Making Better Fulfillment Decisions on the Fly in an Online Retail Environment

JASON ACIMOVIC

Smeal College of Business, Penn State University, acimovic@psu.edu

STEPHEN C. GRAVES

Sloan School of Management, Massachusetts Institute of Technology, sgraves@mit.edu

Additional Authors
Co-authors at our anonymous industrial partner

This Draft: January 7, 2013

Relative to brick-and-mortar retailers, online retailers have the potential to offer more options to their customers, with respect to both inventory as well as service times. To do this entails the management of a distribution network with more decision options than a traditional retailer. The online retailer, not the customer, decides from where items will ship, by what shipping method, and how or whether multiple-item orders will be broken up into multiple shipments. Furthermore, online retailers often carry many low-volume items, and do not stock each item at every warehouse location. One question facing online retailers is this: what is the best way to fulfill each customer's order to minimize average outbound shipping cost? We partner with an online retailer to examine this question. We develop a heuristic that makes fulfillment decisions by minimizing the immediate outbound shipping cost plus an estimate of future expected outbound shipping costs. These estimates are derived from the dual values of a transportation problem. In our experiments on industry data, we capture 36% of the opportunity gap assuming clairvoyance, leading to reductions in outbound shipping costs on the order of 1%. We also characterize the attributes of SKU's most conducive to benefitting from the heuristic.

Key words: Online retailing, inventory management, dynamic allocation, fulfillment policies

1 Introduction

In 2011, sales of items paid for over the internet in the US brought in revenues of \$202 billion (Forrester Research, Inc. 2012c). This number represents a 14% increase in sales over the previous year, and is expected to grow to \$327 billion in 2016, constituting 9% of US retail sales (Forrester Research, Inc. 2012c). Growth rates are similar in Europe, where online revenues in 17 major markets are forecast to grow from \$130 billion USD in 2011 to \$230 billion USD in 2016 (Forrester Research, Inc. 2012b). In China, the growth is much faster; business to consumer online transactions (B2C) in 2011 reached \$38 billion USD, an increase of over 130% from the previous year (Seeking Alpha 2012). If consumer to consumer transactions (C2C) are included with B2C, online retail sales in China are forecast to triple from \$118 billion in 2011 to \$356 billion in 2016 (Forrester Research, Inc. 2012a). The online retail businesses serving this growing customer base operate very differently from brick-and-mortar retailers, and require a new set of tools to run efficiently. Learning to take advantage of these differences and to better manage online operations will become increasingly important as the sector continues to grow.

One important aspect of online retailing non-existent in brick-and-mortar retailing is fulfillment: picking, packing, and shipping orders to individual customers. One element of fulfillment, outbound shipping, can by itself incur significant costs. Based on 10-k statements of several online retailers who report shipping revenue data (Amazon.com, Inc. 2011, 2012; Bluefly, Inc. 2011, 2012; Vitacost.com, Inc. 2011, 2012), outbound shipping revenues (shipping costs charged to the customer) can vary from 3.2% to 6.6% of sales. If we were to extrapolate this industry-wide based on 2016 projected online retail revenue of \$327 billion in the US, it would correspond to outbound shipping revenue of between \$10 and \$21 billion. The actual amount that an online retailer itself *spends* on outbound shipping may exceed what it *charges* customers, especially with the increasing popularity of "free" shipping (Jannarone 2011). For instance, Amazon.com spent 116% and 157% more in outbound shipping than it brought in for 2010 and 2011 respectively (Amazon.com, Inc. 2011, 2012), and it is reasonable to assume that other online retailers offering free shipping are also spending at least what they bring in in revenue. Reducing outbound shipping costs can have a significant impact on total costs. In this paper, we study the impact of

smarter, forward-looking fulfillment decisions on outbound shipping costs in an online retail environment.

In the traditional retail supply chain, vendors typically supply distribution centers which in turn supply retail stores. Customers pick the retail store to visit, and buy items from the inventory on the shelf at the time of their visit. All customers are served immediately, namely, as they check out with their new purchase. In general, assortment is limited by the physical space of the store itself.

Online retailing supply chains, on the other hand, may appear similar to the customer, but actually differ from conventional retailing in several key areas. Instead of a storefront with a backroom, online retailers keep their inventory in fulfillment centers. Within these centers, orders are picked off of shelves, aggregated, packed, and shipped to customers. Additionally, although multiple distribution echelons still exist, the structure is not strictly hierarchical. The distribution network may consist of large fulfillment centers designed to hold a wide variety of stock keeping units (SKU's), as well as small fulfillment centers designed to maximize geographical coverage for the most popular SKU's. Any of these fulfillment centers can serve any customer, and they can even replenish each other.

Besides the structure of the distribution network, online retail supply chains differ from brickand-mortar supply chains in four other ways. First, in online retail supply chains, the online retailer
decides from where to fulfill an order, not the customer. Second, there always exists a time delay
between the placing of an order and the fulfillment in online retail supply chains. Third, in online retail
supply chains, customers usually have an option to choose their service or delivery times, e.g., next day
versus next week, depending on what they are willing to pay. Lastly, the online retail customer often has
access to any item that exists in the network. If the building nearest the consumer is out of a particular
item, the order will be sent from a further location at an additional cost to the seller, but at no additional
cost to the buyer. In contrast for a brick-and-mortar retailer, a stock-out in a particular building can lead
to a customer not buying anything (lost sale), or buying a substitute product, or going to another store to
find the desired product.

These differences suggest that new strategies may be beneficial in managing online retail supply chains that are not applicable or necessary in brick-and-mortar supply chains. Specifically, we focus on how an online retailer should choose the specific facilities from which to fulfill each order in order to minimize average outbound shipping costs, a choice not possible in brick-and-mortar supply chains.

2 Problem Definition

This research grew out of a partnership with a large American-based retailer that sells a broad catalog of physical items online and operates a network of fulfillment centers around the US. Their stock varies in cost and popularity, with some items selling thousands of units in a week, and others selling a dozen units over the course of a year.

Our industrial partner, like many online retailers, makes its fulfillment decisions in real time, both for operational reasons as well as to provide shipment options and delivery commitments to the customer in a timely manner. We assume that our industrial partner makes these decisions myopically: the online retailer fulfills each order the cheapest way possible based on its current inventory position, without accounting for any cost implications for fulfilling future orders.

In this paper, we investigate the extent to which we might improve the performance of the myopic policy with an implementable heuristic. By implementable we imply both computationally tractable and intuitive to the extent necessary both to write flexible code and to sell the idea to business managers. We assess the benefit from making decisions that minimize the sum of the current outbound shipping cost plus an estimate of future expected outbound shipping costs incurred as a result of the new inventory position. What follows is an illustrative example outlining the possible pitfalls of a myopic policy.

Imagine two fulfillment centers (FC's): one in Los Angeles and one in Nashville. The Los Angeles facility has 3 textbooks left in stock, while the Nashville facility has one textbook and 9 CD's in stock. Over the course of the next day, two customers will arrive who each wants his order delivered within three days: one in Dallas wanting a textbook, and one in Washington, DC wanting a textbook and

a CD (although the system is unaware of these customers at the outset of the day). Figure 1 shows the costs of shipping each item or combination of items from each facility to each customer. These costs were retrieved from www.ups.com on March 8, 2011. They represent the cost to send a one pound package to a residential address within a 3-day window. The \$12.12 figure represents the cost to send a two pound package from Nashville to Washington, DC.

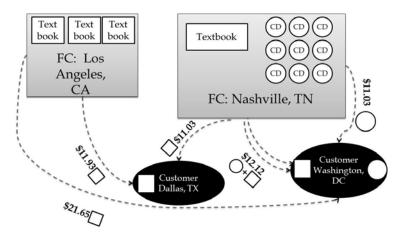


Figure 1: Example of myopic fulfillment with shipping costs

If the Dallas customer arrives first, the online retailer (acting myopically) will ship the textbook from Nashville rather than Los Angeles, saving \$11.93 - \$11.03 = \$0.90. This depletes the textbook inventory at Nashville, and it has only nine CD's remaining. Then the Washington, DC customer arrives, wanting a textbook and a CD. Nashville no longer has the text book; hence, the text book must ship from Los Angeles, and the CD must ship from Nashville, for a total cost of \$21.65 + \$11.03 = \$32.68. The total fulfillment cost for the myopic fulfillment policy (MYO) is \$11.03 + \$32.68 = \$43.71.

If the online retailer could have seen the future, it would have fulfilled the Dallas customer's order from Los Angeles and the Washington, DC customer's order from Nashville, at a total cost of: \$11.93 + \$12.12 = \$24.05, a little over half the cost of the myopic cost. We call this the perfect hindsight policy (PH).

In the above example, both the Dallas and Washington, DC customers paid a premium to receive their orders within 3 days. These premiums did not depend on the cost incurred by the online retailer, so

that it was in the retailer's best interest to fulfill the orders on time as cheaply as possible. Any savings in shipping costs go straight to the bottom line.

We assume that customers have options with respect to how fast they want their items, with shorter delivery times corresponding to higher shipping fees (regardless of the actual fulfillment cost). The online retailer has several options with respect to how to actually ship items to customers. Faster shipping modes incur higher shipping costs on the part of the online retailer. We note that the online retailer need not use a fast shipping mode to serve a customer who requests a short delivery window. If the items in a customer's order are in a facility nearby, the online retailer may use a relatively cheap ship mode, even if the customer requests the items very quickly. Thus, a large savings can be realized not only by shipping items shorter distances, but also by using cheaper *modes* of transportation, namely, choosing trucks over airplanes whenever possible.

The objective of the online retailer is to choose fulfillment centers to serve each customer's request in such a way that minimizes long term average outbound shipping costs. Our contribution in this paper is to develop an order-fulfillment heuristic, demonstrate that it performs well on industry data, and show that it has desirable theoretical properties. The heuristic utilizes dual variables from a transportation linear program, and has the potential to run quickly and be implemented in real-time decision making systems. We also characterize for which types of SKU's the heuristic works best.

3 Literature Review

The relevant literature can be broken into four categories, none of which is specifically related to online retailing: rationing for multiple customer classes, emergency lateral transshipments among multiple depots, dynamic and approximate dynamic programing, and airline network revenue management.

There is a rich literature on rationing inventory in the presence of multiple customer classes, albeit mostly for a single warehouse node. In these cases, customer classes are defined by their priority levels, and each level has a desired fill rate, or service level. For each class, a "support level" is set, such that when the total inventory drops below a customer class' support level, all demand for that class is

backordered. The characteristics of this system are explored in Nahmias and Demmy (1981), building on previous work by Kaplan (1969) and Veinott (1965). In this stream of literature, customers are prioritized, and the inventory system is allowed to either backorder or lose demand for low priority customers in order to fulfill future demand for high priority customers. In online retailing, however, classes are defined by time window requests, and all demand must be satisfied within a requested time window if there is inventory in the system. Otherwise, sales are lost. Additionally, even if good rationing policies could be set for a specific instance of inventory positions in a network of fulfillment centers, this policy would most likely change significantly if the inventory positions in the system changed.

When one fulfillment center serves a customer who lives nearer to another facility, this may be modeled as an emergency lateral transshipment. Oftentimes in emergency lateral transshipment models, the cost of the transfer is high, the lead time is assumed to be negligible, and backorders are allowed. These problems were studied by Lee (1987) and Axsater (1990), who both developed inventory allocation approximations for multi-echelon systems with repairable items. Axsater (2003), for example, develops a decision rule dictating whether to transship or not, or whether to incur the backorder costs. Yang and Qin (2007) discuss a model that utilizes virtual lateral transshipments between two factories. This is similar to online retailing in that inventory need not travel from fulfillment center A to B, then to the customer to be considered a transshipment, but instead may be shipped directly from A to the customer in region B at increased cost. They analyze for a two-factory model both replenishment decisions and fulfillment decisions. However, the approach is complex and limited in the number of fulfillment centers allowed.

Archibald et al. (2009) develop an index heuristic for a multi-location inventory system, inspired by a real problem facing a tire retailer in Scotland and Northern England. For the fulfillment portion of the problem, they estimate the difference in the cost-to-go function if a unit of inventory is depleted. In spirit, this is what we aim to do, but our approaches differ from Archibald et al. These authors estimate the cost-to-go function differences by looking at all pairwise comparison networks – which are easier to analyze – and by using that data to estimate the cost-to-go in the larger system.

Most of the existing emergency lateral transshipment literature deals with optimal inventory allocation policies assuming that myopic fulfillment policies will be used to meet demand. We assume that we are given a (possibly sub-optimal) inventory allocation, and from this allocation we are to determine the best fulfillment policy in order to (effectively) avoid future virtual lateral emergency transshipments.

Generally, the problem of determining an optimal fulfillment policy falls into the broad class of optimal dynamic resource allocation. The system must allocate inventory to customers as soon as they request an item, while simultaneously minimizing future expected costs. While dynamic programming has the ability to solve this class of problems, the dimensionality of the state space prevents obtaining solutions in a reasonable amount of time. Neuro-dynamic programming (D. Bertsekas and J. Tsitsiklis 1998) and approximate dynamic programming (Powell 2007) utilize techniques to estimate the value function in a dynamic program, producing sub-optimal but tractable solutions that perform well in practice (Van Roy et al. 1997; Simao et al. 2009; Maxwell et al. 2010).

For the heuristic we develop, our methodology draws on approximate dynamic programming in that it approximates a value function for future expected cost-to-go. However, our method for estimating the value function is more similar to the revenue management literature. Specifically, we build on their use of a linear programming relaxation assuming future deterministic demand to derive these estimates. A linear program popular in practice first proposed by Simpson (1989) and Williamson (1992) matches flight legs to itineraries (where an itinerary might consist of multiple legs) such that the future expected demand per itinerary is not exceeded and overall revenue is maximized. If an itinerary's revenue does not exceed the sum of the bid prices (equivalent to the dual values of this linear program) of the legs of that itinerary, then that itinerary would not be offered to customers. Talluri and van Ryzin (1998) analyze this regime and show asymptotic optimality in the case of large demand and supply.

However, while similarities exist, the problem we examine diverges from airline network revenue management in several aspects. First, the nature of the choice to be made is different. In airline network revenue management, each itinerary must either be offered or held back in each period. In online

retailing, each customer will be served if possible, and the decision must be made *how* to serve each customer's demand. Additionally, the airline network revenue management problem has a finite horizon: once the airplane takes off, all the inventory disappears. In an online retail environment, the problem has an infinite horizon. Inventory is replenished regularly, and inventory may never be completely drained from the system. Third, in airline network revenue management, the customer purchases a specific itinerary made up of particular legs, and cares how that itinerary is composed. In online retailing, a customer places an order for a set of items, and does not care how those items are delivered. In this sense also, multiple item orders must be considered in online retailing in a way they do not need to be considered in airline network revenue management.

We use similar proof techniques as Talluri and van Ryzin (1998) to show asymptotic properties of the linear program used in our heuristic. Our contribution is to apply the basic principles to a new context (online retailing), to formulate the linear program in a way that approximately accounts for multi-item orders, to show the asymptotic properties of this linear program, and to demonstrate that the resulting heuristic works well in practice.

4 Problem formulation

We can formulate the online retailer order-fulfillment problem as a dynamic program that minimizes the immediate outbound shipping cost plus the resulting future expected cost. Each arriving order requests a set of items to be delivered to a specific location by a due date. The action space includes any minimal subset of fulfillment centers that can satisfy the order; that is, the subset of fulfillment centers can deliver each item in the order by the due date, possibly requiring multiple shipments. Only outbound shipping costs are considered: in general, it is more expensive to ship an item by air than by ground, and it is more expensive to ship a multi-item order in multiple packages than to ship it in a single package from a single fulfillment center. The state of the system is defined by the inventory of every item in every fulfillment center, the timing and quantity of every inbound shipment in the pipeline, as well as the current time. Informally, we can express the optimal value function *J* as:

$$J(S,O) = \min_{u \in U} C(u,O) + E_{\tilde{O}} \left[J\left(f(S,u), \tilde{O}\right) \right]$$

where O is a realized order, S is the system state, u is the order fulfillment decision, C(I) is the cost to fulfill order O from fulfillment center(s) u, f(I) defines how the state S evolves from a given action U, and U is the set of feasible actions. The expectation is taken over all possible future realizations of orders within some fixed time period.

The order O encompasses the region of the customer, the time request of the customer (how fast she wants her items), and the different items requested in the order. The cost function C(t) takes into account the distance items are shipped, the mode by which they are shipped, and whether multiple items were bundled together or not. That is, if a customer ordered multiple items, the function would return one cost if the order were shipped in a single box from Maine and another cost if the order were shipped in two packages, one from Alabama and one from Texas. We assume that every customer request is feasible due to the existence of next day shipping from anywhere to anywhere in the United States, although costs may vary significantly. The function f(t), in determining state evolution, encapsulates not only from where the next order might come, but also when. Because the state includes not only what is on-hand in the system, but also what is on order, it is important to take into account whether the next customer request comes before or after the inventory in the pipeline from a vendor arrives into a specific fulfillment center.

Solving the above dynamic program is intractable. A realistically sized problem with 10 fulfillment centers, n items, a look-ahead period of 10 days, and an inventory level of 10 per item per fulfillment center results in 10^{20n} states. Even with n=1 the state space is very large. In proposing a tractable decision-making framework, we make several approximations:

- 1. The expected value of the cost function can be decomposed into the individual items that make up the order *O*
- 2. The differences in the decomposed cost functions can be approximated by a linear program's dual values (π) associated with inventory constraints of fulfillment centers

3. The system state S for a given item can be approximated by the inventory position (on hand plus pipeline) of each fulfillment center

We note that the assumptions above are not to be taken for granted, especially the first one. The assumption of decomposition assumes much, but we feel is necessary to make this very large problem tractable. As we discuss below, we still approximately take into account multi-item orders even with this first assumption.

Based on these assumptions, we propose the following heuristic to decide the order fulfillment decisions for each order:

$$\hat{u} = \underset{u \in U}{\operatorname{arg\,min}} \quad C(u, O) - \sum_{o \in O} \pi_o^{i(u, o)} \left(X_o \right) \tag{1}$$

where o is a specific item in order O, i(u,o) is the specific fulfillment center that action u assigns to item o, X_o is the inventory position vector for item o, and the π 's are the dual variables associated with the fulfillment center inventory constraints of the transportation problem described below.

In section 5 we describe the linear program from which we obtain these dual values. We report in section 6 our testing of the heuristic on data from our industrial partner.

5 Linear programming heuristic formulation

Here, we describe the linear program (LP) from which we obtain the dual values for the fulfillment heuristic given in (1). The LP itself is a fluid, deterministic approximation of the optimal assignment of inventory to customers for a single SKU.

Even though each linear program represents an estimate of the expected cost-to-go for a single SKU, expected demand for that SKU *along with other items* is still accounted for approximately in the formulation. Because only a limited number of linear programs need to be solved at each decision epoch for each SKU, the resulting LP heuristic has relatively fast performance and is practical to implement.

One beneficial aspect of the linear program is the way it accounts not only for the geographical location of a fulfillment center, but also for the size of the catalogue (number of unique items) at a

fulfillment center. Accounting for both is important. We saw in the example in Figure 1 why geography is important to consider when making fulfillment decisions: qualitatively, inventory should be valued higher at centrally located facilities. We saw in the same example also that inventory should be valued at fulfillment centers that have a lot of other items on-hand. When customers order multiple items at once, considerable savings can be achieved by shipping these orders in a single package. If all else were equal, it would be better to keep a unit of inventory at a large fulfillment center with many kinds of items rather than one that held only a few unique items. The former fulfillment center would have a higher probability of being able to ship a random multi-item order in a single package.

The LP is a transportation problem that matches inventory to expected demand for a specific item. We define the inventory position for the item at each fulfillment center as the current on-hand inventory plus all inbound inventory (on-order or pipeline inventory) over the next n days, where we term n to be the look-ahead period. We denote the system inventory position for the item by the vector X^n , where the i^{th} element corresponds to the i^{th} fulfillment center and represents the i^{th} supply node of the transportation LP. This is a simplification as we assume that we can represent with a single number all of the information about the on-hand and inbound inventory for the next n days, and that we can ignore any inbound information beyond n days.

For the demand nodes, we separate the United States into distinct geographical regions. We further divide each region into several possible customer delivery time options. Then for each pair (region, customer option) we have two demand nodes: one for single-item order demand and one for multi-item order demand. This accounts for the approximate handling of multi-item orders even though each LP represents a single SKU. The single-item demand node represents the demand for the specific SKU when the SKU is ordered by itself; the multi-item demand node represents the demand for the SKU when it is ordered with other items. Thus, one node might be (Chicago, NextDay, Single), while another might be (West Kansas, EightDay, Multi). We specify the model's indices, parameters and variables as follows:

n – Look ahead period in days
 I ∋ i – Set of fulfillment centers (FC's)

 $J \ni j$ – Set of customer regions

 $M \ni m$ – Set of customer delivery time options

 X_i^n — On hand inventory in FC i plus inventory arriving over next n days

Forecasted system daily demand

 λ_m - Proportion of customers of type m requesting multiple items

 ρ_i - Probability FC *i* has 'other items in order'

 $\omega_m \in (0,1]$ - Expected discount of sending a multi-item order in one package

(calculated as the average of one over the number of items in a package)

 α_{im} - Fraction of total demand that is region j, type m

 c_{ijm} – Cost from FC *i* to customer *j* of type *m*

 W_{iim} – Decision variable for flow from FC i to single-item customer (j, m)

 x_{ijm} — Decision variable for *unsplit* flow from FC *i* to multi-item customer (j,m) — Decision variable for *split* flow from FC *i* to multi-item customer (j,m)

The expected demand over the look-ahead period within a specific region *j* and of a given

delivery time m is $\alpha_{jm}dn(1-\lambda_m)$ for single-item orders and $\alpha_{jm}dn\lambda_m$ for multi-item orders.

The transportation problem has a single un-capacitated arc between each supply node and each single-item-order demand node. The cost for this arc represents the shipping cost from the fulfillment center to the customer region by the cheapest mode that will satisfy the delivery time.

The transportation problem has two arcs between each supply node and each multi-item-order demand node. One arc corresponds to satisfying the multi-item order with a single shipment; the second arc corresponds to splitting the multi-item order into multiple shipments. The cost for the multi-item single-shipment (multiple-shipment) arc is ω_m ($2\omega_m$) times the shipping cost from the fulfillment center to the customer region by the cheapest mode that will satisfy the delivery time. As ω is calculated as the average of the inverse of the number of items in a multi-item order, then the fraction ω itself represents the proportion of the shipping cost assigned to each item in the single shipment. When the multi-item order is split into (by assumption) two shipments, we assume that the specific SKU is part of a shipment of size approximately $1/2\omega$ items, and hence 2ω is its proportion of the cost. The single-shipment arc is capacitated to reflect the likelihood that the fulfillment center can fulfill the order with a single shipment. We set the capacity equal to the expected number of multi-item orders that can be fulfilled from a given

fulfillment center, based on that fulfillment center's ability to satisfy the "other items in the order". The capacity for the single-shipment arc for multi-item orders from fulfillment center i to customer region j with delivery time m is $\rho_i \alpha_{jm} dn \lambda_m$, where ρ_i is the probability that fulfillment center i has the "other items" in a random multi-item order. The formulation of the transportation linear program is:

$$\mathbb{C}^{TRANS-LP}\left(\boldsymbol{X}^{n}\right) = \tag{2}$$

$$\min_{x,y,z} \sum_{i,j,m} c_{ijm} w_{ijm} + \sum_{i,j,m} \omega_m c_{ijm} x_{ijm} + \sum_{i,j,m} 2\omega_m c_{ijm} y_{ijm}$$
(2-1)

$$s.t. \quad \sum_{j,m} w_{ijm} + \sum_{j,m} x_{ijm} + \sum_{j,m} y_{ijm} \leq X_i^n \qquad \forall i$$
 (2-2)

$$\sum_{i} w_{ijm} = \alpha_{jm} dn \lambda_{m} \qquad \forall j, m \qquad (2-3)$$

$$\sum_{i} x_{ijm} + \sum_{i} y_{ijm} = \alpha_{jm} dn (1 - \lambda_{m}) \qquad \forall j, m \qquad (2-4)$$

$$x_{ijm} \le \rho_i \alpha_{jm} dn \lambda_m \qquad \forall j, m$$
 (2-5)

$$W_{iim}, X_{ijm}, Y_{ijm} \geq 0 \qquad \forall i, j, m \tag{2-6}$$

The decision variables w, x, and y represent the amount of flow along the arcs for single-item, unsplit multi-item, and split multi-item orders respectively. The objective value minimizes the cost to ship the specific SKU if it is ordered by itself, if it is ordered with other items and ships with those items, and if it is ordered with other items and that order has to be split. Constraints (2-2) ensure that no fulfillment center ships more inventory than it has. Constraints (2-3) and (2-4) require both single item and multi-item demand to be met, while constraints (2-5) limit the number of multi-item orders that can be shipped as a single shipment. The above formulation presumes that supply is sufficient to meet demand; if this is not valid, we scale down the demand in each region so that total supply equals total demand.

We propose that the objective value of the LP in formulation (2) is a good approximation for the disaggregated cost-to-go function J. Thus, the dual values associated with the inventory constraints (2-2) provide an estimate of the marginal value of a unit of inventory at each fulfillment center over the lookahead period. Therefore, we can use these dual values to approximate the differences in the resulting

cost-to-go functions as a function of the chosen fulfillment center. We propose the following approximation

$$\Delta_{i(u,o)}\left(E_{\tilde{o}}\left[J\left(f(S,u),\tilde{O}\right)\right]\right) \approx \pi_o^{i(u,o)}\left(X_o\right) \tag{3}$$

where $\Delta_{i(u,o)}(\cdot)$ denotes the change to the value function from a unit decrease in the inventory of the specific SKU at fulfillment center i chosen by fulfillment decision u for item o, and $\pi_o^{i(u,o)}(X_o)$ is the dual value for the inventory of the specific SKU at fulfillment center i, for given inventory position X_o . As such, the order fulfillment LP heuristic we proposed in equation (1) uses the dual values associated with constraints (2-2).

6 Analysis of heuristic performance

In this section, we test the LP heuristic on data obtained from our industrial partner. We compare its performance to a myopic fulfillment policy and to a perfect hindsight policy that can see all future orders. On this dataset, the perfect hindsight opportunity gap is about 3%, and the LP heuristic captures about a third of this 3%, resulting in a 1% reduction of outbound shipping costs.

We also find that the LP heuristic, in addition to incurring less shipping costs over the course of our simulation, also leaves inventory better balanced at the end of the period, suggesting that even more savings might be realized.

6.1 Analysis assumptions

We make several assumptions that allow the analysis to be tractable. These assumptions were made along with our industrial partner in order to balance accuracy with problem size.

First, we disaggregate SKU's but still approximately account for multi-item orders, we measure proportional improvement on a sample and extrapolate this to the entire dataset, we assume orders ship the day they arrive to the system, and we assume replenishments are exogenous and identical as they were in the actual system.

Ideally, to test the performance of the LP heuristic, one would simulate it using a dataset representing all of the sales that an online retailer faced over a long period of time. However, large online retailers may hold in their warehouses somewhere on the order of one million SKU's, serving millions of orders in a month, and shipping out several times that in number of raw units (assuming many customers request several items per order). Simulating a realistically-sized system would be very computationally intensive. Therefore, we examine a sample of SKU's, and look at all sales over four weeks for each of these SKU's.

We simulate one SKU at a time within our sample of SKU's. We take a SKU in isolation and simulate the performance of the different fulfillment decision-making policies on this product.

Afterwards, we aggregate the results across all SKU's in our sample in order to estimate the impact of a particular fulfillment policy on the entire system.

Examining a sample of individual SKU's is an approximation, not only because of statistical sampling error, but also because of the presence of multi-item orders. Because all customer orders for a specific SKU are examined, some of those orders may include only the specific SKU, some may include the specific SKU and other SKU's in the sample, and others may include the specific SKU and other SKU's not in the sample. In order to account for multi-item orders in our simulation, we alter the set of feasible fulfillment centers for each multi-item order. Based on actual inventory data, for each order k we set a variable k to 1 if fulfillment center k had the "other items in the order" (including every other SKU whether in the sample or not) on-hand the day order k was placed, and 0 otherwise. When performing the simulation, we define the feasible fulfillment centers for order k as those facilities that have positive on-hand inventory for the specific SKU, k and whose associated k as those facilities that have positive on-hand inventory for the specific SKU, k and whose associated k as those facilities that have positive on-hand inventory for the specific SKU, k and whose associated k sequal 1. In determining the shipping costs for the specific order k that are attributable to the specific SKU, we charge k of the cost to send a package, where k is the actual number of items in the order.

If no fulfillment center is feasible – i.e., if no fulfillment center with $Z_{ik} = 1$ has the specific SKU on-hand – then the order must be split. The specific SKU is shipped from a feasible fulfillment center as dictated by a specific policy, and we assign 2/r of the shipping cost to the specific SKU. This assumes

that the order can be fulfilled in two shipments. In our bookkeeping, we keep track only of inventory changes with respect to the specific SKU, not with respect to the "other items in the order". The Z_{ik} 's are fixed *a priori* and are not updated throughout the course of the simulation.

The performance metric of interest to us is proportional improvement. To this end, we ignore physical weight in our outbound shipping cost calculations. For each SKU, at the end of the simulation, we calculate the proportional improvement that a fulfillment strategy had over a myopic policy. We then calculate the overall system proportional improvement by taking an average of the improvements of all of the individual SKU's weighted by their demand.

Additionally, in order to not underrepresent high volume SKU's in our sample, we take a stratified sample of SKU's. High volume items make up a small fraction of SKU's, but a significant fraction of outbound volume, and we stratify our sample by sales volume.

We assume that orders ship as soon as they arrive to the system. In an actual system, this may be suboptimal. However, we believe that forcing orders to ship as soon as they arrive to the system best allows us to cleanly see the impact of implementing different fulfillment policies.

The inbound inventory amounts utilized in the simulation are calculated from actual system data, so that for the perfect hindsight optimization - as well as for the myopic and LP heuristic simulations - inventory arrives as it actually did in the real system. This is an approximation because in reality, a different fulfillment policy would require ordering a different amount of inventory into each fulfillment center. However, we assume that the myopic policy is a good proxy for the actual policy used by our partner, and that the inbound inventory resulting from the actual policy and data is not too far from what the simulated myopic policy would have ordered. If the LP heuristic fulfillment policy would have resulted in significantly different amounts of inbound inventory going into each fulfillment center, we believe that this would only improve the LP heuristic results because for that policy, inventory would then be better positioned.

6.2 Overview of data

Our industrial partner provided us with detailed records of order, shipment, and inventory data over 30 consecutive days of operations. From this, we built a data warehouse of each customer order (the items in a customer's order, the zip code of each customer, the order date of each customer, and the service time request of each customer), the actual fulfillment of each order (whether it was split, from where it shipped, by what method, and at what cost), as well as the on-hand and on-order (or pipeline or inbound) inventory in each warehouse every day. In our analysis and simulations, actual customer order data is used to simulate customer demand. Likewise, to determine fulfillment feasibility, actual on-hand and inbound inventory data were used.

We conducted our analysis on 12 fulfillment centers, spread across the United States. As mentioned above, we pick a random stratified sample of 2639 SKU's which in aggregate sold 1.5 million units over a four week period (for data cleaning purposes, we trim two days off of our 30-day data sample for the simulation). We chose a four week window because it reflects a compromise between an accurate representation of reality and computational tractability. For a sample of SKU's, we simulated the myopic fulfillment policy and also calculated the perfect hindsight cost as if we knew exactly what demand was going to occur in the future. We calculated the average ratio in costs between these two policies on one, two, three, and four weeks of data from our industrial partner. Even though the opportunity gap increased with longer time periods, it increased by smaller margins. Calculating the perfect hindsight cost requires solving an optimization problem which includes decision variables indexed by day. Solving these optimization problems for time periods longer than four weeks seemed less valuable because of the decreasing differences in the opportunity gap, and also required computational power that made testing many policies and tuning the heuristic cumbersome.

We exclude any SKU with sales volume of greater than 1250 per week for computational reasons; in extrapolating our findings, we assume that for SKU's that sold more than 1250 units per week, their improvement is equivalent to SKU's in our sample whose sales were close to 1250 units per week.

In Table 1we list some of the overall characteristic of the SKU's in our sample. One additional simplification is that we assume that each order associated with a single SKU requests exactly one unit of the SKU. In reality, some orders will request multiple units of an SKU, but these instances are relatively rare in this environment.

Table 1: Characteristics of our sample of SKU's

Characteristic	Value
Number of unique SKU's	2639
Number of orders placed	1.52 million
Average number of orders per SKU (in this stratified sample)	576
Average number of orders per SKU per week	144
Per week sales of slowest SKU	~1
Per week sales of fastest SKU	~1250
Number of unique SKU's involved in orders for SKU's in this sample	310,000

We assume that customers have four options with respect to delivery time: Next Day, Second Day, Four Day, and Eight Day. The online retailer has at its disposal four shipping options: Air Next Day, Air Second Day, Premium Ground, and USPS.

To simplify our data analysis, we made a couple of approximations. The cost of each of the online retailer ship options is represented by a linear function which increases with distance. Both the fixed and variable costs increase for each higher priority shipping mode: e.g., Air Next Day has a higher fixed cost and per mile cost than Air Second Day, Ground, or USPS.

We divided the United States into 3-digit zip code prefix regions (Zip3's), resulting in about 1000 customer zones. For our analysis, we approximate the cost of mailing a package from a facility to an address within a Zip3 region as being identical for any address within that region.

We also need to determine which shipping modes are feasible for a given combination of fulfillment center, customer location, and customer service delivery option. For the purposes of our simulation, we approximated the transportation times from point to point with the data in Table 2. We based this table on the empirical delivery times and verified with our industrial partner its accuracy for the purposes of this study, as well as with each carrier's own website.

Table 2: Minimum delivery time required for different ship mode options and distance ranges

Distance from fulfillment	Air 1-Day	Air 2-Day	Premium Ground	USPS delivery
center to customer	delivery time	delivery time	delivery time	time
0-250 miles			1 day	3 days
250-500 miles	1 days 2 days		2 days	4 days
500-750 miles	1 day	2 days	2 days	4 days
750+ miles			3 days	5 days

For instance, in our simulation, a Second Day delivery can be satisfied by Air 1-Day or Air 2-Day from any fulfillment center, and by Premium Ground from any fulfillment center within 500 miles of the customer. USPS service cannot be used, regardless of the fulfillment center location.

From the above data, we create a three dimensional array with elements c_{ijm} , where i represents the fulfillment center, j represents the three-digit zip code prefix of a customer, and m represents the customer's time request. For every i, j, m triplet, there may be up to four feasible shipping options available to the online retailer, or as few as one (where feasibility will be determined by the data in Table 2). The parameter c_{ijm} represents the cheapest of these feasible ship options.

The above approximations allowed our models to be tractable, without sacrificing too much accuracy. All of these assumptions were made in conjunction with our industrial partner, and were thought reasonable.

6.3 Transportation linear program input data

To solve the linear program formulated above, several parameters need to be defined:

n – Look ahead period in days

 X_i^n — On hand inventory in fulfillment center i plus inventory arriving over next n days

d – Forecasted system daily demand

 λ_m - Proportion of customers of type m requesting multiple items

 ρ_i - Probability fulfillment center i has 'other items in order'

 $\omega_m \in (0,1]$ - Expected discount of sending a multi-item order in one package

(calculated as the average of one over the number of items in a package)

 α_{im} - Fraction of total demand that is from region j, and of type m

For α_{jm} , λ_m , ρ_i , and ω_m , we use historical averages based upon all SKU's for which we have records in the sample (so that these parameters are non-SKU dependent). We define \widehat{Z}_{ik} to represent for each order which fulfillment centers have on-hand the "other items" in the order. The element \widehat{Z}_{ik} is 1 if fulfillment center i had the other items from order k on-hand on the day the order was placed, and zero otherwise. Actual data from our industrial partner is queried to populate this matrix. Then, the variables ρ_i are set equal of the percentage of orders for which the fulfillment center has on-hand all the other items

in an order: $\rho_i = \frac{\sum_k \hat{Z}_{ik}}{|\hat{Z}_i|}$, where the Note that we determine \hat{Z}_{ik} for all orders and all SKUs in the sample and hence, ρ_i is SKU independent.

To calculate *d*, we use sales data from the previous month for a given SKU. After this warm start, during the simulation, we update this forecast weekly using exponential smoothing based on observed sales. We set *n* equal to the day in the future with the lowest expected on-hand inventory in the system, based on current inventory, known inbound inventory, and expected system demand for the SKU. If expected on-hand inventory drops below zero, we set *n* equal to the highest indexed day which has positive expected inventory in the system.

In order to improve performance of the transportation LP, we aggregate customer regions in order to reduce the size of the instance. We start out with approximately 900 3-digit zip code prefix regions. We reduce this to 100 region clusters using k-means clustering.

6.4 Comparing fulfillment policies

Using the dataset from our industrial partner, we simulate three fulfillment policies: a myopic (MYO) policy, a policy based on the LP heuristic (LP), and a perfect hindsight (PH) policy.

When making fulfillment decisions, the three fulfillment policies adhere to the following logic: first, for a specific order k, the system attempts to ship all the items in this order in a single package by shipping the SKU from a fulfillment center that houses both the SKU itself as well as the other items in

order k. If shipping the order k in a single package is not possible, the system will ship the SKU from a fulfillment center that houses that SKU while ignoring the other items in the order. This split shipment incurs an additional cost in order to account for the extra package that must be shipped.

Specifically, the cost parameter c_{ik} in the simulation is defined as the contribution of a specific SKU to the cost of sending a package from fulfillment center i to the customer who placed order k.

The parameter c_{ik} and the set of feasible fulfillment centers A may each take on different values depending on the state of the system. If at the time order k is placed, there is a fulfillment center i where both $X_i > 0$ and $Z_{ik} = 1$ (representing the fact that the order can be shipped in a single package), then A is defined as the set of all fulfillment centers satisfying these criteria. We define $c_{ik} \equiv c_{ijm} / r_k$, where r_k is the number of items in order k, and c_{ijm} is the cost to send a package from fulfillment center i to the region j with speed option m which is dictated by the delivery time for order k. If there is no fulfillment center with $X_i > 0$ and $Z_{ik} = 1$, then the order k must be split into multiple shipments. The set k now includes k0 fulfillment centers that have positive inventory of the specific SKU. The cost parameter is then defined as k1 fulfillment centers that have positive inventory of the specific SKU. The cost parameter is

As each demand arrives to the system, both the myopic and LP heuristic policies choose a fulfillment center (FC) from which to fulfill based on the following logic with set A and the cost coefficients as outlined above:

$$FC^{MYO} \equiv \underset{i \in A}{\operatorname{arg \, min}} c_{ik}$$

$$FC^{LP} \equiv \underset{i \in A}{\operatorname{arg\,min}} c_{ik} - \pi^{i} \left(\boldsymbol{X}^{n} \right)$$

where π and X^n are defined as the dual values of the LP and the inventory position vector over n days respectively.

We simulate each of the policies on data from our industrial partner, one SKU at a time. For each SKU, we record actual order data from our industrial partner: the date, location, and time request of each order as well as the "other items" in the order if it was for multiple items. We also capture inventory data from the dataset provided by our industrial partner. We record the starting inventory of each SKU in each fulfillment center as well as inbound inventory. As each order arrives for the particular SKU, a fulfillment decision is made and a unit of inventory is depleted from the corresponding facility.

The total incurred costs for the LP heuristic and myopic policies are defined as C^{LP} and C^{MYO} respectively. The total cost of the perfect hindsight policy is calculated at once in a single optimization problem, and defined as C^{PH} .

During the course of the simulation, we solve the transportation LP only periodically for computational reasons.

6.5 Overall simulation results

The improvement gap is set equal to: $(\mathbf{C}^{\mathbf{MYO}} - \mathbf{C}^{\mathbf{PH}}) / \mathbf{C}^{\mathbf{MYO}}$. The performance of the LP heuristic relative to the myopic policy is defined similarly as: $(\mathbf{C}^{\mathbf{MYO}} - \mathbf{C}^{\mathbf{LP}}) / \mathbf{C}^{\mathbf{MYO}}$. Table 3 shows the results of these simulations.

Table 3:Proportional reduction in outbound shipping costs of perfect hindsight and LP heuristic policies over a myopic one

	Perfect Hindsight over Myopic	LP Heuristic over Myopic
Proportional Improvement (with 95% confidence interval)	$2.94\% \pm 0.14\%$	$1.07\% \pm 0.07\%$

SKU's that are high in sales volume tend to improve more than SKU's that have low sales. In Figure 2, we bucket the SKU's by volume – with the buckets defined by the sampling strata – and plot proportional improvement against this. We see in the figure that although the overall improvement of the LP heuristic is 1.07%, the improvement of the LP heuristic for high volume SKU's is about 2%. Likewise, while the overall perfect hindsight gap is a little under 3%, for high volume SKU's, the gap is more along the lines of almost 4%.

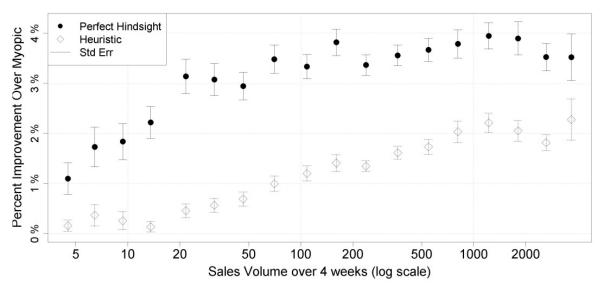


Figure 2: LP heuristic performance vs. volume of sales, bucketed by sample strata

We notice also in Figure 2 that the LP heuristic captures a larger portion of the gap as sales volume increases. For very fast moving SKU's, the transportation LP heuristic is capturing up to 50% of the possible improvement as defined by the perfect hindsight analysis.

The LP heuristic performs better on SKUs with less inventory. Figure 3 shows the proportional improvement from the LP heuristic versus scarcity, where scarcity is defined as the ratio of sales to the total inventory (on-hand plus inbound) that was available over the four week period. This makes intuitive sense. If inventory is very high everywhere, there is not much use to fulfilling smarter, because no facility will run out of inventory anyway. If inventory is very scarce, then it is more likely that several fulfillment centers will run out of inventory and it is more valuable to fulfill smarter.

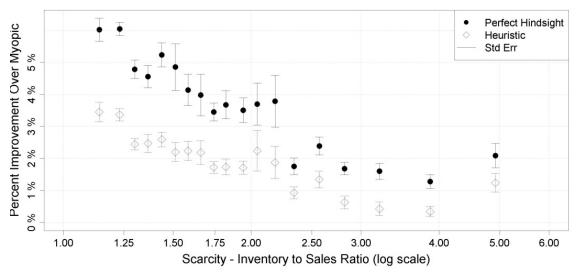


Figure 3: LP heuristic performance vs. inventory scarcity, bucketed into vigintiles (with ratios higher than 6 truncated from plot)

To put this in perspective, let us consider where Wal-Mart would land along the *x*-axis in Figure 3. According to recent financial data (Forbes.com 2012), Wal-Mart's inventory turns are about 8, meaning the cost of goods sold over a year is about 8 times the average inventory on-hand at any given time. This implies that in a month, the cost of goods sold is about two thirds of the on-hand inventory. Taking the inverse of this leads to an inventory-to-sales ratio over the course of 4 weeks of about 1.5. Thus, in this range, according to the above figure, the improvement due to the LP heuristic would be more along the lines of 2%.

We tested this heuristic over a wide variety of scenarios, and observed that the heuristic is robust to a wide variety of conditions including: fulfillment center catalogue disparity, forecast quality, and update frequency. We also perform experiments that show an additional benefit: the ending inventory position is left in a better state under the heuristic policy as compared to a myopic fulfillment policy; that is, the ending inventory is better distributed across the fulfillment centers and reflects the geographical demand distribution. This improvement could potentially lead to additional savings over the next period, in addition to the realized savings in the current period.

7 Asymptotic properties

Under certain conditions, the transportation linear program (TLP) is asymptotically optimal with respect to a perfect hindsight (PH) policy in ratio. This provides some theoretical justification for using the transportation problem to estimate the cost-to-go function in our approximate dynamic program. The overall proof technique is based on Talluri and van Ryzin (1998), Cooper (2002), and an observation from Gallego (1992) as noted by Talluri and van Ryzin. Our contribution is to apply the basic principles from airline network revenue management to a new context (online retailing) and to build on the original formulation of the linear program to approximately account for multi-item orders in this new context.

Some of our assumptions for the analysis are as follows:

- 1. Inventory is stocked and then demand is realized over a specific period of time. We operate in a fixed time horizon environment, with no replenishment of inventory.
- 2. If demand exceeds supply, it can be fulfilled from a dummy fulfillment center with an infinite amount of inventory and with a very high cost to each demand node. This cost is set high enough that it will be used only as a last resort. This is reasonable when "drop shipping" is an option, that is, where the online retailer subcontracts an outside vendor to make a shipment to a customer.

We simplify the problem to look at only one SKU at a time, even if customers order multiple items at once. As discussed above, we recognize this is a strong assumption.

Examining a single SKU, we break customer orders into two groups: those that ordered the specific SKU alone, and those that ordered the specific SKU with other items. A single-item order can be fulfilled from any fulfillment center that has inventory. A multi-item order can be served from any fulfillment center that has the SKU in stock as well as "the other items in the order". Without loss of generality, we drop the subscript *m* denoting the speed with which a customer requested his items. We define additional variables as such:

N,N'	_	Indices for the dummy fulfillment centers (FC's) with infinite inventory and
		high cost (The use of the index N' is useful in the proof of the lemma)
$K_j \ni k$	_	The set of possible individual customer orders in region j
$c_N > c_{ij} \ \forall i, j$	-	The cost from the dummy FC's to any customer
$D_k^{\{0,1\}}$	_	The random variable indicating whether or not customer k placed an order
D_{j}	_	A random variable representing the amount of demand from customer region j
		(Note that $D_j = \sum_{k \in K_j} D_k^{\{0,1\}}$)
$\widetilde{Z}_{ik} \in \{0,1\}$	_	Random variable indicating whether or not FC i has "the other" items on hand
		in the order of customer k . These are iid across k .
		(Note that for single items orders, every FC has "the other" items in the order)
$G_{ij} \equiv \sum_{k \in K_j: D_k^{\{0,1\}} = 1} \widetilde{Z}_{ik} \le D_j$	_	Random variable indicating the number of customers in region j for whom
		FC <i>i</i> has the other items
$\mu_j \equiv E[D_j]$	-	Expected demand in region j over the fixed time horizon
X_{i}	-	Starting inventory in FC i
\mathcal{X}_{ij}	-	Decision variable for inventory sent from i to j
θ	_	Scaling factor

Note that $D_k^{\{0,1\}}$ and D_j are part of the overall random process D. We define additional assumptions as follows:

- 3. \widetilde{Z}_{ik} does not change with different fulfillment assignments (e.g., Perfect Hindsight versus Myopic), and depends solely on the fulfillment center in question. This is a reasonable assumption if there is a very large catalog of items, each of which has approximately equal probability of being ordered with the specific SKU. While this is not strictly true in reality, making this assumption is reasonable and allows the problem to be tractable while still accounting for multi-item orders.
- 4. Whether or not customer *k* places an order is *independent* of whether fulfillment center *i* has the "other items" in that order; the number of customers who place orders in region *j* does not depend on whether or not fulfillment center *i* has the "other items" on-hand.

Additionally, we assume that all costs are non-negative, that is, $c_{ij} \ge 0$ for all i and j. Let D' and G' be the realizations of the random processes D and G respectively. Thus, $v_{PH}(D', G')$ is defined as the

minimum cost possible with perfect hindsight given these realizations, and equals the objective value of the following integer optimization problem:

$$\begin{array}{rcl} v_{PH}(D',G') & \equiv & \displaystyle \min_{x} & \displaystyle \sum_{i,j} c_{ij} x_{ij} \\ & s.t. & \displaystyle \sum_{i} x_{ij} \geq D_{j}^{'} & \forall \ j \\ & \displaystyle \sum_{j} x_{ij} \leq X_{i} & \forall \ i \\ & x_{ij} \leq G_{ij}^{'} & \forall \ i,j \\ & x_{ij} \in \mathbb{Z}_{\geq 0} & \forall \ i,j \end{array}$$

where $\mathbb{Z}_{\geq 0}$ denotes the set of nonnegative integers. Two dummy fulfillment centers N and N' are included in the indexing i, so that the cost of an assignment from N to j would be $c_N x_{Nj}$. The use of two dummy fulfillment centers here has no impact on the solution, but becomes useful in the proof. Likewise, we formulate the linear program and define v_{LP} as such:

$$v_{LP} \equiv \min_{x} \sum_{i,j} c_{ij} x_{ij}$$

$$s.t. \sum_{i} x_{ij} \ge \mu_{j} \quad \forall j$$

$$\sum_{j} x_{ij} \le X_{i} \quad \forall i$$

$$x_{ij} \le \mu_{j} \rho_{i} \quad \forall i, j$$

$$x_{ij} \ge 0 \quad \forall i, j$$

Let θ be a scaling parameter, and let $D^{(\theta)}$ be the random demand process such that:

$$\mu^{(\theta)} = \mu \cdot \theta$$

$$\sigma_j^{D,(\theta)} = \sigma_j^D \cdot \sqrt{\theta}$$

$$\sigma_j^{G,(\theta)} = \sigma_j^G \cdot \sqrt{\theta}$$

Lemma: When inventory and demand are scaled up, the transportation linear program objective value approaches that of the expected value of the perfect hindsight optimization in ratio. Or:

$$\lim_{\theta \to \infty} \frac{E[v_{PH}^{(\theta)}]}{v_{LP}^{(\theta)}} = 1$$

The proof is in the online materials. The outline of the proof is this: to show that the ratio of expected value of the perfect hindsight cost to the objective value of the linear program converges to 1, we sandwich it between two other ratios, each of which converges to 1 as demand and inventory scale up with θ . To obtain the lower bound, we need only to show that the objective value for the scaled linear program is less than the expected value of the cost of the perfect hindsight solution. We show this through Jensen's inequality. To obtain the upper bound on the ratio, we formulate an optimization problem whose cost is more than that of the expected value of the perfect hindsight solution. This is done by adjusting the perfect hindsight optimization problem in a series of steps. The ratio of the cost of this optimization problem to the objective value of the linear program converges to 1 as θ is scaled up. Thus, with the ratio of the expected value of the perfect hindsight cost to the objective value of the linear program bounded above and below by 1, we prove the above lemma.

8 Conclusion

We investigate online fulfillment and inventory data for a large American retailer, and show that a perfect hindsight fulfillment policy can outperform a myopic one by almost 3%, with respect to outbound shipping costs. A possible-to-implement heuristic, based on a transportation problem, captures about a third of this improvement gap by valuing inventory in geographically strategic locations as well as at fulfillment centers with large catalogues. The heuristic performance is robust to a variety of business conditions and leads to an additional benefit of keeping inventory more balanced throughout time.

This heuristic is being tested at our industrial partner on a subset of their SKU's with positive results. Thus far, the results have been favorable, confirming our findings in this paper. Possible full rollout is being considered.

There are several next steps that are worth considering. While the heuristic is most valuable when SKU volume is high, there are still opportunities for improvement when SKU volume is low. This last point suggests one direction for future research: determining implementable policies for low-volume SKU's, where the heuristic performs worst and none of the asymptotic properties apply. Lastly, we

believe there is potentially tremendous opportunity in linking fulfillment *and* replenishment policies, and determining a global optimal combined policy.

Acknowledgements

To maintain author anonymity, we will withhold acknowledgements until the final version.

References

Amazon.com, Inc. 2011. Form 10-k for 2010. http://www.sec.gov/edgar.shtml (last accessed 3 May 2012).

——. 2012. Form 10-k for 2011. http://www.sec.gov/edgar.shtml (last accessed 3 May 2012).

Archibald, T. W., D. Black, and K. D. Glazebrook. 2009. An index heuristic for transshipment decisions in multi-location inventory systems based on a pairwise decomposition. *European Journal of Operational Research* 192 (1):69–78.

Axsater, S. 2003. A New Decision Rule for Lateral Transshipments in Inventory Systems. *Management Science* 49 (9):1168–1179.

———. 1990. Modelling Emergency Lateral Transshipments in Inventory Systems. *Management Science* 36 (11):1329–1338.

Bertsekas, D., and J. Tsitsiklis. 1998. *Neuro-Dynamic Programming*. Belmont, MA: Athena Scientific.

Bluefly, Inc. 2011. Form 10-k for 2010. http://www.sec.gov/edgar.shtml (last accessed 24 September 2012).

_____. 2012. Form 10-k for 2011. http://www.sec.gov/edgar.shtml (last accessed 24 September 2012).

Cooper, W. L. 2002. Asymptotic Behavior of an Allocation Policy for Revenue Management. *Operations Research* 50 (4):720–727.

Forbes.com. 2012. Wal-Mart Stores Inc. Financial Ratios.

http://finapps.forbes.com/finapps/jsp/finance/compinfo/Ratios.jsp?tkr=wmt (last accessed 12 May 2012).

Forrester Research, Inc. 2012a. *Asia Pacific Online Retail Forecast, 2011 To 2016*. April 16, 2012.

——. 2012b. European Online Retail Forecast: 2011 To 2016. February 27, 2012.

——. 2012c. US Online Retail Forecast, 2011 To 2016. February 27, 2012.

Gallego, G. 1992. A minmax distribution free procedure for the (Q, R) inventory model. *Operations Research Letters* 11 (1):55–60.

Jannarone, J. 2011. Free Shipping Puts Retailers in a Box. *Wall Street Journal* 7 December. http://online.wsj.com/article/SB10001424052970204083204577082731595543136.html (last accessed 24 September 2012).

Kaplan, A. 1969. Stock Rationing. Management Science 15 (5):260–267.

Lee, H. L. 1987. A Multi-Echelon Inventory Model for Repairable Items with Emergency Lateral Transshipments. *Management Science* 33 (10):1302–1316.

Maxwell, M. S., M. Restrepo, S. G. Henderson, and H. Topaloglu. 2010. Approximate Dynamic Programming for Ambulance Redeployment. *Informs Journal on Computing* 22 (2):266–281.

Nahmias, S., and W. S. Demmy. 1981. Operating Characteristics of an Inventory System with Rationing. *Management Science* 27 (11):1236–1245.

Powell, W. B. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley-Interscience.

Van Roy, B., D. P. Bertsekas, Y. Lee, and J. N. Tsitsiklis. 1997. A neuro-dynamic programming approach to retailer inventory management. In *Proceedings of the 36th IEEE Conference on Decision and Control*, 4052–4057 vol.4.

Seeking Alpha. 2012. Online Retail Market Booming In China. http://seekingalpha.com/article/453511-online-retail-market-booming-in-china (last accessed 14 September 2012).

Simao, H. P., J. Day, A. P. George, T. Gifford, J. Nienow, and W. B. Powell. 2009. An Approximate Dynamic Programming Algorithm for Large-Scale Fleet Management: A Case Application. *Transportation Science* 43 (2):178–197.

Simpson, R. 1989. *Using Network Flow Techniques to Find Shadow Prices for Market Demands and Seat Inventory Control*. Cambridge, MA: MIT Flight Transportation Library.

Talluri, K., and G. van Ryzin. 1998. An Analysis of Bid-Price Controls for Network Revenue Management. *Management Science* 44 (11-Part-1):1577–1593.

Veinott, A. F. 1965. Optimal Policy in a Dynamic, Single Product, Nonstationary Inventory Model with Several Demand Classes. *Operations Research* 13 (5):761–778.

Vitacost.com, Inc. 2011. Form 10-k for 2010. http://www.sec.gov/edgar.shtml (last accessed 24 September 2012).

______. 2012. Form 10-k for 2011. http://www.sec.gov/edgar.shtml (last accessed 24 September 2012).

Williamson, E. L. 1992. *Airline network seat inventory control: methodologies and revenue impacts*. Cambridge, MA: Massachusetts Institute of Technology, Dept. of Aeronautics & Astronautics, Flight Transportation Laboratory.

Yang, J., and Z. Qin. 2007. Capacitated Production Control with Virtual Lateral Transshipments. *Operations Research* 55 (6):1104–1119.