

e - c o m p a n i o n

ONLY AVAILABLE IN ELECTRONIC FORM

Electronic Companion—"Adaptive Data-Driven Inventory Control with Censored Demand Based on Kaplan-Meier Estimator" by Woonghee Tim Huh, Retsef Levi, Paat Rusmevichientong, and James B. Orlin, *Operations Research*, doi 10.1287/opre.1100.0906.

Appendix A: Alternate Censoring Assumption

In the main body of the paper, we have assumed the observability of an indicator of the form $\delta_t = \mathbf{I}[D_t \leq Y_t]$ in each period t . However, in certain applications, it may not be possible to distinguish the events $[D_t = Y_t]$ and $[D_t > Y_t]$. In this appendix, we address the case where the indicator available in each period t is given by $\tilde{\delta}_t = \mathbf{I}[D_t < Y_t]$. We will describe how to modify the algorithm such that it converges to $\psi^* + 1$, instead of the newsvendor quantile ψ^* .

The main idea behind this modification is to maintain the equivalence of information between the original censoring indicator and the new indicator. We achieve this through an invariant that, for each sample path, the inventory level under the new censoring indicator is exactly one more than the inventory level under the original indicator.

We first describe the modification to the algorithm. Under the new censoring assumption, let $(\tilde{Y}_1, \tilde{Y}_2, \dots)$ denote the sequence of inventory levels and let $(\tilde{\delta}_1, \tilde{\delta}_2, \dots)$ denote the sequence of censoring outcomes, where $\tilde{\delta}_t = \mathbf{I}[D_t < \tilde{Y}_t]$. Since demand is discrete with an integer support, it follows that $\tilde{\delta}_t = \mathbf{I}[D_t \leq \tilde{Y}_t - 1]$. Then, since the censored data is now given in the form of the “less-than” operator, it is possible to obtain the Kaplan-Meier estimator as before. In the initial period, set $\tilde{Y}_0 = Y_0 + 1$. In each period $t \geq 0$, we let the inventory level of the next period, $\tilde{Y}_{t+1}$, be the newsvendor quantile of the Kaplan-Meier empirical distribution in the current period *plus* one.

We now discuss how the modified algorithm is related to the algorithm under the original assumption. For our discussion here, fix a sample path of demand realizations $(D_1, D_2, \dots)$. Let $(Y_1, Y_2, \dots)$ be the sequence of inventory levels under the original censoring assumption, and let $(\delta_1, \delta_2, \dots)$ be the corresponding censoring outcome, where $\delta_t = \mathbf{I}[D_t \leq Y_t]$ for each period t . Then, an inductive argument shows that both $\tilde{Y}_t = Y_t + 1$ and $\tilde{\delta}_t = \mathbf{I}[D_t \leq \tilde{Y}_t - 1] = \mathbf{I}[D_t \leq Y_t] = \delta_t$ hold for each period t . These results imply that, for each period, the empirical distributions constructed by the Kaplan-Meier algorithm under the original censoring assumption and the new assumption are indeed the same.

Now, since the results of this paper show that Y_t converges to ψ^* , where ψ^* is the newsvendor quantile of the true demand distribution, it follows that $\tilde{Y}_t = Y_t + 1$ converges to $\psi^* + 1$.

Appendix B: Approximate Confidence Intervals

In this section, we develop approximate confidence intervals. The main result, stated in Lemma 1, establishes upper and lower bounds for an approximate “confidence interval” around $\bar{F}(y)$ for any $y \leq \psi^*$.

Define $\sigma_y^2 = \text{Var}(\mathbf{I}[D > y])$. For any $y \geq 0$, let $N_y(T)$ denote the number of times that the KM estimator is greater than or equal to y ; that is, $N_y(T) = \#\{t = 1, \dots, T \mid Y_t^{KM} \geq y\}$.

LEMMA 1. *Let $y \leq \psi^*$. Let W be an independent random variable having a normal distribution with mean zero and variance σ_y^2 . For each $a \in \mathbb{R}$,*

$$P \left\{ W \leq a - \frac{\sigma_y}{\bar{F}(y)} \right\} \leq \lim_{T \rightarrow \infty} P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \leq a \right\} \leq P \{W \leq a\} .$$

The proof of Lemma 1 can be found in Appendix B.1. We note that the above result provides an estimate of the confidence interval associated with any value y less than ψ^* . To see this, note that for any $a \leq b$,

$$\begin{aligned} & \lim_{T \rightarrow \infty} P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \in (a, b) \right\} \\ &= \lim_{T \rightarrow \infty} P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \leq b \right\} \\ & \quad - \lim_{T \rightarrow \infty} P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \leq a \right\} . \end{aligned}$$

Since $\sigma_y^2 \leq 1/4$ and $\bar{F}(y) \geq \bar{F}(\psi^*) \geq 1 - r$, we obtain $\sigma_y/\bar{F}(y) \leq 1/(2(1 - r))$, which implies from Lemma 1 that

$$\lim_{T \rightarrow \infty} P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \in (a, b) \right\}$$

is bounded below by

$$P \left\{ a \leq W \leq b - \frac{\sigma_y}{\bar{F}(y)} \right\} \geq P \left\{ a \leq W \leq b - \frac{1}{2(1 - r)} \right\} ,$$

and bounded above by

$$P \left\{ a - \frac{\sigma_y}{\bar{F}(y)} \leq W \leq b \right\} \leq P \left\{ a - \frac{1}{2(1 - r)} \leq W \leq b \right\} .$$

Let us briefly highlight the main challenges in improving the lower bound in Lemma 1. Although we know that $N_y(T)/T$ converges to one almost surely as T increases to infinity, to obtain a tighter lower bound, we need to determine the grow rate of $N_y(T)$. It follows from the proof of Lemma 1 that if $T - N_y(T) = o(\sqrt{T})$ with probability one (equivalently, $\lim_{T \rightarrow \infty} \frac{T - N_y(T)}{\sqrt{T}} = 0$), then we can show that

$$\lim_{T \rightarrow \infty} P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \leq a \right\} = P \{W \leq a\} ,$$

which would give us the exact confidence interval around $\bar{F}(y)$. However, determining the growth rate of the random variable $N_y(T)$ remains an open research question.

B.1. Proof of Lemma 1

We first establish Lemma 2, which is useful in proving Lemma 1.

For any $y \geq 0$, recall that $N_y(T) = \#\{t = 1, \dots, T \mid Y_t^{KM} \geq y\}$ denotes the number of times that the KM estimator is greater than y . Let $t(l)$ be the l^{th} time period this happens.

LEMMA 2. For any $T \geq 1$ and $y \geq 0$,

$$\frac{1}{T} \sum_{t=1}^T \mathbf{I}[D_t > y] \cdot \mathbf{I}[Y_t^{KM} \geq y] \leq \bar{F}_T(y) \leq \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{T} + \frac{T - N_y(T)}{T} .$$

Proof: The lower bound is obtained by applying the same argument as (??) in the proof of Lemma ?? . More specifically, it follows from

$$\bar{F}_T(y) = \prod_{t: Z_{(t)} \leq y} \left(\frac{T-t}{T-t+1} \right)^{\delta_{(t)}} \geq \prod_{t: Z_{(t)} \leq y} \left(\frac{T-t}{T-t+1} \right) = \frac{T - \#\{t: Z_t \leq y\}}{T} .$$

For the upper bound, observe that $\{t \in \{1, \dots, T\} : Z_{(t)} \leq y\}$ is the disjoint union of $\{t \in \{1, \dots, T\} : Y_{(t)} \leq y\}$ and $\{t \in \{1, \dots, T\} : D_{(t)} \leq y < Y_{(t)}\}$. Thus,

$$\begin{aligned} \bar{F}_T(y) &= \prod_{t: Z_{(t)} \leq y} \left(\frac{T-t}{T-t+1} \right)^{\delta_{(t)}} \\ &= \prod_{t: Y_{(t)} \leq y} \left(\frac{T-t}{T-t+1} \right)^{\delta_{(t)}} \prod_{t: D_{(t)} \leq y < Y_{(t)}} \left(\frac{T-t}{T-t+1} \right) \\ &\leq \prod_{t: D_{(t)} \leq y < Y_{(t)}} \left(\frac{T-t}{T-t+1} \right) . \end{aligned}$$

Note that in the above product, the set of factors are distinct and form a subset of $\{(T - k)/(T - k + 1), k = 1, \dots, T\}$. Furthermore, the number of factors is given by $N_y(T) - \sum_{l=1}^{N_y(T)} \mathbf{I}[D_{t(l)} > y]$.

Thus, an upper bound on the above product is

$$\frac{T-1}{T} \cdot \frac{T-2}{T-1} \cdots \frac{T - (N_y(T) - \sum_{l=1}^{N_y(T)} \mathbf{I}[D_{t(l)} > y])}{T - (N_y(T) - \sum_{l=1}^{N_y(T)} \mathbf{I}[D_{t(l)} > y]) - 1} = \frac{T - N_y(T) + \sum_{l=1}^{N_y(T)} \mathbf{I}[D_{t(l)} > y]}{T}$$

completing the proof. ■

Below is the proof of Lemma 1.

Proof:

$$\begin{aligned} \bar{F}_T(y) &\geq \frac{1}{T} \sum_{t=1}^T \mathbf{I}[D_t > y] \cdot \mathbf{I}[Y_t^{KM} \geq y] = \frac{N_y(T)}{T} \cdot \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} \\ \bar{F}_T(y) &\leq \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{T} + \frac{T - N_y(T)}{T} = \frac{N_y(T)}{T} \cdot \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} + \frac{T - N_y(T)}{T}, \end{aligned}$$

which implies that with probability one

$$\begin{aligned} \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) &\geq \frac{N_y(T)}{T} \left(\sqrt{N_y(T)} \left\{ \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} - \bar{F}(y) \right\} \right), \\ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) &\leq \frac{N_y(T)}{T} \left(\sqrt{N_y(T)} \left\{ \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} - \bar{F}(y) \right\} \right) \\ &\quad + \frac{\sqrt{N_y(T)} (T - N_y(T))}{T}. \end{aligned}$$

Therefore, we obtain

$$\begin{aligned} &P \left\{ \frac{N_y(T)}{T} \left(\sqrt{N_y(T)} \left\{ \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} - \bar{F}(y) \right\} \right) + \frac{\sqrt{N_y(T)} (T - N_y(T))}{T} \leq a \right\} \\ &\leq P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \leq a \right\} \\ &\leq P \left\{ \frac{N_y(T)}{T} \left(\sqrt{N_y(T)} \left\{ \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} - \bar{F}(y) \right\} \right) \leq a \right\}. \end{aligned}$$

Note that by Lemma ??, $N_y(T)/T \rightarrow 1$ almost surely as $T \rightarrow \infty$. Moreover, since $D_{t(\ell)}$'s are IID, it follows from the classical Central Limit Theorem that the random variable

$$\sqrt{N_y(T)} \left\{ \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} - \bar{F}(y) \right\}$$

converges weakly to a normal random variable with mean zero and variance σ_y^2 . Thus, by Slutsky's Theorem, the random variable

$$\frac{N_y(T)}{T} \left(\sqrt{N_y(T)} \left\{ \frac{\sum_{\ell=1}^{N_y(T)} \mathbf{I}[D_{t(\ell)} > y]}{N_y(T)} - \bar{F}(y) \right\} \right)$$

converges weakly to a normal random variable with mean zero and variance σ_y^2 . To obtain the desired result, it follows from the proof of Lemma ?? that with probability one

$$\lim_{T \rightarrow \infty} \frac{\sqrt{N_y(T)}(T - N_y(T))}{T} \leq \frac{\sigma_y}{\bar{F}(y)}.$$

Therefore,

$$P \left\{ W \leq a - \frac{\sigma_y}{\bar{F}(y)} \right\} \leq \lim_{T \rightarrow \infty} P \left\{ \sqrt{N_y(T)} \left(\bar{F}_T(y) - \frac{N_y(T)}{T} \bar{F}(y) \right) \leq a \right\} \leq P \{ W \leq a \},$$

which is the desired result. ■

Appendix C: Description of Sales-Driven Policies

- **KM-MYOPIC:** The version of the KM-MYOPIC policy that we test is a slight modification of the policies described in Section ?. (Nevertheless, the analysis presented in Section ? is still applicable for this variant.) Our implementation of the KM-MYOPIC policy maintains the largest uncensored sales data point, and we keep track of the number of times that Y_t^{KM} – the estimated $b/(b+h)$ -quantile under the Kaplan-Meier empirical distribution – equals to the largest data point, and among those time periods, how often we get censored observations. When the proportion of times that we get censored observations from using the largest data point exceeds $h/(2 \cdot (b+h))$, this suggests that our order-up-to levels are too low. When this happens, we double the magnitude of the largest data point and use it as our new order-up-to level.

During our experiments, we also observe that when the lost sales penalty b is large, the long-run average cost under the KM-MYOPIC policy is sensitive to the order-up-to level that we use initially. Since we do not have a lot of data during the initial periods, our implementation of the KM-MYOPIC policy maintains a single fictitious data point during the first 20 periods. This fictitious data point corresponds to the initial order-up-to level used in the first period. We note that this single data point is only used during the first 20 periods of the algorithm.

• **AIM POLICY:** The AIM algorithm is a stochastic gradient-based algorithm that adjusts the inventory level in each period based on the *sales data* observed in the previous period. We provide a brief description of the algorithm below. For more details, the reader is referred to Huh and Rusmevichientong ?. Under the AIM policy, the order-up-to level in period $t + 1$ for the problem instance j is given by

$$y_{j,t+1} = \begin{cases} [y_{j,t} - \epsilon_t h]^+, & \text{if we have any inventory left at the end of period } t, \\ y_{j,t} + \epsilon_t b, & \text{otherwise,} \end{cases}$$

where $\epsilon_t = 10 / (\max\{b, h\} \cdot \sqrt{t})$. This policy can also be implemented without any prior knowledge of the underlying demand distribution. The order-up-to level in each period depends only on the sales data in the previous period, independent of any other information about the underlying demand distribution.

• **BURNETAS-SMITH POLICY:** Under the BURNETAS-SMITH Policy, the order-up-to level in period $t + 1$ for the j^{th} problem instance is given by

$$y_{j,t+1} = \begin{cases} \left(1 - \frac{h}{(b+h) \cdot t}\right) \cdot y_{j,t}, & \text{if we have any inventory left at the end of period } t, \\ \left(1 + \frac{b}{(b+h) \cdot t}\right) \cdot y_{j,t}, & \text{otherwise.} \end{cases}$$

As in the AIM POLICY, the order-up-to levels under the BURNETAS-SMITH POLICY only depend on the sales data in the previous period. The two policies, however, differ in the computation of the order-up-to levels. See ? for more details.

• **CAVE POLICY:** The CAVE POLICY attempts to approximate the objective function through a series of piecewise linear approximations. See ? for more details.

• **UNCENSORED DEMAND BENCHMARK:** Under the UNCENSORED DEMAND BENCHMARK, we observe the realization of uncensored demand in each period. We then set the inventory level to the newsvendor quantile of the empirical distribution based on demand realizations from all previous periods. Thus, for any j^{th} problem instance, the order-up-to level in period $t + 1$ is given by

$$\inf \left\{ x : \frac{1}{t} \sum_{s=1}^t \mathbf{1}[d_{j,s} \leq x] \geq b/(b+h) \right\},$$

where $d_{j,1}, \dots, d_{j,t}$ denote the *uncensored* demand realizations during the first t periods in the j^{th} problem instance. We note that the UNCENSORED DEMAND BENCHMARK is not implementable unless historical uncensored demand data is available.

- NEWSVENDOR BENCHMARK: Under the NEWSVENDOR BENCHMARK, we set the inventory level after ordering in each period to the newsvendor quantile ψ^* of demand distribution, where $\psi^* = \inf \{x : F(x) \geq b/(b+h)\}$.

We note that the NEWSVENDOR BENCHMARK requires knowledge of the underlying demand distribution.

To the best of our knowledge, the AIM, BURNETAS-SMITH and CAVE policies are the only algorithms in the literature that admit provable convergence with censored demand data

Appendix D: Tables and Graphs

Demand Distribution	Back-Ordered Cost (b)	Newsvendor quantile (ψ^*)	Newsvendor Benchmark Cost	Initial Inventory Level	% Difference from Newsvendor Benchmark After 500 Periods				
					Uncensored Demand	CAVE	Burnetas-Smith	AIM	KM-Myopic
Negative Binomial (80,0.5)	3	88	16.38	10	4.34%	58.16%	16.83%	18.45%	11.17%
				20	4.39%	40.68%	7.86%	13.61%	12.62%
				50	2.88%	7.67%	3.24%	4.66%	8.05%
				100	1.28%	81.18%	1.78%	1.68%	4.04%
				150	2.61%	81.66%	2.99%	21.84%	6.58%
				200	4.13%	83.44%	5.81%	121.50%	10.35%
				250	3.70%	83.18%	9.91%	360.55%	12.57%
Negative Binomial (40,0.333)	3	90	20.13	300	3.67%	83.01%	17.98%	666.37%	13.93%
				10	4.34%	46.88%	14.33%	15.84%	11.24%
				20	3.84%	32.82%	6.60%	11.71%	10.43%
				50	2.67%	6.74%	3.07%	4.32%	6.63%
				100	1.69%	64.65%	2.49%	1.89%	3.77%
				150	2.01%	64.92%	2.20%	15.41%	5.19%
				200	2.24%	65.72%	3.34%	89.49%	7.10%
Negative Binomial (80,0.5)	9	96	23.16	250	3.10%	65.27%	7.54%	275.67%	11.05%
				300	5.32%	67.92%	15.59%	522.85%	14.88%
				10	7.60%	134.02%	26.29%	41.60%	24.10%
				20	7.58%	96.71%	12.34%	31.16%	18.91%
				50	4.23%	20.32%	4.59%	9.22%	11.68%
				100	3.28%	187.83%	5.96%	2.26%	7.87%
				150	2.59%	188.55%	3.78%	70.29%	9.18%
Negative Binomial (40,0.333)	9	100	28.61	200	3.31%	190.07%	46.01%	282.44%	15.47%
				250	3.68%	192.59%	157.46%	497.89%	20.19%
				300	5.33%	191.89%	278.31%	714.20%	26.18%
				10	6.90%	110.40%	23.01%	36.70%	20.09%
				20	7.04%	78.24%	11.20%	27.41%	17.06%
				50	5.73%	18.64%	6.01%	10.96%	13.57%
				100	3.31%	0.62%	4.07%	1.62%	10.56%
Negative Binomial (80,0.5)	49	108	32.61	150	4.59%	155.44%	5.44%	45.04%	8.21%
				200	2.82%	152.31%	27.49%	209.56%	11.37%
				250	5.09%	154.25%	108.47%	383.78%	17.76%
				300	4.71%	153.37%	204.88%	557.80%	21.02%
				10	28.53%	554.34%	96.99%	188.85%	88.90%
				20	29.38%	404.31%	46.62%	140.21%	68.50%
				50	17.21%	92.45%	19.05%	45.41%	46.80%
Negative Binomial (40,0.333)	49	114	40.86	100	9.04%	0.62%	16.85%	4.54%	23.14%
				150	10.55%	769.28%	64.43%	97.04%	18.18%
				200	10.61%	767.22%	200.66%	250.08%	26.47%
				250	13.83%	781.80%	337.05%	403.26%	36.14%
				300	10.39%	769.74%	474.05%	557.05%	40.12%
				10	25.18%	456.21%	83.83%	168.07%	74.62%
				20	23.63%	327.18%	40.24%	128.50%	57.84%
Negative Binomial (80,0.5)	99	112	36.19	50	17.56%	81.11%	18.33%	55.49%	45.20%
				100	12.61%	4.44%	35.09%	11.34%	30.12%
				150	11.19%	622.20%	32.55%	57.48%	16.11%
				200	12.00%	618.54%	140.36%	179.80%	23.80%
				250	11.59%	622.73%	249.39%	302.23%	29.41%
				300	11.90%	625.57%	357.86%	424.10%	32.62%
				10	55.19%	1041.31%	188.16%	369.55%	156.98%
Negative Binomial (40,0.33)	99	120	45.67	20	52.03%	750.28%	84.13%	270.61%	124.57%
				50	35.27%	171.33%	34.08%	101.85%	92.62%
				100	17.66%	0.59%	28.45%	9.79%	50.38%
				150	16.81%	1420.43%	70.17%	85.60%	30.16%
				200	19.87%	1413.55%	200.58%	223.79%	33.64%
				250	19.00%	1420.46%	330.61%	361.60%	48.63%
				300	21.62%	1424.85%	460.89%	499.66%	54.57%
Negative Binomial (80,0.5)	199	115	39.29	10	47.98%	823.92%	161.32%	319.08%	150.73%
				20	45.92%	592.49%	73.23%	253.03%	101.55%
				50	32.98%	148.07%	30.02%	112.05%	85.15%
				100	20.06%	2.32%	23.17%	20.48%	52.68%
				150	21.06%	1135.01%	35.40%	47.27%	34.05%
				200	16.52%	1146.30%	138.35%	156.74%	30.27%
				250	18.54%	1144.54%	241.26%	265.82%	33.30%
Negative Binomial (40,0.33)	199	125	50.87	300	18.41%	1145.66%	344.33%	375.06%	35.20%
				10	104.75%	1914.41%	350.00%	694.41%	316.55%
				20	94.98%	1347.48%	160.46%	509.62%	250.36%
				50	68.18%	330.61%	60.04%	209.43%	161.37%
				100	35.33%	2.64%	31.18%	33.54%	101.64%
				150	30.20%	2604.05%	67.49%	74.77%	71.87%
				200	40.54%	2635.11%	191.05%	201.96%	64.90%
Negative Binomial (80,0.5)	199	125	50.87	250	39.37%	2636.37%	315.00%	329.55%	64.62%
				300	39.89%	2648.63%	437.89%	456.07%	76.73%
				10	91.53%	1503.22%	316.35%	612.09%	292.24%
				20	78.31%	1075.06%	144.14%	488.01%	214.25%
				50	56.02%	278.92%	50.06%	216.01%	162.39%
				100	35.16%	3.36%	43.26%	53.55%	120.03%
				150	34.01%	2070.59%	30.02%	35.28%	60.24%
Negative Binomial (40,0.33)	199	125	50.87	200	33.68%	2044.34%	125.10%	133.52%	57.08%
				250	31.62%	2050.85%	220.05%	231.29%	65.72%
				300	33.91%	2076.87%	315.81%	329.85%	64.75%


Table 1 Impact of different initial starting inventory level on the performance of each policy (no warm-up period and censoring corresponds to demand strictly bigger than the inventory).

Coefficient of Variation	Back-Ordered Cost (b)	Newsvendor quantile (ψ^*)	Newsvendor Benchmark Cost	Initial Inventory Level	% Difference from Newsvendor Benchmark After 500 Periods				
					Uncensored Demand	CAVE	Burnetas-Smith	AIM	KM-Myopic
2	49	755	978.23	100	4.76%	31.69%	87.28%	137.09%	9.20%
				200	4.72%	16.27%	51.99%	91.36%	8.46%
				500	4.64%	0.90%	9.33%	13.17%	7.73%
				1000	4.44%	58.49%	5.02%	7.59%	7.74%
				2000	4.72%	57.55%	71.53%	94.21%	8.46%
				3000	4.57%	58.16%	160.96%	197.55%	10.33%
				4000	4.91%	58.91%	247.88%	296.63%	12.90%
	99	973	1207.89	100	8.71%	73.47%	181.61%	273.57%	16.83%
				200	8.85%	43.29%	113.88%	192.59%	15.42%
				500	8.35%	6.53%	32.36%	49.48%	13.38%
				1000	8.25%	124.13%	1.36%	0.11%	12.54%
				2000	8.65%	124.85%	49.24%	58.14%	13.12%
				3000	8.51%	125.23%	125.24%	140.44%	14.89%
				4000	8.69%	124.19%	200.27%	220.58%	17.08%
	199	1200	1445.16	100	16.68%	157.76%	349.88%	522.48%	27.43%
				200	16.17%	95.92%	230.82%	377.82%	24.63%
				500	15.57%	22.60%	76.23%	117.32%	21.36%
				1000	15.71%	0.29%	5.32%	5.03%	21.49%
				2000	15.36%	251.24%	31.30%	34.94%	20.72%
				3000	15.81%	246.33%	92.93%	99.27%	22.63%
				4000	15.40%	248.46%	160.40%	169.06%	23.60%
4	49	1299	2266.41	100	4.47%	23.15%	62.70%	70.01%	8.47%
				200	4.59%	16.29%	44.86%	53.80%	8.57%
				500	4.67%	5.52%	18.38%	23.50%	7.73%
				1000	4.82%	0.24%	2.82%	2.56%	8.38%
				2000	4.61%	32.63%	6.66%	9.11%	8.22%
				3000	4.92%	32.49%	30.33%	38.77%	9.38%
				4000	4.58%	32.93%	58.82%	75.97%	9.88%
	99	1982	3039.46	100	8.09%	69.11%	140.27%	153.87%	16.41%
				200	8.34%	53.44%	103.56%	123.42%	16.77%
				500	8.32%	28.35%	57.56%	71.42%	14.84%
				1000	8.04%	6.56%	18.82%	23.20%	14.39%
				2000	8.10%	82.88%	0.71%	0.01%	14.02%
				3000	8.27%	84.17%	9.67%	11.68%	15.26%
				4000	8.33%	84.68%	30.43%	36.12%	14.53%
	199	2740	3865.14	100	14.91%	149.28%	266.41%	292.86%	30.39%
				200	14.74%	126.93%	211.74%	247.96%	28.33%
				500	14.44%	75.01%	124.43%	153.58%	25.43%
				1000	14.20%	31.34%	56.27%	68.74%	22.79%
				2000	14.88%	1.73%	7.68%	8.09%	23.96%
				3000	15.16%	180.88%	1.19%	0.52%	23.40%
				4000	14.12%	181.33%	10.92%	12.11%	23.01%
10	49	812	4508.75	100	4.08%	0.97%	4.76%	4.85%	4.90%
				200	4.45%	0.65%	3.26%	3.36%	5.08%
				500	4.48%	0.07%	0.72%	0.67%	5.73%
				1000	4.17%	1.51%	0.34%	0.27%	6.21%
				2000	4.10%	1.59%	5.28%	6.08%	6.70%
				3000	4.83%	1.70%	14.61%	16.61%	8.29%
				4000	4.18%	1.60%	24.79%	28.74%	8.17%
	99	2651	7563.44	100	7.04%	13.77%	21.60%	21.90%	12.20%
				200	7.20%	11.63%	20.85%	21.44%	12.34%
				500	7.30%	7.77%	13.73%	14.54%	12.55%
				1000	7.40%	3.34%	6.77%	7.06%	12.60%
				2000	7.09%	0.18%	1.00%	0.91%	12.36%
				3000	7.29%	14.52%	0.42%	0.39%	12.35%
				4000	7.41%	15.82%	2.01%	2.24%	12.87%
	199	5561	11565.13	100	13.10%	47.84%	66.32%	66.91%	25.33%
				200	12.95%	43.98%	59.65%	60.92%	25.32%
				500	13.11%	34.36%	45.59%	47.23%	23.63%
				1000	12.01%	24.01%	31.83%	33.34%	21.55%
				2000	13.07%	12.21%	16.60%	17.47%	21.82%
				3000	12.57%	5.16%	7.84%	8.26%	22.18%
				4000	12.59%	1.16%	2.36%	2.38%	23.15%

Table 2 Performance of various policies under different coefficient of variations and starting initial inventory level. The demand distribution is a negative binomial with mean 100.


% Difference from Newsvendor Benchmark

% Diff from Newsvendor Benchmark Under Different Inventory Policies
50 instances, $b=3$, Initial = 300, Demand $\sim \text{Unif}\{0, 1, \dots, 100\}$
(dash lines correspond to the 95% confidence interval)




(a) $b = 3$

% Diff from Newsvendor Benchmark Under Different Inventory Policies
50 instances, $b=9$, Initial = 300, Demand $\sim \text{Unif}\{0, 1, \dots, 100\}$
(dash lines correspond to the 95% confidence interval)



(b) $b = 9$

% Diff from Newsvendor Benchmark Under Different Inventory Policies
50 instances, $b=49$, Initial = 300, Demand $\sim \text{Unif}\{0, 1, \dots, 100\}$
(dash lines correspond to the 95% confidence interval)



(c) $b = 49$

Figure 1 The relative difference of the running average cost over time compared to the newsvendor benchmark for 50 randomly generated problem instances for lost sales penalty costs b ranging from 3 to 49. The demand distribution is assumed to be a discrete uniform distribution over the set of integers from 0 to 100. Dash lines correspond to 95% confidence intervals.




Figure 2 The log-log plot of the difference between the NEWSVENDOR BENCHMARK and the average cost over time under the KM-MYOPIC policy, with the lost sales penalty cost ranging from 3 to 49. The cost is averaged over 200 problem instances. The demand in each period is uniformly distributed over the set of integers from 0 to 100.






Figure 3 The percentage difference (compared to the newsvendor benchmark) of the running average cost over time under different inventory policies for 50 randomly generated problem instances for a Poisson demand distribution with mean 80 when the newsvendor quantile equals to 0.75, 0.9, and 0.98, respectively.



Negative Binomial(80,0.5) with variance = 160





Negative Binomial(40,0.333) with variance = 240



(a) $r^* = 0.75$ ($b = 3$)



(b) $r^* = 0.9$ ($b = 9$)



(c) $r^* = 0.98$ ($b = 49$)

Figure 4 The running average cost under Negative Binomial demand distributions with increasing standard deviations.






Figure 5 The running average cost under Pareto and Lognormal demand distributions.



Demand Distribution is Pareto(2,40)



Demand Distribution is LogNormal(3.88,1)

Figure 6 The sampled standard deviation of the running average costs over time under different policies for Pareto and lognormal demand distributions with $b = 19$.




Figure 7 The running average cost under Gaussian demand distributions with increasing standard deviations.




Figure 8 Impact of an alternative censoring assumption on the performance of the KM-MYOPIC policy.