Production Lot Sizing with Immediately Observable Random Production Rate

Ilan Adler · Dan Bu · Philip Kaminsky

Dept. of Industrial Engineering and Operations Research

University of California, Berkeley, CA 94720

Abstract

To explore one impact of the information available by adding sensors in a classical production planning setting, we consider a continuous time infinite horizon lot-sizing model where a single product is manufactured on a single machine. Each time manufacturing restarts, a random production rate is realized, and production continues at this rate until the machine is shut down. Although the rate is random and chosen from an arbitrary set of random rates, it is known as soon as production starts due to the availability of sensor information, so this information could be used to determine when to stop production. However, we show that given the objective of minimizing either average cost per unit time or total discounted cost, it is optimal to produce up to the same inventory level regardless of the realized production rate. We also develop heuristics for the multi-product version of this model.

1 Introduction

One characteristic of modern manufacturing processes is the increasing prevalence of sensors. In both discrete and continuous process manufacturing, sensors provide real-time feedback of operating characteristics and system performance, facilitating monitoring, troubleshooting, adjustment, and quality control (see, e.g., **?**). In contrast, canonical operations research manufacturing models, from the economic production quantity model to manufacturing scheduling models, either assume that the problem is deterministic, or that quantities such as production time or production rate are stochastic, and that realizations cannot be measured until after the fact.

In contrast, we are interested in manufacturing setting where sensors allow a firm to resolve this uncertainty more rapidly, and in this paper, motivated by a manufacturing technology used in industries such as biopharmaceutical manufacturing, we consider a variant of a classic operations research production model. Specifically, we consider a production planning problem faced by a firm that meets constant deterministic demand by producing a product on a single machine. We focus on a setting where the production rate on that machine is random and varies from production cycle to production cycle, but is known immediately after the cycle starts. The firm must determine a production strategy in order to minimize setup cost and holding cost.

Of course, if production rate is known and deterministic, this is the well-known economic production quantity (EPQ) model, and the production quantity is easy to find. This suggests the following policy: for each cycle, since the production rate is known as soon as production starts, produce the quantity that would be optimal for this production rate if this was the single deterministic production rate. Surprisingly, this is not the optimal policy for this problem. Indeed, we consider several variants of this setting, with both average and discounted costs, and we show the same surprising result for each case: for any problem instance, it is optimal to produce up to the same level each time, **independent of the realized production rate in that cycle**. In other words, although we are able to observe the production rate immediately after the start of production, we do not alter the level that we are producing up to account for this information. Once again, this is true even though given an instance of this problem with set of possible production rates, if any of those rates was the unique (deterministic) production rate (so that we had a variant of the traditional economic production quantity model), the optimal produce-up-to level would be different depending on the rate, and the optimal cost of operating the system would also be a function of the production rate. Nevertheless, we are able to fully characterize the optimal solution in the average cost case, although to do so we have to develop new proof approaches, as the traditional calculus-based approaches are insufficient for this problem variant.

Our work is related to random yield production planning models, but the majority of random yield production planning models consider settings in which production decisions are made, batches are manufactured, and production yield (and thus production quantity) are determined after manufacturing. We are motivated, however, by a manufacturing technology used in biopharmaceutical manufacturing, called continuous perfusion, where the production yield per unit time –the effective production rate– is random, but can be discovered soon after the start of a production cycle. In traditional biomanufacturing, the initial production step, fermentation, is completed in batches. After a traditional fermentation batch, the yield of the batch can be measured. Perfusion, in contrast, can be viewed as a continuous production run divided into "batches," or production cycles, with a given maximum length. Product is harvested continuously, so the production rate –the rate at which final product is produced or harvested, which is called the "yield rate" in the industry– can be estimated from the start of the batch (or more accurately, the rate curve can be estimated, since in contrast to our model, in practice the yield rate increases, then is stable for a period, and then decreases over the processing time of the batch), and product is collected at that rate during the time that the batch is processed. Note that this is called "yield" because the volume of process output collected per unit time is constant, but the concentration of good product per unit volume varies from batch to batch. Note also that the cost per batch is roughly proportional to a fixed cost per batch, plus a linear function of the size of the batch, since the majority of cost is accrued during downstream processing after fermentation, so there is no per unit cost savings for higher yield rate batches.

This setting gives rise to a variety of interesting production planning issues, and the model we are focusing on in this paper, where production rate is random but constant over the life of a single batch, and known immediately after the start of production, captures a highly stylized version of one of these issues. Of course, the actual problem described above is much more complicated than the setting we consider, since we do not model the stochastic yield curve, with its implied maximum batch length. Nevertheless, our model does capture a key feature of this setting, random production rate, and it demonstrate how in-process sensors might change (and not change) the way that we think about even simple production planning problems. In the next section, we briefly review relevant literature. After that, in Section 3 we introduce our model, and in Section 4 we fully characterize the optimal policy in an average cost setting. In Section 5 we characterize the optimal policy for some versions of our model with discounted costs – surprisingly, the optimal policy is qualitatively the same in this setting. Finally, in Section 6, we briefly introduce a few heuristics for some multi-product versions of this problem.

2 Literature Review

A considerable amount of research focuses on a variety of types of random yield. Random yield can be categorized in a variety of ways; for our purposes, we divide random yield models into imperfect production processes (IPP), in which output is a random function of input, and stochastic production rate models (SPR), such as the models we are focusing on in this paper. The bulk of the related literature has focused on IPP while relatively little attention has been paid to SPR. The "random yield" in imperfect production processes is a result of uncertainty in the relationship between the quantity received or made and the quantity requisitioned, particularly in batch-based manufacturing. The "random yield" in stochastic production rate settings is a result of production rates randomly evolving over time.

Researchers have proposed a variety of approaches to modeling the relationship between inputs and outputs in imperfect production processes. In their comprehensive review, **?** divide the modeling of imperfect production processes roughly into three categories: binomial yield, stochastically proportional yield, and interrupted geometric yield. The first assumes that every unit of production is independent of all other units, and that the creations of good units can be modeled by a Bernoulli process. Thus, the number of good units in a batch of size Q conforms to a binomial distribution. *Stochastically proportional yield* is generalization of the binomial case, and specifies the effective output distribution (or yield rate) with both the mean and variance. The distribution of the fraction of good units is independent of the batch size, but the yields of the individual units are perfectly correlated (as explained in **?**).

The two aforementioned approaches focus on the output distribution, while the *interrupted geometric* model captures a production setting in which the time until a process goes "out of control" is geometric. All units produced prior to this point are assumed to be acceptable and all subsequent units are assumed to be defective. **?** address the problem of the traditional Economic Lot Scheduling Problem with imperfect production. They point out that although most production processes start from an "in-control" state, they may shift to an "out of control" state at a random time and produce defective items until the next production cycle. **?** observe that unit production cost and process quality depend on the production rate, and they extend the model to cases where the production rate is a decision variable.

In contrast to much of this prior work, we are concerned with stochastic production rates, and there is relatively little literature with this focus. **?** study a production-inventory system under stochastic production and demand rates, model this system as a bivariate Markovian stochastic process and derive the limiting distribution of the inventory level. They show that the classical deterministic Economic Order Quantity (EOQ) policy minimizes the long-run average cost if one replaces the deterministic demand rate by the expected demand and production rate in steady state. Also, **?** study a production / inventory system where the unit production time is a random variable. As discussed above, we focus on continuous time models in which production rate is random but is known as soon as a given production cycle hs started.

3 The Model

We consider a continuous time production planning model, in which a single product is manufactured using a single machine that can be started and stopped as needed, in order to meet constant deterministic demand with rate D. We assume that each time the machine is started, production occurs at a random production rate, selected from one of $L + M$ distinct possible production rates, and that demand must be met without backorder. In particular, there are L production rates bigger than demand rate, $D < \mu_1 < \mu_2 < \ldots, m_L$, where μ_i occurs with probability p_i ($i = 1, \ldots, L$); and M production rates smaller than the demand rate, $\theta_1 < \theta_2 < \ldots, < \theta_M < D$ where θ_j occurs with probability q_j ($j = 1, \ldots, M$), such that $\sum_{i=1}^{L} p_i + \sum_{j=1}^{M} q_j = 1$. When the machine is idle, we say that it operates at rate $\theta_0 = 0$.

The cost of production is a constant c per unit regardless of the production rate (as we discussed in the introduction, motivated from our bioproduction example by the fact that expensive downstreaming processing is a function of the number of units produced), and each time production starts, a positive setup cost K is incurred. In addition, inventory can be stored, and a positive inventory cost rate h per unit is charged. Our initial objective is to minimize (almost surely) average cost per unit time. Later, we demonstrate that our key results are robust to the details of the problem setting, by considering the discounted cost objective (which requires a completely different analysis approach, but arrives at a qualitatively similar result).

Our goal in each of these cases is to derive optimal policies, π^*, that determine when to start and stop production in order to minimize production, holding, and fixed costs. We first consider the average cost case, and note that in general, even determining the structure of the optimal policy is non-trivial – especially with potential production rates less than the demand rate, it isn't obvious that a zero-inventory-ordering policy is optimal, so traditional calculus-based approaches to determining optimal production quantities are insufficient. Nevertheless, in the next section, we are able to completely characterize the devise a simple efficient algorithm to obtain a policy that minimizes (almost surely) the expected cost over the infinite horizon.

4 Minimizing Average Cost

As discussed in the previous section, the structure of the optimal policy for this problem is non-trivial to determine, so direct formulation of the problem as a mathematical program leads to a complicated optimization problem. Therefore, our strategy in this section is to first develop several few key properties that are satisfied by any optimal solution, then to introduce and analyze a simpler problem, and then ultimately to connect this simpler problem to the original problem. Using this approach, we arrive at a straightforward (and efficiently computed) optimal policy.

To facilitate the presentation we introduce the following notation:

- $\hat{\mu}_i \triangleq \frac{\mu_i}{(\mu_i - D)}$.
- $\hat{\theta}_j \triangleq \frac{\theta_j}{(D - \theta_j)}$.
- c^* is the almost sure average cost per unit time under the infinite horizon optimal policy π^*.

This notation significantly simplifies subsequent exposition by relating the rate of production to the length of certain simple cycles.

We begin by proving a key property of an optimal solution to this problem – **we always produce-up-to the same inventory level, regardless of the realized production rate**. As we observed in the introduction to the paper, we find this result quite surprising: although we are able to observe the production rate immediately after the start of production, we do not alter the level that we are producing up to account for this information. Again, this is true even though the deterministic produce-up-to level would in general be different for each possible production rate if that were the only rate, and thus, in our optimal solution, the fixed cost for a given cycle divided by the length of that cycle is explicitly not equal to the total holding cost for that cycle divided by the length of the cycle, **although this is in fact an optimality condition for the single rate deterministic case.**

Specifically, we both show below that it always optimal to produce up to the same level, and we derive an expression for this level as a function of the optimal cost per unit time.

Theorem 4.1 *Let* $I^* = \frac{c^*}{h}$*. Suppose the current inventory level is in* $[0, I^*]$*. If the current production rate is above* D*, then it is optimal to continue with that rate until the inventory level reaches* I^* *and then to immediately stop production.*

Observe that this theorem implies an optimal policy for production rates greater than the demand rate, and characterizes the relationship between the optimal cost per unit time and the order-up-to level, but it has nothing to day about operating the operating policy for rates less than the demand rate, or about the long run average optimal cost. Nevertheless, we are able to prove this theorem without knowledge of either the full structure of the optimal policy, or the cost of that policy.

Proof of Theorem 4.1. The theorem follows from the following two claims:

(i) Suppose that the system is producing at any rate μ_i $(i = 1, \ldots, L)$, and that the inventory level is $0 \leq I_i < I^*$. Then, policy π^* calls for continuation of production up to an inventory level of at least I^*.

(ii) Let $\hat{I}$ be the largest inventory level that policy π^* ever reaches. Then, $I^* \geq \hat{I}$.

We now prove these two claims:

(i) Suppose, to the contrary, that in an optimal policy π^*, production at rate μ_i stops at inventory level $0 \leq I_i < I^*$. Consider an alternative policy, $\hat{\pi}$, identical to π^* except for the following modifications:

- Whenever the system is producing at rate μ_i and it reaches the inventory level I_i, rather than stopping, production continues until the inventory level I^*, and then production stops.
- Production is idle until the inventory level falls to I_i.

- At that point, policy π^* is resumed.

Note that the average inventory cost (denoted by c_Δ) over the interval of time where policy $\hat{\pi}$ deviates from policy π^* is $c_\Delta = h(I_i + \frac{1}{2}(c^* - I_i)) = \frac{1}{2}(I_i + c^*) < c^*$. However, this implies that the average cost over the infinite horizon of policy $\hat{\pi}$ is a weighted average of c_Δ and c^*, so it is strictly smaller than c^*, which contradicts the optimality of π^*.

(ii) Recall that $\hat{I}$ is the largest inventory level that policy π^* ever reaches. This means that by definition, production stops when the inventory level reaches $\hat{I}$, and at this point the inventory level decreases for some period of time (either because there is no production, or because production is started with a production rate less than the demand rate). Let $\epsilon > 0$ be smaller than the smallest drop of the inventory level from $\hat{I}$ before policy π^* either starts production if it was idle, or stops production otherwise. In addition, let ϵ be sufficiently small so that there is no change of production rate when the inventory level is in the range $(\hat{I} - \epsilon, \hat{I})$, and that $\hat{I} - \epsilon > I^* = \frac{c^*}{h}$. Now, consider an alternative policy $\hat{\pi}$ that is identical to π^* except for the following modifications:

- Production stops at inventory level $\hat{I} - \epsilon$ instead of level $\hat{I}$.
- At that point, the action which is prescribed by policy π^*, whenever production stops at inventory level $\hat{I}$, is followed.

Note that the average inventory cost (denoted by c_Δ) over the interval of time where policy $\hat{\pi}$ deviates from policy π^* is $c_\Delta = h(\hat{I} - \epsilon + \frac{\epsilon}{2}) > hI^* = c^*$. However, the average cost over the infinite horizon of policy π^* is a weighted average of c_Δ and the average cost of policy $\hat{\pi}$, so since the average cost of $\hat{\pi}$ is at lease as high as c^* (as π^* is optimal), we get that the average cost of policy π^* (which is c^*) is bigger than c^*, a contradiction. □

As a consequence of the preceding theorem, we can charcterize the structure of an optimal policy by considering the following generica cycle. A cycle starts when the inventory level reaches I^*, where there are two possible courses of action: either production is not restarted, and the inventory level starts to fall, or production is restarted, possibly repeatedly, until an acceptable production rate less than the demand rate is realized. Thereafter, as inventory levels decrease, and as long as the inventory level is above zero, the optimal policy can call for a stop followed be redrawing a new production rate. Unless this rate is bigger than the demand rate (when the optimal policy requires using this rate until the inventory level reaches I^*, ending the cycle), the process is the same as at the beginning of the cycle, choosing to either depleting the inventory while the production is idle, or selecting a desired production rate smaller than the demand rate. Eventually (unless a production rate bigger than the demand rate has been drawn), the inventory level reaches zero, thereafter the first drawn production rate which is bigger than the demand rate is selected, and used without interruption until the invrntory level reaches level I^*, where the cycle ends.

Remarks

- Note that if a production rate bigger than the demand rate is drawn at the beginning of the cycle (when the inventory level is I^*), the cycle length is zero.
- Since the probability of drawing a production rate greater than the demand rate is $\sum_{i=1}^{L} p_i$, the expected fixed cost per cycle for any policy that meets the conditions of Theorem 4.1 must be $\frac{K}{\sum_{i=1}^{L} p_i}$. In the remaining of this section we'll denote $\hat{K} \triangleq \frac{K}{\sum_{i=1}^{L} p_i}$.

Finally, we note that our goal of finding a policy that minimizes the infinite horizon average cost can be achieved by simply finding a policy that minimizes the average cost over the cycle presented above. In particular, let T_n^π be the total cost accrued under some stationary policy π during the n-th cycle, and define the renewal reward process $T^\pi(t) = \sum_{n=1}^{N^\pi(t)} T_n^\pi$ associated with the renewal processes, $N^\pi(t),\ t \geq 0$, corresponding to the random variable ℓ_n^π measuring the length of the n-th cycle of policy π. Then, defining $\bar{T}^\pi$ and $\bar{\ell}^\pi$ as the expected total cost and length of a cycle under policy π, and assuming both are finite, we have by the classical Renewal Theorem (see, e.g, **?** Proposition 7.3), that with probability 1, $\lim_{t\to\infty} \frac{T^\pi(t)}{t} = \frac{\bar{T}^\pi}{\bar{\ell}^\pi}$. Hence our task is to find a policy π that minimizes the expected total cost per cycle divided by the expected length of a cycle.

Specifically, an optimal policy needs to determine the optimal order-up-to level I^*. In addition, whenever the level of inventory is below I^*, and a production rate bigger than the demand rate is present, the optimal policy (by Theorem 4.1) calls for continuing production until the end of the cycle when the inventory level reaches I^*. It follows that whenever the inventory level is 0, any production rate smaller than the demand rate is rejected. On the other hand, whenever the level of inventory is greater than zero and the production is either idle or is of rate smaller than the demand rate, a policy should specify whether to stop and then restart production (with another drawn production rate), or possibly let the inventory goes down by the demand rate by not restarting the production. Fortunately, the following observations, which are summarized in a seies of lemmas below, considerably restrict the set of policies that could be optimal. In fact, we use these observations to identify an easily computes simple optimal policy which is presented in Theorem 4.9. For the rest of this section we refer to idle production as production rate θ_0, that is $\theta_0 = 0$ (so $\hat{\theta}_0 = 0$).

Lemma 4.2 *Suppose it is optimal to stop production (at positive inventory level) with rate θ_i $(i = 0, \ldots, M-1)$. Then, it is optimal to resume production (with rate smaller than the demand rate) only with rate θ_k, $(k = i+1, \ldots, M)$.*

Proof. Suppose the production θ_k $(k \leq i)$ is drawn following stopping production rate θ_i (when the inventory level is positive). Obviously the decision to resume production once θ_k is drawn, contradicts the optimality of the decision to stop production while the superior (or same) rate θ_i was available without additional cost. □

As a consequence of this lemma, we have the following corollary establishing the existence of an optimal policy that involves a unique stopping point for each rate θ_i.

Corollary 4.3 *For every $i = 1, \ldots, M$ there exists an optimal inventory level $\alpha_i^* I$ (where $1 \geq \alpha_0^* \geq \alpha_1^* \ldots, \geq \alpha_M^* \geq 0$) such that if production rate θ_i is active at inventory level greater than $\alpha_i^* I$, the production stops at inventory level $\alpha_i^* I$.*

Remark $\alpha_0^* = 1$, calls for drawing production rates until an acceptable rate less than the demand rate is drawn. On the other habd, $\alpha_0^* = 0$, (once the inventory level reaches I^*) calls for the inventory level to decrease, without drawing a production rate. Note the tradeoff regarding the choice of α_0^*. If the next draw is a production rate smaller than the demand rate, $\alpha_0^* = 1$ would have been the better choice as this would prolong the cycle, thereby reducing the average fixed cost per cycle. However, if the next draw is a production rate bigger than the demand rate, $\alpha_0^* = 0$ would have been the better choice as chosing $\alpha_0^* = 1$ would result with a zero lenght cycle, increasing the average fixed cost per cycle.

Thus, our problem is to determine an order-up-to level I, and a policy $\pi = (\alpha_0, \alpha_1, \ldots \alpha_M)$ minimizing the expected holding cost per unit time over a cycle as defined above. We denote by Π the set of all such policies π.

To facilitate the presentation of an optimal policy we introduce the following notation: Let

- $hH(\pi, I)$ be the expected holding cost per cycle, given policy $\pi \in \Pi$ and stop-production inventory level I.
- $L(\pi, I)$ be the expected cycle's length, given policy $\pi \in \Pi$ and stop-production inventory level I.

Our problem can be stated as

$$\mathcal{P} : \min_{\pi \in \Pi, I \geq 0} \frac{\hat{K} + hH(\pi, I)}{L(\pi, I)},$$

Next, we exploit the preceding characterization of an optimal solution to reduce $\mathcal{P}$ to the much simpler problem of maximizing the expected length of a cycle, a problem that (as we will show) can be solved very effectively in linear time. To proceed with the reduction, we first introduce the following two lemmas.

Lemma 4.4 *For any* $\pi \in \Pi$, $L(\pi, I) = IL(\pi, 1)$ *and* $H(\pi, I) = I^2 H(\pi, 1)$.

Proof. This result follows the observation that changing I to γI ($\gamma > 0$) can be considered as change of units of measurement of inventory level which obviously leads to scaling time by γ and inventory cost for any interval of time by γ^2.

Lemma 4.5 $hH(\pi^*, I^*) = \hat{K}$.

Proof. Let π^*, I^* be an optimal solution to problem $\mathcal{P}$. Then, by Lemma **??**

$$I^* = \arg\min_{I \geq 0} \frac{\hat{K} + hH(\pi^*, I)}{L(\pi^*, I)} = \arg\min_{I \geq 0} \frac{\hat{K} + I^2 hH(\pi^*, 1)}{IL(\pi^*, 1)}.$$

Noticing that the objective function above is convex in I, taking the derivative with respect to I, and equating to zero, we get that I^*, satisfies

$$\frac{1}{L(\pi^*, 1)} \left(-\frac{\hat{K}}{(I^*)^2} + hH(\pi^*, 1) \right) = 0$$

which leads to $hH(\pi^*, I^*) = (I^*)^2 hH(\pi^*, 1) = \hat{K}$. □

Next, we show that it is possible to solve $\mathcal{P}$ by solving the following problem:

$$\hat{\mathcal{P}} : \max_{\pi \in \Pi} L(\pi, 1) \text{ subject to } L(\pi, 1) = 2H(\pi, 1). \tag{1}$$

Lemma 4.6 *Suppose π^* is an optimal solution to $\hat{\mathcal{P}}$. Then, $\pi^*, I^* = \sqrt{\frac{2\hat{K}}{hL(\pi^*,1)}}$ is an optimal solution to $\mathcal{P}$.*

Proof. Let π^*, I^* be an optimal solution to $\mathcal{P}$. Then, by Theorem 4.1 , we have

$$I^* = \frac{\hat{K} + hH(\pi^*, I^*)}{hL(\pi^*, I^*)},$$

which allows to write problem $\mathcal{P}$ as

$$\min_{\pi \in \Pi, I \geq 0} I \ \textit{subject to} \ I = \frac{\hat{K} + hH(\pi, I)}{hL(\pi, I)}.$$

So, by Lemma 4.4, it can be written as

$$\min_{\pi \in \Pi, I \geq 0} I \ \textit{subject to} \ I = \frac{\hat{K} + I^2 hH(\pi, 1)}{IhL(\pi, 1)}.$$

Rearranging terms, noting (by Lemma 4.5) that at optimality $\hat{K} = I^2 hH(\pi,^* 1)$, and (since I^2 is monotonically increasing in I) replacing I by I^2 as the objective function, we can replace $\mathcal{P}$ by the equivalent problem

$$\min_{\pi \in \Pi, I \geq 0} I^2 \ \textit{subject to} \ I^2 hL(\pi, 1) = 2\hat{K}, \ \ I^2 hH(\pi, 1) = \hat{K}.$$

Finally, by substituting $I^2 = \frac{2\hat{K}}{hL(\pi,1)}$ we can set the problem above as

$$\min_{\pi \in \Pi, I \geq 0} \frac{2\hat{K}}{L(\pi, 1)} \ \textit{subject to} \ L(\pi, 1) = 2H(\pi, 1) = \hat{K}, \ I = \sqrt{\frac{2\hat{K}}{hL(\pi, 1)}}.$$

Noting that one can first solve the problem above with respect to π when imposing only the first constraint, and then determine I^* as a function of π^*, completes the proof. □

We now proceed to solve a relaxed version of problem $\hat{\mathcal{P}}$, namely

$$\mathcal{Q} : \max_{\pi \in \Pi} L(\pi, 1)$$

In particular, in the following theorem we provide a simple and effectively computed optimal solution π^* to problem $\mathcal{Q}$ which satisfies the dropped constraint $L(\pi^*, 1) = 2H(\pi^*, 1)$ of problem $\hat{\mathcal{P}}$. Thus, considering Lemma 4.6, we get that $\pi^*, I^* = \sqrt{\frac{2\hat{K}}{hL(\pi^*,1)}}$ is an optimal solution to $\mathcal{P}$.

Theorem 4.7 *There exists an optimal solution $\pi^* = (\alpha_0^*, \ldots, \alpha_M^*)$ to problem $\mathcal{Q}$ such that*

(i) $\alpha_M^* = 0$,

(ii) *If for $k = 0, \ldots, M-1$, $\alpha_\ell^* = 0$ for $\ell = k+1, \ldots, M$, then α_k^* is equal to either 0 or 1.*

(iii) $L(\pi^*, 1) = 2H(\pi^*, 1)$.

Proof.

(i) Suppose we had realized production rate θ_M when the inventory level is 1, and that we stop production at inventory level α_M (where $0 \leq \alpha_M \leq 1$). Then, the expected length of the cycle (starting and ending at inventory level 1) is:

$$(1 - \alpha_M)\left(\hat{\theta}_M + \frac{\sum_{i=1}^L p_i \hat{\mu}_i}{\sum_{i=1}^L p_i}\right)$$

so $\alpha_M^* = 0$, as the term in parentheses is positive.

(ii) Now, suppose that $\alpha_\ell^* = 0$ for $\ell = k+1, \ldots, M$, and consider realized production rate θ_k when the inventory level is 1, and that we stop production at inventory level α_k (where $0 \leq \alpha_k \leq 1$). Then, the expected length of the cycle is:

$$\frac{\sum_{i=1}^L p_i}{\sum_{i=1}^L p_i + \sum_{j=\hat{k}+1}^M q_j}(1-\alpha_k)\left(\hat{\theta}_k + \hat{\mu}_i\right) + \frac{\sum_{j=\hat{k+1}}^M q_j}{\sum_{i=1}^L p_i + \sum_{j=\hat{k}}^M q_j}\left((1-\alpha_k)\hat{\theta}_k + \frac{\alpha_k \hat{\theta}_j}{\sum_{j=\hat{k}+1}^M q_j} + \frac{\sum_{i=1}^L p_i \hat{\mu}_i}{\sum_{i=1}^L p_i}\right)$$

$$= \text{Constant } + \alpha_k \left(\frac{\sum_{j=\hat{k}+1}^M q_j \hat{\theta}_j - \sum_{i=1}^L p_i \hat{\mu}_i}{\sum_{i=1}^L p_i + \sum_{j=\hat{k}+1}^M q_j} - \hat{\theta}_k\right).$$

So (setting $\alpha_k^* = 1$ in case of a tie),

$$\alpha_k^* = \begin{cases} 1 & \text{if} \quad \frac{\sum_{j=\hat{k}+1}^{M} q_j \hat{\theta}_j - \sum_{i=1}^{L} p_i \hat{\mu}_i}{\sum_{i=1}^{L} p_i + \sum_{j=\hat{k}+1}^{M} q_j} - \hat{\theta}_k \geq 0 \\ 0 & \text{if} \quad \frac{\sum_{j=\hat{k}+1}^{M} q_j \hat{\theta}_j - \sum_{i=1}^{L} p_i \hat{\mu}_i}{\sum_{i=1}^{L} p_i + \sum_{j=\hat{k}+1}^{M} q_j} - \hat{\theta}_k < 0 \end{cases} \tag{2}$$

is optimal for problem $\mathcal{Q}$.

(iii) Note that (starting at inventory level 1) in any realized cycle, the optimal policy depletes the inventory in a constant rate (θ_j for some $j = 0, 1, \ldots, M$) until reaching 0 inventory level, thereafter increasing the inventory level by another constant rate (μ_j for some $j = 1, \ldots, L$). Thus, the graph of the level of inventory over a cycle is in the form of two consecutive triangles of height I^* intersecting at inventory level 0, implying that $L(\pi^*1) = 2H(\pi^*, 1)$.

□

Corollary 4.8 *There exists $k^* \in \{0, 1, \ldots, M-1\}$ such that*

$$\alpha_j^* = \begin{cases} 0 & \text{if } \ j \geq k^* \\ 1 & \text{if } \ j < k^* \end{cases} \tag{3}$$

is an optimal solution to problem $\mathcal{Q}$.

Proof. Follows directly from Corollary 4.3 and Theorem 4.7.

Remark Note that the structure of the optimal solution for problem $\mathcal{Q}$ implies that if $k^* = 0$, then it is always optimal not to start production unless the inventory level reaches 0. However, if $k^* \geq 1$, then (upon reaching inventory level 1) it is optimal to keep drawing (and possibly rejecting) production rates until a production rate θ_j with $j \geq k^*$ is drawn. Thereafter, the production doesn't stop until the inventory is completely depleted. Once at 0 inventory level, production resumes with the first drawn production rate bigger than the demand rate. Then, by Theorem 4.1, production stops when reaching inventory level of 1, completing the cycle.

The next theorem is the main result of this section.

Theorem 4.9 *For $k = 0, 1, \ldots, M$ let*

$$f(k) \triangleq \frac{\sum_{j=k+1}^{M} q_j \hat{\theta}_j - \sum_{i=1}^{L} p_i \hat{\mu}_i}{\sum_{i=1}^{L} p_i + \sum_{j=k+1}^{M} q_j} - \hat{\theta}_k \ \ (\textit{where} \ \sum_{j=M+1}^{M} q_j \hat{\theta}_j = \sum_{j=M+1}^{M} q_j = 0),$$

and let k^* be the largest k for which $f(k) < 0$.

(a) If $f(0) < 0$, then $\alpha_j^* = 0\ (j = 0, 1, \ldots, M)$,

$$I^* = \sqrt{\frac{2K}{h \sum_{i=1}^{L} p_i \hat{\mu}}}$$

is an optimal policy for problem $\mathcal{P}$.

(b) If $f(0) \geq 0$, then $\alpha_j^* = \begin{cases} 0 & \text{if } j \geq k^* \\ 1 & \text{if } j < k^*, \end{cases}$

$$I^* = \sqrt{\frac{2K(\sum_{i=1}^{L} p_i + \sum_{j=k^*}^{M} q_j)}{h\big[(\sum_{i=1}^{L} p_i)(\sum_{j=k^*}^{M} q_j \hat{\theta}_j) + (\sum_{j=k^*}^{M} q_j)(\sum_{i=1}^{L} p_i \hat{\mu}_i)\big]}}$$

is an optimal policy for problem $\mathcal{P}$.

Proof. Recall that $\pi^* = (\alpha_0^*, \alpha_1^*, \ldots, \alpha_M^*)$ in the statement of the theorem is optimal for problem $\mathcal{Q}$, so by Lemma 4.6 it solves problem $\hat{\mathcal{P}}$ (see (1)) as well. Finally, by Theorem 4.9 , we conclude that π^*, I^* is an optimal solution to $\mathcal{P}$. □

Following the preceding theorem we construct the following simple (and very efficient) algorithm to solve problem $\mathcal{P}$.

Algorithm 4.10

`Solve`$\mathcal{P}$

Input: $\hat{\mu}_1, \ldots, \hat{\mu}_L,\ \hat{\theta}_1, \ldots, \hat{\theta}_M,\ p_1, \ldots, p_L,\ q_1, \ldots, q_M,\ D$

Setup: *For* $k = 0, \ldots, M-1$ *set* $\alpha_k^* = 1,\ \alpha_M^* = 0$

Repeat *until* $\alpha_k^* = 1$ *or* $k = -1$

$k = k - 1,\ \alpha_k^* = 0$

If $\frac{\sum_{j=\hat{k}+1}^{M} q_j \hat{\theta}_j - \sum_{i=1}^{L} p_i \hat{\mu}_i}{\sum_{i=1}^{L} p_i + \sum_{j=\hat{k}+1}^{M} q_j} \geq \hat{\theta}_k,\ \alpha_k^* = 1$

$k^* = k$

If $k^* = 0,\ I^* = \sqrt{\frac{2K(\sum_{i=1}^{L} p_i + \sum_{j=k^*}^{M} q_j)}{h\big[(\sum_{i=1}^{L} p_i)(\sum_{j=k^*}^{M} q_j \hat{\theta}_j) + (\sum_{j=k^*}^{M} q_j)(\sum_{i=1}^{L} p_i \hat{\mu}_i)\big]}}$

If $k^* > 0,\ I^* = \sqrt{\frac{2K}{h \sum_{i=1}^{L} p_i \hat{\mu}}}$

Output: $\alpha_0^*, \alpha_1^*, \ldots, \alpha_M^*,\ I^*$

Using the output from the algoritm above, and assuming that the inventory level at the beginning is 0, we present below a summary of an optimal policy.

1. Start production with the first drawn production rate bigger than the demand rate, stop production when inventory level reaches I^*.

2. If $\alpha_0^* = 0$, keep the production idle; else, realize the first drawn production rate smaller than the demand rate (say θ_k) for which $\alpha_k^* = 0$. When the inventory level reaches 0, stop the production if it is on.

3. Realize the first drawn production rate bigger than the demand rate. When the inventory level reaches I^*, stop the production.

4. Goto step 1.

We conclude this section with several observations.

- We assume that all the available production rates are different than the demand rate D. Otherwise, it is obviously optimal to reject production rates (paying K each time) until the production rate which is equal to D is realized, and then to produce at rate D forever.

- It is interesting to note that we use all production rates greater than the demand rate, but we are selective about using production rates smaller than the demand rate.

- If there are no production rates smaller than the demand rate, the optimal policy is the same as when $\alpha_0^* = 0$,

4.1 Backorder Allowed

In this subsection we extend the model to allow backorder with a positive penalty cost rate of b. We show that, as in the case where no backrder is allowed, we still obtain an optimal solution with a simple structure that uses only one production rate whenever the inventory is at a level strictly between two values $-B^*$ and I^*, which are easy to calculate. We prove our assertion by closely following arguments in the previous section.

- Theorem 4.1 is still valid.

- Defining $B^* \triangleq \frac{c^*}{b}$, we establish the following key result.

Theorem 4.11 *Let $\hat{B}$ be the largest backorder level that policy π^* ever reaches. Then, $B^* \geq \hat{B}$.*

Proof. Suppose that to the contrary, $\hat{B} > B^*$, and consider an alternative policy $\hat{\pi}$ that is identical to π^* except for the following modifications:

1. Production stops at inventory level B^* instead of level $\hat{B}$.
2. At that point, the action which is prescribed by policy π^*, whenever production stops at inventory level $\hat{I}$, is followed.

Note that the average inventory cost (denoted by c_Δ) over the interval of time where policy $\hat{\pi}$ deviates from policy π^* is bigger than c^*. However, the average cost per unit time over the infinite horizon of policy π^* is a weighted average of c_Δ and the average cost of policy $\hat{\pi}$, so since the average cost of $\hat{\pi}$ is at lease as high as c^* (as π^* is optimal), we get that the average cost per unit time of policy π^* (which is c^*) is bigger than c^*, a contradiction. □

- All the discussion and results are valid where $I^* + B^*$ replaces I as the inventory level moves between $-B^*$ and I^*. An essential observation supporting this assertion is that the length of a period where there is a constant production rate as a function of the difference between the inventory levels at the beginning and the end of the period is the same regardless whether the inventory level over the period is positive or negative.
- Noticing that $I^* = hc^*$ and $B^* = b$, which implies that $I^* = \frac{b}{h}B^*$, we get the following theorem which is the analogue of Theorem 4.9 where backorder allowed.

Theorem 4.12 *For $k = 0, 1, \ldots, M$ let*

$$f(k) \triangleq \frac{\sum_{j=k+1}^{M} q_j\hat{\theta}_j - \sum_{i=1}^{L} p_i\hat{\mu}_i}{\sum_{i=1}^{L} p_i + \sum_{j=k+1}^{M} q_j} - \hat{\theta}_k \;\; (where \sum_{j=M+1}^{M} q_j\hat{\theta}_j = \sum_{j=M+1}^{M} q_j = 0),$$

and let k^* be the largest k for which $f(k) < 0$.

(a) If $f(0) < 0$, then $\alpha_j^* = 0$ $(j = 0, 1, \ldots, M)$,

$$I^* = \sqrt{\frac{2K}{(h+b)\sum_{i=1}^{L} p_i\hat{\mu}}}, \;\; B^* = \frac{h}{b}I^*$$

is an optimal policy for problem $\mathcal{P}$.

(b) If $f(0) \geq 0$*, then* $\alpha_j^* = \begin{cases} 0 & \text{if } j \geq k^* \\ 1 & \text{if } j < k^*, \end{cases}$

$$I^* = \sqrt{\frac{2K(\sum_{i=1}^{L} p_i + \sum_{j=k^*}^{M} q_j)}{(h+b)\big[(\sum_{i=1}^{L} p_i)(\sum_{j=k^*}^{M} q_j \hat{\theta}_j) + (\sum_{j=k^*}^{M} q_j)(\sum_{i=1}^{L} p_i \hat{\mu}_i)\big]}}, \quad B^* = \frac{h}{b} I^*$$

is an optimal policy for problem $\mathcal{P}$*.*

5 Discounted Infinite Horizon

In the previous section, we derive the optimal inventory levels with an average cost model to minimize the total cost per unit time. Here, we address the discounted cost version of our model when all production rates are greater than the demand rate and backorder is not allowed. The optimal policy for other discounted versions of the model remains an open question.

? was among the first papers to consider a discounted version of the traditional EOQ problem over an infinite horizon. Later, this approach was adapted to the analysis of similar models in the presence of trade credit, permissible late payment (**?**, **?** and **?**), and deteriorating inventory (**?**).

In light of the long history of the EPQ problem, however, there is surpringly little published research exploring the discounted version of that model. **?** and **?** investigated replenishment policy under permissible delay in payments and cash discount within the EPQ framework. Perhaps the closest model to ours is found in **?**. The model in this paper is essentially a discounted version of the traditional EPQ model with a single production rate, and the authors explore the characteristics of the total cost when the interest rate is perturbed. However, their analytical results are primarily for limiting cases, when the production rate goes to infinity and the interest rate goes to zero.

In contrast, we consider the same model as in the previous section, with random production rates that are observed immediately after production starts, but here, the objective is to minimize expected discounted cost over an infinite horizon. We consider models both with and without backlogging, taking into account a penalty cost and a general discount rate $r > 0$. We focus on a a setting where $\mu_i > D$, for all $i = 1, 2, \ldots, L$ (recall that analysis of discounted cost models with production rates less than the demand rate, as well as models that allow backorder, remains an open question).

As in the case with the average cost objective, since there is no setup time, and since all production rates are greater than the demand rate, there is an optimal production strategy based on a zero-inventory producing policy, where production will not start while there is a positive inventory. For this analysis, we call the period between two consecutive zero inventory levels a cycle. Given the expected discounted cost objective, it is natural to model the problem as minimizing the total discounted cost over the infinite horizon as a renewal process. In particular, when the inventory level reaches zero, and upon observing the (random) production rate μ_i, production starts and continues until the inventory reaches an I_i level, Thereafter, the demand is satisfied from inventory until it runs out, where a new cycle begins. Note that the beginning of a cycle can be viewed as time 0. Thus, the optimal strategy can be characterized as $\boldsymbol{I}^* = (I_1^*, I_2^*, \ldots, I_L^*)$ where I_i^* is the optimal 'produce-up-to' inventory level when a production rate μ_i is observed at the beginning of a cycle. **Surprisingly, in spite of the fact that in this discounted version of the problem, the initial decision seems in some sense more heavily weighted, we are able to show in this section that the property that the optimal 'produce-up-to' inventory levels are all identical regardless of the observed production rate (that is $\boldsymbol{I_i^*} = I^*$, for $i = 1, \ldots, L$) is preserved even when the objective of minimizing the average cost is replaced by the objective of minimizing the expected discounted cost over the infinite horizon, even though in this setting a completely different proof approach is required.**

Recalling that $\tau_i = \frac{I_i}{\mu_i - D}$, $T_i = \tau_i + \frac{I_i}{D}$, the total discounted cost for a cycle starting at time 0 with production rate μ_i can be expressed as

$$f_i(I_i) \triangleq K + h\left\{\int_0^{\tau_i} (\mu_i - D)te^{-rt}dt + \int_{\tau_i}^{T^i} (-Dt + DT^i)e^{-rt}dt\right\} + c\int_0^{\tau_i} \mu_i e^{-rt}dt =$$

$$= K + \frac{1}{r^2}\left(\mu_i(cr+h) - hD + e^{\frac{-rI_i}{\mu_i - D}}(D - \mu_i - rI_i) + e^{\frac{-r\mu_i I_i}{\mu_i - D}}\left(D + e^{\frac{rI_i}{D}}(-D + rI_i)\right)\right).$$

Suppose that starting at the second cycle we use a strategy whose total expected discounted cost over the infinite horizon is S. Then, given that we have μ_i as the production rate in the first cycle, and using I_i as the level of inventory when production is stopped and never resumed until the inventory level is 0, the total expected discounted cost over the infinite horizon, starting at time 0, can be expressed as

$$g_i(I_i, S) \triangleq f_i(I_i) + e^{-rT_i}S = f_i(I_i) + e^{-r\frac{\mu_i I_i}{(\mu_i - D)D}}S.$$

Theorem 5.1 *Suppose $S > \frac{cD}{r}$. Then, for $i = 1, \ldots, L$, the unique solution $I_i(S)$ of the minimiza-*

tion problem $\min_{I_i \geq 0} g_i(I_i, S)$ *is*

$$I_i(S) = \frac{D}{r} ln \left(\frac{Dh + Sr^2}{Dh + Dcr} \right). \tag{4}$$

Proof. Observing that

$$\begin{aligned} \frac{\partial g_i(I_i, S)}{\partial I_i} &= \frac{-De^{\frac{rI_i}{D-\mu_i}}(h + cr)\mu_i + e^{\frac{rI_i\mu_i}{D^2 - D\mu_i}}\left(Dh + Sr^2\right)\mu_i}{Dr(D - \mu_i)} \\ &= \frac{\mu_i e^{-\frac{rI_i}{\mu_i - D}}}{Dr(\mu_i - D)}\left[D(h + cr) - e^{-\frac{rI_i}{D}}(Dh + Sr^2)\right], \end{aligned}$$

we get that the unique solution to the first order condition equation $\frac{\partial g_i(I_i,S)}{\partial I_i} = 0$, whenever $S > \frac{cD}{r}$, is $I_i(S)$ (see 4). Since D, h, c, r are all positive parameters,

$$\frac{\partial g_i(I_i, S)}{\partial I_i} < 0 \quad \text{for} \quad 0 \leq I_i < I_i(S), \quad \text{and} \quad \frac{\partial g_i(I_i, S)}{\partial I_i} > 0 \text{ for } \ I_i(S) < I_i,$$

so $I_i(S)$ is the unique global optimal point of $g_i(I_i, S)$. □

Now, for $S > \frac{cD}{r}$, let $F(S) = \sum_{i=1}^{L} p_i g_i(I_i(S), S)$. It is clear that the optimal value S^* of the total discounted cost of the model presented in this section has to satisfy $S^* = F(S^*)$. The next lemma is the key to showing that we can efficiently find S^* (to any level of approximation) by a binary search. As an input for such search, we need to identify a lower bound $\underline{S}$ for S^*, as well as an upper bound $\bar{S}$ for S^*. Observing that the discounted cost of producing continuously with production D (which is $\frac{cD}{r}$) is smaller than the discounted cost (when backorder is not allowed) of any policy; we have $\underline{S} = \frac{cD}{r}$. Since the discounted cost of any feasible policy is larger than S^*, we notice first that the discounted cost of (starting at time 0) continually producing at rate μ_i (which is feasible policy if μ_i is available) is $K - \frac{hD}{r^2} + \left(\frac{c}{r} + \frac{h}{r^2}\right)\mu_i$. Thus we get the following upper bound,

$$\bar{S} = \sum_{i=1}^{L} \left[K - \frac{hD}{r^2} + \left(\frac{c}{r} + \frac{h}{r^2}\right)\mu_i\right] p_i = K - \frac{hD}{r^2} + \left(\frac{c}{r} + \frac{h}{r^2}\right)\sum_{i=1}^{L} \mu_i p_i$$

Lemma 5.2 *Let* $\underline{S} = \frac{cD}{r}$ *and* $\bar{S} = K - \frac{hD}{r^2} + \left(\frac{c}{r} + \frac{h}{r^2}\right)\sum_{i=1}^{L} \mu_i p_i$.

(i) $F(\underline{S}) > \underline{S}$.

(ii) $F(\bar{S}) < \bar{S}$.

(iii) For $S > \underline{S}$, $0 < \frac{\partial F(S)}{\partial S} < 1$.

Proof.

(i) $F(\underline{S}) = F(\frac{cD}{r}) = \sum_{i=1}^{L} p_i g_i(I_i(\frac{cD}{r}), \frac{cD}{r}) = \sum_{i=1}^{L} p_i(K + \frac{cD}{r}) = K + \frac{cD}{r} > \frac{cD}{r} = \underline{S}$.

(ii) $F(\bar{S}) = F(K - \frac{hD}{r^2} + \left(\frac{c}{r} + \frac{h}{r^2}\right) \sum_{i=1}^{L} \mu_i p_i)$.

$F(\bar{S}) - \bar{S} = \sum_{i=1}^{L} -p_i(\mu_i - D)\frac{h+cr}{r^2} \left(\frac{h\mu_i + c\mu_i r + Kr^2}{Dh + crD}\right)^{-\frac{D}{\mu_i - D}} < 0$

$F(\bar{S}) < \bar{S}$

(iii) $\frac{\partial F(S)}{\partial S} = \sum_{i=1}^{L} p^i \frac{\partial g_i(I_i(S), S)}{\partial S} = \sum_{i=1}^{L} p^i \left(\frac{Dh + r^2 S}{Dh + cDr}\right)^{-\frac{\mu_i}{\mu_i - D}}$.

However, since $S > \frac{cD}{r}$

$$0 < \left(\frac{Dh + r^2 S}{Dh + cDr}\right) < 1$$

by the assumptions of the model

$$\frac{\mu_i}{\mu_i - D} > 0$$

Moreover, $0 \leq p_i \leq 1$. Thus

$$0 < \frac{\partial F(S)}{\partial S} < 1.$$

□

An immediate consequence of the preceding lemma is that the unique S^* satisfying $F(S^*) = S^*$ is the total cost of the optimal policy for the model with discounted cost and no backorders. The optimal “produce-up-to” level (at which point production stops until inventory level falls to 0) is

$$I_i(S^*) = \frac{D}{r} ln \left(\frac{Dh + S^* r^2}{Dh + Dcr}\right)$$

Thus, even in the discounted cost case, we still get the property that the optimal produce-up-to level is independent of the realized production rate μ_i.

Remark Although if we extend the model to allow backorder at a cost π per unit time, the optimal policy remains an open question, we observe that in this case it isn't immediately obvious that it is optimal to ever start production, nor is it clear that even if production is started, it is optimal to raise inventory to a positive level. If it is optimal to produce up to a positive level, we have shown that the optimal policy is simialr in structure to the ”no backorder” case.

6 Multiple Products Model

Our ultimate interest lies in the multi-product version of this single-machine lot sizing/sequencing problem. The research on related multi-product single-machine lot sizing and sequencing starts from the traditional Economic Lot Scheduling Problem (ELSP), which assumes a constant, predetermined production rate of perfect quality. Typically, costs include setup cost, production cost, and holding cost, and the goal is to determine a production strategy that minimizes long run average cost. In this section, we focus on the long run average cost, as is common in the literature, and on a setting where all production rates are greater than the demand rate and backorder is not allowed. In principle, similar heuristics could be developed for these alternative settings. For the ELSP (without setup times), the necessary and sufficient condition for a cyclic policy to be feasible is that the total production time does not exceed the total available time, i.e. $\sum_i \sigma_i/T_i \leq 1$, where σ_i is the processing time, and T_i is the cycle length (**?**).

? points out that contributions to the ELSP are typically either analytical approaches that achieve the optimum of a restricted versions of the original problem, or heuristics that achieve good solutions of the original problem. The most elementary approaches to the ELSP guarantee feasibility at the outset by imposing some constraints(s) on the cycle times, and then optimize individual cycle durations subject to the imposed constraints. Among these, two approaches seem most prevalent: the Common Cycle (CC) approach (**?**) and Basic Period (BP) approach (**?**). The CC approach first assumes a common cycle T that can accommodate the production of the required amount of each item exactly once, and then optimizes the cycle T^* such that the total cost per unit time is minimized. In contrast, the BP method admits different cycles for different items but constrains each cycle T_i of item i be an integer multiple n_i of a basic period W, where one basic period is long enough to accommodate the production of a single cycle of each of the items. Both of these approaches give a feasible upper bound on the ELSP problem – the BP method is less constrained, obviously leading to a tighter bound.

Our multi-product problem is equivalent to the Economic Lot Scheduling Problem (ELSP) (**?**) but with the addition of stochastic production rates. One alternative is to modify existing heuristics for this NP-hard (**?**) problem to account for the stochastic production rates. We present modified versions of the CC and BP approaches below. Note that in contrast to the single product case, these approaches need to make explicit use of the fact that one can observe the production rate.

Consider a setting with multiple products $i = 1, 2, \ldots n$, each with L^i possible production rates μ_{ij} with respective probabilities r^{ij}, $\sum_{j=1}^{L^i} r^{ij} = 1$. For each product i, there is a setup cost K_i, holding cost per unit time h_i and production cost per unit c_i. If we define $\rho^{ij} = \frac{D_i}{\mu_{ij}}$ and $\widetilde{\rho^i} = \frac{D_i}{\min_j\{\mu_{ij}\}}$, a sufficient condition for the existence of a feasible policy is $\sum_{i=1}^{n} \widetilde{\rho_i} \leq 1$.

We first present adaptions of CC and BP, and then present a novel heuristic based on our observations in the single product case.

6.1 Adapted Common Cycle Approach (ACC)

A classical approach from the literature, the Common Cycle approach, constrains the cycle length T to be the same for each product, where T can accommodate the production of each item at least once. We adapt the CC approach into our scenario. Note that if the condition $\sum_{i=1}^{N} \widetilde{\rho_i} \leq 1$ is satisfied, any T is feasible. Following the same development as in Section **??**, the total cost per unit time for product i is:

$$AC_i = \frac{K_i}{T} + h_i D_i (1 - E\left[\rho^i\right]) \frac{T}{2}$$

where $E\left[\rho^i\right] = \sum_{j=1}^{L^i} r^{ij} \rho^{ij}$, and thus total cost per unit time over all products is

$$\min_{T} \quad AC = \sum_{i=1}^{n} \left\{ \frac{K_i}{T} + h_i D_i (1 - E\left[\rho^i\right]) \frac{T}{2} + c_i D_i \right\} \tag{5}$$

which is convex in T. To minimize AC, we set its derivative with respect to T equal to zero, and obtain that $T^* = \sqrt{\frac{2 \sum_{i=1}^{n} K_i}{\sum_{i=1}^{n} h_i D_i (1 - \mathbb{E}[\rho_i])}}$.

Given T^*, $Q_i = D_i T^*$ of each product is sequentially produced, where the time between production starts for each product i is T^*, and the production time and produce-up-to level for product i depends on the realized production rate, i.e. $\tau^{ij} = D_i T^* / \mu_{ij}$.

6.2 Adapted Basic Period Approach (ABP)

Similarly, we can adapt the basic period heuristic. The basic period heuristic allows different cycle lengths for each product subject to the restriction that each cycle length has to be an integer multiple of a basic period W, i.e. $T_i = m_i W$. W is chosen so that it can accommodate production of each product, which guarantee feasibility (**?**).

Adapting BP for our problem and following the approach outlined above, the cost per unit time for item i:

$$AC_i(m_i, W) = \frac{K_i}{m_i W} + h_i D_i(1 - \mathbb{E}\left[\rho^i\right])\frac{m_i W}{2}. \tag{6}$$

where $AC_i(m_i, W)$ is a function of the cycle length $m_i W$. The best W and $\{m_1, m_2, ...\}$ for this heuristic are found by solving the following constrained optimization problem:

$$\begin{aligned} \min_{m_i, W} \quad & \sum_{i=1}^{n} AC_i(m_i, W) \\ \text{s.t.} \quad & \sum_{i=1}^{n} m_i \widetilde{\rho}_i \leq 1 \\ & m_i = 0, 1, 2 \cdots \end{aligned} \tag{7}$$

Note that as is typical for this type of approach, constraint (7) ensures that the total production time of all products cannot exceed W even at the slowest set of production rates.

Constraint (7) is sufficient but not necessary. Since any product i with $m_i > 1$ will not be produced in every base period (but instead in every m_i base periods), there is no need to have sufficient capacity in each base period to make a cycle's worth of each product. For the ELSP, **?** extended the Basic Period approach to account for this observation, and developed a systematic approach for generating a feasible schedule.

For details, see **?**. We adapt this heuristic – denoted as ABP-H – for our problem.

6.3 Produce-up-to the Same Level

In Sections 4 and 5, we show that for the single product model it is optimal (for average cost or discounted cost objectives) to raise inventory to a single target level independent of the observed production rate. We are thus motivated to develop a heuristic for the multiple-product case where inventory for each product is raised to a single product-specific maximum level independent of production rate. Implementing this approach, it is unnecessary to observe production rates when production starts – it is sufficient to identify the time at which inventory hits its maximum level. In the multiple product case, however, because it takes different amounts of time to produce up to a given level depending on the realized production rate, in general, a zero inventory ordering policy will not be feasible. We address this issue by developing a class of Fixed Idle Time (FIT) heuristics for this problem, in which we cycle through the production of each product in a given sequence, produce each product up to a single product-specific level regardless of the realized the production rate, and then insert a constant fixed amount of idle time into the schedule before restarting production of the next product in the sequence (so that in general, inventory level of a particular product will not be at zero when production of that product is restarted). In Appendix A, we present details of two variants of this FIT heuristic.

6.4 Computational Experiments

We completed a series of computational tests to compare the effectiveness of ACC, ABP-H, and FIT policy based on ABP-H. For ACC and ABP-H, we can easily assess the expected cost per unit time. This calculation is much more complicated for the FIT heuristic, however, so we use simulation to assess the cost of the FIT heuristic.

We complete a series of experiments in which we vary the fixed cost, holding cost and relative production rates, and compare the performance of the heuristics. The parameters are selected as follows: We have four products (indexed $i \in I = \{1, 2, 3, 4\}$ and two production rates ($j \in J = \{1, 2\}$). Demand is generated $D_i = 20 + 10i$, so demand varies as a function of product index i, and the probability of a given production rate j for product i, r^{ij}, is generated from the uniform distribution $U(0, 1)$, and rescaled such that $\sum_{j=1}^{2} r^{ij} = 1$. Fixed costs are randomly generated from $K_i \sim 100 + C \cdot U(0, 1) \cdot i$, where constant C is chosen from the set $\{100, 1000\}$, holding costs take on two values $h = \{1, 100\}$, and production rates are varied as follows: $\mu_{ij} = AD_i + Bj$, where A, B are constants, $A = \{4, 10, 100\}$, $B = \{10, 100\}$. Observe that A is a demand rate multiplier, while B controls the difference between production rates for a given product. These parameters imply that demand is selected from a range of $(20, 60)$ with increments of 10 between consecutive products, production rates are at least 4 times demand rates, thus ensuring feasibility (recall that

η	ACC	ABP-H	FIT
Average	1.0551	1.0422	1.2385

Table 1: Overall Performance of ACC, ABP-H and FIT

the feasibility condition for the multi-product single machine production problem is $\sum_{i=1}^{N} \frac{D_i}{\mu_{ij}} \leq 1$).

We can calculate the optimal production cycle time $\bar{T}_i^*$ of individual item – the optimal production cycle time if product i is the only product – and obtain a lower bound by summing corresponding individual product cost AC_i, so that

$$LB = \sum_{i=1}^{n} AC_i(\bar{T}_i^*).$$

Thus we test the performances of the three heuristics with respect to the lower bound for each possible combination of $(K_i, h, \mu_{ij}), i \in I, j \in J$, a total of $2 \times 2 \times 3 \times 2 = 24$ possibilities. For each combination (K_i, h, μ_{ij}), we consider 50 realizations of the random parameters, calculate the expected cost of applying ACC and ABP-H based heuristics for those realizations, and simulate for 100 periods the FIT heuristic.

We define the following performance measure for each heuristic:

$$\eta = \frac{Cost - LowerBound}{LowerBound}$$

which measures the percentage distance from the lower bound, where a lower η value indicates better heuristic performance. Averaging over all 24 combinations of parameter sets, Table 1 summarizes the heuristics' performance. Observe that for the selected parameters, ACC and ABP-H perform quite similarly on average while FIT doesn't perform as well. Next, we explore the impact of problem parameters on heuristic performance.

6.4.1 The Impact of Fixed Costs

To explore the impact of the magnitude of fixed costs on algorithm performance, we average across parameters except for fixed costs in Table 2. Observe that ACC and ABP-H perform similarly, although ABP-H seems to outperform ACC slightly, and as the setup costs become more distinct this becomes more apparent. Regardless of the setup costs, however, both heuristics outperform

η	$K_i \sim 100 + 100i \cdot U(0,1)$ ACC	ABP-H	FIT	$K_i \sim 100 + 1000i \cdot U(0,1)$ ACC	ABP-H	FIT
Avg.	1.0260	1.0244	1.2111	1.0843	1.0600	1.2659
Med.	1.0241	1.0230	1.2170	1.0674	1.0503	1.2559
Max	1.1126	1.1064	1.3954	1.3059	1.2789	1.6220
Min	0.9367	0.9366	1.0121	0.9526	0.9505	1.0415

Table 2: Statistics of η value under various fixed cost

η	$h = 1$ ACC	ABP-H	FIT	$h = 100$ ACC	ABP-H	FIT
Avg.	1.0565	1.0432	1.2358	1.0538	1.0412	1.2412
Med.	1.0481	1.0392	1.2382	1.0434	1.0340	1.2347
Max	1.2106	1.1900	1.5037	1.2080	1.1953	1.5136
Min	0.9364	0.9353	1.0334	0.9528	0.9518	1.0202

Table 3: Statistics of η value under various holding costs

FIT by about 20% on average, although they both need to make explicit use of knowledge of the production rates.

6.4.2 The Impact of Holding Cost

We average over parameters other than holding cost in Table 3. Holding cost seems to have little impact on heuristic performance.

6.4.3 The Impact of Production Rate

Finally, we explore the impact of different relative production rate, separating results by μ_{ij} values in Table 4. As the relative production rate increases, the performance of FIT approaches, or even exceeds, that of the other two heuristics, despite the fact that it doesn't require knowledge of production rates. We also explore the impact of disparity in different production rates in Table 5. From this table, we see that when $\mu_{ij} = 4D_i + Bj$, a bigger B will generally result in better performance of FIT, since FIT performs better when production rates increase relative to demand rate. When $\mu_{ij} = 10D_i + Bj$, the performance of FIT approaches that of ACC and ABP-H, and a bigger B leads to worse performance of FIT, while when $\mu_{ij} = 100D_i + Bj$, a bigger B leads to a better performance of FIT.

Overall, when production rates are relatively low, ACC and ABP-H outperform FIT. But when production rates are significantly higher than demand rates, FIT performs well, and provides a

	η	ACC	ABP-H	FIT
$\mu_{ij} = 4D_i + Bj$	Avg.	1.0529	1.0436	1.5695
	Med.	1.0491	1.0394	1.5855
	Max	1.2359	1.2752	2.0995
	Min	0.8750	0.8718	1.0709
$\mu_{ij} = 10D_i + Bj$	Avg.	1.0601	1.0435	1.0977
	Med.	1.0479	1.0370	1.0868
	Max	1.1998	1.1743	1.2380
	Min	0.9586	0.9586	1.0123
$\mu_{ij} = 100D_i + Bj$	Avg.	1.0524	1.0396	1.0482
	Med.	1.0402	1.0334	1.0370
	Max	1.1922	1.1284	1.1885
	Min	1.0002	1.0002	0.9973

Table 4: Statistics of η under various production rates

		$B = 10$			$B = 100$		
	η	ACC	ABP-H	FIT	ACC	ABP-H	FIT
$\mu_{ij} = 4D_i + Bj$	Avg.	1.0446	1.0362	1.7223	1.0612	1.0509	1.4168
	Med.	1.0416	1.0323	1.7533	1.0566	1.0465	1.4178
	Max	1.2664	1.3636	2.4065	1.2053	1.1868	1.7924
	Min	0.8160	0.8127	1.0912	0.9340	0.9309	1.0506
$\mu_{ij} = 10D_i + Bj$	Avg.	1.0550	1.0400	1.0706	1.0652	1.0470	1.1248
	Med.	1.0451	1.0319	1.0617	1.0507	1.0421	1.1119
	Max	1.1956	1.1768	1.2019	1.2040	1.1717	1.2742
	Min	0.9471	0.9471	1.0108	0.9702	0.9702	1.0138
$\mu_{ij} = 4D_i + 10j$	Avg.	1.0546	1.0397	1.0502	1.0502	1.0394	1.0462
	Med.	1.0403	1.0323	1.0366	1.0402	1.0346	1.0375
	Max	1.2416	1.1390	1.2388	1.1427	1.1178	1.1382
	Min	0.9997	0.9997	0.9947	1.0008	1.0007	0.9998

Table 5: Statistics of η under various B values

feasible yet simple production scheme, which is helpful especially when we can not observe realized production rates.

7 Conclusions and Future Research

In this paper, motivated both generally by a desire to better understand the impact of sensor-derived information in a classical production setting, and specifically by the perfusion processes employed in biotechnology manufacturing, we introduced a novel continuous time production model that captures a random production rate that is known as soon as a production cycle starts. We found that in the single product case, with both average and discounted cost objectives, even when this production rate is known, it is surprisingly optimal to produce up to one unique inventory level, and keep the same lowest back-order position if backorder allowed (although we were only able to prove the backorder result in the average cost case), regardless of the realized production rate. Motivated by this observation, for the multi-product case we propose a novel fixed idle time heuristic policy – FIT; we also adapt common heuristic approaches for ELSP in the literature to this setting, and compare ACC and ABP-H with our novel policy. While those policies outperform FIT in general, FIT is useful and relatively effective when we are unable to track the production rate (or have limited capacity for storage).

We are currently developing discrete time MDP models that capture more of the characteristics of perfusion production, and we are currently developing heuristics based on these MDP models for the multi-product case. Overall, modern manufacturing, with a vast amount of sensor-generated data, is a rich source of interesting problems, and we hope to address a variety of these in the future.

Acknowledgment

This work was supported by industry members of the I/UCRC Center for Excellence in Logistics and Distribution, and by the National Science Foundation under Grant No. 1067994.

A The FIT Class of Heuristics

In this appendix, we develop a multi-product policy where inventory for each product is raised to a single product-specific maximum level independent of production rate. Any such heuristic needs to address several key issues:

Determining the production sequence. Instead of sequentially producing each of the products, it may make sense to have a more complex production sequence, where some products are produced more frequently than others.

Determining the produce-up-to level for each product. For each product, a produce up to level must be selected so that even at the slowest production rate, there is time to produce up to the target inventory level before the inventory level of other products reaches zero.

Determining when to start each production cycle. In the multi-product case, since production must be started in time to ensure that production of other products can also be started in time to prevent stock-outs. Therefore, any feasible solution where inventories are raised to the same level for each product for each cycle might not be a zero inventory ordering policy. Furthermore, any such policy must determine the start time for each cycle.

In general, simultaneously optimizing all three of these decision parameters is extremely challenging – indeed, ensuring that a set of parameters leads to a feasible solution is a challenge. However, we can ensure feasibility by 1) adopting the production sequence and maximum inventory levels (given the slowest production rate) from either ACC or ABP-H; 2) by employing what we call a **fixed idle time policy** to determine production start times. We detail this approach below, first starting with the ACC based solution, and later the ABP-H based solution.

A.0.1 Fixed Idle Time Policy (FIT)

.

The ACC-based Approach

Our starting point for this solution is the ACC solution described in Section 6.1. Given this solution, we assess the following:

The production sequence: In the ACC approach, we cycle through the products, producing each product once in the cycle – we adopt the same approach in the ACC version of the FIT policy.

The produce-up-to level: Here, for each product, we produce up to the inventory level achieved during the slowest production rate in the ACC approach. To calculate this, recall that the optimal common cycle time is

$$T^* = \sqrt{2\sum_{i=1}^{n} K_i \Big/ \sum_{i=1}^{n} h_i D_i (1 - \mathbb{E}[\rho_i])}$$

Thus, the maximum inventory level for product i given production rate μ_{ij} is

$$H_{ij} = T^* \cdot \frac{D_i(\mu_{ij} - D_i)}{\mu_{ij}},$$

so for product i, we produce up to the inventory level

$$\Theta_i = \min_j \{H_{ij}\} \tag{8}$$

The production start time: For any ACC solution, production in a cycle can be arranged so that production of all products is sequential, and then there is some (possibly zero) idle time before production restarts. The length of this idle time will vary, depending on realized production rates during the cycle, and will be smallest when each product is produced at its slowest rate. We determine this minimum possible idle time, and in the FIT heuristic, insert this amount of idle time after producing each of the products once. To determine this, we calculate the maximum possible processing time based on the slowest rate for each product

$$\tau_{i,max} = \frac{\Theta_i}{\min_j \{\mu_{ij}\} - D_i}$$

The minimum possible idle time is therefore

$$\Delta = T^* - \sum_i \tau_{i,max} \tag{9}$$

Thus, in the FIT heuristic, we produce each product in turn up to level Θ_i, insert time Δ, and then start over. Note that this will not in general be a zero inventory producing policy.

In Figure 1 we illustrate for a two-product case (where each product has a slow and a fast

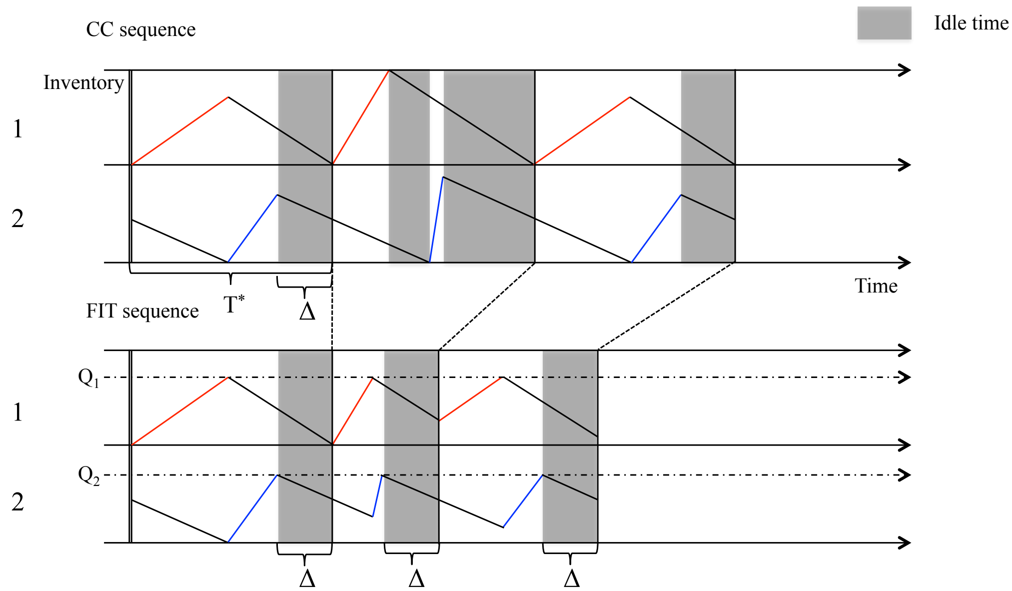

Figure 1: Sample production schedule of ACC and FIT

production rate) the ACC sequence and the corresponding FIT sequence. We illustrate a sample path where production rates in the first three periods are slow, fast, and then slow, and shade the time during which the machine is idle. Observe that the FIT sequence leads to some shorter production cycles, resulting in more frequent production. Given this schedule (which we argue in the Appendix will always be feasible), we can then search over possible T^* values to further reduce costs.

The ABP-H-based Approach

We can similarly adapt the Hassler version of the BP heuristic (ABP-H). Recall that the ABP-H solution consists of a basic period W and a set of integer multiples of the basic period $\boldsymbol{m} = (m_1, m_2, \cdots, m_n)$, where if $m_i = 1$, product i is produced every base period, if $m_i = 2$, product i is produced every second base period, if $m_i = 3$, product i is produced every third base period, etc. Starting from the ABP-H solution described in section 6.2, we can develop a version of the FIT heuristic as follows:

The production sequence: Here, we adopt the same production sequence as in the ABP-H heuristic, noting that depending on the multiplier, a product may appear more than one time in the sequence. For instance, in the three product case if $\boldsymbol{m} = (1, 1, 2)$, then the production sequence will be $1231212312\cdots$.

The produce-up-to level: Here, given basic period W, the corresponding maximum inventory levels are $H_{ij} = W \cdot \frac{D_i(\mu_{ij}-D_i)}{\mu_{ij}}$ $\forall i$, and thus the produce-up-to inventory levels are $\Theta_i = \min_j\{H_{ij}\}$.

The production start time: Recall that we can generate a production schedule in ABP-H from the production multipliers $\boldsymbol{m} = (m_1, m_2, \cdots, m_n)$. The ABP-H solution can be viewed as a series of subcycles making up a cycle, where each subcycle corresponds to a basic period, and the cycle corresponds to the time where the sequence restarts. Let $\mathcal{K}$ denotes the least common multiple of the m_i's, and let t denote the index of a sub-cycle where $t \in \{1, 2, \cdots, \mathcal{K}\}$. In any given subcycle t, for all of the products i that we produce in that subcycle, we produce up to Θ_i, and then append an idle time equal to

$$\Delta_t = W - \sum_i \tau_{i,max} \cdot \mathbb{1}_i^t \tag{10}$$

where the maximum possible processing time based on the slowest rate for each product is $\tau_{i,max} = \frac{\Theta_i}{\min\{\mu_{ij}\}-D_i}$, and binary parameter $\mathbb{1}_i^t$ is equal to 1 when we produce i in the subcycle t and 0 otherwise. Observe that in general we can have different fixed idle times in different subcycles.

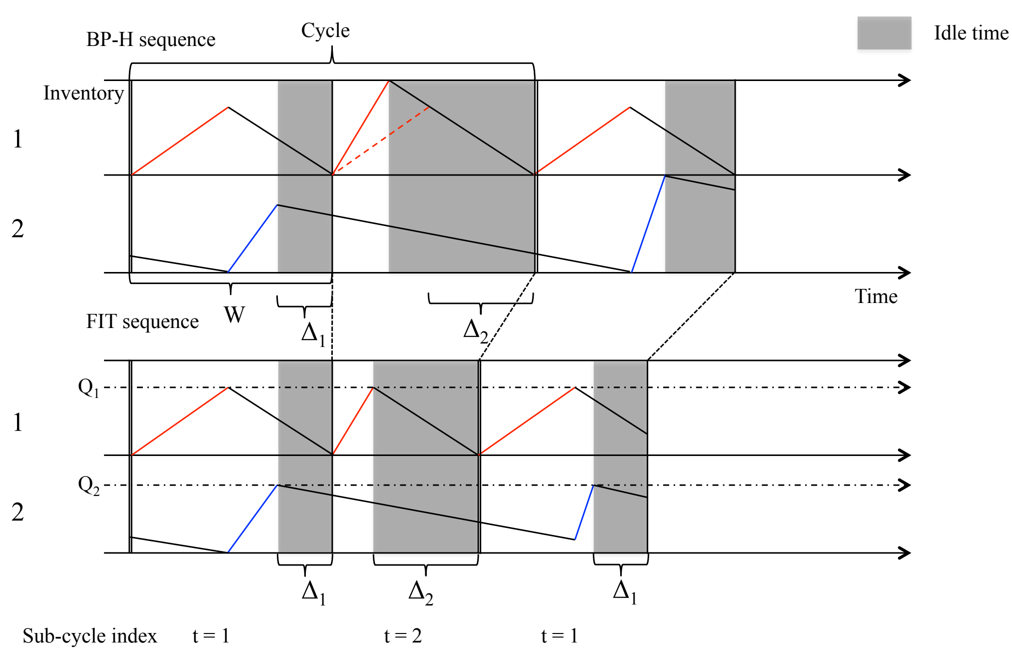

Figure 2: Sample production schedule of ABP-H and FIT

Thus, in this version of FIT, we produce the appropriate products in each subcycle up to level Θ_i, insert the appropriate idle time Δ_t given the subcycle we are producing, and then start the next subcycle. In Figure 2 we illustrate for a two-product case the ABP-H sequence and the corresponding FIT sequence. In this example, the vector of periods numbers is $\boldsymbol{m} = (m_1, m_2) = (1, 2)$, and the production rate sequence over the first three basic periods is {slow, fast, slow} for product 1, and is {slow, fast} for product 2. Note that the idle time vector $\vec{\Delta} = (\Delta_1, \Delta_2)$ can be pre-calculated.

Below, we prove that the FIT policy based on ACC or ABP-H is always feasible.

THEOREM C.1. *The FIT policy based on ACC or ABP-H is always feasible.*
Proof. To demonstrate the feasibility, it is sufficient to show that Θ_i *can always satisfy demand.* For any product i, a production cycle ends when inventory level equals Θ_i. For feasibility, inventory level Θ_i must be sufficient to meet demand for i from this *production-end* time until the next *production-start* time – we denote this interval R_i, and illustrate this in Figure 3, which is based on the example in Figure 2. Observe that R_i consists of two components: the realized processing

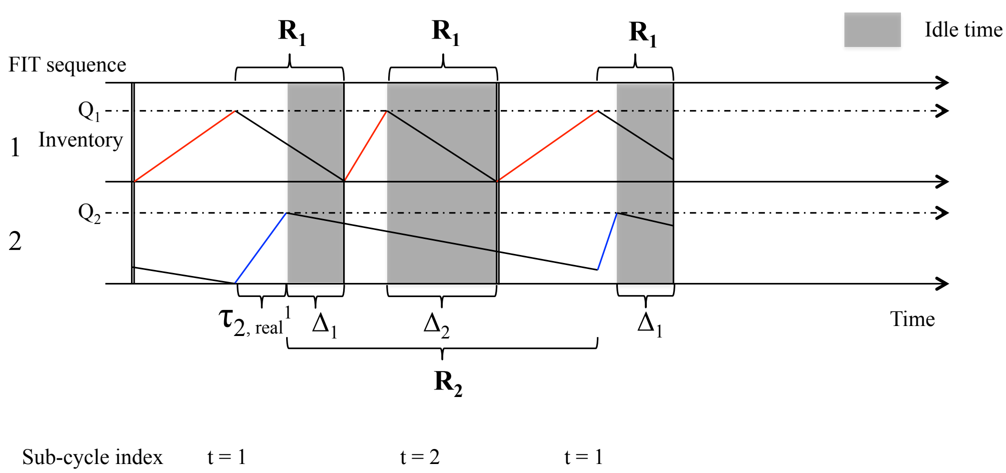

Figure 3: An example of R_2

time of other products, and the inserted idle time Δ_t. The realized processing time of product i in sub-cycle t, denoted as $\tau_{i,real}^t$, is a function of the realized production rate of i and the inventory level. Thus in this figure, $R_1 = \tau_{2,real}^1 + \Delta_1$ or $R_1 = \Delta_2$, depending on the sub-cycle.

To derive an expression for R_i, for any product i, denote the sub-cycle that contains the *production-end* time sub-cycle t', and the sub-cycle that contains the *production-start* time subcycle t''. Then

$$R_i = \begin{cases} \sum_{k=t'}^{t''} (\sum_{j \neq i} \tau_{j,real}^k \cdot \mathbb{1}_j^k + \Delta_k \cdot \mathbb{1}_\Delta^k) & \text{if } t' < t'' \\ (\sum_{k=t'}^{\mathcal{K}} + \sum_{k=1}^{t''})(\sum_{j \neq i} \tau_{j,real}^k \cdot \mathbb{1}_j^k + \Delta_k \cdot \mathbb{1}_\Delta^k) & \text{if } t' \geq t'' \end{cases} \tag{11}$$

where if $t' < t''$, R_i is within one single cycle, while if $t' \geq t''$ then R_i stretches over two cycles, and two indicators variables are:

$$\mathbb{1}_j^k = \begin{cases} 1 & \text{if } j \text{ is produced in sub-cycle } k \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbb{1}_\Delta^k = \begin{cases} 1 & \text{if } \Delta_k \text{ is inserted} \\ 0 & \text{otherwise} \end{cases}$$

Thus, a sufficient condition for feasibility is that:

$$\frac{\Theta_i}{D} \geq R_i \tag{12}$$

Finally, observe that *the sufficient feasibility condition is satisfied in FIT*. Since we adopt the Θ_i (lowest maximum inventory level) from either ACC and ABP-H, which both select Θ_i values so that demand can be satisfied even the slowest production rates are realized, so that $\forall i$,

$$\frac{\Theta_i}{D} \geq \begin{cases} \sum_{k=t'}^{t''} (\sum_{j \neq i} \tau_{j,max}^k \cdot \mathbb{1}_j^k + \Delta_k \cdot \mathbb{1}_\Delta^k) & \text{if } t' < t'' \\ (\sum_{k=t'}^{\mathcal{K}} + \sum_{k=1}^{t''})(\sum_{j \neq i} \tau_{j,max}^k \cdot \mathbb{1}_j^k + \Delta_k \cdot \mathbb{1}_\Delta^k) & \text{if } t' \geq t'' \end{cases}$$

And since the maximum possible processing time is no less than the realized processing time,

$$\tau_{j,max}^t \geq \tau_{j,real}^t$$

Therefore

$$\frac{\Theta_i}{D} \geq R_i$$

so the sufficient feasibility condition is satisfied in our new policy FIT. □