Inventory Systems with Stochastic Demand and Supply: Properties and Approximations

Amanda J. Schmitt
Center for Transportation and Logistics
Massachusetts Institute of Technology
Cambridge, MA, USA

Lawrence V. Snyder
Dept. of Industrial and Systems Engineering
Lehigh University
Bethlehem, PA, USA

Zuo-Jun Max Shen
Dept. of Industrial Engineering and Operations Research
University of California
Berkeley, CA, USA

December 12, 2009

ABSTRACT

We model a retailer whose supplier is subject to complete supply disruptions. We combine discrete-event uncertainty (disruptions) and continuous sources of uncertainty (stochastic demand or supply yield), which have different impacts on optimal inventory settings. This prevents optimal solutions from being found in closed form. We develop a closed-form approximate solution by focusing on a single stochastic period of demand or yield. We show how the familiar newsboy fractile is a critical trade-off in these systems, since the optimal base-stock policies balance inventory holding costs with the risk of shortage costs generated by a disruption.

1 Introduction

Supply disruptions can have drastic impacts on firms who fail to protect against them. Traditional inventory models focus on demand uncertainty and design the system to best mitigate that risk. However, the effects of supply disruptions can have very different implications for system design. In the past five years, there has been an explosion of research on inventory and supply chain models with supply disruptions. In this paper, we examine a single-stage inventory system with disruptions and introduce an effective approximation for systems with both disruptions and demand or yield uncertainty. These results constitute a set of tools that will be useful for future research on inventory models with supply disruptions.

We examine optimal base-stock inventory policies using infinite-horizon, periodic-review models, for a single supplier whose single retailer is subject to stochastic disruptions. Due to the complexity of mixing discrete and continuous distributions in modeling, it is complicated to analyze supply disruptions in combination with either demand uncertainty or yield uncertainty. We examine both cases in this paper.

Our first model considers stochastic demand. We develop an approximate technique to set base-stock levels since the model cannot be solved in closed-form. This technique determines approximately how many periods' worth of demand should be kept in stock (based on the expected duration of disruptions and relative weight of penalty costs), and then sets safety stock levels accordingly. It essentially assumes that all other periods experience demand equal to their mean, or deterministic demand, so we call this the Single Stochastic Period (SSP) approximation. We compare the SSP performance to the optimal solution (found numerically), and demonstrate that on average it generates a cost increase of 0.17% and outperforms approximations based on uniform or triangular distributions for demand.

We next consider supply disruptions and yield uncertainty together, and apply the SSP approximation. We show that it performs very well in this system, with an average cost increase of 0.03%.

There are many benefits of having a closed-form approximate solution, such as the one we develop, for a problem which would otherwise require numerical optimization. A closed-form solution clearly demonstrates the sensitivity of solutions to input parameters. It can also be embedded into more complicated models to add tractability. Closed-form approximations are also useful tools in practice, since they are easier to implement and use on an ongoing basis.

The remainder of the paper is outlined as follows: In Section 2 we review relevant literature for the topic. In Section 3 we introduce our approach for modeling disruptions and review a result that will be used later in the paper. We present approximate solutions for problems with disruptions and stochastic demand in Section 4, and with disruptions and yield uncertainty in Section 5. We summarize our findings in Section 6.

2 Literature Review

Supply uncertainty is typically modeled as complete disruptions, where supply halts entirely, or as yield uncertainty, where the supply quantity received varies stochastically. One of the first authors to discuss the impact of yield uncertainty was Silver [1976], who analyzes how to modify the EOQ order quantity in order to cope with variance in receipt quantities. He considers variance that is either independent of order quantity or directly proportional to it. Yano and Lee [1995] provide an extensive review of the papers on yield uncertainty models. They stress that supply uncertainty is very complex and since closed-form solutions are often unachievable, valid heuristics must be further studied and developed.

Many papers on supply uncertainty focus on supply becoming completely unavailable in the case of disruptions. Parlar and Berkin [1991] analyze the EOQ model with disruptions. Berk and Arreola-Risa [1994] published a correction to Parlar and Berkin's model, addressing logic errors regarding the occurrence of stock-outs and associated costs. Snyder [2009] introduces a closed-form approximation for the problem, and Qi et al. [2007] extend the model to include disruptions at both the supplier and retailer. Parlar and Perry [1995, 1996] extend it to include fixed costs and multiple suppliers.

Some of the other contributors to the supply disruption field who do not focus on the EOQ model include Gupta [1996], who considers fixed lead times in variable supply models and evaluates approximate methodologies for a (Q, r) system with lost sales, Parlar [1997], who considers a (Q, r) system with backordering, and Song and Zipkin [1996], who consider variable lead times and variable order quantities with a dynamic programming approach. Güllü et al. [1997] examine dynamic deterministic demand over finite-horizon and non-stationary disruption probabilities, and relate the optimal base-stock level to the newsboy fractile. Dada et al. [2007] extend the stochastic-demand newsboy model to include multiple unreliable suppliers. Snyder and Shen [2006] simulate inventory systems with either supply disruptions or demand uncertainty to study how the two sources of uncertainty can cause different inventory level and placement decisions to be optimal. Schmitt and Singh [2009] also use simulation to test the impact of different types of uncertainty, including disruptions, stochastic demand, and discrete jumps in demand. They test the advantages of various mitigation strategies for each of these risks.

Generally papers focus on *either* yield uncertainty or supply disruptions. Chopra et al. [2007] model both, analyzing the costs involved in bundling the variance from these two distinct sources in a single-period setting. They compare complete disruptions to additive yield uncertainty and stress the importance of correctly identifying and analyzing the types of stochasticity in supply. Schmitt and Snyder [2009] also consider a system with both yield uncertainty and supply disruptions, extending the analysis to an infinite-horizon setting. They demonstrate the importance of considering the long-term impact of disruptions through multiple-period analysis.

Tomlin [2006] discusses three general strategies for coping with supply disruptions: inventory control, sourcing, and acceptance. Inventory control strategies involve ordering and stocking decisions and can be considered mitigating, proactive techniques. Sourcing strategies are contingency plans and can be reactive to an actual shortage or used proactively in planning for a potential shortage, and involve back-up supplier usage. Acceptance means choosing not to proactively mitigate disruptions. Tomlin formulates an infinite-horizon, periodic-review base-stock system when both an unreliable supplier (subject to disruptions) and/or a more expensive, perfectly reliable supplier are available. He proves that single sourcing is optimal when the firm is risk-neutral and demand is either stochastic or deterministic, but that if the firm is risk-averse, dual-sourcing is often optimal. He presents a formula, a special case of which is presented in Theorem 1 in Section 3 below, for the optimal base-stock level under deterministic demand and random disruptions.

Snyder and Tomlin [2008] examine how inventory systems can be developed to take advantage of advanced warnings of disruptions. They consider a system where the disruption profile changes over time; advanced warning of disruptions can change their anticipated probability of occurrence. They conclude that a threat advisory system can be extremely beneficial and allows increased cost savings especially when the disruption probabilities are significantly different in different states, and that tight capacity reduces the benefit of a threat advisory system. Where Snyder and Tomlin consider how changing disruption probabilities should impact mitigation behavior, Golany et al. [2009] consider how mitigation behavior can impact disruption probabilities. They indicate that when disruptions are "strategic" (caused intentionally, e.g. by terrorists), then mitigation at a location may reduce the likelihood of that location being targeted. They advocate distributing mitigation across a network in order to lower the likelihood overall.

Contracting and financial investment are other methods of mitigating supply risk. Babich [2008] discusses financial investment in a supplier to prevent a supply disruption caused by supplier bankruptcy. He considers profit-sharing as one method of improving the supplier's financial stability. Yang et al. [2009] consider a manufacturer who may pay to increase its information on how reliable a supplier may be, in order to better decide whether to use an alternate supply source. Wagner et al. [2009] provide empirical evidence from the automotive industry that suppliers' disruption risks are typically positively correlated, and Babich et al. [2007] include correlation in their model; they discuss how pricing competition between multiple unreliable suppliers can lead to a diversified supply base and less disruption impact risk for a retailer. When suppliers' disruption correlation is negative, then then can compete based on reliability, whereas when their disruption correlations are positive they must compete more based on cost.

In this paper, we contribute the following tools for analysis of inventory systems subject to supply disruptions:

- Exact and approximate expected cost functions when supply is disrupted and demand is stochastic
- A closed-form approximation for the optimal base-stock level when supply is disrupted and demand is stochastic
- A closed-form approximation for the optimal base-stock level when supply is disrupted and supply yield is stochastic

Throughout the paper, we focus on how the familiar newsboy fractile is a critical value

Table 1: Disrupted-Supply System Notation

Notation	*Definition*
S	Base-stock level for the system
i	Index representing being in a state with i disruptions in a row
π_i	Probability of being in state i
$F(\cdot)$	Cumulative distribution function for the disruption states
d	Deterministic demand at the retailer per period
p	Penalty cost per item per period
h	Holding cost per item per period

in systems with supply uncertainty, as it is for demand uncertainty, since the optimal base-stock policies balance the costs of over-stocking with the risk of costs due to supply shortage from disruptions.

3 Supply Disruptions and Deterministic Demand

This section serves to introduce our method of modeling supply disruptions and discuss known optimal policies for coping with them. We consider deterministic demand and supply yield, but relax these assumptions in subsequent sections.

A full list of the model parameters and variables are given in Table 1. The order of events in a period is as follows: At the beginning of a period, the retailer orders up to the base-stock level, S. It then receives material instantaneously in a non-disrupted period, or receives nothing in a disrupted period. Demand is satisfied from on-hand inventory, and any unsatisfied demand is backordered. Holding or penalty costs are then assessed for the positive or negative ending inventory level. The decision variable for the model is the base-stock level.

Disruptions are modeled using an infinite-state discrete-time Markov chain (DTMC), where the state represents the number of consecutive disrupted periods. We denote π_i as the probability of being in state i and assume that π_i are decreasing with i. We define $F(i) = \sum_{j=0}^{i} \pi_j$ as the *cdf* of this distribution.

The following Proposition gives results that are a special case of the 2-supplier model presented by Tomlin [2006].

Theorem 1 (Tomlin, 2006)**.** *For a retailer with deterministic demand and supply disruptions with* pmf π_i *and* cdf $F(i)$*,*

(a) the expected cost per period is given by

$$E[C] = \sum_{i=0}^{\infty} \pi_i [h(S-(i+1)d)^+ + p((i+1)d-S)^+] \tag{1}$$

(b) the expected cost given in (1) is convex.

(c) $S^* = jd$*, where* j *is the smallest integer such that* $F(j-1) \geq \frac{p}{p+h}$.

Proof: Follows from Tomlin [2006].

Theorem 1 indicates that at optimality, the system will *not* stock out $\frac{p}{p+h}$ percent of periods. These results parallel that of the familiar newsboy model with deterministic supply but stochastic demand, where the optimal order quantity for that model also provides a solution that causes the system to not stock out in $\frac{p}{p+h}$ percent of periods.

4 Supply Disruptions and Stochastic Demand

We now consider a retailer subject to both supply disruptions (but deterministic yield) and random demand. In Section 4.1 we establish the exact expected cost function. In Section 4.2 we develop an approximation that assumes that the demand is stochastic in at most one period per order cycle and deterministic otherwise. We demonstrate how, by choosing which period is treated as stochastic, one also chooses the base-stock level, and we show how to find the optimal choice of the stochastic period. We justify this approach mathematically, using properties of the demand distribution, as well as numerically, in our computational study. We call this the Single Stochastic Period (SSP) approximation. In Section 4.6 we establish other approximations and compare those approaches to the optimal and to the SSP approximation.

We add stochastic demand to the model with *pdf* $m(x)$, where X has a mean of μ per period and variance of σ^2. For i periods of demand, we denote the *pdf* as $m_i(x_i)$ (with no subscript implying $i = 1$); X_i has a mean of $i\mu$ and variance of $i\sigma^2$.

4.1 Exact Model

The expected cost in a period with a successful delivery is:

$$h\int_{-\infty}^{S}(S-x)m(x)dx + p\int_{S}^{\infty}(x-S)m(x)dx \tag{2}$$

Define a cycle as the time between successful deliveries. If the supply is disrupted for $i-1$ periods in a row (thus the cycle is in its i^{th} period), the costs are:

$$h\int_{-\infty}^{S}(S-x_i)m_i(x_i)dx_i + p\int_{S}^{\infty}(x_i-S)m_i(x_i)dx_i \tag{3}$$

Thus we have the following expected cost function:

$$c(S) = \sum_{i=1}^{\infty}\pi_{i-1}\left(h\int_{-\infty}^{S}(S-x_i)m_i(x_i)dx_i + p\int_{S}^{\infty}(x_i-S)m_i(x_i)dx_i\right) \tag{4}$$

Although the summand is similar to the newsboy cost function, the effective demand distribution is different for each term of the sum, and therefore the sum does not collapse into a single newsboy function. The cost in (4) cannot be minimized in closed form.

For the remainder of the paper, we assume the demand is normally distributed. Our analysis holds for any demand distribution with the appropriate loss function substituted into the results below and other minor changes made. For a single-period demand we have $X \sim N(\mu, \sigma^2)$, with $\phi(\cdot)$ and $\Phi(\cdot)$ denoting the *pdf* and *cdf* of the distribution, respectively. For multiple periods, $X_i \sim N(i\mu, i\sigma^2)$. We can then simplify the cost as given in the following proposition.

Proposition 2. *The expected cost for a single retailer subject to normally distributed demand and supply disruptions is convex and is equal to:*

$$c(S) = \sum_{i=1}^{\infty}\pi_{i-1}\left(h(S-i\mu) + \sigma\sqrt{i}(p+h)G\left(\frac{S-i\mu}{\sigma\sqrt{i}}\right)\right) \tag{5}$$

where $G(r) = \int_r^{\infty}(v-r)\phi(v)dv$ represents the standard normal loss function.
Proof: See Appendix, Section A.1.

The derivative of (5) is given in the proof of Proposition 2 as

$$\frac{d}{dS}\,c(S) = \sum_{i=1}^{\infty}\pi_{i-1}\left((p+h)\Phi\left(\frac{S-i\mu}{\sigma\sqrt{i}}\right) - p\right) \tag{6}$$

Seeking a minimizer for (5), we set (6) equal to zero and find:

$$\sum_{i=1}^{\infty} \pi_{i-1} \Phi \left(\frac{S - i\mu}{\sigma\sqrt{i}} \right) = \frac{p}{p+h} \tag{7}$$

Since $\Phi\left(\frac{S-i\mu}{\sigma\sqrt{i}}\right)$ appears in (7) and we cannot separate S from an infinite number of i terms, a closed form minimizer for equation (5) cannot be found directly. Thus we look to approximate this cost with the approximation in the following sections.

Proposition 3. *When the base-stock level S is set optimally for a system with supply disruptions and stochastic demand, the type-1 service level (i.e., the probability that all demands in a given period will be met from stock) equals $\frac{p}{p+h}$.*
Proof: See Appendix, Section A.2.

Thus the familiar optimal service level from the classic newsboy problem and from Theorem 1 is still optimal when both stochastic demand and supply disruptions are present.

4.2 Single Stochastic Period Approximation

Let S be fixed and call I the approximate number of periods' of mean demand stocked in the base-stock level; $S \cong I\mu$. We show here that by approximating the loss function terms for $i \neq I$, we can approximate the demand as deterministic for all periods of a cycle other than the I^{th} (where a cycle is defined as the time between successful deliveries). Thus we call this the Single Stochastic Period (SSP) approximation.

The loss function term involving I, $G\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right)$, is significant since $S \cong I\mu$ and $G(r)$ cannot be approximated easily for small $|r|$. However for $j \neq I$, the term in parenthesis for the loss function, $\left(\frac{S-j\mu}{\sigma\sqrt{j}}\right)$, is either relatively large or small. Thus we can develop approximations for $G\left(\frac{S-j\mu}{\sigma\sqrt{j}}\right)$.

It is well known that $G(r) = \phi(r) - r(1 - \Phi(r))$ [Zipkin, 2000]. Thus if r is a large positive number, corresponding to $j < I$, $G(r) \cong 0$. If r is a large negative number ($j > I$), $G(r) \cong -r$. We use this to write out the SSP cost approximation. Take the sum $\sum_{j=1}^{I-1}$ to equal 0 if $I = 1$. The SSP cost approximation is:

$$\begin{aligned}\widetilde{c}(S) &= \sum_{j=1}^{I-1} \pi_{j-1}\left(h(S-j\mu)\right) + \pi_{I-1}\left(h(S-I\mu) + \sigma\sqrt{I}(p+h)G\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right)\right) + \\ &\quad \sum_{j=I+1}^{\infty} \pi_{j-1}\left(h(S-j\mu) - \sigma\sqrt{j}(p+h)\left(\frac{S-j\mu}{\sigma\sqrt{j}}\right)\right) \\ &= \sum_{j=1}^{I-1} \pi_{j-1}\left(h(S-j\mu)\right) + \pi_{I-1}\left(h(S-I\mu) + \sigma\sqrt{I}(p+h)G\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right)\right) - \\ &\quad \sum_{j=I+1}^{\infty} \pi_{j-1}(p(S-j\mu))\end{aligned} \tag{8}$$

Note that the demand terms in the summations are deterministic; our approximation of the loss terms for periods not equal to I means that, in effect, we assume the demand equals the mean $(j\mu)$ for $j \neq I$.

Although (8) holds for any I, we show in Section 4.3 that, for given values of the input parameters, there is a unique value of I that is valid (in a sense defined below). Therefore, when differentiating $\widetilde{c}(S)$ to optimize it, we may treat I as a constant.

4.3 SSP Solution

We find the solution for (8) in two parts below. The first is the exact solution when the derivative of (8) yields a well-defined solution, and the second is for when it does not.

4.3.1 Well-Defined Solution

Recall that $F(\cdot)$ is the *cdf* of the disruption distribution, and let $F(r) = 0$ for $r < 0$. The solution to (8) is given in the following Proposition.

Proposition 4. *Given the approximate cost in (8) for a retailer subject to uncertain demand and disrupted supply, for fixed I, the $\widetilde{S}$ that minimizes $\widetilde{c}(S)$ is*

$$\widetilde{S} = I\mu + \sigma\sqrt{I}\Phi^{-1}\left(\frac{\frac{p}{p+h} - F(I-2)}{\pi_{I-1}}\right) \tag{9}$$

Proof: See Appendix, Section A.3.

We establish properties for I such that (9) is well defined in the following proposition.

Proposition 5. *There exists at most one I such that both $F(I-2) < \frac{p}{p+h}$ and $F(I-1) > \frac{p}{p+h}$, and the argument to Φ^{-1} in equation (9) is in $(0,1)$ iff both of these inequalities hold.*
Proof: See Appendix, Section A.4.

We discuss how to find the solution when the argument to Φ^{-1} in equation (9) is not in $(0,1)$ in the next section. We had previously taken I as fixed in terms of the input parameters, and Proposition 5 confirms this; I is $I = j+1$ for the minimum j such that $F(j) > \frac{p}{p+h}$. Note that this is exactly the solution given in the deterministic demand case in Theorem 1, with the exception that this is a strict inequality and Theorem 1 also holds for $F(j) = \frac{p}{p+h}$.

4.3.2 Balanced Stock Solution

Suppose $F(j) = \frac{p}{p+h}$ for some j; then by Proposition 5 the solution given in (9) is not well defined for any I; essentially we do not know whether to set $I = j+1$ or $j+2$. We derive an alternate method to set $\widetilde{S}$ for that case in this section.

Let $I_1 = j+1$. Plugging this into (9) gives:

$$\begin{aligned}\widetilde{S}_1 &= (j+1)\mu + \sigma\sqrt{j+1}\Phi^{-1}\left(\frac{\frac{p}{p+h} - F(j-1)}{\pi_j}\right) = (j+1)\mu + \sigma\sqrt{j+1}\Phi^{-1}\left(\frac{\pi_j}{\pi_j}\right) \\ &= (j+1)\mu + \sigma\sqrt{j+1}\Phi^{-1}(1) \qquad (10)\end{aligned}$$

Since $\Phi^{-1}(1) = \infty$, the right-hand side of (10) equals $(j+1)\mu + \infty$. Replace ∞ with M, denoting a very large number, and we have

$$\widetilde{S}_1 = (j+1)\mu + M \qquad (11)$$

Now let $I_2 = j+2$. Plugging this into (9) gives:

$$\widetilde{S}_2 = (j+2)\mu + \sigma\sqrt{j+2}\Phi^{-1}\left(\frac{\frac{p}{p+h} - F(j)}{\pi_{j-1}}\right) = (j+2)\mu + \sigma\sqrt{j+2}\Phi^{-1}(0) \qquad (12)$$

Since $\Phi^{-1}(0) = -\infty$, the right-hand side of (12) equals $(j+2)\mu - \infty$. Again replace ∞ with M and we have

$$\widetilde{S}_2 = (j+2)\mu - M \qquad (13)$$

We now set (11)=(13):

$$(j+1)\mu + M = (j+2)\mu - M \Leftrightarrow M = \frac{\mu}{2} \tag{14}$$

Plugging this in for M guides us to split the difference between the I_1 and I_2 solutions. We call this the balanced solution to our approximatation, since it balances the solutions for $I = j+1$ or $j+2$. Thus if $F(j) = \frac{p}{p+h}$ for some j (we will refer to these points as "jump points"), we set $I = j+1$ and the solution as $\widetilde{S} = \mu(I + \frac{1}{2})$.

4.3.3 Final Solution

We combine the solutions for $\widetilde{S}$ given in Sections 4.3.1 and 4.3.2 above as follows.

$$\widetilde{S} = \begin{cases} \mu\left(I + \frac{1}{2}\right), & \text{if there exists } I \text{ such that } F(I-1) = \frac{p}{p+h}; \\ I\mu + \sigma\sqrt{I}\Phi^{-1}\left(\frac{\frac{p}{p+h} - F(I-2)}{\pi_{I-1}}\right), & \text{for the smallest } I \text{ such that } F(I-1) > \frac{p}{p+h}. \end{cases} \tag{15}$$

This solution yields $\widetilde{S}$ that can be greater than or less than the optimal S^*. The following Proposition describes the approximate solution's behavior based on the input parameters.

Proposition 6. *$\widetilde{S}$ is increasing with μ, increasing with $\frac{p}{p+h}$ for fixed I, and can be either increasing or decreasing with σ.*
Proof: See Appendix, Section A.5.

We discuss the difference between the optimal and approximate costs in the following section, and present numerical results and comparisons for the costs and solutions in Section 4.5.

4.4 Difference Between the Exact and SSP Approximate Cost

We want to evaluate the difference between the true cost and the SSP approximate cost in order to examine the extent to which each input affects the approximation's accuracy. The following Proposition compares the approximate and exact costs.

Proposition 7. *The difference $c(S) - \widetilde{c}(S)$ is always positive, and is given by:*

$$c(S) - \widetilde{c}(S) = \sigma(p+h)\left(\sum_{j=1}^{I-1} \pi_{j-1}\sqrt{j}G\left(\frac{S - j\mu}{\sigma\sqrt{j}}\right) + \sum_{j=I+1}^{\infty} \pi_{j-1}\sqrt{j}G\left(\frac{j\mu - S}{\sigma\sqrt{j}}\right)\right) \tag{16}$$

This difference approaches ∞ *as* σ *approaches* ∞*.*

Proof: See Appendix, Section A.6.

Thus the approximate cost function always underestimates costs. Unfortunately there is no fixed bound on the approximation error for a given S, as indicated by the fact that (16) approaches ∞ as σ approaches ∞. However we show in the following section that while extreme error cases may occur, typically the SSP approximation performs extremely well.

4.5 SSP Numerical Evaluation

We examine the approximation performance in three parts: first we test the solution over a wide range of random inputs to see its general performance. Next we see how specific input parameters affect its performance. Finally we evaluate the approximation specifically at the jump points, where $F(j) = \frac{p}{p+h}$ for some j.

We specify the DTMC for the disruption states using two parameters: α, the probability of a disrupted period following a non-disrupted period (failure probability), and β, the probability of a non-disrupted period following a disrupted period (recovery probability).

4.5.1 Varying All Inputs

We created 1000 data sets, setting $\mu = 100$ and $h = 1$ and generating the other inputs randomly. We drew the newsboy fractile, $\frac{p}{p+h}$, uniformly from [0.5,0.95], the failure probability α from [0.0,0.5], the recovery probability β from [0.1,1], and σ from [0,33.33]. While in reality a failure probability of 50% and recovery probability of only 10% may not be entirely reasonable, we selected these wide ranges in order to thoroughly evaluate the SSP approximation's performance.

In general, the approximation performed extremely well. The average percent error in the absolute value of the difference between the exact and approximate base-stock solutions, that is, $\frac{|\widetilde{S}-S^*|}{S^*}$, was 1.1%, and the average cost increase, $\frac{c(\widetilde{S})-c(S^*)}{c(S^*)}$, was 0.17%. (We calculated S^* numerically using Excel Solver.) The worst three cost increases came when either σ was high or both σ and p were high, generating cost increases of 11%, 25%, and 62%. In all other cases the cost error was less than 7%, and in 99.1% of the cases it was less than 2%. The error was not clearly related to α and β within the ranges tested, but it increased as

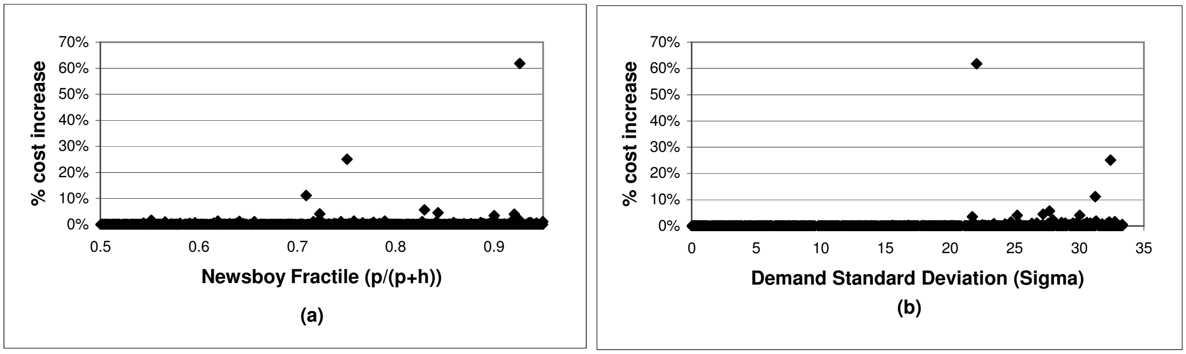

Figure 1: Cost Increase vs. (a) $\frac{p}{p+h}$ and (b) σ

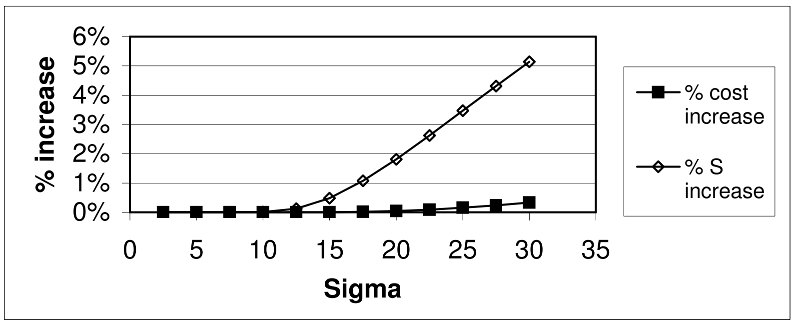

Figure 2: Approximate Solution Percent Cost Increase for Increasing Sigma

the newsboy fractile and σ increase, as shown in Figure 1. With the exception of a few high observations when σ or p are high, there is no obvious trend in the graphs because the error is consistently close to zero.

4.5.2 Varying σ

We tested the solution for increasing σ values, since this is one of the inputs that increases the difference in the costs, as given in (16). We fixed $\mu = 100$, $h = 1$, $p = 20$, $\alpha = 0.2$, and $\beta = 0.4$, and found that as σ increases the approximation error increases. This occurs because larger demand variances causes the approximations we made for the loss function terms, $G(\cdot)$, to be less accurate. A graph of the base-stock solution error and cost increase error of the approximate solution is given in Figure 2.

For this evaluation, the approximation performs very well; the average cost increase for this data sets was 0.07% and maximum was 0.33%. However, some extreme error cases may still occur, as exemplified by the single unusual case in Section 4.5.1 when the cost increase was 62%.

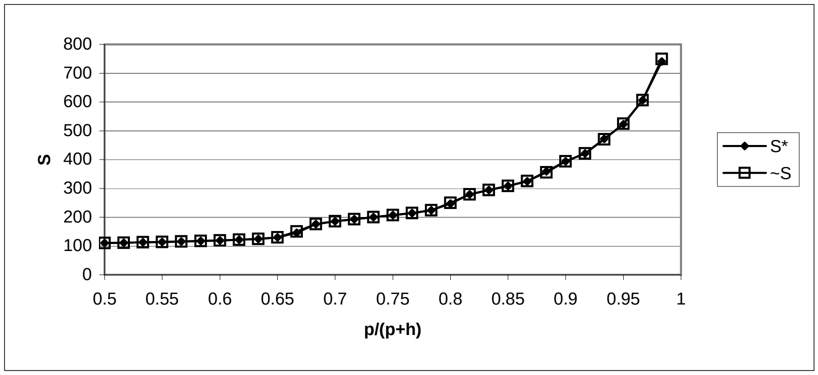

Figure 3: Base-stock Solutions for Increasing Newsboy Fractiles

4.5.3 Varying $\frac{p}{p+h}$

We also tested the approximation performance as the newsboy fractile, $\frac{p}{p+h}$, increases. From Proposition 3, we know that the newsboy fractile is the optimal service level for the system. We fixed σ at 15 and maintained all other inputs as given in Section 4.5.2. We noted that the 1000 random inputs tested in Section 4.5.1 never required the first solution for $\widetilde{S}$ from (15) to be used, where $F(j) = \frac{p}{p+h}$ for some j. For this data set, that solution is used twice, for $j = 1$ and 2. Figure 3 compares the optimal and approximate S solutions.

Clearly the two solutions match very closely, as it is difficult to see the approximate solution behind the optimal values. The two cases where $F(j) = \frac{p}{p+h}$ occur at $\frac{p}{p+h} = \frac{2}{3}$ and $\frac{4}{5}$, yielding approximate S solutions that are 2.8% and 1.1% higher than the optimal solution, respectively. However the cost increases at those points are only 0.003% and 0.002%. The average increase in cost for all $\widetilde{S}$ solutions for this data set was 0.003%.

4.5.4 Varying the Disruption Parameters

We also tested the approximation performance as the disruption parameters change. We reset $p = 20$ and kept $\sigma = 15$, then set $\beta = 0.4$ and increased the failure probability, α, and produced Figure 4. Next we set $\alpha = 0.2$ and increased the recovery probability, β, to produce Figure 5. Both figures show how well the SSP solution matches the optimal solution; the average absolute value of the error in the base-stock solution was 0.5% for Figure 4 (maximum absolute value of the error of 3.2%), and 0.6% for Figure 5 (maximum of 3.6%). In general, the error in using the SSP approximation is less sensitive to changes in the disruption parameters than to changes in the demand parameters. This is because the

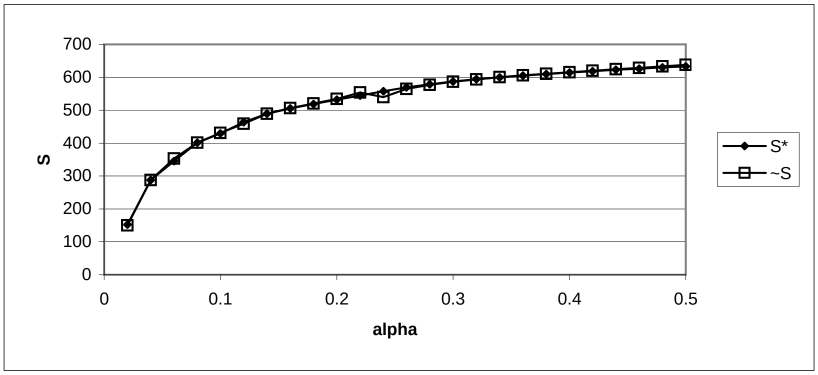

Figure 4: Base-stock Solutions for Increasing Failure Probabilities

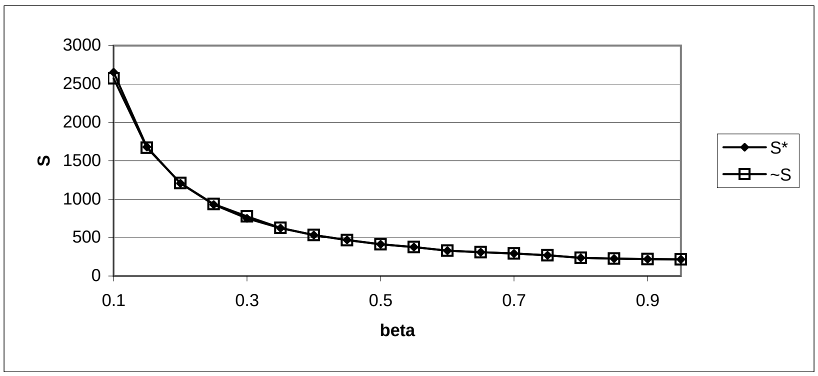

Figure 5: Base-stock Solutions for Increasing Recovery Probabilities

SSP does not approximate the distribution for the disruptions at all, but it does approximate the demand for multiple periods into a single stochastic period.

The figures also demonstrate the extreme inventory levels that may be necessary if disruptions are very likely or expected to be very long. When β is very low, meaning disruptions are very long in length, significant quantities of inventory must be carried. For example, when $\beta = 0.1$ (meaning disruptions average 10 periods in length), the optimal solution indicates that over 25 extra periods' worth of demand should be kept on stock to protect against disruptions.

4.5.5 Results when $F(j) = \frac{p}{p+h}$

None of the 1000 original data sets tested in Section 4.5.1 generated a solution that had $F(j) = \frac{p}{p+h}$ (which we refer to as jump points), requiring the first solution for $\widetilde{S}$ from (15) to be used. This demonstrates the low probability of such a case occurring randomly. To more thoroughly explore the approximate solution at the jump points, we generated 1000 new random data set where we forced this to occur. We kept all random data already generated

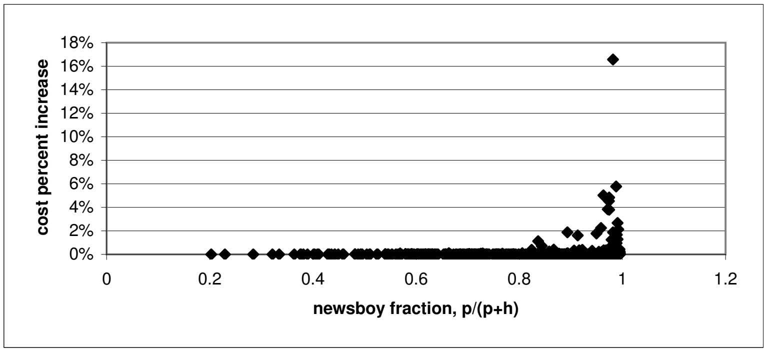

Figure 6: Cost Increase at Newsboy Jump Point Solutions ($p < 200$)

with the exception of p. The penalty cost was determined so that $F(j) = \frac{p}{p+h}$, where the appropriate j was determined so that I was between 1 and 10, I being set randomly and uniformly. We found that for many random data sets generated, $\frac{p}{p+h}$ approached 1 in order to make j as high as was required, and this often made the exact system model unstable (the optimal solution approaches ∞). Thus we generated an excess of random data sets and chose the first 1000 such that $\frac{p}{p+h} \leq 0.999999$ for stability.

The approximation error generally increases with $\frac{p}{p+h}$. When the fractile is unrestricted, we found one instance where the cost error reached 358%. However, when the penalty cost is less than 200 times the holding cost $\left(\frac{p}{p+h} < 0.995\right)$, the average cost error is 0.1% and all but one cost error is less than 6%. Figure 6 demonstrates this. The average cost increase for all newsboy fractiles was 1.3%. Note that while the cost error for this data set is higher on average and at extremes than for the original data set, the likelihood of these cases occurring is also smaller; while $F(j)$ is determined by the characteristics of the supply process, $\frac{p}{p+h}$ is determined independently by the cost structure and it is not likely that these values would be exactly equal.

4.6 Alternate Approximations

The SSP approximation is limited to examining just one stochastic demand period per order cycle since including more than one does not yield a closed-form solution for $\widetilde{S}$. S appears inside multiple $\Phi(\cdot)$ terms, so an inverse cannot be taken to solve for $\widetilde{S}$. We can include more stochastic terms if we approximate the normal distribution with another distribution. In this section, we use the uniform and triangular distributions, as supported by Scherer et al.

[2003], and numerically compare their performance to the SSP approximation in Section 4.7. We first introduce two simple approximations where one of the sources of uncertainty is ignored.

4.6.1 Simple Approximations

Two simple approximate solutions can be found by either assuming that demand stochasticity can be ignored ($\sigma = 0$) or that supply disruptions can be ignored ($\alpha = 0$). The first uses the solution given in Theorem 1:

$$S_{\sigma=0} = jd, \text{ where } j \text{ is the smallest integer such that } F(j-1) \geq \frac{p}{p+h} \tag{17}$$

The second uses the classic newsboy solution for the base-stock level:

$$S_{\alpha=0} = \mu + \sigma\Phi^{-1}\left(\frac{p}{p+h}\right) \tag{18}$$

We compare these solutions with the use of approximations that account for both sources of uncertainty in our numerical evaluations.

4.6.2 Uniform Approximation

If we approximate the normal *cdf* terms in (6) with the uniform *cdf*, then we can include an infinite number of stochastic demand terms since the uniform *cdf* is linear and therefore more tractable. Scherer et al. [2003] propose using a uniform distribution with mean $i\mu$ and range $[i\mu - \sigma\sqrt{3i}, i\mu + \sigma\sqrt{3i}]$ to match the first and second moments (mean and variance) of a normal distribution with mean $i\mu$ and standard deviation $\sigma\sqrt{i}$. This leads to an approximation for the *cdf* of $F_U(x) = \frac{x - i\mu + \sigma\sqrt{3i}}{2\sigma\sqrt{3i}}$. Substituting this into the first derivative of the exact cost, (6), leads to the following approximate solution.

Proposition 8. *If the demand distribution $N(i\mu, \sigma\sqrt{i})$ is approximated with $U(i\mu - \sigma\sqrt{3i}, i\mu + \sigma\sqrt{3i})$, then for a retailer subject to uncertain demand and disrupted supply, the optimal solution to the approximation is*

$$S_U = \frac{1}{\left[\sum_{i=1}^{\infty} \frac{\pi_{i-1}}{\sqrt{i}}\right]}\left(\mu \sum_{i=1}^{\infty} \pi_{i-1}\sqrt{i} + 2\sigma\sqrt{3}\left(\frac{p}{p+h} - \frac{1}{2}\right)\right) \tag{19}$$

Proof: See Appendix, Section A.7.

For the behavior of this solution, we have the following Proposition.

Proposition 9. *S_U is increasing with μ and σ, increasing with $\frac{p}{p+h}$ if $p > h$, and decreasing with $\frac{p}{p+h}$ if $h > p$,.*

Proof: See Appendix, Section A.8.

4.6.3 Triangular Approximation

Suppose instead of considering a single stochastic period, we also consider one period above and below I as well. We can do this if we approximate the normal *cdf* with that of the triangular, since the triangular distribution is piecewise linear. The triangular distribution has been shown to approximate the normal distribution better than the uniform [Scherer et al., 2003].

If we include 3 loss terms instead of just 1 as we did in the SSP cost, (8), the approximate 3-term cost is:

$$\begin{aligned}
\widetilde{c_3}(S) = & \sum_{j=1}^{I-2} \pi_{j-1} \left[h(S - j\mu)\right] + \\
& \pi_{I-2} \left[h(S - (I-1)\mu) + \sigma\sqrt{I-1}(p+h)G\left(\frac{S-(I-1)\mu}{\sigma\sqrt{I-1}}\right)\right] + \\
& \pi_{I-1} \left[h(S - I\mu) + \sigma\sqrt{I}(p+h)G\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right)\right] + \\
& \pi_{I} \left[h(S - (I+1)\mu) + \sigma\sqrt{I+1}(p+h)G\left(\frac{S-(I+1)\mu}{\sigma\sqrt{I+1}}\right)\right] + \\
& \sum_{j=I+2}^{\infty} \pi_{j-1} \left[h(S - j\mu) - \sigma\sqrt{j}(p+h)\left(\frac{S-j\mu}{\sigma\sqrt{j}}\right)\right]
\end{aligned} \quad (20)$$

Note that this assumes $I \geq 2$; if $I \leq 1$, then the term involving π_{I-2} and the first summation are zero. The first summation is also zero if $I = 2$.

This yields the following derivative:

$$\begin{aligned}
\frac{d}{dS}\widetilde{c_3}(S) = & (h+p)F(I-3) - p + (p+h)\left[\pi_{I-2}\Phi\left(\frac{S-(I-1)\mu}{\sigma\sqrt{I-1}}\right) + \right. \\
& \left. \pi_{I-1}\Phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right) + \pi_I\Phi\left(\frac{S-(I+1)\mu}{\sigma\sqrt{I+1}}\right)\right]
\end{aligned} \quad (21)$$

This is where an approximation for the $\Phi(\cdot)$ terms is needed. To approximate the normal distribution with the triangular distribution, we use the following triangular *cdf* as proposed by Scherer et al. [2003]:

$$T(x) = \begin{cases} 0, & x < t_1; \\ \frac{(x-t_1)^2}{(t_3-t_1)(t_2-t_1)}, & t_1 \le x \le t_2; \\ 1 - \frac{(t_3-x)^2}{(t_3-t_1)(t_3-t_2)}, & t_2 < x \le t_3; \\ 1, & x > t_3. \end{cases} \tag{22}$$

where t_1, t_2, and t_3 represent the minimum, mean, and maximum of the possible values for x. To approximate the normal *cdf* with (22), where μ_i and σ_i^2 are the mode and variance of the normal distribution being approximated (so for i demand periods, $\mu_i = i\mu$ and $\sigma_i = \sigma\sqrt{i}$), we use $t_2 = \mu_i$, $t_1 = \mu_i - \sigma_i\sqrt{6}$, and $t_3 = \mu_i + \sigma_i\sqrt{6}$. Thus for the triangular approximation, we have:

$$T(S) = \begin{cases} 0, & S < (\mu_i - \sigma_i\sqrt{6}); \\ \frac{(S-\mu_i+\sigma_i\sqrt{6})^2}{12\sigma_i^2}, & (\mu_i - \sigma_i\sqrt{6}) \le S \le \mu_i; \\ 1 - \frac{(\mu_i+\sigma_i\sqrt{6}-S)^2}{12\sigma_i^2}, & \mu < S \le (\mu_i + \sigma_i\sqrt{6}); \\ 1, & S > (\mu_i + \sigma_i\sqrt{6}). \end{cases} \tag{23}$$

An issue arises around the mean; we do not know whether S is less than, greater than, or equal to $I\mu$, so we are unsure how to exactly approximate the middle term of (21), $\Phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right)$, with the triangular distribution (we do not know how to choose between the 2^{nd} and 3^{rd} cases of (23)). We move forward by testing both cases.

Also, unlike the SSP approximation, I cannot be determined by checking the condition in Proposition 5. Thus we assume that the best I to use here is that given by the deterministic-demand solution in Theorem 1, where I is the minimum I such that $F(I-1) \ge \frac{p}{p+h}$.

Since the triangular *cdf* involves the square of the S term, the solution is in terms of the quadratic equation coefficients. The quadratic coefficients for the two cases which solve the approximation are given in the following Proposition.

Proposition 10. *The solution for the triangular 3-term approximation is given as*

$$S_t = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a} \tag{24}$$

where either, for case 1, $(\mu_I - \sigma_I\sqrt{6}) \le S \le \mu_I$,

$$a = \frac{p+h}{12\sigma^2}\left[\frac{\pi_{I-2}}{I-1} - \frac{\pi_{I-1}}{I} - \frac{\pi_I}{I+1}\right] \tag{25}$$

$$b = \frac{-2(p+h)}{12\sigma^2}\left[\pi_{I-2}\left(\mu - \frac{\sigma\sqrt{6}}{\sqrt{I-1}}\right) - \pi_{I-1}\left(\mu + \frac{\sigma\sqrt{6}}{\sqrt{I}}\right) - \pi_I\left(\mu + \frac{\sigma\sqrt{6}}{\sqrt{I+1}}\right)\right] \quad (26)$$

$$c = (p+h)(F(I-3) + \pi_{I-1} + \pi_I) - p + \frac{p+h}{12\sigma^2}\left[\pi_{I-2}\left((I-1)\mu^2 - 2\mu\sigma\sqrt{6(I-1)} + 6\sigma^2\right) - \pi_{I-1}\left(I\mu^2 + 2\mu\sigma\sqrt{6I} + 6\sigma^2\right) - \pi_I\left((I+1)\mu^2 + 2\mu\sigma\sqrt{6(I+1)} + 6\sigma^2\right)\right] \quad (27)$$

or, for case 2, $\mu < S \leq (\mu_I + \sigma_I\sqrt{6})$,

$$a = \frac{p+h}{12\sigma^2}\left[\frac{\pi_{I-2}}{I-1} + \frac{\pi_{I-1}}{I} - \frac{\pi_I}{I+1}\right] \quad (28)$$

$$b = \frac{-2(p+h)}{12\sigma^2}\left[\pi_{I-2}\left(\mu - \frac{\sigma\sqrt{6}}{\sqrt{I-1}}\right) + \pi_{I-1}\left(\mu - \frac{\sigma\sqrt{6}}{\sqrt{I}}\right) - \pi_I\left(\mu + \frac{\sigma\sqrt{6}}{\sqrt{I+1}}\right)\right] \quad (29)$$

$$c = (p+h)(F(I-3) + \pi_I) - p + \frac{p+h}{12\sigma^2}\left[\pi_{I-2}\left((I-1)\mu^2 - 2\mu\sigma\sqrt{6(I-1)} + 6\sigma^2\right) + \pi_{I-1}\left(I\mu^2 - 2\mu\sigma\sqrt{6I} + 6\sigma^2\right) - \pi_I\left((I+1)\mu^2 + 2\mu\sigma\sqrt{6(I+1)} + 6\sigma^2\right)\right] \quad (30)$$

Proof: See Appendix, Section A.9.

In applying this approximation, we solve for all S_t values and choose the solution which yields the lower expected exact cost.

4.7 Approximation Comparisons

We evaluated the Uniform and Triangular approximation techniques using the same set of 1000 random data sets tested in Section 4.5.1. A summary of the cost error results (approximation solution cost increase above the optimal) is given in Table 2. No approximation performed better than the SSP approximation. When compared to the optimal solution, the Uniform approximation had an average absolute value of the percent error of 34.5% for S and 14.6% for expected cost. The Triangular approximation had average absolute value of the percent errors of 12.1% and 8.2% for S and the expected cost, respectively. There were occasional observations where the Uniform or Triangular approximations outperformed the SSP approximation, but the averages and percentiles for the errors make it clear that the SSP is a more reliable approximation overall. With an average absolute value of error of 1.1% for S and 0.2% for expected cost, the SSP solution clearly outperforms the alternate

Table 2: Approximation Cost Increase Error Results

Approximation	SSP	Uniform	Triangular	$\sigma = 0$	$\alpha = 0$
Average Cost Error	0.2 %	14.6 %	8.2 %	44.5 %	32.45 %
Maximum Cost Error	61.8 %	126.6 %	211.9 %	1539.3 %	390.2 %
Percent < 1% error	97.6 %	15.2 %	48.7 %	45.0 %	18.1 %
Percent < 5% error	99.6 %	33.9 %	70.5 %	55.8 %	39.4 %

approximations; it is better able to capture the stochasticity of the demand with its single stochastic period than the Triangular or Uniform approximations are able to do with three or all stochastic periods.

We also include the cost error for two simple approximations discussed in Section 4.6.1. The first, labeled $\sigma = 0$, assumes deterministic demand and uses the solution given in (17). This solution has the worst average and maximum cost errors of all solutions. The second simple approximation ignores disruptions, labeled $\alpha = 0$; while not as bad as the $\sigma = 0$ solution, the error for this solution is still very high and few solutions have a small error (only 18.1% have less than 1% cost error). The cost error tends to be greater for the $\sigma = 0$ solution when disruption probabilities are low, since optimal costs are also lower in those cases. Low disruption probabilities mean the optimal base-stock level is relatively close to the mean, but the $\sigma = 0$ solution never stocks less than 2 periods' worth of demand and incurs heavy holding costs in those cases. In contrast, the $\alpha = 0$ solution performs very poorly when disruption probabilities are higher. Clearly there is a need to address the stochasticity of both demand and supply in setting inventory base-stock levels.

Figure 7 shows the percent cost error for all 1000 observations, in increasing order of error (sorted separately for each approximation), for each of these approximations. The simple approximations, $\sigma = 0$ and $\alpha = 0$, have error values that extend above the upper limits of the graph. Despite only being able to incorporate a single stochastic demand period into its model, the SSP approximation performs very well and clearly outperforms the other approximations. With 97.6% of its solutions (for the general random data) providing a cost error of less than 1%, and 99.6% proving a cost error less than 5%, it is a practical option for setting base-stock levels in a system with both supply disruptions and stochastic demand.

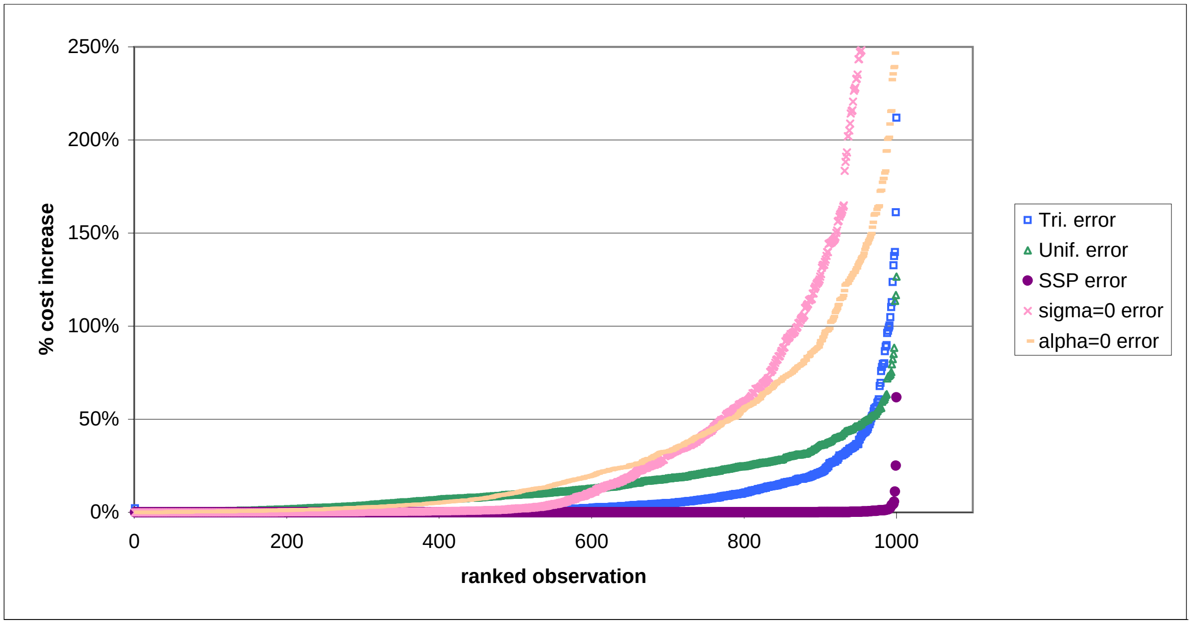

Figure 7: Approximation Cost Error Comparison

5 Supply Disruptions and Stochastic Supply Yield

We now consider another combination of discrete and continuous uncertainty in an inventory system: supply disruptions and yield uncertainty. We model demand as deterministic, equal to d per period, with $id = i$ times d (i periods of demand). We consider additive yield uncertainty, assuming that the quantity received from the supplier is normally distributed with a mean equal to the quantity ordered and standard deviation of σ_y (independent of the order quantity). Note that this means deliveries could be either greater or less than that ordered, but this could be approximately adjusted by adjusting the mean. We do not explicitly include unit costs in our model, so either the supplier or buyer could be held accountable for any excess units delivered. The reader is referred to Yano and Lee [1995] for a review of more complex models of yield uncertainty. The use of additive yield (as opposed to yield that is proportional to the order size) is justified by Chopra et al. [2007] as being realistic in the case where contracts are based on production batches but the exact yield is stochastic (e.g., flu vaccines or semiconductors). This assumption also allows us to formulate the SSP model and produce useful insights that would not be achievable with a proportional yield formulation. Schmitt and Snyder [2009] show that the following is the expected cost for this system:

$$c_y(S) = \sum_{i=1}^{\infty} \left[\pi_{i-1} \left(p(id - S) + (p+h)\sigma_y G\left(\frac{id - S}{\sigma_y}\right)\right)\right] \tag{31}$$

While no closed-form optimal expression can be found for (31), we use it to determine the optimal service level in Proposition 11.

Proposition 11. *When the base-stock level S is set optimally for a system subject to supply disruptions and additive yield, the type-1 service level (i.e., the probability that all demands in a given period will be met from stock) equals $\frac{p}{p+h}$.*
Proof: See Appendix, Section A.10.

We apply the SSP approximation approach by considering only one stochastic period from (31) and approximating the rest. We present this model and numerical results in the following sections.

5.1 Approximate Cost Formulation

Our approximation is that $I \approx \frac{S}{d}$, $G\left(\frac{jd-S}{\sigma_y}\right) \approx 0$ for $j > I$ and $G\left(\frac{jd-S}{\sigma_y}\right) \approx \left(\frac{id-S}{\sigma_y}\right)$ for $j < I$. We formulate our approximate cost as follows.

$$\begin{aligned}
\widetilde{c_y}(S) &= \sum_{j=1}^{I-1} \left[\pi_{j-1} \left(p(jd - S) - (p+h)\sigma_y \left(\frac{jd - S}{\sigma_y}\right)\right)\right] + \\
&\quad \pi_{I-1} \left(p(Id - S) + (p+h)\sigma_y G\left(\frac{Id - S}{\sigma_y}\right)\right) + \sum_{j=I+1}^{\infty} \left[\pi_{j-1}\left(p(jd - S)\right)\right] \\
&= \sum_{j=1}^{I-1} \pi_{j-1}\left(h(S - jd)\right) + \pi_{I-1}\left(p(Id - S) + \sigma_y(p+h)G\left(\frac{Id - S}{\sigma_y}\right)\right) - \\
&\quad \sum_{j=I+1}^{\infty} \pi_{j-1}(p(S - jd))
\end{aligned} \tag{32}$$

5.2 Approximate Model Solution

In solving this system, we again found that I is determined by the input parameters to be the minimum I such that $F(I-1) > \frac{p}{p+h}$. Proof of this is omitted since it follows the same approach as the proof of Proposition 5. We take a derivative of (32) to find the exact optimal solution when this inequality holds. For the case when $F(I-1) = \frac{p}{p+h}$ for some I, we again

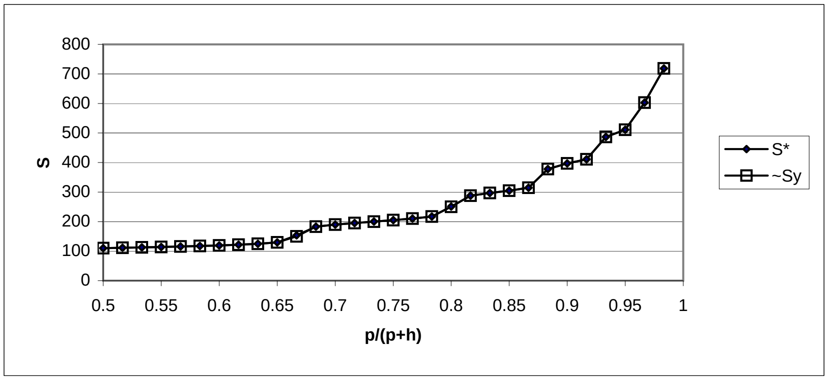

Figure 8: Base-stock Solutions for Increasing Newsboy Fractiles with Uncertain Yield

applied the balanced stock argument from Section 4.3.2 to determine the base-stock level. This leads to the following solution to the system.

Proposition 12. *Given the approximate cost in (32) for a retailer subject to disrupted supply and yield uncertainty, the optimal base-stock level for this model is:*

$$\widetilde{S_y} = \begin{cases} d\left(I + \frac{1}{2}\right), & \textit{if there exists } I \textit{ such that } F(I-1) = \frac{p}{p+h}\textit{; else} \\ Id - \sigma_y \Phi^{-1}\left(\frac{F(I-1) - \frac{p}{p+h}}{\pi_{I-1}}\right), & \textit{for the smallest } I \textit{ such that } F(I-1) > \frac{p}{p+h}. \end{cases} \tag{33}$$

Proof: Follows the same approach as the proofs of Propositions 4 and 5.

5.3 Numerical Evaluation

In order to see how the approximation performs for both cases when $F(j)$ does and does not equal $\frac{p}{p+h}$ for all j, we again tested data with inputs $\alpha = 0.2$, $\beta = 0.4$, $\mu = 100$, $\sigma = 15$, and $h = 1$ for increasing $\frac{p}{p+h}$, which generates $F(I-1) = \frac{p}{p+h}$ for $I = 1$ and 2. A graph of the optimal and approximate S solutions is given in Figure 8. Since the approximate solution is barely visible behind the optimal solution, clearly the approximation performs very well. The highest cost increase from any of the approximate solutions shown in Figure 8 is 0.00065%.

We also tested 1000 random data sets from Section 4.5.1, making the σ from that data set equal to the σ_y for this system. The SSP approximation performed very well; the average cost increase was 0.02%, with a maximum occurrence of 5.9% increase. The percent with a cost error less than 1% was 99.7%.

In order to again test the case where $F(j) = \frac{p}{p+h}$, we also tested the data set from Section 4.5.5 where we force this case to occur. Recall that this case never occurred in the

1000 original random sets (so the likelihood of it occurring naturally is low), but we generated p values which made this case occur. The approximation had an average cost increase of 0.3% for this data, with a maximum increase of 65%. For the data where $p < 200$, the average cost increase is only 0.1% and the maximum observed is 5.1%.

Both the average and extreme error cases for the SSP approximation performance for this system were better than the system with demand uncertainty. Since the standard deviation of the yield uncertainty is not proportional to $\sqrt{I}$ (as it is for the demand uncertainty case), this makes the approximation more accurate for larger I for this system. This helps improve the accuracy of the SSP performance for the yield uncertainty system.

6 Conclusions

The SSP approach performs well for modeling a system subject to supply disruptions. It provides a closed-form base-stock solution, which is valuable for researchers and practitioners alike. Researchers may embed it in larger models, or examine the impact of input parameters. Practitioners can more easily implement and update a closed-form solution.

The results of this paper also demonstrate that supply disruptions can have significant negative impact on a retailer if it has not proactively protected itself against them. We have examined three cases with supply disruptions in this paper: (1) deterministic demand and deterministic supply yield, (2) deterministic demand and stochastic, additive supply yield, and (3) deterministic supply yield and stochastic demand. It is interesting to note the impact of disruptions on a retailer in these cases if it does not proactively mitigate them. Since case (1) is entirely deterministic, it would carry no safety stock and disruptions would have the largest impact in this system. In the absence of disruptions and with equal standard deviation (either on the demand or the supply yield), cases (2) and (3) would stock the same safety stock quantity. Therefore they would be equally affected by disruptions.

Note, however, that safety stock maintained to protect against regular demand or yield variability is only a fraction of a single period of demand. Thus if disruptions are moderate in duration (greater than a single period), all three cases would suffer shortages of full demand quantities. Clearly disruptions must be protected against, regardless of the other sources of uncertainty which are already mitigated in the system.

Throughout this paper, we have presented several properties of a system subject to supply disruptions in order to allow retailers to establish best practices for inventory management in such a setting. We have shown how the optimal base-stock level can be determined by the familiar newsboy fractile. When demand or yield are stochastic, we presented a closed-form Single Stochastic Period approximate solution that yields very good results. The results from this paper can help firms proactively and cost effectively protect against supply disruption risk.

7 Acknowledgements

This research was supported in part by National Science Foundation grants DGE-9972780, DMI-0522725, and DMI-0621433. This support is gratefully acknowledged. We are also thankful for the helpful suggestions provided by anonymous referees.

References

S. Axsäter. *Inventory Control*. Kluwer Academic Publishers, Boston, MA, first edition, 2000.

V. Babich. Independence of capacity ordering and financial subsidies to risky suppliers. Working paper, Dept. of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, 2008.

V. Babich, A.N. Burnetas, and P. H. Ritchken. Competition and diversification effects in supply chains with supplier default risk. *Manufacturing & Service Operations Management*, 9(2):123–146, 2007.

E. Berk and A. Arreola-Risa. Note on "Future supply uncertainty in EOQ models". *Naval Research Logistics*, 41:129–132, 1994.

S. Chopra and P. Meindl. *Supply Chain Management*. Pearson Prentice Hall, Upper Saddle River, NJ, second edition, 2004.

S. Chopra, G. Reinhardt, and U. Mohan. The importance of decoupling recurrent and disruption risks in a supply chain. *Naval Research Logistics*, 54(5):544–555, 2007.

M. Dada, N. Petruzzi, and L. Schwarz. A newsvendor's procurement problem when suppliers are unreliable. *Manufacturing & Service Operations Management*, 9(1):9–32, 2007.

B. Golany, E. H. Kaplan, A. Marmur, and U. G. Rothblum. Nature plays with dice - terrorists do not: Allocating resources to counter strategic versus probabilistic risks. *European Journal of Operational Research*, 192:198–208, 2009.

R. Güllü, E. Önol, and N. Erkip. Analysis of a deterministic demand production/inventory system under nonstationary supply uncertainty. *IIE Transactions*, 29:703–709, 1997.

D. Gupta. The (Q,r) inventory system with an unreliable supplier. *INFOR*, 34(2):59–76, 1996.

M. Parlar. Continuous-review inventory problem with random supply interruptions. *European Journal of Operational Research*, 99:366–385, 1997.

M. Parlar and D. Berkin. Future supply uncertainty in EOQ models. *Naval Research Logistics*, 38:107–121, 1991.

M. Parlar and D. Perry. Analysis of a (Q,r,T) inventory policy with deterministic and random yields when future supply is uncertain. *European Journal of Operational Research*, 84:431–443, 1995.

M. Parlar and D. Perry. Inventory models of future supply uncertainty with single and multiple suppliers. *Naval Research Logistics*, 43:191–210, 1996.

L. Qi, Z. J. M. Shen, and L. V. Snyder. A continuous-review inventory model with disruptions at both supplier and retailer. Working paper, Forthcoming in *Production and Operations Management*, 2007.

W. T. Scherer, T. A. Pomroy, and D. N. Fuller. The triangular density to approximate the normal density: decision rules-of-thumb. *Reliability Engineering and System Safety*, 82: 331–341, 2003.

A. J. Schmitt and M. Singh. A quantitative analysis of disruption risk in a multi-echelon supply chain. Working paper, MIT Center for Transportation and Logistics, Cambridge, MA, 2009.

A. J. Schmitt and L. V. Snyder. Infinite-horizon models for inventory control under yield uncertainty and disruptions. Working paper, P.C. Rossin College of Engineering and Applied Sciences, Lehigh University, Bethlehem, PA, 2009.

E. A. Silver. Establishing the order quantity when the amount received is uncertain. *INFOR*, 14(1):32–39, 1976.

L. V. Snyder. A tight approximation for a continuous-review inventory model with supplier disruptions. Working paper, P.C. Rossin College of Engineering and Applied Sciences, Lehigh University, Bethlehem, PA, 2009.

L. V. Snyder and Z. J. M. Shen. Supply and demand uncertainty in multi-echelon supply chains. Working paper, P.C. Rossin College of Engineering and Applied Sciences, Lehigh University, Bethlehem, PA, 2006.

L. V. Snyder and B. Tomlin. Inventory management with advanced warning of disruptions. Working paper, P.C. Rossin College of Engineering and Applied Sciences, Lehigh University, Bethlehem, PA, 2008.

J.-S. Song and P. H. Zipkin. Inventory control with information about supply conditions. *Management Science*, 42(10):1409–1419, 1996.

B. Tomlin. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science*, 52(5):639–657, 2006.

S. M. Wagner, C. Bode, and P. Koziol. Supplier default dependencies: Empirical evidence from the automotive industry. *European Journal of Operational Research*, 199(1):150–161, 2009.

Z. Yang, G. Aydin, V. Babich, and D. R. Bell. Supply disruptions, asymmetric information and a backup production option. *Management Science*, 55(2):192–209, 2009.

C. A. Yano and H. L. Lee. Lot sizing with random yields: A review. *Operations Research*, 43(2):311–334, 1995.

P. H. Zipkin. *Foundations of Inventory Management*. McGraw-Hill Higher Education, Boston, MA, first edition, 2000.

A Appendix

A.1 Proof of Proposition 2, Section 4.1

In Part 1 of this proof we establish the formula for $c(S)$, and in Part 2 we establish its convexity.

A.1.1 Formulation of $c(S)$

We use the following formulas:

$$\int_r^\infty (v-r)f(v)dv = \sigma G(z) \tag{34}$$

$$\int_{-\infty}^r (r-v)f(v)dv = \sigma z + \sigma G(z) \tag{35}$$

where $z = \frac{r-\mu}{\sigma} \sim N(0,1)$ and $G(\cdot)$ is the standard normal loss function. These follow immediately from standard sources [e.g., Axsäter, 2000, Chopra and Meindl, 2004, Zipkin, 2000]. Thus we have:

$$\begin{aligned} c(S) &= \sum_{i=1}^{\infty} \pi_{i-1} \left[h\left(\sigma\sqrt{i} \left(\frac{S-i\mu}{\sigma\sqrt{i}} \right) + \sigma\sqrt{i} G\left(\frac{S-i\mu}{\sigma\sqrt{i}} \right) \right) + p\sigma\sqrt{i} G\left(\frac{S-i\mu}{\sigma\sqrt{i}} \right) \right] \\ &= \sum_{i=1}^{\infty} \pi_{i-1} \left(h(S-i\mu) + \sigma\sqrt{i}(p+h) G\left(\frac{S-i\mu}{\sigma\sqrt{i}} \right) \right) \end{aligned} \tag{36}$$

A.1.2 Convexity of $c(S)$

Using $G'(x) = \Phi(x) - 1$ [Axsäter, 2000], where $\Phi(\cdot)$ and $\phi(\cdot)$ are the standard normal *cdf* and *pdf*, respectively:

$$\begin{aligned} \frac{d}{dS}\, c(S) &= \sum_{i=1}^{\infty} \pi_{i-1} \left(h + (p+h) \left(\Phi\left(\frac{S-i\mu}{\sigma\sqrt{i}} \right) - 1 \right) \right) \\ &= \sum_{i=1}^{\infty} \pi_{i-1} \left((p+h)\Phi\left(\frac{S-i\mu}{\sigma\sqrt{i}} \right) - p \right) \end{aligned} \tag{37}$$

$$\frac{d^2}{dS^2}\, c(S) = \sum_{i=1}^{\infty} \pi_{i-1} \left(\frac{p+h}{\sigma\sqrt{i}} \right) \phi\left(\frac{S-i\mu}{\sigma\sqrt{i}} \right) \geq 0 \tag{38}$$

Therefore $c(S)$ is convex.

□

A.2 Proof of Proposition 3, Section 5

If a disruption has lasted i periods ($i \geq 0$), then the probability that all demands in the current period are met from stock is given by $\Phi\left(\frac{S-i\mu}{\sigma\sqrt{i}}\right)$. Therefore, the left-hand side of (7) is the steady-state probability of not stocking out in a given period. Since the optimal S satisfies (7), the result follows.

□

A.3 Proof of Proposition 4, Section 4.3

The derivative of (8) is:

$$\begin{aligned} \frac{d}{dS}\widetilde{c}(S) &= \sum_{j=1}^{I-1} \pi_{j-1} h + \pi_{I-1}\left(h + (p+h)\left(\Phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right) - 1\right)\right) - \sum_{j=I+1}^{\infty} \pi_{j-1} p \\ &= hF(I-2) - p\left(1 - F(I-1)\right) + (p+h)\pi_{I-1}\Phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right) - p\pi_{I-1} \\ &= hF(I-2) + p(F(I-1) - \pi_{I-1}) - p + (p+h)\pi_{I-1}\Phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right) \\ &= (p+h)F(I-2) - p + (p+h)\pi_{I-1}\Phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right) \qquad (39) \end{aligned}$$

and

$$\frac{d^2}{dS^2}\widetilde{c}(S) = \frac{(p+h)\pi_{I-1}}{\sigma\sqrt{I}}\phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right) \qquad (40)$$

Since this is non-negative for all S, the cost function is convex, and we can set the first derivative equal to zero to solve for the optimal solution $\widetilde{S}$ to the approximation:

$$\Phi\left(\frac{S-I\mu}{\sigma\sqrt{I}}\right) = \frac{p-(p+h)F(I-2)}{(p+h)\pi_{I-1}} \qquad (41)$$

$$\widetilde{S} = I\mu + \sigma\sqrt{I}\Phi^{-1}\left(\frac{\frac{p}{p+h} - F(I-2)}{\pi_{I-1}}\right) \qquad (42)$$

□

A.4 Proof of Proposition 5, Section 4.3

Since the disruption *cdf* is an increasing function, at most one I can satisfy the conditions given in Proposition 5. Clearly, the $(\cdot)$ term in (9) is > 0 iff $F(I-2) < \frac{p}{p+h}$. It is also < 1

iff $F(I-1) > \frac{p}{p+h}$:

$$\frac{\frac{p}{p+h} - F(I-2)}{\pi_{I-1}} < 1 \Leftrightarrow \frac{p}{p+h} < \pi_{I-1} + F(I-2) \Leftrightarrow \frac{p}{p+h} < F(I-1) \tag{43}$$

□

A.5 Proof of Proposition 6, Section 4.3.3

We have two solutions for $\widetilde{S}$ in (15):

$$\widetilde{S}_1 = \mu\left(I + \frac{1}{2}\right) \tag{44}$$

if there exists I such that $F(I-1) = \frac{p}{p+h}$, and

$$\widetilde{S}_2 = I\mu + \sigma\sqrt{I}\Phi^{-1}\left(\frac{\frac{p}{p+h} - F(I-2)}{\pi_{I-1}}\right) \tag{45}$$

for the smallest I such that $F(I-1) > \frac{p}{p+h}$ if such an I exists.

Note that I is independent of μ. The derivative of (44) with respect to μ is $\left(I + \frac{1}{2}\right)$; therefore the first solution for $\widetilde{S}$ is increasing with μ. The derivative of (45) with respect to μ is I; therefore the first solution for $\widetilde{S}$ is also increasing with μ.

$\widetilde{S}_1$ is clearly increasing with $\frac{p}{p+h}$ since I is increasing with $\frac{p}{p+h}$. For $\widetilde{S}_2$, for a fixed I, the term inside the parentheses for Φ^{-1} is increasing with $\frac{p}{p+h}$, therefore Φ^{-1} is increasing and $\widetilde{S}$ is increasing with $\frac{p}{p+h}$.

$\widetilde{S}_1$ is independent of σ. $\widetilde{S}_2$ can be increasing or decreasing, since its derivative with respect to σ is $\sqrt{I}\Phi^{-1}\left(\frac{\frac{p}{p+h}-F(I-2)}{\pi_{I-1}}\right)$, which is positive if $\left(\frac{\frac{p}{p+h}-F(I-2)}{\pi_{I-1}}\right) > \frac{1}{2}$ and negative if $\left(\frac{\frac{p}{p+h}-F(I-2)}{\pi_{I-1}}\right) < \frac{1}{2}$.

□

Note: Since increasing $\frac{p}{p+h}$ can also potentially increase I, thus increasing $F(I-2)$ and decreasing the term inside the parenthesis for Φ^{-1}, we cannot conclusively say that $\widetilde{S}$ is increasing as I is increasing. We have found this to be true in numerical evaluations, though.

A.6 Proof of Proposition 7, Section 4.4

From (5) and (8), we have:

$$
\begin{aligned}
c(S) - \tilde{c}(S) &= \sum_{j=1}^{I-1} \pi_{j-1} \left(\sigma \sqrt{j}(p+h) G\left(\frac{S - j\mu}{\sigma\sqrt{j}}\right)\right) + \\
&\quad \sum_{j=1+1}^{\infty} \pi_{j-1} \left((p+h)(S - j\mu) + \sigma\sqrt{j}(p+h) G\left(\frac{S - j\mu}{\sigma\sqrt{j}}\right)\right) \\
&= (p+h) \left[\sum_{j=1}^{I-1} \pi_{j-1} \sigma \sqrt{j} G\left(\frac{S - j\mu}{\sigma\sqrt{j}}\right) + \right. \\
&\quad \left. \sum_{j=I+1}^{\infty} \pi_{j-1} \left((S - j\mu) + \sigma\sqrt{j} G\left(\frac{S - j\mu}{\sigma\sqrt{j}}\right)\right)\right] \qquad (46)
\end{aligned}
$$

Using $G(-r) = r + G(r)$ [Axsäter, 2000] (which holds for all $x \in \mathbb{R}$), we can reduce the second summation terms as follows:

$$
\begin{aligned}
(S - j\mu) + \sigma\sqrt{j} G\left(\frac{S - j\mu}{\sigma\sqrt{j}}\right) &= (S - j\mu) + (j\mu - S) + \sigma\sqrt{j} G\left(\frac{j\mu - S}{\sigma\sqrt{j}}\right) \\
&= \sigma\sqrt{j} G\left(\frac{j\mu - S}{\sigma\sqrt{j}}\right) \qquad (47)
\end{aligned}
$$

Thus:

$$
c(S) - \tilde{c}(S) = \sigma(p+h) \left(\sum_{j=1}^{I-1} \pi_{j-1} \sqrt{j} G\left(\frac{S - j\mu}{\sigma\sqrt{j}}\right) + \sum_{j=I+1}^{\infty} \pi_{j-1} \sqrt{j} G\left(\frac{j\mu - S}{\sigma\sqrt{j}}\right)\right) \qquad (48)
$$

Since all the terms in (48) are positive, this difference is always positive.

Then

$$
\begin{aligned}
\sigma &\rightarrow \infty \Rightarrow \\
\lim_{\sigma\to\infty} c(S) - \tilde{c}(S) &= (p+h)G(0) \left(\sum_{j=1}^{I-1} \pi_{j-1}\sqrt{j} + \sum_{j=I+1}^{\infty} \pi_{j-1}\sqrt{j}\right) \lim_{\sigma\to\infty} \sigma \qquad (49)
\end{aligned}
$$

Since $G(0) = 0.3989$, this limit equals ∞.

□

A.7 Proof of Proposition 8, Section 4.6.2

Substituting the uniform distribution into (5) and taking the derivative yields the following:

$$
\frac{d}{dS}\, c_U(S) = \sum_{i=1}^{\infty} \pi_{i-1} \left((p+h) \left(\frac{S - i\mu + \sigma\sqrt{3i}}{2\sigma\sqrt{3i}} \right) - p \right) \qquad (50)
$$

Since the coefficient of S in 50 is positive, we conclude that the function is still convex and we can optimize by setting this derivative equal to zero. Therefore,

$$\frac{d}{dS}\, c_U(S) = 0 \Leftrightarrow \tag{51}$$

$$\sum_{i=1}^{\infty} \pi_{i-1} \left(\frac{S_U - i\mu + \sigma\sqrt{3i}}{2\sigma\sqrt{3i}} \right) = \frac{p}{p+h} \Leftrightarrow \tag{52}$$

$$\frac{S_U}{2\sigma\sqrt{3}} \sum_{i=1}^{\infty} \frac{\pi_{i-1}}{\sqrt{i}} = \frac{\mu}{2\sigma\sqrt{3}} \sum_{i=1}^{\infty} \pi_{i-1}\sqrt{i} + \frac{p}{p+h} - \frac{1}{2} \tag{53}$$

$$S_U = \frac{1}{\sum_{i=1}^{\infty} \frac{\pi_{i-1}}{\sqrt{i}}} \left(\mu \sum_{i=1}^{\infty} \pi_{i-1}\sqrt{i} + 2\sigma\sqrt{3} \left(\frac{p}{p+h} - \frac{1}{2} \right) \right) \tag{54}$$

□

A.8 Proof of Proposition 9, Section 4.6.2

The derivative of (19) with respect to μ is $\left[\sum_{i=1}^{\infty} \pi_{i-1}\sqrt{i}\right] \Big/ \left[\sum_{i=1}^{\infty} \frac{\pi_{i-1}}{\sqrt{i}}\right] > 0$; therefore S_U is increasing with μ.

The derivative of (19) with respect to σ is $\left[2\sqrt{3}\left(\frac{p}{p+h} - \frac{1}{2}\right)\right] \Big/ \left[\sum_{i=1}^{\infty} \frac{\pi_{i-1}}{\sqrt{i}}\right] > 0$; therefore S_U is also increasing with σ.

Finally, the derivative of (19) with respect to $\frac{p}{p+h}$ is $\left[2\sigma\sqrt{3}\right] \Big/ \left[\sum_{i=1}^{\infty} \frac{\pi_{i-1}}{\sqrt{i}}\right] > 0$; therefore S_U is increasing with $\frac{p}{p+h}$ if $p > h$ and decreasing if $p < h$.

□

A.9 Proof of Proposition 10, Section 4.6.3

We present the triangular approximation of the derivative below and solve for the quadratic equation coefficients for both cases. Again note that if $I \leq 2$, $F(I-3) = 0$ and that if $I = 1$, then all terms involving π_{I-2} must be made equal to zero.

A.9.1 Approximation if $S \geq I\mu$: case 1

$$\begin{aligned}
\frac{d}{ds}\widehat{c}_3(S_{\geq I\mu}) &= (p+h)F(I-3) - p + \\
&(p+h)\left[\frac{\pi_{I-2}}{12(I-1)\sigma^2}\Big(S - \Big[(I-1)\mu - \sigma\sqrt{6(I-1)}\Big]\Big)^2 + \right.\\
&\pi_{I-1}\left(1 - \frac{1}{12I\sigma^2}\Big(\Big[I\mu + \sigma\sqrt{6I}\Big] - S\Big)^2\right) + \\
&\left.\pi_I\left(1 - \frac{1}{12(I+1)\sigma^2}\Big(\Big[(I+1)\mu + \sigma\sqrt{6(I+1)}\Big] - S\Big)^2\right)\right] \\
&= (p+h)(F(I-3) + \pi_{I-1} + \pi_I) - p + \\
&\frac{p+h}{12\sigma^2}\left[\frac{\pi_{I-2}}{I-1}\Big(S^2 - 2S\Big[(I-1)\mu - \sigma\sqrt{6(I-1)}\Big] + \right.\\
&\Big[(I-1)\mu - \sigma\sqrt{6(I-1)}\Big]^2\Big) - \\
&\frac{\pi_{I-1}}{I}\Big(S^2 - 2S\Big[I\mu + \sigma\sqrt{6I}\Big] + \Big[I\mu + \sigma\sqrt{6I}\Big]^2\Big) - \\
&\left.\frac{\pi_I}{I+1}\Big(S^2 - 2S\Big[(I+1)\mu + \sigma\sqrt{6(I+1)}\Big] + \Big[(I+1)\mu + \sigma\sqrt{6(I+1)}\Big]^2\Big)\right]
\end{aligned} \tag{55}$$

This yields us the following coefficients for the quadratic equation, $\frac{-b \pm \sqrt{b^2-4ac}}{2a}$:

$$a = \frac{p+h}{12\sigma^2}\left[\frac{\pi_{I-2}}{I-1} - \frac{\pi_{I-1}}{I} - \frac{\pi_I}{I+1}\right] \tag{56}$$

$$b = \frac{-2(p+h)}{12\sigma^2}\left[\pi_{I-2}\left(\mu - \frac{\sigma\sqrt{6}}{\sqrt{I-1}}\right) - \pi_{I-1}\left(\mu + \frac{\sigma\sqrt{6}}{\sqrt{I}}\right) - \pi_I\left(\mu + \frac{\sigma\sqrt{6}}{\sqrt{I+1}}\right)\right] \tag{57}$$

$$\begin{aligned}
c &= (p+h)(F(I-3) + \pi_{I-1} + \pi_I) - p + \\
&\frac{p+h}{12\sigma^2}\Big[\pi_{I-2}\Big((I-1)\mu^2 - 2\mu\sigma\sqrt{6(I-1)} + 6\sigma^2\Big) - \\
&\pi_{I-1}\Big(I\mu^2 + 2\mu\sigma\sqrt{6I} + 6\sigma^2\Big) - \pi_I\Big((I+1)\mu^2 + 2\mu\sigma\sqrt{6(I+1)} + 6\sigma^2\Big)\Big]
\end{aligned} \tag{58}$$

A.9.2 Triangular Approximation if $S \leq I\mu$: case 2

$$
\begin{aligned}
\frac{d}{ds}\widehat{c_3}(S_{\leq I\mu}) &= (p+h)F(I-3) - p + \\
& (p+h)\left[\frac{\pi_{I-2}}{12(I-1)\sigma^2}\Big(S - \Big[(I-1)\mu - \sigma\sqrt{6(I-1)}\Big]\Big)^2 + \right. \\
& \frac{\pi_{I-1}}{12I\sigma^2}\Big(S - \Big[I\mu - \sigma\sqrt{6I}\Big]\Big)^2 + \\
& \left.\pi_I\left(1 - \frac{1}{12(I+1)\sigma^2}\Big(\Big[(I+1)\mu + \sigma\sqrt{6(I+1)}\Big] - S\Big)^2\right)\right] \\
&= (p+h)(F(I-3) + \pi_I) - p + \\
& \frac{p+h}{12\sigma^2}\left[\frac{\pi_{I-2}}{I-1}\Big(S^2 - 2S\Big[(I-1)\mu - \sigma\sqrt{6(I-1)}\Big] + \right. \\
& \Big[(I-1)\mu - \sigma\sqrt{6(I-1)}\Big]^2\Big) + \\
& \frac{\pi_{I-1}}{I}\Big(S^2 - 2S\Big[I\mu - \sigma\frac{\pi_I}{I+1}\Big(S^2 - 2S\Big[(I+1)\mu + \sigma\sqrt{6(I+1)}\Big] + \\
& \left.\Big[(I+1)\mu + \sigma\sqrt{6(I+1)}\Big]^2\Big)\right] \qquad (59)
\end{aligned}
$$

This yields us the following coefficients for the quadratic equation, $\frac{-b\pm\sqrt{b^2-4ac}}{2a}$:

$$a = \frac{p+h}{12\sigma^2}\left[\frac{\pi_{I-2}}{I-1} + \frac{\pi_{I-1}}{I} - \frac{\pi_I}{I+1}\right] \qquad (60)$$

$$b = \frac{-2(p+h)}{12\sigma^2}\left[\pi_{I-2}\left(\mu - \frac{\sigma\sqrt{6}}{\sqrt{I-1}}\right) + \pi_{I-1}\left(\mu - \frac{\sigma\sqrt{6}}{\sqrt{I}}\right) - \pi_I\left(\mu + \frac{\sigma\sqrt{6}}{\sqrt{I+1}}\right)\right] \qquad (61)$$

$$
\begin{aligned}
c &= (p+h)(F(I-3)+\pi_I) - p + \frac{p+h}{12\sigma^2}\left[\pi_{I-2}\left((I-1)\mu^2 - 2\mu\sigma\sqrt{6(I-1)} + 6\sigma^2\right) + \right. \\
& \left.\pi_{I-1}\left(I\mu^2 - 2\mu\sigma\sqrt{6I} + 6\sigma^2\right) - \pi_I\left((I+1)\mu^2 + 2\mu\sigma\sqrt{6(I+1)} + 6\sigma^2\right)\right] \qquad (62)
\end{aligned}
$$

□

A.10 Proof of Proposition 11, Section 5

Setting the derivative of (31) equal to zero leads to [Schmitt and Snyder, 2009]:

$$\sum_{i=1}^{\infty}\pi_{i-1}\Phi\left(\frac{id - S}{\sigma_y}\right) = \frac{p}{p+h} \qquad (63)$$

Then the same argument applies as given in the proof of Proposition 3, Section A.2.

□