Balancing Learning and Economies of Scale: The Case of Adaptive Clinical Trials

Adam Fleischhacker • Yao Zhao

Department of Supply Chain Management & Marketing Sciences
Rutges Business School - Newark and New Brunswick
Rutgers, The State University of New Jersey
180 University Avenue, Newark, NJ 07102, USA
afleisch@pegasus.rutgers.edu • yaozhao@andromeda.rutgers.edu

Prior to the start of an adaptive clinical trial, demand for an investigational drug can be highly uncertain. Both recommended dosages and patient recruitment can fluctuate in response to early trial results. While initial demand forecasts can be very wrong, the factors influencing future demand can be learned during the trial. To take advantage of this learning, intra-trial batches can be produced, but at the expense of scale economies. Using various learning curves, we study this balance between learning and economies of scale in a finite horizon inventory model with fixed production costs and two production options: The pre-trial batch and the intra-trial batch. We characterize the optimal policy for both production batches in regards to optimally scheduling and sizing production. Through analytical and numerical studies, we develop insights on the impact of fixed costs, learning rates, and penalty costs on the value of the intra-trial batch, the timing of the intra-trial batch, and the size of the pre-trial batch.

1. Introduction

Before an investigational drug can obtain FDA approval, the drug must be proven both safe and effective in humans. Unfortunately, only one in five drugs that enter this human testing hurdle of the FDA approval process, also known as clinical trials, actually obtains approval (DiMasi 2003). This failure to obtain approval does not always mean that the drug is truly unsafe or ineffective; it may simply be the test was not setup for success. Possibly, a different dosage or a different length of treatment may have led to a better outcome for the clinical trial. In an effort to achieve more successes, the pharmaceutical industry is moving towards adaptive clinical trials which allow for flexibility in how a trial is conducted (Lesko 2007). For example, an adaptive approach allows for different dosages during the same trial and to weed out the dosages that prove toxic or non-therapeutic as trial data is collected and analyzed. This is in stark contrast to the traditional, more rigid approach to clinical trials where a dosage, a patient population, a length of treatment, and specific measures of success are chosen all prior to commencement of the trial.

Of course, the purely rigid approach is a supply manager's dream because demand is more certain. However, in an adaptive trial, more uncertainty about demand is introduced. For example, if during an adaptive trial the recommended dosage of an investigational treatment goes from 10mg to 20mg, the demand has just been effectively doubled. As a result, we see that the increased probability of a successful trial for FDA approval comes at the expense of the supply chain. Many of the key parameters used for forecasting supply are now subject to change during the trial.

To ensure supply for the start of a trial, clinical supply managers are forced to "forecast based on a limited number of variables before the cost of the supply chain overshadows the risk of excluded variables impacting supplies." (Hall 2008). Given more time, all of the variables impacting clinical supply requirements become less fluid and a better forecast of demand can be made. Given the need to start clinical production before demand is known leads to two possible strategies. If we assume a sufficient shelf life of the investigational drug, one supply strategy is to produce enough material in the initial production run to accommodate any potential scenario for clinical trial demand. Alternatively, a supply manager can break production into two production runs. The first production run is made to ensure enough supply is available to start the trial. The second production run ensures enough supply to end the trial and is made with greater precision in forecasted demand. For example, a clinical trial, with globally distributed testing centers, will start with enough supply to satisfy a launch of the trial in the United States. Once launched in the United States, international testing centers are opened and most demand for these sites will be satisfied using supply from a second production run.

Supply from a second production run is committed to with greater certainty as to dosage and study enrollment parameters. A trial that begins with a range of possible dosages, also called treatment arms, also begins with high demand uncertainty. However, as the trial continues patient allocation to dosages becomes concentrated on the most promising of the initial dosages (see e.g. Krams, Lees, Hacke, Grieve, Orgogozo, and Ford 2003). Thus, the longer the trial continues, the longer the supply manager learns about demand and the more certain he can be about the remaining supply requirements.

Unfortunately, reducing the costs of supply-demand mismatches can only be realized by increasing the number of production runs. The additional runs, while costly, can be planned for using a more accurate forecast of trial demand. By incurring multiple setups, the high fixed costs of physical drug production and the associated quality control activities are amortized over less inventory and economies of scale are sacrificed. The essence of this paper is to understand this sacrifice of economies of scale and to optimally balance the benefits of learning against higher costs for physical supply.

In this paper, we construct a model to plan drug supply for adaptive clinical trials. While minimizing inventory costs, the model explicitly incorporates the learning that takes place during a clinical trial. At the beginning of a clinical trial, both dosage and patient recruitment requirements are very uncertain. As a trial is underway and continues, the ideal dosage and the desired patient enrollment are zeroed in and a more accurate forecast can be made. Various rates of demand learning are considered through the use of learning curves that are a function of time. In addition to demand learning, the model incorporates setup costs (i.e., fixed costs in production) for both pre-trial and intra-trial production. Thus, the model enables us to weigh the benefit of demand uncertainty reduction against the cost of incurring a second setup cost for production. The objective of this paper is to study the impact of fixed costs and demand learning rates on the value and the optimal timing of demand learning.

We organize the rest of the paper as follows. We review the related literature in §2. The model and analysis are presented in §3. A numerical study is presented in §4. Finally, we summarize the paper and discuss future research directions in §5.

2. Literature Review

While our model is motivated by challenges of the pharmaceutical industry, it is both applicable and similar to the problems faced by production planners of many short lifecycle products with uncertain demand. In one of the most cited works on production planning for short lifecycle products, Fisher and Raman (1996) pioneered an approach to reduce stock-out and markdown costs for a set of seasonal fashion products based on the idea that sufficiently short manufacturing lead times allow for a portion of production to occur after some initial demand is observed. This initial demand could then be used to better forecast actual demand for a set of new products. They described this logic as follows:

> The dramatic improvement in forecast accuracy after observing only 20% of ini-

> tial demand suggests a strategy for reducing the cost of too much or too little inventory: commit to a modest amount of initial inventory for each product, observe initial demand, and then produce an additional amount of each product based on improved forecasts.

Consistent with the above logic, the authors present a two-stage model that incorporates second stage forecast uncertainty reduction, which we will also refer to as *demand learning*, through the correlation of demand between observed first stage demand and total expected demand. Through this correlation, the authors captured how the second-stage production enables a more accurate match of supply and demand. As a result of decreased supply-demand mismatch costs, it is shown that use of the authors' methodology at Sport Obermeyer can potentially quadruple profits.

As our work in this paper is highly motivated by the vein of research in Fisher and Raman's seminal work, we highlight three key distinctions between our model and theirs. First, we allow the second production run to occur at different time points and thus, the scheduling of the second production run is also a decision variable. Second, we include setup costs directly whereas their work includes minimum lot size requirements as a proxy for these costs. Lastly, we study the impact of learning rate on the system performance by incoporate learning curves in the demand learning process. This enable us to connect the amount of learning with the amount of time we have available to learn. Combined, these three distinctions allow us to weigh the benefits of demand learning against the fixed costs of introducing a second production run.

If we abstract from adaptive clinical trials, the general idea we seek to model is that forecasts and planning decisions can be improved based on observation of a full or partial season's demand. Within this abstraction, we can find papers that address a similar problem to ours. For example, Parlar and Weng (1997) research a two-period production decision where a second production run is possible. However, this second production run occurs after the realization of demand and thus, full learning has occurred. As an example of partial learning, Eppen and Iyer (1997) study a large catalog retailer's decisions of how much of a women's fashion product to order and then how much to divert to outlet stores upon observation of a portion of demand. They continue a long-line of studies that utilize bayesian updating for forecast revision. A much earlier paper recognizing the importance of adaptively revising forecasts is given by Murray and Silver (1966). They employ a Bayesian methodology for updating an unknown sales probability of an item based on a known amount of potential customers.

In our paper, we employ a learning curve approach to study the effect that different rates of demand learning have on the optimal first period batch size and second period production quantity. As detailed in the survey by Yelle (1979), learning curve applications have extended far beyond the more traditional applications of modeling the decrease in per unit manufacturing costs or the increase in labor productivity due to organizational experience. To our knowledge, this paper is among the first to apply a learning curve model that predicts forecast uncertainty as decreasing with the log of time allotted to observe demand. While the application is new, previous studies suggest the applicability of its use. For example, Bitran, Haas, and Matsuo (1986)'s study of production planning at a consumer electronics company notes the reduction in the coefficient of variation (CV) between forecasts that are made in January ($CV = 1$), April ($CV = 0.5$), and October ($CV = 0.2$) and actual sales for the Christmas season. As opposed to using a learning curve to model this reduction in forecast uncertainty, the author's simply assume forecast error in each period is normally distributed with a known and decreasing standard deviation over time.

Even though our use of learning curves to model forecast uncertainty reduction is unique, the use of learning curves to understand the benefits of learning in production planning is not. Terwiesch and Bohn (2001) study whether learning to ramp up the yield of a supply process and thus, delaying time to market and foregoing early demand, is preferable to producing first to satisfy early market demand and then investing in learning to achieve greater yield. Using a dynamic programming approach, the authors prescribe that manager's experiment early during production ramp up to increase learning early on. Despite this early learning taking time from production when prices are at their highest, the authors argue that this is the right time to devote production capacity towards engineering trials and efforts to improve yield and production rates.

In another interesting study of learning and production, Kornish and Keeney (2008) study the tradeoff of learning against capacity. Their model, motivated by the annual decision of which strains of flu to vaccinate against, address "a trade-off between quantity (producing more) and quality (produce a more effective vaccine because you know more)." By delaying the commitment to which flu strains will be targeted by the produced vaccine, a more effective vaccine can be made, but there is less time for production. The optimal policy is discussed and its implications for choosing to commit to production or wait for more information are detailed in their study.

The inclusion of a second production run, or mid-season replenishment option if in a retail setting, for short lifecycle products is a logical extension to the classical newsboy model. Lau and Lau (1997) study a mid-season replenishment possibility, but do not include set-up costs and only uniformly distributed demand is addressed in their model. While they consider making the time of the replenishment a decision variable prior to the selling season, their use of modeling each period's demand as a uniform distribution restricts this possibility. Milner and Rosenblatt (2002) study of the buyer's perspective when making a supply contract for a short life-cycle product is more amenable to making the timing of the second production run a decision variable. While the focus of their two-period model is on the contract form, they have a secondary contribution which is to consider the duration of the first and second period as a decision variable. Fisher, Rajaram, and Raman (2001) explicitly consider time as a decision variable when planning mid-season replenishment and use a heuristic solution to help a catalog retailer. More recently, Li, Chand, Dada, and Mehta (2009) have relaxed many of the assumptions in Fisher et al. (2001) and yield more structural results on the form of the optimal policy. Our paper differs from all of these mid-season replenishment models through the inclusion of setup costs and learning curves in our model.

The inclusion of setup costs in a production planning model for short lifecycle products with forecast updating has also been addressed in the literature. For example, the previously mentioned study by Bitran, Haas, and Matsuo (1986) includes the effect of setup costs for a family of products in their model. In a more recent study, Weng (2004) includes setup costs in his model to study the effects of those costs on coordination of ordering quantities between manufacturer and retailer. They find that as setup costs increase, the importance of a coordinating contract also increases. Our inclusion of both setup costs and learning allows us to find a balance between sacrificing economies of scale with additional production and benefiting from production made with less demand uncertainty.

3. The Two-Period Model

In this section, we first introduce the notation and the model in Subsection 3.1. We then characterize the optimal ordering policy in Subsection 3.2. Next we introduce the demand learning model in Subsection 3.3 and finally we study the impact of fixed cost in Subsection 3.4.

3.1 Notation and Model

Let the planning horizon be $[0, T]$ and assume zero lead time. We will use the following notation throughout the paper:

- t: Length of the learning period (i.e. the first period) where $0 \leq t < T$. In the case of one production batch, $t = 0$.
- D: Total demand in period $[0, T]$.
- D_1: Demand in the 1st period, $[0, t]$.
- D_2: Demand in the 2nd period, $[t, T]$.
- $D_2|_{D_1=\xi}$: Demand in the 2nd period $[t, T]$ given $D_1 = \xi$ where ξ is the realization.
- x_1: Production quantity for the first period and made available at time 0.
- x_2: Production quantity for the second period and made available at time t.
- κ: Setup cost of production.
- π_b: Backorder penalty per unit short after the first period that is ultimately satisfied.
- π_s: Shortage penalty per unit of unmet demand at the end of the time horizon.
- r: Overage penalty (destruction/recycle cost) per unit leftover item at the end of the time horizon.
- y_2: Order up to level for the second period.

To avoid trivial cases and ensure a realistic model, we make the following assumption.

Assumption 1 *$\pi_s > \pi_b > c \geq 0$, $\pi_s - \pi_b > c$ and demand is non-negative.*

Assuming zero initial inventory, we let $f_1(t)$ be the optimal expected inventory cost for the two-period problem with t being the duration of the learning period and let $\delta(x_t)$ be the indicator function of $x_t > 0$. Then, the optimal cost is expressed as

$$f_1(t) = \min_{x_1 \geq 0} \left\{ \delta(x_1)\kappa + cx_1 + \pi_b E_{D_1}\left[(D_1 - x_1)^+\right] + E_{D_1}\Big[f_2\left(x_1 - D_1, t\right)|_{D_1}\Big] \right\}, \tag{3.1}$$

where $f_2(I,t)|_{D_1}$ represents the optimal expected inventory costs for the second period conditioning on D_1, given a starting inventory of $I = x_1 - D_1$ and the length of the first period, t.

$$\begin{aligned} f_2(I,t)|_{D_1=\xi} \quad = \quad & \min_{x_2 \geq 0} \Big\{ \delta(x_2)\kappa + cx_2 - \pi_b \left(-I - x_2\right)^+ + \\ & \pi_s E_{D_2|_{D_1=\xi}} \left[(D_2|_{D_1=\xi} - I - x_2)^+ \right] + r E_{D_2|_{D_1=\xi}} \left[(I + x_2 - D_2|_{D_1=\xi})^+ \right] \Big\}. \end{aligned} \tag{3.2}$$

Thus, the optimal first period cost is the sum of the first period setup costs, first period variable production costs, first period backorder penalty costs, and the expected second period costs. The second period costs consist of a setup cost for a second period production, variable production costs for the second production run, a rebate on backorder penalties charged in the first period that turn out to be lost sales (i.e. second period production does not satisfy the unmet demand of the first period), a lost sales penalty, and destruction costs. Note that in adding an additional replenishment option, we also must introduce in intra-period shortage penalty. We consider this a backorder penalty which is much less costly than the shortage penalty charged at the end of the horizon. The second period shortage penalty, since there is no additional recourse for additional replenishment, is analogous to a lost sales penalty.

In our analysis of the timing of the first period, three points are worthy of mention. First, our model explicitly excludes lead time in the consideration. Second, our model requires that the timing of the second production run be scheduled in advance of the season. Third, our model excludes the substantial risk of the trial being halted prior to the end of the time horizon. This risk, which we call failure risk, is due to the possibility of a trial showing that a drug is unsafe or ineffective prior to the conclusion of the trial.

For certain types of clinical drug supply the zero lead time assumption may be untenable. However, for other types this assumption is valid. For example, when NeoRx Corporation outsourced clinical trial supply to International Isotopes Inc., purchase orders were only placed one week in advance and rolling forecasts were provided for 3 months of future demand (NeoRX 10-Q Filed on May 9, 2000.). These durations are much shorter than the overall duration of the trial, which may take as long as 2-3 years. Please note that these purchase orders are for batches produced after the first batch. The lead time on the first batch may still be lengthy as manufacturing facilities are configured for initial production. Once production facilities and processes are in place, lead time on additional batches can be much shorter than that of the first batch.

The timing of the intra-trial batch being determined prior to commencement of the actual trial is consistent with outsourcing contracts for clinical supply where the availability of manufacturing capacity is reserved in advance. It is also worth noting that Li, Chand, Dada, and Mehta (2009) have found little value to dynamically determining the timing of the second production run. We do not consider the effects of cancellation fees on the supply manager's decisions, although we believe this might be an interesting area for future research.

Lastly, the risk of failure risk, although important, is excluded from our analysis and reflects the typical supply philosophy of planning for success during a clinical trial. If failure were to occur during time $(0, t)$, it is as if the second production option would go unutilized. By applying a discount factor to the cost of producing in the second period, this aspect of clinical trials could be captured. To keep our analysis to the balancing of economies of scale and uncertainty reduction, we propose that inclusion of failure risk may be an interesting area of future research.

3.2 Optimal Ordering Policy

It is more convenient to use $y_2 = I + x_2$ and thus Eq. (3.2) becomes,

$$f_2(I,t)|_{D_1=\xi} = -cI + \min_{y_2 \geq I} \Big\{ \delta(y_2 - I)\kappa + cy_2 - \pi_b \left(-y_2\right)^+ + L_\xi(y_2, t) \Big\}, \tag{3.3}$$

where $L_\xi(y,t) = \pi_s E_{D_2|_{D_1=\xi}} \left[(D_2|_{D_1=\xi} - y)^+ \right] + r E_{D_2|_{D_1=\xi}} \left[(y - D_2|_{D_1=\xi})^+ \right]$.

To further analyze Eq. (3.3), we note that we can either produce (i.e. $x_2 > 0$) or not produce,

$$\begin{aligned} f_2(I,t)|_{D_1=\xi} \quad = \quad & -cI + \min \Big\{ \min_{y_2 > I} \Big\{ \kappa + cy_2 - \pi_b \left(-y_2\right)^+ + L_\xi(y_2, t) \Big\}, \\ & cI - \pi_b \left(-I\right)^+ + L_\xi(I, t) \Big\}. \end{aligned} \tag{3.4}$$

Observation 1 *If $y_2 > I$, then the optimal order-up-to level for the second period $y_2^* \geq 0$.*

Proof. If $I \geq 0$, $y_2^* \geq 0$ by definition. If $I < 0$, consider a $y_2 \leq 0$. By Assumption 1 and Eq. (3.3),

$$\begin{aligned} \kappa + cy_2 - \pi_b \left(-y_2\right)^+ + L_\xi(y_2, t) \quad &= \quad \kappa + cy_2 - \pi_b(-y_2) + \pi_s E_{D_2|_{D_1=\xi}} \left[D_2|_{D_1=\xi} - y_2 \right] \\ &= \quad \kappa + cy_2 + E_{D_2|_{D_1=\xi}} \left[\pi_s D_2|_{D_1=\xi} + (\pi_s - \pi_b)(-y_2) \right]. \end{aligned}$$

By Assumption 1, the cost function is decreasing in y_2 for $y_2 \leq 0$. Thus $y_2^* \geq 0$. □

By Observation 1, Eq. (3.4) can be reduced to,

$$f_2(I,t)|_{D_1=\xi} = -cI + \min\left\{\min_{y_2>I}\left\{\kappa + cy_2 + L_\xi(y_2,t)\right\}, cI - \pi_b\,(-I)^+ + L_\xi(I,t)\right\}. \quad (3.5)$$

The following observation shows that the second period cost function is not convex and thus we cannot directly apply the classical result of (s, S) policy (e.g., see Zipkin (2000, Section 9.5)) to this problem.

Observation 2 *The second period cost function $cI - \pi_b\,(-I)^+ + L_\xi(I,t)$ is not convex in I, but it is unimodular in I and approaches infinity as $I \to \pm\infty$.*

Proof. First, we note that $-\pi_b\,(-I)^+$ is concave in I. For $I < 0$, it follows by Assumption 1 that the second period cost function reduces to

$$cI - \pi_b(-I) + \pi_s E_{D_2|_{D_1=\xi}}[D_2|_{D_1=\xi} - I] = \pi_s E_{D_2|_{D_1=\xi}}[D_2|_{D_1=\xi}] - (\pi_s - \pi_b - c)I,$$

which is clearly convex in I. For $I \geq 0$, the second period cost function reduces to $cI+L_\xi(I,t)$, which is also convex in I. However, the left derivative of the cost function at $I = 0$ equals $-(\pi_s - \pi_b - c) < 0$ (by Assumption 1), which is greater than the right derivative of the cost function at $I = 0$, $-(\pi_s - c)$. Thus, the cost function is not convex in I for $I \in (-\infty, \infty)$.

To show the cost function is unimodular, we note that it is convex and decreasing in I for $I \in (-\infty, 0]$. Because it is also convex in I for $I \in [0, \infty)$, it must be unimodular.

Finally, as $I \to -\infty$, the slope of the cost function is $-(\pi_s - \pi_b - c)$; as $I \to \infty$, the slope approaches $c + r$. The proof is now completed. □

Now we are ready to identify the optimal ordering policy for the second period given $D_1 = \xi$. Let $S_2(\xi)$ be the smallest global minimizer of $cy + L_\xi(y,t)$, and $s_2(\xi)$ be the largest I (but smaller than $S_2(\xi)$) such that $cI - \pi_b(-I)^+ + L_\xi(I,t) = \kappa + cS_2(\xi) + L_\xi(S_2(\xi),t)$. Indeed, $S_2(\xi) = \Phi^{-1}_{D_2|_{D_1=\xi}}\left(\dfrac{\pi_s - c}{\pi_s + c}\right) \geq 0$, where $\Phi_{D_2|_{D_1=\xi}}(\cdot)$ is the probability density function of $D_2|_{D_1=\xi}$. $s_2(\xi)$ must exist by the unimodularity and asymptotic properties shown in Observation 2.

Theorem 1 *The optimal ordering policy for the second period is a (s, S) type of policy depending on $D_1 = \xi$, where $s = s_2(\xi)$ and $S = S_2(\xi)$. In other words, if the beginning inventory position $I < s_2(\xi)$, we order up to $S_2(\xi)$; otherwise, we do not order.*

Proof. The proof follows directly from the definition of $s_2(\xi), S_2(\xi)$ and Observation 2. □

Note that $s_2(\xi)$ and $S_2(\xi)$ are dependent on $D_1 = \xi$ but independent of second period starting inventory I.

By Theorem 1, we can write $f_2(I,t)|_{D_1=\xi}$ as follows,

$$f_2(I,t)|_{D_1=\xi} = -cI + \begin{cases} \kappa + cS_2(\xi) + L_\xi(S_2(\xi),t), & I \leq s_2(\xi) \\ cI - \pi_b(-I)^+ + L_\xi(I,t), & I > s_2(\xi). \end{cases} \tag{3.6}$$

We now show $f_2(I,t)|_{D_1=\xi}$ is κ-convex for any ξ. By Zipkin (2000, Section 9.5), we have the following definition.

Definition 1 *We call a function $f(x)$ κ-convex if for any x, and nonnegative u and v, $f(x)$ satisfies*

$$f(x) + v\frac{f(x) - f(x-u)}{u} \leq f(x+v) + \kappa.$$

Lemma 1 *$f_2(I,t)|_{D_1=\xi}$ is κ-convex in I for any ξ.*

Proof. For simplicity, we drop ξ from our notation without causing confusion. We also define $\tilde{f}(I,t) = f_2(I,t) + cI$. Clearly if $\tilde{f}(I,t)$ is κ-convex, $f_2(I,t)$ is also κ-convex.

If $s_2 \geq 0$, then

$$\tilde{f}_2(I,t) = \begin{cases} \kappa + cS_2 + L(S_2,t), & I \leq s_2 \\ cI + L(I,t), & I > s_2. \end{cases}$$

By Zipkin (2000, Section 9.5), $\tilde{f}_2(I,t)$ and thus $f_2(I,t)$ is κ-convex.

If $s_2 < 0$, we first consider any $I \leq s_2$. By the definition of the (s_2, S_2) policy, we must have $\tilde{f}_2(I,t) + v\frac{\tilde{f}_2(I,t) - \tilde{f}_2(I-u,t)}{u} = \tilde{f}_2(I,t) = \tilde{f}_2(S_2,t) + \kappa \leq \tilde{f}_2(x,t) + \kappa$ for all x. Next we consider $I \geq S_2$. The κ-convexity inequality must hold because $\tilde{f}_2(I,t)$ is convex and increasing. Finally, we consider $s_2 < I \leq S_2$. Note that $\tilde{f}_2(I,t)$ is decreasing for $I < S_2$ by Observation 2, thus $\tilde{f}_2(I,t) + v\frac{\tilde{f}_2(I,t) - \tilde{f}_2(I-u,t)}{u} \leq \tilde{f}_2(I,t) \leq \tilde{f}_2(S_2) + \kappa \leq \tilde{f}_2(I+v) + \kappa$. The proof is now completed. □

Theorem 2 *The optimal ordering policy for the first period is a (s,S) type of policy.*

Proof. By Lemma 1, $f_2(I,t)|_{D_1=\xi}$ is κ-convex. By Lemma 9.5.1 of Zipkin (2000), $E_{D_1}[f_2(x_1 - D_1,t)|_{D_1}]$ is also κ-convex, and so is $cx_1 + \pi_b E_{D_1}\left[(D_1 - x_1)^+\right] + E_{D_1}\left[f_2\left(x_1 - D_1, t\right)|_{D_1}\right]$. By Theorem 9.5.2 of Zipkin (2000), the proof is completed. □

Let (s_1, S_1) be the optimal (s,S) policy for the first period. Thus, if $s_1 > 0$, then we produce up to S_1 in the first period. Otherwise, we do not produce.

3.3 Demand Learning Model

The initial belief of the total demand during the planning horizon is that D is normally distributed with mean μ and standard deviation σ. We assume dependent demand in periods 1-2: D_1 and D_2 ($D = D_1 + D_2$). Following Fisher and Raman (1996) and Fisher, et al. (2001), we assume (D_1, D_2) follows a bivariate normal distribution with correlation coefficient $\rho(t)$, where $\rho(t)$ depends on the amount of learning that can take place by time t. In Section 4, we borrow methodology from the learning curves literature to model this function. The marginal distribution D_1 is also normal with mean μ_1 and standard deviation σ_1. We assume that for $t \in [0, T]$,

$$E(D_1) = \mu_1 = \alpha(t)\mu, \qquad \sigma^2(D_1) = \sigma_1^2 = \beta(t)\sigma^2,$$

where $\alpha(t)$ and $\beta(t)$ are fractions increasing from 0 to 1 as t increases from 0 to T. For example, $\alpha(t) = t/T$ and $\beta(t) = t/T$. Then the marginal distribution of D_2 is normal with

$$E(D_2) = \mu_2 = (1 - \alpha(t))\mu, \qquad \sigma(D_2) = \sigma_2 = -\rho(t)\sigma_1 + \sqrt{\rho^2(t)\sigma_1^2 + \sigma^2 - \sigma_1^2}.$$

Conditioning on $D_1 = \xi$, D_2 follows a normal distribution with the following parameters (Fisher and Raman 1996):

$$\mu_2(\xi) = \mu_2 + \rho(t)\sigma_2 \frac{\xi - \mu_1}{\sigma_1}, \quad \sigma_2(\xi) = \sigma_2\sqrt{1 - \rho^2(t)}.$$

It is easy to see that given t, as ρ increases, σ_2 decreases and thus $\sigma_2(\xi)$ decreases.

We model demand learning through correlation between first and second period demand. Through correlation, a fraction of the variance in the second period's demand is explained by the realization of the first period's demand. Mathematically, this fraction is simply $\rho(t)^2$ and the fraction of variance that remains unexplained in the second period is $1 - \rho(t)^2$. Consistent with this mathematical interpretation, we will model demand learning as a reduction in unexplained variance. Intuitively, the fraction of unexplained variance in second period demand should be close to one early in the time horizon and closer to zero at the end of the horizon. To study different rates of learning, we will assume that the amount of learning, more specifically the reduction in uncertainty surrounding second period demand, $1 - \rho(t)^2$, follows a power law form that was introduced as a learning curve model by Wright (1936). Modifying Wright's learning curve to represent a supply manager's ability to remove uncertainty in the variance of second period demand, we have $1 - \rho(t)^2 = \left(\frac{T-t}{T}\right)^\gamma$ and thus

$$\rho(t) = \sqrt{1 - \left(\frac{T-t}{T}\right)^\gamma} \tag{3.7}$$

where $\gamma \geq 0$ is the shape parameter of the learning curve. When $\gamma = 1$ the amount of learning is linear in time, when $\gamma < 1$, learning is slow and the uncertainty parameter, $1 - \rho(t)^2$, will be a concave function of time. At $\gamma = 0$, there is no learning and demand is independent across periods. Lastly, when $\gamma > 1$, learning occurs more rapidly and $1 - \rho(t)^2$ is a convex function of time.

3.4 Impact of The Fixed Cost

Intuitively, as the fixed cost, κ, increases, the two-period model reduces to the newsvendor model. This intuition is confirmed by the following proposition.

Proposition 1 *S_1 tends to the newsvendor quantity as $\kappa \to \infty$.*

Proof. For simplicity, we drop the dependence on D_1 for (s_2, S_2) without causing confusion. Suppose we produce in the first period, the total cost function can be expressed as follows,

$$\begin{aligned}\kappa \quad + \quad & \min_{x_1 > 0} \Big\{ cx_1 + \pi_b E_{D_1}[(D_1 - x_1)^+] + \\ & + E_{D_1}[-c(x_1 - D_1) + \kappa + cS_2 + L_{D_1}(S_2, t) | D_1 \geq x_1 - s_2] \\ & + E_{D_1}[-\pi_b (D_1 - x_1)^+ + L_{D_1}(x_1 - D_1, t) | D_1 < x_1 - s_2] \Big\}.\end{aligned}$$

As $\kappa \to \infty$, $s_2 \to -\infty$ (by Observation 2) for each realization of D_1. Thus, the cost function tends to

$$\kappa + \min_{x_1 > 0} \Big\{ cx_1 + E_{D_1}[L_{D_1}(x_1 - D_1, t)] \Big\},$$

where $E_{D_1}[L_{D_1}(x_1 - D_1, t)] = E_{D_1}[E_{D_2}[\pi_s(D_1 + D_2 - x_1)^+ + r(x_1 - D_1 - D_2)^+ | D_1]] = E_{D_1+D_2}[\pi_s(D_1 + D_2 - x_1)^+ + r(x_1 - D_1 - D_2)^+]$. The last equality comes from the definition of conditional expectation. Note that $E_{D_1+D_2}[\pi_s(D_1 + D_2 - x_1)^+ + r(x_1 - D_1 - D_2)^+]$ represents the cost function of the newsvendor model without the second period, the proof is now completed. $\square$

In general, as the fixed cost κ increases, S_1 will be more likely used to cover both D_1 and D_2, and thus S_1 typically increases.

4. Numerical Analysis

The objective of this section is to quantify the effects of setup costs, learning rates, and penalty costs on the value of the second production option, the optimal timing of learning, and the optimal first batch size. The penalty costs include the overage penalty cost, $r + c$, and underage penalty costs, $\pi_s - c$ and $\pi_b - c$.

As a starting point for the analysis, a baseline problem is created where we assume mean demand over the time horizon is $1,000$ units and the standard deviation of demand is 300 units. We assume demand for the time horizon is normally distributed and since this distribution is divisible, we can mathematically divide demand between two periods. Thus, the fraction, $\frac{t}{T}$, represents the percentage of total demand expected to occur in the first period. The parameters of our baseline model, in the absence of demand learning (i.e. $\gamma = 0$ in Eq. 3.7), are shown in Table 1.

Parameter	Description	Baseline Value
t	Length of first (learning) period.	5
T	Length of planning horizon.	10
γ	Rate of demand learning.	0
$D(0,t)$	First Period Demand $N(\mu, \sigma)$	$N(1000 * \frac{5}{10}, \sqrt{300^2 * \frac{5}{10}})$
$D(t,T)$	Second Period Demand $N(\mu, \sigma)$	$N(1000 * \frac{10-5}{10}, \sqrt{300^2 * (1 - \frac{5}{10}})$
κ	Setup cost of a production run	0
π_b	Backorder penalty	20
π_s	End of horizon shortage penalty	50
c	Variable production cost	2
r	Destruction Cost	1

Table 1: Baseline Parameter Values

In Figure 1, we compare the expected costs of our baseline model with the expected costs of a newsvendor model (i.e. a single production run at $t = 0$). For our baseline model, we have arbitrarily scheduled an additional replenishment option mid-way through the planning horizon. As can be seen from the graph, the additional replenishment opportunity leads to a greater than 10% reduction in costs. It is also interesting to note that the expected costs of the baseline model are less sensitive than the newsvendor model to first period order quantity. At this point, one might conclude that a mid-season replenishment can reduce expected costs and lead to a decision that is less sensitive to model parameters. However, the comparison we made ignores several key components that constitute the motivation for this study. First, setup costs are zero and when producing clinical trial drug supply, large setup costs are a reality that must be accommodated. And second, our comparison fails to account for demand learning and the freeing of the variable, t, so that this additional production can be optimally scheduled.

We now introduce setup costs into our analysis without considering demand learning. From the proof of Proposition 1, we know that as setup costs increase the two-period model reduces to the newsvendor model. We can see this effect by comparing Figure 2 which includes a setup cost ($\kappa = \$1,600$) to Figure 1 which assumes there are zero fixed costs when producing. With setup costs introduced, we can see that the value of an additional replenishment midway through the planning horizon yields minimal savings of 2% of the newsvendor solution's inventory costs.

Counteracting this decrease in value from our baseline model, we can free the scheduling of the second replenishment and show the effects of this scheduling on expected costs. To analyze this, we graph expected costs as a function of the first period length (t) as seen in Figure 3. For every choice of first period length, t, the optimal first period order quantity has been numerically determined. We see the optimal first period length is close to the entire planning horizon ($t \approx 8.5$) and not the arbitrarily chosen mid-horizon production ($t = 5$). As might be expected, the mid-horizon production case with setup costs leads to reduced savings over the newsvendor solution as compared to an available second production run without setup costs. Interestingly, we see that when $t = T$, the solution still outperforms the newsvendor solution. In this case, the value of the additional replenishment (when optimally planned) is purely derived from replacing the expected end of horizon shortage penalty with a backorder penalty in cases of high demand. In contrast, when $t = 8.5$, some of the value provided by this model as compared to the neswvendor solutions is from the ability to effectively match supply and demand for a portion of the horizon.

While setup costs have reduced some of the benefit of the intra-season replenishment option, demand learning creates greater incentives to plan an additional replenishment and counter-balance the costs imposed by an additional setup. Inituitively, faster rate of demand learning (higher γ in Eq. 3.7) encourages earlier scheduling of the potential second production run. To see this in our example, we now analyze the baseline model with setup costs for various rates of demand learning. We pick various values of our learning parameter, γ, to represent different rates of uncertainty reduction and plot the expected costs of our baseline model with setup costs ($\kappa = \$1,600$) and learning in Figure 4. We see from this graph that faster learning leads to both an earlier scheduling for the second production run and larger

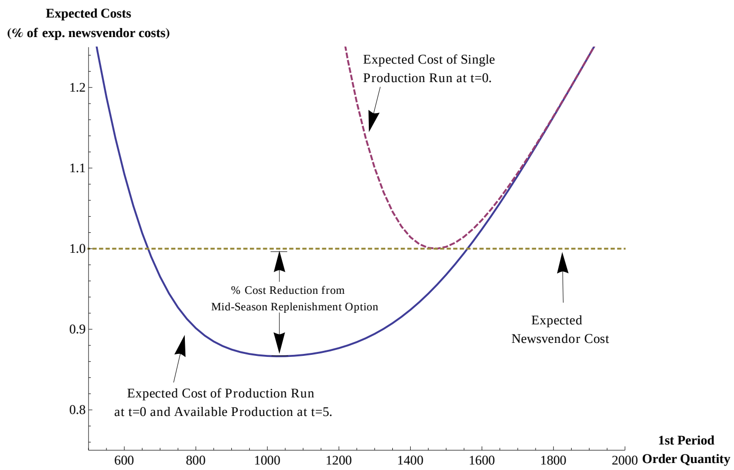

Figure 1: Expected Costs when $\kappa = \$0$

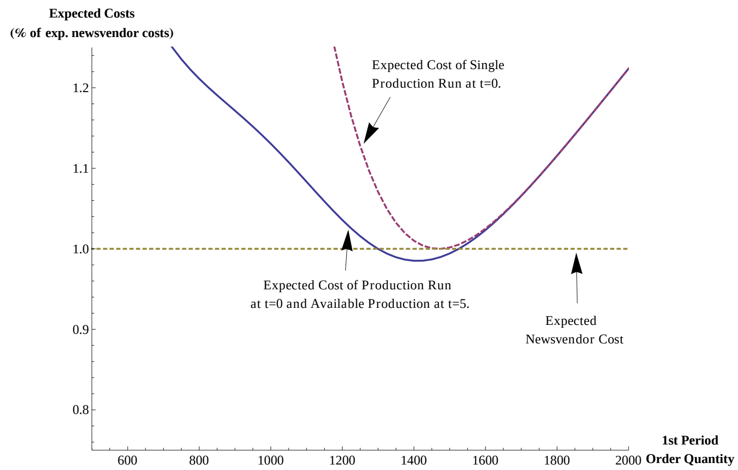

Figure 2: Expected Costs when $\kappa = \$1,600$

cost reductions versus the single newsvendor production. More importantly, the cost benefit of learning has effectively nullified the substantial setup cost. We see that when the rate of learning is simply linear in time ($\gamma = 1$), we can achieve savings of greater than 10% over the single production solution. This linear learning rate would yield a correlation coefficient of 0.82 at the optimal duration of the first period ($t \approx 6.75$) which is consistent with the correlations used in (Fisher, Rajaram, and Raman 2001).

While the scheduling decision is important, it is not made in isolation. The optimal supply strategy will simultaneously consider the timing of the second production and the sizing of the first batch. In Figure 5, we analyze the interplay of production scheduling and optimal first batch size for various setup costs and linear learning ($\gamma = 1$). In the absence of setup costs (i.e. $\kappa = 0$) and when replenishment is planned after observing a small fraction (roughly 10%) of demand, we observe that the optimal first period batch is less than half the newsvendor batch size. However, with even modest setup costs of \$100, the optimal batch sizing this early in the season is much closer to the newsvendor quantity. This is a key observation that the suggestion of a "modest amount of initial inventory" (Fisher and Raman (1996)) is less appropriate when setup costs are factored into the decision making. The benefits of ample inventory, including avoiding additional setup costs and first period backorder costs, outweigh the benefits of uncertainty reduction afforded by a second production.

In the absence of setup costs and in the presence of demand learning, a second production run is a likely event. For our baseline model with linear learning and zero setup costs, the second production run is optimally scheduled at around $t \approx 5.5$ and it is expected that 71%

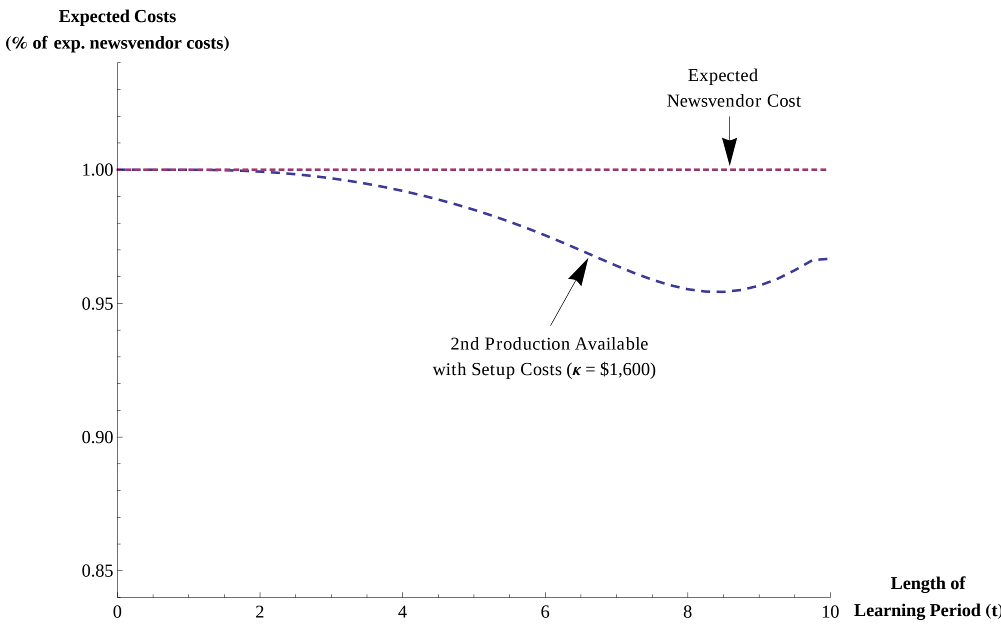

Figure 3: Expected Costs Versus Timing of Second Replenishment with Setup Costs

of the time the production run will be utilized. The other 29%, demand is so low in the first 55% of the planning horizon that a second production is not needed. Even though overage risk is present, this risk is offset by having enough inventory to avoid intra-period backorder costs.

As soon as we introduce setup costs, we also introduce a notion of economies of scale in production. A manager's expectation of producing more than once reflects his willingness to sacrifice scale economies to achieve savings. The tradeoff between sacrificing scale economies to better match supply and demand is summarized in Table 2. We see from this table that setup costs significantly decrease the probability of a second production. For example, in the case of linear learning ($\gamma = 1$), the introduction of setup costs of $1,600 reduces the likelihood of producing a second time from 71% to 21%. Further increases in setup costs drastically reduce the likelihood of producing a second time. From a planning perspective, mid-season replenishment in the presence of high setup costs is really an emergency supply option for cases of extremely high demand.

Even though the likelihood of producing a second time can be small, the value of this option remains significant in the presence of learning. This can be seen in Table 3 which shows the expected savings over the newsvendor model when optimally scheduling potential replenishment. From this table, we can see that with linear learning and setup costs of $1,600, a 10.7% reduction in costs can be expected by just having a resupply option available. From

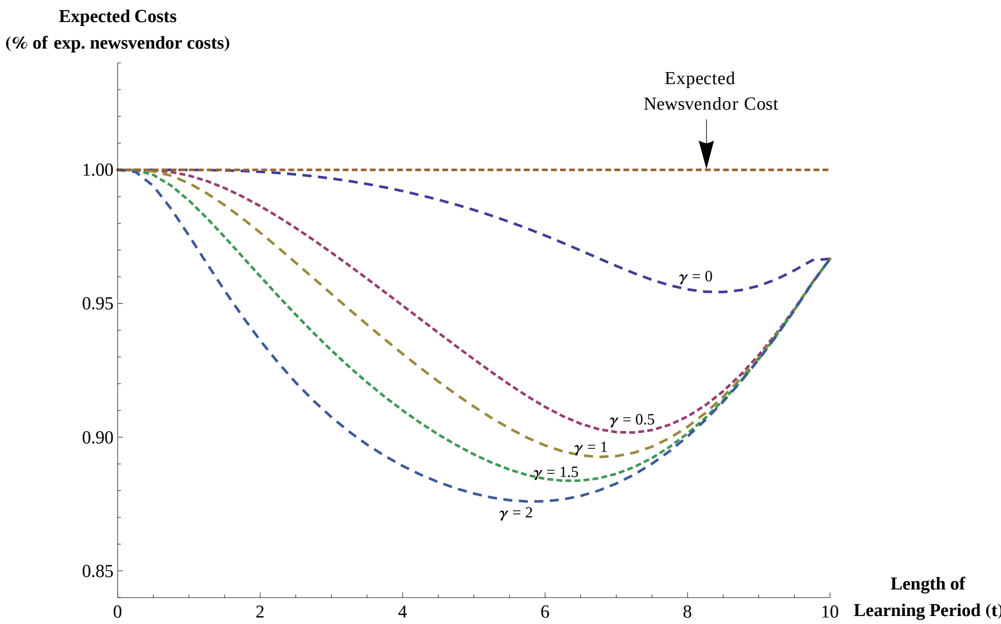

Figure 4: Expected Costs Versus Timing of Second Replenishment with Setup Costs and Demand Learning

Table 2, we know that this resupply option will only be exercised about 21% of the time. Digging deeper into Table 3, we see that even with higher levels of setup costs (e.g. $3,200), savings of greater than 5% are achievable. While in a pharmaceutical setting, these cost reductions are significant, in a retail setting Fisher et al. (2001) show how much smaller cost savings can translate into big gains in profitability.

	Learning Rate (γ)				
κ	**0**	**0.5**	**1**	**2**	**4**
-	66.7%	66.3%	70.6%	74.1%	77.4%
100	50.2%	59.4%	64.5%	69.1%	73.6%
200	41.1%	52.1%	57.5%	63.0%	68.1%
400	30.5%	43.3%	48.8%	54.4%	60.6%
800	19.8%	31.9%	36.8%	40.0%	47.5%
1,600	10.4%	18.7%	21.4%	24.4%	27.2%
3,200	4.8%	9.2%	10.2%	11.0%	11.9%
6,400	1.6%	3.4%	3.6%	3.8%	4.0%
12,800	0.1%	0.7%	0.7%	0.8%	0.8%

Table 2: Probability of Mid-Season Replenishment for Various Setup Costs and Learning Rates

Another consideration in this mid-season replenishment environment is how penalty costs

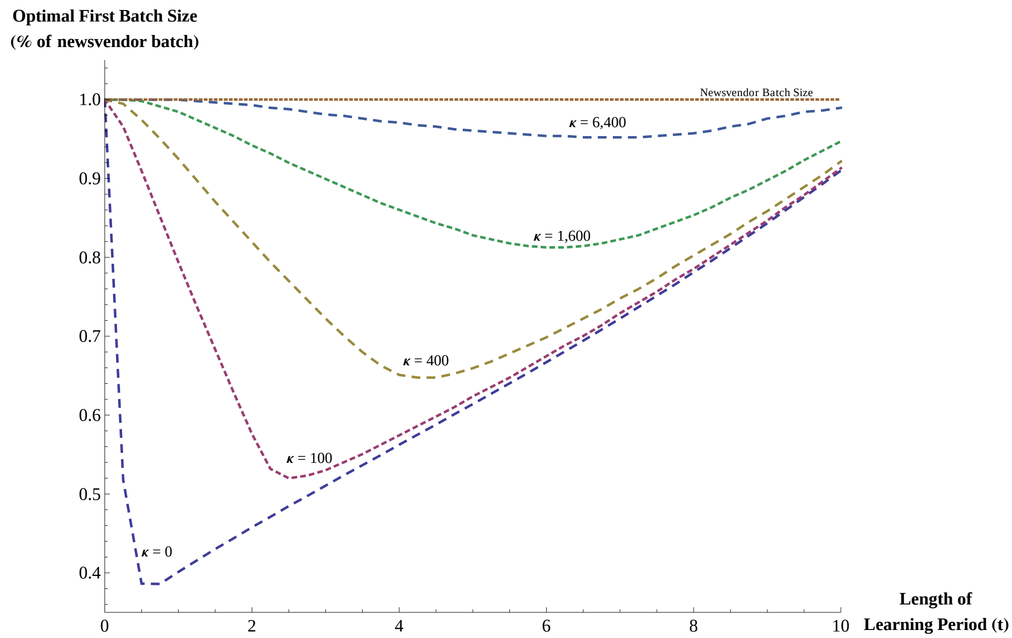

Figure 5: Optimal First Batch Size Versus Timing of Second Replenishment with Setup Costs and Demand Learning

(i.e. overage, lost sales, and intra-period backorder penalties) impact our decisions of replenishment timing and first period batch size. And even more importantly, how do changes in these parameters affect the magnitude of savings over a simpler newsvendor solution? The clinical trial supply environment is driven by a fear of delaying a trial due to insufficient supply and intuitively, one would think increasing underage penalties (π_s or π_b) would lead to greater expected savings of a second production. In studying this numerically, we surprisingly find the advantage of having an intra-season replenishment option is not dramatically improved by dramatically increased underage penalties. For example, our numerical study has found that doubling the two underage penalties of our baseline model with setup costs ($\kappa = \$1,600, \pi_b = 40, \pi_s = 100$) only increases expected savings over the newsvendor solution an additional 2.8% from 10.7% to 13.5%. Further increases to these underage penalties, as shown in Table 4, yield similarly modest results with the reason being that avoiding these underage penalties is relatively inexpensive; overage costs are only \$3 which is small in comparison to the end of horizon shortage penalty of \$50 of our our baseline model. Basically, it is cheap to hedge against having too little inventory by simply producing more.

In studying the effect of changes to the overage penalty, we find that increasing the overage penalty leads to greater jumps in savings magnitude than increasing underage costs. For example, if we look at the case where $\pi_b = 40, \pi_s = 100$, and $\kappa = \$1,600$, we find that

	Learning Rate (γ)				
κ	**0**	**0.5**	**1**	**2**	**4**
-	16.5%	28.8%	31.7%	34.8%	38.2%
100	14.7%	26.5%	29.1%	32.1%	35.3%
200	13.1%	24.4%	26.9%	29.7%	32.7%
400	10.9%	20.9%	23.1%	25.5%	28.0%
800	7.8%	15.8%	17.4%	19.1%	20.8%
1,600	4.6%	9.8%	10.7%	11.6%	12.4%
3,200	1.9%	4.5%	4.9%	5.2%	5.5%
6,400	0.4%	1.3%	1.4%	1.5%	1.6%
12,800	0.0%	0.2%	0.2%	0.2%	0.2%

Table 3: Table of Savings for Various Setup Costs and Learning Rates

π_b / π_s	% Savings
20 / 50	10.7
40 / 100	13.5
80 / 200	15.8
160 / 400	17.7
320 / 800	19.35

Table 4: Expected Savings Over a Newsvendor Solution for Increasing Underage Costs

savings of greater than 50% over the newsvendor solution can be expected when overage costs are 64% of the end of horizon shortage costs. A selection of overage cost penalty values and expected savings are shown in Table 5. Less dramatic than the changes in savings, both optimal batch sizes and optimal timing values fall within small ranges as the overage penalty is adjusted. Optimal batch sizes fall between 71% and 79% of their respective newsvendor optimal order sizes and optimal timing for the second production run is between 62% - and 68% of the planning horizon.

While increases to both overage and underage penalties will always increase savings over a newsvendor solution, the increase in percentage savings can be both small and large. Percentage savings increases per dollar of increased penalty cost are fastest when overage and underage costs are highly unbalanced and increases are made to the lower of the two costs. Conversely, when increases are made to the higher of the two costs, only marginal benefits will be realized. In a pharmaceutical setting where underage costs far exceed overage costs, the observation on the effects of increasing overage penalties suggests that increases in variable production costs, which effectively reduce the lost sales penalty and increase the overage penalty, will greatly increase the attractiveness of intra-season replenishment.

r	% Savings
1	13.5
2	16.9
4	22.4
8	29.9
16	38.1
32	45.4
64	50.7
128	53.9

Table 5: Expected Savings Over a Newsvendor Solution for Increasing Destruction Costs

In summary, our numerical study pursued an understanding of the value of an intra-season replenishment over the newsvendor solution. The value that is created depends on two key decisions in planning for this intra-season option. First, how is our first batch size decision affected by the presence of a resupply option and second, when should the potential resupply be planned for. These two decisions, and the potential to achieve meaningful savings are affected by a multitude of parameters and for simplicity, we now summarize our findings in Table 6.

Parameter↑	*relative to newsvendor solution* % Savings $\left(\frac{f_1(0)-f_1(t_{opt})}{f_1(0)}\right)$	*relative to newsvendor solution* 1st Batch Size (%) $\left(\frac{x_1}{x_{newsv}}\right)$	1st Batch Size (x_1)	1st Period Length (t)
Learning (γ)	↑	↓	↓	↓
Setup Costs (κ)	↓	↑	↑	↑↓
Backorder Penalty (π_b)	↓	↓	↓	↓
Shortage Penalty (π_s)	↑	↑↓	↑	↑
Destruction Costs (r)	↑	↑↓	↓	↓

Table 6: Effect of Parameter Increases on Performance and Decision Variables

5. Conclusions

The inclusion of setup costs and learning curves in our study leads us to many conclusions that add to the body of literature dealing with an additional replenishment option for products with short lifecylces. When setup costs are present, the first batch remains large as compared to a newsvendor batch. As opposed to a smaller batch, the inventory helps to potentially avoid incurring setup costs a second time and avoiding intra-season backorder penalties.

Selection of learning period length is driven by both changes in setup costs and learning. We observe that increasing setup costs will initially increase the optimal learning period length and then decrease it. At lower levels of setup costs, when these costs increase, it is advantageous to have a longer learning period to permit greater observation of demand and a more certain second period forecast. Eventually, further increases to setup costs decrease the learning period to avoid backorder costs in the cases of extremely high demand that would actually warrant incurring a second setup.

Our study is the first to look at the tradeoff between sacrificing economies of scale by planning for multiple batches and benefiting from demand learning so that a better match of supply and demand can be made. We have found that the ideal conditions for consideration of an additional production run are when setup costs are low, learning is fast, and both overage and underage penalties are significant. In certain examples, we find savings to exceed 50%. Admittedly, these high-value examples are less applicable to clinical trials where underage costs far exceed overage costs. However, these high value examples are realistic when extending this model to a fashion environment or other short lifecycle products where underage and overage costs are not so lop-sided.

Acknowledgments

This research is supported by the grant CMMI-0747779 from the National Science Foundation.

References

Bitran, Gabriel R., Elizabeth A. Haas, Hirofumi Matsuo. 1986. Production planning of style goods with high setup costs and forecast revisions. *Operations Research* **34**(2) 226–236. URL `http://www.jstor.org/stable/170817`.

DiMasi, J. A. 2003. The price of innovation: new estimates of drug development costs. *Journal of Health Economics* **22**(4–5) 469–479.

Eppen, Gary D., Ananth. V. Iyer. 1997. Improved fashion buying with bayesian updates. *Operations Research* **45**(6) 805–819. URL `http://www.jstor.org/stable/172066`.

Fisher, Marshall, Kumar Rajaram, Ananth Raman. 2001. Optimizing inventory replenishment of retail fashion products. *Manufacturing & Service Operations Management* **3**(3) 230–241.

Fisher, Marshall, Ananth Raman. 1996. Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research* **44**(1) 87–99. URL `http://www.jstor.org/stable/171907`.

Hall, Catherine. 2008. Clinical supplies on the front lines: Adapting to trial demand. *Clinical Supply Management for Pharmaceuticals Conference*. Pfizer Global Research and Development, Philadelphia, PA.

Kornish, Laura J., Ralph L. Keeney. 2008. Repeated Commit-or-Defer Decisions with a Deadline: The Influenza Vaccine Composition. *OPERATIONS RESEARCH* **56**(3) 527–541. doi:10.1287/opre.1070.0436. URL `http://or.journal.informs.org/cgi/content/abstract/56/3/527`.

Krams, Michael, Kennedy R. Lees, Werner Hacke, Andrew P. Grieve, Jean-Marc Orgogozo, Gary A. Ford. 2003. Acute Stroke Therapy by Inhibition of Neutrophils (ASTIN): An Adaptive Dose-Response Study of UK-279,276 in Acute Ischemic Stroke. *Stroke* **34**(11) 2543–2548. doi:10.1161/01.STR.0000092527.33910.89. URL `http://stroke.ahajournals.org/cgi/content/abstract/34/11/2543`.

Lau, H-S., A. H-L. Lau. 1997. A semi-analytical solution for a newsboy problem with mid-period replenishment. *The Journal of the Operational Research Society* **48**(12) 1245–1253. URL `http://www.jstor.org/stable/3010755`.

Lesko, L. J. 2007. Paving the critical path: How can clinical pharmacology help achieve the vision? *Clinical Pharmacology & Therapeutics* **81** 170–177.

Li, Jian, Suresh Chand, Maqbool Dada, Shailendra Mehta. 2009. Managing Inventory Over a Short Season: Models with Two Procurement Opportunities. *MANUFACTURING SERVICE OPERATIONS MANAGEMENT* **11**(1) 174–184. doi:10.1287/msom.1070.0204. URL `http://msom.journal.informs.org/cgi/content/abstract/11/1/174`.

Milner, Joseph M., Meir J. Rosenblatt. 2002. Flexible supply contracts for short life-cycle goods: The buyer's perspective. *Naval Research Logistics* **49**(1) 25–45.

Murray, Jr., George R., Edward A. Silver. 1966. A bayesian analysis of the style goods inventory problem. *Management Science* **12**(11) 785–797. URL http://www.jstor.org/stable/2627744.

NeoRX, NeoRx. 10-Q Filed on May 9, 2000. Clinical manufacture and supply agreement. *Poniard Pharmeceuticals (formerly NeoRX Inc.) 10-Q* .

Parlar, Mahmut, Z. Kevin Weng. 1997. Designing a firm's coordinated manufacturing and supply decisions with short product life cycles. *Management Science* **43**(10) 1329–1344. URL http://www.jstor.org/stable/2634409.

Terwiesch, Christian, Roger E. Bohn. 2001. Learning and process improvement during production ramp-up. *International Journal of Production Economics* **70**(1) 1 – 19. doi:DOI: 10.1016/S0925-5273(00)00045-1. URL http://www.sciencedirect.com/science/article/B6VF8-4292KN9-1/2/24fb1d300a143f780744292819d0d6d5.

Weng, Z. Kevin. 2004. Coordinating order quantities between the manufacturer and the buyer: A generalized newsvendor model. *European Journal of Operational Research* **156**(1) 148 – 161. URL http://www.sciencedirect.com/science/article/B6VCT-4BRR7D3-D/2/c372c5dba77daf3df5b07be16b99c83e.

Wright, T. 1936. Factors affecting the cost of airplanes. *Journal of Aeronautical Sciences* **3**(4) 122–128.

Yelle, Louis E. 1979. The learning curve: Historical review and comprehensive survey. *Decision Sciences* **10**(2) p302 – 328. URL https://login.proxy.libraries.rutgers.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=buh&AN=5002434&site=ehost-live.