Appendix A: Proofs

A.1. Proof of Lemma 1

Proof. The proof relies on the observation that swapping adjacent stores when the earlier store has a higher unreliability factor results in a lowering of the expected cost.

Assume for the sake of a contradiction that $r_m = r_{\sigma^*(l)} > r_{\sigma^*(l+1)} = r_n$.

$$\begin{aligned} r_m > r_n &\implies \frac{r_m}{1-\phi_n} - \frac{\phi_n \cdot r_m}{1-\phi_n} > \frac{r_n}{1-\phi_m} - \frac{\phi_m \cdot r_n}{1-\phi_m} \\ &\implies (1-\phi_m)\cdot r_m + \phi_m \cdot (1-\phi_n)\cdot r_n > (1-\phi_n)\cdot r_n + \phi_n \cdot (1-\phi_m)\cdot r_m \end{aligned} \tag{18}$$

Expanding the cost of the optimal policy, we get

$$\begin{aligned} \mathcal{V}(\sigma^*) = &\left[\sum_{l=1}^{q-1}\left(\prod_{l'=1}^{l-1}\phi_{\sigma^*(l')}\right)\cdot(1-\phi_{\sigma^*(l)})\cdot r_{\sigma^*(l)}\right] \\ &+ \left(\prod_{l=1}^{q-1}\phi_{\sigma^*(l)}\right)\cdot(1-\phi_m)\cdot r_m + \left(\prod_{l=1}^{q}\phi_{\sigma^*(l)}\right)\cdot(1-\phi_n)\cdot r_n \\ &+ \left[\sum_{l=q+2}^{L}\left(\prod_{l'=1}^{l-1}\phi_{\sigma^*(l')}\right)\cdot(1-\phi_{\sigma^*(l)})\cdot r_{\sigma^*(l)}\right] + \left[\left(\prod_{l=1}^{L}\phi_{\sigma^*(l)}\right)\cdot d\right] \end{aligned} \tag{19}$$

Examining the middle two terms corresponding to the l^{th} and $(l+1)^{st}$ stores we get

$$\begin{aligned} (19) &= \left(\prod_{l=1}^{q-1}\phi_{\sigma^*(l)}\right)\cdot\Big[(1-\phi_m)\cdot r_m + \phi_m\cdot(1-\phi_n)\cdot r_n\Big] \\ &> \left(\prod_{l=1}^{q-1}\phi_{\sigma^*(l)}\right)\cdot\Big[(1-\phi_n)\cdot r_n + \phi_n\cdot(1-\phi_m)\cdot r_m\Big] \qquad \text{using (18)} \end{aligned} \tag{20}$$

Interchanging the trial positions of stores m and n would yield a policy with cost term (19) in $\mathcal{V}(\sigma^*)$ replaced by (20). The cost of this policy is lower than $\mathcal{V}(\sigma^*)$, giving the desired contradiction. □

A.2. Proof of Lemma 3

Table 3 Notations

$\mathcal{Q}^O_{jk}$	:	inventory level for the k^{th} online order to be tried at store j.
$t'_{ij}(q)$	:	expected trial cost at store j with inventory level q
		$t'_{ij}(q) = b_j + (1-\phi_j(q))\cdot s_{ij} + \phi_j(q)\cdot d_i$

Recall

$$t_{ijk}(\mathcal{Q}^P_j) = \sum_{q=\mathcal{Q}^P_j-(k-1)}^{\mathcal{Q}^P_j} \mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right]\cdot t'_{ij}(q)$$

1. We observe from the definition of $t_{ijk}(\mathcal{Q}^P_j)$ that it is a convex combination of expected fulfillment costs at different feasible inventory levels.

2. The expected fulfillment cost at inventory level q, $t'_{ij}(q)$ is the sum of a constant b_j and convex combination of s_{ij} and d_i. Since $s_{ij} < d_i$ and $\phi_j(q)$ is a non-increasing function (assumption 1), it follows that the expected trial cost $t'_{ij}(q)$ is non-increasing in q.

3. $\mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right]$ is the probability that the leftover inventory is q after the trial of $k-1$ online orders. The $k-1$ trials can be decomposed into two events (i) successful attempts of $\mathcal{Q}^P_j - q$ online orders one each at inventory level $\left\{q+1, \ldots, \mathcal{Q}^P_j\right\}$ and (ii) failed attempts of the remaining $(k-1)-(\mathcal{Q}^P_j - q)$ online orders in the set of inventory levels $\left\{q, \ldots, \mathcal{Q}^P_j\right\}$. The probability of event (ii) decreases with k because when more items are independently tried at same pick failure probability levels, more items are likely to fail.

$$\mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right] = \underbrace{\left(\prod_{q'=q+1}^{\mathcal{Q}^P_j} [1-\phi_j(q)]\right)}_{\text{(i):independent of } k} \cdot \underbrace{\mathbf{Pr}\left[(k-1)-(\mathcal{Q}^P_j - q) \text{ failures}\right]}_{\text{(ii):non-increasing in } k} \tag{21}$$

Therefore, $\mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right]$ is non-increasing in k.

4. We are now ready to compare $t_{ijk}(\mathcal{Q}^P_j)$ and $t_{ij(k-1)}(\mathcal{Q}^P_j)$.

$$\begin{aligned}
t_{ijk}(\mathcal{Q}^P_j) &= \sum_{q=\mathcal{Q}^P_j-(k-1)}^{Q^P_j} \mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right] \cdot t'_{ij}(q) \\
&= \sum_{q=\mathcal{Q}^P_j-(k-2)}^{Q^P_j} \mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right] \cdot t'_{ij}(q) \;+\; \mathbf{Pr}\left[\mathcal{Q}^O_{jk} = Q^P_j - (k-1)\right] \cdot t'_{ij}(Q^P_j - (k-1)) \\
&\overset{\textbf{Step 2}}{\geq} \sum_{q=Q^P_j-(k-2)}^{Q^P_j} \left(\mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right] + \underbrace{\left[\mathbf{Pr}\left[\mathcal{Q}^O_{j(k-1)} = q\right] - \mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right]\right]}_{\geq 0 \;\; \textbf{(Step 3)}} \right) \cdot t'_{ij}(q) \\
&\quad + \underbrace{\left(\mathbf{Pr}\left[\mathcal{Q}^O_{jk} = Q^P_j - (k-1)\right] + \sum_{q=Q^P_j-(k-2)}^{Q^P_j} \mathbf{Pr}\left[\mathcal{Q}^O_{jk} = q\right] \right)}_{=1 \;\; \textbf{(Step 1)}} \cdot t'_{ij}(Q^P_j - (k-1)) \\
&\quad - \underbrace{\left(\sum_{q=Q^P_j-(k-2)}^{Q^P_j} \mathbf{Pr}\left[\mathcal{Q}^O_{j(k-1)} = q\right] \right)}_{=1 \;\; \textbf{(Step 1)}} \cdot t'_{ij}(Q^P_j - (k-1)) \\
&= \sum_{q=Q^P_j-(k-2)}^{Q^P_j} \mathbf{Pr}\left[\mathcal{Q}^O_{j(k-1)} = q\right] \cdot t'_{ij}(q) \\
&= t_{ij(k-1)}(Q^P_j)
\end{aligned}$$

□

A.3. Proof of Lemma 4

We have the following objective function to be minimized.

$$G(S) = \mathbb{E}_{\mathscr{D}^O,\mathscr{D}^P}\left[\sum_{k=1}^{\min\left\{\mathscr{D}^O,S,\mathcal{Q}^P\right\}} t_k(\mathcal{Q}^P) + c\cdot\left[\min\left\{\mathscr{D}^O,S\right\} - \min\left\{\mathscr{D}^O,S,\mathcal{Q}^P\right\}\right]\right]$$
$$+\,\mathbb{E}_{\mathscr{D}^O,\mathscr{D}^P}\left[p\cdot\min\left\{\mathscr{D}^O - \min\left\{\mathscr{D}^O,S\right\},\mathcal{Q}^P - \min\left\{\mathscr{D}^O,S,\mathcal{Q}^P\right\}\right\}\right]$$

Consider the derivative of the objective function. Recall that $\overline{Q}$ is the initial store inventory and $\mathcal{Q}^P = \max(0,\overline{Q}-\mathscr{D}^P)$.

$$\begin{aligned}
G'(S) &= \mathbf{Pr}\left[\mathscr{D}^O \geq S\right]\mathbf{Pr}\left[\mathscr{D}^P \leq \overline{Q}-S\right]\cdot\mathbb{E}_{\mathscr{D}^P}\left[t_S(\mathcal{Q}^P)\right]\\
&\quad + c\cdot\mathbf{Pr}\left[\mathscr{D}^O \geq S\right]\cdot\left(1-\mathbf{Pr}\left[\mathscr{D}^P \leq \overline{Q}-S\right]\right)\\
&\quad - p\cdot\mathbf{Pr}\left[\mathscr{D}^O \geq S\right]\cdot\mathbf{Pr}\left[\mathscr{D}^P \leq \overline{Q}-S\right]\\
&= \mathbf{Pr}\left[\mathscr{D}^O \geq S\right]\cdot F(\overline{Q}-S)\cdot\mathbb{E}_{\mathscr{D}^P}\left[t_S(\mathcal{Q}^P)\right]\\
&\quad + c\cdot\mathbf{Pr}\left[\mathscr{D}^O \geq S\right]\cdot\left[1-F(\overline{Q}-S)\right] - p\cdot\mathbf{Pr}\left[\mathscr{D}^O \geq S\right]\cdot F(\overline{Q}-S)\\
&= \mathbf{Pr}\left[\mathscr{D}^O \geq S\right]\cdot\left[\left(\mathbb{E}_{\mathscr{D}^P}\left[t_S(\mathcal{Q}^P)\right]-p-c\right)\cdot F(\overline{Q}-S)+c\right]
\end{aligned}$$

Here F is the cdf of $\mathscr{D}^P$. Note that the first derivative, $G'(S)$ is independent of the distribution of online orders.

Let us consider S^* such that $G'(S^*) = 0$. Let $\delta > 0$.

$$\begin{aligned}
G'(S^*+\delta) &= \mathbf{Pr}\left[\mathscr{D}^O \geq S^*+\delta\right]\cdot\left(\left[\mathbb{E}_{\mathscr{D}^P}\left[t_{S^*+\delta}(\mathcal{Q}^P))\right]-p-c\right]\cdot F(\overline{Q}-S^*-\delta)+c\right)\\
&\geq \mathbf{Pr}\left[\mathscr{D}^O \geq S^*+\delta\right]\cdot\left(\left[\mathbb{E}_{\mathscr{D}^P}\left[t_{S^*}(\mathcal{Q}^P))\right]-p-c\right]\cdot F(\overline{Q}-S^*)+c\right) \qquad (22)\\
&= 0 \qquad G(S^*)=0\\
G'(S^*-\delta) &= \mathbf{Pr}\left[\mathscr{D}^O \geq S^*-\delta\right]\cdot\left(\left[\mathbb{E}_{\mathscr{D}^P}\left[t_{S^*-\delta}(\mathcal{Q}^P)\right]-p-c\right]\cdot F(\overline{Q}-S^*+\delta)+c\right)\\
&\leq \mathbf{Pr}\left[\mathscr{D}^O \geq S^*-\delta\right]\cdot\left(\left[\mathbb{E}_{\mathscr{D}^P}\left[t_{S^*}(\mathcal{Q}^P)\right]-p-c\right]\cdot F(\overline{Q}-S^*)+c\right) \qquad (23)\\
&= 0 \qquad G(S^*)=0
\end{aligned}$$

(22) and (23) rely on $F(\cdot)$ being a cdf and t_k being an increasing function in k (Lemma 3). Since, $G'(S)$ is decreasing when $S < S^*$ and increasing when $S > S^*$, S^* is a global minimum. □

A.4. Proof of Corollary 1

Proof. Let $\tilde{S} = \overline{Q} - F^{-1}\left(\frac{c}{p+c}\right)$ be the optimal threshold for the one store one zone problem without pick failure, as shown in Karp (2017). Let us further assume that $b = 0$ and $s = 0$, so that we are in the same setup as Karp (2017), though the corollary holds for any positive values of b and s.

$$
\begin{aligned}
G'(\tilde{S}) &= \mathbf{Pr}\left[\mathscr{D}^O \geq \tilde{S}\right] \cdot \left(\left[\mathbb{E}\left[t_{\tilde{S}}(\mathcal{Q}^P)\right] - p - c\right] \cdot F(\overline{Q} - \tilde{S}) + c \right) \\
&= \mathbf{Pr}\left[\mathscr{D}^O \geq \tilde{S}\right] \cdot \left(\left[\mathbb{E}\left[t_{\tilde{S}}(\mathcal{Q}^P)\right] - p - c\right] \cdot \frac{c}{p+c} + c \right) && \text{definition of } \tilde{S} \\
&= \mathbf{Pr}\left[\mathscr{D}^O \geq \tilde{S}\right] \cdot \left(\mathbb{E}\left[t_{\tilde{S}}(\mathcal{Q}^P)\right] \cdot \frac{c}{p+c} \right) \\
&> 0
\end{aligned}
$$

Since, $G'(S)$ is decreasing when $S < S^*$ and increasing when $S > S^*$, as shown in Lemma 4, it follows that $S^* < \tilde{S}$. □

Note that our single store problem reduces to the single store problem in Karp (2017) if b (the picking trial cost), s (the shipping cost) and $\phi(q)$ (the pick failure probability) are set to 0.

A.5. Proof of Theorem 5

Proof adapted from the proof of (Karp 2017, Theorem 2). We know from Theorem 3 that the linear program (14) can be solved using a min-cost circulation network. We first state a lemma about the supermodularity of objective value of min cost flow as a function of the supplies. Using this lemma, we proceed to prove unimodality of the objective function.

LEMMA 10 **(Karp (2017), Lemma 2)**. *In a minimum-cost single-commodity flow problem with multiple sources, one sink, and integer supplies, demands, and capacities, the objective value of a minimum cost feasible flow as a function of the supplies at the source nodes is supermodular.*

Let S^* be an optimal solution. We will consider two cases of threshold S, (i) $S < S^*$ and (ii) $S > S^*$. For each of the cases, let us consider realizations of demands $\mathscr{D}_i^O = D_i^O$ and $\mathscr{D}_j^P = D_j^P$. Let the online demand realizations D_i^O be such that $\sum_{i \in I} D_i^O \geq S$, because for other realizations, the objective value $G(S)$ does not change with respect to S. Therefore the total number of orders accepted $\sum_{i \in [I]} A_i^O(S)$ is S.

The min-cost flow problem has multiple sources with supplies. Let $V(S) = [A_1^O(S), ..., A_I^O(S)]$ be the supplies at zones with realized demands D_i^0, D_j^P at threshold S. We say $V(S_1) \subseteq V(S_2)$ if $A_i^O(S_1) \leq A_i^O(S_2)$ for $i \in [I]$.

Case 1: [$\boldsymbol{S \geq S^*}$] We observe that $V(S^*) \subseteq V(S^*+1) \subseteq V(S) \subseteq V(S+1)$. Now by Lemma 10 (supermodularity of $G(S)$), $G(S+1) - G(S) \geq G(S^*+1) - G(S^*)$). This implies that $G(S+1) - G(S)$ is non-decreasing for $S > S^*$

Case 2: [$\boldsymbol{S < S^*}$] We observe that $V(S-1) \subseteq V(S) \subseteq V(S^*-1) \subseteq V(S^*)$. Now by the Lemma 10 (supermodularity of $G(S)$), $G(S) - G(S-1) \leq G(S^*) - G(S^*-1)$). This implies that $G(S+1) - G(S)$ is non-increasing for $S < S^*$

Since, $G(S+1) - G(S)$ is non-increasing when $S < S^*$ and non-decreasing when $S > S^*$, S^* is a global minimum. □

Appendix B: Extracting model parameters from data and sampling orders

We discuss in detail the process of extraction of the model parameters from the data provided by our partner solutions provider. After the estimation of the distribution of model parameters from data, we describe the sampling of an order using the estimated parameter distributions.

1. *Distribution of number of stores for an SKU.* We consider the distribution of store-counts of SKUs, i.e., the number of stores that carry each SKU and made available online for SFS. To do this, we examined the inventory counts of the entire store network for every single SKU of the five retailers. For each SKU, we counted the number of locations where it is available to be fulfilled from and aggregated these counts in a $\log - \log$ plot in Figure 7.

Figure 7 Log-log plot of the distribution of the number of stores in which a typical SKU is available for fulfillment segregated by retailer

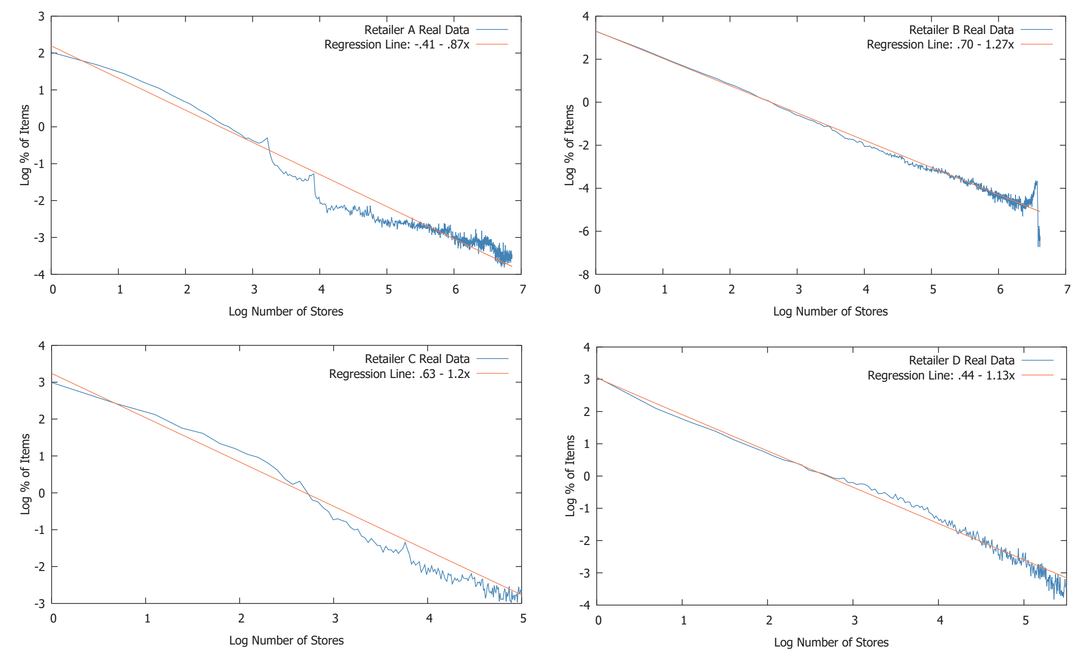

The clear straight line plots between the log of percentage of SKUs and the log of number of stores as shown in Figure 7 demonstrate a power law degree distribution for the store availability of SKUs . Suppose γ_r^k is the *percentage* of SKUs having k stores for retailer r. The regression computed the best fit for $\log \gamma_r^k = \delta_r + \eta_r \log k$.

Table 4 **Power law degree parameters for each retailer**

Retailer	A	B	C	D
δ	10.486	11.832	10.195	9.641
η	-0.871	-1.268	-1.2	-1.134

The power law and the long-tail nature of the SKU availability in stores is a major motivation for our work for two reasons: (i) An SKU can be present in multiple stores and therefore it is important to find the right store to fulfill it in order. (ii) Very few copies of the SKU are likely to be present in each store; hence there is considerable chance that we are unable to find the SKU in the store resulting in multiple tries. To model and sample from the fitted distributions, we truncate the power law at 500 stores and normalize to construct a synthetic distribution.

2. *Distribution of Shipping Costs.* USPS classifies the origin-destination distance into 8 zones, with zone 1 and zone 8 representing the shortest and longest distance respectively. Each zone has a shipping cost given as a function of the weight of the SKU (in lbs.) shipped. We use a publicly available data set of all the US zip codes along with their estimated population produced by the census (**?**), to come up with the distribution of zones for a store, given a random customer location. It is typically the case that stores are distributed close to population centers. Therefore, to simulate the customer-store zone distribution, we randomly sample pairs of (origin, destination) zip codes (using census data) and then use the USPS zone charts website (**?**) to determine the shipping zone for that pair of zip codes. The resulting distribution of zones is shown in Figure 8.

Figure 8 **Percentage of stores in different zones with respect to a typical customer**

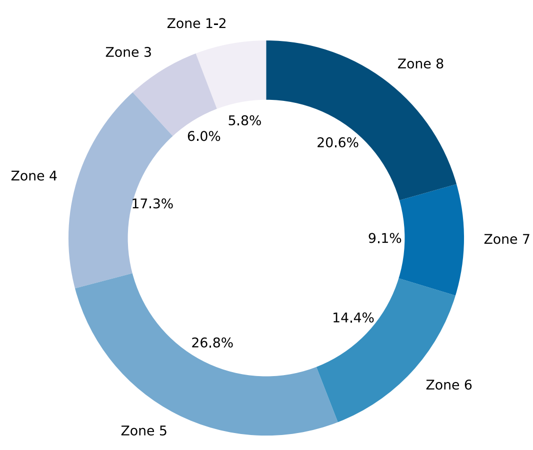

Figure 8 represents the distribution of zones of stores from a customer location, generated at random using census data. Assuming each SKU weighs 4lbs., the shipping cost (s_j for store j) is determined from the USPS Ground website (**?**) for a given zone in our simulations.

3. *Distribution of Labor Costs.* To sample labor costs, we use a representative retailer's per-store volume and per-store minimum wage. Figure 9 represents the distribution of store volume and store minimum wage provided by the representative retailer.

Figure 9 Minimum hourly wage distribution for a representative retailer

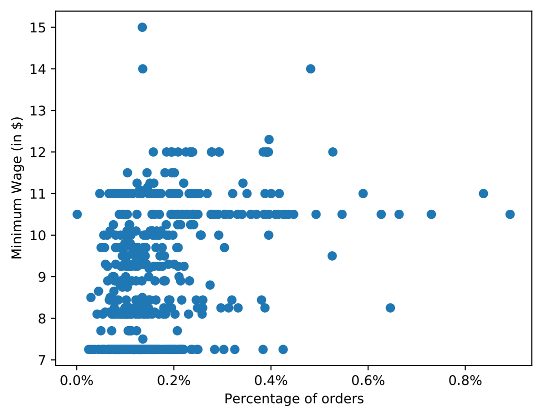

Note. A dot in the plot represents a store, with its x-coordinate showing the percentage of total retailer's volume which that store handles, and the y-coordinate represents the minimum wage at that store location.

To sample a store's labor cost, we sample a store from the list of stores for this retailer where the probability of picking a store is proportional to the store's sales volume. Once the store is picked we estimate the labor cost from the minimum wage data of the store based on its location.

These labor cost estimates are upper bounds on real costs. Minimum wage is not representative of actual handling costs due to sunk labor costs and the fact that typically 4-5 items can be processed within an hour. We scale these costs *down* by 4 to be realistic but scaling them by any other reasonable value does not impact the relative gains of our methods.

4. *Inventory Distribution & Pick Failure Probability.* By far the most important feature in predicting the pick failure probability is the location of the store itself. We received data from the provider where each row consists of an SKU that is ordered, the store that was attempted with its inventory of that SKU and the pick success/failure information.

We set the pick failure probability function at store j to be a simple step function of inventory so that there is constant pick failure ϕ_j up to a certain level of inventory $\widetilde{Q}_j$ and no pick failure beyond that. The fulfillment costs can be illustrated as

$$t_{ijk}(Q_j^P) = \begin{cases} b_j + (1-\phi_j)\cdot s_{ij} + \phi_j \cdot d_i & \text{if } \; k \geq \left(Q_j^P - \widetilde{Q}_j\right)^+ \\ b_j + s_{ij} & \text{otherwise} \end{cases}$$

Figure 10 shows the fit of the pick failure probabilities as a function of inventory levels at three stores.

Figure 10 One-step fit of pick failure probabilities at 3 stores as a function of the inventory

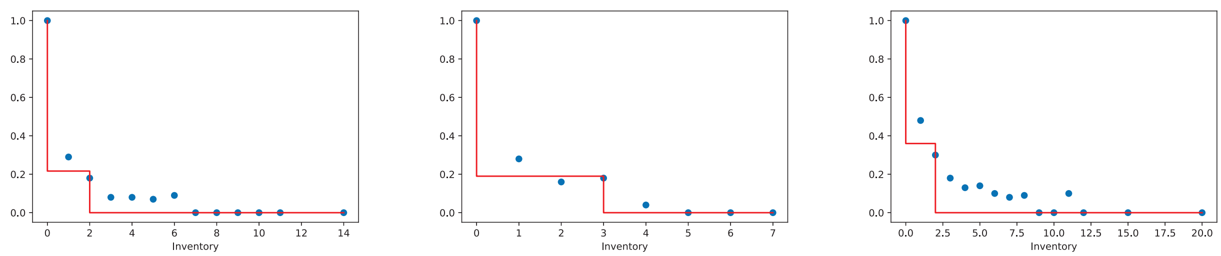

Sampling an order. We generate a random order for each retailer for the single-order fulfillment models (Sections 2 and 3) using the following procedure. Note that pick failure probabilities, labor costs and shipping costs at the store sampled for an SKU are all generated independently of each other.

1. *Number of stores that carry an SKU* (J). The number of stores that carry an SKU is sampled using a power law distribution with parameter values taken from Table 4 for each retailer. For example, if $J(B)$ is the random variable for the number of stores that carry an SKU at retailer B, then $\mathbf{Pr}\left[J(B)=k\right] \propto \exp(11.832)\cdot k^{-1.268}$ for $k = 1, ..., 500$.
2. *Pick failure probability at store j, (ϕ_j).* For each retailer, we pick the top 5 SKUs that have the highest volume of pick failure. For each of these 5 SKUs, we take the union of all stores that feature those SKUs. For each store in this collection, we fit the best 1-step pick failure probability function.

   To sample a store's pick failure probability (for any SKU),

   (a) We sample the store j itself from the list of selected stores in a way that the probability of picking that store is proportional to that store's sales volume.

   (b) *Sampling inventory* $\left(\overline{Q}_j\right)$. For each store, we sample the inventory of the store $\overline{Q}_j$ from a multinomial distribution modeling the inventory level.

(c) *Generating pick failure probability* (ϕ_j, ϕ_{jl}) For the static pick failure model (Section 2), the pick failure probability ϕ_j is simply obtained by applying the one-step pick failure function to $\overline{Q}_j$. For the dynamic pick failure model (Section 3), we set the physical demand at store j, $\mathscr{D}_j^P$ at each stage of fulfillment, to be $Poisson(0.6 * \overline{Q}_j)$.

$$\phi_{jl} = \mathbb{E}\left[\phi_j\left(\overline{Q}_j - \sum_{l'=1}^{l} \mathscr{D}_j^P\right)\right]$$

where $\phi_j(\cdot)$ is the one step pick failure probability distribution function.

3. *Labor cost at store j* (b_j). As described previously, we sample the store with a probability proportional to its volume and assign all SKUs carried by that store a labor cost equal to the minimum wage at that store location divided by a scaling factor of 4.

4. *Shipping cost at store j* (s_j). For each store associated with the order, we sample its zone independently from the distribution in Figure 8 and then assign the shipping cost from USPS website (**?**). We assume each item weighs 4lbs.

5. *Delayed cancellation cost* (d). The delayed cancellation cost is set to $25 per order at any store.

For multi-order models (Sections 4 and 5), we fix the number of stores, J, to 10 and the number of customer shipping zones, I, to 7. The initial inventory $\overline{Q}_j$ at store j is sampled from the multinomial of inventory levels at stores. The physical demand $\mathscr{D}_j^P$ for store j is $Poisson(0.6 * \overline{Q}_j)$. The online demand from zone i, $\mathscr{D}_i^O$ is $Poisson(0.5 * z_i * \sum_j \overline{Q}_j)$, where z_i is the proportion of orders from zone i calculated as shown in Figure 8. The cancellation cost c and lost-sales cost p are set to $15 and $10 per order respectively

Appendix C: Additional Computational Results

Sensitivity to delayed cancellation cost and the number of stages in single-order fulfillment model under dynamic pick failure probabilities (Section 3). We examine the sensitivity of the cost savings obtained from modeling pick failure to the parameters of our models. We study the variation of savings obtained by our efficient dynamic programming algorithm with respect to the baseline greedy algorithm which is optimal when pick failure is not accounted for. We analyze the savings for our single-order fulfillment model under dynamic pick failure probabilities (Section 3). Figure 11a shows that the average savings increases as the delayed cancellation cost increases. This is because the greedy algorithm is oblivious to pick failure and therefore does not account for the delayed cancellation cost. Figure 11b shows that the savings decrease as the number of stages that an order is tried increases. In other words, the value of modelling pick failure in high when the number of stages of fulfillment is low.

Figure 11 Sensitivity of savings to delayed cancellation cost (d) and number of stages (L)

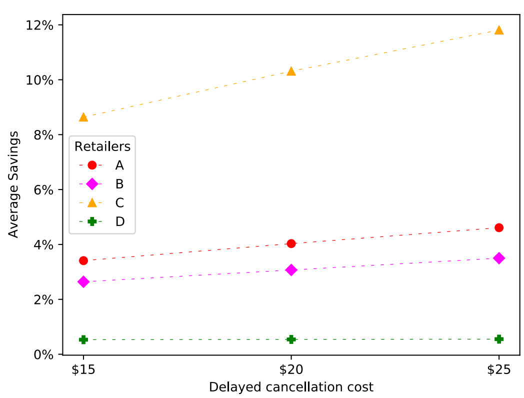

(a) Average savings across delayed cancellation costs ($L = 3$)

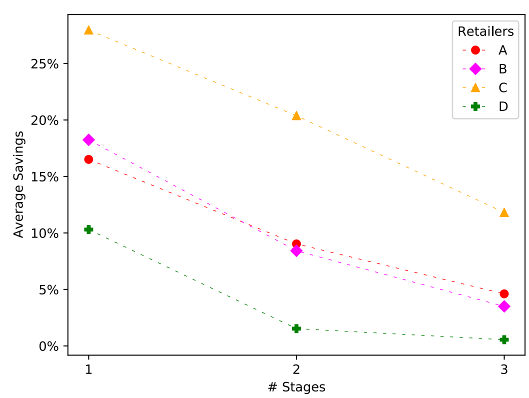

(b) Average savings across number of stages ($d = \$25$)

Sensitivity to physical demand in multi-order fulfillment model (Section 4). We vary the means of the physical demands ($\mathbb{E}\left[\mathscr{D}_j^P\right]$) as a percentage of available inventory at stores at the beginning of the day ($\overline{Q}_j$).

Figure 12 **Variation of savings with respect to physical demand as a fraction of inventory when** $d = \$25$ **and** $L = 3$

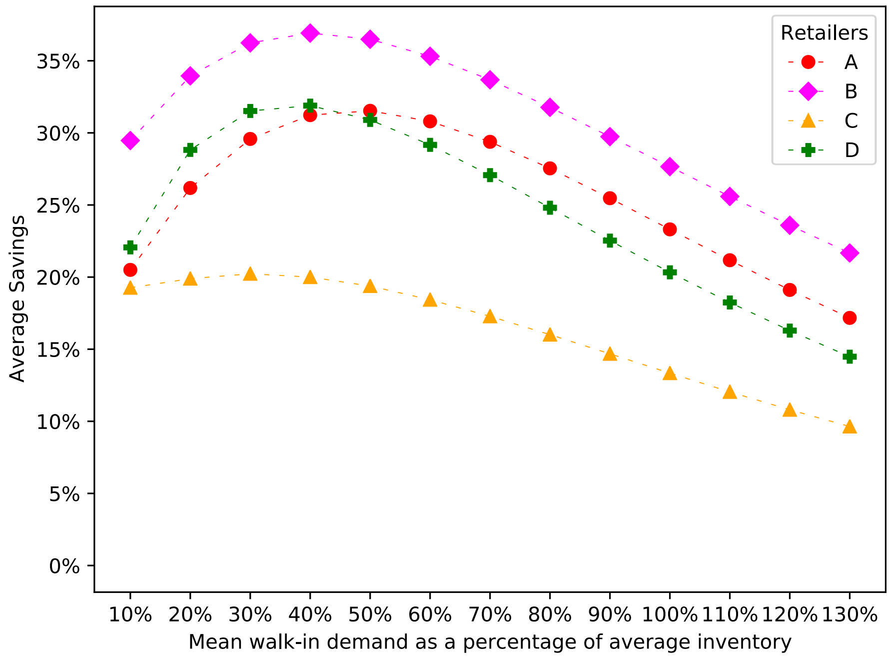

We observe in Figure 12 that the savings peak when the means of the physical demands are around 40% of initial inventory at the stores. When the physical demand is low, the value of modeling pick failure drops because of lower pick failures at higher inventory levels. On the other hand when the physical demand is significant with respect to the inventory at stores, the optimal policy that takes into account pick failure is relatively ineffective since the store demand depletes inventory much earlier and the delayed cancellation cost becomes a dominant portion of the fulfillment costs.

Comparison of expected numvber of trials of our policies. In addition to comparing costs of our optimal policies with respect to the benchmark policies, we compare the expected number of trials for our optimal policy against the benchmark policies for models in Sections 2 and 3 in Figure 13.

Figure 13 **Improvement in average expected number of trials**

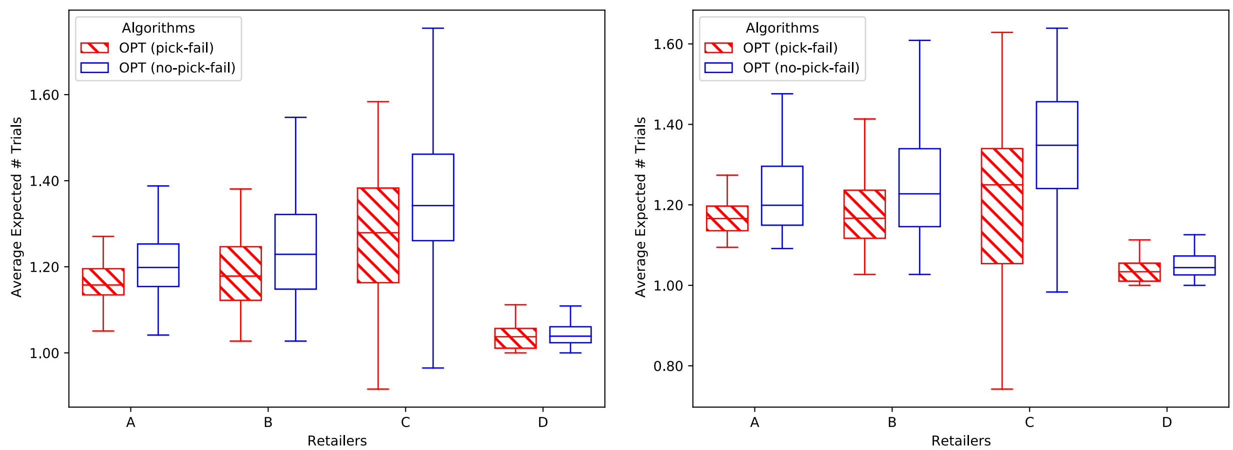

(a) Single-order fulfillment under static pick failure probabilities (Section 2)

(b) Single-order fulfillment under dynamic pick failure probabilities (Section 3)

We use an example to explain the computation of expected number of trials for a policy. Let's consider a 3-stage problem with a fulfillment policy with stores, say store 1, store 2 and store 3 in order. Let $(\phi_1, \phi_2, \phi_3) = (0.2, 0.6, 0.5)$. The expected number of trials for fulfillment policy $[1,2,3]$ is $1*(1-0.2)+2*0.2*(1-0.6)+3*0.2*0.6*(1-0.5)$.