Electronic Companion: Managing Multi-Tier Inventory Networks With Expediting Under Normal and Disrupted Modes

Yanyang Zhao, John R. Birge, Levi DeValve, Robert R. Inman

Appendix A: Example From Introduction

The FZ heuristic developed by Federgruen and Zipkin (1984c) and Gallego et al. (2007) is described in detail in Section C.1.1. On the example given in the Introduction, this policy solves a newsvendor model with holding cost 1 and backlog cost 10 and a normal demand distribution with mean 12 and variance 4. This amounts to finding the $10/(10+1) \approx 91\%$ fractile of this demand, which is about 17. The heuristic then allocates all available inventory to the retailers in each period by solving the stochastic program (9) (details in Section C.1.1), which pushes all inventory to the retailers in this case since holding costs are the same at the warehouse and retailers (Gallego et al., 2007). The policies we test on this example are the FZ hueristic with prioritized expediting described in Section C.1.1 and our SP^{L+l} stochastic programming base-stock policy from Section 3.2. We also test two alternative policies to provide more context on the cost improvements obtained by our approach. The first is a version of the FZ heuristic but with no expediting (all else is held the same), which we call FZ^{ne}. The final heuristic we test is a modified version of SP^{L+l} that assumes there is no probability of disruption, i.e., $\alpha = 0$, which we denote by $\text{SP}^{L+l}_{\text{nd}}$. We simulated these policies for 200 sample paths, each with 365 days, and the results are summarized in Table 10.

The SP^{L+l} policy is the only one that directly plans for disruptions and holds central inventory, and it has statistically significant lower cost than all other policies (the number of asterisks indicates the number of statistically significant cost comparisons to the other policies at a 95% confidence level) and keeps backlogs at the lowest level. The FZ^{ne} heuristic is clearly dominated, as it has over 1 unit of backlog per retailer per day, which is equal to the average demand per retailer per day, demonstrating the obvious result that expediting should be taken advantage of. The $\text{SP}^{L+l}_{\text{nd}}$ policy, while achieving lower cost than either FZ heuristic, still has high backlogs, and leaves about 2.6% extra cost on the table by not taking into account the chance of disruption. In this example, it is the increased and centralized inventory of the SP^{L+l} policy that allows it to clear backlogs more effectively than the other policies.

Table 10 Simulation for Introduction Example

Policy	System BSL	Central BSL	Cost	Backlog/day/ret.	Expedite/day/ret.
FZ^{ne}	17	0	29.14	1.01	0
FZ	17	0	27.15^{*}	0.65	0.19
$\text{SP}^{L+l}_{\text{nd}}$	18	0	25.74^{**}	0.59	0.16
SP^{L+l}	24	4	25.08^{***}	0.32	0.10

Appendix B: Technical Proofs

B.1. Further Explanation of SP Objective and Constraints

Before presenting the formal proofs, we first provide more explanation of the objective and constraints in the stochastic program (5). For the objective, g_3 provides a lower bound on the cost of a policy over $\tau+1$ periods: the first two terms on the first line (with cost coefficients of h_0 and h_i) represent the holding costs at the warehouse and retailers, respectively. The coefficient of $\tau+1$ on the first term reflects the fact that inventory held at the warehouse at the end of the third stage must have been there for the entire course of the final $\tau+1$ periods. The third term on the first line represents the cost of expediting during the third stage, which is normalized by l to account for the duration of the third stage. The second line represents backlogging cost, where the first term inside the large parentheses represents the backlogging cost incurred in period $L+\hat{l}$ from demand over the first $L+\hat{l}$ periods (i.e., from just demand $\mathbf{D}^1_{\hat{l}} + \mathbf{D}^3_{\hat{l}}$), and the second term represents the backlogging cost incurred during the final τ periods from demand over all $L+\hat{l}+\tau$ periods. The second term includes a novel characterization of these costs in terms of the disruption duration τ. Then, g_2 represents the third stage recourse cost of local fulfillment and expediting after the realization of third stage demand, and g_1 represents the second stage expected recourse cost of allocating and expediting after the realization of second stage demand, while $C_{\hat{l}}$ represents the expected cost of the initial inventory and backlog variables.

The third stage constraints in the definition of g_2 capture the requirements that all allocation and expediting must use warehouse inventory, all local fulfillment must use local inventory, and only realized demand (plus any initial backlog) can be fulfilled, respectively. The second stage constraints in the definition of g_1 capture the requirements that all allocation and expediting must use warehouse inventory, and second stage expediting can only fulfill second stage demand.

Finally, we further explain the intuition behind the demand distributions defined for the stochastic program. In particular, we define $\mathbf{D}^1_{\hat{l}}$ and $\mathbf{D}^2_{\hat{l}}$ separately instead of just defining their sum $\mathbf{D}^1_{\hat{l}} + \mathbf{D}^2_{\hat{l}}$ (and similarly for $\mathbf{D}^3_{\hat{l}}$ and $\mathbf{D}^4_{\hat{l}}$), since we can use these separate variables to construct the necessary periods of demand as follows: $\mathbf{D}^1_{\hat{l}} + \mathbf{D}^3_{\hat{l}} = \sum_{j=1}^{L+\hat{l}} \mathbf{D}(j)$, and $\mathbf{D}^2_{\hat{l}} + \mathbf{D}^4_{\hat{l}} = \sum_{j=L+\hat{l}+1}^{L+\hat{l}+\tau} \mathbf{D}(j)$, i.e., $\mathbf{D}^1_{\hat{l}} + \mathbf{D}^3_{\hat{l}}$ represents demand over the first $L+\hat{l}$ periods, and $\mathbf{D}^2_{\hat{l}} + \mathbf{D}^4_{\hat{l}}$ represents demand over the next τ periods. To see that these expressions are valid, note that we always have either $\mathbf{D}^2_{\hat{l}} = 0$ (if $l \geq \tau$) or $\mathbf{D}^3_{\hat{l}} = 0$ (if $l \leq \tau$). Thus, if $l \geq \tau$ then $\mathbf{D}^1_{\hat{l}} = \sum_{j=1}^{L+\hat{l}-l+\tau} \mathbf{D}(j)$ and $\mathbf{D}^3_{\hat{l}} = \sum_{j=L+\hat{l}-l+\tau+1}^{L+\hat{l}} \mathbf{D}(j)$, while if $l \leq \tau$ then $\mathbf{D}^1_{\hat{l}} = \sum_{j=1}^{L+\hat{l}} \mathbf{D}(j)$ and $\mathbf{D}^3_{\hat{l}} = 0$, so in either case we have $\mathbf{D}^1_{\hat{l}} + \mathbf{D}^3_{\hat{l}} = \sum_{j=1}^{L+\hat{l}} \mathbf{D}(j)$. A similar case argument verifies that $\mathbf{D}^2_{\hat{l}} + \mathbf{D}^4_{\hat{l}} = \sum_{j=L+\hat{l}+1}^{L+\hat{l}+\tau} \mathbf{D}(j)$. To summarize, our stochastic program uses $\mathbf{D}^1_{\hat{l}} + \mathbf{D}^2_{\hat{l}}$ as the second stage demand, and $\mathbf{D}^3_{\hat{l}} + \mathbf{D}^4_{\hat{l}}$ as the third stage demand, as well as $\mathbf{D}^1_{\hat{l}} + \mathbf{D}^3_{\hat{l}}$ and $\mathbf{D}^2_{\hat{l}} + \mathbf{D}^4_{\hat{l}}$ to calculate costs.

B.2. Proof of Theorem 1

In our proof of Theorem 1, we assume $\hat{l}=0$ and therefore drop dependence on $\hat{l}$ for convenience. Before proving the theorem, we first present the following lemma.

Lemma 2. *For an integer $a \geq 1$ and reals X, Y with $X \geq 0$, we have $\sum_{j=1}^{a}\left(\frac{jX}{a}-Y\right)^{+} \geq \left(\left(\frac{a+1}{2}\right)X-aY\right)^{+}$.*

Proof. First note that the left hand side is non-negative by definition, so if the right hand side evaluates to zero, the claim follows. Thus, we show the claim for the case when the positive part on the right hand side is non-binding. Further, note that if $\frac{j}{a}X-Y \leq 0$ for all $1 \leq j \leq a$, then we have $X-Y \leq 0$ (from the case $j=a$), from which it follows $\frac{a+1}{2}X-aY \leq a(X-Y) \leq 0$ (because $\frac{a+1}{2} \leq a$ for $a \geq 1$), implying that both sides evaluate to zero. Therefore, consider the remaining case where $\frac{j}{a}X-Y \geq 0$ for some $1 \leq j \leq a$, and let j' be the smallest such j (and note that $\frac{j}{a}X-Y \geq 0$ for all $j \geq j'$ since $X \geq 0$). Then, note that either $j'=1$, or $\frac{j'-1}{a}X-Y \leq 0$.

Now, the left hand side evaluates to $\sum_{j=j'}^{a}\left(\frac{jX}{a}-Y\right)=\frac{X}{2a}\left(a(a+1)-(j'-1)j'\right)-(a-j'+1)Y=\left(\frac{a+1}{2}\right)X-aY+(j'-1)\left(Y-\frac{Xj'}{2a}\right)$, so that to establish the claim we must show the final term is larger than zero. It equals zero if $j'=1$, and otherwise, if $\frac{j'-1}{a}X-Y \leq 0$ then $(j'-1)\left(Y-\frac{Xj'}{2a}\right) \geq (j'-1)\left(\frac{j'-1}{a}X-\frac{Xj'}{2a}\right)=\frac{X(j'-1)(j'-2)}{2a} \geq 0$, where the final inequality follows from the fact that j' is an integer larger than 1. Q.E.D.

Proof of Theorem 1. First, observe that on any sample path, for any $T \geq 1$, we have $\sum_{t=1}^{T}\sum_i f_i y_{i,t} \geq \sum_{t=1}^{T}\sum_i \frac{f_i}{l}\sum_{s=t-l+1}^{t} y_{i,s}$, where the inequality follows from the convention that $y_{i,s}=0$ for any $s \leq 0$, and observing that the summation on the right hand side undercounts the $y_{i,t}$ variables for periods at the end of the time horizon, i.e., $T-l+1 \leq t \leq T$. Thus, to establish a lower bound on the long run average cost in (4), we prove a lower bound on the following cost

$$\limsup_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left[h_0(I_t+IT_t)+\sum_i\left(h_iX_{i,t}+b_iB_{i,t}+\frac{f_i}{l}\sum_{s=t-l+1}^{t}y_{i,s}\right)\right]. \tag{11}$$

Next we perform a sample path cost accounting; i.e., on certain sets of sample paths we will assign certain costs to a specific period to account for the costs incurred by a feasible policy. In particular, we fix a period t and assign costs to this period based on the evolution of the state r_t as follows.

$$C(t,r_t,\tau_{t+1})=\begin{cases}0, & r_t>0\\ \sum_{u=t+L}^{t+L+\tau_{t+1}} h_0(I_u+IT_u)+\sum_i\left(h_iX_{i,u}+b_iB_{i,u}+\frac{f_i}{l}\sum_{s=u-l+1}^{u}y_{i,s}\right), & r_t=0\end{cases}$$

Observe that each term in (11) is only counted once by $C(t,r_t,\tau_{t+1})$ on each sample path, because if $r_t>0$ then no costs are accounted for by period t, while if $r_t=0$ then $C(t,r_t,\tau_{t+1})$ accounts for costs during all periods from $t+L$ (when an order from period t arrives at the warehouse) to $t+L+\tau_{t+1}$ (the last period before the next order after period t arrives at the warehouse). If $\tau_{t+1}=0$, it is clear that the costs from period $t+L$ are only counted once in period t by $C(t,r_t,\tau_{t+1})$. Meanwhile, if $r_t=0$ and $\tau_{t+1}>0$, then a disruption began in period $t+1$, which will last τ_{t+1} periods, i.e., we will have $r_s>0$ for all periods $s\in\{t+1,\dots,t+\tau_{t+1}\}$, meaning no costs are accounted by $C(s,r_s,\tau_{s+1})$ during these periods (by definition), and all costs from period $t+L$ to $t+L+\tau_{t+1}$ are counted once in period t by $C(t,r_t,\tau_{t+1})$. Thus, by (11), it is clear that the long run average cost is at least as large as $\limsup_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\mathbb{E}\left[C(t,r_t,\tau_{t+1})\right]$, and we thus establish the claim by showing the following for all t:

$$\mathbb{E}\left[C(t,r_t,\tau_{t+1})\right] \geq \frac{\underline{C}}{1+\mathbb{E}[\tau]} \tag{12}$$

By the evolution equations, in any period $u \geq t$ we have

$$\begin{aligned}
I_u+IT_u &= I_t+\sum_{s=t-L+1}^{u-L}x_s-\sum_{s=t+1}^{u}\sum_i y_{i,s}-\sum_{s=t+1}^{u-l}\sum_i z_{i,s},\\
X_{i,u} &= X_{i,t}+\sum_{s=t-l+1}^{u-l}z_{i,s}-\sum_{s=t+1}^{u}w_{i,s}\\
B_{i,u} &= B_{i,t}+\sum_{s=t+1}^{u}\left(D_{i,s}-w_{i,s}-y_{i,s}\right)
\end{aligned}$$

Next, we consider the case when $r_t=0$, and use the evolution equations above to write the costs in $C(t,r_t,\tau_{t+1})$ in terms of decisions made from period t onward:

$$\begin{aligned}
&C(t,r_t=0,\tau_{t+1}), \qquad\qquad (13)\\
&=\sum_{u=t+L}^{t+L+\tau_{t+1}}\left(h_0(I_u+IT_u)+\sum_i\left(h_iX_{i,u}+b_iB_{i,u}+\frac{f_i}{l}\sum_{s=u-l+1}^{u}y_{i,s}\right)\right),\\
&=\sum_{u=t+L}^{t+L+\tau_{t+1}}\left(h_0\left(I_t+\sum_{s=t-L+1}^{u-L}x_s-\sum_{s=t+1}^{u}\sum_i y_{i,s}-\sum_{s=t+1}^{u-l}\sum_i z_{i,s}\right)^{+}+\sum_i\frac{f_i}{l}\sum_{s=u-l+1}^{u}y_{i,s}\right)\\
&\quad+\sum_{u=t+L}^{t+L+\tau_{t+1}}\sum_i\left(h_i\left(X_{i,t}+\sum_{s=t-l+1}^{u-l}z_{i,s}-\sum_{s=t+1}^{u}w_{i,s}\right)^{+}+b_i\left(B_{i,t}+\sum_{s=t+1}^{u}\left(D_{i,s}-w_{i,s}-y_{i,s}\right)\right)^{+}\right),
\end{aligned}$$

$$
\begin{aligned}
&\geq h_0(1+\tau_{t+1})\left(I_t+\sum_{s=t-L+1}^{t}x_s-\sum_{s=t+1}^{t+L+\tau_{t+1}}\sum_i y_{i,s}-\sum_{s=t+1}^{t+L-l+\tau_{t+1}}\sum_i z_{i,s}\right)^+\\
&\quad+\sum_i h_i\left(X_{i,t}+\sum_{s=t-l+1}^{t+L-l+\tau_{t+1}}z_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}w_{i,s}\right)^+ +\sum_i\frac{f_i}{l}\sum_{s=t+L-l+1+\tau_{t+1}}^{t+L+\tau_{t+1}}y_{i,s}\\
&\quad+\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}(w_{i,s}+y_{i,s})\right)^+ \qquad (14)\\
&\quad+\sum_{u=t+L+1}^{t+L+\tau_{t+1}}\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{u}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}(w_{i,s}+y_{i,s})\right)^+, \qquad (15)
\end{aligned}
$$

where all the positive parts in the second equality follow from the non-negativity of the associated inventory and backlog variables (so the positive parts are non-binding here), and the first line after the inequality corresponds to the warehouse holding cost in the expression before the inequality (i.e., the term starting with multiplication by h_0) and follows from noting $x_s\geq 0$, $y_{i,s}\geq 0$, and $z_{i,s}\geq 0$, the second line after the inequality corresponds to the retailer holding cost in the expression before the inequality (i.e., the term starting with multiplication by h_i) for $u=t+L+\tau_{t+1}$ plus the expediting cost (i.e., the term starting with multiplication by f_i/l) and follows from $y_{i,s}\geq 0$, the third line after the inequality (i.e., (14)) corresponds to the backlogging cost in the expression before the inequality (i.e., the term starting with multiplication by b_i) for $u=t+L$ and follows from $w_{i,s}\geq 0$ and $y_{i,s}\geq 0$, and the final line after the inequality (i.e., (15)) corresponds to the backlogging cost in the expression before the inequality for the remaining $u>t+L$ and follows from $w_{i,s}\geq 0$ and $y_{i,s}\geq 0$.

Now consider the following two sigma algebras:

$$
\begin{aligned}
\mathcal{S}_t&=\sigma\{x_s,w_{i,s},y_{i,s},z_{i,s},\forall i,s\leq t-1,D_{i,s},\forall i,s\leq t,r_t\},\\
\mathcal{A}_t&=\sigma\left\{x_s,w_{i,s},y_{i,s},z_{i,s},\forall i,s\leq t-1,D_{i,s},\forall i,s\leq t+L,r_t,r_{t+1},\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s},\forall i\right\},
\end{aligned}
$$

where the first, $\mathcal{S}_t$, denotes all information available before the ordering, shipping, fulfillment, and expediting decisions are made on day t. The second, $\mathcal{A}_t$, includes all information in $\mathcal{S}_t$ (ie., $\mathcal{S}_t\subseteq\mathcal{A}_t$), but also includes the demand realizations over the next L periods, the cumulative demand realizations over the following τ_{t+1} periods, and the disruption state r_{t+1}. The σ-algebra $\mathcal{A}_t$ will be used to simplify the cost expression with an intermediate application of the conditional Jensen's inequality. To this end, consider a single element, $u\in\{t+L+1,\ldots,t+L+\tau_{t+1}\}$, from the sum in (15) for a given retailer i in the case when $\tau_{t+1}\geq 1$ (since otherwise there is no such sum), and we will consider its expected value conditioned on $\mathcal{A}_t$:

$$
\begin{aligned}
&\mathbb{E}\left[\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}+\sum_{s=t+L+1}^{u}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}(w_{i,s}+y_{i,s})\right)^+\middle|\mathcal{A}_t\right],\\
&\geq\left(\mathbb{E}[B_{i,t}|\mathcal{A}_t]+\sum_{s=t+1}^{t+L}\mathbb{E}[D_{i,s}|\mathcal{A}_t]+\sum_{s=t+L+1}^{u}\mathbb{E}[D_{i,s}|\mathcal{A}_t]-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}[w_{i,s}+y_{i,s}|\mathcal{A}_t]\right)^+,\\
&=\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}+\sum_{s=t+L+1}^{u}\frac{\sum_{s'=t+L+1}^{t+L+\tau_{t+1}}D_{i,s'}}{\tau_{t+1}}-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}[w_{i,s}+y_{i,s}|\mathcal{A}_t]\right)^+,
\end{aligned}
$$

where the inequality is conditional Jensen's, and the equality follows from $\mathbb{E}[B_{i,t}|\mathcal{A}_t]=B_{i,t}$ and $\mathbb{E}[D_{i,s}|\mathcal{A}_t]=D_{i,s}$ for $s\in\{t+1,\ldots,t+L\}$ (since these random variables are measurable with respect to $\mathcal{A}_t$), and the fact that

$$
\mathbb{E}[D_{i,s}|\mathcal{A}_t]=\frac{\sum_{s'=t+L+1}^{t+L+\tau_{t+1}}D_{i,s'}}{\tau_{t+1}} \qquad (16)
$$

for all $s\in\{t+L+1,\ldots,t+L+\tau_{t+1}\}$, which we verify with the following argument. Observe that for any $s,s'\in\{t+L+1,\ldots,t+L+\tau_{t+1}\}$, we have $\mathbb{E}[D_{i,s}|\mathcal{A}_t]=\mathbb{E}\left[D_{i,s'}|\mathcal{A}_t\right]$, by the assumption that demand is i.i.d. across time, and this coupled with the fact that $\mathbb{E}\left[\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}|\mathcal{A}_t\right]=\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}$, (since the sum is measurable with respect to $\mathcal{A}_t$) implies (16).

Next observe that for a given i we have

$$
\begin{aligned}
&\sum_{u=t+L+1}^{t+L+\tau_{t+1}}\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}+\sum_{s=t+L+1}^{u}\frac{\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}}{\tau_{t+1}}-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}[w_{i,s}+y_{i,s}|\mathcal{A}_t]\right)^+,\\
&=\sum_{v=1}^{\tau_{t+1}}\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}+v\frac{\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}}{\tau_{t+1}}-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}[w_{i,s}+y_{i,s}|\mathcal{A}_t]\right)^+,\\
&\geq\left(\frac{\tau_{t+1}+1}{2}\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}+\tau_{t+1}\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}[w_{i,s}+y_{i,s}|\mathcal{A}_t]\right)\right)^+,
\end{aligned}
$$

where the inequality follows from Lemma 2. Thus, we have:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{u=t+L+1}^{t+L+\tau_{t+1}}\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{u}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}(w_{i,s}+y_{i,s})\right)^+\middle|\mathcal{S}_t,r_t=0\right], \qquad (17)\\
&=\mathbb{E}\left[\mathbb{E}\left[\sum_{u=t+L+1}^{t+L+\tau_{t+1}}\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{u}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}(w_{i,s}+y_{i,s})\right)^+\middle|\mathcal{A}_t\right]\middle|\mathcal{S}_t,r_t=0\right],\\
&\geq\mathbb{E}\left[\sum_i b_i\left(\frac{\tau_{t+1}+1}{2}\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}+\tau_{t+1}\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}+-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}\left[w_{i,s}+y_{i,s}|\mathcal{A}_t\right]\right)\right)^+\middle|\mathcal{S}_t,r_t=0\right].
\end{aligned}$$

Then, we observe that all terms in the lines preceding (15) (after the inequality), can also use the conditional Jensen's inequality to lower bound their expected costs in a similar way. For example consider the sum in (14); conditional Jensen's inequality implies

$$\begin{aligned}
&\mathbb{E}\left[\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}(w_{i,s}+y_{i,s})\right)^+\middle|\mathcal{S}_t,r_t=0\right], \qquad (18)\\
&=\mathbb{E}\left[\mathbb{E}\left[\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}(w_{i,s}+y_{i,s})\right)^+\middle|\mathcal{A}_t\right]\middle|\mathcal{S}_t,r_t=0\right],\\
&\geq\mathbb{E}\left[\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}\left[w_{i,s}+y_{i,s}|\mathcal{A}_t\right]\right)^+\middle|\mathcal{S}_t,r_t=0\right],
\end{aligned}$$

where, again, the final line follows from the conditional Jensen's inequality and the fact that $B_{i,t}$ and $\sum_{s=t+1}^{t+L}D_{i,s}$ are measurable with respect to $\mathcal{A}_t$.

Using similar bounds on the lines preceding (14) (after the inqequality), and combining them with the inequality beginning in (13) and ending in (15), and with the inequalities beginning in (17) and (18) gives

$$\begin{aligned}
&\mathbb{E}\left[C(t,r_t,\tau_{t+1})|\mathcal{S}_t,r_t=0\right],\\
&\geq\mathbb{E}\left[h_0(1+\tau_{t+1})\left(I_t+\sum_{s=t-L+1}^{t}x_s+\mathbb{E}\left[-\sum_{s=t+1}^{t+L+\tau_{t+1}}\sum_i y_{i,s}-\sum_{s=t+1}^{t+L-l+\tau_{t+1}}\sum_i z_{i,s}\middle|\mathcal{A}_t\right]\right)^+\middle|\mathcal{S}_t,r_t=0\right]\\
&\quad+\mathbb{E}\left[\sum_i h_i\left(X_{i,t}+\sum_{s=t-l+1}^{t}z_{i,s}+\mathbb{E}\left[\sum_{s=t+1}^{t+L-l+\tau_{t+1}}z_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}w_{i,s}\middle|\mathcal{A}_t\right]\right)^+\middle|\mathcal{S}_t,r_t=0\right]\\
&\quad+\mathbb{E}\left[\sum_i\frac{f_i}{l}\sum_{s=t+L-l+1+\tau_{t+1}}^{t+L+\tau_{t+1}}\mathbb{E}\left[y_{i,s}|\mathcal{A}_t\right]+\sum_i b_i\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}\left[w_{i,s}+y_{i,s}|\mathcal{A}_t\right]\right)^+\middle|\mathcal{S}_t,r_t=0\right]\\
&\quad+\mathbb{E}\left[\sum_i b_i\left(\frac{\tau_{t+1}+1}{2}\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}+\tau_{t+1}\left(B_{i,t}+\sum_{s=t+1}^{t+L}D_{i,s}-\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}\left[w_{i,s}+y_{i,s}|\mathcal{A}_t\right]\right)\right)^+\middle|\mathcal{S}_t,r_t=0\right],\\
&\geq \underline{C},
\end{aligned}$$

where the final inequality follows by observing that $\sum_{s=t+1}^{t+L}D_{i,s}$ and $\sum_{s=t+L+1}^{t+L+\tau_{t+1}}D_{i,s}$ have the same distribution as $D_i^1+D_i^3$ and $D_i^2+D_i^4$, respectively, and by constructing the following feasible solution for the stochastic program (5): first stage variables $I=I_t+\sum_{s=t-L+1}^{t}x_s$, $X_i=X_{i,t}+\sum_{s=t-l+1}^{t}z_{i,s}$, and $B_i=B_{i,t}$ for all i, second stage variables $z_i=\sum_{s=t+1}^{t+L-l+\tau_{t+1}}\mathbb{E}\left[z_{i,s}|\mathcal{A}_t\right]$ and $y_i^2=\sum_{s=t+1}^{t+L-l+\tau_{t+1}}\mathbb{E}\left[y_{i,s}|\mathcal{A}_t\right]$ for all i, and third stage variables $w_i=\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}\left[w_{i,s}|\mathcal{A}_t\right]$ and $y_i^3=\sum_{s=t+L-l+1+\tau_{t+1}}^{t+L+\tau_{t+1}}\mathbb{E}\left[y_{i,s}|\mathcal{A}_t\right]$ for all i. To see that these variables are feasible for the information sets revealed by the second and third stage demand in the stochastic program, note that $\sum_{s=t+1}^{t+L+\tau_{t+1}}\mathbb{E}\left[w_{i,s}|\mathcal{A}_t\right]$ and $\sum_{s=t+L-l+1+\tau_{t+1}}^{t+L+\tau_{t+1}}\mathbb{E}\left[y_{i,s}|\mathcal{A}_t\right]$ are adapted to the full realization of demand over periods $t+1,\ldots,t+L+\tau_{t+1}$, and thus the third stage variables, w_i and y_i^3, are appropriately adapted to the the full realization of second and third stage demand, $\mathbf{D}^1+\mathbf{D}^2+\mathbf{D}^3+\mathbf{D}^4$, in the stochastic program. Meanwhile, $\sum_{s=t+1}^{t+L-l+\tau_{t+1}}\mathbb{E}\left[z_{i,s}|\mathcal{A}_t\right]$ and $\sum_{s=t+1}^{t+L-l+\tau_{t+1}}\mathbb{E}\left[y_{i,s}|\mathcal{A}_t\right]$ are adapted to the demand over periods $t+1,\ldots,t+L-l+\tau_{t+1}$, and thus the second stage variables z_i and y_i^2 are appropriately adapted to the realization of second stage demand, $\mathbf{D}^1+\mathbf{D}^2$, in the stochastic program. Finally, we demonstrate feasibility of these variables for the constraints of the stochastic program. By the constraint that $0\leq I_{t+L+\tau_{t+1}}$ and the evolution equations, we must have

$$0\leq I_{t+L+\tau_{t+1}},=I_t+\sum_{s=t-L+1}^{t+\tau_{t+1}}x_s-\sum_{s=t+1}^{t+L+\tau_{t+1}}(y_{i,s}+z_{i,s}),=I_t+\sum_{s=t-L+1}^{t}x_s-\sum_{s=t+1}^{t+L+\tau_{t+1}}(y_{i,s}+z_{i,s}),$$

where the second inequality follows since $x_s = 0$ for all $s \in \{t+1,\dots,t+\tau_{t+1}\}$ since the system is disrupted during this time. Therefore, we have

$$\sum_{s=t+1}^{t+L-l+\tau_{t+1}} z_{i,s} + y_{i,s} \leq \sum_{s=t+1}^{t+L-l+\tau_{t+1}} z_{i,s} + \sum_{s=t+1}^{t+L+\tau_{t+1}} y_{i,s} \leq \sum_{s=t+1}^{t+L+\tau_{t+1}} z_{i,s} + y_{i,s} \leq I_t + \sum_{s=t-L+1}^{t} x_s,$$

where the first two inequalities follow from $y_{i,s}, z_{i,s} \geq 0$. After taking expectations, these inequalities imply both $\sum_i z_i + y_i^2 \leq I$ (i.e., the first of the second stage constraints in the stochastic program (5)) and $\sum_i z_i + y_i^2 + y_i^3 \leq I$ (i.e., the first of the third stage constraints in the stochastic program (5)). Similar arguments starting with the constraints $X_{i,t+L+\tau_{t+1}} \geq 0$, $B_{i,t+L-l+\tau_{t+1}} \geq 0$, and $B_{i,t+L+\tau_{t+1}} \geq 0$ demonstrate feasibility for the remaining second and third stage constraints.

Taking a final expectation gives

$$\begin{aligned}
\mathbb{E}\left[C(t,r_t,\tau_{t+1})\right] &= \mathbb{E}\left[\mathbb{E}\left[C(t,r_t,\tau_{t+1})|\mathcal{S}_t, r_t\right]\right], \\
&= \mathbb{P}\left(r_t = 0\right)\mathbb{E}\left[\mathbb{E}\left[C(t,r_t,\tau_{t+1})|\mathcal{S}_t, r_t = 0\right]\right] + \mathbb{P}\left(r_t > 0\right)\mathbb{E}\left[\mathbb{E}\left[C(t,r_t,\tau_{t+1})|\mathcal{S}_t, r_t > 0\right]\right], \\
&= \frac{\mathbb{E}\left[\mathbb{E}\left[C(t,r_t,\tau_{t+1})|\mathcal{S}_t, r_t = 0\right]\right]}{1+\mathbb{E}\left[\tau\right]}, \\
&\geq \frac{\underline{C}}{1+\mathbb{E}\left[\tau\right]},
\end{aligned}$$

where the third equality follows from $C(t,r_t,\tau_{t+1}) = 0$ for $r_t > 0$ and the stationary distribution characterization of $\mathbb{P}\left(r_t = 0\right)$. Q.E.D.

B.3. Proof of Theorem 2

Before commencing with the proof of Theorem 2, we first discuss a few important implications of the result. First, it verifies the intuition that decentralization should be preferred when the expediting cost, f, is large relative to the backlog cost, b. Naturally, it also shows that shorter disruption times, τ, favor decentralization. More interestingly, however, Theorem 2 also implies that as either demand or supply grow large, decentralization becomes a better option. This is intuitive for large supply, a (since plentiful supply is better sent to the the retailers at a lower cost), but perhaps less so for large demand, λ. To explain the intuition with large λ, we consider two cases. First, when supply and demand are relatively balanced, i.e., $\lambda \approx a$, backlog in the decentralized system occurs toward the end of the time horizon, incurring only a small backlog cost, while the centralized system incurs an expediting cost that grows linearly in λ. Second, when demand far outweighs supply $\lambda >> a$, most of the retailers would be able to use all their inventory, and so it is best for them to store it locally to avoid the expediting cost.

This highlights an important point: decentralization only suffers a disadvantage when inventory at one retailer would have been better used by another retailer, i.e., one retailer stocks out while another has excess. When demand and supply are relatively balanced, this disadvantage is limited by the fact that backlog occurs toward the end of the horizon, and when demand is very large, this disadvantage has a low probability of occurrence.

Next, we begin our preparation for the proof of Theorem 2 with a characterization of the expected cost of the centralized and decentralized systems in Section 4. The advantage of using the Poisson process approximation of periodic demand is that we can write these costs in a simple closed form, presented in the next Lemma, which is amenable to further analysis.

Lemma 3. *The centralized system's expected total cost is*

$$\begin{aligned}
& b\tau\left(\frac{n(\lambda-a)^2+a}{2\lambda}\mathbb{P}\left(D^n \geq na\right) + \frac{n(\lambda-a)-1}{2}\mathbb{P}\left(D^n = na-1\right)\right) \\
& \quad + fn\left(\lambda\mathbb{P}\left(D^n < na\right) + a\mathbb{P}\left(D^n > na\right)\right).
\end{aligned} \tag{19}$$

The decentralized system's expected total cost is

$$bn\tau\left(\frac{(\lambda-a)^2+a}{2\lambda}\mathbb{P}\left(D_1 \geq a\right) + \frac{\lambda-a-1}{2}\mathbb{P}\left(D_1 = a-1\right)\right). \tag{20}$$

Lemma 3 presents exact expressions for the expected cost of each policy, and so can be used to evaluate and compare costs directly for any particular problem parameters. However, to gain further insight into the tradeoffs involved, we also derive a simple sufficient condition (Theorem 2) for decentralization to be preferred.

Proof of Lemma 3. We first find the expected backlog cost of an individual retailer i for Poisson D with rate λ and inventory a using the law of total expectation. Conditioning on there being $D = d$ arrivals of the Poisson process at an retailer, the arrivals are uniformly distributed over the time horizon T (Durrett, 1999). Thus, using order statistics of uniform sample of size d and assuming that a is integral, the expected backlog cost is

$$\mathbb{E}\left[B|D=d\right] = bT\sum_{k=a+1}^{d}\left(1-\frac{k}{d+1}\right) = \frac{bT}{d+1}\sum_{j=1}^{d-a} j = \frac{bT(d-a)^+(d-a+1)}{2(d+1)} = bT\left(\frac{1}{2}d + \frac{a(a+1)}{2(d+1)} - a\right)\mathbb{1}_{\{d \geq a\}},$$

where the first equality follows from a re-indexing with $j = d+1-k$ and reversing the order of the sum, and the second equality is a standard formula for the sum of consecutive integers, with the positive part coming from noting that the initial sum is 0 if $d \leq a$ (i.e., if there are fewer arrivals than a, there is no backlog).

The law of total expectation dictates the expected backlog cost of retailer i for Poisson D with rate λ and inventory a is $\mathbb{E}[B] = \mathbb{E}\left[\mathbb{E}\left[B|D\right]\right]$. We take expectations of the individual parts of the backlog expression,

$$\begin{aligned}
&\frac{1}{2}\mathbb{E}D\mathbb{1}_{\{D\geq a\}} = \frac{1}{2}\sum_{k=a}^{\infty}\frac{k\lambda^k e^{-\lambda}}{k!} = \frac{\lambda}{2}\sum_{k=a}^{\infty}\frac{\lambda^{k-1}e^{-\lambda}}{(k-1)!} = \frac{\lambda}{2}\sum_{j=a-1}^{\infty}\frac{\lambda^j e^{-\lambda}}{j!} = \frac{\lambda}{2}\mathbb{P}\left(D\geq a-1\right),\\
&\frac{a(a+1)}{2}\mathbb{E}\frac{1}{D+1}\mathbb{1}_{\{D\geq a\}} = \frac{a(a+1)}{2}\sum_{k=a}^{\infty}\frac{\lambda^k e^{-\lambda}}{(k+1)k!} = \frac{a(a+1)}{2\lambda}\sum_{k=a}^{\infty}\frac{\lambda^{k+1}e^{-\lambda}}{(k+1)!} = \frac{a(a+1)}{2\lambda}\mathbb{P}\left(D\geq a+1\right),\\
&a\mathbb{E}\mathbb{1}_{\{D\geq a\}} = a\mathbb{P}\left(D\geq a\right).
\end{aligned}$$

Therefore the expected backlog cost for Poisson D with rate λ and inventory a is

$$\begin{aligned}
&bT\left(\frac{\lambda}{2}\mathbb{P}\left(D\geq a-1\right)+\frac{a(a+1)}{2\lambda}\mathbb{P}\left(D\geq a+1\right)-a\mathbb{P}\left(D\geq a\right)\right),\\
&= bT\left(\left(\frac{\lambda^2+a(a+1)-2a\lambda}{2\lambda}\right)\mathbb{P}\left(D\geq a\right)+\frac{\lambda}{2}\mathbb{P}\left(D=a-1\right)-\frac{a(a+1)}{2\lambda}\mathbb{P}\left(D=a\right)\right),\\
&= bT\left(\left(\frac{(\lambda-a)^2+a}{2\lambda}\right)\mathbb{P}\left(D\geq a\right)+\frac{\lambda-a-1}{2}\mathbb{P}\left(D=a-1\right)\right).
\end{aligned}\tag{21}$$

For the decentralized system, the total system cost is simply the sum of n retailers backlog cost in (21), which completes the proof of the second part of Lemma 3. Next, for the centralized system, it sees aggregate Poisson demand $D^n = \sum_i D_i$ with rate $n\lambda$ and has total inventory of na, so its expected backlog cost is $bT\left(\left(\frac{n(\lambda-a)^2+a}{2\lambda}\right)\mathbb{P}\left(D^n\geq na\right)+\frac{n(\lambda-a)-1}{2}\mathbb{P}\left(D^n=na-1\right)\right)$.

Then, the expected fulfillment cost (with unit cost f) for the centralized system is

$$\begin{aligned}
f\mathbb{E}\min(D^n,na) &= f\left(\mathbb{E}D\mathbb{1}_{\{D^n\leq na\}}+\mathbb{E}a\mathbb{1}_{\{D^n>na\}}\right),\\
&= f\left(\sum_{k=0}^{na}\frac{k(n\lambda)^k e^{-(n\lambda)}}{k!}+a\mathbb{P}\left(D^n>na\right)\right),\\
&= f\left((n\lambda)\sum_{k=1}^{na}\frac{(n\lambda)^{k-1}e^{-(n\lambda)}}{(k-1)!}+a\mathbb{P}\left(D^n>na\right)\right),\\
&= f\left((n\lambda)\sum_{j=0}^{na-1}\frac{(n\lambda)^j e^{-(n\lambda)}}{j!}+a\mathbb{P}\left(D^n>na\right)\right),\\
&= fn\left(\lambda\mathbb{P}\left(D^n<na\right)+a\mathbb{P}\left(D^n>na\right)\right).
\end{aligned}$$

Adding the expedited fulfillment cost to the backlog cost completes the proof of the first part of Lemma 3. Q.E.D.

Next, in order to prove Lemma 1 from Section 4, we first define a few gamma functions and relate them to the Poisson distribution. The gamma function and lower incomplete gamma function are defined for real α and β as $\Gamma(\alpha) = \int_0^\infty t^{\alpha-1}e^{-t}dt$, $\gamma(\alpha,x) = \int_0^x t^{\alpha-1}e^{-t}dt$. For an integer $\alpha\geq 1$ the gamma function satisfies $\Gamma(\alpha) = (\alpha-1)!$ (Press et al., 1992). Following Press et al. (1992), the reverse cumulative distribution function of a Poisson with rate λ can be written in terms of the gamma functions as follows $\mathbb{P}\left(X_i\geq a\right) = \frac{\gamma(a,\lambda)}{\Gamma(a)}$. Therefore, noting that Lemma 1 is trivially true when $a=0$ (since both probabilities equal 1), we can prove the lemma by proving the following for $a\geq 1$ and $\lambda\geq a+1$:

$$\frac{\gamma(na,n\lambda)}{\Gamma(na)}\geq\frac{\gamma(a,\lambda)}{\Gamma(a)}.\tag{22}$$

To the best of our knowledge, this analysis of the incomplete gamma function is novel, as the existing literature proves the inequality in the opposite direction for the case when $\lambda = a$ (Van der Vaart, 1961), and thus the analysis here may be of independent interest. The main idea of the analysis is to show that when λ is strictly larger than a, scaling both λ and a increases the Poisson tail probability, and the requirement that $\lambda\geq a+1$ appears tight when a grows large. The key building block to prove this result is the following bound comparing the integrand to a partial integral involved in the gamma function.

Lemma 4. *For positive integers $\alpha,\beta\in\mathbb{Z}_+$ such that $\alpha\geq\beta\geq 1$, and a positive real $x\in\mathbb{R}_+$ such that $x\geq\beta+1$, we have*

$$\left(\frac{x\alpha}{\beta}\right)^\alpha e^{-\frac{xa}{\beta}}\leq\int_{\frac{x\alpha}{\beta}}^{\frac{x(\alpha+1)}{\beta}}t^\alpha e^{-t}dt.$$

Proof. First, for generic limits of integration $y<z$, integration by parts gives $\int_y^z t^\alpha e^{-t}dt = y^\alpha e^{-y}-z^\alpha e^{-z}+\alpha\int_y^z t^{\alpha-1}e^{-t}dt$, from which it is clear that $y^\alpha e^{-y}\leq\int_y^z t^\alpha e^{-t}dt$ is equivalent to $z^\alpha e^{-z}\leq\alpha\int_y^z t^{\alpha-1}e^{-t}dt$. Thus, to prove the lemma we prove the equivalent claim that

$$\left(\frac{x(\alpha+1)}{\beta}\right)^\alpha e^{-\frac{x(\alpha+1)}{\beta}}\leq a\int_{\frac{xa}{\beta}}^{\frac{x(\alpha+1)}{\beta}}t^{\alpha-1}e^{-t}dt.\tag{23}$$

To prove this, we first establish a few facts about the function $f(t) = t^{\alpha-1}e^{-t}$. First, we show that $f(t)$ is decreasing in t for $t\geq\alpha-1$. To see this consider the derivative $f'(t) = t^{\alpha-1}e^{-t}(\frac{\alpha-1}{t}-1)$, which is negative when $t\geq\alpha-1$.

Next, we show for $\alpha-1 \le t \le \alpha-1+\sqrt{\alpha-1}$ that $f(t)$ is concave, and for $t \ge \alpha-1+\sqrt{\alpha-1}$, $f(t)$ is convex. To see this, consider the second derivative $f''(t) = t^{\alpha-3}e^{-t}(t^2-2(\alpha-1)t+(\alpha-1)^2-(\alpha-1))$, whose sign is determined by the polynomial $t^2-2(\alpha-1)t+(\alpha-1)^2-(\alpha-1)$ since $t^{\alpha-3}e^{-t}$ is always non-negative for $t \ge 0$. The polynomial $t^2-2(\alpha-1)t+(\alpha-1)^2-(\alpha-1)$ is a convex quadratic (since the leading coefficient is positive) and hence is negative between its roots and positive otherwise. The two roots are given by the quadratic formula $t = \alpha-1 \pm \sqrt{\alpha-1}$, so for $\alpha-1 \le t \le \alpha-1+\sqrt{\alpha-1}$ we have $f''(t) \le 0$ and hence $f(t)$ is concave, while for $t \ge \alpha-1+\sqrt{\alpha-1}$ we have $f''(t) \ge 0$ and hence $f(t)$ is convex.

Next, we progress toward proving (23) by proving a lower bound on the integrand $f(t) = t^{\alpha-1}e^{-t}$ between generic limits of integration $y < z$ with $y \ge \alpha-1$. To do this we define two auxiliary functions $g(t) = \frac{(f(z)-f(y))t+f(y)z-yf(z)}{z-y}$, and $h(t) = f(z)+f'(z)(t-z)$, where $g(t)$ represents the secant line to f through y and z, and $h(t)$ represents the tangent line to f at z. Then we claim that for $t \in [y,z]$ we have

$$f(t) \ge \min(g(t), h(t)). \tag{24}$$

First, we note that if $z \le \alpha-1+\sqrt{\alpha-1}$ then $f(t)$ is concave for all $t \in [y,z]$ and so $f(t) \ge g(t)$ follows directly, while if $y \ge \alpha-1+\sqrt{\alpha-1}$ then $f(t)$ is convex for all $t \in [y,z]$ and so $f(t) \ge h(t)$ follows directly.

Thus, it remains to consider the case $y \le \alpha-1+\sqrt{\alpha-1} \le z$. For $t \in [\alpha-1+\sqrt{\alpha-1}, z]$, the convexity of f implies $f(t) \ge h(t)$, so the claim holds in this region. Next consider the region $[y, \alpha-1+\sqrt{\alpha-1}]$ and consider two cases. First, consider if $h(y) \le f(y)$, and recall that the convexity of f in the region $[\alpha-1+\sqrt{\alpha-1}, z]$ means that $h(\alpha-1+\sqrt{\alpha-1}) \le f(\alpha-1+\sqrt{\alpha-1})$. Thus, the concavity of f in the region $[y, \alpha-1+\sqrt{\alpha-1}]$ implies that $h(t) < f(t)$ for all t in this region, and the claim holds. Otherwise, if $h(y) \ge f(y)$, recall that $g(y) = f(y)$ by definition, so $h(y) \ge g(y)$. Further, since $h(z) = g(z)$ by definition, we must have $h(t) \ge g(t)$ for all $t \le z$ (since h and g are lines in the plane, which can only cross once). Therefore, we have $g(\alpha-1+\sqrt{\alpha-1}) \le h(\alpha-1+\sqrt{\alpha-1}) \le f(\alpha-1+\sqrt{\alpha-1})$, and thus, by the concavity of f we have $g(t) \le f(t)$ for all $t \in [y, \alpha-1+\sqrt{\alpha-1}]$. This completes the proof of the claim in (24).

From (24) we have that $\int_y^z f(t)dt \ge \int_y^z \min(g(t),h(t))dt = \min\left(\int_y^z g(t)dt, \int_y^z h(t)dt\right)$, where the equality follows since we have $h(z) = g(z)$ by definition, so we either have $h(t) \ge g(t)$ for all $t \in [y,z]$ or $h(t) \le g(t)$ for all $t \in [y,z]$, since h and g are lines in the plane and can only cross once. Thus, we will establish (23) by proving it is true when we replace the integrand on the right hand side with either lower bound of h or g.

For the lower bound h, after evaluating the integral (23) becomes $\left(\frac{x(\alpha+1)}{\beta}\right)^{\alpha} e^{-\frac{x(\alpha+1)}{\beta}} \le \frac{\alpha x}{2\beta}\left(1+\frac{x}{\beta}+\frac{2}{\alpha+1}\right)\left(\frac{x(\alpha+1)}{\beta}\right)^{\alpha-1} e^{-\frac{x(\alpha+1)}{\beta}}$, from which the exponential terms can be canceled on each side, as well as an $(x/\beta)^{\alpha}(\alpha+1)^{\alpha-1}$ term, to give the equivalent expression $\alpha+1 \le \frac{\alpha}{2}\left(1+\frac{x}{\beta}+\frac{2}{\alpha+1}\right)$. Subtracting α from both sides and combining the fractions on the right gives the equivalent expression $1 \le \frac{\alpha((x-\beta)(\alpha+1)+2\beta)}{2\beta(\alpha+1)}$, and multiplying both sides by $2\beta(\alpha+1)$ and canceling common terms gives the equivalent expression $2\beta \le \alpha(x-\beta)(\alpha+1)$, which follows from $x \ge \beta+1$, $\alpha \ge \beta$, and $\alpha+1 \ge \beta+1 \ge 2$, and thus completes the verification of (23) for the lower bound h.

For the lower bound g, after evaluating the integral (23) becomes $\left(\frac{x(\alpha+1)}{\beta}\right)^{\alpha} e^{-\frac{x(\alpha+1)}{\beta}} \le \frac{\alpha x}{2\beta}\left(\left(\frac{xa}{\beta}\right)^{\alpha-1} e^{-\frac{xa}{\beta}} + \left(\frac{x(\alpha+1)}{\beta}\right)^{\alpha-1} e^{-\frac{x(\alpha+1)}{\beta}}\right)$, from which the $(x/\beta)^{\alpha}$ terms cancel on each side, and multiplying each side by the exponential term $e^{\frac{x(\alpha+1)}{\beta}}$ gives the equivalent expression $(\alpha+1)^{\alpha} \le \frac{\alpha}{2}\left(\alpha^{\alpha-1}e^{\frac{x}{\beta}} + (\alpha+1)^{\alpha-1}\right)$. Multiplying by 2 on each side and collecting the $\alpha+1$ terms on the left hand side gives the equivalent expression $\left(2-\frac{\alpha}{\alpha+1}\right)(\alpha+1)^{\alpha} \le \alpha^{\alpha}e^{\frac{x}{\beta}}$, and rearranging the leading fraction and dividing by α^{α} gives the equivalent expression $\left(\frac{\alpha+2}{\alpha+1}\right)\left(1+\frac{1}{\alpha}\right)^{\alpha} \le e^{\frac{x}{\beta}}$. Here we note $x \ge \beta+1$ and $\alpha \ge \beta$ imply $1+\frac{1}{\alpha} \le \frac{x}{\beta}$, so the inequality is satisfied if we have $\left(\frac{\alpha+2}{\alpha+1}\right)\left(1+\frac{1}{\alpha}\right)^{\alpha} \le e^{1+\frac{1}{\alpha}}$, which follows from the standard inequality $1+\frac{1}{\alpha} \le e^{\frac{1}{\alpha}}$ (i.e., the elementary inequality $1+t \le e^t$ for $t = 1/a$). To see this note that $1+\frac{1}{\alpha} \le e^{\frac{1}{\alpha}}$ implies $\left(1+\frac{1}{\alpha}\right)^{\alpha} \le e$, and also implies $e^{\frac{1}{\alpha}} \ge 1+\frac{1}{\alpha} = \frac{\alpha+1}{\alpha} \ge \frac{\alpha+2}{\alpha+1}$. Putting these two inequalities together gives $\left(\frac{\alpha+2}{\alpha+1}\right)\left(1+\frac{1}{\alpha}\right)^{\alpha} \le e\left(e^{\frac{1}{\alpha}}\right) = e^{1+\frac{1}{\alpha}}$, which verifies (23) for the lower bound g, and thus completes the proof. Q.E.D.

Using Lemma 4, we can then prove the following recursive bound for the incomplete gamma function.

Lemma 5. *For positive integers $\alpha, b \in \mathbb{Z}_+$ such that $\alpha \ge b \ge 1$, and a positive real $x \in \mathbb{R}_+$ such that $x \ge \beta+1$, we have*

$$a\gamma\left(\alpha, x\frac{\alpha}{\beta}\right) \le \gamma\left(\alpha+1, x\left(\frac{\alpha+1}{\beta}\right)\right).$$

Proof. Integration by parts on the integral representation of the incomplete gamma function gives the following

$$\alpha\gamma\left(\alpha, x\frac{\alpha}{\beta}\right) = \left(\frac{xa}{\beta}\right)^{\alpha} e^{-\frac{xa}{\beta}} + \int_0^{\frac{xa}{\beta}} t^{\alpha}e^{-t}dt, \le \int_{\frac{xa}{\beta}}^{\frac{x(\alpha+1)}{\beta}} t^{\alpha}e^{-t}dt + \int_0^{\frac{xa}{\beta}} t^{\alpha}e^{-t}dt = \gamma\left(\alpha+1, x\left(\frac{\alpha+1}{\beta}\right)\right),$$

where the inequality follows from Lemma 4. Q.E.D.

With Lemma 5, the proof of Lemma 1 now follows with a simple recursion.

Proof of Lemma 1. We prove (22) by establishing the following claim through induction on the integers $j \ge 0$

$$\frac{\Gamma(a+j)}{\Gamma\alpha(a)}\gamma(a,\lambda) \le \gamma\left(a+j, \lambda\left(\frac{a+j}{a}\right)\right), \tag{25}$$

from which (22) follows by letting $j=(n-1)a$. For $j=0$, (25) trivially holds with equality, while for $j=1$ by Lemma 5, letting $\alpha=\beta=a$ and $x=\lambda$ we have $a\gamma(a,\lambda)\le\gamma\left(a+1,\lambda\left(\frac{a+1}{a}\right)\right)$, from which (25) follows since $a=\Gamma(a+1)/\Gamma(a)$. Now assume (25) holds for $j\ge1$ and consider $j+1$, for which we have:

$$\frac{\Gamma(a+j+1)}{\Gamma\alpha(a)}\gamma(a,\lambda)=\frac{(a+j)\Gamma(a+j)}{\Gamma(a)}\gamma(a,\lambda)\le(a+j)\gamma\left(a+j,\lambda\left(\frac{a+j}{a}\right)\right)\le\gamma\left(a+j+1,\lambda\left(\frac{a+j+1}{a}\right)\right),$$

where the first equality follows from $\Gamma(a+j+1)=(a+j)\Gamma(a+j)$, the inequality follows from the induction hypothesis, and the second equality follows from Lemma 5. Q.E.D.

Now, following Lemma 1, we divide our analysis of decentralized and centralized costs below into the two cases of $\lambda\ge a+1$ and $\lambda\le a+1$.

Lemma 6. *For Poisson rate $a\le\lambda\le a+1$, $n\ge2$ retailers, the centralized system cost is lower bounded by*

$$bn\tau\frac{(\lambda-a)^2}{4\lambda}+fna(1-2e^{-2}),\tag{26}$$

while the decentralized system cost is upper bounded by

$$bn\tau\left(1-e^{-2}\right)\left(\frac{a+(\lambda-a)^2}{2\lambda}\right).\tag{27}$$

Proof. For $a\le\lambda\le a+1$, the centralized system cost in (19) can be lower bounded by

$$\begin{aligned}&bT\left(\frac{n(\lambda-a)^2+a}{2\lambda}\mathbb{P}\left(D^n\ge na\right)+\frac{n(\lambda-a)-1}{2}\mathbb{P}\left(D^n=na-1\right)\right)+fn\left(\lambda\mathbb{P}\left(D^n<na\right)+a\mathbb{P}\left(D^n>na\right)\right)\\ \ge\ &bT\left(\frac{n(\lambda-a)^2+a}{2\lambda}\mathbb{P}\left(D^n\ge na\right)-\frac{1}{2}\mathbb{P}\left(D^n=na-1\right)\right)+fn\left(\lambda\mathbb{P}\left(D^n<na\right)+a\mathbb{P}\left(D^n>na\right)\right),\end{aligned}$$

where the inequality follows because $\lambda\ge a$. Observe that $\mathbb{P}\left(D^n=na-1\right)=\frac{(n\lambda)^{na-1}}{(na-1)!e^{n\lambda}}=\frac{na}{n\lambda}\frac{(n\lambda)^{na}}{(na)!e^{n\lambda}}=\frac{a}{\lambda}\mathbb{P}\left(D^n=na\right)$. Then, $bT\left(\frac{a}{2\lambda}\mathbb{P}\left(D^n\ge na\right)-\frac{1}{2}\mathbb{P}\left(D^n=na-1\right)\right)=bT\left(\frac{a}{2\lambda}\mathbb{P}\left(D^n>na\right)\right)\ge0$. Therefore, the centralized system cost can be lower bounded by

$$\begin{aligned}bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}\left(D^n\ge na\right)+fn\left(\lambda\mathbb{P}\left(D^n<na\right)+a\mathbb{P}\left(D^n>na\right)\right)&\ge bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}\left(D^n\ge na\right)+fna\left(\mathbb{P}\left(D^n<na\right)+\mathbb{P}\left(D^n>na\right)\right)\\&=bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}\left(D^n\ge na\right)+fna\left(1-\mathbb{P}\left(D^n=na\right)\right),\end{aligned}$$

where the first inequality follows from $\lambda\ge a$. We introduce the following lemmas to lower bound the remaining terms.

Lemma 7. *For all $a\ge1$ and for any $\lambda\ge a$, $\mathbb{P}\left(D_1\ge a\right)\ge\frac{1}{2}$.*

Proof. For a fixed rate $\hat{\lambda}$ and for all $a\le\hat{\lambda}$, $\mathbb{P}\left(D_1\ge a\right)\ge\mathbb{P}\left(D_1\ge\hat{\lambda}\right)$ by the definition of reverse cumulative distribution function. Therefore, it suffices to show that for $\lambda=a$, $\mathbb{P}\left(D_1\ge\lambda\right)\ge\frac{1}{2}$, which is true, shown in (5.5) in Van der Vaart (1961). Q.E.D.

As an immediate consequence of Lemma 7, $\mathbb{P}\left(D^n\ge na\right)\ge\frac{1}{2}$. Next, we want to upper bound the term $\mathbb{P}\left(D^n=na\right)$.

Lemma 8. *$\mathbb{P}\left(D^n=na\right)$ for D^n with rate $n\lambda$, where $\lambda\ge a\ge1$, is decreasing in λ.*

Proof. With rate $n\lambda$ $\mathbb{P}\left(D^n=na\right)=\frac{(n\lambda)^{na}e^{-n\lambda}}{(na)!}=\frac{n^{na}}{(na)!}\lambda^{na}e^{-n\lambda}$. Taking the derivative with respect to λ yields $\left[\frac{n^{na}}{(na)!}\lambda^{na}e^{-n\lambda}\right]'=\frac{n^{na}}{(na)!}\left[ne^{-n\lambda}\lambda^{na-1}\left(a-\lambda\right)\right]\le0$ because $\lambda\ge a$. Q.E.D.

Lemma 8 shows that for fixed n,a, the probability mass function $\mathbb{P}\left(D^n=na\right)$ is greater when λ is closer to a. Therefore, if $a\le\lambda\le a+1$, we set $\lambda=a$ to upper bound $\mathbb{P}\left(D^n=na\right)$.

Lemma 9. *$\mathbb{P}\left(D^n=na\right)$ with rate na is decreasing in na, where $n\ge2$, $a\ge1$, and n and a are integers.*

Proof. Let $x=na$. We claim that $\mathbb{P}\left(X=x\right)$ where X is a Poisson random variable with rate x is decreasing in x, where x is an integer. By the Poisson probability mass function, $\mathbb{P}\left(X=x\right)=\frac{x^x}{x!e^x}$, this is equivalent to showing, $\frac{(x+1)^{x+1}}{(x+1)!e^{x+1}}-\frac{x^x}{x!e^x}<0$. The left-hand-side is equivalent to $\frac{1}{x!e^x}\left[\frac{(x+1)^x}{e}-x^x\right]$. Following Taylor series expansion of $e^y>1+y$, we have $1+\frac{1}{x}<e^{\frac{1}{x}}\implies x+1<xe^{\frac{1}{x}}\implies(x+1)^x<x^xe\implies\frac{(x+1)^x}{e}<x^x$. Hence, $\frac{(x+1)^x}{e}-x^x<0$, completing the proof. Q.E.D.

Lemma 9 further implies that we want to set na to its minimum value which happens at $n=2$ and $a=1$. Hence, for $a\le\lambda\le a+1$,

$$bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}\left(D^n\ge na\right)+fna\left(1-\mathbb{P}\left(D^n=na\right)\right)\ge bnT\frac{(\lambda-a)^2}{2\lambda}\frac{1}{2}+fna(1-\frac{2^2}{2!e^2})=bnT\frac{(\lambda-a)^2}{2\lambda}\frac{1}{2}+fna(1-2e^{-2}).$$

Next, for $a\le\lambda\le a+1$ the decentralized system cost in (20) can be upper bounded by

$$bnT\left(\frac{(\lambda-a)^2+a}{2\lambda}\mathbb{P}\left(D_1\ge a\right)+\frac{\lambda-a-1}{2}\mathbb{P}\left(D_1=a-1\right)\right)\le bnT\left(\frac{(\lambda-a)^2+a}{2\lambda}\mathbb{P}\left(D_1\ge a\right)\right),$$

where the inequality follows from the fact that $\lambda-a-1\le0$. Next, for $a<\lambda\le a+1$, observe that $\mathbb{P}\left(D_1\ge a\right)$ is maximized at $\lambda=a+1$. The following lemma helps upper bound $\mathbb{P}\left(D_1\ge a\right)$.

Lemma 10. *$\mathbb{P}(D_1 \geq a)$ with rate $a+1$ is decreasing in a for $a \geq 1$.*

Proof. Using the gamma and lower incomplete gamma functions, the reverse cumulative distribution funciton of a Poisson with rate $\lambda = a+1$ can be written as $\mathbb{P}(D_1 \geq a) = \frac{\gamma(a,a+1)}{\Gamma(a)}$, where $\gamma(a,a+1) = \int_0^{a+1} t^{a-1}e^{-t}dt$. We show that for all $a \geq 1$, $\frac{\gamma(a,a+1)}{\Gamma(a)} > \frac{\gamma(a+1,a+2)}{\Gamma(a+1)}$. We first show that for $t > a$, function $e^{-t}t^a$ is decreasing in t. Observe that

$$\frac{a}{at}\left[e^{-t}t^a\right] = e^{-t}t^{a-1}(a-t) < 0. \tag{28}$$

Then, (28) implies $e^{-(a+1)}(a+1)^a > e^{-t}t^a$, for all $t \in [a+1, a+2]$. Therefore, integrating both sides of the equation with respect to t, we obtain $e^{-(a+1)}(a+1)^a > \int_{a+1}^{a+2} e^{-t}t^a dt$. Next, integration by parts gives $a\int_0^{a+1} e^{-t}t^{a-1}dt = e^{-(a+1)}(a+1)^d + \int_0^{a+1} e^{-t}t^a dt > \int_{a+1}^{a+2} e^{-t}t^a dt + \int_0^{a+1} e^{-t}t^a dt = \int_0^{a+2} e^{-t}t^a dt$. Therefore, by definition of the incomplete gamma function, $a\gamma(a,a+1) > \gamma(a+1,a+2)$, which implies $\frac{\gamma(a,a+1)}{\Gamma(a)} > \frac{\gamma(a+1,a+2)}{\Gamma(a+1)}$. Q.E.D.

By Lemma 10, $\mathbb{P}(D_1 \geq a)$ with rate $\lambda = a+1$ is decreasing in a, so $\mathbb{P}(D_1 \geq a)$ is maximized at $a = 1$. Therefore, $\mathbb{P}(D_1 \geq a) \leq 1 - e^{-2}$, and the decentralized system cost is upper bounded by $bnT\left(\frac{(\lambda-a)^2+a}{2\lambda}(1-e^{-2})\right)$, which completes the proof of Lemma 6. Q.E.D.

Next, for the other case $\lambda \geq a+1$, we have the following result.

Lemma 11. *For Poisson rate $\lambda \geq a+1$, $n \geq 2$ retailers, the centralized system cost is lower bounded by*

$$bn\tau\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D_1 \geq a) + fna(1 - \frac{32}{3}e^{-4}), \tag{29}$$

while the decentralized system cost is upper bounded by

$$bn\tau\left(\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D_1 \geq a) + \frac{a}{2\lambda} + \frac{1}{2\sqrt{2\pi e}}\right). \tag{30}$$

Proof. For $\lambda \geq a+1$, by following the same steps in proof of Lemma 6, the centralized system cost in (19) can be lower bounded by $bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D^n \geq na) + fna(1 - \mathbb{P}(D^n = na))$. By Lemma 1, $\mathbb{P}(D^n \geq na) \geq \mathbb{P}(D_1 \geq a)$ for $\lambda \geq a+1$. Next, we upper bound $\mathbb{P}(D^n = na)$. By Lemma 8, if $a+1 \leq \lambda$, we set $\lambda = a+1$ to upper bound $\mathbb{P}(D^n = na)$. Then, observe that with a Poisson distribution with rate $n(a+1)$, $\mathbb{P}(D^n = na) \leq \mathbb{P}(D^n = n(a+1)) \leq \max_{n(a+1)} \mathbb{P}(D^n = n(a+1))$. Lemma 9 shows that $\mathbb{P}(D^n = n(a+1))$ is maximized at the minimal value of $n(a+1)$, which is at $n = 2$ and $a+1 = 2$. Hence, for $\lambda \geq a+1$, the centralized system cost is lower bounded by $bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D^n \geq na) + fna(1 - \mathbb{P}(D^n = na)) \geq bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D_1 \geq a) + fna(1 - \frac{4^4}{4!e^4}) = bnT\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D_1 \geq a) + fna(1 - \frac{32}{3}e^{-4})$. Next, for $\lambda \geq a+1$, the decentralized cost in (20) can be upper bounded by $bnT\left(\frac{(\lambda-a)^2+a}{2\lambda}\mathbb{P}(D_1 \geq a) + \frac{\lambda-a-1}{2}\mathbb{P}(D_1 = a-1)\right) \leq bnT\left(\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D_1 \geq a) + \frac{a}{2\lambda} + \frac{\lambda-a-1}{2}\frac{\lambda^{a-1}}{(a-1)!e^\lambda}\right)$ where $\mathbb{P}(D_1 \geq a)$ in the second term is upper bounded by 1, and $\mathbb{P}(D_1 = a-1) = \frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ is the Poisson probability mass function in the last term. The lemma below is helpful, showing $\frac{\lambda-a-1}{2}\frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ can be bounded by a constant.

Lemma 12. *For $\lambda \geq a \geq 1$, $\frac{\lambda-a-1}{2}\frac{\lambda^{a-1}}{(a-1)!e^\lambda} \leq \frac{1}{2\sqrt{2\pi e}}$.*

Proof. Using the standard Stirling's approximation (Dutkay et al., 2013), we lower bound the factorial $(a-1)! \geq \sqrt{2\pi(a-1)}(\frac{a-1}{e})^{a-1}$, and as a consequence, $\frac{\lambda-a-1}{2}\frac{\lambda^{a-1}}{(a-1)!e^\lambda}$ can be upper bounded by

$$\frac{1}{2\sqrt{2\pi}}\frac{\lambda-a-1}{\sqrt{(a-1)}}\left(\frac{\lambda}{a-1}\right)^{a-1}e^{-(\lambda-a+1)}. \tag{31}$$

Since the function in (31) is divergent for $a = 1$. We divide this proof into two cases. We first show for large a, (31) can be upper bounded by a constant. Then, we show for small a, the same constant remains as a valid upper bound. The following two lemmas help show that (31) is upper bounded by a constant for all a.

Lemma 13. *The function in* (31) *is maximized at $\lambda^* = a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}$, for $\lambda \geq a \geq 1$.*

Proof. The partial derivative with respect to λ is

$$\frac{\partial}{\partial\lambda}\left[\frac{1}{2\sqrt{2\pi}}\frac{\lambda-a-1}{\sqrt{(a-1)}}\left(\frac{\lambda}{a-1}\right)^{a-1}e^{-(\lambda-a+1)}\right] = \frac{1}{2\sqrt{2\pi}(a-1)^{d-\frac{1}{2}}e^{\lambda-a+1}}\lambda^{d-2}\left[-\lambda^2 + (2a+1)\lambda + 1 - a^2\right],$$

which setting equal to zero yields critical points $\lambda_1 = a + \frac{1}{2} - \sqrt{a + \frac{5}{4}}$ and $\lambda_2 = a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}$. Observe that λ_1 is not in the domain $\lambda \geq a$ because $\lambda_1 < a + \frac{1}{2} - \sqrt{1 + \frac{5}{4}} = d - 1 < d$. λ_2, therefore, is the only candidate. Next, observe that the coefficient $\frac{1}{2\sqrt{2\pi}(a-1)^{d-\frac{1}{2}}e^{\lambda-a+1}}\lambda^{d-2}$ is non-negative, and the leading coefficient of the quadratic term λ^2 is negative. This implies that the derivative is positive between λ_1 and λ_2 and negative before λ_1 and after λ_2. Therefore, λ_2 is a maximum of the function over the range $\lambda >= a$. Q.E.D.

Next, We substitute $\lambda^* = a + \frac{1}{2} + \sqrt{a + \frac{5}{4}}$ in (31).

Lemma 14. *For $a \geq 5$,*

$$\frac{1}{2\sqrt{2\pi}} \frac{\sqrt{a+\frac{5}{4}}-\frac{1}{2}}{\sqrt{(a-1)}} \left(\frac{a+\frac{1}{2}+\sqrt{a+\frac{5}{4}}}{a-1}\right)^{a-1} e^{-(\frac{3}{2}+\sqrt{a+\frac{5}{4}})} \leq \frac{1}{2\sqrt{2\pi e}}. \tag{32}$$

Proof. We first show $\frac{\sqrt{a+\frac{5}{4}}-\frac{1}{2}}{\sqrt{(a-1)}} \leq 1$ for $a \geq 5$. Setting $\frac{\sqrt{a+\frac{5}{4}}-\frac{1}{2}}{\sqrt{(a-1)}} \leq 1$, we obtain $\sqrt{a+\frac{5}{4}}-\frac{1}{2} \leq \sqrt{a-1} \implies a+\frac{6}{4}-\sqrt{a+\frac{5}{4}} \leq a-1 \implies 5 \leq a$. Because $\frac{\sqrt{a+\frac{5}{4}}-\frac{1}{2}}{\sqrt{(a-1)}} \leq 1$ for large a, what is left is to show $\left(\frac{a+\frac{1}{2}+\sqrt{a+\frac{5}{4}}}{a-1}\right)^{a-1} e^{-(\frac{3}{2}+\sqrt{a+\frac{5}{4}})} \leq \frac{1}{\sqrt{e}}$. Equivalently, we want to show

$$\frac{a+\frac{1}{2}+\sqrt{a+\frac{5}{4}}}{a-1} \leq e^{\frac{1+\sqrt{a+\frac{5}{4}}}{a-1}}. \tag{33}$$

Taylor series expansion of e^x gives that $e^x \geq 1+x+\frac{x^2}{2}$ for $x \geq 0$. Substituting $x = \frac{1+\sqrt{a+\frac{5}{4}}}{a-1}$, we obtain

$$e^{\frac{1+\sqrt{a+\frac{5}{4}}}{a-1}} \geq 1+\frac{1+\sqrt{a+\frac{5}{4}}}{a-1}+\frac{(\frac{1+\sqrt{a+\frac{5}{4}}}{a-1})^2}{2} = \frac{2a^2+2a\sqrt{a+\frac{5}{4}}-a+\frac{9}{4}}{2(a-1)^2} \geq \frac{2a^2+2(a-2)\sqrt{a+\frac{5}{4}}-a-1}{2(a-1)^2} = \frac{a+\frac{1}{2}+\sqrt{a+\frac{5}{4}}}{a-1}$$

where the last equality follows from dividing both the numerator and denominator by $2(a-1)$. Q.E.D.

Lemma 14 implies that for $a \geq 5$, $\frac{\lambda-a-1}{2}\frac{\lambda^{a-1}}{(a-1)!e^{\lambda}} \leq \frac{1}{2\sqrt{2\pi e}} \approx 0.121$. We numerically validate that this upper bound holds for $4 \geq a \geq 1$. We set $\lambda(a) = a+\frac{1}{2}+\sqrt{a+\frac{5}{4}}$ and evaluate expression $f(a) = \frac{\lambda(a)-a-1}{2}\frac{\lambda(a)^{a-1}}{(a-1)!e^{\lambda(a)}}$ for small a: $f(1) = 0.0249$, $f(2) = 0.0379$, $f(3) = 0.0464$, $f(4) = 0.0525$. The maximum value is $0.0525 < \frac{1}{2\sqrt{2\pi e}}$, completing the proof of Lemma 12. Q.E.D.

Lemma 12 implies that, for $a+1 \leq \lambda$, the decentralized system cost is upper bounded by $bnT\left(\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D_1 \geq a)+\frac{a}{2\lambda}+\frac{1}{2\sqrt{2\pi e}}\right)$, which completes the proof of Lemma 11. Q.E.D.

With the bounds of Lemmas 6 and 11 in hand, we now are ready to prove Theorem 2.

Proof of Theorem 2. For $a \leq \lambda \leq a+1$, setting (27) less than (26) yields

$$fa(1-2e^{-2}) \geq b\tau\left(\frac{a}{2\lambda}(1-e^{-2})+\frac{(\lambda-a)^2}{2\lambda}\left(1-e^{-2}-\frac{1}{2}\right)\right). \tag{34}$$

To derive a sufficient condition, we can further upper bound $\frac{(\lambda-a)^2}{2\lambda}$. Observe that the derivative of $\frac{(\lambda-a)^2}{2\lambda}$ with respect to λ is $\frac{1}{2}-\frac{a^2}{2\lambda^2} \geq 0$, for $a \leq \lambda \leq a+1$. Therefore, $\frac{(\lambda-a)^2}{2\lambda} \leq \frac{(a+1-a)^2}{2(a+1)} \leq \frac{1}{4}$, where the last inequality follows from $a \geq 1$. Hence, (34) is implied by the sufficient condition $fa(1-2e^{-2}) \geq b\tau\left(\frac{a}{2\lambda}(1-e^{-2})+\frac{1}{4}\left(1-e^{-2}-\frac{1}{2}\right)\right)$, which simplifies to

$$f \geq b\tau\left(\frac{1-e^{-2}}{2(1-2e^{-2})\lambda}+\frac{1}{8a}\right) \approx b\tau\left(\frac{0.593}{\lambda}+\frac{0.125}{a}\right). \tag{35}$$

For $\lambda \geq a+1$, observe that the term $bn\tau\frac{(\lambda-a)^2}{2\lambda}\mathbb{P}(D_i \geq a)$ cancels in (29) and (30). Then, setting (30) less than (29) yields $fa(1-\frac{32}{3}e^{-4}) \geq b\tau\left(\frac{a}{2\lambda}+\frac{1}{2\sqrt{2\pi e}}\right)$, and simplifying we obtain

$$f \geq b\tau\left(\frac{1}{2(1-\frac{32}{3}e^{-4})\lambda}+\frac{1}{2\sqrt{2\pi e}(1-\frac{32}{3}e^{-4})a}\right) \approx b\tau\left(\frac{0.621}{\lambda}+\frac{0.150}{a}\right). \tag{36}$$

Finally, observe that both constants above λ and a in (36) are greater than those in (35), so the cost condition in (36) is a sufficient condition for all $\lambda \geq a \geq 1$. This completes the proof. Q.E.D.

B.4. One RDC, one dealer in normal mode.

In this section we demonstrate in a simple model that even when the manufacturer's end customer is an independent dealer who also holds inventory, it remains optimal for the manufacturer to use a base-stock policy. This further suggests (in addition to our analysis of Section 3) that the class of base-stock policies are effective for the automaker's setting. The model we analyze in this section considers a single manufacturer location (the RDC) serving a single independent dealer. The dealer receives inventory from the RDC and serves external demand.

Let D_t denote the random external demand at the dealer on day t (we assume D_t is i.i.d. across t). Let $X_t^{\mathcal{D}}$ and $B_t^{\mathcal{D}}$ denote dealer's inventory and backlog levels at time t. To minimize its own costs, we assume the dealer orders according to a base-stock policy in every period, and let $S^{\mathcal{D}}$ denote its base-stock level. Let O_t and v_t denote the dealer's order to the RDC and its demand fulfillment at time t. Let w_t denote what the RDC sends to the dealer at the end of day t, and we assume the fulfillment w_t does not show up until the beginning of day $t+1$, at which point it clears the dealer's backlog. We assume there is no quicker expediting option here for simplicity. Let L denote the manufacturer's lead time for replenishing the RDC. We assume, for simplicity, that the manufacturer can fulfill any order from the RDC within the lead time (i.e., the manufacturer ordering policy from an external supplier is not our concern here). Finally, let X and B represent the RDC's inventory and backlog.

The evolution equations for the RDC and dealer are

$$X_t = X_{t-1} + x_{t-L} - w_t, \qquad B_t = B_{t-1} + O_t - w_t, \qquad X^{\mathcal{D}}_{t+1} = X^{\mathcal{D}}_t + w_t - v_{t+1}, \qquad B^{\mathcal{D}}_{t+1} = B^{\mathcal{D}}_t + D_{t+1} - v_{t+1}. \tag{37}$$

The objective we consider in this model consists of the RDC holding cost and two levels of backlog cost: unit cost b at the RDC level and $b^{\mathcal{D}}$ at the dealer level, respectively. This is to model the scenario in which the manufacturer often bears additional cost when a customer is actively waiting for a part at the dealer (eg. daily rental car cost). The manufacturer needs to select a policy π to minimize long run average costs

$$C_\pi = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T}\left(hX_t + bB_t + b^{\mathcal{D}}B^{\mathcal{D}}_t\right)\right]. \tag{38}$$

Next we show that the manufacturer can use a base-stock policy to minimize (38).

Theorem 3. *There exists a base-stock policy that is long-run optimal for Problem* (38).

We present the correct base-stock level X^* at the RDC below, which is an extension of the newsvendor formula.

$$\mathbb{P}\left(D^L \le X^*\right) = \frac{b + b^{\mathcal{D}}\mathbb{P}\left(D^L + D^1 \ge X^* + S^{\mathcal{D}},\ D^L \ge X^*\right)}{h+b},$$

where D^L is the random demand over the RDC's lead time, and D^1 is the one-day demand over the RDC's fulfillment time to the dealer. The left-hand-side, $\mathbb{P}(D^L \le X^*)$, represents the RDC's optimal service level for a lead time L, which would typically equal $b/(h+b)$ in a traditional newsvendor formula. The numerator on the right-hand-side, compared to the traditional newsvendor formula, has an extra penalty term with coefficient $b^{\mathcal{D}}$ which increases the service level due to the extra dealer's backlog cost. The coefficient $b^{\mathcal{D}}$ is moderated by the joint probability of two events. The first event, $D^L + D^1 \ge X^* + S^{\mathcal{D}}$, is intuitive (since it represents that the demand exceeds the supply of both the RDC and dealer, so the dealer carries backlog), but perhaps the second event $D^L \ge X^*$ is less intuitive. To interpret this, consider the complementary condition $D^L \le X^*$, which together with $D^L + D^1 \ge X^* + S^{\mathcal{D}}$ implies $D^1 \ge S^{\mathcal{D}}$. In this case the manufacturer should not take the dealer's backlog cost into consideration when designing its own base-stock policy, because $D^1 \ge S^{\mathcal{D}}$ implies the one-day demand exceeds the dealer's base-stock level, and the manufacturer cannot control this through its own policy. Despite this issue, Theorem 3 demonstrates that a base-stock policy can remain effective with an independent dealer, and we therefore focus on this policy class in our simulations. We omit the proof of Theorem 3 in the appendix due to the page limit constraint of the competition, but we provide a high-level outline of the proof here. We first formulate a stochastic program of this one-RDC-one-dealer problem by relaxing the time constraint. Then, we show that the optimal solution from this stochastic program provides a lower bound of Problem (38). Next, we characterize the optimal solution as well as the optimal objective value in tractable forms. Finally, we prove that, with a correct base-stock level, the cost of a base-stock policy, naturally serving as an upper bound of Problem (38), matches the lower bound from the stochastic program. This optimality result demonstrates that a base-stock policy is plausible in our industrial partner's setting.

Appendix C: More on Simulation Details

C.1. Details on Newsvendor-Based Policies

C.1.1. The FZ Heuristic. The FZ heuristic is a natural candidate policy in our setting because it prescribes a simple and intuitive base-stock policy (which fits naturally with our industrial partner's desire to minimize alterations to their operations), and it is straightforward to adapt to our setting with expediting and disruptions. In particular, when holding costs at the warehouse and retailers are the same and demand is Poisson (both of which we assume in our simulations), Federgruen and Zipkin (1984c); Gallego et al. (2007) recommend a simple heuristic for setting a system-wide base-stock level by solving a single newsvendor model with a normal demand approximation: $S^{\mathrm{FZ}} \in \arg\min_S \mathbb{E}\left[h(S-\tilde{D})^+ + b(\tilde{D}-S)^+\right]$, where $\tilde{D}$ is a normal random variable with mean $(L+l)\sum_i \mathbb{E}[D_i]$ and variance $L\sum_i \mathrm{Var}[D_i] + \left(\sum_i \sqrt{l\mathrm{Var}[D_i]}\right)^2$ (see detailed discussions in Federgruen and Zipkin, 1984c; Zipkin, 2000; Gallego et al., 2007 for further explanation of the demand approximation). The heuristic uses S^{FZ} as a system-wide base-stock level for ordering (i.e., orders to bring the warehouse echelon inventory position up to S^{FZ} in every period using the same logic as (8)). To adapt this policy to our setting, we use the same myopic expediting policy as prescribed for our SP^L policy described in Section 3.2 (as this provides a fair comparison between the policies): before allocation, myopically expedite to clear as much backlog as possible up to the warehouse on-hand inventory. Then, we allocate inventory to the retailers using the heuristic of Federgruen and Zipkin (1984c); Gallego et al. (2007): use a stochastic program to myopically minimize the expected holding and backlog cost when the allocation arrives at the retailers (i.e., after the lead time l), which in our context amounts to solving (9) with the retailer base-stock level constraints, $X_{i,t} + IT^-_{i,t} + z_{i,t} \le S_i$, removed. Gallego et al. (2007) note that if $h_i = h_0$ for all i (as is the case in our simulations), then this policy sends all inventory to the retailers and holds none centrally at the warehouse.

C.1.2. The NV Heuristic. The FZ heuristic of the previous section serves as a naive benchmark since its base-stock levels are not specifically designed to take into account expediting or disruption. The NV heuristic is more carefully calibrated to take both these features into account. The heuristic first solves for preliminary retailer base-stock levels using expediting as the cost of underage (since expediting is used first before demand is backlogged). It then solves for a warehouse base-stock level accounting for demand in excess of the retailer's preliminary base-stock levels, then ultimately resolves for final retailer base stock levels given the warehouse base stock level (assuming excess warehouse inventory is used for expediting). The intuition for this final resolve is to take the quantity of centralized inventory into account when deciding on retailer base-stock levels. This simple heuristic is easy to compute as it relies on newsvendor calculations (i.e., finding distribution percentiles). The heuristic's design also qualitatively improves upon the FZ heuristic by intentionally doing the following: 1) holding inventory at the central warehouse, 2) directly incorporating expediting costs, and 3) directly incorporating disruptions. To state the

Table 11 Average and worst case cost ratios of SP^{L}-N, SP^{L+l}-N, and SP^{avg}-N compared to SP^{min}-N

	SP^{L}-N/SP^{min}-N		SP^{L+l}-N/SP^{min}-N		SP^{avg}-N/SP^{min}-N	
Disr. Prob., α	Avg.	Max.	Avg.	Max.	Avg.	Max
0.005	1.48	1.91	1.01	1.08	1.08	1.29
0.01	1.01	1.04	1.09	1.24	1.03	1.09
0.02	1.01	1.02	1.04	1.13	1.01	1.05
Overall	1.17	1.91	1.05	1.24	1.04	1.29

heuristic, let $\check{D}_i = D^1_{i,l} + D^2_{i,l} + D^3_{i,l} + D^4_{i,l}$ denote the demand for retailer i over $L+l+\tau$ periods (recalling the definitions of $\mathbf{D}^1_{\hat{l}}$, $\mathbf{D}^2_{\hat{l}}$, $\mathbf{D}^3_{\hat{l}}$, $\mathbf{D}^4_{\hat{l}}$ for $\hat{l}=l$ from Section 3.1.1). Also we use $\tilde{b}=\sum_i b_i/n$ to denote the average backlog cost. Then, we determine the base-stock levels in an iterative fashion as follows (see Appendix C.2 for further intuitive explanations):

1. Calculate preliminary retailer i base stock level S'_i:

$$S'_i \in \arg\min_{S_i} \mathbb{E}[h_i(S_i - \check{D}_i)^+ + \frac{f_i}{l}(\check{D}_i - S_i)^+]. \tag{39}$$

2. Let $D_0 = \sum_i (\check{D}_i - S'_i)^+$, then calculate local warehouse base stock level S'_0:

$$S'_0 \in \arg\min_{S_0} \mathbb{E}[h_0(S_0 - D_0)^+ + \tilde{b}(D_0 - S_0)^+]. \tag{40}$$

3. Assume $f_1 - b_1 <= f_2 - b_2 \ldots$, iterate through $i = 1, 2, \ldots$ for final base stock level:

$$S_i \in \arg\min_{S_i} \mathbb{E}\left[h_i(S_i - \check{D}_i)^+ + \frac{f_i}{l}\min\left(\left(S_0 - \sum_{j<i}(\check{D}_j - S_j)^+\right)^+, (\check{D}_i - S_i)^+\right) + b_i\left(\check{D}_i - S_i - \left(S_0 - \sum_{j<i}(\check{D}_j - S_j)^+\right)^+\right)^+\right]. \tag{41}$$

The final base-stock levels we use from the heuristic are S_i for each retailer $i \geq 1$ and $S_0 = S'_0 + \sum_i S_i$ for the warehouse (recalling S'_0 represents the local base-stock level, while S_0 is the echelon base-stock level in line with the other heuristics in the paper).

C.2. Simulations of Section 5

Here we provide more details and context on the simulations of Section 5 for the base model. First, a brief comment on the appropriateness of the FZ heuristic for our setting is warranted. Despite its long history, the FZ remains one of the leading heuristics proposed in the distribution systems literature. Indeed, to the best of our knowledge, the most competitive heuristics for centralized control of a distribution system are the Lagrangian relaxation heuristics of Kunnumkal and Topaloglu (2008, 2011) and Federgruen et al. (2018) (which themselves are generalizations of the FZ relaxation heuristic). These heuristics suffer from a few draw-backs in our setting: i) they do not implement a straightforward base-stock policy, but instead solve dynamic programming sub-problems, ii) they are more computationally intensive (i.e., harder to implement, Kunnumkal and Topaloglu, 2011; Federgruen et al., 2018), and iii) they are developed for the finite horizon problem in the literature. Thus, to compare with a practically relevant policy for our industrial partner's setting, we focus on the FZ heuristic.

We next provide further intuitive explanation of the NV heuristic. Step 1 is similar to the first step of the RO heuristic in Rong et al. (2017), but uses the expediting cost instead of the backlog cost, as expediting is the system's first line of defense against underage. In step 2, D_0 represents the uncertain demand that the warehouse needs to plan for over $L+l+\tau$ periods (over and above the retailer's inventory positions). This is like RO step 2 from Rong et al. (2017), but takes into account the potential for disruption demand over a random τ periods (which the warehouse needs to cover). Finally, step 3 is motivated by RO step 3 in Rong et al. (2017) since it resolves the retailer's local base stock level assuming the warehouse's base stock level, but it does it in an iterative fashion, assuming expediting in increasing cost order.

C.2.1. Simulation implementation. All simulations were run for 364 periods (for a year's worth of daily operations), after a burn-in period to allow for pipeline orders to be initiated. For each simulation scenario, the policies were all tested using common random numbers for all stochastic model components (demand, disruptions, etc.) so that confidence intervals for cost comparisons can be constructed using paired tests. The initial disruption state is randomly drawn according to the stationary distribution of r_t characterized in (1). Costs were averaged across the 364 periods, then across 100 separate simulations for each instance. For each instance, all stochastic programs were solved with a sample average approximation using 20,000 samples from the relevant demand distributions. All optimization problems were solved with Gurobi version 10.0. Calculating base-stock levels accounts for most of the variation in simulation run time (discussed in more detail in Appendix C.2.5), with times varying from a few seconds (FZ and NV) to a few minutes (SP policies).

We also note that in this section, since we model stationary demand distributions at the retailers, all base-stock policies compute stationary base-stock levels over time, since the policies only take into account the distribution of demand and the (constant over time) cost parameters (in Section 6 we consider non-stationary demand and therefore non-stationary base-stock levels, see Appendix C.3.1 for further discussion).

C.2.2. Comparison of Stochastic Program Base-Stock Policies. Here we test the performance of the SP base-stock policies relative to each other. To make this comparison, we first define SP^{min} as the minimum cost policy among the three SP base-stock policies on a given instance, i.e., $SP^{min} = \min(SP^{L}, SP^{L+l}, SP^{avg})$. To focus solely on the performance of the base-stock levels, we make this comparison without implementing any special allocation heuristic during disruption, i.e., we consider the N disruption allocation policy. The results are summarized in Table 11. The table reports the ratio of average

Table 12 Average and worst case cost ratios of SP$^{\text{avg}}$-N, FZ-N, and NV-N compared to LB

	SP$^{\text{avg}}$-N/LB		FZ-N/LB		NV-N/LB	
α	Avg.	Max.	Avg.	Max.	Avg.	Max
0.005	1.47	1.71	1.52	1.80	1.51	1.80
0.01	1.49	1.61	1.58	1.91	1.56	1.65
0.02	1.71	1.83	2.17	2.74	1.87	2.09
All	1.56	1.83	1.76	2.74	1.65	2.09

Table 13 Runtime of solving base-stock levels for SP^L, SP^{L+l}, FZ, NV

n	SP^L	SP^{L+l}	SP$^{\text{avg}}$	FZ	NV
4	55s	45s	100s	0.3ms	0.68s
8	86s	87s	173s	0.2ms	1.52s
16	138s	137s	275s	0.2ms	2.82s
All	93s	89s	182s	0.2ms	1.71s

daily costs between the SP$^{(\cdot)}$-N policies and SP$^{\min}$-N, in the average and worst case across disruption probabilities. Overall, the performance of SP^{L+l}-N and SP$^{\text{avg}}$-N dominate that of SP^L-N, providing a 10% and 11% improvement, respectively, from SP^L on average. From the breakdown, although SP^L-N performs relatively competitively compared with SP$^{\text{avg}}$-N at larger disruption probabilities, the performance of SP^L-N suffers when $\alpha = 0.005$. We suspect this is because SP^L-N underestimates the base-stock levels (this underestimation becomes less of a problem when the number of disruption periods increases). Due to the consistent performance of both SP$^{\text{avg}}$-N and SP^{L+l}-N in terms of both average and maximum cost ratios, we focus on these two base-stock policies, SP$^{\text{avg}}$ and SP^{L+l}, in the remainder of the synthetic simulations when considering additional disruption allocation policies (other than N) and when comparing to FZ and NV base-stock policies.

C.2.3. A Stochastic Program for Disruption Allocation A more nuanced alternative to the all-or-nothing centralization vs. decentralization comparison of Section 4 is to consider partially centralizing some inventory, while leaving the remaining inventory at the retailers. Such a policy can be implemented in a similar way to the centralization policy described in Section 4: at the beginning of a disruption we can calculate new retailer base-stock levels to keep in place during the disruption, if these are lower than in normal mode then more inventory will be kept centrally over the disruption. This obviously requires more involved computation, but such an approach can be characterized compactly with a stochastic program, similar in spirit to the approach in Section 3.1.2.

When a disruption begins in period t, let τ be the disruption length and $IP_i = X_{i,t-1} - B_{i,t-1} + \sum_{s=t-l}^{t-1} z_{i,s}$ be the inventory position of retailer i. Then let $\hat{L} = \min(\tau, L)$ denote the number of periods of outstanding supplier orders that will arrive during the disruption, and let $IP_0 = I_{t-1} + \sum_{s=t-\hat{L}}^{t-1} x_s$ denote the total inventory that will be available at the warehouse over the course of the disruption. Also, let $\bar{\mathbf{D}} = \sum_{j=1}^{\tau} \mathbf{D}(j)$ denote the demand over τ periods of the disruption. Then we solve the following two stage stochastic program to decide on the target inventory levels of the warehouse and RDCs.

$$\min_{I,\mathbf{X}\geq 0} \left\{ \mathbb{E}_{\bar{\mathbf{D}}} \left[g(I, \mathbf{X}, \bar{\mathbf{D}}) \right] \,\middle|\, I + \sum_i X_i = IP_0 \right\} \tag{42}$$

$$g(I, \mathbf{X}, \bar{\mathbf{D}}) = \min_{\mathbf{y}\geq 0} \left\{ \sum_i f_i y_i + \frac{\tau+1}{2} \sum_i b_i \left(\bar{D}_i - IP_i - X_i - y_i \right)^+ \,\middle|\, \sum_i y_i \leq I \right\}.$$

Letting I^* and X_i^* denote the optimal first stage solution, we set new base-stock levels for allocation during the disruption as $S_0 = I^*$, $S_i = X_i^* + IP_i$. The intuition for this stochastic program is as follows. First, similar to Section 4, we drop the holding cost accounting as we are ordering no new inventory during the disruption. The variables I and X_i represent the inventory we aim to hold at the warehouse and at retailer i, respectively, and y_i denotes the amount of expedited fulfillment to retailer i during the disruption. The objective then accounts for the expediting cost and the backlogging cost (using an approximation similar to the one developed in Section 3.1.2).

The final disruption allocation policy uses the stochastic program (42) to set base-stock levels during disruptions. This program is solved at the start of the disruption, and the resulting base-stock levels are implemented for the duration of the disruption. We call this policy DSP (for disruption SP).

C.2.4. Comparison with Lower Bound. Here we benchmark our SP$^{\text{avg}}$-N policy relative to an optimal policy, by comparing with the stochastic program lower bound characterized in Theorem 1. We let LB denote this lower bound, and now consider the ratio of the cost of SP$^{\text{avg}} - N$ over LB (which provides an upper bound on the performance relative to an optimal policy). For comparison's sake, we also consider the ratio of the cost of FZ-N and NV-N over LB. These ratios are summarized in Table 12, with average and worst case values broken down by the disruption probability. As expected from the earlier results, the SP$^{\text{avg}}$-N policy is closer to optimal than the FZ-N and NV-N policies, with a relative error less than 56% on average and 83% in the worst case tested, compared to 76% and 174% errors for FZ-N and 65% and 109% for NV-N. Put another way, for these problem instances our policy reduces the gap between the benchmark FZ-N policy and an optimal policy by close to 30% on average and over 50% in the worst case.

As evidenced by the table, a gap between SP$^{\text{avg}}$-N and optimal may still remain, but we note that, given the complex nature of distribution systems, approximations and bounds in the literature often suffer errors of similar magnitude (e.g., Federgruen and Zipkin, 1984b,a; Gallego et al., 2007; Federgruen et al., 2018 all report gaps on the order of 10%-80% for various sets of problem instances). Further, we conjecture in our case that the gap may be more due to the looseness of the lower bound, as opposed to the SP$^{\text{avg}}$-N policy itself. In deriving Theorem 1, a number of relaxations had to be made simultaneously in order to obtain a tractable stochastic program (including relaxing non-anticipativity constraints, scaling expediting costs for lead times, approximating costs over disruption intervals, and aggregating demand across lead times); it is an interesting direction for future research to explore further whether some of these can be tightened.

To explain further, the simulation ratios partly stem from the use of leadtime demand for expediting decisions. Roughly speaking, in our model the expediting cost can be thought of as a type of preliminary backlog cost, incurred for demand unfulfilled by local inventory (when central inventory is available). The difference is that this cost is only incurred in one period (since it clears the backlog when it is incurred), whereas the true backlog cost may persist for many periods. This difference makes the cost accounting between the lower bound and heuristic policy more challenging. For example, consider a simple single-echelon inventory model with equal holding and backorder costs (i.e., $h = b$). For leadtime $L = 1$, the optimal policy is to order up to the median demand μ in each period. For normally distributed demand in each period with standard deviation σ, the average backlog cost per period is $b\sigma\mathcal{L}(0)$, where $\mathcal{L}(0)$ is the standard normal loss function $\mathcal{L}(z)$ at $z = 0$. When demand is aggregated over L periods, the order-up-to quantity is $L\mu$ and the total backlog cost over L periods is $b\sigma\sqrt{L}\mathcal{L}(0)$. But for expediting, this cost should be translated back to a per-period cost of $b\sigma\mathcal{L}(0)/\sqrt{L}$. In reality, costs are, however, accumulated for expediting in each period and at least on the order of costs with $L = 1$, suggesting a gap between the heuristic and this part of the lower bound on the order of $\sqrt{L}$. The ratio is less for other parts of the objective, and so overall consistent with the results in Table 12.

C.2.5. Runtime Summaries. Here we discuss the runtime of simulating the different policies. We first note that our simulation approach involves two steps. For each policy, we first solve for the base-stock levels and then use those calculated base-stock levels as input in the simulation to estimate the system cost of each policy. Running the actual simulation takes approximately 267 seconds (per instance) for a system with four retailers, 343 seconds for eight retailers, and 422 seconds for 16 retailers. These numbers are fairly consistent across the SP^{L+l}, SP^{avg}, FZ, and NV policies (all for the N disruption allocation heuristic). Therefore, Table 13 only reports the runtime for solving the base-stock levels for each policy prior to running the actual simulation. As expected, all policies (except FZ) require more time to solve as the number of retailers increases. It is worth noting that FZ solves almost instantly, primarily because it only solves for a system-wide base-stock level using a simple newsvendor heuristic. This approach has two drawbacks: (1) it does not take disruptions into consideration, and (2) it does not provide individual base-stock levels for retailers (which, in practice, raises concerns for our industrial partner, as their inventory planning application requires retailer-specific stocking levels). The NV base-stock policy, meanwhile, solves relatively quickly and can perform well when the disruption probability is low, but our industrial partner believes the increased computation time of the SP base-stock policy is well worth the improved performance when the risk of disruption is higher. This is especially true in their practical context, where the base-stock levels are typically updated on a weekly basis so the computation time is not constrained to be on the order of seconds.

C.3. More details on data and simulations for our industrial partner

In this section, we provide more details on our parameter estimation procedure. From our industrial partner, we are given a year-long data on 23 repair parts. We are informed that the data representing their normal mode of operations as the data year is pre-COVID and the system did not experience major supply or technology disruption. It is worth noting that all data provided are masked but represent the realistic setting.

Poisson demand. To estimate weekly demand for each part, we first select a partial demand not affected by dealers' base-stock policies and then apply the James-Stein estimator (Stein, 1981) to estimate each week's overall demand. For each part, we identify a small group of dealers who only placed customer orders. This unique ordering behavior implies these dealers do not carry inventory (equivalently zero base-stock level), and we are able to observe the uncensored demand trend from them. We first summarize weekly customer-order demand from these dealers on the RDC level. We then deseasonalize the weekly data points by what we call a "dynamic moving average" method. For RDC i in week t with demand $D_{t,i}$, its demand is detrended to be the average of

$$D'_{t,i} = \frac{\sum_{t-r}^{t-1} D_{t,i} + D_{t,i} + \sum_{t+1}^{t+r} D_{t,i}}{2r+1}$$

where $r = 6$ in our detrend method. To apply this method for all 52 weeks, we *wrap around* the year and make the whole year a cycle. In addition to the RDC-level data, we can adopt the same detrend method to the national-level data and obtain the deseasonalized national trend (D'_t) for each part. For a part, we define

$$N_{t,i} = D'_t \frac{\sum_t D'_{t,i}}{\sum_t D'_t}$$

to be the scaled national estimate for RDC i. The intuition behind using the James-Stein estimation is that it takes into account both the local trend $D'_{t,i}$ as well as the scaled national trend $N_{t,i}$ to provide a more accurate and reasonable estimate for weekly demand. By calculating the Stein's parameter, we obtain a convex combination of the two trends, denoted by $J_{t,i}$.

Last, we split the raw, weekly aggregated demand (customer and normal orders at the national level) proportionally to $J_{t,i}$'s weekly weight. Denote this scaled RDC demand as $\tilde{J}_{t,i}$. To split demand to dealer level, we split $\tilde{J}_{t,i}$ proportionally to each dealer's yearly demand compared to the belonged RDC's yearly demand to obtain dealer's weekly demand. Finally, we divide the weekly demand by seven to obtain customer's daily arrival rate λ_t at a dealer.

Dealers' review policies. In conjunction with our industrial partner, we model four major types of dealers for each part in our simulation. The first type is those who do not hold inventory. These dealers are easily to be identified as in data they only place customer orders, implying they do not carry any inventory and hence their base-stock level is zero. The second type of dealers are those who carry inventory and conduct daily review. The third and fourth types are dealers who review their inventory every seven days on average. We categorize the third type as those following a review interval of Binomial(n, p) where $np = 7$ and the fourth type as those following a review length of Poisson (7). From the raw data and estimated demand distribution, we determine for each dealer, whether it belongs to the second, the third, or the fourth group by finding the review policy that gives the maximum likelihood (MLE).

We introduce a few parameters before detailing our MLE approach. From data, for each part-dealer pair, we first calculate inter-order times. Let T_k denote the number of days for the kth inter-order time. Let λ_{T_k} denote the cumulative Poisson rate for the kth inter-order interval. Last, let $\bar{\lambda}_{T_k} = \frac{\lambda_{T_k}}{T_k}$ denote the average daily arrival rate during the kth interval.

For daily review, the log likelihood is given as follows

$$LL_{\text{daily}} = \sum_k \log\left(e^{-\bar{\lambda}_{T_k}(T_k - 1)}(1 - e^{-\bar{\lambda}_{T_k}})\right).$$

Table 14 Unit labor cost in different categories

Weight (lb)	Picking & Packing	Storing	Cross-docking
≤ 1	\$0.42	\$0.83	\$0.83
10	\$1.25	\$2.50	\$2.50
100	\$5.00	\$10.00	\$10.00

Table 15 Avg. cost ratios of SP base-stock with each disruption heuristic compared to SP-MIN

	SP-(·)/SP-MIN			
Disr. Alloc.	N	DFI	MDFI	DSP
Avg. cost ratio	1.04	1.12	1.10	1.10

For dealers with Binomial(n,p) where $np=7$, we adopt a grid search and a Fibonacci search to search for n^* that gives the highest Binomial likelihood. We refer to Nocedal and Wright (1999) for implementing search algorithms. The log likelihood of Binomial(n,p) is

$$LL_{\text{Bin}}(n,p)=\sum_k \log\left(\sum_{i=1}^{T_k}\sum_{z=0}^{T_k}\Pr^{\text{B}}\Big(z;(i-1)n,p\Big)\Pr^{\text{B}}\Big(T_k-z;n,p\Big)e^{-\bar{\lambda}z}(1-e^{-\bar{\lambda}(T_k-z)})\right),$$

where $\Pr^{\text{B}}(m;n,p)=\binom{n}{m}p^m(1-p)^{n-m}$ represents the Binomial probability mass function. Last, the log likelihood of Poisson(7) is similar to the Binomial case because of the Poisson approximation of Binomial for large n:

$$LL_{\text{Pois}}(7)=\sum_k \log\left(\sum_{i=1}^{T_k}\sum_{z=0}^{T_k}\Pr^{\text{P}}\Big(z;7(i-1)\Big)\Pr^{\text{P}}\Big(T_k-z;7\Big)e^{-\bar{\lambda}z}(1-e^{-\bar{\lambda}(T_k-z)})\right),$$

where $\Pr^{\text{P}}(m;\lambda)=\frac{\lambda^m e^{-\lambda}}{m!}$ represents the Poisson probability mass function.

For each part-dealer pair, we calculate the log likelihood of each review policy (daily, Binomial(n,p), and Poisson(7)) and assign the policy with the maximum likelihood to that dealer.

Dealers' base-stock levels. In conjunction with our industrial partner, for dealers who hold inventory, we infer their base-stock levels using the following method. First, we determine each part-dealer pair's service level β by setting it equal to

$$\beta=\frac{\text{number of normal orders}}{\text{number of normal orders}+\text{customer orders}}.$$

Recall that each dealer with a review policy either reviews daily or reviews weekly on average. In conjunction with our industrial partner, we add 5 days to each dealer's review period length for a "lead time" estimate. This approach is standard in the automaker's inventory calculation. So each dealer for each part will have a lead time length of either 6 or 12 days. Multiplying the lead time length with the averaged daily demand (i.e. sum of Poisson demand over a year divided by 365) gives the aggregate Poisson demand rate. Finally, the smallest number that guarantees the dealer's service level β in the aggregate Poisson demand distribution is the dealer's base stock level for this part. In our simulation, most of dealers for most parts have base-stock level of 1, which again is confirmed by our industrial partner that this is a standard practice.

Initial RDC inventory. From the raw data, we have information of month-end inventory on hand for each part-RDC pair. For simplicity, we average 12 month-end inventories and use it as the initial part-RDC's on-hand inventory when a simulation starts. Observe that the initial inventory can be arbitrary, as long as there is a sufficient burn-in period, which is what we also have implemented.

Lead Times. We estimate both supplier-to-warehouse lead time and warehouse-to-retailer lead times for each part. For each origin-destination pair, we estimate the mean as well as the standard deviation of lead times from a year-long dataset. In conjunction with our industrial partner, we assume warehouse-to-retailer and supplier-to-warehouse lead times are normally distributed. To fully specify the resulting adjustments to the model of Section 2 required to incorporate these lead times, we introduce additional notation. Let L'_t denote a nominal lead time random variable for period t, independent and identically distributed across time with a normal distribution. Then, to enforce the standard sequential lead time model (i.e., no order crossing, Zipkin, 2000), we let the true lead time at time t be $L_t=\max(L'_t, L_{t-1}-1)$ so that $t+L_t$ is increasing in t. Similarly, for retailer i let $l'_{i,t}$ denote a nominal lead time random variable for period t, independent and identically distributed across time with its own normal distribution, and let the true lead time at time t be $l_{i,t}=\max(l'_{i,t}, l_{i,t-1}-1)$. Then, the inventory evolution equations of (2) are updated as follows

$$I_t=I_{t-1}+\sum_{s|s+L_s=t}x_s-\sum_i(y_{i,t}+z_{i,t}),\qquad X_{i,t}=X_{i,t-1}+\sum_{s|s+l_{i,s}=t}z_{i,s}-w_{i,t}, \tag{43}$$

Partner RDCs. The automaker has one warehouse and 15 RDCs across the contiguous United States and Canada. 10 of the 15 RDCs have one or two partner RDCs based on geographic proximity. These partner RDCs act as backup options for fulfillment. We provide the following masked RDC IDs along with their partner RDC pairs for demonstration. RDC List: 1,2,3,4,5,6,7,8,9,10,11,12,13,14,15. Partner RDC sets: {1,3,7}, {2,11,13}, {9,12}, {14,15}.

Costs. For each unit of inventory that passes through the central warehouse, we charge a unit cross-docking fee. For each unit of inventory that is stored at the warehouse, we charge a unit storing fee. For each unit of those stored inventory used for RDC replenishment or expedited shipments, we charge a unit picking and packing fee. These disguised but realistic unit costs depend on repair parts' weights, which is shown in the table below.

C.3.1. Adapting Policies to the Practice-Based Context Here we describe how we adapt the the policies developed for the more stylized models of Sections 3 and 4 to accommodate the additional problem features present in our industrial partner's setting. We first describe the base-stock policies, then the disruption allocation policies.

Base-Stock Policies. We consider three base-stock policies: our industrial partner's incumbent policy I (which sets independent RDC base-stock levels, see section 6.1.5 for more details), our SP policy introduced in Section 3, and the NV heuristic introduced in Section 3.3. We next describe the implementation of the SP and NV policies in our industrial partner's setting. For the SP policy, we first recall that the simulations in section 5 lead us to focus on the SP^{L+l} policy as the best balance of performance and computation for these practice based simulations. We thus focus on adapting the SP^{L+l} to this setting, and to simplify notation we simply refer to this policy as SP in this section. To implement the SP base-stock policy, we adapt our stochastic program from Section 3.2 to accommodate the additional problem features of the automaker's practice-based context in order to design an effective base-stock policy. First, we take into account the fact that the automaker's demand is non-stationary by setting time-dependent base-stock levels that vary with the forecast of future demand. In particular, in a period t where an order can be placed, we set base-stock levels by solving a stochastic program considering demand over several periods of future demand, starting from period $t+1$. When deciding how many future periods of demand to include in the stochastic program, we recall a few additional features of the problem. Because the automaker places supplier orders at weekly intervals, we account for uncertain demand during periods in between successive orders by including a parameter $p=6$ to model this interval, and include an extra p periods of demand in the stochastic program. We also recall that the automaker's lead-times are stochastic, and may have different means across retailers. To account for this, we could define a compound sampling process for demand where we draw a random lead time then draw random demand over this lead time. However, the coefficient of variation of the lead time in the automaker's data is quite low in general (about 0.043 on average), and so we found that for the purposes of the stochastic program, this compound sampling process introduces unnecessary computational burden without much benefit over using the mean lead times. Thus, for defining the stochastic program in this section, we let L denote the mean lead time from the external supplier (rounded, if necessary, to the nearest integer), and l_i denote the mean lead time from the warehouse to RDC i. Finally, let $D_{i,t}$ denote a random variable distributed as the total end customer demand for retailer i on day t, and let $\mathbf{D}_t$ denote the vector of these demands. Then, for a given day t, we define the following demand vectors:

$$\hat{\mathbf{D}}_t^1 = \sum_{s=t+1}^{t+L+\min(l_i,p+\tau)} \mathbf{D}_s, \qquad \hat{\mathbf{D}}_t^2 = \sum_{s=t+L+\min(l_i,p+\tau)+1}^{t+L+p+\tau} \mathbf{D}_s, \qquad \hat{\mathbf{D}}_t^3 = \sum_{s=t+L+p+\tau+1}^{t+L+\max(l_i,p+\tau)} \mathbf{D}_s, \qquad \hat{\mathbf{D}}_t^4 = \sum_{s=t+L+\max(l_i,p+\tau)+1}^{t+L+l_i+p+\tau} \mathbf{D}_s,$$

which we will use as inputs for the stochastic program. We highlight that for RDC i we consider demand over $L+l_i+p+\tau$ periods in total. We start with $L+l_i$ periods as this is the amount of time it takes for an order to arrive at RDC i, and our experiments in Section 5 suggest that this is effective. We include another $p+\tau$ periods, since this is the (random) amount of time until another order can be placed.

With these demand vectors defined, we are now ready to specify the SP base-stock policy. On a given day t when we can place a supplier order, we let $\hat{\mathscr{D}}_t^2 = (\hat{\mathbf{D}}_t^1, \hat{\mathbf{D}}_t^2)$ denote second stage demand and $\hat{\mathscr{D}}_t^3 = (\hat{\mathbf{D}}_t^3, \hat{\mathbf{D}}_t^4)$ denote third stage demand, and we solve the stochastic program $\mathcal{SP}(\hat{\mathscr{D}}_t^2, \hat{\mathscr{D}}_t^3)$, i.e., the stochastic program (5) with these new demand vectors as inputs. Then, we set base-stock levels according to (6) and (7) using the optimal solution of $\mathcal{SP}(\hat{\mathscr{D}}_t^2, \hat{\mathscr{D}}_t^3)$ and follow the base-stock policy of Section 3.2, and we denote this policy in this section as SP. The NV base-stock policy follows a similar adaptation to the practice based setting: the same recursive process described in Section 3.3 is used to calculate base-stock levels, but the demand is updated to use the $\hat{\mathbf{D}}_t^j$ for $j=1,\ldots,4$ distributions above in order to incorporate the review period and different lead times for each RDC. In particular, the recursive optimization process described in Steps 1-3 in Appendix C.1.2 remain the same. The main change is in updating the demand: in period t the demand $\check{D}_i$ is replaced with $\grave{D}_{i,t} = \hat{D}_{i,t}^1 + \hat{D}_{i,t}^2 + \hat{D}_{i,t}^3 + \hat{D}_{i,t}^4$, which represents the first $L+l_i+p+\tau$ periods of demand after period t, and demand D_0 in Step 2 is replaced with $\grave{D}_{0,t} = \sum_i (\grave{D}_{i,t} - S'_{i,t})^+$. Further, in (39) and (41), the leadtime l is updated to the RDC specific leadtime l_i.

We note that the SP and NV policies also follow the same partner RDC fulfillment logic as the automaker's incumbent I policy (see Section 6.1.5 and Appendix C.3 for more details), in order to maintain a consistent comparison. As a final note on the base-stock policies we highlight that all policies we consider, I, NV, and SP, compute non-stationary base-stock levels. This is driven by the fact that demand in our industrial partner's system is non-stationary over time, meaning base-stock levels need to be adjusted accordingly. Practically, this is implemented in each policy by varying the demand distribution input according to the order interval considered. Specifically, while policy I maintains constant service level targets (see Section 6.1.5) over time, these are implemented on different demand distributions in each order interval, leading to non-stationary base-stock levels. The NV and SP policies are updated for the demand distribution in each order interval and have non-stationary base-stock levels.

Disruption Allocation Policies The disruption allocation policies we test in our industrial partner's context are straightforward applications of the N, DFI, MDFI, and DSP policies already described in Section 5.2. The N policy simply does keeps the same base-stock levels in both disrupted and undisrupted states, while the DFI policy sets all RDC base-stock levels to 0 during a disruption. The MDFI policy modifies the DFI policy to maintain the undisrupted base-stock levels during a disruption if condition (10) is met at the beginning of the disruption. Finally, for DSP we solve a slightly modified version of the SP (42): when a disruption happens in period t, let $\bar{\mathbf{D}}_t = \sum_{s=1}^{\tau+p} \mathbf{D}_s$ denote the random demand vector over the next $\tau+p$ periods (covering the random disruption length plus a review period until the next ordering opportunity), and we solve the stochastic program (42) with $\bar{D}$ replaced by $\bar{D}_t$. We then follow the same DSP disruption allocation policy using the optimal solution of this SP described in Section C.2.3.

C.3.2. Testing Disruption Allocation Policies in the Practice Based Context Here compare the disruption allocation heuristics, and for a default we use the SP base-stock policy. For each simulation scenario, let SP-MIN denote the minimum-cost disruption heuristic used by the SP policy, i.e., SP-MIN = min(SP-N,SP-DFI,SP-MDFI,SP-DSP). Table 15 compares the performance of each disruption heuristic using the same set of base-stock levels provided by the SP before running the simulation. The observations here are similar to those from the synthetic simulation: the MDFI heuristic performs on par with the more complex DSP heuristic, while the N heuristic (allocating inventory during disruptions the same as in normal times) still dominates overall, providing a 5-7% improvement over the other options. This further demonstrates that our SP approach inherently performs well across both normal times and disruptions because it already considers disruptions when calculating optimal base-stock levels. In the remainder of this section, we focus on the default SP-N policy and compare its performance with the automaker's incumbent policy I-DFI, the modified incumbent policy I-MDFI, and the newsvendor base-stock policy NV-N.

Electronic Companion References

Durrett, Richard. 1999. *Essentials of Stochastic Processes*, vol. 1. Springer.

Dutkay, Dorin Ervin, Constantin P. Niculescu, Florin Popovici. 2013. Stirling's formula and its extension for the Gamma function. *The American Mathematical Monthly* **120**(8) 737–740.

Federgruen, Awi, C Daniel Guetta, Garud Iyengar. 2018. Two-echelon distribution systems with random demands and storage constraints. *Naval Research Logistics (NRL)* **65**(8) 594–618.

Federgruen, Awi, Paul Zipkin. 1984. Computational issues in an infinite-horizon, multiechelon inventory model. *Operations Research* **32**(4) 818–836.

Gallego, Guillermo, Özalp Özer, Paul Zipkin. 2007. Bounds, heuristics, and approximations for distribution systems. *Operations Research* **55**(3) 503–517.

Kunnumkal, Sumit, Huseyin Topaloglu. 2008. A duality-based relaxation and decomposition approach for inventory distribution systems. *Naval Research Logistics (NRL)* **55**(7) 612–631.

Kunnumkal, Sumit, Huseyin Topaloglu. 2011. Linear programming based decomposition methods for inventory distribution systems. *European Journal of Operational Research* **211**(2) 282–297.

Nocedal, Jorge, Stephen J. Wright. 1999. *Numerical Optimization*. Springer.

Press, William H, Brian P Flannery, Saul A Teukolsky, William T Vetterling. 1992. Incomplete gamma function, error function, chi-square probability function, cumulative Poisson function. *Numerical Recipes in Fortran*. Cambridge University Press, New York City, 209–215.

Rong, Ying, Zümbül Atan, Lawrence V Snyder. 2017. Heuristics for base-stock levels in multi-echelon distribution networks. *Production and Operations Management* **26**(9) 1760–1777.

Stein, Charles M. 1981. Estimation of the mean of a multivariate normal distribution. *The Annals of Statistics* 1135–1151.

Van der Vaart, H.R. 1961. Some extensions of the idea of bias. *The Annals of Mathematical Statistics* **32** 436–447.