{(P-k)^2}{P} \right] ; \tag{23}$$

(b) the order variance is

$$\sigma_{o,k}^2 = \sigma_d^2/P \text{ for all } k. \tag{24}$$

The inventory variance of the STOUT-E policy (23) is greater than that of STOUT. Note that $\bar{\sigma}_{o,P}$ is identical for STOUT and STOUT-E. As a result, the realized capacity cost will be the same, despite the difference in overtime strategy. Thus for any relative weighting of inventory and capacity costs, the STOUT policy dominates the STOUT-E policy. Notably, the STOUT-E inventory variance is sometimes decreasing in k, and is minimized when $k = P/2$ for even P, or when

14

$k = (P \pm 1)/2$ for odd P. See Figure 4 for an example with $P = 5$. If one uses a constant safety stock, the heteroskedastic inventory variance will cause the availability to fluctuate. Availability is defined as the probability of satisfying all demand from stock in a period, $S_1 = \Phi(i_k^*/\sigma_{i,k})$. Figure 5 illustrates the availability fluctuations when the safety stock is a constant based on the standard deviation of the end-of-cycle inventory, $\sigma_{i,P}^2$. Although not optimal, constant safety stocks are common practice in industry. When safety stocks are set to minimize the expected inventory costs in each period, the availability is constant over the cycle.

Proposition 3. *The STOUT-E policy has a minimum cost when $\lambda \in [\lambda_{P-1}, \lambda_P]$.*

Note, while the STOUT-E policy has an optimal P^* given by the same approach as the STOUT policy, $P^*|_{\text{STOUT-E}} \geq P^*|_{\text{STOUT}}$ due to the increased inventory variance.
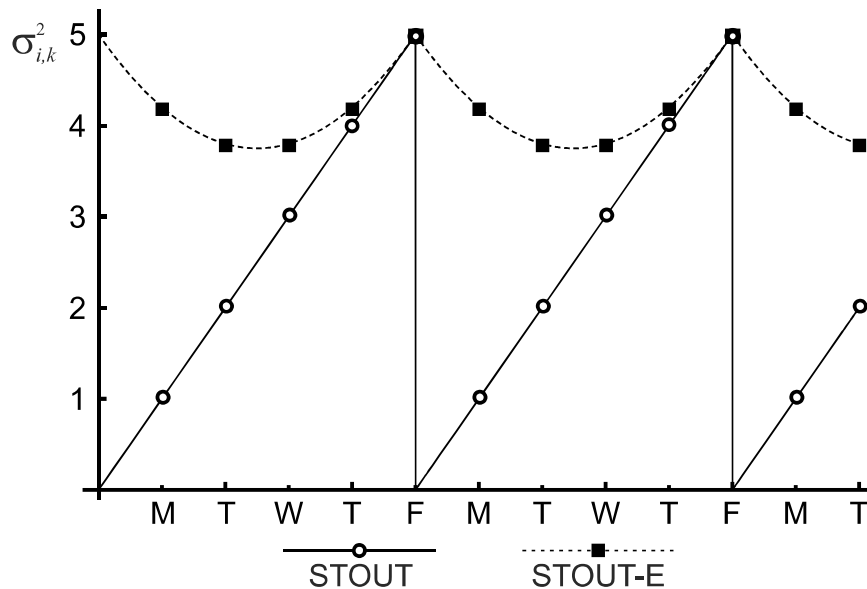

Figure 4: The heteroskedasticity of inventory levels depends on the overtime strategy used when $P = 5, L = 0$.

4. The staggered proportional order-up-to policy (SPOUT)

We have seen that the STOUT policy can exploit the order cycle length P to strike a favourable balance between the different cost drivers. The balance between inventory and capacity costs can also be managed by a proportional order-up-to policy (POUT) [28], which corrects a fraction α of the inventory position's error each time reordering takes place. This policy has a long history

Figure 5: Availability fluctuations following a (suboptimal) constant safety stock setting.

in the non-staggered setting, see for example, Deziel and Eilon [8], John et al. [19], Dejonckheere et al. [7], and Balakrishnan et al. [1]. This replenishment policy was implemented for monthly planning at the Eastman Kodak Company [28], for weekly planning at Lexmark [10], and for the thrice daily planning at Tesco [25]. These examples do not specify how overtime work was allocated; for this policy, we assume that all of the overtime work is done as soon as the plan is released to production; perhaps by extending or adding a shift, as Proposition 1 continues to hold in this setting.

To obtain the SPOUT policy, we add a proportional control parameter, α, to the first period of the STOUT policy. Using the definition $x_0^* = x_P^* - \mu P$; the SPOUT policy is

$$o_{t,k} = \begin{cases} x_1^* - x_0^* + \alpha\left(x_0^* - x_{t,0}\right), & \text{when } k = 1, \\ x_k^* - x_{k-1}^*, & \text{otherwise.} \end{cases} \tag{25}$$

The SPOUT orders are stationary if $0 \le \alpha < 2$, and the inventory is stationary if $0 < \alpha < 2$ [9]. The variances of this policy are as follows:

Lemma 5.

(a) The inventory variance under the SPOUT policy is

$$\sigma_{i,k}^2 = \sigma_d^2\left[k + L + \frac{P\left(1-\alpha\right)^2}{\alpha\left(2-\alpha\right)}\right]; \tag{26}$$

(b) the variance of the orders is

$$\sigma_{o,k}^2 = \begin{cases} \frac{\sigma_d^2\alpha P}{2-\alpha}, & \text{when } k = 1, \\ 0, & \text{otherwise.} \end{cases} \tag{27}$$

16

As SPOUT is a generalization of STOUT, several properties are inherited: The inventory variance increases linearly with σ_d^2, k, L and P. However, it is a convex decreasing function of α. The order variance (27) decreases with P, and increases with α. SPOUT cannot achieve a lower inventory variance than STOUT (they are identical when $\alpha = 1$), but it can always achieve a lower order variance. Thus, the SPOUT policy can always perform at least as well as the STOUT policy in the presence of both inventory and capacity costs.

4.1. Finding the optimal smoothing setting α^*

We have established that the STOUT policy can smooth production by extending the order cycle length. SPOUT also has this property, but additionally can smooth production via its feedback parameter α. To compare these approaches, we shall find the optimal α for the non-staggered ($P = 1$) SPOUT policy, and compare this to the STOUT policy with an optimal order cycle length. To obtain the optimal α for $P = 1$, we differentiate the total cost function under the SPOUT policy, $C_1|_{\text{SPOUT}}$, with respect to α,

$$\frac{dC_1|_{\text{SPOUT}}}{d\alpha} = \frac{(\alpha + \lambda - 1)(2\alpha\lambda - \alpha - \lambda + 1)}{\lambda^2} - (2 - \alpha)\alpha^3 L. \qquad (28)$$

From (28), we may observe that the optimal α^* is independent of the mean and standard deviation of demand. In addition, we obtain:

Corollary 2. *The optimal smoothing parameter α^* is (a) a decreasing function of L, and (b) a decreasing function of λ.*

Setting $dC_1|_{\text{SPOUT}}/d\alpha = 0$, and solving for α leads to a very large expression when $L > 0$, as it is a quartic function of α. However, the trivial case $L = 0$ provides $\alpha^* = 1 - \lambda$, with λ defined in (20). Figure 6 shows the results of numerically solving for α^* at different values of L, which confirms that $\alpha^* \leq 1 - \lambda$, and that the required damping increases with the lead time. For the special case $\{P = 1, L = 0\}$ the optimal total cost is $C_1^* = \psi\sigma_d\sqrt{1 - \lambda^2} + u\mu_d$. Whenever $\lambda > 0$, this cost is lower than the minimum cost obtainable via STOUT. This results from Proposition 4, which uses the following notation for brevity: A SPOUT(P, α) policy is defined as a SPOUT policy with arbitrary order cycle P and smoothing parameter α, with $C|_{\text{SPOUT}(P,\alpha)}$ representing the total cost (18). When comparing two policy settings, say SPOUT(P, α) and SPOUT(Q, β), we assume that the remaining variables $(\lambda, \psi, L, \mu, \sigma_d)$ are constant and identical between the configurations compared.

Proposition 4. *There exists a variable α, such that the expected cost of a SPOUT$(1, \alpha)$ policy never exceeds the expected cost of a SPOUT(P, β) policy, i.e. $C|_{SPOUT(1,\alpha)} \leq C|_{SPOUT(P,\beta)}$.*

Remark Proposition 4 shows it is more economical to embed production smoothing in the order policy via SPOUT than to acquire smoothing by manipulating the order cycle length. Proposition 4 also shows that the SPOUT policy is sufficient to avoid a capacity cost trap from short planning cycles in the absence of fixed planning costs.

17

Figure 6: Optimal values for the feedback controller α for the SPOUT policy when $P = 1$.

Corollary 3. *The smoothing parameter $\alpha^* \to 1$ as $P \to \infty$.*

Remark Corollary 3 implies that when P becomes large the SPOUT policy degenerates into the STOUT policy.

4.2. The staggered proportional policy with equal overtime

Just as the STOUT policy has an equal-overtime variant, so can one be identified for SPOUT; we study it for the same reasons. We define the SPOUT-E policy as

$$o_{t,k} = x_k^* - x_{k-1}^* + \alpha P^{-1} \left(x_0^* - x_{t,0} \right). \tag{29}$$

The variances required to calculate costs and availability are given below.

Lemma 6. *For the SPOUT-E policy,*

(a) the inventory variance is

$$\sigma_{i,k}^2 = \sigma_d^2 \left[k + L + \frac{(P - \alpha k)^2}{\alpha P (2 - \alpha)} \right]; \qquad (30)$$

(b) the order variance is

$$\sigma_{o,k}^2 = \frac{\sigma_d^2 \alpha}{P (2 - \alpha)}. \qquad (31)$$

Proposition 1 still holds, indicating equal capacity costs between SPOUT and SPOUT-E, but the inventory cost of the latter is higher, rendering the equal-overtime policy inferior to SPOUT. When k is specified, SPOUT-E has an inventory variance no less than STOUT-E. The minimal inventory occurs when $k = P/2$ for even P, or when $k = (P \pm 1)/2$ for odd P, just as we observed with the STOUT-E policy. Corollary 3 also holds for SPOUT-E.

The propositions presented thus far describe how these ordering policies operate, how they compare to each other, and how they should be configured in terms of order cycle length, smoothing, and safety stock settings. We shall now put these into context with a numerical study, where the features suggested by the propositions are highlighted.

5. Numerical study

To further explore the consequences of the four policies, consider the set-up $\{\mu = 10,\ \sigma_d = 1,\ L = 5,\ b = 9,\ h = 1,\ u = 40,\ v = 60\}^2$. With these settings, the optimal SPOUT[-E] setting is $\{P^* = 1,\ \alpha^* = 0.06\}$, following Proposition 4 and (28). The total cost of each strategy is illustrated in Figure 7, where safety stocks and capacity levels have been set to minimise per period costs, and α^* has been optimized numerically for STOUT[-E] configurations where $P > 1$.

The SPOUT policy gives the lowest total cost, regardless of P (Proposition 4). The cost advantage that can be gained from improving the ordering policy depends on P, as the adoption of production smoothing has a significant economic impact when P is small. When P is large, savings can instead be realized by changing the overtime strategy so that overtime production is collected to the start of the order cycle (Proposition 1).

The STOUT[-E] policy does not smooth production, and therefore suffers from impaired efficiency. Furthermore, when these policies are used, the order cycle length plays an important role in the balancing of capacity and inventory

---

[2]It seems reasonable to assume the $v = 1.5u$, reflecting the practice that overtime work to remunerated on a time-and-a-half basis. The ratio $b = 9h$ ensures that 90% availability is achieved when the safety stock has been set to minimize inventory costs. The relationship between $h = 1$ and $u = 40$ reflects that per period inventory holding costs are 2.5% of the production cost in regular hours.

Figure 7: Total variability cost under the four policies for different order cycle lengths.

costs. If the order cycle is too long, inventory costs dominate; if it is too short, capacity costs inflate. The STOUT policy has an optimal P at $P^* = 23$ (Proposition 2). The STOUT-E has an optimal P at $P^* = 17$ (Proposition 3).

When the order cycle is short, the production strategies without smoothing (STOUT[-E]) suffer from high costs, while both of the smoothing policies perform better (Proposition 4). In this case, the economic potential of smoothing production across cycles is greater than that of the changing overtime strategy (smoothing within cycles). The opposite holds true when the order cycle is long, as smoothing then has less impact on the total cost (Proposition 1). In these cases, a smart allocation of overtime within the cycle is more important.

Irrespective of demand variability and cost factors, the production control policy should collect the inventory corrections to a short period at the beginning of the order cycle. It is then desirable to react in a moderate but timely fashion to keep the overtime costs in check, hence the need for production smoothing (Propositions 1 and 4).

We shall now detail how the four policies are applied in practice.

Example 1. *Suppose we are to place orders for a system with the same parameter settings as the preceding numerical example with the additional setting of $P = 5$ and an initial observed inventory position of $x_{0,0} = 47$. As machine precision should be used, the numerical calculations below are truncated instead of rounded. The calculations below are for the STOUT policy; corresponding calculations for the other policies appear in Table 5.*

To calculate the orders in a cycle we take the following steps:

1. Determine the target inventory positions, x_k^*, in the cycle. From (7) we

obtain $x_k^* = \mu(k + L) + \sigma_{i,k}\Phi^{-1}\left[b/(b + h)\right] = 10(k + 5) + \sqrt{k + 5}\Phi^{-1}(0.9)$. For the first period, inserting $k = 1$ gives $x_1^* = 63.13$. Increment k in the same calculation to obtain the remaining values of x_k^*. For policies other than STOUT, $\sigma_{i,k}$ is calculated differently: for STOUT-E, use (23) for SPOUT, use (26); and for SPOUT-E, use (30).

2. Obtain the deterministic production requirement by calculating the difference between the target inventory positions of consecutive periods. Take $x_k^* - x_{k-1}^*$ for every value of k. For $k = 1$ this is $x_1^* - x_0^* = 63.13 - 54.05 = 9.08$, continuing with $k = 2$, $x_2^* - x_1^* = 10.25$, and so forth. For each period k (and for every policy), order the deterministic production requirement.

3. Calculate the deficit between the target and the actual inventory position, $x_0^* - x_{t,0} = 63.13 - 47 = 16.13$. Depending on the policy in use,

   STOUT: Add the entire deficit to the first order $(o_{t,1})$.

   STOUT-E: Divide the deficit by P; add to every order $(o_{t,1}$ to $o_{t,P})$.

   SPOUT: Multiply the deficit by α; add to the first order.

   SPOUT-E: Multiply the deficit by α; divide by P, add the resulting amount to every order.

Complete calculations for all strategies appear in Table 2.

Example 2. *To find P^* under STOUT:*

1. Identify the necessary cost parameters. In this numerical study, they are: $b = 9$, $h = 1$, $u = 40$, and $v = 60$.

2. Calculate λ by first calculating ψ using (19) as a preliminary step before calculating (20). In this case $\psi = 60\varphi\left[\Phi^{-1}(0.33)\right] + 10\varphi\left[\Phi^{-1}(0.9)\right] = 23.57$. Then $\lambda = 60\varphi\left[\Phi^{-1}(0.33)\right]/23.57 = 0.9255$.

3. Inspect the $L = 5$ line in Figure 3 to see that $\lambda = 0.9255$ corresponds to $P^* = 23$. If P^* cannot be identified on the graph, enumerate λ_P from unity to some large value. When there are two successive λ_P values such that $\lambda_{P-1} \le \lambda < \lambda_P$ then P^* has been found. In this example, $\lambda_{23} \le \lambda < \lambda_{24}$, or equivalently $0.92409 \le 0.9255 < 0.927538$, demonstrating that $P^* = 23$.

Remark. Mathematica enumerated all values of λ_P for $P < 100$ and $L < 100$ in less than 1 second on an Intel i7-4600U CPU @ 2.10GHz. Should this enumeration be computationally expensive, one may search for a value λ_P such that $\lambda < \lambda_P$; then perform a binary search.

Example 3. *Using the procedures outlined above we have collected the results for all policies under two sets of cost parameters in Table 4 for comparative purposes. As the inventory holding and backlog costs increases, λ decreases and α^* increases in the SPOUT[-E] policies (as per Figure 6) and the optimal cycle length, P^*, in the STOUT[-E] policies decrease (as per Propositions 2 and 3).*

Period t		5	6	7	8	9	10	11
Index k		5	1	2	3	4	5	1
STOUT	x_k^*	–	63.13	73.39	83.62	93.84	104.05	–
	$x_k^* - x_{k-1}^*$	–	9.08	10.25	10.23	10.21	10.20	–
	$x_0^* - x_{0,0}$	–	7.05	↓	↓	↓	↓	–
	$o_{0,k}$	–	16.13	10.25	10.23	10.21	10.20	–
STOUT-E	x_k^*	–	63.88	73.80	83.80	93.88	104.05	–
	$x_k^* - x_{k-1}^*$	–	9.83	9.91	10	10.08	10.16	–
	$(x_0^* - x_{0,0})/P$	–	1.41	1.41	1.41	1.41	1.41	–
	$o_{0,k}$	–	11.24	11.32	11.41	11.49	11.57	–
SPOUT[a]	x_k^*	–	64.77	74.94	85.10	95.26	105.41	–
	$x_k^* - x_{k-1}^*$	–	9.35	10.16	10.16	10.15	10.15	–
	$\alpha(x_0^* - x_{0,0})$	–	8.41	↓	↓	↓	↓	–
	$o_{0,k}$	–	17.77	10.16	10.16	10.15	10.15	–
SPOUT-E[b]	x_k^*	–	65.45	75.44	85.44	95.45	105.47	–
	$x_k^* - x_{k-1}^*$	–	9.98	9.99	10	10.00	10.01	–
	$\alpha(x_0^* - x_{0,0})/P$	–	1.69	1.69	1.69	1.69	1.69	–
	$o_{0,k}$	–	11.67	11.68	11.69	11.70	11.70	–

Dashes (–) refer to values unrelated to the present ordering decision. [a] $\alpha^* = 0.217944$
[b] $\alpha^* = 0.211445$; numerically optimized.

Table 3: Calculating the orders to be received in periods 6–10, when $L = 5, P = 5$.

6. Managerial implications

We have considered three approaches to production smoothing: (1) extending the order cycle length; (2) deferring overtime or idling to another cycle; and (3) smoothing overtime work or idling within the current cycle. Manipulating the order cycle length allows managers to control the balance between inventory and capacity costs. Increasing the cycle length decreases capacity costs at the expense of increased inventory costs. If the STOUT policy is used, cycle length increases are an indirect way to implement production smoothing. On the other hand, if SPOUT is used and properly configured, the order cycle should be set to unity (i.e. daily planning) and all smoothing should be effected through a well-tuned feedback parameter.

The option to defer overtime to future cycles is the preferred way of smoothing, as it leads to a lower total cost than smoothing via the order cycle length. Higher capacity costs relative to inventory costs require more smoothing, as do production systems where the order cycle length is reduced.

The third kind of smoothing, level allocation within a cycle, is entirely detrimental. An hour of overtime costs the same regardless of the weekday it

Costs	$\{h=1, b=9, u=40, v=60\}$		$\{h=10, b=90, u=40, v=60\}$	
Scaling factors	$\{\lambda=0.9255, \psi=23.571\}$		$\{\lambda=0.5541, \psi=39.3658\}$	
Policy	P^*	α^*	P^*	α^*
SPOUT[-E]	1	0.0600	1	0.2993
STOUT	23	—	4	—
STOUT-E	17	—	2	—

Table 4: Comparison of our staggered policies as the inventory cost increase

is worked, so this kind of smoothing has no effect on capacity costs. Inventory costs are however minimized by correcting all deviations as soon as possible, i.e. doing all overtime or idling at the start of each cycle.

In summary, a well-tuned proportional policy is best, where the order cycle is as short as possible, and where all overtime is done at the start of the cycle. Ideally, the order cycle length should be unity, removing staggering altogether.

In practice, we often find that companies use a STOUT or STOUT-E policy with order cycles of a week or more, as these policies are natural extensions of the ubiquitous order-up-to policy. The ideal state is however a combination of smoothing and a short order cycle length. To reach this state, the appropriate course of action is first to implement smoothing and proper overtime allocation (via SPOUT), and thereafter to reduce the order cycle length. This sequence ensures that the total cost of production decreases as one implements the changes. Should one reduce the order cycle length before smoothing is implemented, capacity costs may soar as the smoothing effect of the order cycle length is lost: We may call this the *capacity cost trap* associated with the order cycle length. This, along with the preferred improvement path is illustrated in Figure 8.

Alternatively, current planning practice may be to chase daily demand in a pure pull or pass-on-orders mode. This is equivalent to STOUT with unit cycle length (i.e. OUT) and is optimal in the presence of inventory only costs. In the presence of capacity costs, level scheduling (with the level recalculated every P periods via the STOUT-E policy) dominates the OUT policy. This strategy is, in turn, dominated by level scheduling and accounting for all stochastic deviations as quickly as possible via the STOUT policy, which may have a different planning cycle length. However, the level schedule route is in fact a dead end; it is sufficient to keep the unit cycle length and smooth the demand variability with the SPOUT policy.

7. Concluding remarks

Cyclical planning affords production systems new ways for production smoothing. In particular, the length of the order cycle has an intrinsic smoothing effect that appears even when a conventional staggered OUT policy is used. The other type of smoothing comes from reallocating overtime within the cycle, but this is ineffective. Instead, excessive overtime or idling should be shifted across cycles, into the future.

Inventory and capacity cost

Figure 8: Avoid the capacity cost trap by improving the order policy before reducing the order cycle length.

These findings rely on two fundamental assumptions: (1) that demand is i.i.d., and (2) that the capacity cost includes an installed (fixed) capacity cost and the opportunity to pay overtime for output beyond the installed base. When the first assumption is relaxed, as with autocorrelated demand, a different overtime allocation within the same cycle may be beneficial [16]. The second assumption, of piecewise linear capacity costs, is central to the conclusions, but it only applies to production contexts where the marginal cost of overtime is fixed, meaning that we effectively can extend a shift and pay per minute of overtime work.

Although this applies to the companies that inspired the study, other companies may require that a whole shift of overtime be worked, even if the required production is only one unit. Other companies have an overtime cost that is the same as for normal production. Then smoothing becomes a non-issue, and the STOUT policy prevails. In short, the capacity cost function determines how one should approach production smoothing, and it is not only determined by the process type, but varies between countries and industrial contexts.

In some situations, peak demands could be met by a subcontractor with an effectively limitless capacity, in others only a limited amount of overtime is available in-house; there will also be cases where *hard* capacity limits prohibit the production of peak orders. In general terms, the impact of hard capacity constraints is complex with few analytical results. Ponte et al. [24] used simulation

24

to show the capacity constraint had a smoothing effect on the orders with the OUT policy; furthermore the available capacity could be considered to be a decision variable, balancing the cost of inventory and order variability. One might conjecture that this behaviour would also be observed in a staggered delivery setting; we leave this for future work.

Our ultimate recommendation to managers is that they do not consider the order cycle as just another lead time, but as a lever for balancing the variability of inventory and production. We recommend reducing the order cycle length only if a policy capable of smoothing across cycles is in place.

Acknowledgments

Appendix A. Technical proofs

Proof of Lemma 1

This is essentially a newsvendor problem. The expected inventory cost in period $t + k + L$ is

$$\mathbb{E}\left[j\left(i_{t,k}\right)\right] = h\mathbb{E}\left[x_{t,k} - D_{t,k}\right] + (b+h) \int\limits_{i_{t,k}<0} \left(D_{t,k} - x_{t,k}\right) d\mathbb{P}. \qquad (A.1)$$

Differentiating with respect to our decision variable $x_{t,k}$, setting the result equal to zero, and solving for $x_{t,k}$ gives the optimality criterion $h/(b+h) = \mathbb{P}(i_{t,k} < 0)$. Applying $F_{i,k}^{-1}(\cdot)$ to both sides gives (6). ■

Proof of Lemma 2

Taking the expectation of (10) over $o_{t,k}$ gives

$$\mathbb{E}\left[a(o_{t,k})\right] = uz_k + v \int\limits_{o_{t,k}>z_k} \left(o_{t,k} - z_k\right) d\mathbb{P}. \qquad (A.2)$$

Differentiating with respect to z_k, setting this equal to zero, and solving for z_k gives $(v-u)/v = \mathbb{P}(o_{t,k} \leq z_k)$. Applying $F_{o,k}^{-1}(\cdot)$ to both sides gives (11), completing the proof. ■

Proof of Lemma 3

(a) Taking the variance of (4) provides

$$\sigma_{i,k}^2 = \text{var}(x_{t,k}) + \text{var}(D_{t,k}) = \sigma_d^2(k+L), \qquad (A.3)$$

as demand is i.i.d., and therefore future demand is uncorrelated with the inventory position. For this policy $x_{t,k} = x_k^*$ is constant over time (for each k), providing $\text{var}(x_{t,k}) = 0$.

(b) From (14), (4), and the definition $o_{t,k} = r_{t+k+L}$, we obtain

$$o_{t+P,1} = x_1^* - i_{t+P} = x_1^* - \left(x_P^* - \sum_{n=1}^{P} d_{t+n} \right), \qquad (A.4)$$

which leads to $\sigma_{o,1}^2 = \sigma_d^2 P$. For the remaining periods with $k > 1$, $o_{t,k} = \mathbb{E}[o_{t,k}] = x_k^* - x_{k-1}^*$, and therefore $\sigma_{o,k}^2 = 0$, completing the proof. ∎

Proof of Proposition 1

Let $q_k \geq 0, \sum_{k=1}^{P} q_k = 1$ represent the allocation of overtime or idling within each cycle, such that $o_{t,k} = q_k \sum_{n=1}^{P} o_{t,n}$. The variance is $\sigma_{o,k}^2 = q_k^2 \text{var}\left(\sum_{n=1}^{P} o_{t,n}\right)$ which leads to

$$\bar{\sigma}_{o,P} = \frac{1}{P} \sum_{k=1}^{P} \sqrt{q_k^2 \text{var}\left(\sum_{n=1}^{P} o_{t,n}\right)} = \frac{1}{P} \sqrt{\text{var}\left(\sum_{n=1}^{P} o_{t,n}\right)}, \qquad (A.5)$$

revealing that q_k has no influence on the capacity cost. To minimize the inventory cost J_P, we minimize the inventory variance of each period. It is obtained through (4) as $\sigma_{i,k}^2 = \text{var}(x_{t,k}) + \text{var}(D_{t,k})$, of which only $\text{var}(x_{t,k})$ may be influenced. As $x_{t,k} = x_{t,0} + \sum_{m=1}^{k} q_m \sum_{n=1}^{k} o_{t,n}$, we find that $\text{var}(x_{t,k})$ is minimized when $q_1 = 1$, which also minimizes $\bar{\sigma}_{i,P}$, J_P, and hence C_P. ∎

Proof of Corollary 1

(a) Is obvious.

(b) It is sufficient to show that $\forall P$, $\sum_{k=1}^{P} \sqrt{k+L} > P$. Without further consequences, consider the case when $L = 0$. The first addend of the sum is $\sqrt{k=1} = 1$; subsequent addends are increasing and $\forall P$, $\sum_{k=1}^{P} \sqrt{k+L} > P$. Therefore $\bar{\sigma}_{i,P}$ is increasing in P. ∎

Proof of Proposition 2

Let λ_P be the point at which we are indifferent between the choice of P or $P+1$, occurring when $C_P(\lambda_P) = C_{P+1}(\lambda_P)$. Solving for λ_P gives (21), which is equivalent to

$$\lambda_P = 1 - \left[1 + \left(\sqrt{P} + \frac{P}{\sqrt{P+1}} \right) \frac{\sigma_{i,P+1} - \bar{\sigma}_{i,P}}{\sigma_d} \right]^{-1}. \qquad (A.6)$$

From this expression, it is clear that λ_P is increasing in P. Therefore, P minimises costs under the STOUT policy when $\lambda \in [\lambda_{P-1}, \lambda_P]$. ∎

Proof of Lemma 4

(a) From (4) and (22) we see that $x_{t,P} = x_P^*$. Consequently, $x_{t,0} = x_P^* - \sum_{n=1}^{P} d_{t-P+n}$. The inventory position can then be expressed as

$$x_{t,k} = x_k^* - \frac{P-k}{P} \sum_{n=1}^{P} (d_{t-n} - \mu), \tag{A.7}$$

which provides the variance of the inventory position,

$$\mathrm{var}\,(x_{t,k}) = \sigma_d^2 \frac{(P-k)^2}{P}. \tag{A.8}$$

Recall that future demand is uncorrelated with the inventory position:

$$\sigma_{i,k}^2 = \mathrm{var}(x_{t,k}) + \mathrm{var}\,(D_{t,k}) = \sigma_d^2 \left[k + L + \frac{(P-k)^2}{P} \right], \tag{A.9}$$

and the first part of the proof is complete.

(b) Note that $x_{t,0} = x_0^* - \sum_{n=1}^{P} (d_{t-P+n} - \mu)$, which when inserted in (22) provides

$$o_{t,k} = x_k^* - x_{k-1}^* + \sum_{n=1}^{P} \frac{d_{t-P+n} - \mu}{P}. \tag{A.10}$$

Taking the variance of (A.10) gives $\sigma_{o,k}^2 = \sigma_d^2/P$, completing the proof. ∎

Proof of Proposition 3

This follows by showing that λ_P is increasing. From (21), we observe that λ_P is increasing if $\Delta_i \geq \Delta_o$. Since $\Delta_o|_{\text{STOUT}} = \Delta_o|_{\text{STOUT-E}}$, it remains to be shown that $\Delta_i|_{\text{STOUT-E}} \geq \Delta_i|_{\text{STOUT}}$. The latter inequality can be rearranged as

$$\bar{\sigma}_{i,P+1}|_{\text{STOUT-E}} - \bar{\sigma}_{i,P+1}|_{\text{STOUT}} \geq \bar{\sigma}_{i,P}|_{\text{STOUT-E}} - \bar{\sigma}_{i,P}|_{\text{STOUT}}. \tag{A.11}$$

As this inequality compares successive values of P, the inequality holds if $\bar{\sigma}_{i,P}|_{\text{STOUT-E}} - \bar{\sigma}_{i,P}|_{\text{STOUT}}$ is increasing in P. The right hand side of (A.11) is an average of square root terms: $P^{-1} \left(\sum_{k=1}^{P} \sqrt{k + L + \frac{(P-k)^2}{P}} - \sqrt{k+L} \right)$. This can be seen to be increasing by differencing the term $\frac{(P-k)^2}{P}$, providing $\frac{(P+1-k)^2}{P+1} - \frac{(P-k)^2}{P} = \frac{P^2+P-k^2}{P(P+1)}$, which is positive for $P \geq k$. As the average of an increasing function is also increasing, the inequality in (A.11) holds. ∎

Proof of Lemma 5

(a) First, we express $x_{t,0}$ in terms of $x_{t-P,0}$

$$x_{t,0} = x_{t-P,0} + \sum_{n=1}^{k} o_{t-P,n} - d_{t-P+n}$$

$$= x_{t-P,0} + x_P^* - x_0^* + \alpha\left(x_0^* - x_{t-P,0}\right) - \sum_{n=1}^{P} d_{t-P+n} \qquad \text{(A.12)}$$

$$= (1-\alpha)\, x_{t-P,0} + \alpha x_0^* - \sum_{n=1}^{P} \varepsilon_{t-P+n}.$$

Continuing the recursion q cycles back gives

$$x_{t,0} = (1-\alpha)^q\, x_{t-qP,0} + \sum_{m=1}^{q} (1-\alpha)^{m-1}\left[\alpha x_0^* - \sum_{n=1}^{P}\varepsilon_{t-mP+n}\right]. \qquad \text{(A.13)}$$

When $q \to \infty$ we obtain

$$x_{t,0} = x_0^* - \sum_{n=1}^{P}\sum_{m=1}^{\infty} (1-\alpha)^{m-1}\,\varepsilon_{t-mP+n}, \qquad \text{(A.14)}$$

which reveals the expectation $\mathbb{E}\left[x_{t,0}\right] = x_0^*$. As $x_{t,k} = x_{t,0} + \sum_{n=1}^{k} o_{t,n}$

$$x_{t,k} = x_k^* - \sum_{n=1}^{P}\sum_{m=1}^{\infty} (1-\alpha)^{m}\,\varepsilon_{t-mP+n}, \qquad \text{(A.15)}$$

taking the variance of $x_{t,k}$ and adding the variance of lead-time demand gives

$$\sigma_{i,k}^2 = \sigma_d^2\left[k + L + \frac{P\left(1-\alpha\right)^2}{\alpha\left(2-\alpha\right)}\right],$$

completing this part of the proof.

(b) Inserting (A.14) in (25) provides

$$o_{t,1} = x_1^* - x_0^* + \alpha\sum_{n=1}^{P}\sum_{m=1}^{\infty} (1-\alpha)^{m-1}\,\varepsilon_{t-mP+n}. \qquad \text{(A.16)}$$

Taking the variance gives

$$\sigma_{o,1}^2 = \frac{\sigma_d^2\alpha P}{2-\alpha}.$$

For $k \neq 1$, the orders are constant, therefore $\forall k > 1$, $\sigma_{o,k}^2 = 0$, and the proof is complete. ∎

Proof of Corollary 2

(a) Differencing the first-order criterion (28) with respect to L provides

$$\alpha^3(\alpha - 2),$$ (A.17)

which is negative and decreasing in α, meaning that α^* must be reduced when incrementing L for the first-order condition to be maintained.

(b) Define the functions $f(\alpha) = \bar{\sigma}_{o,1}$ and $g(\alpha) = \bar{\sigma}_{i,1}$. Then, the total cost $C_1 \propto \lambda f(\alpha) + (1 - \lambda)g(\alpha)$, provides the first-order condition for α,

$$\lambda f'(\alpha) + (1 - \lambda)g'(\alpha) = 0.$$ (A.18)

Rearranging the first-order condition for λ yields

$$\lambda = \frac{-g'(\alpha)}{f'(\alpha) - g'(\alpha)}.$$ (A.19)

Taking the derivative of λ w.r.t. α yields,

$$\frac{d\lambda}{d\alpha} = \frac{\sigma_{o,1}(1 + L\alpha(3(1 - \alpha) + \alpha^2)}{(\alpha - 2)\sqrt{\alpha}\sigma_{i,1}(\sigma_{o,1}(1 - \alpha) + \sqrt{\alpha^3}\sigma_{i,1})^2},$$ (A.20)

which is always negative, indicating λ is decreasing in α; equivalently α is decreasing in λ. ∎

Proof of Proposition 4

Consider two SPOUT policy configurations, SPOUT$(1, \alpha)$, and SPOUT(P, β), for the same underlying demand variance, lead time, and cost parameters. Assuming that β and P are arbitrary, we shall set α such that the capacity cost is equal for the two configurations, which results in a higher inventory cost for the (P, β) configuration.

Setting the capacity cost to be equal, $A|_{SPOUT(1,\alpha)} = A|_{SPOUT(P,\beta)}$ using (13), provides

$$\frac{1}{P}\sqrt{\frac{P\beta}{(2 - \beta)}} = \sqrt{\frac{\alpha}{2 - \alpha}},$$ (A.21)

which results in the following expression for α,

$$\alpha = \frac{2\beta}{\beta + P(2 - \beta)}.$$ (A.22)

Using (A.22) and (26), we can express the difference in the inventory variances for a given k, $\Delta_P = \sigma_{i,1}^2|_{SPOUT(1,\alpha)} - \sigma_{i,k}^2|_{SPOUT(P,\beta)}$, as

$$\frac{(P - 1)[\beta + (3\beta - 4)P]}{4P(\beta - 2)} - k.$$ (A.23)

Differencing (A.23) by taking $\Delta_{P+1} - \Delta_P$ gives

$$\frac{\beta + P(P+1)(3\beta - 4)}{4P(\beta - 2)(P+1)}, \tag{A.24}$$

which is the increase in inventory variance experienced when P increases for a fixed k. On the permissible parameter range $\{0 < \beta \leq 2, P \in \mathbb{Z}^+\}$, the supremum of (A.24) is $\frac{1}{2}$, as $\beta \to 0$, which may be verified by differentiating (A.24) with respect to β. As $\sigma_{i,1}^2|_{\text{SPOUT}(1,\alpha)} = \sigma_{i,1}^2|_{\text{SPOUT}(1,\beta)}$, the supremum of (A.24) implies that $\sigma_{i,1}^2|_{\text{SPOUT}(1,\alpha)} - \sigma_{i,k}^2|_{\text{SPOUT}(P,\beta)} \leq k - 1/2$, and by extension that $\sigma_{i,1}^2|_{\text{SPOUT}(1,\alpha)} - \sigma_{i,2}^2|_{\text{SPOUT}(P,\beta)} \leq k - P/2$. This further implies that $C|_{\text{SPOUT}(1,\alpha)} \leq C|_{\text{SPOUT}(P,\beta)}$ if for positive m,

$$\sum_{k=1}^{P} \sqrt{m + 1/2 + P/2} \leq \sum_{k=1}^{P} \sqrt{m + k}, \tag{A.25}$$

which is true as a consequence of Jensen's inequality. Note, here m is a constant term related to the costs and the normal distribution, see (9) and (13). This completes the proof. ∎

Proof of Corollary 3

The total cost C is a weighted sum of $\bar{\sigma}_{o,P}$ and $\bar{\sigma}_{i,P}$. First we observe that $\lim_{P\to\infty} \bar{\sigma}_{o,P} = 0$, eliminating the influence of $\bar{\sigma}_{o,P}$. Remaining is $\bar{\sigma}_{i,P}$, which can be minimized by minimizing $\sigma_{i,k}^2$, providing $\arg\min_\alpha \sigma_{i,P}^2 = 1$. ∎

Proof of Lemma 6

(a) In Lemma 5(a), replacing the SPOUT orders (25) with the SPOUT-E orders (29), produces the same expression for $x_{t,0}$, which is (A.14). As $x_{t,k} = x_{t,0} + \sum_{n=1}^{k} o_{t,n}$ we obtain

$$x_{t,k} = x_k^* - x_0^* + \frac{\alpha k}{P}\left(x_0^* - x_{t,0}\right)$$
$$= x_k^* - \sum_{n=1}^{P} \sum_{m=1}^{\infty} \frac{P - \alpha k}{P}\left(1 - \alpha\right)^{m-1} \varepsilon_{t-mP+n}. \tag{A.26}$$

Taking the variance of $x_{t,k}$ and adding the variance of lead-time demand, $\sigma_d^2(k + L)$, gives (30) completing this part of the proof.

(b) Inserting (A.14) into (29) provides

$$o_{t,k} = x_k^* - x_{k-1}^* + \sum_{n=1}^{P} \sum_{m=1}^{\infty} \frac{\alpha}{P}\left(1 - \alpha\right)^{m-1} \varepsilon_{t-mP+n}. \tag{A.27}$$

Taking the variance gives (31), completing the proof. ∎

Appendix B. Table of notation

Symbol	Domain	Description
$\cdot\|_{\text{POLICY}}$	—	Any variable conditional on the POLICY
α	$[0,2)$	The proportional smoothing parameter
α^*	$[0,2)$	An optimal proportional smoothing parameter
$\delta_i,\ \delta_o$	$\mathbb{R}$	The change in $\bar{\sigma}$ from incrementing P
λ	$[0,1]$	The relative weighting of capacity costs to inventory costs
λ_P	$[0,1]$	The relative weighting of capacity costs to inventory costs in a policy with cycle length P
μ	$\mathbb{R}$	Mean demand
$\Phi(\cdot)$	$\mathbb{R}$	CDF of the standard normal distribution
$\Phi(\cdot)^{-1}$	$\mathbb{R}$	Inverse CDF of the standard normal distribution
ψ	$\mathbb{R}$	A cost scaling factor
$\sigma^2_{i,k}$	$\mathbb{R}$	Variance of the inventory observed after the k^{th} receipt in a cycle
$\sigma^2_{o,k}$	$\mathbb{R}$	Variance of the k^{th} production order in a cycle
σ^2_d	$\mathbb{R}$	Variance of demand
$\bar{\sigma}_{i,P}$	$\mathbb{R}$	Average standard deviation of the inventory level
$\bar{\sigma}_{o,P}$	$\mathbb{R}$	Average standard deviation of production orders
ε_t	$\mathbb{R}$	Random deviation from the mean of period t's demand
$\varphi(\cdot)$	$\mathbb{R}$	Probability density function of the standard normal distribution
$a(o_{t,k})$	$\mathbb{R}$	Single-period capacity cost for order quantity $o_{t,k}$ with installed capacity z_k
A_P	$\mathbb{R}$	Expected long-run capacity cost per period
b	$\mathbb{R}^+$	Inventory backlog cost per unit and period
C_P	$\mathbb{R}$	The expected long-run cost of capacity and inventory in a policy with cycle length P
$D_{t,k}$	$\mathbb{R}$	Total demand from order placement to receipt
d_t	$\mathbb{R}$	Demand in period t
$F^{-1}_{i_{t,k}}$	$\mathbb{R}$	Inverse CDF of $i_{t,k}$
$F^{-1}_{o_{t,k}}$	$\mathbb{R}$	Inverse CDF of $o_{t,k}$
h	$\mathbb{R}^+$	Inventory holding cost per unit and period
i^*_k	$\mathbb{R}$	Optimal target inventory level, when the k^{th} order is received for a given policy
$i_{t,k}$	$\mathbb{R}$	On-hand inventory recorded as received in period $t+k+L$
i_t	$\mathbb{R}$	On-hand inventory in period t
$j(i_{t,k})$	$\mathbb{R}$	The single-period inventory cost for inventory level $i_{t,k}$
J_P	$\mathbb{R}$	Expected long-run inventory cost per period
k	$\mathbb{Z}^+$	Index of production order, as in the k^{th} placed in period t
L	$\mathbb{Z}_0$	The lead time from order release until produced and accounted for as inventory
$o_{t,k}$	$\mathbb{R}$	The production amount ordered for the k^{th} order placed in period t

Symbol	Domain	Description
P	$\mathbb{Z}^+$	Order cycle length, i.e. periods between two successive ordering occasions
q_k	$\mathbb{R}$	The fractional allocation of a cycle's overtime or idling to the k^{th} order
r_t	$\mathbb{R}$	Receipts from production, recorded as inventory in period t
t	$\mathbb{N}$	Time period
u	$\mathbb{R}^+$	Capacity cost under regular production
v	$\mathbb{R}^+$	Unit capacity cost under overtime production
x_k^*	$\mathbb{R}$	Optimal inventory position associated with the k^{th} order of a cycle for a given policy
$x_{t,0}$	$\mathbb{R}$	The inventory position in period t before any order is placed
$x_{t,k}$	$\mathbb{R}$	The inventory position in period t after placing the k^{th} order
z	$\mathbb{R}$	Capacity reserved at the regular price
z_k	$\mathbb{R}$	Capacity reserved at the regular price for the k^{th} order in a cycle

Table B.5: Summary of notation

References

[1] Balakrishnan, A., Geunes, J., and Pangburn, M. S. (2004) Coordinating Supply Chains by Controlling Upstream Variability Propagation. *Manufacturing & Service Operations Management*, vol. 6(2), pp. 163–183.

[2] Bradley, J. R. and Conway, R. W. (2003) Managing cyclic inventories. *Production and Operations Management*, vol. 12(4), pp. 464–479.

[3] Burbidge, J. (1983) Five golden rules to avoid bankruptcy. *Production Engineer*, vol. 62(10), pp. 13–14.

[4] Chiang, C. (2001) Order splitting under periodic review inventory systems. *International Journal of Production Economics*, vol. 70(1), pp. 67–76.

[5] Chiang, C. (2009) A periodic review replenishment model with a refined delivery scenario. *International Journal of Production Economics*, vol. 118(1), pp. 253–259.

[6] Churchman, C. W., Ackoff, R. L., and Arnoff, E. L. (1957) *Introduction to Operations Research.* John Wiley & Sons, New York.

[7] Dejonckheere, J., Disney, S. M., Lambrecht, M. R., and Towill, D. R. (2003) Measuring and avoiding the bullwhip effect: A control theoretic approach. *European Journal of Operational Research*, vol. 147(3), pp. 567–590.

[8] Deziel, D. P. and Eilon, S. (1967) A linear production-inventory control rule. *Production Engineer*, vol. 46(2), pp. 93–104.

[9] Disney, S. M. (2008) Supply chain aperiodicity, bullwhip and stability analysis with Jury's inners. *IMA Journal of Management Mathematics*, vol. 19(2), pp. 101–116.

[10] Disney, S. M., Hoshiko, L., Polley, L., and Weigel, C. (2013) *Removing bullwhip from Lexmark's toner operations*. In *Proceedings of the 24th Annual Conference of the Production and Operations Management Society*. Denver, CO. May 3rd – 6th, 10 pages.

[11] Flynn, J. (2000) Selecting T for a periodic review inventory model with staggered deliveries. *Naval Research Logistics (NRL)*, vol. 47(4), pp. 329–352.

[12] Flynn, J. (2001) Selecting review periods for a coordinated multi-item inventory model with staggered deliveries. *Naval Research Logistics (NRL)*, vol. 48(5), pp. 430–449.

[13] Flynn, J. (2008) An effective heuristic for the review period in an inventory model with staggered deliveries and normal demands. *European Journal of Operational Research*, vol. 186(2), pp. 671–680.

[14] Flynn, J. and Garstka, S. (1990) A dynamic inventory model with periodic auditing. *Operations Research*, vol. 38(6), pp. 1089–1103.

[15] Flynn, J. and Garstka, S. (1997) The optimal review period in a dynamic inventory model. *Operations Research*, vol. 45(5), pp. 736–750.

[16] Hedenstierna, C. P. T. and Disney, S. M. (2016) Inventory performance under staggered deliveries and autocorrelated demand. *European Journal of Operational Research*, vol. 249(3), pp. 1082–1091.

[17] Hedenstierna, P. (2009) *Modelling the Benefits of Supply Chain Segmentation: An Empirical Study in a Global Fast Moving Consumer Goods (FMCG) Company*. Master's thesis, Cranfield University, Cranfield.

[18] Hosoda, T. and Disney, S. M. (2012) On the replenishment policy when the market demand information is lagged. *International Journal of Production Economics*, vol. 135(1), pp. 458–467.

[19] John, S., Naim, M., and Towill, D. R. (1994) Dynamic analysis of a WIP compensated decision support system. *International Journal of Manufacturing System Design*, vol. 1(4), pp. 283–297.

[20] Kaku, B. K. and Krajewski, L. J. (1995) Period Batch Control in group technology. *International Journal of Production Research*, vol. 33(1), pp. 79–99.

[21] Modigliani, F. and Hohn, F. E. (1955) Production planning over time and the nature of the expectation and planning horizon. *Econometrica, Journal of the Econometric Society*, vol. 23(1), pp. 46–66.

[22] Nathan, J. and Venkataraman, R. (1998) Determination of master production schedule replanning frequency for various forecast window intervals. *International Journal of Operations & Production Management*, vol. 18(8), pp. 767–777.

[23] Ohno, T. (1988) *Toyota Production System: Beyond Large-Scale Production.* Productivity Press, New York.

[24] Ponte, B., Wang, X., Fuente, D. d. l., and Disney, S. M. (2017) Exploring nonlinear supply chains: the dynamics of capacity constraints. *International Journal of Production Research*, vol. 55(14), pp. 4053–4067.

[25] Potter, A. and Disney, S. M. (2010) *Removing bullwhip from the Tesco supply chain.* In *Production and Operations Management Society Annual Conference, May 7th – 10th.* Vancouver, Canada. Paper No. 015-0397, 19 pages.

[26] Prak, D., Teunter, R., and Riezebos, J. (2015) Periodic review and continuous ordering. *European Journal of Operational Research*, vol. 242(3), pp. 820–827.

[27] Shingo, S. (1989) *A Study of the Toyota Production System: From an Industrial Engineering Viewpoint.* Productivity Press, New York.

[28] Simon, H. A. and Holt, C. C. (1954) The control of inventories and production rates-a survey. *Journal of the Operations Research Society of America*, vol. 2(3), pp. 289–301.

[29] Sloan, A. P. (1963) *My Years with General Motors.* Doubleday, New York.