Electronic Companion for The Data-Driven Censored Newsvendor Problem

Appendix A: Useful Facts

A.1. Closed-form expression for difference in newsvendor costs

We repeatedly rely on the following closed-form expression for the difference between the newsvendor cost under two arbitrary ordering quantities q_1 and q_2.

PROPOSITION EC.1. *For any distribution $F \in \mathcal{G}$, $q_1, q_2 \in [0, M]$:*

$$C_F(q_1) - C_F(q_2) = b(q_2 - q_1) + (b+h)\mathbb{E}_F[(q_1 - D)\mathbb{1}\{D \le q_1\} - (q_2 - D)\mathbb{1}\{D \le q_2\}]. \tag{EC.1}$$

Proof. By definition, for any q:

$$\begin{aligned} C_F(q) &= b\mathbb{E}_F[(D-q)^+] + h\mathbb{E}_F[(q-D)^+] \\ &= b\mathbb{E}_F[(D-q)\mathbb{1}\{D > q\}] + h\mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\}] \\ &= b\mathbb{E}_F[(D-q)(1-\mathbb{1}\{D \le q\})] + h\mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\}] \\ &= b\mathbb{E}_F[D-q] - b\mathbb{E}_F[(D-q)\mathbb{1}\{D \le q\}] + h\mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\}] \\ &= b\mathbb{E}_F[D-q] + (b+h)\mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\}]. \end{aligned} \tag{EC.2}$$

Applying (EC.2) to q_1 and q_2, we obtain:

$$\begin{aligned} C_F(q_1) - C_F(q_2) &= b\mathbb{E}_F[D-q_1] - b\mathbb{E}_F[D-q_2] + (b+h)\mathbb{E}_F[(q_1-D)\mathbb{1}\{D \le q_1\}] \\ &\quad - (b+h)\mathbb{E}_F[(q_2-D)\mathbb{1}\{D \le q_2\}] \\ &= b(q_2-q_1) + (b+h)\mathbb{E}_F[(q_1-D)\mathbb{1}\{D \le q_1\} - (q_2-D)\mathbb{1}\{D \le q_2\}]. \end{aligned}$$

□

A.2. On the relationship between λ and M

We implicitly rely on the following fact in all of our results.

PROPOSITION EC.2. *Suppose $G^-(\lambda) < \rho$. Then, $\lambda \le q_G^\star \le M$.*

Proof. We argue that $G^-(\lambda) < \rho$ implies $\lambda \le q_G^\star$ via the converse. Namely, suppose $\lambda > q_G^\star$. Recall, by definition of $q_G^\star$, $G(q_G^\star) \ge \rho$, which implies there exists $q < \lambda$ such that $G(q) \ge \rho$, i.e., $G^-(\lambda) \ge \rho$. Note that $q_G^\star \le M$ by assumption, and we obtain the claim. □

Appendix B: Omitted Proofs

B.1. Section 3 Omitted Proofs

B.1.1. Proof of Proposition 1

Proof. Suppose first that $\Delta > 0$. For any policy π, $\sup_{F\in\mathcal{F}(\lambda;G)}\{C_F(q^\pi) - C_F(q_F^\star)\} \geq \Delta > 0$, by definition of Δ. Therefore, for all $\epsilon \in (0,\Delta)$, $\sup_{F\in\mathcal{F}(\lambda;G)}\left\{C_F(q^\pi) - C_F(q_F^\star)\right\} > \epsilon$, which implies that the problem is unidentifiable.

The proof of the converse is a corollary of our main algorithmic result, in which we design a policy π such that $\sup_{F\in\mathcal{F}(\lambda;G)}\left\{C_F(q^\pi) - C_F(q_F^\star)\right\} = \widetilde{O}(1/\sqrt{N})$ with probability at least $1 - O(1/\sqrt{N})$ when $\Delta = 0$ (see Theorem 2). Taking $N \to \infty$ then implies that the problem is identifiable. □

B.1.2. Proof of Lemma 1

Proof. By definition, $q_G^\star = \inf\{q : G(q) \geq \rho\}$. If $G^-(\lambda) \geq \rho$, there must exist $q < \lambda$ such that $G(q) \geq \rho$, which implies $q_G^\star < \lambda$.

Consider now $q_F^\star$. Again, by definition,

$$\begin{aligned} q_F^\star &= \inf\{q : F(q) \geq \rho\} \\ &= \inf\{q : G(q)\mathbb{1}\{q < \lambda\} + (\mathbb{P}_G(D < \lambda) + \mathbb{P}_F(\lambda \leq D < q))\mathbb{1}\{q \geq \lambda\} \geq \rho\}, \end{aligned}$$

where the second equality uses the fact that $F \in \mathcal{F}(\lambda;G)$, hence $F(x) = G(x)\ \forall\, x < \lambda$. Again, using the fact that $G^-(\lambda) \geq \rho$, there must be $q < \lambda$ such that $G(q) \geq \rho$. Hence, the above infimum is necessarily achieved at $q_G^\star$. □

B.1.3. Proof of Lemma 2

Proof. By Proposition EC.1, we have:

$$\begin{aligned} &\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) \\ &= \sup_{\tilde{q}\in[\lambda,M]} \sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} b(\tilde{q}-q) + (b+h)\mathbb{E}_F[(q-D)\mathbb{1}\{D\leq q\} - (\tilde{q}-D)\mathbb{1}\{D\leq\tilde{q}\}] \\ &= \sup_{\tilde{q}\in[\lambda,M]} b(\tilde{q}-q) + (b+h)\left(\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D\leq q\} - (\tilde{q}-D)\mathbb{1}\{D\leq\tilde{q}\}]\right). \end{aligned} \tag{EC.3}$$

Here, the supremum is over all distributions $F \in \mathcal{F}(\lambda;G)$ such that $\tilde{q} = q_F^\star \geq \lambda$, since $q_F^\star < \lambda \implies F^-(\lambda) \geq \rho$. However, since $F \in \mathcal{F}(\lambda;G)$, $F^-(\lambda) = G^-(\lambda) < \rho$, a contradiction.

We analyze (EC.3) depending on the following three cases: (i) $q > \tilde{q}$, (ii) $q < \tilde{q}$, and (iii) $q = \tilde{q}$.

(i) $q > \tilde{q}$: In this case, the worst-case distribution F that has $\tilde{q}$ as its newsvendor solution maximizes:

$$\mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\} - (\tilde{q}-D)\mathbb{1}\{D \le \tilde{q}\}] = \mathbb{E}_F[(q-\tilde{q})\mathbb{1}\{D \le \tilde{q}\} + (q-D)\mathbb{1}\{\tilde{q} < D \le q\}]. \quad \text{(EC.4)}$$

Observe that $q - D < q - \tilde{q}$ for all $D \in (\tilde{q}, q]$. Hence, for any distribution F such that $\mathbb{P}_F(D \in (\tilde{q}, q]) > 0$, there exists a strictly worse distribution that sets $\mathbb{P}_F(D \in (\tilde{q}, q]) = 0$, moving all of the remaining mass to $\tilde{q}$. (Note that such a distribution will still have $q_F^\star = \tilde{q}$.) For any such F, then, (EC.4) $= (q-\tilde{q})\mathbb{P}_F(D \le \tilde{q})$. Using this in (EC.3):

$$\begin{aligned}
&\sup_{\tilde{q}\in[\lambda,q)} b(\tilde{q}-q) + (b+h)\left(\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\} - (\tilde{q}-D)\mathbb{1}\{D \le \tilde{q}\}]\right) \\
&= \sup_{\tilde{q}\in[\lambda,q)} b(\tilde{q}-q) + (b+h)(q-\tilde{q})\left(\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{P}_F(D \le \tilde{q})\right) \\
&= \sup_{\tilde{q}\in[\lambda,q)} b(\tilde{q}-q) + (b+h)(q-\tilde{q}), \qquad \text{(EC.5)}
\end{aligned}$$

achieved by letting F be such that $\mathbb{P}_F(D \le \tilde{q}) = 1$, since $q > \tilde{q}$ by assumption. Further simplifying, we obtain:

$$\text{(EC.5)} = \sup_{\tilde{q}\in[\lambda,q)} h(q-\tilde{q}) = h(q-\lambda). \quad \text{(EC.6)}$$

(ii) $q < \tilde{q}$: Suppose first that $q < \lambda$. In this case, for a fixed $\tilde{q}$:

$$\begin{aligned}
&\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\} - (\tilde{q}-D)\mathbb{1}\{D \le \tilde{q}\}] \\
&= \sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G[(\tilde{q}-D)\mathbb{1}\{D < \lambda\}] - \mathbb{E}_F[(\tilde{q}-D)\mathbb{1}\{\lambda \le D < \tilde{q}\}]. \quad \text{(EC.7)}
\end{aligned}$$

Since $\tilde{q} - D > 0$ for all $D \in [\lambda, \tilde{q})$, a similar argument as above establishes that any worst-case distribution $F \in \mathcal{F}(\lambda; G)$ necessarily sets $\mathbb{P}_F(D \in [\lambda, \tilde{q})) = 0$. Using this fact above, we obtain:

$$\begin{aligned}
&\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\} - (\tilde{q}-D)\mathbb{1}\{D \le \tilde{q}\}] \\
&= \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G[(\tilde{q}-D)\mathbb{1}\{D < \lambda\}].
\end{aligned}$$

Further using this in (EC.3), we then have:

$$\sup_{\tilde{q}\in[\lambda,M]} b(\tilde{q}-q) + (b+h)\left(\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D \le q\} - (\tilde{q}-D)\mathbb{1}\{D \le \tilde{q}\}]\right)$$

$$
\begin{aligned}
&= \sup_{\tilde{q}\in[\lambda,M]} b(\tilde{q}-q) + (b+h)(\mathbb{E}_G[(q-D)\mathbb{1}\{D\le q\}] - \mathbb{E}_G[(\tilde{q}-D)\mathbb{1}\{D<\lambda\}]) \\
&= -bq + (b+h)(\mathbb{E}_G[(q-D)\mathbb{1}\{D\le q\}] + \mathbb{E}_G[D\mathbb{1}\{D<\lambda\}]) \\
&\quad + \sup_{\tilde{q}\in[\lambda,M]} \left\{\tilde{q}(b-(b+h)G^-(\lambda))\right\}, \qquad \text{(EC.8)}
\end{aligned}
$$

where the supremum is attained at $\tilde{q}=M$, since $G^-(\lambda)<\rho \implies b-(b+h)G^-(\lambda)>0$. Therefore:

$$
\text{(EC.8)} = b(M-q) + (b+h)(\mathbb{E}_G[(q-D)\mathbb{1}\{D\le q\}] - \mathbb{E}_G[(M-D)\mathbb{1}\{D<\lambda\}]). \qquad \text{(EC.9)}
$$

Now, if $q\ge\lambda$, for a fixed $\tilde{q}$ we have:

$$
\begin{aligned}
&\mathbb{E}_F[(q-D)\mathbb{1}\{D\le q\} - (\tilde{q}-D)\mathbb{1}\{D\le\tilde{q}\}] \\
&= \mathbb{E}_F[(q-\tilde{q})\mathbb{1}\{D\le q\} - (\tilde{q}-D)\mathbb{1}\{q<D\le\tilde{q}\}] \\
&= \mathbb{E}_F[-(\tilde{q}-q)\mathbb{1}\{D\le q\} - (\tilde{q}-D)\mathbb{1}\{q<D\le\tilde{q}\}] \\
&= \mathbb{E}_F[-(\tilde{q}-q)\mathbb{1}\{D<\lambda\} - (\tilde{q}-q)\mathbb{1}\{\lambda\le D\le q\} - (\tilde{q}-D)\mathbb{1}\{q<D\le\tilde{q}\}] \\
&= -(\tilde{q}-q)G^-(\lambda) - (\tilde{q}-q)\mathbb{P}_F(\lambda\le D\le q) - \mathbb{E}_F\left[(\tilde{q}-D)\mathbb{1}\{q<D\le\tilde{q}\}\right]. \qquad \text{(EC.10)}
\end{aligned}
$$

Since $0<\tilde{q}-D<\tilde{q}-q$ for all $D\in(q,\tilde{q})$, the supremum of the above is achieved for F such that $\mathbb{P}_F(\lambda\le D<\tilde{q})=0$. Then,

$$
\begin{aligned}
&\sup_{\tilde{q}\in[q,M]} b(\tilde{q}-q) + (b+h)\left(\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D\le q\} - (\tilde{q}-D)\mathbb{1}\{D\le\tilde{q}\}]\right) \\
&= \sup_{\tilde{q}\in[q,M]} b(\tilde{q}-q) - (b+h)(\tilde{q}-q)G^-(\lambda) \\
&= \sup_{\tilde{q}\in[q,M]} (\tilde{q}-q)(b-(b+h)G^-(\lambda)) \\
&= (M-q)(b-(b+h)G^-(\lambda)), \qquad \text{(EC.11)}
\end{aligned}
$$

again, since $G^-(\lambda)<\rho$ by assumption.

(iii) $q=\tilde{q}$: In this case, $C_F(q)-C_F(\tilde{q})=0$ for all $F\in\mathcal{F}(\lambda;G)$.

Putting these three cases together and applying them to (EC.3), we have:

$$
\begin{aligned}
&\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) \\
&= \begin{cases} b(M-q) + (b+h)(\mathbb{E}_G[(q-D)\mathbb{1}\{D\le q\}] - \mathbb{E}_G[(M-D)\mathbb{1}\{D<\lambda\}]) & \text{if } q<\lambda \\ \max\{h(q-\lambda), (b-(b+h)G^-(\lambda))(M-q)\} & \text{if } q\ge\lambda. \end{cases}
\end{aligned}
$$

Noting that $h(q-\lambda)>(b-(b+h)G^-(\lambda))(M-q) \iff q>q_G^\dagger$, with equality at $q_G^\dagger$, we obtain the result.

□

B.1.4. Proof of Theorem 1

Proof. We proceed separately for each regime.

Case I: $G^-(\lambda) \geq \rho$. By Lemma 1, $q_F^\star = q_G^\star < \lambda$ for all $F \in \mathcal{F}(\lambda; G)$. Therefore, for all $F \in \mathcal{F}(\lambda; G)$, $q \in [0, M]$:

$$C_F(q) - C_F(q_F^\star) = C_F(q) - C_F(q_G^\star).$$

Letting $q = q_G^\star$, this implies that $C_F(q_G^\star) - C_F(q_F^\star) = 0$ for all $F \in \mathcal{F}(\lambda; G)$. Using the fact that $\sup_{F \in \mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) \geq 0$ for all q by definition of $q_F^\star$, we then have $\Delta := \inf_q \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) = 0$, precisely achieved at $q^\Delta = q_G^\star$.

Case II: $G^-(\lambda) < \rho$. Recall, $q_G^\dagger := \frac{bM + h\lambda - (b+h)G^-(\lambda)M}{(b+h)(1-G^-(\lambda))}$ by definition. We analyze $\inf_{q \in [0,M]} \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star)$ by partitioning the proof into three cases: (i) $q < \lambda$, (ii) $q \in [\lambda, q_G^\dagger]$, and (iii) $q > q_G^\dagger$.

By Lemma 2, for $q < \lambda$:

$$\sup_{F \in \mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) = b(M-q) + (b+h)\left[\mathbb{E}_G\left[(q-D)\mathbb{1}\{D \leq q\}\right] - (M-D)\mathbb{1}\{D < \lambda\}\right]\Big].$$

Hence, we have:

$$\begin{aligned}
&\inf_{q<\lambda} \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) \\
&= \inf_{q<\lambda} bM - (b+h)\mathbb{E}_G[(M-D)\mathbb{1}\{D<\lambda\}] - bq + (b+h)\mathbb{E}_G[(q-D)\mathbb{1}\{D \leq q\}], \qquad \text{(EC.12)}
\end{aligned}$$

where (EC.12) uses the fact that $F \in \mathcal{F}(\lambda; G)$ for the first expectation. Now, the newsvendor cost can equivalently be written as

$$C_G(q) = b(\mathbb{E}_G[D] - q) + (b+h)\mathbb{E}_G[(q-D)\mathbb{1}\{D \leq q\}]. \qquad \text{(EC.13)}$$

Plugging this formulation into (EC.12), we obtain:

$$\inf_{q<\lambda} \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) = bM - (b+h)\mathbb{E}_G[(M-D)\mathbb{1}\{D<\lambda\}] - b\mathbb{E}_G[D] + \inf_{q<\lambda} C_G(q). \qquad \text{(EC.14)}$$

Recall, $G^-(\lambda) < \rho$ implies $q_G^\star \geq \lambda$. Hence, by convexity of $C_G(q)$, $C_G(\lambda) \leq C_G(q)$ for all $q < \lambda$. Then:

$$\begin{aligned}
&\inf_{q<\lambda} \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) \\
&= bM - (b+h)\mathbb{E}_G[(M-D)\mathbb{1}\{D<\lambda\}] - b\mathbb{E}_G[D] + C_G(\lambda) \\
&= bM - (b+h)\mathbb{E}_G[(M-D)\mathbb{1}\{D<\lambda\}] - b\lambda + (b+h)\mathbb{E}_G[(\lambda-D)\mathbb{1}\{D \leq \lambda\}],
\end{aligned}$$

where the second equality follows from (EC.13). Simplifying, we finally obtain:

$$\inf_{q<\lambda} \sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) = b(M-\lambda) - (b+h)\mathbb{E}_G\Big[(M-\lambda)\mathbb{1}\{D<\lambda\}\Big]$$
$$= (M-\lambda)(b-(b+h)G^-(\lambda)). \tag{EC.15}$$

Suppose now $q \in \left[\lambda, q_G^\dagger\right]$. Again, by Lemma 2:

$$\min_{q\in[\lambda,q_G^\dagger]} \sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) = \min_{q\in[\lambda,q_G^\dagger]} (M-q)(b-(b+h)G^-(\lambda))$$
$$= (M-q_G^\dagger)(b-(b+h)G^-(\lambda)), \tag{EC.16}$$

where the final equality follows from the fact that the right-hand side of (EC.16) is decreasing in q, since $G^-(\lambda) < \rho$. It is easy to derive that $q_G^\dagger \geq \lambda$ for $G^-(\lambda) < \rho$. Comparing (EC.15) and (EC.16), this implies:

$$\min_{q\in[\lambda,q_G^\dagger]} \sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) \leq \inf_{q<\lambda} \sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star).$$

Hence, setting $q = q_G^\dagger$ (weakly) dominates setting $q < \lambda$.

We conclude by considering $q \in (q_G^\dagger, M]$. By Lemma 2:

$$\inf_{q\in(q_G^\dagger,M]} \sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star) = \inf_{q\in(q_G^\dagger,M]} h(q-\lambda) = h(q_G^\dagger - \lambda).$$

By definition of $q_G^\dagger$,

$$h(q_G^\dagger - \lambda) = (M-q_G^\dagger)(b-(b+h)G^-(\lambda)) = \frac{h(b-(b+h)G^-(\lambda)(M-\lambda)}{(b+h)(1-G^-(\lambda))}.$$

Putting this together with (EC.16), we conclude that $q^\Delta = q_G^\dagger$, with

$$\Delta = \frac{h(b-(b+h)G^-(\lambda)(M-\lambda)}{(b+h)(1-G^-(\lambda))}.$$

□

B.1.5. Proof of Proposition 3

Proof. Recall that, in this regime, by Lemma 1, $q_F^\star = q_G^\star < \lambda$ for all $F \in \mathcal{F}(\lambda;G)$. Applying this fact to Proposition EC.1, for any $F \in \mathcal{F}(\lambda;G)$ we have:

$$C_F(q) - C_F(q_F^\star) = b(q_G^\star - q) + (b+h)\mathbb{E}_F\left[(q-D)\mathbb{1}\{D\leq q\} - (q_G^\star - D)\mathbb{1}\left\{D\leq q_G^\star\right\}\right]. \tag{EC.17}$$

If $q < \lambda$, since $F \in \mathcal{F}(\lambda;G)$ and $q_F^\star < \lambda$, the above evaluates to:

$$C_F(q) - C_F(q_F^\star) = b(q_G^\star - q) + (b+h)\mathbb{E}_G\left[(q-D)\mathbb{1}\{D\leq q\} - (q_G^\star - D)\mathbb{1}\left\{D\leq q_G^\star\right\}\right]$$
$$= C_G(q) - C_G(q_G^\star),$$

which proves the first case.

Suppose now that $q \geq \lambda$. Note that only the second term of (EC.17) has dependence on F. Again, since $q_G^\star < \lambda$, for all $F \in \mathcal{F}(\lambda; G)$ we have:

$$\begin{aligned}
&\mathbb{E}_F\left[(q-D)\mathbb{1}\{D \leq q\} - (q_G^\star - D)\mathbb{1}\left\{D \leq q_G^\star\right\}\right] \\
&= \mathbb{E}_G[(q-D)\mathbb{1}\{D < \lambda\}] + \mathbb{E}_F[(q-D)\mathbb{1}\{\lambda \leq D < q\}] - \mathbb{E}_G\left[(q_G^\star - D)\mathbb{1}\left\{D \leq q_G^\star\right\}\right] \\
&\leq \mathbb{E}_G[(q-D)\mathbb{1}\{D < \lambda\}] + (q-\lambda)\mathbb{P}_F(D = \lambda) - \mathbb{E}_G\left[(q_G^\star - D)\mathbb{1}\left\{D \leq q_G^\star\right\}\right], \qquad \text{(EC.18)}
\end{aligned}$$

where the inequality is attained by the distribution F that places all remaining mass on λ. Re-arranging (EC.18), we obtain:

$$\begin{aligned}
&\sup_{F \in \mathcal{F}(\lambda;G)} \mathbb{E}_F\left[(q-D)\mathbb{1}\{D \leq q\} - (q_G^\star - D)\mathbb{1}\left\{D \leq q_G^\star\right\}\right] \\
&= \mathbb{E}_G[(q-\lambda)\mathbb{1}\{D < \lambda\}] + \mathbb{E}_G[(\lambda - D)\mathbb{1}\{D < \lambda\}] \\
&\qquad + (q-\lambda)(1 - G^-(\lambda)) - \mathbb{E}_G\left[(q_G^\star - D)\mathbb{1}\left\{D \leq q_G^\star\right\}\right] \\
&= (q-\lambda) + \mathbb{E}_G[(\lambda - D)\mathbb{1}\{D < \lambda\}] - \mathbb{E}_G\left[(q_G^\star - D)\mathbb{1}\left\{D \leq q_G^\star\right\}\right].
\end{aligned}$$

Plugging this back into (EC.17), the second case is shown. □

B.2. Worst-Case Nature of Two-Point Mass Distributions

B.2.1. Proof of Proposition 2

Proof. This fact is a corollary of the derivation of Regret(q) in the proof of Lemma 2. To show this, recall our three cases: (i) $q > \tilde{q} \geq \lambda$, (ii) $q < \lambda \leq \tilde{q}$, and (iii) $\lambda \leq q < \tilde{q}$, where $\tilde{q} := q_F^\star$.

(i) $q > \tilde{q} \geq \lambda$: In this case, we showed that any worst-case distribution F is such that $\tilde{q} = \lambda$, and $\mathbb{P}_F(D \leq \lambda) = 1$. Therefore, the worst-case distribution places the entirety of the remaining mass at λ.

(ii) $q < \lambda \leq \tilde{q}$: In this case, we argued that any worst-case distribution F is such that $\mathbb{P}_F(\lambda \leq D < \tilde{q}) = 0$ and $\tilde{q} = M$. This is satisfied by a distribution F that places the entirety of the remaining mass at M.

(iii) $\lambda \leq q < \tilde{q}$: Here we similarly argued that any worst-case distribution F is such that $\mathbb{P}_F(\lambda \leq D < \tilde{q}) = 0$ and $\tilde{q} = M$. Again, this is achieved by a distribution F that places the entirety of the remaining mass on M.

□

B.3. Proof of Theorem 2

Proof. Proposition EC.3 below upper bounds our algorithm's regret by (i) the minimax risk Δ, and (ii) the cost difference between q^{alg} and q^Δ, for any distribution $F \in \mathcal{F}(\lambda; G)$. We defer its proof to Appendix B.3.1.

PROPOSITION EC.3. *The regret of Algorithm 1 is upper bounded by:*

$$\text{Regret}(q^{\text{alg}}) \leq \Delta + \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^{\text{alg}}) - C_F(q^\Delta). \qquad \text{(EC.19)}$$

Hence, the remainder of the proof focuses on bounding the second term. Before doing so, we show a series of useful facts. First off, Lemma EC.1 establishes Lipschitzness of the newsvendor cost function. Its proof can be found in Appendix B.3.2.

LEMMA EC.1. *For any distribution F, and ordering quantities q, q':*

$$|C_F(q) - C_F(q')| \le \max\{b, h\}|q - q'|.$$

Moreover, by Fact 1, $\mathbb{1}\left\{s_i^{\text{off},\lambda} < \lambda\right\} = \mathbb{1}\left\{d_i^{\text{off}} < \lambda\right\}$ for all $i \in [N]$. We use this to prove Lemma EC.2, which establishes that $\widehat{G}^-(\lambda)$ and $G^-(\lambda)$ are close with constant probability, below. We defer its proof to Appendix B.3.3.

LEMMA EC.2. *Let $\mathcal{E} = \left\{|\widehat{G}^-(\lambda) - G^-(\lambda)| < \zeta\right\}$. Then, $\mathbb{P}_G(\mathcal{E}) \ge 1 - \delta$.*

We moreover use Fact 1 to show that $q^\star_{\widehat{G}}$ is precisely equal to the actual SAA of $q^\star_G$. To formalize this, we introduce some additional notation. Let $\widehat{G}^{\text{uc}}(x) = \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{d_i^{\text{off}} \le x\right\}$ be the empirical cdf of the (true) uncensored demand, and $q^\star_{\widehat{G}^{\text{uc}}} = \inf\{x : \widehat{G}^{\text{uc}}(x) \ge \rho\}$ be the corresponding newsvendor quantile. We repeatedly rely on the following lemma in the remainder of the proof. Its proof can be found in Appendix B.3.4.

LEMMA EC.3. *Suppose $\widehat{G}^-(\lambda) \ge \rho$. Then, $q^\star_{\widehat{G}} = q^\star_{\widehat{G}^{\text{uc}}} < \lambda$.*

Having established these useful facts, we may begin with the proof of our upper bound. We prove each of the four regret bounds separately, beginning with the two "extreme" cases of $G^-(\lambda) \ge \rho + 2\zeta$ and $G^-(\lambda) < \rho - 2\zeta$. Moreover, in the remainder of the proof we condition on the good event $\mathcal{E}$ defined in Lemma EC.2.

Case I: $G^-(\lambda) \ge \rho + 2\zeta$. Under event $\mathcal{E}$, $\widehat{G}^-(\lambda) > \rho + \zeta$, which implies $q^{\text{alg}} = q^\star_{\widehat{G}}$ by construction. Since $\widehat{G}^-(\lambda) \ge \rho$, by Lemma EC.3, $q^\star_{\widehat{G}} = q^\star_{\widehat{G}^{\text{uc}}} < \lambda$. By Proposition 3, then,

$$\text{Regret}(q^{\text{alg}}) \le C_G(q^\star_{\widehat{G}^{\text{uc}}}) - C_G(q^\star_G).$$

Hence, bounding the regret of q^{alg} in this region reduces to the question of bounding the regret of the SAA of $q^\star_G$.

The following lemma, whose proof is adapted from Chen and Ma (2026) (see Appendix B.3.5), allows us to bound this SAA error.

LEMMA EC.4. *Suppose $G^-(\lambda) \ge \rho$, and $\widehat{G}^-(\lambda) \ge \rho$. With probability at least $1 - \delta$,*

$$C_G(q^\star_{\widehat{G}^{\text{uc}}}) - C_G(q^\star_G) \le \lambda(b + h)\sqrt{\frac{\log(2/\delta)}{2N}}.$$

Taking a union bound over the event $\mathcal{E}$ and this latter good event, we obtain the first case in the theorem statement.

Case II: $G^-(\lambda) < \rho - 2\zeta$. By Theorem 1, $q^\Delta = q^\dagger_G$ in this case. Moreover, under event $\mathcal{E}$, $\widehat{G}^-(\lambda) < \rho - \zeta$, and our algorithm outputs $q^{\text{alg}} = q^\dagger_{\widehat{G}}$. Then, by Proposition EC.3:

$$\text{Regret}(q^{\text{alg}}) \le \Delta + \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^\dagger_{\widehat{G}}) - C_F(q^\dagger_G) \le \Delta + \max\{b,h\}|q^\dagger_{\widehat{G}} - q^\dagger_G|, \tag{EC.20}$$

where the second inequality follows from the fact that C_F is $\max\{b,h\}$-Lipschitz (Lemma EC.1). It suffices then to bound the error in our algorithm's estimate of $q^\dagger_G$. Lemma EC.5, whose proof can be found in Appendix B.3.6, leverages the fact that $|G^-(\lambda) - \widehat{G}^-(\lambda)| < \zeta$ under event $\mathcal{E}$ to bound this error.

LEMMA EC.5. *Under event* $\mathcal{E}$, $|q^\dagger_G - q^\dagger_{\widehat{G}}| \le \frac{M-\lambda}{1-\rho}\zeta$.

Plugging this back into (EC.20), we obtain:

$$\text{Regret}(q^{\text{alg}}) \le \Delta + \frac{\max\{b,h\}(M-\lambda)}{1-\rho}\zeta, \tag{EC.21}$$

completing the proof of this case in the theorem statement.

We now turn our attention to the remaining case, where $G^-(\lambda) \in [\rho - 2\zeta, \rho + 2\zeta)$, further partitioning the analysis based on whether $G^-(\lambda) \ge \rho$.

Case III: $G^-(\lambda) \in [\rho, \rho + 2\zeta)$. By Theorem 1, $q^\Delta = q^\star_G$, and $\Delta = 0$. Moreover, given $\mathcal{E}$, $\widehat{G}^-(\lambda) \in (\rho - \zeta, \rho + 3\zeta)$.

Suppose first that $\widehat{G}^-(\lambda) \in [\rho + \zeta, \rho + 3\zeta)$. Then, $q^{\text{alg}} = q^\star_{\widehat{G}} < \lambda$. Identical arguments as those used for Case I yield:

$$\text{Regret}(q^{\text{alg}}) \le C_G(q^\star_{\widehat{G}^{\text{uc}}}) - C_G(q^\star_G) \le \lambda(b+h)\zeta. \tag{EC.22}$$

Suppose now that $\widehat{G}^-(\lambda) \in (\rho - \zeta, \rho + \zeta)$. In this case, $q^{\text{alg}} = \lambda$. Using the fact that $q^\Delta = q^\star_G < \lambda$ and $\Delta = 0$, by Proposition EC.3, we have:

$$\text{Regret}(q^{\text{alg}}) \le \Delta + \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^{\text{alg}}) - C_F(q^\Delta) = \sup_{F \in \mathcal{F}(\lambda;G)} C_F(\lambda) - C_F(q^\star_G).$$

Proposition 3 and Lemma 1 together imply that:

$$\begin{aligned}
&\sup_{F \in \mathcal{F}(\lambda;G)} C_F(\lambda) - C_F(q^\star_G) \\
&= b(q^\star_G - \lambda) + (b+h)\mathbb{E}_G\left[(\lambda - D)\mathbb{1}\{D < \lambda\} - (q^\star_G - D)\mathbb{1}\left\{D \le q^\star_G\right\}\right] \\
&= b(q^\star_G - \lambda) + (b+h)\mathbb{E}_G\left[(\lambda - q^\star_G)\mathbb{1}\{D < \lambda\} + (q^\star_G - D)\mathbb{1}\{D < \lambda\} - (q^\star_G - D)\mathbb{1}\left\{D \le q^\star_G\right\}\right] \\
&= b(q^\star_G - \lambda) + (b+h)\mathbb{E}_G\left[(\lambda - q^\star_G)\mathbb{1}\{D < \lambda\} + (q^\star_G - D)\mathbb{1}\left\{q^\star_G < D < \lambda\right\}\right] \\
&\le (\lambda - q^\star_G)\Big((b+h)G^-(\lambda) - b\Big) \\
&= (b+h)(\lambda - q^\star_G)(G^-(\lambda) - \rho).
\end{aligned}$$

Using the fact that $q^\star_G \in [0, \lambda)$ for $G^-(\lambda) \geq \rho$, and $G^-(\lambda) \leq \rho + 2\zeta$ by assumption, we obtain the final bound of:

$$\text{Regret}(q^{\text{alg}}) \leq 2\lambda(b+h)\zeta.$$

We proceed to our final case.

Case IV: $G^-(\lambda) \in [\rho - 2\zeta, \rho)$. By Theorem 1, $q^\Delta = q^\dagger_G$. Moreover, under $\mathcal{E}$, $\widehat{G}^-(\lambda) \in (\rho - 3\zeta, \rho + \zeta)$.

Suppose first that $\widehat{G}^-(\lambda) \in (\rho - 3\zeta, \rho - \zeta)$. In this case, the algorithm outputs $q^{\text{alg}} = q^\dagger_{\widehat{G}}$. By the same arguments as those used in Case II, we obtain:

$$\text{Regret}(q^{\text{alg}}) \leq \Delta + \frac{\max\{b,h\}(M-\lambda)}{(1-\rho)}\zeta.$$

Suppose now that $\widehat{G}^-(\lambda) \in [\rho - \zeta, \rho + \zeta)$. In this case, $q^{\text{alg}} = \lambda$. Again, leveraging Proposition EC.3, we have:

$$\text{Regret}(q^{\text{alg}}) \leq \Delta + \sup_{F \in \mathcal{F}(\lambda;G)} C_F(\lambda) - C_F(q^\dagger_G) \leq \Delta + \max\{b,h\}|\lambda - q^\dagger_G|, \tag{EC.23}$$

where the final inequality follows from Lipschitzness of the newsvendor cost function (Lemma EC.1). Plugging in the definition of $q^\dagger_G$, algebra gives us that:

$$\begin{aligned}
|\lambda - q^\dagger_G| &= \frac{(M-\lambda)(b-(b+h)G^-(\lambda))}{(b+h)(1-G^-(\lambda))} \\
&\leq \frac{(M-\lambda)(b-(b+h)(\rho-2\zeta))}{(b+h)(1-\rho)} \\
&= \frac{2(M-\lambda)(b+h)\zeta}{h} \\
&= \frac{2(M-\lambda)}{1-\rho}\zeta,
\end{aligned}$$

where the first inequality uses the fact that $G^-(\lambda) \in [\rho - 2\zeta, \rho)$, and the subsequent equalities plug in the definition of $\rho = b/(b+h)$. Plugging this bound back into (EC.23), we obtain:

$$\text{Regret}(q^{\text{alg}}) \leq \Delta + \max\{b,h\}\frac{2(M-\lambda)}{1-\rho}\zeta,$$

concluding the proof of the theorem. □

B.3.1. Proof of Proposition EC.3

Proof. Adding and subtracting $C_F(q^\Delta)$ from the definition of regret, we have:

$$\begin{aligned}
\text{Regret}(q^{\text{alg}}) &= \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^{\text{alg}}) - C_F(q^\Delta) + C_F(q^\Delta) - C_F(q^\star_F) \\
&\leq \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^{\text{alg}}) - C_F(q^\Delta) + \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^\Delta) - C_F(q^\star_F) \\
&= \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^{\text{alg}}) - C_F(q^\Delta) + \Delta,
\end{aligned}$$

where the final equality uses the definition of $\Delta = \sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^\Delta) - C_F(q^\star_F)$. □

B.3.2. Proof of Lemma EC.1

Proof. By definition:

$$|C_F(q) - C_F(q')| = \left| \mathbb{E}_F\left[b\big((D-q)^+ - (D-q')^+\big) + h\big((q-D)^+ - (q'-D)^+\big)\right] \right|.$$

Without loss of generality, suppose $q < q'$. Consider first the term $b((D-q)^+ - (D-q')^+)$. We have:

$$(D-q)^+ - (D-q')^+ = \begin{cases} q'-q & \text{if } D \geq q' \\ 0 & \text{if } D \leq q \\ D-q & \text{if } D \in (q, q'). \end{cases}$$

Similarly:

$$(q-D)^+ - (q'-D)^+ = \begin{cases} 0 & \text{if } D \geq q' \\ q-q' & \text{if } D \leq q \\ D-q' & \text{if } D \in (q, q'). \end{cases}$$

Putting these two together, we have:

$$b\big((D-q)^+ - (D-q')^+\big) + h\big((q-D)^+ - (q'-D)^+\big) = \begin{cases} b(q'-q) & \text{if } D \geq q' \\ h(q-q') & \text{if } D \leq q \\ b(D-q) + h(D-q') & \text{if } D \in (q, q') \end{cases}$$

$$\implies \left| \mathbb{E}_F\left[b\big((D-q)^+ - (D-q')^+\big) + h\big((q-D)^+ - (q'-D)^+\big)\right] \right| \leq \max\{b,h\}|q-q'|.$$

□

B.3.3. Proof of Lemma EC.2

Proof. Since $\mathbb{1}\left\{s_i^{\text{off},\lambda} < \lambda\right\} = \mathbb{1}\left\{d_i^{\text{off}} < \lambda\right\}$ for all i, we have:

$$\widehat{G}^-(\lambda) = \frac{1}{N} \sum_{i \in [N]} \mathbb{1}\left\{d_i^{\text{off}} < \lambda\right\} \implies \mathbb{E}_G[\widehat{G}^-(\lambda)] = G^-(\lambda).$$

Hence, by Hoeffding's inequality (Boucheron et al. 2013):

$$\mathbb{P}_G(|\widehat{G}^-(\lambda) - G^-(\lambda)| \geq \zeta) \leq 2\exp(-2N\zeta^2).$$

Letting $\zeta = \sqrt{\frac{\log(2/\delta)}{2N}}$, we obtain the result. □

B.3.4. Proof of Lemma EC.3

Proof. We first show that $\widehat{G}^-(\lambda) \geq \rho$ implies $q^\star_{\widehat{G}^{\text{uc}}} < \lambda$. Again, since $\mathbb{1}\left\{s_i^{\text{off},\lambda} < \lambda\right\} = \mathbb{1}\left\{d_i^{\text{off}} < \lambda\right\}$,

$$\frac{1}{N}\sum_{i=1}^{N} \mathbb{1}\left\{d_i^{\text{off}} < \lambda\right\} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}\left\{s_i^{\text{off},\lambda} < \lambda\right\} = \widehat{G}^-(\lambda) \geq \rho,$$

by assumption. Then, it must be that $\widehat{G}^{\text{uc}}(x) \geq \rho$ for some $x < \lambda$, which implies that $q^\star_{\widehat{G}^{\text{uc}}} < \lambda$.

We use this to show that $q^\star_{\widehat{G}} = q^\star_{\widehat{G}^{\text{uc}}}$. By definition, $q^\star_{\widehat{G}^{\text{uc}}}$ satisfies:

$$\begin{cases} \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{d_i^{\text{off}} \le x\right\} < \rho & \forall\, x < q^\star_{\widehat{G}^{\text{uc}}} \\ \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{d_i^{\text{off}} \le x\right\} \ge \rho & \forall\, x \ge q^\star_{\widehat{G}^{\text{uc}}}. \end{cases}$$

Fix $x \ge q^\star_{\widehat{G}^{\text{uc}}}$. Since $s_i^{\text{off},\lambda} \le d_i^{\text{off}}$ for all i,

$$\frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} \ge \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{d_i^{\text{off}} \le x\right\} \ge \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{d_i^{\text{off}} \le q^\star_{\widehat{G}^{\text{uc}}}\right\} \ge \rho,$$

by definition of $q^\star_{\widehat{G}^{\text{uc}}}$.

Now, fix $x < q^\star_{\widehat{G}^{\text{uc}}}$. Since $x < q^\star_{\widehat{G}^{\text{uc}}} < \lambda$, if $d_i^{\text{off}} \ge \lambda$, $s_i^{\text{off},\lambda} = \lambda$, which then implies that $\mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} = 0$. We moreover have in this case that $\mathbb{1}\left\{d_i^{\text{off}} \le x\right\} = 0$. Putting these two together, we have that $\mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} = \mathbb{1}\left\{d_i^{\text{off}} \le x\right\}$ if $d_i^{\text{off}} \ge \lambda$.

If $d_i^{\text{off}} < \lambda$, on the other hand, $s_i^{\text{off},\lambda} = d_i^{\text{off}}$ by definition. Therefore, $\mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} = \mathbb{1}\left\{d_i^{\text{off}} \le x\right\}$ whenever $d_i^{\text{off}} < \lambda$.

Since $\mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} = \mathbb{1}\left\{d_i^{\text{off}} \le x\right\}$ whenever $x < q^\star_{\widehat{G}^{\text{uc}}}$, we obtain:

$$\frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} = \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{d_i^{\text{off}} \le x\right\} < \rho \quad \forall\, x < q^\star_{\widehat{G}^{\text{uc}}}.$$

Putting these two facts together, we have:

$$\begin{cases} \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} < \rho & \forall\, x < q^\star_{\widehat{G}^{\text{uc}}} \\ \frac{1}{N}\sum_{i\in[N]} \mathbb{1}\left\{s_i^{\text{off},\lambda} \le x\right\} \ge \rho & \forall\, x \ge q^\star_{\widehat{G}^{\text{uc}}}, \end{cases}$$

which implies that $q^\star_{\widehat{G}} = q^\star_{\widehat{G}^{\text{uc}}}$. □

B.3.5. Proof of Lemma EC.4

Proof. The proof is modified from Chen and Ma (2026).

By the Dvoretzky–Kiefer–Wolfowitz–Massart inequality (DKW) inequality (Massart 1990):

$$\mathbb{P}\left[\sup_{a\ge 0} |\widehat{G}^{\text{uc}}(a) - G(a)| \le \sqrt{\frac{\log(2/\delta)}{2N}}\right] \ge 1 - 2\exp\left(-2N\left(\sqrt{\frac{\log(2/\delta)}{2N}}\right)^2\right) = 1 - \delta.$$

Therefore with probability at least $1-\delta$:

$$\sup_{a\ge 0} |\widehat{G}^{\text{uc}}(a) - G(a)| \le \sqrt{\frac{\log(2/\delta)}{2N}}. \tag{EC.24}$$

We condition our analysis on the above event, partitioning the proof based on (i) $q^\star_{\widehat{G}^{\text{uc}}} \le q^\star_G$, and (ii) $q^\star_{\widehat{G}^{\text{uc}}} > q^\star_G$.

Case I: $q^\star_{\widehat{G}^{\text{uc}}} \le q^\star_G$. By Proposition EC.1:

$$\begin{aligned}
&C_G(q^\star_{\widehat{G}^{\text{uc}}}) - C_G(q^\star_G) \\
&= b(q^\star_G - q^\star_{\widehat{G}^{\text{uc}}}) + (b+h)\mathbb{E}_G\left[(q^\star_{\widehat{G}^{\text{uc}}} - D)\mathbb{1}\left\{D \le q^\star_{\widehat{G}^{\text{uc}}}\right\} - (q^\star_G - D)\mathbb{1}\left\{D \le q^\star_G\right\}\right] \\
&= b(q^\star_G - q^\star_{\widehat{G}^{\text{uc}}}) + (b+h)\mathbb{E}_G\left[(q^\star_{\widehat{G}^{\text{uc}}} - q^\star_G)\mathbb{1}\left\{D \le q^\star_{\widehat{G}^{\text{uc}}}\right\} - (q^\star_G - D)\mathbb{1}\left\{q^\star_{\widehat{G}^{\text{uc}}} < D \le q^\star_G\right\}\right] \\
&\le (b+h)(q^\star_G - q^\star_{\widehat{G}^{\text{uc}}})(\rho - G(q^\star_{\widehat{G}^{\text{uc}}})), \qquad \text{(EC.25)}
\end{aligned}$$

where the upper bound follows from the fact that $(q^\star_G - D)\mathbb{1}\left\{q^\star_{\widehat{G}^{\text{uc}}} < D \le q^\star_G\right\} \ge 0$, and $b = (b+h)\rho$ by definition. Moreover:

$$\rho - G(q^\star_{\widehat{G}^{\text{uc}}}) = \rho - \widehat{G}^{\text{uc}}(q^\star_{\widehat{G}^{\text{uc}}}) + \widehat{G}^{\text{uc}}(q^\star_{\widehat{G}^{\text{uc}}}) - G(q^\star_{\widehat{G}^{\text{uc}}}) \le \sup_{a\ge 0}|\widehat{G}^{\text{uc}}(a) - G(a)|, \qquad \text{(EC.26)}$$

where the inequality follows from $\widehat{G}^{\text{uc}}(q^\star_{\widehat{G}^{\text{uc}}}) \ge \rho$, by definition of $q^\star_{\widehat{G}^{\text{uc}}}$.

Using this fact in (EC.25), we have:

$$\begin{aligned}
C_G(q^\star_{\widehat{G}^{\text{uc}}}) - C_G(q^\star_G) &\le (b+h)(q^\star_G - q^\star_{\widehat{G}^{\text{uc}}})(\rho - G(q^\star_{\widehat{G}^{\text{uc}}})) \\
&\le (b+h)\cdot|q^\star_G - q^\star_{\widehat{G}^{\text{uc}}}|\cdot\sup_{a\ge 0}|\widehat{G}^{\text{uc}}(a) - G(a)|.
\end{aligned}$$

Case II: $q^\star_{\widehat{G}^{\text{uc}}} > q^\star_G$. Again, by Proposition EC.1:

$$\begin{aligned}
&C_G(q^\star_{\widehat{G}^{\text{uc}}}) - C_G(q^\star_G) \\
&= b(q^\star_G - q^\star_{\widehat{G}^{\text{uc}}}) + (b+h)\mathbb{E}_G\left[(q^\star_{\widehat{G}^{\text{uc}}} - D)\mathbb{1}\left\{D \le q^\star_{\widehat{G}^{\text{uc}}}\right\} - (q^\star_G - D)\mathbb{1}\left\{D \le q^\star_G\right\}\right] \\
&= -\rho(b+h)(q^\star_{\widehat{G}^{\text{uc}}} - q^\star_G) + (b+h)\mathbb{E}_G\left[(q^\star_{\widehat{G}^{\text{uc}}} - q^\star_G)\mathbb{1}\left\{D \le q^\star_G\right\} + (q^\star_{\widehat{G}^{\text{uc}}} - D)\mathbb{1}\left\{q^\star_G < D < q^\star_{\widehat{G}^{\text{uc}}}\right\}\right] \\
&\le (b+h)(q^\star_{\widehat{G}^{\text{uc}}} - q^\star_G)\left(-\rho + G^-\left(q^\star_{\widehat{G}^{\text{uc}}}\right)\right), \qquad \text{(EC.27)}
\end{aligned}$$

where the inequality follows from $q^\star_{\widehat{G}^{\text{uc}}} - D < q^\star_{\widehat{G}^{\text{uc}}} - q^\star_G$ for all $D > q^\star_G$.

For all $q < q^\star_{\widehat{G}^{\text{uc}}}$:

$$G(q) - \rho = G(q) - \widehat{G}^{\text{uc}}(q) + \widehat{G}^{\text{uc}}(q) - \rho \le \sup_{a\ge 0}|\widehat{G}^{\text{uc}}(a) - G(a)|,$$

where the inequality follows from the fact that $\widehat{G}^{\text{uc}}(q) < \rho$ for all $q < q^\star_{\widehat{G}^{\text{uc}}}$, by definition of $q^\star_{\widehat{G}^{\text{uc}}}$. We similarly use this to obtain the final upper bound:

$$C_G(q^\star_{\widehat{G}^{\text{uc}}}) - C_G(q^\star_G) \le (b+h)\cdot|q^\star_{\widehat{G}^{\text{uc}}} - q^\star_G|\cdot\sup_{a\ge 0}|\widehat{G}^{\text{uc}}(a) - G(a)|.$$

Applying (EC.24) to both cases, and using the fact that both $q^\star_G \le \lambda$ and $q^\star_{\widehat{G}^{\text{uc}}} \le \lambda$ under the assumption that $G^-(\lambda) \ge \rho$ and $\widehat{G}^-(\lambda) \ge \rho$, we obtain the claim. □

B.3.6. Proof of Lemma EC.5

Proof. Consider the function $Q^\dagger(x) = \frac{bM+h\lambda-(b+h)Mx}{(b+h)(1-x)}$. Observe that $q^\dagger_G = Q^\dagger(G^-(\lambda))$, and $q^\dagger_{\widehat{G}} = Q^\dagger(\widehat{G}^-(\lambda))$ by definition. For any $x \in [0, \rho-\zeta]$:

$$\frac{dQ^\dagger}{dx} = \frac{h(\lambda - M)}{(b+h)(1-x)^2} \implies \left|\frac{dQ^\dagger}{dx}\right| \le \frac{h(M-\lambda)}{(b+h)(1-\rho+\zeta)^2}.$$

Hence $Q^\dagger$ is $\frac{h(M-\lambda)}{(b+h)(1-\rho+\zeta)^2}$-Lipschitz, which implies

$$|q^\dagger_G - q^\dagger_{\widehat{G}}| \le \frac{h(M-\lambda)}{(b+h)(1-\rho+\zeta)^2}|G^-(\lambda) - \widehat{G}^-(\lambda)| \le \frac{h(M-\lambda)}{(b+h)(1-\rho+\zeta)^2}\zeta = \frac{(1-\rho)(M-\lambda)}{(1-\rho+\zeta)^2}\zeta$$

under event $\mathcal{E}$. Using the fact that since $\rho < 1$ and $\zeta > 0$, $(1-\rho)^2 \le (1-\rho+\zeta)^2$, which completes the proof of the claim. □

B.4. Proof of Corollary 2

Proof. For ease of notation, we let $R^{\text{id}} = 2\lambda(b+h)\zeta$ be the upper bound on the regret of RCN derived in Theorem 2; we also let $R^{\text{id}}_{\max} = \max_{q\in[0,M]} \sup_{F\in\mathcal{F}(\lambda;G)} \text{Regret}(q)$ when $G^-(\lambda) \ge \rho$. Similarly, let $R^{\text{ui}} = \Delta + 2\max\left\{\frac{b}{h}, 1\right\}(M-\lambda)(b+h)\zeta$ and $R^{\text{ui}}_{\max} = \max_{q\in[0,M]} \sup_{F\in\mathcal{F}(\lambda;G)} \text{Regret}(q)$ when $G^-(\lambda) < \rho$. Applying the bounds from Theorem 2, we have:

$$\mathbb{E}_G[\text{Regret}(q^{\text{alg}})] \le \begin{cases} R^{\text{id}} + R^{\text{id}}_{\max}\cdot 2\delta & \text{if } G^-(\lambda) \ge \rho \\ R^{\text{ui}} + R^{\text{ui}}_{\max}\cdot 2\delta & \text{if } G^-(\lambda) < \rho. \end{cases} \tag{EC.28}$$

We first bound $R^{\text{id}}_{\max}$. By Proposition 3, if $q < \lambda$, $\text{Regret}(q) = C_G(q) - C_G(q^\star_G)$. By Lemma EC.1:

$$C_G(q) - C_G(q^\star_G) \le \max\{b,h\}|q - q^\star_G| \le (b+h)\max\{\rho, 1-\rho\}\lambda.$$

If $q \ge \lambda$, by Proposition 3,

$$\begin{aligned}
\text{Regret}(q) &= b(q^\star_G - q) + (b+h)\left[(q-\lambda) + \mathbb{E}_G\left[(\lambda - D)\mathbb{1}\{D < \lambda\} - (q^\star_G - D)\mathbb{1}\left\{D \le q^\star_G\right\}\right]\right] \\
&= bq^\star_G + hq - (b+h)\lambda + (b+h)\mathbb{E}_G\left[(\lambda - D)\mathbb{1}\{D < \lambda\} - (q^\star_G - D)\mathbb{1}\left\{D \le q^\star_G\right\}\right] \\
&\le bq^\star_G + hq + (b+h)\left(-\lambda + \lambda G^-(\lambda) - q^\star_G\rho - \mathbb{E}_G\left[D\mathbb{1}\left\{D \in (q^\star_G, \lambda)\right\}\right]\right).
\end{aligned}$$

Using the fact that $(b+h)\rho q^\star_G = bq^\star_G$ and $\mathbb{E}_G\left[D\mathbb{1}\left\{D \in (q^\star_G, \lambda)\right\}\right] \ge 0$, we obtain:

$$\text{Regret}(q) \le hM - (b+h)\lambda(1 - G^-(\lambda)) \le hM.$$

Putting these two bounds together, we obtain:

$$R^{\text{id}}_{\max} \le (b+h)\max\{\max\{\rho, (1-\rho)\}\lambda, (1-\rho)M\} = (b+h)\max\{\rho\lambda, (1-\rho)M\}.$$

We now bound $R^{\text{ui}}_{\max}$. By Lemma 2, for $q < \lambda$,

$$\text{Regret}(q) = b(M-q) + (b+h)\left[\mathbb{E}_G\left[(q-D)\mathbb{1}\{D \le q\} - (M-D)\mathbb{1}\{D < \lambda\}\right]\right] \le bM,$$

where the inequality follows from the fact that $-bq + (b+h)\mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}]$ is decreasing in q for $q < \lambda < q^\star_G$.

For $q \in \left[\lambda, q^\dagger_G\right]$,

$$\text{Regret}(q) = (b - (b+h)G^-(\lambda))(M-q) \le (b - (b+h)G^-(\lambda))(M-\lambda).$$

Finally, for $q > q^\dagger_G$,

$$\text{Regret}(q) = h(q-\lambda) \le h(M-\lambda).$$

Putting these three bounds together, we obtain:

$$R^{\text{ui}}_{\max} \le (b+h)\max\{\rho M, (\rho - G^-(\lambda))(M-\lambda), (1-\rho)(M-\lambda)\} = (b+h)\max\{\rho M, (1-\rho)(M-\lambda)\}.$$

We conclude by plugging these bounds back into (EC.28). Since $\zeta = \sqrt{\log(2/\delta)/2N} = c'\sqrt{\log N/N}$ for some $c' > 0$ and $\delta = c/\sqrt{N}$, if $G^-(\lambda) \ge \rho$, we have:

$$\mathbb{E}_G[\text{Regret}(q^{\text{alg}})] \le 2c'(b+h)(\lambda + \max\{\rho\lambda, (1-\rho)M\})\sqrt{\log N/N}.$$

If $G^-(\lambda) < \rho$:

$$\mathbb{E}_G[\text{Regret}(q^{\text{alg}})] \le \Delta + 2c'(b+h)\left(\max\left\{\frac{\rho}{1-\rho}, 1\right\}(M-\lambda) + \max\{\rho M, (1-\rho)(M-\lambda)\}\right)\sqrt{\log N/N}.$$

Renaming $c' = 2c'$, we obtain the result. □

B.5. Proof of Theorem 3

Proof. We consider an instance for which $\lambda \in (0,1)$, and present the set of "hard" distributions associated with each regime:

(i) strictly unidentifiable regime:

$$G^{\text{ui}}_0 = \text{Ber}(1-\rho+\delta^{\text{ui}}_0), \quad G^{\text{ui}}_1 = \text{Ber}(1-\rho+\delta^{\text{ui}}_0+\delta^{\text{ui}}_1),$$

where $\delta^{\text{ui}}_0 = \min\left\{\frac{\rho}{2}, \frac{1-\rho}{2}\right\}$ and $\delta^{\text{ui}}_1 = \min\left\{\frac{\rho}{4}, \frac{3\rho-1}{4}, \frac{1}{2}\sqrt{\frac{1-\rho}{N}}\right\}$. For these distributions, $G^{\text{ui},-}_0(\lambda) = \rho - \delta^{\text{ui}}_0$, and $G^{\text{ui},-}_1(\lambda) = \rho - (\delta^{\text{ui}}_0 + \delta^{\text{ui}}_1)$.

(ii) knife-edge regime:

$$G^{\text{ke}}_0 = \text{Ber}(1-\rho+\delta^{\text{ke}}), \quad G^{\text{ke}}_1 = \text{Ber}(1-\rho-\delta^{\text{ke}}),$$

where $\delta^{\text{ke}} = \min\left\{\frac{\rho}{2}, \frac{1-\rho}{2}, \frac{1}{4}\sqrt{\frac{1-\rho}{N}}\right\}$. Here, $G^{\text{ke},-}_0(\lambda) = \rho - \delta^{\text{ke}}$, $G^{\text{ke},-}_1(\lambda) = \rho + \delta^{\text{ke}}$.

(*iii*) strictly identifiable regime: G_0^{id}, G_1^{id} are respectively defined by cdf's:

$$G_0^{\mathrm{id}}(x) = \begin{cases} 0 & \forall\, x < 0 \\ \rho - \delta^{\mathrm{id}} & \forall\, x \in [0, H) \\ 1 & \forall\, x \geq H, \end{cases} \qquad G_1^{\mathrm{id}}(x) = \begin{cases} 0 & \forall\, x < 0 \\ \rho + \delta^{\mathrm{id}} & \forall\, x \in [0, H) \\ 1 & \forall\, x \geq H, \end{cases}$$

where $\delta^{\mathrm{id}} = \min\left\{\frac{\rho}{2}, \frac{1-\rho}{2}, \frac{1}{4}\sqrt{\frac{1-\rho}{N}}\right\}$, and $H = \lambda/2$. In this case, $G_0^{\mathrm{id},-}(\lambda) = G_1^{\mathrm{id},-}(\lambda) = 1$.

Note that all instances are effectively uncensored. For $k \in \{\mathrm{ui}, \mathrm{ke}\}$, this holds because $s^{\mathrm{off},\lambda} = \lambda \implies d_i^{\mathrm{off}} = 1$. For $k = \mathrm{id}$, this holds because $d_i^{\mathrm{off}} < \lambda$ for all i, which then implies that $d_i^{\mathrm{off}} = s_i^{\mathrm{off},\lambda}$ for all i.

For all $k \in \{\mathrm{id}, \mathrm{ui}, \mathrm{ke}\}$:

$$\sup_{G \in \mathcal{G}^k} \mathbb{E}_G\left[\sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) - \Delta_G\right] \geq \sup_{G \in \{G_0^k, G_1^k\}} \mathbb{E}_G\left[\sup_{F \in \mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) - \Delta_G\right]. \tag{EC.29}$$

For clarity, in the remainder of the proof we make clear the dependence of Δ and q^Δ on the underlying demand distribution G.

Recall, by Theorem 1, for any distribution G, the minimax optimal ordering quantity and minimax risk are respectively given by:

$$q_G^\Delta = \begin{cases} q_G^\star & \text{if } G^-(\lambda) \geq \rho \\ q_G^\dagger & \text{if } G^-(\lambda) < \rho \end{cases} \qquad \Delta_G = \begin{cases} 0 & \text{if } G^-(\lambda) \geq \rho \\ \frac{h(b-(b+h)G^-(\lambda))(M-\lambda)}{(b+h)(1-G^-(\lambda))} & \text{if } G^-(\lambda) < \rho, \end{cases} \tag{EC.30}$$

with $\frac{h(b-(b+h)G^-(\lambda))(M-\lambda)}{(b+h)(1-G^-(\lambda))} = (b-(b+h)G^-(\lambda))(M - q_G^\dagger) = h(q_G^\dagger - \lambda)$ for G such that $G^-(\lambda) < \rho$ (see proof of Theorem 1).

Lemma EC.6 provides the minimax optimal ordering quantity for the three sets of distributions described above. We defer its proof to Appendix B.5.1.

LEMMA EC.6. *The following holds, in each regime:*

(*i*) *strictly unidentifiable regime:*

$$q_G^\Delta = q_G^\dagger \quad \forall\, G \in \{G_0^{\mathrm{ui}}, G_1^{\mathrm{ui}}\}$$

(*ii*) *knife-edge regime:*

$$q_{G_0^{\mathrm{ke}}}^\Delta = q_{G_0^{\mathrm{ke}}}^\dagger, \quad q_{G_1^{\mathrm{ke}}}^\Delta = q_{G_1^{\mathrm{ke}}}^\star = 0$$

(*iii*) *strictly identifiable regime:*

$$q_{G_0^{\mathrm{id}}}^\Delta = q_{G_0^{\mathrm{id}}}^\star = H, \quad q_{G_1^{\mathrm{id}}}^\Delta = q_{G_1^{\mathrm{id}}}^\star = 0.$$

Lemma EC.7 next formalizes the idea that minimizing the worst-case regret in excess of Δ_G reduces to the problem of estimating q_G^Δ. We defer its proof to Appendix B.5.2.

LEMMA EC.7. *For $k \in \{\text{id}, \text{ui}, \text{ke}\}$, $G \in \{G_0^k, G_1^k\}$:*

$$\sup_{F \in \mathcal{F}(\lambda; G)} C_F(q^\pi) - C_F(q_F^\star) - \Delta_G \geq \begin{cases} (b+h)\delta_0^k \, |q^\pi - q_G^\Delta| & \text{if } k = \text{ui} \\ (b+h)\delta^k \, |q^\pi - q_G^\Delta| & \text{if } k \in \{\text{ke}, \text{id}\}. \end{cases}$$

Applying Lemma EC.7 to (EC.29), it suffices to lower bound the worst-case absolute difference between q^π and q_G^Δ. Namely:

$$\sup_{G \in \mathcal{G}^k} \mathbb{E}_G \left[\text{Regret}(q^\pi) - \Delta_G \right] \geq (b+h)\delta_0^k \sup_{G \in \{G_0^k, G_1^k\}} \mathbb{E}_G \left[|q^\pi - q_G^\Delta| \right] \quad \text{if } k = \text{ui} \tag{EC.31}$$

$$\sup_{G \in \mathcal{G}^k} \mathbb{E}_G \left[\text{Regret}(q^\pi) - \Delta_G \right] \geq (b+h)\delta^k \sup_{G \in \{G_0^k, G_1^k\}} \mathbb{E}_G \left[|q^\pi - q_G^\Delta| \right] \quad \text{if } k \in \{\text{ke}, \text{id}\}. \tag{EC.32}$$

We next lower bound $\sup_{G \in \{G_0^k, G_1^k\}} \mathbb{E}_G \left[|q^\pi - q_G^\Delta| \right]$ via a reduction to hypothesis testing. To formalize this, we introduce some additional notation. Let $\mathcal{D}^N$ be the set of all possible demand samples. For $k \in \{\text{id}, \text{ui}, \text{ke}\}$, let $\Psi^k : \mathcal{D}^N \mapsto \{G_0^k, G_1^k\}$ denote a mapping from the observed demand samples to a prediction of the underlying demand distribution. We moreover let $G_0^{k,N}$ and $G_1^{k,N}$ be the joint distributions of the N independent demand samples generated under G_0^k and G_1^k, respectively.

By Proposition 9.2.1. in Duchi (2024):

$$\begin{aligned} &\sup_{G \in \{G_0^k, G_1^k\}} \mathbb{E}_G \left[|q^\pi - q_G^\Delta| \right] \\ &\quad \geq \frac{1}{2} \left| q_{G_0^k}^\Delta - q_{G_1^k}^\Delta \right| \inf_{\Psi^k} \left\{ \mathbb{P}_{G_1^k} \left(\Psi^k(d^{\text{off}}) = G_0^k \right) + \mathbb{P}_{G_0^k} \left(\Psi^k(d^{\text{off}}) = G_1^k \right) \right\} \\ &\quad \geq \frac{1}{4} \underbrace{\left| q_{G_0^k}^\Delta - q_{G_1^k}^\Delta \right|}_{(I)} \underbrace{\exp\left(-d_{KL}(G_1^{k,N} \,||\, G_0^{k,N}) \right)}_{(II)}, \end{aligned} \tag{EC.33}$$

where $d_{KL}(G_1^{k,N} \,||\, G_0^{k,N})$ denotes the Kullback-Leibler (KL) divergence, and the second inequality follows from the Bretagnolle-Huber inequality (Lattimore and Szepesvári 2020).

Equation (EC.33) demonstrates the main drivers of regret in our setting. In particular, any algorithm must incur higher regret if (i) the minimax optimal ordering quantities under G_0^k and G_1^k are far apart (term (I)), or (ii) G_0^k and G_1^k are distributionally "close enough" that any hypothesis test that aims to distinguish between the two incurs high error rate (term (II)). Lemmas EC.8 and EC.9 respectively lower bound these two component terms for all regimes. We defer their proofs to Appendix B.5.3 and B.5.4, respectively.

LEMMA EC.8. *For $k \in \{\text{ui}, \text{ke}, \text{id}\}$:*

$$\left| q_{G_0^k}^\Delta - q_{G_1^k}^\Delta \right| \geq \begin{cases} \delta_1^{\text{ui}} \cdot \frac{h(M-\lambda)}{b+h} & \text{if } k = \text{ui} \\ \lambda & \text{if } k = \text{ke} \\ H & \text{if } k = \text{id}. \end{cases}$$

LEMMA EC.9. *For $k \in \{\text{ui}, \text{ke}, \text{id}\}$:*

$$\exp\Big(-d_{KL}(G_1^{k,N} \,||\, G_0^{k,N})\Big) \geq \begin{cases} \exp\left(-2N \cdot \frac{(\delta_1^k)^2}{1-\rho}\right) & \text{if } k = \text{ui} \\ \exp\left(-8N \cdot \frac{(\delta^k)^2}{1-\rho}\right) & \text{if } k \in \{\text{ke}, \text{id}\}. \end{cases}$$

Applying these two lemmas to (EC.33), and plugging back into (EC.31) and (EC.32), we obtain:

$$\sup_{G \in \mathcal{G}^k} \mathbb{E}_G\left[\text{Regret}(q^\pi) - \Delta_G\right] \geq \begin{cases} \frac{1}{4}h(M-\lambda)\delta_0^{\text{ui}}\delta_1^{\text{ui}} \exp\left(-2N \cdot \frac{(\delta_1^{\text{ui}})^2}{1-\rho}\right) & \text{if } k = \text{ui} \\ \frac{1}{4}\lambda(b+h)\delta^{\text{ke}} \exp\left(-8N \cdot \frac{(\delta^{\text{ke}})^2}{1-\rho}\right) & \text{if } k = \text{ke} \\ \frac{1}{4}H(b+h)\delta^{\text{id}} \exp\left(-8N \cdot \frac{(\delta^{\text{id}})^2}{1-\rho}\right) & \text{if } k = \text{id}. \end{cases}$$

For $\delta_0^{\text{ui}} = \min\left\{\frac{\rho}{2}, \frac{1-\rho}{2}\right\}$ and $\delta_1^{\text{ui}} = \min\left\{\frac{\rho}{4}, \frac{3\rho-1}{4}, \frac{1}{\sqrt{2 \cdot 2N/(1-\rho)}}\right\}$, we obtain, for $k = \text{ui}$:

$$\begin{aligned} &\sup_{G \in \mathcal{G}^k} \mathbb{E}_G\left[\text{Regret}(q^\pi) - \Delta_G\right] \\ &\geq \frac{h(M-\lambda)}{8} \min\{\rho, 1-\rho\} \cdot \min\left\{\frac{\rho}{4}, \frac{3\rho-1}{4}, \frac{\sqrt{1-\rho}}{2\sqrt{N}}\right\} \cdot e^{-1/2} \\ &\geq \frac{h(M-\lambda)\sqrt{1-\rho}\min\{\rho, 1-\rho\} \cdot \min\{\rho, 3\rho-1\} \cdot e^{-1/2}}{64\sqrt{N}}, \end{aligned}$$

where the second inequality follows from the fact that $\min\{x, y\} \geq xy$ for $x, y \in [0, 1]$.

For $\delta^{\text{ke}} = \min\left\{\frac{\rho}{2}, \frac{1-\rho}{2}, \frac{1}{\sqrt{2 \cdot 8N/(1-\rho)}}\right\}$, we obtain, for $k = \text{ke}$:

$$\begin{aligned} &\sup_{G \in \mathcal{G}^k} \mathbb{E}_G\left[\sup_{F \in \mathcal{F}(\lambda; G)} C_F(q^\pi) - C_F(q_F^\star) - \Delta_G\right] \\ &\geq \frac{\lambda(b+h)}{4} \min\left\{\frac{\rho}{2}, \frac{1-\rho}{2}, \frac{1}{4}\sqrt{\frac{1-\rho}{N}}\right\} \cdot e^{-1/2} \\ &\geq \frac{\lambda(b+h)\sqrt{1-\rho}\min\{\rho, 1-\rho\}e^{-1/2}}{32\sqrt{N}}. \end{aligned}$$

Finally, using the same instantiation of δ^{id} as δ^{ke}, and setting $H = \lambda/2$, we obtain:

$$\begin{aligned} &\sup_{G \in \mathcal{G}^k} \mathbb{E}_G\left[\sup_{F \in \mathcal{F}(\lambda; G)} C_F(q^\pi) - C_F(q_F^\star) - \Delta_G\right] \\ &\geq \frac{H(b+h)\sqrt{1-\rho}\min\{\rho, 1-\rho\}e^{-1/2}}{32\sqrt{N}} \\ &= \frac{\lambda(b+h)\sqrt{1-\rho}\min\{\rho, 1-\rho\}e^{-1/2}}{64\sqrt{N}}. \end{aligned}$$

□

B.5.1. Proof of Lemma EC.6

Proof. We proceed regime-by-regime.

Case I: Strictly unidentifiable regime. This follows from the fact that $G^-(\lambda) < \rho$ for all $G \in \{G_0^{\text{ui}}, G_1^{\text{ui}}\}$.

Case II: Knife-edge regime. This follows from the fact that $G_0^{\text{ke},-}(\lambda) < \rho$, and $G_1^{\text{ke},-}(\lambda) \geq \rho$. Moreover, since G_1^{ke} is Bernoulli, it must then be that $q^\star_{G_1^{\text{ke}}} = 0$.

Case III: Strictly identifiable regime. This follows from the fact that $G^-(\lambda) \geq \rho$ for all $G \in \{G_0^{\text{id}}, G_1^{\text{id}}\}$. For $G = G_0^{\text{id}}$, $\mathbb{P}_G(D = 0) < \rho$, which implies that $q^\star_G = H$. For $G = G_1^{\text{id}}$, $\mathbb{P}_G(D = 0) \geq \rho$, which implies $q^\star_G = 0$. □

B.5.2. Proof of Lemma EC.7

Proof. We prove the lemma separately for each regime. In all cases, we rely repeatedly on the closed-form characterizations of the minimax risk established in Lemma 2 and Proposition 3, bounding the loss based on the location of q^π relative to λ.

Case I: Strictly unidentifiable regime. By Lemma EC.6, for all $G \in \{G_0^{\text{ui}}, G_1^{\text{ui}}\}$, $q^\Delta_G = q^\dagger_G$.

Suppose first that $q^\pi < \lambda$. By Lemma 2, we have:

$$\begin{aligned}\sup_{F \in \mathcal{F}(\lambda; G)} C_F(q^\pi) - C_F(q^\star_F) &= b(M - q^\pi) + (b+h)\Big(q^\pi \mathbb{P}_G(D=0) - M\mathbb{P}_G(D=0)\Big) \\ &= (M - q^\pi)(b - (b+h)G^-(\lambda)),\end{aligned}$$

since $\lambda \in (0,1)$ and G is Bernoulli. Subtracting $\Delta_G = (b - (b+h)G^-(\lambda))(M - q^\dagger_G)$ on both sides, we have:

$$\begin{aligned}\sup_{F \in \mathcal{F}(\lambda; G)} C_F(q^\pi) - C_F(q^\star_F) - \Delta_G &= (M - q^\pi)(b - (b+h)G^-(\lambda)) - (b - (b+h)G^-(\lambda))(M - q^\dagger_G) \\ &= (q^\dagger_G - q^\pi)(b - (b+h)G^-(\lambda)) \\ &\geq (q^\dagger_G - q^\pi)(b+h)\delta_0^{\text{ui}},\end{aligned}$$

where the final inequality follows from the fact that $G^-(\lambda) \leq \rho - \delta_0^{\text{ui}}$ for $G \in \{G_0^{\text{ui}}, G_1^{\text{ui}}\}$, and $q^\dagger_G \geq \lambda > q^\pi$.

Consider now the case where $q^\pi \in [\lambda, q^\dagger_G]$. By Lemma 2, for $G^-(\lambda) < \rho$, we have:

$$\begin{aligned}\sup_{F \in \mathcal{F}(\lambda; G)} C_F(q^\pi) - C_F(q^\star_F) - \Delta_G &= (b - (b+h)G^-(\lambda))(M - q^\pi) - (b - (b+h)G^-(\lambda))(M - q^\dagger_G) \\ &= (b - (b+h)G^-(\lambda))(q^\dagger_G - q^\pi) \\ &\geq (b+h)\delta_0^{\text{ui}}(q^\dagger_G - q^\pi),\end{aligned}$$

where again we used the fact that $G^-(\lambda) \leq \rho - \delta_0^{\text{ui}}$ for $G \in \{G_0^{\text{ui}}, G_1^{\text{ui}}\}$.

Finally, if $q^\pi \in (q^\dagger_G, M]$, using the alternative representation of $\Delta_G = h(q^\dagger_G - \lambda)$, by Lemma 2:

$$\sup_{F \in \mathcal{F}(\lambda; G)} C_F(q^\pi) - C_F(q^\star_F) - \Delta_G = h(q^\pi - \lambda) - h(q^\dagger_G - \lambda) = h(q^\pi - q^\dagger_G).$$

Combining the three cases and using the fact that $\delta_0^{\mathrm{ui}} < 1-\rho \implies (b+h)\delta_0^{\mathrm{ui}} < h$, we obtain the lower bound for the strictly unidentifiable regime.

Case II: Knife-edge regime. If $G = G_0^{\mathrm{ke}}$, since $G_0^{\mathrm{ke},-}(\lambda) < \rho$, by the same arguments as above:

$$\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) - \Delta_G \geq (b+h)\delta^{\mathrm{ke}}|q^\pi - q_G^\dagger|. \tag{EC.34}$$

If $G = G_1^{\mathrm{ke}}$, by (EC.30), $\Delta_G = 0$, since $G_1^{\mathrm{ke},-}(\lambda) \geq \rho$ by construction. Moreover, by Lemma EC.6, $q_G^\Delta = q_G^\star = 0$.

Suppose first that $q^\pi < \lambda$. Applying Proposition 3 to $G = G_1^{\mathrm{ke}}$, we have:

$$\begin{aligned}\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) &= -bq^\pi + (b+h)q^\pi \mathbb{P}_G(D=0)\\ &= -bq^\pi + (b+h)q^\pi(\rho+\delta^{\mathrm{ke}})\\ &= q^\pi(b+h)\delta^{\mathrm{ke}},\end{aligned}$$

where the first equality simplified the expression in Proposition 3 using the fact that $\lambda \in (0,1)$, and $q_G^\star = 0$.

Suppose now that $q^\pi \geq \lambda$. Again, by Proposition 3:

$$\begin{aligned}\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) &= -bq^\pi + (b+h)(q^\pi - \lambda + \lambda\mathbb{P}_G(D=0))\\ &= hq^\pi - \lambda(b+h)\mathbb{P}_G(D=1)\\ &\geq q^\pi\Big(h - (b+h)(1-\rho-\delta^{\mathrm{ke}})\Big)\\ &= q^\pi(b+h)\delta^{\mathrm{ke}}.\end{aligned}$$

Hence, in both cases:

$$\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) \geq (q^\pi - q_G^\star)(b+h)\delta^{\mathrm{ke}}. \tag{EC.35}$$

Since $q_G^\Delta = q_G^\dagger$ for $G = G_0^{\mathrm{ke}}$, and $q_G^\Delta = q_G^\star$ for $G = G_1^{\mathrm{ke}}$, putting (EC.34) and (EC.35) together we obtain:

$$\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) - \Delta_G \geq (b+h)\delta^{\mathrm{ke}}|q^\pi - q^\Delta| \quad \forall\, G \in \{G_0, G_1\}.$$

Case III: Strictly identifiable regime. By Lemma EC.6, $q_{G_0^{\mathrm{id}}}^\Delta = q_{G_0^{\mathrm{id}}}^\star = H$, and $q_{G_1^{\mathrm{id}}}^\Delta = q_{G_1^{\mathrm{id}}}^\star = 0$. Moreover, for $G \in \{G_0^{\mathrm{id}}, G_1^{\mathrm{id}}\}$, $\Delta_G = 0$, by (EC.30).

Suppose first that $G = G_0^{\mathrm{id}}$. By Proposition 3, if $q^\pi < \lambda$:

$$\begin{aligned}&\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star)\\ &= b(H-q^\pi) + (b+h)\Big(q^\pi\mathbb{P}_G(D=0) + (q^\pi - H)\mathbb{1}\{H\leq q^\pi\}\mathbb{P}_G(D=H) - H\mathbb{P}_G(D=0)\Big)\end{aligned}$$

$$
\begin{aligned}
&= (H - q^\pi)(b - (b+h)\mathbb{P}_G(D=0) - (b+h)\mathbb{P}_G(D=H)\mathbb{1}\{H \le q^\pi\}) \\
&= (b+h)(H-q^\pi)\Big(\delta^{\mathrm{id}} - (1-\rho+\delta^{\mathrm{id}})\mathbb{1}\{H \le q^\pi\}\Big) \\
&= (b+h)(1-\rho)(q^\pi - H)\mathbb{1}\{H \le q^\pi\} + (b+h)(H-q^\pi)\delta^{\mathrm{id}}\mathbb{1}\{H > q^\pi\} \\
&\ge (b+h)|H-q^\pi|\delta^{\mathrm{id}} = (b+h)|q_G^\Delta - q^\pi|\delta^{\mathrm{id}}, \qquad \text{(EC.36)}
\end{aligned}
$$

where the inequality follows from the fact that $\delta^{\mathrm{id}} < 1-\rho$.

If $q^\pi \ge \lambda$, on the other hand, by Proposition 3:

$$
\begin{aligned}
&\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) \\
&= b(H-q^\pi) + (b+h)\Big(q^\pi - \lambda + \lambda\mathbb{P}_G(D=0) + (\lambda - H)\mathbb{P}_G(D=H) - H\mathbb{P}_G(D=0)\Big) \\
&= b(H-q^\pi) + (b+h)(q^\pi - H) \\
&= h(q^\pi - H) \\
&\ge (b+h)|H-q^\pi|\delta^{\mathrm{id}} = (b+h)|q_G^\Delta - q^\pi|\delta^{\mathrm{id}}, \qquad \text{(EC.37)}
\end{aligned}
$$

where the inequality follows from the fact that $q^\pi \ge \lambda > H$, and $\delta^{\mathrm{id}} < 1-\rho$.

Putting (EC.36) and (EC.37) together, we obtain the lower bound for $G = G_0^{\mathrm{id}}$.

Suppose now that $G = G_1^{\mathrm{id}}$. By Proposition 3, if $q^\pi < \lambda$:

$$
\begin{aligned}
\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) &= -bq^\pi + (b+h)\Big(q^\pi\mathbb{P}_G(D=0) + (q^\pi - H)\mathbb{P}_G(D=H)\mathbb{1}\{H \le q^\pi\}\Big) \\
&= -bq^\pi + (b+h)\Big(q^\pi(\rho+\delta^{\mathrm{id}}) + (q^\pi - H)\mathbb{1}\{H \le q^\pi\}(1-\rho-\delta^{\mathrm{id}})\Big) \\
&= (b+h)\Big(q^\pi\delta^{\mathrm{id}} + (q^\pi - H)(1-\rho-\delta^{\mathrm{id}})\Big)\mathbb{1}\{H \le q^\pi\} \\
&\quad + (b+h)\delta^{\mathrm{id}}q^\pi\mathbb{1}\{H > q^\pi\} \\
&\ge (b+h)\delta^{\mathrm{id}}q^\pi = (b+h)\delta^{\mathrm{id}}|q^\pi - q_G^\Delta|, \qquad \text{(EC.38)}
\end{aligned}
$$

where the inequality follows from the fact that $\delta^{\mathrm{id}} < 1-\rho$.

If $q^\pi \ge \lambda$, on the other hand, by Proposition 3:

$$
\begin{aligned}
\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^\pi) - C_F(q_F^\star) &= -bq^\pi + (b+h)\Big[(q^\pi - \lambda) + \lambda\mathbb{P}_G(D=0) + (\lambda - H)\mathbb{P}_G(D=H)\Big] \\
&= hq^\pi - H(b+h)(1-\rho-\delta^{\mathrm{id}}) \\
&\ge \Big(h - (b+h)(1-\rho-\delta^{\mathrm{id}})\Big)q^\pi \\
&= (b+h)\delta^{\mathrm{id}}q^\pi = (b+h)\delta^{\mathrm{id}}|q^\pi - q_G^\Delta|, \qquad \text{(EC.39)}
\end{aligned}
$$

where the inequality follows from the fact that $q^\pi \ge \lambda > H$.

Putting (EC.38) and (EC.39) together, we obtain the lower bound for $G = G_1^{\mathrm{id}}$. □

B.5.3. Proof of Lemma EC.8

Proof. We proceed case-by-case.

Case I: Strictly unidentifiable regime. By Lemma EC.6, in this setting, for $G \in \{G_0^{\text{ui}}, G_1^{\text{ui}}\}$,

$$q_G^{\Delta} = q_G^{\dagger} = \frac{h\lambda + (b - (b+h)G^-(\lambda))M}{(b+h)(1-G^-(\lambda))}.$$

As in the proof of Theorem 2, define $Q_G^{\dagger}(\cdot)$ as:

$$Q_G^{\dagger}(x) = \frac{h\lambda + (b-(b+h)x)M}{(b+h)(1-x)}.$$

Recall, $G_0^{\text{ui},-}(\lambda) = \rho - \delta_0^{\text{ui}}$, and $G_1^{\text{ui},-}(\lambda) = \rho - \delta_0^{\text{ui}} - \delta_1^{\text{ui}}$. Hence, by the mean value theorem, for some $c \in (\rho - \delta_0^{\text{ui}} - \delta_1^{\text{ui}}, \rho - \delta_0^{\text{ui}})$:

$$\begin{aligned} Q_G^{\dagger}(\rho - \delta_0^{\text{ui}}) - Q_G^{\dagger}(\rho - (\delta_0^{\text{ui}} + \delta_1^{\text{ui}})) = \delta_1^{\text{ui}} \cdot \frac{dq_G^{\dagger}}{dx}\bigg|_{x=c} = \delta_1^{\text{ui}} \cdot \frac{h(\lambda - M)}{(b+h)(1-c)^2} \\ \implies \left|Q_G^{\dagger}(\rho - \delta_0^{\text{ui}}) - Q_G^{\dagger}(\rho - (\delta_0^{\text{ui}} + \delta_1^{\text{ui}}))\right| \geq \delta_1^{\text{ui}} \cdot \frac{h(M-\lambda)}{b+h}, \end{aligned} \tag{EC.40}$$

thus providing the bound for the unidentifiable regime.

Case II: Knife-edge regime. By Lemma EC.6, in this setting $q_{G_0^{\text{ke}}}^{\Delta} = q_{G_0^{\text{ke}}}^{\dagger}$, and $q_{G_1^{\text{ke}}}^{\Delta} = q_{G_1^{\text{ke}}}^{\star} = 0$. Hence:

$$|q_{G_0^{\text{ke}}}^{\Delta} - q_{G_1^{\text{ke}}}^{\Delta}| = \frac{h\lambda + (b-(b+h)(\rho - \delta^{\text{ke}}))M}{(b+h)(1-\rho+\delta^{\text{ke}})} \geq \frac{h\lambda + (b-(b+h)\rho)M}{(b+h)(1-\rho)} = \lambda,$$

where the inequality follows from the fact that $Q_G^{\dagger}(x)$ is decreasing in x, and $\rho - \delta^{\text{ke}} \leq \rho$ for all $\delta \geq 0$.

Case III: Strictly identifiable regime. This follows from the fact that, by Lemma EC.6, $q_{G_0^{\text{id}}}^{\Delta} = H$ and $q_{G_1^{\text{id}}}^{\Delta} = 0$. □

B.5.4. Proof of Lemma EC.9

Proof. Since we are in the uncensored setting and demand samples are drawn i.i.d., we have that $d_{KL}(G_1^{k,N} \,||\, G_0^{k,N}) = N d_{KL}(G_1^k \,||\, G_0^k)$. Moreover, across all regimes, G_0^k and G_1^k are distributions with support over two point masses. Hence, we can use reverse Pinsker's inequality to upper bound the KL divergence from $G_1^{k,N}$ to $G_0^{k,N}$ (Sason (2015), Eqn. 10). Namely, for any two distributions G_0^k, G_1^k with support on the same two point masses a and b, with $a > b$, we have:

$$d_{KL}(G_1^k \,||\, G_0^k) = 2 \cdot \frac{(p-q)^2}{q},$$

where p and q respectively denote the mass G_1^k and G_0^k place on a. We apply this to each regime, to obtain:

Case I: Strictly unidentifiable regime.

$$\exp\left(-d_{KL}(G_1^{\text{ui},N} \,||\, G_0^{\text{ui},N})\right) \geq \exp\left(-2N \cdot \frac{(\delta_1^{\text{ui}})^2}{1-\rho+\delta_0^{\text{ui}}}\right) \geq \exp\left(-2N \cdot \frac{(\delta_1^{\text{ui}})^2}{1-\rho}\right).$$

Case II: Knife-edge regime.

$$\exp\left(-d_{KL}(G_1^{\mathrm{ke},N} \,||\, G_0^{\mathrm{ke},N})\right) \geq \exp\left(-2N \cdot \frac{(2\delta^{\mathrm{ke}})^2}{1-\rho+\delta^{\mathrm{ke}}}\right) \geq \exp\left(-8N \cdot \frac{(\delta^{\mathrm{ke}})^2}{1-\rho}\right).$$

Case III: Strictly identifiable regime.

$$\exp\left(-d_{KL}(G_1^{\mathrm{id},N} \,||\, G_0^{\mathrm{id},N})\right) \geq \exp\left(-2N \cdot \frac{(2\delta^{\mathrm{id}})^2}{1-\rho+\delta^{\mathrm{id}}}\right) \geq \exp\left(-8N \cdot \frac{(\delta^{\mathrm{id}})^2}{1-\rho}\right).$$

□

Appendix C: Analysis of RCN$^+$

THEOREM EC.1. *Fix $\delta \in (0,1)$, and let $\zeta_k = \sqrt{\frac{\log(2(K-1)/\delta)}{2N_k}}$ for $k \neq K$ and $\zeta_K = \sqrt{\frac{\log(2/\delta)}{2N_k}}$. With probability at least $1-3\delta$, Algorithm 2 outputs an ordering quantity q^{alg} such that:*

(i) if $G^-(q_k^{\mathrm{off}}) \geq \rho + 2\zeta_k$ for some $k \in [K]$:

$$\mathrm{Regret}(q^{\mathrm{alg}}) \leq \lambda(b+h)\sqrt{\frac{\log(2/\delta)}{2\sum_{k\in\mathcal{U}_1} N_k}},$$

where $\mathcal{U}_1 = \{k : G^-(q_k^{\mathrm{off}}) \geq \rho + 2\zeta_k\}$.

(ii) if $G^-(\lambda) \geq \rho$ and $G^-(q_k^{\mathrm{off}}) < \rho + 2\zeta_k$ for all $k \in [K]$:

$$\mathrm{Regret}(q^{\mathrm{alg}}) \leq \lambda(b+h)\max\left\{2\zeta_K, \sqrt{\frac{\log(2/\delta)}{2\min_{k\in\mathcal{U}_2} N_k}}\right\},$$

where $\mathcal{U}_2 = \{k : G^-(q_k^{\mathrm{off}}) \geq \rho\}$.

(iii) if $G^-(\lambda) \in [\rho - 2\zeta_K, \rho)$:

$$\mathrm{Regret}(q^{\mathrm{alg}}) \leq \Delta + 2\max\left\{\frac{b}{h}, 1\right\}(M-\lambda)(b+h)\zeta_K$$

(iv) if $G^-(\lambda) < \rho - 2\zeta_K$:

$$\mathrm{Regret}(q^{\mathrm{alg}}) \leq \Delta + \max\left\{\frac{b}{h}, 1\right\}(M-\lambda)(b+h)\zeta_K.$$

Proof. Abusing notation, we define $q^{\dagger}_{\widehat{G}} = \frac{bM+h\lambda-(b+h)\widehat{G}_K^-(\lambda)M}{(b+h)(1-\widehat{G}_K^-(\lambda))}$. We also let $N_{\mathrm{est}} = \sum_{k\in\mathcal{U}_{\mathrm{est}}} N_k$, $\widehat{G}(x) = \frac{1}{N_{\mathrm{est}}}\sum_{k\in\mathcal{U}_{\mathrm{est}}}\sum_{i\in[N_k]} \mathbb{1}\left\{s_{ki}^{\mathrm{off}} \leq x\right\}$, and $\widehat{G}^{\mathrm{uc}}(x) = \frac{1}{N_{\mathrm{est}}}\sum_{k\in\mathcal{U}_{\mathrm{est}}}\sum_{i\in[N_k]} \mathbb{1}\left\{d_{ki}^{\mathrm{off}} \leq x\right\}$.

Recall, by Proposition EC.3,

$$\mathrm{Regret}(q^{\mathrm{alg}}) \leq \Delta + \sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^{\mathrm{alg}}) - C_F(q^{\Delta}). \tag{EC.41}$$

Therefore, as in the proof of Theorem 2, we focus on bounding $\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q^{\mathrm{alg}}) - C_F(q^{\Delta})$. We will rely heavily on the following facts in the remainder of our proof.

LEMMA EC.10. *Let $\mathcal{E} = \left\{|\widehat{G}_k^-(q_k^{\mathrm{off}}) - G^-(q_k^{\mathrm{off}})| < \zeta_k \text{ for all } k \in [K]\right\}$. Then, $\mathbb{P}_G(\mathcal{E}) \geq 1-2\delta$.*

Lemma EC.10 is an analog of Lemma EC.2, where we additionally union bound over all selling seasons $k \in [K]$. It similarly follows from the fact that $\mathbb{1}\{s_{ki}^{\text{off}} < q_k^{\text{off}}\} = \mathbb{1}\{d_{ki}^{\text{off}} < q_k^{\text{off}}\}$ for all $k \in [K]$, $i \in [N_k]$. We omit the proof of this fact as such.

With these facts in hand, we prove each of the regret bounds separately, beginning with the two "extreme" cases of $G^-(q_k^{\text{off}}) \geq \rho + 2\zeta_k$ for some $k \in [K]$ and $G^-(\lambda) < \rho - 2\zeta_K$. Moreover, in the remainder of the proof we condition on the good event $\mathcal{E}$ defined in Lemma EC.10.

Case I: $G^-(q_k^{\text{off}}) \geq \rho + 2\zeta_k$ *for some* $k \in [K]$. In this case, the following facts hold:

1. $G^-(\lambda) \geq \rho$. Therefore, $q_G^\star < \lambda$.
2. Under event $\mathcal{E}$:
   (a) For all $k \in \mathcal{U}_1$, $\widehat{G}_k^-(q_k^{\text{off}}) \geq \rho + \zeta_k$. Hence, $\mathcal{U}_{\text{est}} \neq \emptyset$, and $q^{\text{alg}} = \inf\left\{x \mid \widehat{G}(x) \geq \rho\right\} < \min_{k \in \mathcal{U}_{\text{est}}} q_k^{\text{off}}$, where the strict inequality holds because $\widehat{G}_k^-(q_k^{\text{off}}) > \rho$ for all $k \in \mathcal{U}_{\text{est}}$.
   (b) For all $k \in \mathcal{U}_{\text{est}}$, $G^-(q_k^{\text{off}}) \geq \rho$. Therefore, $q_G^\star < \min_{k \in \mathcal{U}_{\text{est}}} q_k^{\text{off}}$.

Since $G^-(\lambda) \geq \rho$, we can apply Proposition 3, obtaining:

$$\text{Regret}(q^{\text{alg}}) = C_G(q^{\text{alg}}) - C_G(q_G^\star). \tag{EC.42}$$

Suppose first that $q^{\text{alg}} \leq q_G^\star$. By the same arguments as those used to establish (EC.25), we have:

$$\begin{aligned} C_G(q^{\text{alg}}) - C_G(q_G^\star) &\leq (b+h)(q_G^\star - q^{\text{alg}})(\rho - G(q^{\text{alg}})) \\ &\leq (b+h)\lambda(\widehat{G}(q^{\text{alg}}) - G(q^{\text{alg}})), \end{aligned}$$

where we use the fact that $\widehat{G}(q^{\text{alg}}) \geq \rho$ by definition, and moreover $\max\{q_G^\star, q^{\text{alg}}\} \leq \lambda \implies q_G^\star - q^{\text{alg}} \leq \lambda$.

Now, since $q^{\text{alg}} \leq q_G^\star$ and $q_G^\star < \min_{k \in \mathcal{U}_{\text{est}}} q_k^{\text{off}}$, $q^{\text{alg}} < \min_{k \in \mathcal{U}_{\text{est}}} q_k^{\text{off}}$. Hence, $\mathbb{1}\{s_{ki}^{\text{off}} \leq q^{\text{alg}}\} = \mathbb{1}\{d_{ki}^{\text{off}} \leq q^{\text{alg}}\}$ for all $k \in \mathcal{U}_{\text{est}}$, $i \in [N_k]$. Therefore, $\widehat{G}(q^{\text{alg}}) = \widehat{G}^{\text{uc}}(q^{\text{alg}})$. Plugging this into the above, we obtain:

$$C_G(q^{\text{alg}}) - C_G(q_G^\star) \leq (b+h)\lambda(\widehat{G}^{\text{uc}}(q^{\text{alg}}) - G(q^{\text{alg}})).$$

Applying the DKW inequality, we have that with probability at least $1-\delta$,

$$C_G(q^{\text{alg}}) - C_G(q_G^\star) \leq (b+h)\lambda\sqrt{\frac{\log(2/\delta)}{2N_{\text{est}}}}.$$

Suppose now that $q^{\text{alg}} > q_G^\star$. In this case, by the same arguments as those used to establish (EC.27), we have:

$$\begin{aligned} C_G(q^{\text{alg}}) - C_G(q_G^\star) &\leq (b+h)(q^{\text{alg}} - q_G^\star)(G^-(q^{\text{alg}}) - \rho) \\ &\leq (b+h)(q^{\text{alg}} - q_G^\star)\left(\lim_{x \to q^{\text{alg}}} \widehat{G}^{\text{uc}}(x) + \sqrt{\frac{\log(2/\delta)}{2N_{\text{est}}}} - \rho\right) \end{aligned}$$

$$\leq (b+h)(q^{\text{alg}} - q^\star_G)\left(\lim_{x \to q^{\text{alg}}} \widehat{G}(x) + \sqrt{\frac{\log(2/\delta)}{2N_{\text{est}}}} - \rho\right)$$
$$\leq (b+h)\lambda\sqrt{\frac{\log(2/\delta)}{2N_{\text{est}}}},$$

where the second inequality applies the DKW inequality, the third inequality uses the fact that $\mathbb{1}\left\{s^{\text{off}}_{ki} \leq x\right\} \geq \mathbb{1}\left\{d^{\text{off}}_{ki} \leq x\right\}$ for all x, the fourth inequality uses the fact that $\lim_{x \to q^{\text{alg}}} \widehat{G}(x) < \rho$ by definition of q^{alg}, and the final inequality uses the fact that $q^{\text{alg}} \leq \lambda$.

Using the fact that $k \in \mathcal{U}_1 \implies k \in \mathcal{U}_{\text{est}}$, we have that $N_{\text{est}} \geq \sum_{k \in \mathcal{U}_1} N_k$. Applying this bound to the above concludes the proof of the first case.

Case II: $G^-(\lambda) < \rho - 2\zeta_K$. In this case, by Theorem 1, $q^\Delta = q^\dagger_G$. Moreover, under event $\mathcal{E}$, our algorithm outputs $q^{\text{alg}} = q^\dagger_{\widehat{G}}$. To see why this is the case, observe that, under $\mathcal{E}$, $\widehat{G}^-_K(\lambda) < \rho - \zeta_K$. It remains to establish that $\widehat{G}^-_k(q^{\text{off}}_k) < \rho + \zeta_k$ for all $k \leq K-1$. Suppose for contradiction that $\widehat{G}^-_k(q^{\text{off}}_k) \geq \rho + \zeta_k$ for some $k \leq K-1$. Then, under $\mathcal{E}$, we must have $G^-(q^{\text{off}}_k) \geq \rho$, a contradiction, since $G^-(\lambda) < \rho - 2\zeta_K$ by assumption.

Hence, the same analysis as that of Case II in the proof of Theorem 2 applies, and we obtain:

$$\text{Regret}(q^{\text{alg}}) \leq \Delta + \frac{\max\{b,h\}(M-\lambda)}{1-\rho}\zeta_K, \tag{EC.43}$$

completing the proof of the second case.

Having analyzed the two extreme cases, we proceed to the "hard" regime, i.e., $G^-(\lambda) \in [\rho - 2\zeta_K, \rho + 2\zeta_K)$.

Case III: $G^-(\lambda) \in [\rho, \rho + 2\zeta_K)$, $G^-(q^{\text{off}}_k) < \rho + 2\zeta_k$ *for all* $k \leq K-1$. By Theorem 1, $q^\Delta = q^\star_G$, and $\Delta = 0$. Moreover, given $\mathcal{E}$, $\widehat{G}^-_k(q^{\text{off}}_k) \in (\rho - \zeta_k, \rho + 3\zeta_k)$ for all $k \in \mathcal{U}_2$, and $\widehat{G}^-(q^{\text{off}}_k) < \rho + \zeta_k$ for all $k \notin \mathcal{U}_2$.

a. Suppose first that $\widehat{G}^-_k(q^{\text{off}}_k) \in [\rho + \zeta_k, \rho + 3\zeta_k)$ for some $k \in \mathcal{U}_2$. Then, $\mathcal{U}_{\text{est}} \neq \emptyset$, and identical arguments as those used for Case I yield:

$$\text{Regret}(q^{\text{alg}}) \leq (b+h)\lambda\sqrt{\frac{\log(2/\delta)}{2N_{\text{est}}}}. \tag{EC.44}$$

Using the fact that $N_{\text{est}} \geq \min_{k \in \mathcal{U}_2} N_k$ in this case, we obtain:

$$\text{Regret}(q^{\text{alg}}) \leq (b+h)\lambda\sqrt{\frac{\log(2/\delta)}{2\min_{k \in \mathcal{U}_2} N_k}}. \tag{EC.45}$$

b. Suppose now that $\widehat{G}^-_k(q^{\text{off}}_k) \in (\rho - \zeta_k, \rho + \zeta_k)$ for all $k \in \mathcal{U}_2$. In this case, $q^{\text{alg}} = \lambda$. Using the same arguments as those used in Case III of the proof of Theorem 2, we have:

$$\text{Regret}(q^{\text{alg}}) \leq (b+h)(\lambda - q^\star_G)(G^-(\lambda) - \rho).$$

Using the fact that $q^\star_G \in [0, \lambda)$ for $G^-(\lambda) \geq \rho$, and $G^-(\lambda) \leq \rho + 2\zeta_K$ by assumption, we obtain the final bound of:

$$\text{Regret}(q^{\text{alg}}) \leq 2\lambda(b+h)\zeta_K.$$

Case IV: $G^-(\lambda) \in [\rho - 2\zeta_K, \rho)$. By Theorem 1, $q^\Delta = q^\dagger_G$. Moreover, under $\mathcal{E}$, $\widehat{G}^-_K(\lambda) \in (\rho - 3\zeta_K, \rho + \zeta_K)$, and $\widehat{G}^-_k(q^{\text{off}}_k) < \rho + \zeta_k$ for all $k \leq K-1$.

a. Suppose first that $\widehat{G}^-_K(\lambda) < \rho - \zeta_K$. In this case, the algorithm outputs $q^{\text{alg}} = q^\dagger_{\widehat{G}}$. By the same arguments as those used in Case II, we obtain:

$$\text{Regret}(q^{\text{alg}}) \leq \Delta + \frac{\max\{b,h\}(M-\lambda)}{(1-\rho)}\zeta_K.$$

b. Suppose now that $\widehat{G}^-_K(\lambda) \geq \rho - \zeta_K$. In this case, $q^{\text{alg}} = \lambda$. By the same arguments as those used in Case IV of the proof of Theorem 2, we have:

$$\text{Regret}(q^{\text{alg}}) \leq \Delta + \max\{b,h\}\frac{2(M-\lambda)}{1-\rho}\zeta_K.$$

□

λ	44.50	57.21	69.93	82.64	95.36	108.07	120.79	133.50
True SAA				0.0 (0.0%)	0.0 (0.1%)	0.0 (0.1%)	0.0 (0.1%)	0.0 (0.1%)
RCN	4.5 (3.0%)	5.4 (5.0%)	11 (26%)	0.2 (0.4%)	0.1 (0.2%)	0.1 (0.2%)	0.1 (0.1%)	0.1 (0.2%)
Censored SAA	184 (120%)	74 (68%)	8.8 (21%)	0.1 (0.2%)	0.1 (0.1%)	0.1 (0.2%)	0.1 (0.1%)	0.0 (0.1%)
KM	184 (120%)	74 (68%)	8.8 (21%)	0.0 (0.1%)	0.0 (0.1%)	0.1 (0.1%)	0.0 (0.1%)	0.0 (0.1%)
Naive SAA	184 (120%)	80 (74%)	21 (49%)	12 (31%)	9.0 (24%)	7.3 (19%)	5.6 (15%)	4.8 (13%)
Subsample SAA	209 (137%)	97 (89%)	27 (63%)	10.0 (27%)	3.9 (10%)	2.7 (7.2%)	2.3 (6.2%)	2.3 (6.2%)

(a) Uniform, $q_G^\star = 74$

λ	92.07	118.38	144.68	170.99	197.30	223.60	249.91	276.22
True SAA		0.1 (0.1%)	0.1 (0.1%)	0.1 (0.1%)	0.1 (0.1%)	0.1 (0.1%)	0.1 (0.1%)	0.1 (0.1%)
RCN	11 (23%)	0.4 (0.3%)	0.7 (0.7%)	0.3 (0.3%)	0.2 (0.2%)	0.3 (0.3%)	0.2 (0.2%)	0.2 (0.2%)
Censored SAA	13 (27%)	0.2 (0.2%)	0.2 (0.2%)	0.2 (0.2%)	0.2 (0.2%)	0.1 (0.1%)	0.2 (0.2%)	0.1 (0.1%)
Naive SAA	34 (70%)	19 (17%)	13 (12%)	7.6 (6.9%)	7.5 (6.7%)	5.5 (5.0%)	3.3 (2.9%)	1.4 (1.3%)
KM	13 (27%)	0.2 (0.2%)	0.2 (0.2%)	0.2 (0.2%)	0.2 (0.1%)	0.2 (0.1%)	0.2 (0.1%)	0.1 (0.1%)
Subsample SAA	46 (95%)	26 (23%)	18 (17%)	13 (12%)	11 (9.7%)	8.3 (7.5%)	6.4 (5.7%)	4.5 (4.1%)

(b) Exponential, $q_G^\star = 110.90$

λ	46	59.14	72.29	85.43	98.57	111.71	124.86	138
True SAA					0.0 (0.1%)	0.0 (0.2%)	0.0 (0.1%)	0.0 (0.2%)
RCN	0.0 (0.0%)	0.5 (0.2%)	2.4 (1.4%)	1.0 (6.8%)	0.0 (0.2%)	0.0 (0.3%)	0.0 (0.3%)	0.0 (0.3%)
Censored SAA	628 (300%)	590 (297%)	380 (219%)	0.6 (4.4%)	0.0 (0.2%)	0.0 (0.3%)	0.0 (0.2%)	0.0 (0.2%)
KM	628 (300%)	590 (297%)	380 (219%)	0.6 (4.4%)	0.0 (0.2%)	0.0 (0.3%)	0.0 (0.2%)	0.0 (0.3%)
Naive SAA	628 (300%)	590 (297%)	380 (219%)	3.6 (26%)	3.0 (26%)	2.8 (25%)	2.2 (19%)	2.1 (18%)
Subsample SAA	628 (300%)	593 (298%)	383 (221%)	2.6 (18%)	0.1 (0.5%)	0.1 (1.1%)	0.2 (1.6%)	0.2 (1.5%)

(c) Poisson, $q_G^\star = 86$

Table EC.1 Impact of λ on policy performance. Here, $\rho = 0.75$. Values to the left of the thick vertical line correspond to the unidentifiable regime, where we report Regret$(q^\pi) - \Delta$ and $\mathcal{R}^{ui}(q^\pi)\%$; values to the right of the thick vertical line correspond to the identifiable regime, where we report Vanilla-Regret(q^π) and $\mathcal{R}^{id}(q^\pi)\%$.

Appendix D: Computational Experiments: Additional Details

D.1. Simulation Information

Computing Infrastructure: The experiments were conducted on a personal computer with an Apple M2, 8-core processor and 16.0GB of RAM.

D.2. Synthetic Experiments: Supplemental Results

D.2.1. Impact of the Observable Boundary. Tables EC.1 and EC.2 show the impact of λ on policy performance, for $\rho \in \{0.75, 0.98\}$. Trends are identical to those identified for $\rho = 0.9$. Note that, for $\rho = 0.98$, all non-robust benchmarks exhibit extremely high relative regret for small values of λ, as the worst-case distribution forces these policies to incur exorbitant underage costs in the unidentifiable regime.

λ	44.50	57.21	69.93	82.64	95.36	108.07	120.79	133.50
True SAA						0.0 (0.0%)	0.0 (0.0%)	0.0 (0.1%)
RCN	4.7 (1.7%)	6.9 (2.7%)	6.9 (2.9%)	11 (5.4%)	115 (100%)	0.0 (0.1%)	0.1 (0.1%)	0.1 (0.3%)
Censored SAA	7165 (2651%)	5102 (2001%)	3335 (1401%)	1605 (751%)	116 (101%)	1.2 (2.4%)	0.9 (1.9%)	2.5 (5.1%)
KM	7165 (2651%)	5102 (2001%)	3335 (1401%)	1605 (751%)	116 (101%)	0.0 (0.1%)	0.0 (0.1%)	0.1 (0.2%)
Naive SAA	7165 (2651%)	5102 (2001%)	3335 (1401%)	1605 (751%)	116 (101%)	1.2 (2.4%)	1.3 (2.7%)	1.3 (2.7%)
Subsample SAA	7207 (2666%)	5135 (2014%)	3377 (1419%)	1629 (762%)	121 (105%)	0.4 (0.9%)	0.7 (1.3%)	0.7 (1.5%)

(a) Uniform, $q_G^\star = 97$

λ	92.07	118.38	144.68	170.99	197.30	223.60	249.91	276.22
True SAA								
RCN	6.2 (2.8%)	7.6 (4.0%)	7.6 (4.8%)	7.7 (6.0%)	69 (71%)	111 (163%)	48 (118%)	11 (59%)
Censored SAA	3232 (1482%)	1956 (1038%)	1138 (720%)	626 (490%)	316 (324%)	140 (206%)	51 (124%)	13 (71%)
Naive SAA	3232 (1482%)	1956 (1038%)	1138 (720%)	626 (490%)	316 (324%)	140 (206%)	55 (136%)	30 (169%)
KM	3232 (1482%)	1956 (1038%)	1138 (720%)	626 (490%)	316 (324%)	140 (206%)	49 (120%)	11 (61%)
Subsample SAA	3326 (1525%)	2053 (1090%)	1246 (788%)	737 (576%)	424 (435%)	246 (362%)	144 (352%)	101 (564%)

(b) Exponential, $q_G^\star = 312.96$

λ	46	59.14	72.29	85.43	98.57	111.71	124.86	138
True SAA						0.1 (0.6%)	0.2 (0.7%)	0.1 (0.5%)
RCN	0.0 (0.0%)	0.5 (0.2%)	2.5 (1.0%)	8.1 (3.7%)	3.1 (16%)	0.2 (1.0%)	0.2 (1.1%)	0.2 (1.0%)
Censored SAA	13397 (4900%)	12652 (4857%)	9578 (3887%)	2717 (1226%)	3.3 (16%)	0.6 (2.5%)	0.9 (4.2%)	1.9 (8.4%)
KM	13397 (4900%)	12652 (4857%)	9578 (3887%)	2717 (1226%)	3.2 (16%)	0.2 (1.0%)	0.2 (1.1%)	0.2 (0.9%)
Naive SAA	13397 (4900%)	12652 (4857%)	9578 (3887%)	2717 (1226%)	4.8 (24%)	1.7 (7.8%)	1.7 (7.5%)	1.5 (6.7%)
Subsample SAA	13397 (4900%)	12680 (4868%)	9589 (3892%)	2722 (1229%)	5.2 (26%)	0.3 (1.4%)	0.4 (1.7%)	0.5 (2.4%)

(c) Poisson, $q_G^\star = 99$

Table EC.2 Impact of λ on policy performance. Here, $\rho = 0.98$. Values to the left of the thick vertical line correspond to the unidentifiable regime, where we report Regret$(q^\pi) - \Delta$ and $\mathcal{R}^{ui}(q^\pi)$%; values to the right of the thick vertical line correspond to the identifiable regime, where we report Vanilla-Regret(q^π) and $\mathcal{R}^{id}(q^\pi)$%.

λ	59.23					177.69				
σ	20	25	30	35	40	20	25	30	35	40
True SAA						0.1 (0.2%)	0.1 (0.1%)	0.1 (0.1%)	0.1 (0.2%)	0.1 (0.2%)
RCN	2.1 (0.9%)	2.6 (1.1%)	3.1 (1.4%)	3.2 (1.4%)	3.0 (1.3%)	0.1 (0.4%)	0.1 (0.3%)	0.2 (0.3%)	0.2 (0.4%)	0.2 (0.3%)
Censored SAA	1760 (751%)	1620 (697%)	1512 (656%)	1428 (623%)	1363 (598%)	0.1 (0.3%)	0.1 (0.3%)	0.3 (0.5%)	0.3 (0.5%)	0.3 (0.4%)
KM	1760 (751%)	1620 (697%)	1512 (656%)	1428 (623%)	1363 (598%)	0.1 (0.3%)	0.1 (0.2%)	0.2 (0.3%)	0.2 (0.3%)	0.2 (0.2%)
Naive SAA	1760 (751%)	1620 (697%)	1512 (656%)	1428 (623%)	1363 (598%)	2.6 (7.4%)	3.8 (8.5%)	4.3 (8.1%)	5.4 (8.9%)	7.5 (11%)
Subsample SAA	1773 (756%)	1636 (704%)	1533 (665%)	1453 (635%)	1393 (612%)	1.0 (2.8%)	1.4 (3.2%)	1.4 (2.7%)	2.5 (4.1%)	3.3 (4.7%)

Table EC.3 Impact of σ on policy performance. Here, $\rho = 0.9$. At $\lambda = 59.23$, the problem is unidentifiable for all values of σ; in this case we report Regret$(q^\pi) - \Delta$ and $\mathcal{R}^{ui}(q^\pi)$%. At $\lambda = 177.69$, the problem is identifiable for all values of σ; in this case we report Vanilla-Regret(q^π) and $\mathcal{R}^{id}(q^\pi)$%.

D.2.2. Impact of Demand Variability. In Table EC.3 we illustrate the impact of demand variability for additional values of λ, fixing $\rho = 0.9$. (Results are identical for $\rho \in \{0.75, 0.98\}$.)

When $\lambda = 59.23$, the problem is easily unidentifiable across all values of σ. Our algorithm achieves less than 1.5% relative regret in this regime, with all other algorithms achieving over 598% relative regret. When $\lambda = 177.69$, on the other hand, the problem is easily identifiable for all values of σ, with RCN achieving a relative regret of less than 0.5% across all such values.

D.2.3. Vanilla Regret Comparison in the "Clearly Unidentifiable" Regime. While RCN achieves the optimal worst-case regret guarantee in the unidentifiable regime, a natural question is whether its strong performance extends to *vanilla regret* (see Eq. (3)) in this regime, evaluated with respect to the true distribution G. In this section we provide numerical support for the idea that vanilla regret is not the meaningful regret metric in the unidentifiable regime. In particular, our results will illustrate the intuitive fact that the performance of all non-robust algorithms is poor in the unidentifiable regime, but improves as λ increases. By contrast, the performance of RCN is highly dependent on the gap between $q^\star_G$ and $q^\dagger_G$, which in turn depends on λ and the tightness of the upper bound M.

Setup. We consider the uniform, exponential, and Poisson distributions described in Section 6.2. Across all experiments, we let $h = 1$ and $b = 9$ (i.e., $\rho = 0.9$). We moreover let $M = 100$ in the case of the uniform distribution, and $M = 325$ for the exponential and Poisson distributions. The sampling of s^{off} and the order levels is identical to that of Section 6.2.1, where we set $N = 500$. We report the vanilla regret of the five policies of interest in the "clearly unidentifiable" regime, where $\lambda \ll q^\star_G$, in Table EC.4. We study the knife-edge regime, where $\lambda \approx q^\star_G$ in Appendix D.2.4. Finally, we omit the "clearly identifiable" regime, where $\lambda \gg q^\star_G$, as this regime is already covered by Table 1.

Results. For the uniform distribution (Table EC.4a), it is easy to show that $q^\star_G = q^\dagger_G$ across all values of λ. Since RCN outputs an empirical estimate of $q^\dagger_G$ in this setting, it achieves near-zero vanilla regret across all values of λ. In the worst case, its additive vanilla regret is 0.9 (corresponding to a relative vanilla regret of 2.1%) when $\lambda = 82.64$. In contrast, the other policies perform very poorly for small values of λ, with a relative vanilla regret of 300% in the worst case, achieved by Subsample SAA. Since Censored SAA, KM and Naive SAA output λ in almost all replications in the unidentifiable regime, their performance improves as λ increases, thereby approaching $q^\star_G$. Similarly, the performance of Subsample SAA improves, since the fraction of samples for which $d_{ki} < q^{\text{off}}_k$ increases with λ.

Under the exponential distribution (Table EC.4b), RCN continues to perform well in the unidentifiable regime, despite the fact that $q^\star_G \neq q^\dagger_G$ ($q^\star_G = 184.21$ and $q^\dagger_G \in \{251.32, 234.17, 214.81, 194.24\}$ for increasing unidentifiable values of λ). Its cost exceeds the optimal newsvendor cost by 12% in the worst case, with improved performance as λ increases and $q^\dagger_G$ approaches $q^\star_G$. By contrast, the remaining policies exhibit significantly higher regret (up to 79% for Subsample SAA), as for the uniform distribution.

Under the Poisson distribution (Table EC.4c), the difference between $q^\star_G$ and $q^\dagger_G$ is significantly higher than in the previous two cases ($q^\star_G = 92$ versus $q^\dagger_G \in \{297.10, 298.14, 293.32, 234.82\}$ for increasing unidentifiable λ). The gap between $q^\dagger_G$ and $q^\star_G$ — explained by the very loose upper bound M on $q^\star_G$ — results in the poor performance of RCN; in the best case ($\lambda = 85.43$), its cost is over 9 times the optimal newsvendor cost. Finally, we note that when $\lambda = 46$, our algorithm — which outputs an estimate of $q^\dagger_G$ — outperforms all baselines. This is despite the fact that $q^\dagger_G$ is over three times as high as $q^\star_G$ for this value of λ, and can be explained by the fact that $\rho = 0.9$, meaning that lost sales are far more heavily penalized than overage.

Consequently, over-ordering relative to $q_G^\star$ mitigates expected cost more effectively than under-ordering by outputting λ, which the baselines do. As λ increases and approaches $q_G^\star$, this advantage diminishes, since λ approaches $q_G^\star$ more rapidly than $q_G^\dagger$.

λ	44.50	57.21	69.93	82.64
True SAA	0.1 (0.2%)	0.1 (0.2%)	0.1 (0.2%)	0.1 (0.2%)
RCN	0.0 (0.1%)	0.1 (0.1%)	0.1 (0.1%)	0.9 (2.1%)
Censored SAA	101 (225%)	52 (115%)	19 (41%)	2.1 (4.7%)
KM	101 (225%)	52 (115%)	19 (41%)	2.1 (4.7%)
Naive SAA	101 (225%)	52 (115%)	19 (41%)	5.2 (12%)
Subsample SAA	135 (300%)	85 (188%)	45 (100%)	19 (42%)

(a) Uniform, $q_G^\star = 89$

λ	92.07	118.38	144.68	170.99
True SAA	0.6 (0.3%)	0.6 (0.3%)	0.5 (0.3%)	0.5 (0.3%)
RCN	21 (12%)	13 (7.0%)	9.6 (5.2%)	1.3 (0.7%)
Censored SAA	81 (44%)	37 (20%)	12 (6.4%)	1.6 (0.9%)
KM	81 (44%)	37 (20%)	12 (6.4%)	1.5 (0.8%)
Naive SAA	81 (44%)	37 (20%)	24 (13%)	26 (14%)
Subsample SAA	146 (79%)	103 (56%)	74 (40%)	54 (29%)

(b) Exponential, $q_G^\star = 184.21$

λ	46	59.14	72.29	85.43
True SAA	0.0 (0.2%)	0.0 (0.1%)	0.0 (0.2%)	0.0 (0.2%)
RCN	201 (1251%)	202 (1258%)	197 (1228%)	138 (857%)
Censored SAA	290 (1805%)	172 (1069%)	63 (389%)	4.6 (29%)
KM	290 (1805%)	172 (1069%)	63 (389%)	4.6 (29%)
Naive SAA	290 (1805%)	172 (1069%)	63 (389%)	4.6 (29%)
Subsample SAA	290 (1806%)	179 (1114%)	65 (405%)	8.5 (53%)

(c) Poisson, $q_G^\star = 92$

Table EC.4 Comparison of policies' vanilla regret in the "clearly unidentifiable" regime. Here, $\rho = 0.9$.

D.2.4. Algorithm Performance in the "Knife-Edge" Regime. We next closely study the impact of N and λ on the performance of RCN, Censored SAA, and KM in the knife-edge regime, where λ lies near $q^\star_G$. (We omit Subsample SAA and Naive SAA, given their clear suboptimality.) We consider four equally spaced values of $\lambda \in [0.95q^\star_G, 1.05q^\star_G]$. For each value of λ, we report algorithm performance for $N \in \{100, 250, 500\}$.

Comparison of Worst-Case Regret. We first compare the worst-case regret of our three policies in Table EC.5. Recall, for $\lambda < q^\star_G$, we report Regret$(q^\pi) - \Delta$, while in the identifiable regime we report Vanilla-Regret(q^π).

Impact of λ. We first fix $N = 500$ and compare algorithm performance as λ increases. In the case of the uniform and exponential distributions (Tables EC.5a and EC.5b), we observe that RCN slightly outperforms KM and Censored SAA for $\lambda < q^\star_G$, though the difference in regret is quite small. This is because, for all of these values of λ, RCN classifies the problem as knife-edge (thereby outputting λ) in the vast majority of replications. KM and Censored SAA similarly output λ in the majority of replications. RCN obtains a slight edge due to the minority of replications in which it correctly classifies the problem as unidentifiable, in which case it outputs an estimate of $q^\dagger_G$. (For instance, this occurs in 30% of replications when $\lambda = 84.55$ in the case of the uniform distribution.) We moreover observe an improvement in all three policies as they approach the boundary. This is because the distance between $q^\dagger_G$ and λ is decreasing in λ when $G^-(\lambda) < \rho$; hence outputting λ in the majority of replications is less costly. While we observe an improvement in all policies as λ increases in the case of the Poisson distribution as well (Table EC.5c), for $\lambda = 87.40$ we remark that the regret of RCN is less than one-tenth that of KM and Censored SAA. This follows from the fact that our algorithm correctly classifies the problem as unidentifiable in 99% of replications.

To the right of the observable boundary, while KM and Censored SAA are able to correctly output an estimate of $q^\star_G$, RCN continues to output λ in a majority of replications, for both the uniform and exponential distributions. However, given that λ is close to $q^\star_G$ in this regime, it still exhibits strong performance, with a slight decrease as λ increases, and the inability to classify the instance becomes more costly as a result. The Poisson distribution exhibits slightly different behavior as λ increases. In this case, when $\lambda = 96.60$ the algorithm correctly classifies the problem as identifiable in 100% of replications, resulting in a *decrease* in vanilla regret as λ increases.

Impact of N. For fixed λ, the worst-case regret of RCN weakly decreases in N in this regime. In the case of the uniform and exponential distributions, this is driven by the improvement in classification rates when $\lambda < q^\star_G$. For instance, under the uniform distribution, when $\lambda = 84.55$, the fraction of replications in which RCN correctly classifies the problem as unidentifiable increases from 4% when $N = 250$ to 30% when $N = 500$. The improvement is less pronounced when $\lambda > q^\star_G$; the reason for this is that an increase in N does not have a significant impact on the classification rate of RCN in this regime. These results contrast with the Poisson distribution, where, at both "knife-edge extremes" of $\lambda \in \{87.4, 96.6\}$, an increase in N

has a pronounced impact on the classification rate. In particular, for $\lambda = 87.4$, RCN goes from correctly classifying the problem as unidentifiable in 18% of replications when $N = 100$, to 99% of replications when $N = 500$. Similarly, when $\lambda = 96.6$ the classification rate improves from 0% to 100%; here, however, there is no correspondingly large decrease in vanilla regret, since outputting λ already achieves relatively low regret. The improvement in classification rate is not as significant for $\lambda \in \{90.35, 93.53\}$; here, we conjecture that we would have to increase N significantly (e.g., $N = 1000$) in order to observe a decrease in regret. Finally, we briefly remark on the impact of N on the performance of KM and Censored SAA when $\lambda < q^\star_G$. (When $\lambda > q^\star_G$, the fact that an increase in sample size improves vanilla regret is an expected and well-studied phenomenon.) This improvement follows from the fact that, as N increases, the fraction of replications in which both algorithms output λ increases (as $\widehat{G}^-(\lambda) > \rho$ becomes an increasingly rare event). Since $q^\dagger_G > \lambda$ in this regime, outputting λ more frequently results in a performance improvement.

In summary, these results show that, while the knife-edge regime is the most difficult in terms of classification, all algorithms exhibit strong performance in the majority of tested instances. The reason for this is that these algorithms frequently output λ in this setting (albeit for different reasons); by definition of the knife-edge regime, this quantity is guaranteed to be close to the optimum (be it $q^\dagger_G$ or $q^\star_G$). Our results for the Poisson distribution, on the other hand, establish that the strong performance of KM and Censored SAA in this regime is not robust to the true underlying demand distribution.

Comparison of Vanilla Regret. We next compare the vanilla regret of all algorithms in Table EC.6. The results are intuitive, given the above discussion. (We do not discuss results for $\lambda > q^\star_G$ below, given that they are identical to those reported in Table EC.5.)

Fixing $N = 500$, all algorithms output λ in the majority of replications when $\lambda < q^\star_G$ (except in the case of the Poisson distribution, for $\lambda = 87.4$). Since λ is close to $q^\star_G$ in these settings, all algorithms exhibit low vanilla regret; this performance only improves as λ increases and approaches $q^\star_G$. The exception to this is $\lambda = 87.4$ under the Poisson distribution where, as discussed above, RCN correctly classifies the problem as unidentifiable in the majority of replications. For this specific instance, $q^\dagger_G \approx 203$; hence, $|q^\dagger_G - q^\star_G| \gg |q^\star_G - \lambda|$, which explains why the vanilla regret of RCN is higher than that of KM and Censored SAA.

Finally, the vanilla regret of RCN weakly decreases as N increases (again, except for the isolated instance under the Poisson distribution). This result is slightly unintuitive, since RCN outputs an estimate of $q^\dagger_G$ *more* frequently as N increases. However, for these instances, $q^\dagger_G$ is actually *closer* to $q^\star_G$ than λ is; hence, worst-case and vanilla regret of RCN are directionally aligned. This is not the case for $\lambda = 87.4$ under the Poisson distribution, where $q^\dagger_G \gg q^\star_G$, as discussed above.

Algorithm	N	84.55	87.51	90.48	93.45
RCN	100	2.6 (50%)	0.4 (19%)	0.0 (0.1%)	0.8 (1.7%)
	250	2.6 (50%)	0.4 (19%)	0.0 (0.1%)	0.8 (1.7%)
	500	2.1 (42%)	0.4 (19%)	0.0 (0.1%)	0.7 (1.6%)
KM	100	2.7 (52%)	0.6 (27%)	0.3 (0.7%)	0.5 (1.0%)
	250	2.6 (50%)	0.5 (22%)	0.1 (0.2%)	0.2 (0.3%)
	500	2.6 (50%)	0.4 (20%)	0.0 (0.1%)	0.1 (0.1%)
Censored SAA	100	2.7 (52%)	0.7 (31%)	0.5 (1.1%)	0.6 (1.2%)
	250	2.6 (50%)	0.5 (22%)	0.1 (0.2%)	0.2 (0.4%)
	500	2.6 (50%)	0.4 (20%)	0.1 (0.1%)	0.1 (0.1%)

(a) Uniform, $q^\star_G = 89$

Algorithm	N	174.96	181.10	187.24	193.38
RCN	100	2.8 (17%)	1.1 (20%)	0.0 (0.0%)	0.4 (0.2%)
	250	2.1 (13%)	0.3 (5.5%)	0.0 (0.0%)	0.4 (0.2%)
	500	2.1 (13%)	0.3 (5.5%)	0.0 (0.0%)	0.4 (0.2%)
KM	100	3.2 (20%)	1.9 (35%)	1.2 (0.7%)	2.1 (1.1%)
	250	2.4 (15%)	0.8 (15%)	0.5 (0.3%)	0.6 (0.4%)
	500	2.3 (14%)	0.6 (12%)	0.2 (0.1%)	0.3 (0.2%)
Censored SAA	100	3.8 (23%)	3.1 (59%)	1.9 (1.0%)	2.8 (1.5%)
	250	2.7 (16%)	1.0 (19%)	0.8 (0.4%)	0.8 (0.4%)
	500	2.3 (14%)	0.7 (13%)	0.3 (0.2%)	0.4 (0.2%)

(b) Exponential, $q^\star_G = 184.21$

Algorithm	N	87.40	90.46	93.53	96.60
RCN	100	101 (86%)	9.6 (23%)	0.3 (2.1%)	1.9 (12%)
	250	30 (26%)	8.7 (21%)	0.3 (2.1%)	0.7 (4.5%)
	500	9.7 (8.2%)	9.6 (23%)	0.3 (2.0%)	0.0 (0.3%)
KM	100	117 (99%)	8.8 (21%)	0.2 (1.1%)	0.3 (1.7%)
	250	117 (99%)	8.8 (21%)	0.1 (0.6%)	0.1 (0.5%)
	500	117 (99%)	8.7 (21%)	0.0 (0.3%)	0.0 (0.3%)
Censored SAA	100	117 (99%)	8.8 (21%)	0.2 (1.1%)	0.3 (1.8%)
	250	117 (99%)	8.8 (21%)	0.1 (0.6%)	0.1 (0.5%)
	500	117 (99%)	8.7 (21%)	0.1 (0.3%)	0.0 (0.3%)

(c) Poisson, $q^\star_G = 92$

Table EC.5 Impact of λ and N on policy performance in the knife-edge regime. Here, $\rho = 0.9$. Values to the left of the thick vertical line correspond to the unidentifiable regime, where we report Regret$(q^\pi) - \Delta$ and $\mathcal{R}^{ui}(q^\pi)$%; values to the right of the thick vertical line correspond to the identifiable regime, where we report Vanilla-Regret(q^π) and $\mathcal{R}^{id}(q^\pi)$%.

Algorithm	N	84.55	87.51	90.48	93.45
RCN	100	1.2 (2.7%)	0.2 (0.4%)	0.0 (0.1%)	0.8 (1.7%)
	250	1.2 (2.6%)	0.2 (0.4%)	0.0 (0.1%)	0.8 (1.7%)
	500	0.9 (1.9%)	0.2 (0.4%)	0.0 (0.1%)	0.7 (1.6%)
KM	100	1.3 (2.9%)	0.4 (0.8%)	0.3 (0.7%)	0.5 (1.0%)
	250	1.2 (2.7%)	0.2 (0.5%)	0.1 (0.2%)	0.2 (0.3%)
	500	1.2 (2.7%)	0.2 (0.5%)	0.0 (0.1%)	0.1 (0.1%)
Censored SAA	100	1.3 (3.0%)	0.4 (1.0%)	0.5 (1.1%)	0.6 (1.2%)
	250	1.2 (2.7%)	0.3 (0.6%)	0.1 (0.2%)	0.2 (0.4%)
	500	1.2 (2.7%)	0.2 (0.5%)	0.1 (0.1%)	0.1 (0.1%)

(a) Uniform, $q_G^\star = 89$

Algorithm	N	174.96	181.10	187.24	193.38
RCN	100	0.7 (0.4%)	0.2 (0.1%)	0.0 (0.0%)	0.4 (0.2%)
	250	0.4 (0.2%)	0.0 (0.0%)	0.0 (0.0%)	0.4 (0.2%)
	500	0.4 (0.2%)	0.0 (0.0%)	0.0 (0.0%)	0.4 (0.2%)
KM	100	1.5 (0.8%)	1.5 (0.8%)	1.2 (0.7%)	2.1 (1.1%)
	250	0.7 (0.4%)	0.4 (0.2%)	0.5 (0.3%)	0.6 (0.4%)
	500	0.6 (0.3%)	0.2 (0.1%)	0.2 (0.1%)	0.3 (0.2%)
Censored SAA	100	2.1 (1.1%)	2.7 (1.5%)	1.9 (1.0%)	2.8 (1.5%)
	250	0.9 (0.5%)	0.6 (0.3%)	0.8 (0.4%)	0.8 (0.4%)
	500	0.6 (0.3%)	0.3 (0.2%)	0.3 (0.2%)	0.4 (0.2%)

(b) Exponential, $q_G^\star = 184.14$

Algorithm	N	87.40	90.46	93.53	96.60
RCN	100	26 (164%)	1.4 (8.9%)	0.3 (2.1%)	1.9 (12%)
	250	93 (576%)	0.1 (0.8%)	0.3 (2.1%)	0.7 (4.5%)
	500	108 (675%)	1.9 (12%)	0.3 (2.0%)	0.0 (0.3%)
KM	100	2.0 (12%)	0.2 (1.3%)	0.2 (1.1%)	0.3 (1.7%)
	250	2.0 (12%)	0.2 (0.9%)	0.1 (0.6%)	0.1 (0.5%)
	500	2.0 (12%)	0.1 (0.8%)	0.0 (0.3%)	0.0 (0.3%)
Censored SAA	100	2.0 (12%)	0.2 (1.4%)	0.2 (1.1%)	0.3 (1.8%)
	250	2.0 (12%)	0.2 (0.9%)	0.1 (0.6%)	0.1 (0.5%)
	500	2.0 (12%)	0.1 (0.8%)	0.1 (0.3%)	0.0 (0.3%)

(c) Poisson, $q_G^\star = 92$

Table EC.6 Impact of λ and N on policy performance for vanilla regret in the knife-edge regime. Here, $\rho = 0.9$. Across all values of λ we report Vanilla-Regret(q^π) and $\mathcal{R}^{id}(q^\pi)$%.

λ	44.50	57.21	69.93	82.64	95.36	108.07	120.79	133.50
RCN	4.0 (1.7%)	6.7 (3.3%)	7.7 (4.5%)	27 (27%)	0.1 (0.2%)	0.1 (0.2%)	0.1 (0.1%)	0.1 (0.2%)
KM	1033 (450%)	653 (320%)	341 (200%)	70 (70%)	0.1 (0.1%)	0.1 (0.1%)	0.0 (0.1%)	0.1 (0.1%)
Robust KM	4.0 (1.7%)	6.7 (3.3%)	7.7 (4.5%)	27 (27%)	0.1 (0.2%)	0.1 (0.1%)	0.0 (0.1%)	0.1 (0.1%)

(a) Uniform, $q^\star_G = 89$

λ	92.07	118.38	144.68	170.99	197.30	223.60	249.91	276.22
RCN	6.7 (4.2%)	6.9 (6.0%)	21 (29%)	4.3 (19%)	1.0 (0.5%)	7.4 (4.0%)	3.9 (2.1%)	0.8 (0.5%)
KM	345 (216%)	148 (128%)	45 (64%)	4.5 (19%)	0.6 (0.3%)	0.6 (0.3%)	0.5 (0.3%)	0.5 (0.3%)
Robust KM	6.7 (4.2%)	6.9 (6.0%)	21 (29%)	4.3 (19%)	1.0 (0.5%)	6.8 (3.7%)	2.9 (1.6%)	0.5 (0.5%)

(b) Exponential, $q^\star_G = 184.21$

λ	46	59.14	72.29	85.43	98.57	111.71	124.86	138
RCN	0.0 (0.0%)	0.5 (0.2%)	2.4 (1.1%)	7.1 (4.7%)	0.1 (0.3%)	0.0 (0.3%)	0.1 (0.3%)	0.1 (0.4%)
KM	2260 (900%)	2131 (891%)	1542 (697%)	247 (165%)	0.1 (0.3%)	0.0 (0.3%)	0.1 (0.3%)	0.1 (0.4%)
Robust KM	0.0 (0.0%)	0.5 (0.2%)	2.4 (1.1%)	7.1 (4.7%)	0.1 (0.3%)	0.0 (0.3%)	0.1 (0.3%)	0.1 (0.4%)

(c) Poisson, $q^\star_G = 92$

Table EC.7 Comparison of RCN, KM, and Robust KM algorithms for $\rho = 0.9$. Values to the left of the thick vertical line correspond to the unidentifiable regime, where we report Regret$(q^\pi) - \Delta$ and $\mathcal{R}^{ui}(q^\pi)$%; values to the right of the thick vertical line correspond to the identifiable regime, where we report Vanilla-Regret(q^π) and $\mathcal{R}^{id}(q^\pi)$%.

D.2.5. Robustifying the Kaplan-Meier Estimator In this section we demonstrate the flexibility of our robust framework by implementing Robust KM, a "robust" variant of the KM algorithm. Robust KM is identical to RCN in all regimes except the "likely identifiable" regime, where it outputs the ρ-quantile of the KM estimator, instead of the SAA estimate used by RCN.

We compare the performance of KM, Robust KM, and RCN in Table EC.7 for various values of λ, as described in Section 6.2.1, and letting $\rho = 0.9$ and $N = 500$. For the five smallest values of λ, Robust KM and RCN coincide across all demand families, far outperforming KM. This behavior is expected, since both algorithms output the same classification (either "knife-edge" or "likely unidentifiable") in these regimes. Moreover, conditional on an instance being classified in either of these two ways, the algorithms output the same quantity. For larger values of λ, we observe that Robust KM slightly outperforms RCN, with the most pronounced improvement in the exponential case, for $\lambda = 223.60$ and $\lambda = 249.91$. The reason for this is that, for these values of λ, while the classification of the two algorithms may be the same, when the instance is classified as "likely identifiable," Robust KM outputs the ρ-quantile of the KM estimator, which slightly outperforms the censored SAA estimate in the identifiable regime. Finally, we remark that, in the case of the exponential distribution, KM outperforms Robust KM throughout the identifiable regime. This is the cost of robustness: in this regime, Robust KM makes classification errors in which it outputs λ instead of the KM estimator.

Overall, these results highlight the flexibility of our framework as a "robustification" tool to adapt optimal algorithms for the uncensored setting to the censored setting.

Appendix E: On The Value of Information: The Case of Globally Well-Separated Distributions

In this section we study how to leverage additional distributional information to improve decision-making in the face of censored data. We do so by considering the setting where the demand distribution is *globally well-separated*, and the decision-maker knows a lower bound $\gamma > 0$ on its pdf. For this setting, our goal is to quantify (i) the *value of information*, measured as the reduction in minimax risk when additional well-separatedness information is available, and (ii) the resulting *sample complexity improvements* of a modified version of RCN that explicitly exploits this information to detect whether the problem instance is identifiable or not. In the remainder of this section, we let Δ^{WS} and $q^{\Delta^{\mathsf{WS}}}$ respectively denote the minimax risk and minimax optimal ordering quantity under the well-separatedness condition.

E.1. Characterizing the Minimax Risk

We first characterize Δ^{WS} and $q^{\Delta^{\mathsf{WS}}}$ in this setting. Under this additional distributional information, the ambiguity set is given by:

$$\mathcal{F}(\lambda; G) = \left\{F \in \mathcal{G} \mid F(x) = G(x)\ \forall\, x < \lambda,\ f(x) \geq \gamma \text{ on its support}, q_F^\star \leq M\right\}. \tag{EC.46}$$

The following lemma will be useful in our derivation of $q^{\Delta^{\mathsf{WS}}}$. We defer its proof to Appendix E.1.1.

LEMMA EC.11. *For all $F \in \mathcal{F}(\lambda; G)$, F is a distribution with bounded support, with $f(x) = 0$ for all $x > \frac{1}{\gamma}$. Moreover, if $G^-(\lambda) \leq \rho$, $q_F^\star \leq \lambda + \frac{\rho - G^-(\lambda)}{\gamma}$.*

Noting that $G^-(\lambda) \geq \lambda\gamma \implies \lambda + \frac{\rho - G^-(\lambda)}{\gamma} \leq \frac{\rho}{\gamma} \leq \frac{1}{\gamma}$, we let $\widetilde{M} = \min\left\{M, \lambda + \frac{\rho - G^-(\lambda)}{\gamma}\right\}$ and modify the definition of $\mathcal{F}(\lambda; G)$ in (EC.46) to be such that $q_F^\star \leq \widetilde{M}$ when $G^-(\lambda) \leq \rho$.

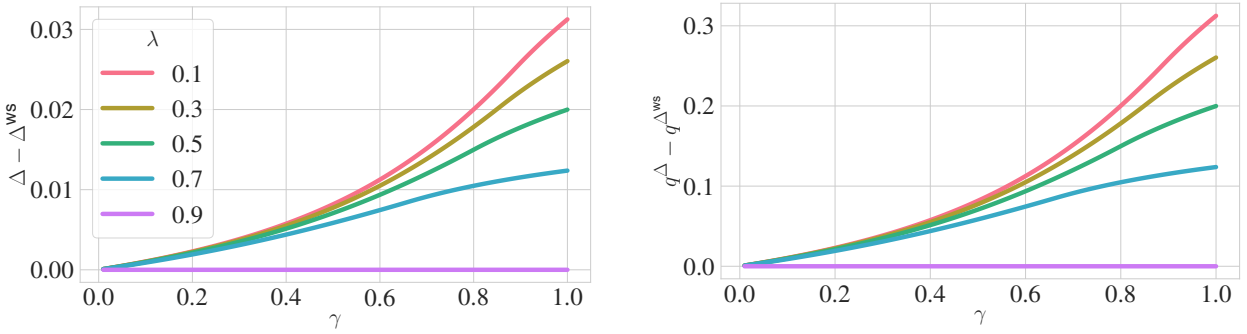
THEOREM EC.2. *Consider a data-driven censored newsvendor problem with well-separated demand distribution G and observable boundary λ. Then, given $\gamma > 0$:*

1. *If $G^-(\lambda) \geq \rho$, $q^{\Delta^{\mathsf{WS}}} = q_G^\star < \lambda$, and $\Delta^{\mathsf{WS}} = 0$.*
2. *If $G^-(\lambda) < \rho$,*

$$q^{\Delta^{\mathsf{WS}}} = \lambda + \frac{1 - G^-(\lambda) - \sqrt{(1 - G^-(\lambda))^2 + (\rho - G^-(\lambda) - \gamma(\widetilde{M} - \lambda))^2 - (\rho - G^-(\lambda))^2}}{\gamma} \geq \lambda.$$

Moreover,

$$\Delta^{\mathsf{WS}} = (b + h)(1 - \rho) \cdot \frac{1 - G^-(\lambda) - \sqrt{(1 - G^-(\lambda))^2 + (\rho - G^-(\lambda) - \gamma(\widetilde{M} - \lambda))^2 - (\rho - G^-(\lambda))^2}}{\gamma}.$$



(a) $\Delta - \Delta^{\mathsf{WS}}$ vs. γ (b) $q^{\Delta} - q^{\Delta^{\mathsf{WS}}}$ vs. γ

Figure EC.1 Dependence of $q^{\Delta} - q^{\Delta^{\mathsf{WS}}}$ and $\Delta - \Delta^{\mathsf{WS}}$ on γ, for $D \sim \mathbf{Unif}[0,1]$. Here, we assume $M = 1$, $b = 9$ and $h = 1$.

To gain further insight into the value of distributional information in this setting, in Figure EC.1a we plot $\Delta - \Delta^{\mathsf{WS}}$ as a function of γ, when $D \sim \text{Unif}[0,1]$ and $b = 9, h = 1$ so $\rho = 0.9$. We observe that $\Delta - \Delta^{\mathsf{WS}}$ is increasing in γ, for fixed λ. This phenomenon is intuitive, since the size of the ambiguity set decreases as γ increases; this supports the interpretation of γ as the quality of the decision-maker's additional information. We additionally observe that, for fixed γ, $\Delta - \Delta^{\mathsf{WS}}$ is concave decreasing in λ. This shows that the decision-maker has the most to gain from additional distributional information when the dataset is highly censored and therefore relatively uninformative.

Figure EC.1b further illustrates that γ not only affects the worst-case regret of the decision-maker, but also her minimax optimal ordering quantity. In particular, we observe that $q^{\Delta^{\mathsf{WS}}}$ is *lower* than q^{Δ} for all values of λ and γ. Moreover, for fixed λ, the difference between q^{Δ} and $q^{\Delta^{\mathsf{WS}}}$ is increasing in γ. The reason for this is also intuitive: as γ increases, the decision-maker is guaranteed that a larger amount of mass resides close to λ, thereby reducing the likelihood of large realizations of D and lowering the minimax optimal ordering quantity as a result.

Proof of Theorem EC.2. The proof of the first fact is identical to the more general setting we consider in Theorem 1; we omit it as such. In the remainder of the proof we focus on the case where $G^-(\lambda) < \rho$. The proof proceeds in the same fashion as that of Theorem 1. Namely, we first provide closed-form expressions for the worst-case regret over $\mathcal{F}(\lambda; G)$, given $q \le \widetilde{M}$.

LEMMA EC.12.

$$\text{Regret}(q) = \begin{cases} (b+h)\Big(\rho(\widetilde{M}-q) + \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G\big[(\widetilde{M}-D)\mathbb{1}\{D < \lambda\}\big] - \frac{\gamma}{2}(\widetilde{M}-\lambda)^2\Big) & \text{if } q < \lambda \\ (b+h)\max\big\{(1-\rho)(q-\lambda), (\widetilde{M}-q)\rho - G^-(\lambda) - \gamma(q-\lambda)\big) - \frac{\gamma}{2}(\widetilde{M}-q)^2\big\} & \text{if } q \ge \lambda. \end{cases}$$

We defer the proof of the lemma to Appendix E.1.2. We use these expressions to compute the minimax optimal ordering quantity $q_G^{\dagger}$ in this case. For all $q < \lambda$, we have:

$$\frac{d}{dq}\Big[\rho(\widetilde{M}-q) + \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G\big[(\widetilde{M}-D)\mathbb{1}\{D < \lambda\}\big] - \frac{\gamma}{2}(\widetilde{M}-\lambda)^2\Big] = -\rho + G(q) < 0,$$

since $G^-(\lambda) < \rho$. Therefore:

$$\begin{aligned}\inf_{q<\lambda} \text{Regret}(q) &= (b+h)\Big(\rho(\widetilde{M}-\lambda) + \mathbb{E}_G[(\lambda-D)\mathbb{1}\{D\le\lambda\}] - \mathbb{E}_G\big[(\widetilde{M}-D)\mathbb{1}\{D<\lambda\}\big] - \frac{\gamma}{2}(\widetilde{M}-\lambda)^2\Big)\\ &= (b+h)\Big(\rho(\widetilde{M}-\lambda) - (\widetilde{M}-\lambda)G^-(\lambda) - \frac{\gamma}{2}(\widetilde{M}-\lambda)^2\Big)\\ &= (b+h)\Big((\widetilde{M}-\lambda)(\rho-G^-(\lambda)) - \frac{\gamma}{2}(\widetilde{M}-\lambda)^2\Big) = \text{Regret}(\lambda).\end{aligned}$$

Thus we have, $q_G^\dagger = \arg\min_{q\ge\lambda} \max\big\{(1-\rho)(q-\lambda), (\widetilde{M}-q)(\rho-G^-(\lambda)-\gamma(q-\lambda)) - \frac{\gamma}{2}(\widetilde{M}-q)^2\big\}$. Equivalently, $q_G^\dagger$ is the smallest solution to:

$$\begin{aligned}&(1-\rho)(q-\lambda) = (\widetilde{M}-q)(\rho-G^-(\lambda)-\gamma(q-\lambda)) - \frac{\gamma}{2}(\widetilde{M}-q)^2\\ \iff & q_G^\dagger = \lambda + \frac{1-G^-(\lambda) - \sqrt{(1-G^-(\lambda))^2 + (\rho-G^-(\lambda)-\gamma(\widetilde{M}-\lambda))^2 - (\rho-G^-(\lambda))^2}}{\gamma} \ge \lambda,\end{aligned}$$

where the final solution simply follows from algebra; we omit it as such. □

E.1.1. Proof of Lemma EC.11

Proof. To see that $f(x) = 0$ for all $x > \frac{1}{\gamma}$, note that $F(x) \ge \gamma x$ for all x in the support of F. Therefore, $F(x) = 1$ for some $x \le \frac{1}{\gamma}$, which proves the claim.

For the lower bound on $q_F^\star$, since $q_F^\star \ge \lambda$ whenever $G^-(\lambda) \le \rho$, we have:

$$q_F^\star = \inf\{q>0 : F(q) \ge \rho\} \le \inf\{q>0 : G^-(\lambda) + \gamma(q-\lambda) \ge \rho\}.$$

Rearranging the above inequality yields the result.

□

E.1.2. Proof of Lemma EC.12

Proof. As in the proof of Lemma 2 (see Eq. (EC.3)), it suffices to analyze:

$$\begin{aligned}&\sup_{F\in\mathcal{F}(\lambda;G)} C_F(q) - C_F(q_F^\star)\\ &= \sup_{\tilde{q}\in[\lambda,\widetilde{M}]} b(\tilde{q}-q) + (b+h)\left(\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D\le q\} - (\tilde{q}-D)\mathbb{1}\{D\le\tilde{q}\}]\right)\\ &= (b+h)\left(\sup_{\tilde{q}\in[\lambda,\widetilde{M}]} \rho(\tilde{q}-q) + \left(\sup_{\substack{F\in\mathcal{F}(\lambda;G):\\ q_F^\star=\tilde{q}}} \mathbb{E}_F[(q-D)\mathbb{1}\{D\le q\} - (\tilde{q}-D)\mathbb{1}\{D\le\tilde{q}\}]\right)\right) \qquad \text{(EC.47)}\end{aligned}$$

Note that, for $\tilde{q} = q$, (EC.47) $= 0$. Therefore, we partition our analysis into two cases: $\tilde{q} < q$ and $\tilde{q} > q$.

(i) $q > \tilde{q} \ge \lambda$: In this case, by the proof of Lemma 2 (Eq. (EC.4)), the worst-case distribution F that has $\tilde{q}$ as its newsvendor solution maximizes:

$$\mathbb{E}_F[(q-\tilde{q})\mathbb{1}\{D\le\tilde{q}\} + (q-D)\mathbb{1}\{\tilde{q}<D\le q\}]. \qquad \text{(EC.48)}$$

Observe that $q - D < q - \tilde{q}$ for all $D \in (\tilde{q}, q]$. Hence, any worst-case distribution F should minimize $\mathbb{P}_F(D \in (\tilde{q}, q])$, which occurs by letting $\tilde{q}$ be the maximum of the support of the distribution. Therefore:

$$\begin{aligned}&(\text{EC.48}) = (q - \tilde{q})\\ \implies &\text{Regret}(q) = (b+h) \sup_{\tilde{q} \in [\lambda, \widetilde{M}]} (1-\rho)(q-\tilde{q}) = (b+h)(1-\rho)(q-\lambda).\end{aligned}$$

(ii) $q < \lambda \le \tilde{q}$. As in the proof of Lemma 2 (Eq. (EC.7)), for a fixed $\tilde{q}$, it suffices to find the distribution that maximizes:

$$\sup_{\substack{F \in \mathcal{F}(\lambda; G):\\ q_F^\star = \tilde{q}}} \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G[(\tilde{q}-D)\mathbb{1}\{D < \lambda\}] - \mathbb{E}_F[(\tilde{q}-D)\mathbb{1}\{\lambda \le D < \tilde{q}\}]. \quad \text{(EC.49)}$$

Suppose first that $\tilde{q} > \lambda$. Since $\tilde{q} - D > 0$ for all $D \in [\lambda, \tilde{q})$, a similar argument as above establishes that any worst-case distribution $F \in \mathcal{F}(\lambda; G)$ necessarily sets $f(x) = \gamma$ for all $x \in [\lambda, \tilde{q})$. Using this fact above, we obtain:

$$\begin{aligned}(\text{EC.49}) &= \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G[(\tilde{q}-D)\mathbb{1}\{D < \lambda\}] - \gamma \int_\lambda^{\tilde{q}} (\tilde{q} - x)dx\\ &= \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G[(\tilde{q}-D)\mathbb{1}\{D < \lambda\}] - \frac{\gamma}{2}(\tilde{q}-\lambda)^2.\end{aligned}$$

Plugging this back into (EC.47), we have:

$$(\text{EC.47}) = (b+h)\left(\sup_{\tilde{q} \in [\lambda, \widetilde{M}]} \rho(\tilde{q}-q) + \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G[(\tilde{q}-D)\mathbb{1}\{D < \lambda\}] - \frac{\gamma}{2}(\tilde{q}-\lambda)^2\right).$$

By the first-order condition, the unconstrained supremum of the above is achieved at the smallest value of $\tilde{q}$ such that:

$$\rho - G^-(\lambda) - \gamma(\tilde{q} - \lambda) \le 0 \iff \tilde{q} \ge \lambda + \frac{\rho - G^-(\lambda)}{\gamma}.$$

Since $\tilde{q} \le \widetilde{M} \le \frac{\rho - G^-(\lambda)}{\gamma}$, the constrained supremum is attained at $\widetilde{M}$, and

$$\text{Regret}(q) = (b+h)\left(\rho(\widetilde{M}-q) + \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G\left[(\widetilde{M}-D)\mathbb{1}\{D < \lambda\}\right] - \frac{\gamma}{2}(\widetilde{M}-\lambda)^2\right). \quad \text{(EC.50)}$$

Now, if $\tilde{q} = \lambda$:

$$\text{Regret}(q) = (b+h)\left(\rho(\lambda - q) + \mathbb{E}_G[(q-D)\mathbb{1}\{D \le q\}] - \mathbb{E}_G[(\lambda - D)\mathbb{1}\{D < \lambda\}]\right). \quad \text{(EC.51)}$$

Comparing (EC.50) and (EC.51), setting $\tilde{q} > \lambda$ is optimal if and only if:

$$(\rho - G^-(\lambda))(\widetilde{M} - \lambda) - \frac{\gamma}{2}(\widetilde{M}-\lambda)^2 \ge 0 \iff \rho - G^-(\lambda) - \frac{\gamma}{2}(\widetilde{M}-\lambda) \ge 0.$$

The above holds, since by definition $\widetilde{M} \le \lambda + \frac{\rho - G^-(\lambda)}{\gamma}$, and $\rho - G^-(\lambda) > 0$.

(iii) $\lambda \le q < \tilde{q}$: In this case, by the proof of Lemma 2 (Eq. (EC.10)), it suffices to find a distribution F that maximizes:

$$-(\tilde{q}-q)G^-(\lambda) - (\tilde{q}-q)\mathbb{P}_F(\lambda \le D \le q) - \mathbb{E}_F\Big[(\tilde{q}-D)\mathbb{1}\{q < D \le \tilde{q}\}\Big]. \tag{EC.52}$$

Since $0 < \tilde{q} - D < \tilde{q} - q$ for all $D \in (q, \tilde{q})$, the supremum of the above is achieved for F such that $f(x) = \gamma$ for all $x \in [\lambda, \tilde{q})$. In this case, we have:

$$\begin{aligned}(\text{EC.52}) &= -(\tilde{q}-q)G^-(\lambda) - \gamma(\tilde{q}-q)(q-\lambda) - \gamma\int_q^{\tilde{q}}(\tilde{q}-x)dx \\ &= -(\tilde{q}-q)(G^-(\lambda) + \gamma(q-\lambda)) - \frac{\gamma}{2}(\tilde{q}-q)^2.\end{aligned}$$

Plugging this back into (EC.47), we have:

$$(\text{EC.47}) = (b+h)\left(\sup_{\tilde{q}\in[\lambda,\widetilde{M}]} \rho(\tilde{q}-q) - (\tilde{q}-q)(G^-(\lambda)+\gamma(q-\lambda)) - \frac{\gamma}{2}(\tilde{q}-q)^2\right).$$

By the first-order condition, the unconstrained supremum of the above is attained at the smallest $\tilde{q}$ such that:

$$\rho - (G^-(\lambda)+\gamma(q-\lambda)) - \gamma(\tilde{q}-q) \le 0 \iff \tilde{q} \ge q + \frac{\rho - (G^-(\lambda)+\gamma(q-\lambda))}{\gamma} = \lambda + \frac{\rho - G^-(\lambda)}{\gamma}.$$

Again, using the fact that $\tilde{q} \le \widetilde{M} \le \lambda + \frac{\rho - G^-(\lambda)}{\gamma}$, we obtain that the constrained supremum is attained at $\tilde{q} = \widetilde{M}$, and

$$\text{Regret}(q) = (b+h)\left((\widetilde{M}-q)(\rho - (G^-(\lambda)+\gamma(q-\lambda)) - \frac{\gamma}{2}(\widetilde{M}-q)^2\right).$$

Putting these three cases together, we obtain the result. □

E.2. Practical Performance Improvements in the Knife-Edge Regime

We next numerically investigate how the well-separated condition can generate sample complexity gains in the knife-edge regime, where testing for identifiability is the hardest. To do so, we introduce a new algorithm, which we denote RCN$^{\text{WS}}$. Similar to RCN, RCN$^{\text{WS}}$ proceeds hierarchically, by first testing the identifiability regime, and then estimating the appropriate $q^{\Delta^{\text{WS}}}$. However, its identifiability test leverages the well-separatedness condition via the following observations:

1. If $\gamma\lambda \ge \rho$, by the well-separatedness condition, $G^-(\lambda) \ge \lambda\gamma \ge \rho$, and the problem is guaranteed to be identifiable. In this case, no statistical test is needed.
2. If $\widehat{G}^-(\lambda) > \rho$, $q^\star_{\widehat{G}} < \lambda$. Then, by the well-separatedness condition, with high probability:

$$G^-(\lambda) \ge G^-(q^\star_{\widehat{G}}) + \gamma(\lambda - q^\star_{\widehat{G}}) \ge \widehat{G}^-_K(q^\star_{\widehat{G}}) - \zeta_K + \gamma(\lambda - q^\star_{\widehat{G}}) \ge \rho - \zeta_K + \gamma(\lambda - q^\star_{\widehat{G}}).$$

Thus, if $\gamma(\lambda - q^\star_{\widehat{G}}) \ge \zeta_K$, $G^-(\lambda) \ge \rho$, and the problem is identifiable.

ALGORITHM 3: Robust Censored Well-Separated Newsvendor (RCN$^{\text{ws}}$)

Input: Censored demand samples s^{off}, confidence terms $\zeta_k, k \in [K]$, well-separated parameter γ

Output: Ordering quantity q^{alg}

For all $k \in [K]$, compute censored SAA of $G^-(q_k^{\text{off}})$: $\widehat{G}_k^-(q_k^{\text{off}}) = \frac{1}{N_k}\sum_{i\in[N_k]} \mathbb{1}\{s_{ki}^{\text{off}} < q_k^{\text{off}}\}$.

Define $\widehat{G}^-(\lambda) = \widehat{G}_K^-(q_K^{\text{off}})$.

Compute censored SAA of $q_G^\star$: $q_{\widehat{G}}^\star = \inf\left\{x \mid \widehat{G}(x) \geq \rho\right\}$, where $\widehat{G}(x) = \frac{1}{N_K}\sum_{i\in[N_K]} \mathbb{1}\{s_{iK}^{\text{off}} \leq x\}$.

if $\gamma\lambda \geq \rho$ or $\widehat{G}^-(\lambda) \geq \rho + \zeta_K$ or $\widehat{G}_k^-(q_k^{\text{off}}) \geq \rho + \zeta_k - \gamma(\lambda - q_k^{\text{off}})$ for some $k \in [K]$ **then** `// Likely identifiable`

  Let $q^{\text{alg}} = q_{\widehat{G}}^\star$.

else if $\widehat{G}^-(\lambda) > \rho$ **then** `// Borderline: use separation test`

  if $\gamma(\lambda - q_{\widehat{G}}^\star) \geq \zeta_K$ **then** `// Separation overcomes noise`

    Let $q^{\text{alg}} = q_{\widehat{G}}^\star$. `// Likely identifiable`

  else

    Let $q^{\text{alg}} = \lambda$. `// Knife-edge`

else if $\widehat{G}^-(\lambda) \in (\rho - \zeta, \rho]$ **then** `// Knife-edge`

  Let $q^{\text{alg}} = \lambda$.

else `// Likely unidentifiable`

  Compute empirical estimate of $q_G^\dagger$:

$$q_{\widehat{G}}^\dagger = \frac{\widetilde{M}(\rho - \widehat{G}^-(\lambda) + \gamma\lambda) + \lambda(1-\rho) - \frac{\gamma}{2}(\widetilde{M}^2 + \lambda^2)}{1 - \widehat{G}^-(\lambda)} \quad \text{where } \widetilde{M} = \min\left\{M, \frac{1}{\gamma}, \lambda + \frac{\rho - \widehat{G}^-(\lambda)}{\gamma}\right\}. \tag{EC.53}$$

  Let $q^{\text{alg}} = q_{\widehat{G}}^\dagger$.

return q^{alg}

3. If $\widehat{G}_k^-(q_k^{\text{off}}) \geq \rho + \zeta_k - \gamma(\lambda - q_k^{\text{off}})$ for some $k \in [K]$, we similarly have:

$$G^-(\lambda) \geq G^-(q_k^{\text{off}}) + \gamma(\lambda - q_k^{\text{off}}) \geq \widehat{G}_k^-(q_k^{\text{off}}) - \zeta_k + \gamma(\lambda - q_k^{\text{off}}) \geq \rho,$$

   which again guarantees identifiability.

We use these three observations in our modified algorithm, presented in Algorithm 3.

E.2.1. Experimental setup. Across all experiments we let $h = 1$ and $b = 9$, so that $\rho = 0.9$. We consider the following demand distributions:

- *Continuous Uniform*: Demand is drawn from a continuous uniform distribution over $[0, 100]$. We let $M = 100$, and $\gamma = 1/100$.
- *Truncated Exponential*: Demand is drawn from a truncated exponential distribution with rate $1/80$ with support over $[0, 325]$. We let $M = 325$, and $\gamma = \frac{1}{80}e^{-325/80}$.

We moreover consider $K = 2$ historical ordering quantities and vary λ as in Section 6.2.1.

E.2.2. Results. We demonstrate the improvements in classifying the problem as identifiable by comparing the performance of RCN$^{\text{ws}}$ to that of RCN in the knife-edge regime where $\lambda \in [q_G^\star, 1.05q_G^\star]$. Additionally,

λ	90.89	91.79	92.69	93.59	94.49
True SAA	0.1 (0.2%)	0.1 (0.2%)	0.1 (0.1%)	0.1 (0.2%)	0.1 (0.2%)
RCN	0.1 (0.2%)	0.2 (0.4%)	0.4 (0.9%)	0.7 (1.5%)	0.7 (1.5%)
RCN$^{\text{WS}}$	0.1 (0.3%)	0.1 (0.2%)	0.1 (0.2%)	0.1 (0.3%)	0.1 (0.3%)

(a) Uniform, $q^\star_G = 89$

λ	186.69	190.63	194.58	198.52	202.47
True SAA	0.2 (0.1%)	0.2 (0.1%)	0.3 (0.2%)	0.3 (0.2%)	0.2 (0.1%)
RCN	0.1 (0.0%)	0.2 (0.1%)	0.5 (0.3%)	1.1 (0.6%)	1.8 (1.1%)
RCN$^{\text{WS}}$	0.1 (0.0%)	0.2 (0.1%)	0.5 (0.3%)	1.1 (0.6%)	1.8 (1.1%)

(b) Truncated Exponential, $q^\star_G = 184.14$

Table EC.8 Vanilla-Regret(q^π) vs. λ. The relative vanilla regret $\mathcal{R}^{id}(q^\pi)\%$ is shown in parentheses.

in order to isolate the *sample complexity* gains of RCN$^{\text{WS}}$, we modify RCN to instead estimate the value of $q^{\Delta^{\text{WS}}}$ associated with the well-separated setting. Our results can be found in Table EC.8. The performance of the two algorithms in the likely unidentifiable and likely identifiable regimes is similar, and omitted as such.

For the Continuous Uniform distribution (Table EC.8a), RCN$^{\text{WS}}$ yields modest but consistent performance gains over RCN across all tested values of λ. Specifically, for $\lambda \in \{90.89, 91.79, 92.69\}$, RCN$^{\text{WS}}$ achieves a 0% misclassification rate thanks to the "extra conditions" for identifiability, and therefore correctly outputs an estimate of $q^\star_G$. RCN, on the other hand, misclassifies the problem as *knife-edge* in 100% of replications, therefore outputting λ. For $\lambda = 93.59$ and $\lambda = 94.49$, RCN$^{\text{WS}}$ applies the identifiability condition in 90% and 60% of runs respectively (again with no misclassifications), while RCN continues to misclassify the problem at those same rates.

For the Truncated Exponential distribution (Table EC.8b), RCN and RCN$^{\text{WS}}$ achieve identical performance. This behavior arises because the chosen value of γ severely underestimates the local probability density around λ, rendering the additional identifiability conditions in RCN$^{\text{WS}}$ inactive. This highlights the importance of having tight bounds on γ in order to achieve sample complexity gains from this additional distributional information.