

Supporting hurricane inventory management decisions with consumer demand estimates



Douglas J. Morrice^{a,*}, Paul Cronin^a, Fehmi Tanrisever^b, John C. Butler^a

^a McCombs School of Business, The University of Texas at Austin, Austin, TX 78712, USA

^b Faculty of Business Administration, Bilkent University, Bilkent, Ankara 06800, Turkey

ARTICLE INFO

Article history:

Accepted 4 May 2016

Available online 9 July 2016

Accepted by: Mikko Ketokivi

Keywords:

Humanitarian operations

Hurricanes

Forecasting

Inventory management

Disruption management

ABSTRACT

Matching supply and demand can be very challenging for anyone attempting to provide goods or services during the threat of a natural disaster. In this paper, we consider inventory allocation issues faced by a retailer during a hurricane event and provide insights that can be applied to humanitarian operations during slow-onset events. We start with an empirical analysis using regression that triangulates three sources of information: a large point-of-sales data set from a Texas Gulf Coast retailer, the retailer's operational and logistical constraints, and hurricane forecast data from the National Hurricane Center (NHC). We establish a strong association between the timing of the hurricane weather forecast, the forecasted landfall position of the storm, and hurricane sales. Storm intensity is found to have a weaker association on overall inventory decisions. Using the results of the empirical analysis and the NHC forecast data, we construct a state-space model of demand during the threat of a hurricane and develop an inventory management model to satisfy consumer demand prior to a hurricane making landfall. Based on the structure of the problem, we model this situation as a two-stage, two-location inventory allocation model from a centralized distribution center that balances transportation, shortage and holding costs. The model is used to explore the role of recourse, i.e., deferring part of the inventory allocation until observing the state of the hurricane as it moves towards landfall. Our approach provides valuable insights into the circumstances under which recourse may or may not be worthwhile in any setting where an anticipated extreme event drives consumer demand.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Pre-positioning of critical supplies for an anticipated disaster like a hurricane, typhoon, or some other predictable event is an important humanitarian problem (Rawls and Turnquist, 2010; Salmeron and Apte, 2010). In this paper, we take advantage of a detailed dataset from a Texas Gulf Coast retailer that allows us to investigate how consumers respond to hurricane forecasts and how the retailer should manage hurricane inventories to meet consumer needs. Retailers of hurricane supplies play an integral role in hurricane disaster management. Hurricanes as devastating and costly as Sandy and Katrina focused a lot of public attention on retailers' planning and preparation as well as their responsiveness

to these major disasters (see, for example, Assimon, 2009; Horwitz, 2009; Banjo, 2012; Target, 2016). Consequently, the focus of this paper is on key problems faced by a retailer providing supplies to customers preparing for such disasters: estimating hurricane demand and using this information to improve inventory management before a hurricane strikes.

As a hurricane advances across the Atlantic Ocean and into the Gulf of Mexico, consumers on the Texas Gulf Coast are faced with decisions regarding when to purchase supplies before the hurricane makes landfall. The retailer also faces difficult choices in terms of inventory management: if it reacts too quickly, it risks sending inventory to the wrong location, but if it waits to learn more about the location of the storm's landfall there is a risk that inventory shipments will not arrive in time to serve pre-landfall demand increases. Both of these outcomes result in unmet demand and lost sales. Unmet demand is particularly critical because it has the potential to exacerbate a humanitarian crisis. Even if no crisis materializes, shortages can still damage the retailer's reputation and lead to a loss of customer goodwill.

* Corresponding author.

E-mail addresses: Douglas.Morrice@mcombs.utexas.edu (D.J. Morrice), Paul.Cronin@utexas.edu (P. Cronin), Tanrisever@bilkent.edu.tr (F. Tanrisever), John.Butler2@mcombs.utexas.edu (J.C. Butler).

There is a great deal of uncertainty in a hurricane's behavior, but, as we demonstrate in this paper, consumers do react to hurricane forecasts and purchase some products in a predictable fashion. In turn, retailers can use these patterns to establish inventory management policies to better serve consumers and meet their own goals. Using a large point-of-sale data set, operational and logistical constraints faced by the retailer, and hurricane forecast information from the National Hurricane Center (NHC), we build a state-space model for hurricane demand in the spirit of [Song and Zipkin \(1993, 1996\)](#). We calculate the state-space transition probabilities and then use regression to estimate demand in each state. To our knowledge, this is the first paper to estimate empirically hurricane demand and then determine the best inventory management policy.

Around this demand model, we develop an inventory management model for satisfying consumer demand prior to a hurricane's landfall. The retailer serves two separate regions out of a central distribution center (DC). Each region consists of many demand locations (i.e., retail stores), but since stores are clustered in each region around a population center some distance from the DC, we consider each region a point location relative to the DC. In addition, the DC releases inventory over at most two stages without the possibility of transshipment due to practical logistical and safety constraints. To match the retailer's logistical constraints, we use 5- and 3-day forecasts to define the two stages of the inventory management system. Thus, we model this situation as a two-stage, two-location inventory allocation model from a centralized DC. It is important to note that our modeling approach is not limited to a two-stage, two-location model. In other words, more stages and locations can be incorporated into our model if the requirements of the problem warrant it. Of course, more stages and more locations increase the number of states in the state-space model making it more complex to solve and requiring more data for accurate demand estimates.

The inventory model is designed to determine when and how much to allocate to each location prior to a hurricane strike by trading off the benefits of waiting to gain more precise demand information against the risk of responding too late to the hurricane demand and running short of pre-strike inventories at the stores. Additionally, the inventory model is used to evaluate the benefits of a full recourse policy over a no recourse policy. Recourse refers to the retailer's ability to make an initial allocation to each region and then make a second allocation decision after gathering more information about the forecasted landfall of the storm, including the possibility of learning that the storm is no longer a threat. This comparison is important to the retailer because additional costs (e.g., extra transportation) and extra safety issues (e.g., sending out personnel and equipment) are associated with the full recourse policy where decisions are made closer to a storm's landfall.

We contribute to two streams in the disaster inventory management literature. First, we know of no other study that explicitly estimates the linkage between storm forecast data and pre-strike consumer demand as opposed to making stylized or simplifying assumptions about this relationship. This empirical analysis reveals that the consumers' perception of and reaction to hurricane risk should be key drivers of the retailer's operating decisions. In particular, the demand risk of the retailer is not always proportional to the uncertainty of the hurricane's eventual landfall location. In addition, this analysis reveals two additional things. First, weather forecasts available to consumers through the news media contain a significant amount of predictive information on hurricane demand. Second, while we are able to establish a strong association between the timing of the hurricane weather forecast, the forecasted landfall position of the storm and hurricane sales, we observed that the hurricane storm intensity was not as influential.

As a second contribution, we construct an inventory management model based on our empirical analysis to explore subtle and complex relationships among the determinants of the value of a recourse action. Our inventory model informs managers about when to follow a recourse policy and partially defer inventory allocations to gain more precise information about the possible path of the hurricane. Since simplifying assumptions of the previous literature do not directly apply to our model and the optimal inventory policy is not myopic, we solve a two-stage stochastic program with recourse and show how the following parameters impact the value of recourse: (1) initial DC inventories, (2) shortage costs relative to holding costs, (3) transportation costs relative to holding costs (4) forecast of the hurricane's landfall location, and (5) hurricane path uncertainty. We find high initial DC inventories motivate the managers to defer inventory allocations though the effect exhibits decreasing marginal returns. Uncertainty in the hurricane's path has a similar effect and motivates the managers to wait more before making large inventory allocations. During the threat of a hurricane, shortage costs tend to dominate as a reflection of the retailer's efforts to help mitigate a humanitarian crisis and maintain customer goodwill. As a result, we analyze the impact of increased relative shortage cost on the value of recourse. Counter to intuition, relative shortage cost does not have a monotone effect on the value of recourse. Its impact depends on the level of DC inventories as well as the initial forecast of the hurricane's landfall location. Not surprisingly, increasing transportation costs dampen the value of recourse. Lastly, we find the initial forecast of hurricane landfall location is a key driver of the value of recourse which ranges from 5% to more than 40%.

Our work allows the retailer to be more effective at hurricane inventory management and humanitarian relief because it provides a more rigorous evaluation of recourse. Historically, the retailer had used a single allocation, no recourse approach, and had moved to experimenting with ad hoc two-stage allocation policies. We were able to provide improved policies and demonstrate the value of recourse under differing circumstances. Ultimately, our work, in conjunction with other initiatives, enabled the retailer to participate in a statewide initiative to improve hurricane disaster preparedness in the state of Texas ([Texas Emergency Management Digest, 2009](#)), particularly for residents of the Texas Gulf Coast. Such public-private collaborations have become common during natural disasters as chronicled in [Raths \(2010\)](#). This article describes partnership arrangements between retailers and state officials in California and Texas. In Texas, retailers work directly with state officials and the state emergency operations center during hurricanes and other natural disasters to provide humanitarian relief.

2. Literature review

Our work is related to [Regnier and Harr \(2006\)](#) who develop a Markov model for hurricane path analysis and show how decision makers can benefit from the option to wait for updated hurricane forecasts when managing evacuations. [Regnier \(2008\)](#) applies a stochastic model of hurricane paths in a dynamic decision-making framework for evacuation policies. Our paper differs from these studies in several respects. First, we take the perspective of a retailer seeking to allocate hurricane inventory to its retail stores to meet pre-strike demand. Second, in addition to the hurricane's path, customer response to the evolution of the hurricane's location and intensity is stochastic in our model. We present an empirical analysis of customer demand to identify the linkage between the hurricane's intensity and forecasted path, and retail demand which is used as an input to our inventory model. Finally, unlike an evacuation decision which is difficult to adjust, inventory

allocations can be revised by a retailer after obtaining a more precise hurricane path forecast. Accordingly, our decision model is a two-stage stochastic program with recourse.

Lodree and Taskin (2009) and Taskin and Lodree (2011) use a Bayesian decision model with hurricane wind speed updates to manage inventories in a serial supply chain. The latter paper generalizes the former by considering a multiple, rather than a single, retailer supply chain and uses the NHC's official tropical cyclone wind speed probabilities. Both papers assume the existence of a function that links demand to the storm intensity and include only one inventory decision with no recourse. We estimate the link between the evolution of the hurricane forecast and consumer demand from retail data. In addition, we explicitly account for the logistical and operational constraints of the motivating retailer that result in a two-stage inventory allocation model with recourse. Hence, our work is the first in the literature to integrate single-depot, multi-location inventory management with logistical constraints under hurricane threat and an empirical model of hurricane demand. Our paper also moves the focus of the inventory management problem from managing hurricane risk to inferring customers' *perceived* hurricane risk by examining the impact of observable factors on actual consumer demand. Understanding how customers respond provides opportunities to manage their perceptions and predict or mitigate the effects on demand.

Salmeron and Apte (2010), Rawls and Turnquist (2010), and Lodree et al. (2012) consider the problem of pre-positioning supplies for surges in hurricane-induced demand for products and assets. These papers develop two-stage stochastic programs, but the stages are pre-strike and post-strike hurricane. While this approach facilitates aggregate planning, especially for disaster relief agencies, one of the most challenging and important problems faced by a retailer is allocating limited inventory to fulfill demand as a storm approaches. Motivated by the logistical constraints of our retailer, we focus on multiple pre-strike decisions. Further, we link demand with factors associated with hurricane forecasts rather than rely on scenario generation that uses simple averages.

Our inventory model is closely related to the literature on single-warehouse multiple-retailer systems. In an early study, Eppen (1979) considers a multi-location newsboy problem where each location features normal demand, linear holding costs and linear shortage costs. Eppen and Schrage (1981) and Federgruen and Zipkin (1984) consider a single DC-multiple location problem with random demands at each location that incorporates fixed order lead time, backorder costs, holding costs, and DC ordering costs. They minimize the total system cost over a finite number of time periods. Federgruen and Zipkin show that the resulting large dimensional dynamic programming model can be systematically approximated by a single-location inventory model and generalize the results of Eppen and Schrage. The most informative work for our model, McGavin et al. (1997), characterizes inventory balancing policies for distributing one product from a single DC to multiple locations over time.

Given that our demand function is determined by the stochastic evolution of the hurricane over time, the simplifying assumptions of the previous literature do not apply directly to our model. In particular, since we have non-identical demands at each location with lost sales and no backordering, the optimal inventory policy is not myopic (McGavin et al., 1997) and requires solving a two-stage model.

3. A demand model

The inventory management system currently utilized by the retailer has two distinct modes: regular and hurricane operations. Under regular operations, stores replenish according to a pull

system using point-of-sale scanner data. During hurricane events, demand is highly volatile and conditions can change rapidly. Consequently, the retailer switches to a “war room” approach; operating in a centralized fashion pushing inventory out to stores. This is common amongst retailers of hurricane supplies (Banjo, 2012).

3.1. A state-space model for hurricane demand

During a hurricane event, the National Weather Service provides forecast updates every six hours. Since 2003, each forecast includes 5 and 3-day forecast cones. Fig. 1 depicts a forecast for Hurricane Emily on July 15, 2005 at 5 a.m. EDT from the Tropical Cyclone Advisory Archive (NHC Graphics Archive., 2016). It contains the 5 and 3-day forecast cones for the hurricane's projected path. The retailer's management monitors these forecasts to make inventory allocation decisions. Its customers have access to the same forecasts through the news media to inform purchasing decisions in the face of a storm threat.

The retailer is most concerned about stores located along the Texas Gulf Coast because these stores and communities are most likely to be impacted by hurricane demand. Rather than focusing on demand at individual stores, a few factors suggest that it is sufficient to focus on two main regions defined by clusters of three-digit zip codes as shown in Fig. 2. First, there are two main population clusters (markets) along the Texas Gulf Coast serviced by the retailer. Region 1 includes the south Texas market from McAllen to Corpus Christi, and Region 2 includes the southeast Texas market containing Houston, Beaumont, and Galveston. Second, the retailer considers each of these regions as a cluster and dispatches hurricane inventory from a single DC in south-central Texas along major roadways within its logistics network accordingly.

The retailer monitors all hurricanes that form in the Atlantic. The retailer makes inventory allocations at two main threat events that align with the weather forecast: i) the first time a 5-day forecast cone intersects at least one region and ii) the first time a 3-day forecast cone intersects at least one region. The retailer does not consider 1 or 2-day forecasts as inventory triggers because it does not believe it can deliver shipments safely and efficiently when a hurricane is expected to make landfall within 1–2 days.

To develop a state-space model for this decision-making process, we examined all fifteen Atlantic storms from the Tropical Cyclone Advisory Archive (NHC Graphics Archive., 2016) that threatened the Texas Gulf Coast from 2003 to 2008 (see Table 1). This time period was chosen because 5-day forecasts became available to the public for the first time in 2003 (NHC Forecast Verification., 2016) and because the retailer featured a consistent set of hurricane inventory offerings during this time frame. The “Forecast Date” is the date on which either a 5-day or 3-day forecast cone (specified in column 3) first intersects one or more of the regions. The “Forecasted Region of Impact” indicates whether any part of the forecast cone intersects Region 1, Region 2, or both regions. The “Current Intensity” measures the current intensity of the storm on the forecast date according to the Saffir-Simpson Hurricane Wind Scale (NHC Saffir-Simpson, 2016). The NHC first provided the “Forecasted Intensity” in 2004, which is the intensity of the storm forecasted over the retailer region(s) either 3 or 5 days in advance. As reported, it is less specific than the Saffir-Simpson scale and only has two categories in the data set: tropical storm/depression and hurricane.

Table 2 describes the variables that define the state-space model and will be used in subsequent regression models. The “Type of Forecast” (TOF) variable has three categories: i) when the first forecast cone to intersect at least one retailer region is the 5-day forecast cone (TOF = 1), ii) when the first forecast cone to



Fig. 1. 5 and 3-day weather forecast for Hurricane Emily (2005).

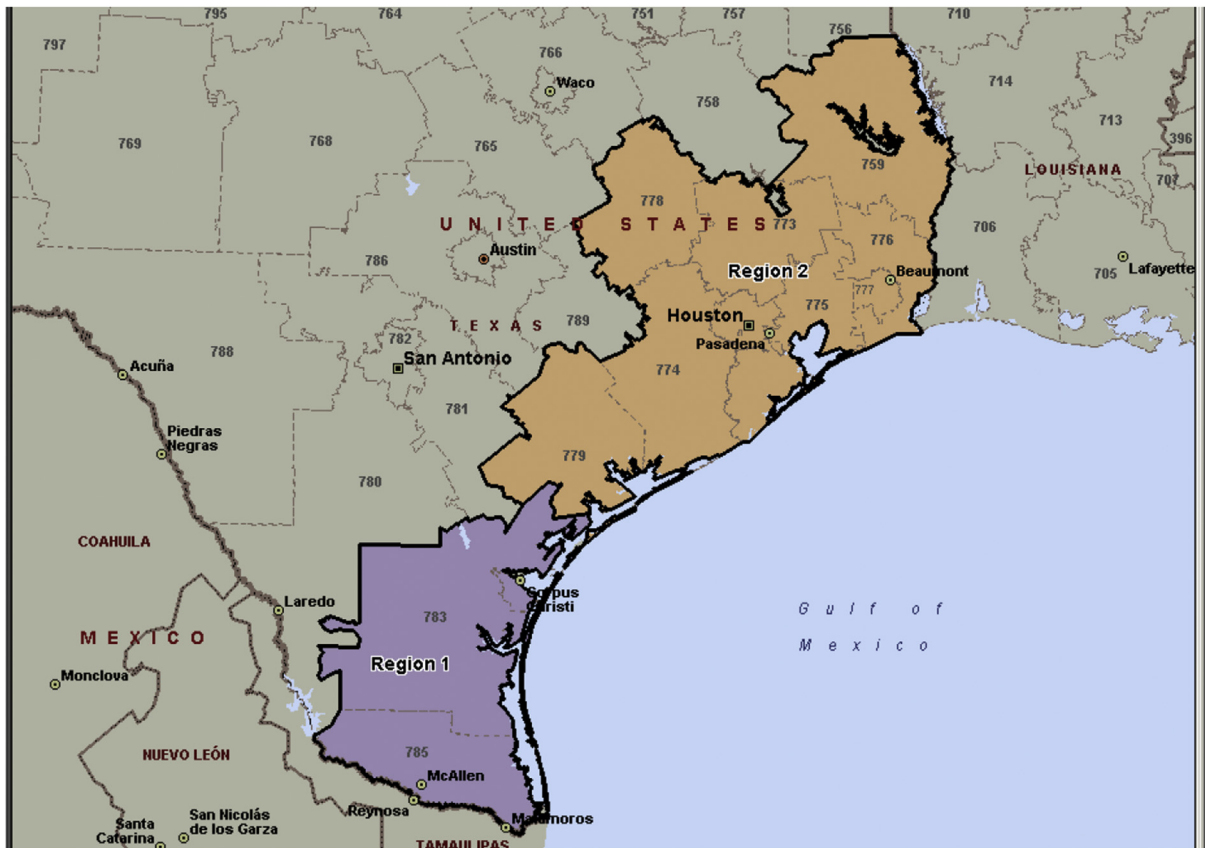


Fig. 2. Regions along the Texas Gulf Coast defined by three-digit zip code clusters.

intersect at least one retailer region is a 3-day forecast cone (TOF = 2), and iii) when a 3-day forecast cone that follows an intersecting 5-day forecast cone first intersects at least one of

region (TOF = 3). Inspection of Table 1 reveals why the Type of Forecast variable was defined in this way. For six storms, the 5-day forecast cone was the first forecast cone to intersect at least one of

Table 1
Threatening storms between 2003 and 2008 along with their characteristics.

Storm name	Forecast date	5-day/ 3-day	Forecasted Region of Impact	Current intensity ^a	Forecasted Intensity ^b
Bill	06/29/2003	3-Day	2	T.S.	N/A
Claudette	07/08/2003	5-Day	Both	T.S.	N/A
	07/10/2003	3-Day	Both	T.S.	N/A
Erika	08/14/2003	3-Day	1	T.S.	N/A
Grace	08/30/2003	3-Day	Both	T.D.	N/A
Ivan	09/22/2004	3-Day	Both	T.D.	T.S./T.D.
Cindy	07/03/2005	3-Day	Both	T.S.	T.S.
Emily	07/14/2005	5-Day	1	Category 2	Hurricane
	07/16/2005	3-Day	1	Category 4	Hurricane
Rita	09/18/2005	5-Day	Both	T.D.	Hurricane
	09/20/2005	3-Day	Both	Category 1	Hurricane
Dean	08/17/2007	5-Day	Both	Category 2	Hurricane
Erin	08/14/2007	3-Day	Both	T.D.	T.S./T.D.
Humberto	09/12/2007	3-Day	2	T.D.	T.S.
Dolly	07/20/2008	3-Day	1	T.S.	T.S.
Eduardo	08/03/2008	3-Day	Both	T.D.	T.S.
Gustav	08/27/2008	5-Day	2	T.S.	Hurricane
	08/30/2008	3-Day	2	Category 2	Hurricane
Ike	09/07/2008	5-Day	2	Category 4	Hurricane
	09/09/2008	3-Day	Both	Category 1	Hurricane

^a T.S.: Tropical Storm. T.D.: Tropical Depression. Both are weaker than a hurricane. Category 1 is the weakest hurricane and Category 5 is the strongest hurricane.

^b N/A: Forecast Intensity data was not available prior to 2004.

Table 2
Description of the state-space model and regression categorical variables.

Variable	Value	Description
Type of Forecast	TOF = 1	First forecast cone to intersect at least one retailer region is the 5-day forecast cone
	TOF = 2	First forecast cone to intersect at least one retailer region is a 3-day forecast cone
	TOF = 3	A 3-day forecast cone that follows a 5-day forecast cone first intersects at least one retailer region
Forecasted Region of Impact	FRI = 1	When the forecast cone intersects Region 1 only
	FRI = 2	When the forecast cone intersects Region 2 only
	FRI = 3	When the forecast cone intersects both Regions 1 and 2
Current Intensity	CI = 1	Storm is a tropical storm/tropical depression
	CI = 2	Storm is a Category 1 or 2 Hurricane
	CI = 3	Storm is a Category 3, 4, or 5 Hurricane
Forecasted Intensity	FI = 1	Storm is forecast to be a tropical storm/depression
	FI = 2	Storm is forecast to be a hurricane
Region	R = 1	Region under consideration is 1
	R = 2	Region under consideration is 2
Product	PRODUCT = 1	The product under consideration is product 1
	PRODUCT = 2	The product under consideration is product 2

the retailer regions. Five of these six storms eventually had a 3-day forecast cone intersect at least one of the retailer regions; the other storm (Dean in 2007) moved off in a different direction and ceased to threaten regions 1 or 2. For the other nine storms in Table 1, the first forecast cone to intersect at least one of the retailer regions was a 3-day cone.

The Forecasted Region of Impact (FRI), Current Intensity (CI), and Forecasted Intensity (FI) variables in Table 2 were introduced in Table 1. The “Region” variable is included to control for possible demographic differences between the two regions. It is important to note that the Forecasted Region of Impact variable can be interpreted as a measure of uncertainty about where the storm will make landfall. When FRI = 3 and the hurricane could impact either region, there is more uncertainty for consumers as to whether or not they will be impacted by the storm. We use three categories for Current Intensity, rather than the five associated with the Saffir-Simpson Hurricane Wind Scale, to be consistent with the retailer’s categorization of storms when making its inventory management decisions.

Lastly, the model contains a Product variable to accommodate two products with differing characteristics. Product 1 had lower volume and more variable regular sales but was generally assumed

to be a more important staple during a natural disaster. In contrast, Product 2 had higher volume and relatively predictable regular sales but was assumed to be less vital during a crisis situation. By analyzing multiple products simultaneously, we pool more observations into a single model and demonstrate that patterns found in the data due to a hurricane threat are similar across different products.

Note that some of the state variables in Table 2 are stage dependent, e.g. the FRI for the 5-day forecast may be different than the FRI for the 3-day forecast, as we discuss in Section 3.3. Further, the categorical variables are transformed into dummy variables in the regression model but are referred to as categorical variables in the paper for ease of exposition.

3.2. Estimating state dependent hurricane demand

To develop a hurricane demand model, we examine more than three hundred thousand records from the retailer’s daily point-of-sales data for Products 1 and 2 from 2003 to 2008. In this analysis, we use these sales data as a proxy for demand. While we recognize the limitations of this approach because sales is “right censored demand”, sales becomes a better proxy for demand during

hurricane stocking (which is our main focus), because the retailer tries to avoid stock outs to help mitigate a potential humanitarian crisis and maintain customer goodwill by pushing inventory out to the stores.

We regard sales as having two components: regular and hurricane sales. Regular sales include all sales during regular operations when there is no hurricane event. By contrast, hurricane sales include all sales during the threat of a hurricane. In the subsections that follow we develop a baseline demand model that explains regular demand for each product and then a hurricane sales model that explains the excess demand as a function of the state of the hurricane (type of forecast, current intensity, and forecasted region of impact).

3.2.1. A regression model of regular sales

To account for regular sales, we fit various regression models to the sales data for each product using SAS (SAS, 2016). We excluded all observations occurring during the threat of the hurricanes given in Table 1 to establish a better baseline demand for the regular sales model. The best multiple linear regression model for predicted regular sales that we found has the following form:

$$\text{Sales}_{s,t} = \alpha + \beta_{\text{price}_{s,t}} \text{Price}_{s,t} + \sum_{s=2}^n \beta_{\text{store}_s} \text{Store}_s + \sum_{m=2}^{12} \beta_{\text{month}_m} \text{Month}_m + \beta_{\text{trend}} \text{Trend} \quad (1)$$

where the notational definitions are provided in Table 3.

The model in (1) captures seasonality, trend, and store characteristics. It has adjusted- R^2 values of approximately 0.59 and 0.78 for Products 1 and 2, respectively, which is consistent with the retailer's belief that during normal operations, sales of Product 1 are more volatile than Product 2.

3.2.2. A regression model of hurricane sales

We use model (1) to predict the regular sales during the time periods of threatening hurricanes in Table 1. The regression residuals from this step are treated as hurricane sales for stores in Regions 1 and 2. For the 15 storms in Table 1, we generate 40 residual sales observations for each product. These are the residual sales that occur in each region for the 20 times when either the 5-day forecast or 3-day forecast intersect at least one of the regions. More specifically, for each of the six storms with a 5-day forecast, we sum the daily residuals from model (1) by region and product to get a residual sales observation. For example, the 5-day forecast residual sales observation for the storm Claudette is the sum of model (1) residuals by region and product on 7/8/2003 (when the 5-day forecast is issued) and 7/9/2003. We calculated the 3-day forecast residual sales observations in a similar manner. For example, the 3-day forecast residual sales observation for Claudette is the sum of model (1) residuals by region and product from 7/10/

Table 4

Residual regression hurricane sales model.

Parameter	Coefficient	P-value	t-statistic
Intercept	8.6145	<0.0001	55.85
Product ₁ *TOF ₂	−1.0625	<0.0001	−4.84
FRI ₃	1.3351	<0.0001	6.23
FRI ₃ *TOF ₂	−1.3633	<0.0001	−5.22
FRI ₁ *R ₁	1.8293	<0.0001	6.16
CI ₁ *TOF ₁	−1.5811	<0.0001	−6.25
CI ₂ *TOF ₁	−0.9194	0.0022	−3.17

Sample size = 80. Adj- R^2 = 0.6124. Shapiro-Wilk Test: p-value = 0.4968. White Test: p-value = 0.2125.

2003 (when the 3-day forecast is issued) to 7/15/2003 when the storm actually made landfall (not shown in Table 1). These datasets are aggregated for the two products into a joint model with 80 observations.

For the hurricane sales regression model, we use the six predictor variables defined in Table 2. The categorical variables were transformed to binary variables (e.g. FRI = 1 and FRI = 2 become FRI₁ and FRI₂, respectively) for regression purposes. Our regression analysis includes main and two-way interaction effects. Rather than attempt to fully enumerate all combinations of dummy variables, we started with combinations of variables based on some initial guidance from the retailer. The model selection techniques employed were stepwise regression (SAS, 2016), MAXR (SAS, 2016), and the BayesVarSel package (Garcia-Donato and Forte, 2014) in R that analyzes the variable selection problem from a Bayesian perspective. Our goal is to find models that yield the highest adjusted- R^2 values, but we limit the number of explanatory variables in the model to avoid over-fitting and multi-collinearity issues. In addition to the main effects listed in Table 2, we consider selected interaction effects, primarily associated with main effects variables that appear to have predictive power.

Eq. (2) contains the model that emerged from our analysis, and the coefficient estimates and regression results are presented in Table 4. Since the hurricane residual sales data are highly right-skewed for both products, we transform the data using a natural logarithmic transformation. For technical reasons and in order to protect sensitive retailer information, we modify the residual sales data for each product by applying an order-preserving transformation before applying the logarithmic transformation so that \hat{y} in (2) is the natural logarithm of transformed residual sales. The model has an adjusted R^2 of 0.6124. We tested the standard regression assumptions of normality, homoscedasticity, statistical independence, and linearity using White's test (SAS, 2016), the Shapiro-Wilk test (SAS, 2016), and visual inspection of the residuals versus fitted plot, respectively. We did not detect any significant departures from these assumptions. Visual inspection of the kernel density estimation (KDE) plot and quantile-quantile (QQ) plot also did not give reason to reject normality.

Table 3

Variable definitions for the regression model for regular sales.

Variable	Description
Sales _{s,t}	The predicted daily sales (in dollars) for store <i>s</i> on day <i>t</i> .
Price _{s,t}	Average daily price for the product for store <i>s</i> calculated on day <i>t</i> . Price information at point of sale was not available. Calculated as sales/units. For stores with no sales on a particular day, the median sale price for that store in that month was used to avoid division by zero.
Store _s	Dummy variables capturing the fixed effects for <i>n</i> stores. Store _s = 1 when the observation's store number is store <i>s</i> , otherwise the variable is set to zero.
Month _m	Dummy variables capturing monthly seasonality effects for 12 months. Month _m = 1 if the observation occurs during month <i>m</i> , otherwise the variable is set to zero.
Trend	Monthly trend variable to capture changes in product demand over time. Initially set to 1 for the first month of the analysis (January 2003) and incremented by 1 for each of the 72 months of the analysis. This variable is designed to capture any secular trends in sales.

$$\hat{y} = \alpha + \beta_1 \text{Product}_1 * \text{TOF}_2 + \beta_2 \text{FRI}_3 + \beta_3 \text{FRI}_3 * \text{TOF}_2 + \beta_4 \text{FRI}_1 * \text{R}_1 + \beta_5 \text{CI}_1 * \text{TOF}_1 + \beta_6 \text{CI}_2 * \text{TOF}_1 \quad (2)$$

We identify several key insights from the model parameters and coefficients. The strongest lift in sales is due to forecasted landfall position of the storm. In particular, significant sales lift results when the storm is forecasted to impact Region 1 alone or when the storm is forecasted to impact both regions. The latter is mitigated only when the storm originates 3 days from landfall since there is less time for consumers to react.

It is somewhat surprising that storm intensity does not appear to be a driving force in this model. While the last two coefficients in Table 4 indicate that there are downward adjustments for weaker intensity storms, these are only for 5-day forecasts ($\text{TOF} = 1$). Perhaps the best explanation of this result comes from managers in charge of hurricane operations for the retailer. From their experience, customers are not so concerned about the size and strength of the storm when they make a purchase to prepare for a storm. In other words, customers simply “stock up” if they feel threatened by any sizable storm, since even a tropical storm can result in significant flooding or wind damage. Further, since storms tend to intensify quickly when they either reach or form in the Gulf (Regnier, 2008), customers may recognize that the intensity of the storm can change rapidly and unpredictably and place less emphasis on current storm intensity when making their purchases. Forecasted Intensity did not turn out to be statistically significant in our model, possibly to due to the limited granularity (tropical storm or hurricane).

The model in Table 4 indicates that both products exhibit similar sales volumes during 5-day (TOF_1) and ensuing 3-day forecasts (TOF_3), reinforcing the retailer's assumption that Product 1, the product with lower volume regular sales, is purchased in relatively higher volumes during these types of storms. Recall that we tested a large number of interaction terms and no combination of TOF_1 , TOF_3 and Product was identified as statistically significant. However, this does not hold for new 3-day forecasts (TOF_2) where consumer purchasing more closely mirrors non-storm buying leading to a negative coefficient on the Product 1 term for new 3-day forecasts ($\text{Product}_1 * \text{TOF}_2 = -1.0625$). Again, this may be due to the fact that consumers have less time to react when the first time a region is threatened is a three day forecast.

3.3. State transition probabilities for demand model

In this section, we provide estimates of the state transition probabilities which are key inputs to the two-stage inventory allocation model. Since we are concerned with examining the value of the recourse option for the retailer, we focus our discussion on the transition from a 5-day storm forecast to a subsequent 3-day forecast. These two forecasts represent the two stages in our inventory allocation model (i.e. from $\text{TOF} = 1$ to $\text{TOF} = 3$). The variables included in the regression model in Table 4 i.e., the Type of Forecast (TOF), Forecasted Region of Impact (FRI), and Current Intensity (CI), define the state space. Further, we need to estimate the probability of the different states and the probability of transitioning between these states.

Table 1 contains the data that allows us to estimate the state transition probabilities shown in Fig. 3. We incorporate a superscript for each variable to make it clear whether the state is associated with the first or second stage. From these data, there is a 40 percent chance that the first forecast of a threatening storm is a 5-day forecast, i.e., $\hat{P}(\text{TOF}^{(1)} = 1) = 2/5$ and a 60 percent chance that a 3-day forecast is the initial threat, i.e., $\hat{P}(\text{TOF}^{(1)} = 2) = 3/5$. These

probabilities are shown in the first level of branches in the probability tree in Fig. 3. The second level of branches in Fig. 3 contains the conditional probabilities $\hat{P}(\text{FRI}^{(1)} = i | \text{TOF}^{(1)} = j)$ for $i = 1, 2, 3$ and $j = 1, 2$, derived from columns 3 and 4 of Table 1. The third level of branches in the probability tree in Fig. 3 provides the conditional probabilities $\hat{P}(\text{CI}^{(1)} = i | \text{TOF}^{(1)} = j)$ for $i = 1, 2, 3$ and $j = 1, 2$, derived from Table 1 columns 3 and 5. From column 3 of Table 1, we estimate a 5/6 chance that a 3-day forecast threat follows a 5-day forecast threat, i.e., $\hat{P}(\text{TOF}^{(2)} = 3 | \text{TOF}^{(1)} = 1) = 5/6$; otherwise the storm no longer threatens either region with probability 1/6. These probabilities are shown in Fig. 3 on the fourth level of the tree.

Only five storms in Table 1 make the state transition from 5- to 3-day Type of Forecast. While we do not have observations for every possible Forecasted Region of Impact state transition, we can make some general observations from these data that allow us to get estimates of the conditional probabilities as shown in left panel of Table 5. Specifically, in four out of the five storms, the Forecasted Region of Impact stayed the same, and in one case the Forecasted Region of Impact expanded from one region to both regions (i.e., moving from $\text{FRI}^{(1)} = 1$ or 2 to $\text{FRI}^{(2)} = 3$). Additionally, the Forecasted Region of Impact never changed from a threat to both regions to a threat to a single region (i.e., moving from $\text{FRI}^{(1)} = 3$ to $\text{FRI}^{(2)} = 1$ or 2) nor did it switch from one region to the other (i.e., moving from $\text{FRI}^{(1)} = 2$ to $\text{FRI}^{(2)} = 1$ or from $\text{FRI}^{(1)} = 1$ to $\text{FRI}^{(2)} = 2$). The fifth level branches on the probability tree in Fig. 3 contain conditional probabilities derived from these general observations.

The Current Intensity state transitions also exhibit particular behaviors as shown in the right panel of Table 5. Recall that CI has 3 levels (1 = tropical storm/depression, 2 = weak (category 1 or 2) hurricane, 3 = strong (category 3–5) hurricane). Hurricanes were observed to have increased or decreased in intensity from the 5-day to 3-day forecast but the transitions were never extreme enough to jump from $\text{CI}^{(1)} = 1$ to $\text{CI}^{(2)} = 3$ or to weaken so substantially that $\text{CI}^{(1)} = 3$ transitioned to $\text{CI}^{(2)} = 1$. Of the five storms that had 5-day and 3-day forecasts, one had a constant intensity, three had increased intensity, and one had a lower intensity by the 3-day forecast. So a storm that starts with intensity 1 will either remain the same ($\hat{P}(\text{CI}^{(2)} = 1 | \text{CI}^{(1)} = 1) = \frac{1/5}{4/5} = 1/4$) or increase to intensity 2. A storm with intensity 2 can decrease, stay the same, or increase. A storm with intensity 3 can remain the same or decrease to 2 only. We use this information to model the transition probabilities for the intensity shown on the fifth level of Fig. 3.

4. Dynamic inventory allocation model

In this section, we utilize the hurricane demand model developed in Section 3 in a decision-making framework for inventory management under the threat of a hurricane that is triggered when a 5-day forecast cone intersects at least one of the two regions. We present a two-stage inventory allocation model where the stages are distinguished by 5- and 3-day forecasts of the hurricane. The retailer starts with a fixed amount of hurricane inventories at the DC and decides how much inventory to allocate to each demand region at each stage of the model as depicted in Fig. 3. The retail demand depends on the evolution of the hurricane over the planning horizon as described in Section 3.

We consider the second stage allocations, i.e., after observing both the 5 and 3-day forecasts, as a recourse action. The option to make a second stage allocation is of significant managerial concern since such an action requires substantial logistics costs and operational challenges. The value of this recourse action informs the managers about “when to pull the trigger”, i.e., when to make a single allocation at stage one, versus making an initial stage one

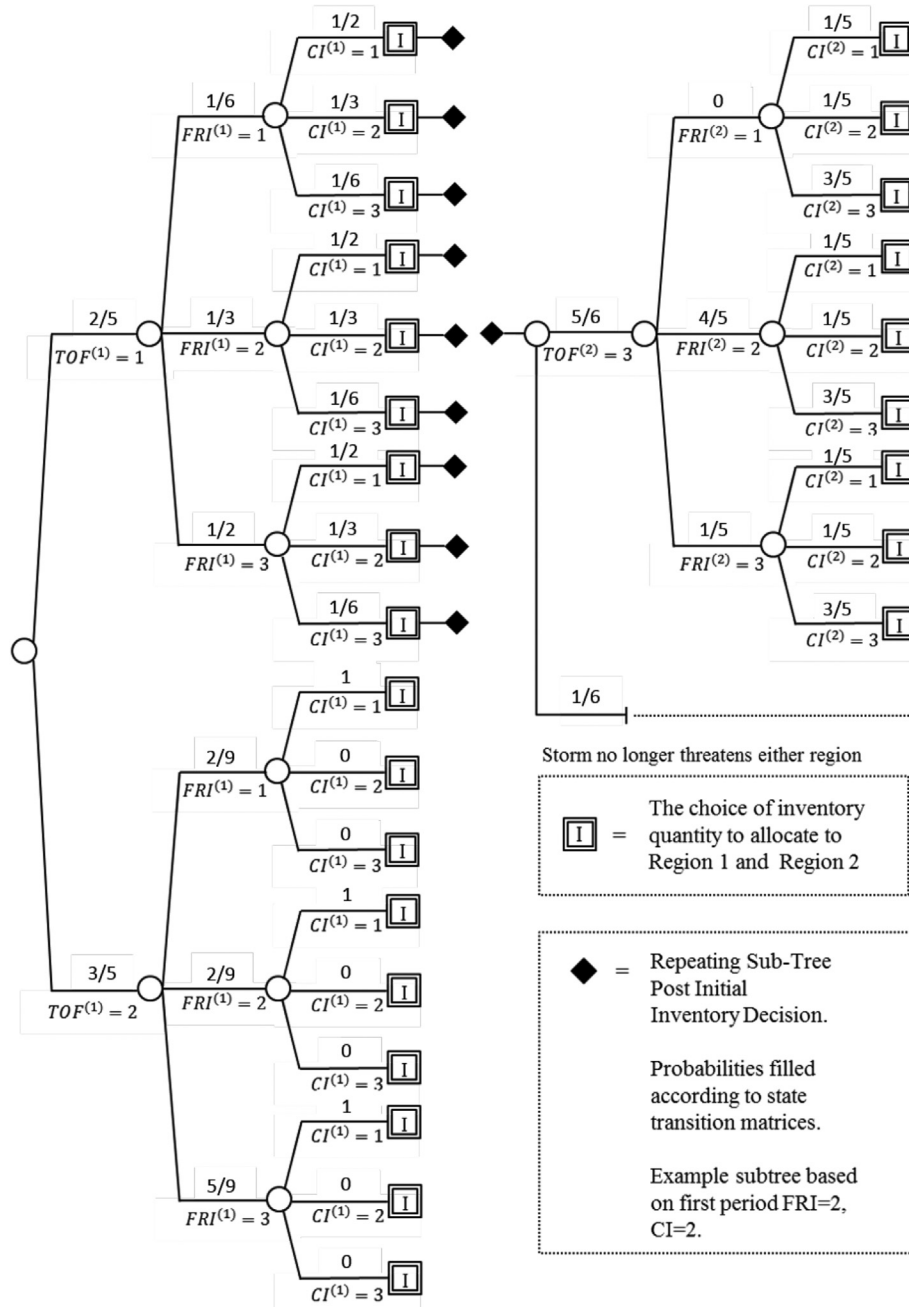


Fig. 3. Decision tree for state space demand model.

allocation and then making a second allocation after observing the 3-day forecast.

Evolution of the hurricane is a key modeling parameter in our model as the demand critically depends on this factor. Forecasted

Table 5

Stage 1 to stage 2 transition probability matrices for forecasted region of impact (left) and current intensity (right).

Forecasted Region of Impact				Current Intensity			
Stage 1 FRI	Stage 2 FRI			Stage 1 CI	Stage 2 CI		
	1	2	3		1	2	3
1	4/5	0	1/5	1	1/4	3/4	0
2	0	4/5	1/5	2	1/5	1/5	3/5
3	0	0	1	3	0	1/2	1/2

Region of Impact (FRI), Current Intensity (CI), and Type of Forecast (TOF) define the state of the storm. The recourse model starts with $TOF = 1$ (5-day forecast) and then either transitions to $TOF = 3$ (3-day forecast following a 5-day forecast) or is no longer a threat. These two stages are indexed by $t = \{1, 2\}$. We define the state of the storm as χ_t where $t = \{1, 2\}$ represents the stages just described. The combination of Forecasted Region of Impact and Current Intensity characterize nine potential storm scenarios faced by the retailer (i.e. $FRI = \{1, 2, 3\} \times CI = \{1, 2, 3\}$). The storm state is a function of these nine scenarios, but for ease of notation we will use χ_t unless we need to explicitly identify a particular scenario. We also define some additional notation in Table 6.

At $t = 1$: The retailer first observes the current state of the hurricane χ_1 which determines the demand distributions $\bar{d}_{1,1}(\chi_1)$

Table 6
Inventory model notation for time $t = 1, 2$, and region $r = 1, 2$.

Notation	Definition
$\tilde{d}_{t,r}(\chi_t)$	Demand in region r , in stage t , given the current state of the hurricane χ_t ; $\tilde{d}_t = [\tilde{d}_{t,1}, \tilde{d}_{t,2}]$
$F_{d_{t,r}}(\cdot)$	Cumulative distribution function of demand $d_{t,r}$.
$x_{t,r}$	Initial inventory position in region r in stage t ; $x_t = [x_{t,1}, x_{t,2}]$
$y_{t,r}$	Echelon inventory position after inventory allocation in region r in stage t ; $y_t = [y_{t,1}, y_{t,2}]$
c_r	Transportation cost per unit allocated in region r
I_t	Inventory position at the DC in stage t
s	Cost of shortage per unit demand
h	Holding cost per unit inventory

and $\tilde{d}_{1,2}(\chi_1)$ in regions 1 and 2, respectively, for the first stage. At the same time, the retailer observes the current level of inventories at the DC (I_1) and level of inventories at each region ($x_{1,1}$ and $x_{1,2}$), and decides how much inventory to allocate to each region. This results in echelon inventory levels of $y_{1,1}$ and $y_{1,2}$ in regions 1 and 2, respectively. At the end of the first stage, demand, stock out amounts, and excess inventories are observed. Excess demand $[y_{1,r} - d_{1,r}(\chi_1)]^-$ for $r = 1, 2$ is lost and the excess inventories, $x_{2,r} = [y_{1,r} - d_{1,r}(\chi_1)]^+$ for $r = 1, 2$ are carried over to the next stage. The retailer has allocated $y_{1,r} - x_{1,r}$ to each region and incurs a total transportation cost of $\sum_{r=1}^2 c_r (y_{1,r} - x_{1,r})$ in the first period. Given the time frame of at most a few days, inventory holding costs are negligible in the first stage.

At $t = 2$: The retailer observes the new state of the hurricane χ_2 which determines the demand distribution in each location, $\tilde{d}_{2,1}(\chi_2)$ and $\tilde{d}_{2,2}(\chi_2)$ for the second stage. At the same time, the retailer observes the current level of inventories at the DC (I_2) and the regional inventories ($x_{2,r}$ for $r = 1, 2$), and determines how much inventory to allocate to each region. This results in echelon inventory levels of $y_{2,1}$ and $y_{2,2}$ in regions 1 and 2, respectively. At the end of the second stage, demand and the stock out amounts are observed and holding and penalty costs are paid as usual.

The two-stage model can be expressed mathematically as

$$V_1(I_1, x_1, \chi_1) = \min_{y_1} \sum_{r=1}^2 c_r (y_{1,r} - x_{1,r}) + E_{\tilde{d}_1} \left[\sum_{r=1}^2 L_{1,r}(y_{1,r}) \right] + E_{\tilde{d}_1} E_{\tilde{\chi}_2 | \chi_1} [V_2(I_2, x_2, \tilde{\chi}_2)] \quad (3)$$

$$s.t. \sum_{r=1}^2 (y_{1,r} - x_{1,r}) \leq I_1 \quad (3.a)$$

$$y_{1,r} \geq x_{1,r} \quad \forall r = 1, 2 \quad (3.b)$$

$$I_2 = I_1 - \sum_{r=1}^2 (y_{1,r} - x_{1,r}) \quad (3.c)$$

where the above functions are defined as,

$$x_{2,r} = [y_{1,r} - d_{1,r}(\chi_1)]^+ \quad \forall r = 1, 2 \quad (3.d)$$

$$L_{1,r}(y_{1,r}) = s[y_{1,r} - \tilde{d}_{1,r}(\chi_1)]^- \quad \forall r = 1, 2 \quad (3.e)$$

$$V_2(I_2, x_2, \chi_2) = \min_{y_2} \sum_{r=1}^2 c_r (y_{2,r} - x_{2,r}) + E_{\tilde{d}_2} \left[\sum_{r=1}^2 L_{2,r}(y_{2,r}) \right] \quad (3.f)$$

$$s.t. \sum_{r=1}^2 (y_{2,r} - x_{2,r}) \leq I_2 \quad (3.g)$$

$$y_{2,r} \geq x_{2,r} \quad \forall r = 1, 2 \quad (3.h)$$

and,

$$L_{2,r}(y_{2,r}) = s[y_{2,r} - \tilde{d}_{2,r}(\chi_2)]^- + h[y_{2,r} - \tilde{d}_{2,r}(\chi_2)]^+ \quad \forall r = 1, 2 \quad (3.i)$$

In the first stage, the retailer makes the allocation decision to minimize the total cost of transportation, expected lost sales in the first stage and the expected cost of managing inventory in the second stage (i.e., the cost-to-go function), as shown in Eq. (3). Constraint (3.a) limits the total allocation (the echelon level $y_{1,r}$ less the initial inventory position $x_{1,r}$ over both regions) to the available inventory at the DC (I_1) and Eq. (3.b) ensures a non-negative allocation level. Constraint (3.c) balances the inventory at the DC. The starting inventory level for period 2 in each region $x_{2,r}$ is set via Eq. (3.d) to the excess inventory level remaining after demand is realized in period 1, or to zero if demand exceeded the echelon inventory level in the region. The function $L(\cdot)$ depicted in Eq. (3.e) and (3.i) describes the loss function in stages 1 and 2, respectively, due to excess inventories and shortages.

In the second stage, the retailer observes the revised state of the hurricane and inventory levels, and then allocates inventories to minimize the inventory costs including the cost of shipment, shortage and excess (Eq. 3.f). Eq. (3.g) ensures that the total allocation amount is limited by the available inventory at the DC and (3.h) guarantees that each allocation amount is non-negative.

Since the demand is state dependent, a myopic policy is not optimal in our case and a closed-form optimal solution is intractable. We utilized the standard error of the residual regression model (Table 4) and the assumption of normal errors to determine the demand distribution in each state. These continuous distributions were discretized and then featured in the linear programming equivalent of our formulation (Birge and Louveaux, 1997, page 156), which was solved numerically with GAMS (GAMS, 2016).

For comparison, a model without the recourse option is considered to explore the value of recourse to the retailer. In the non-recourse model, the retailer cannot change the echelon levels in the second period; all allocations must be made upfront after observing only the 5-day storm forecast. We solve both models and explore the value of the recourse (VOR) option in the numerical analysis section that follows.

5. Numerical analysis

The optimal value with recourse, V_R , gives the minimum inventory management cost with a recourse option which is compared to the optimal cost without recourse, $V_{NR} \geq V_R$. We define the percentage value of recourse as $\Delta VOR = 100 \cdot (V_{NR} - V_R) / V_{NR}$. Based on feedback from the retailer, we consider the following cost parameter settings: $h = 1$, $s = \{1, 5, 10, 20\}$ and

$c_r = \{0.05, 0.10, \dots, 0.50\}$. The DC initial inventory level is also evaluated for each product: $I_1 = \{20,000, 40,000, \dots, 300,000\}$ for Product 1 and $I_1 = \{20,000, 40,000, \dots, 600,000\}$ for Product 2. The sensitivity analysis on these parameter is conducted for all nine storms scenarios: $FRI = \{1, 2, 3\} \times CI = \{1, 2, 3\}$.

Fig. 4 provides the results of our numerical analysis for Product 1. The results for Product 2 are very similar since both products exhibit analogous sales patterns during 5-day and ensuing 3-day forecasts as we discussed in our analysis of the model in Table 4. In Fig. 4, we assume a nominal transportation cost of $c_r = 0.1$ and observe how the Forecasted Region of Impact (FRI), Current Intensity (CI), DC initial inventory level (I_1), relative shortage penalty (s), and transition probabilities all play important roles in determining the VOR. We will discuss these impacts and the various effects these parameters have on the value of the recourse action in the remainder of this section. In Section 5.4, we conduct a sensitivity analysis on the transportation costs.

5.1. Initial inventory level at the DC

At low levels of initial inventory at the DC, the VOR increases monotonically with the DC initial inventory level but reaches a limit beyond which it does not increase. The inventory level at which the VOR stabilizes is roughly consistent across all storm intensity levels for a given Forecasted Region of Impact. This happens because more initial inventory at the DC provides the retailer with more flexibility in exercising its recourse option after observing the state of the hurricane in the second stage. However, after I_1 reaches a certain point, the curves in Fig. 4 flatten out since I_1 is no longer a binding constraint on the recourse action.

5.2. Shortage costs and Forecasted Region of Impact

The effect of shortage costs on the VOR is non-monotonic and is dependent on DC Initial Inventory Level, Forecasted Region of Impact, the state transition probabilities, and to a lesser extent Current Intensity. When the DC Initial Inventory Level is low, there is less opportunity for a recourse action and as shortage costs increase there is more incentive for the retailer to allocate its limited

amount of inventories in the first stage to avoid higher shortage costs. As the DC Initial Inventory Level increases, the situation becomes more complex because there are a number of countervailing forces at work. Notice in Fig. 4 that there are essentially two possibilities as we increase the DC Initial Inventory Level: either the benefit of recourse diminishes (for the most part) as shortage costs increase for all values of I_1 (as in the $FRI = 3$ case) or a cross-over occurs when DC Initial Inventory Level gets large enough that the benefit of recourse switches from decreasing in shortage costs to increasing in shortage costs (when $FRI = 1$ or $FRI = 2$).

For the case where the benefit of recourse diminishes as shortage costs increase for all values of I_1 , there is high probability that the demand in both stages will be high and the second stage demand is less variable. Fig. 5 illustrates the first stage demand distributions for $FRI = 3$ in the top cell and the potential second stage demand distributions in the bottom cell. The mean demand levels calculated for Fig. 5 (and other figures that follow) come from the model presented in Table 4. From the decision tree in Fig. 3 we know that a storm threatening both regions initially will either continue affecting both regions with probability 5/6 or move away with probability 1/6, so the retailer knows that the second stage demand is likely to be very high. With a specific high demand distribution likely, there is less risk in sending out more inventory in the first stage in order to avoid higher shortage costs in both stages. As a result of these large demands and the reduced uncertainty surrounding the second stage distribution, the initial allocation in the recourse policy gravitates towards the no-recourse policy. As the shortage costs increase, the risk of a stock-out is magnified in both stages, and the retailer continues to increase its initial allocation towards the no-recourse allocation level. This in turn reduces the value of the recourse action.

As we saw in Fig. 5 and in Table 5, a storm that threatens both Region 1 and Region 2 will continue to target both regions if the storm transitions to a 3-day forecast. Contrast this with the first stage scenarios of $FRI = 1$ or $FRI = 2$ in Table 5, where a storm threatens Region 1 or Region 2 only, respectively. In both of these cases, the storm continues to affect the respective region alone or transitions to affect both regions. This adds a layer of uncertainty for the retailer who is now concerned with not just whether a

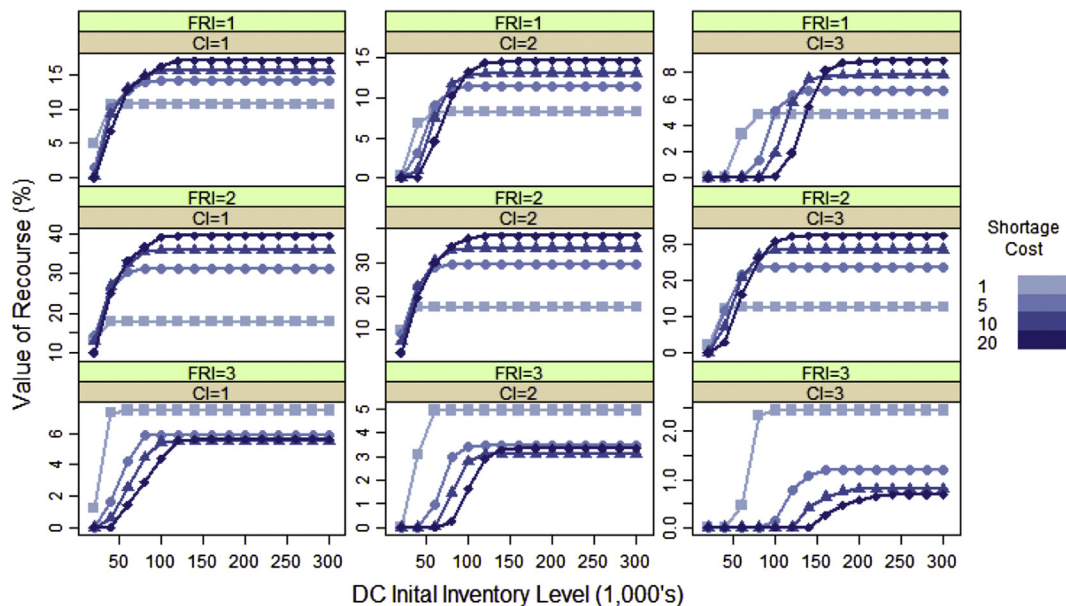


Fig. 4. Product 1 value of recourse (%) vs. DC initial inventory level (I_1) by shortage costs with transportation costs (c_r) set at a nominal level of 0.1.

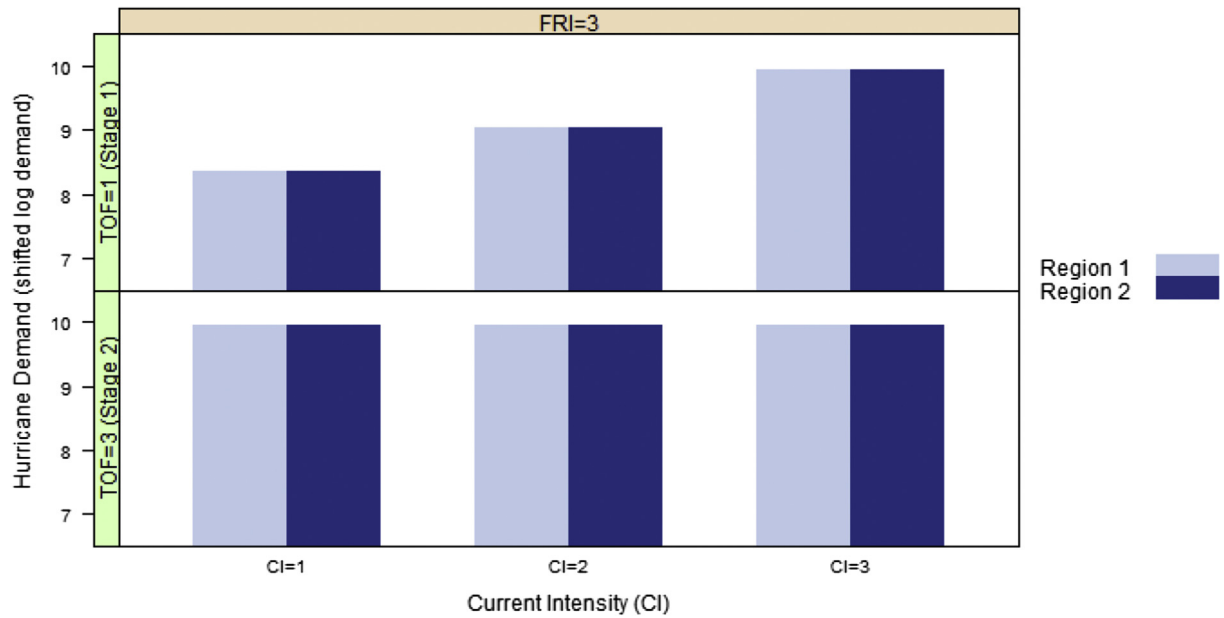


Fig. 5. Mean hurricane demand by state of storm and region for FRI = 3 in both stages.

storm will continue to threaten the Gulf, but also which demand distribution it may face.

Fig. 6 illustrates the situation when the storm is threatening Region 2 only (FRI = 2) in the first stage (note that the right column of the figure is a repeat of Fig. 5 and that a storm threatening both regions in the first stage will not threaten only Region 2 in the second stage). In the first stage, one of the distributions depicted in the upper left cell of Fig. 6 will be realized with the 5-day forecast. From Table 5, if the storm subsequently transitions to a 3-day forecast then there is a 4/5 chance the retailer will face the distributions in the lower left cell of Fig. 6 and a 1/5 chance the storm will be forecast to hit both Region 1 and Region 2 and face the

distributions in the lower right cell of the figure. Given the log scale for the hurricane demand, the difference between these two sets of distributions is not inconsequential. If the retailer prepares for the most likely scenario (FRI = 2 in the second stage as well) then a storm that changes to threaten both regions will leave the retailer with a tremendous shortage. Therefore, the retailer benefits tremendously from the ability to wait and see which storm forecast is realized in the second stage. This drives the high level of recourse value in the FRI = 2 cases of Fig. 4. As the shortage penalty increases, this additional information becomes only more valuable since it can decrease the shortage risk faced by the retailer.

Similar arguments regarding the cross-over nature of the




Fig. 6. Mean hurricane demand by state of storm and region for FRI = 2 and 3 for both stages.

shortage penalty's effect on VOR apply to the $FRI = 1$ case as well. However, it is clear from Fig. 4 that recourse is overall less valuable than in the $FRI = 2$ case. This is due to the fact that the first stage demand in Region 2 is similar to the $FRI = 2$ case but the demand in Region 1 is significantly larger (recall the model in Table 4 and see Fig. 7). A storm that initially threatens only Region 1 and continues to threaten the Gulf will continue to target only Region 1 with probability 4/5 or target both regions with probability 1/5. In this case the likely scenario is that the demand in Region 1 will be very high and the demand in Region 2 will remain comparatively lower (see the lower left cell of Fig. 7). There is still a shortage risk in Region 2 if the retailer's initial allocation is to satisfy the most likely scenario and instead $FRI = 3$ results in the second stage, however that risk is less likely than in the $FRI = 2$ case previously discussed. Therefore, the recourse option, while still valuable, is less so than in the $FRI = 2$ (Region 2 only) case.

We observed that the benefit of the recourse action varied greatly by Forecasted Region of Impact. Notice in Fig. 4 that VOR levels for $FRI = 2 > FRI = 1 > FRI = 3$ with the levels for $FRI = 2$ reaching over 40% compared to the case of $FRI = 3$ where the VOR barely exceeds 5%. In other words, when the storm is targeting only one region in the first stage, there is far more benefit of a recourse option than when there is more uncertainty over the expected landfall location.

5.3. Current intensity

As we discussed following Table 4, Current Intensity affects the demand distributions only for the first stage. As the Current Intensity increases from $CI = 1$ to $CI = 3$ the first stage demand increases for all Forecasted Region of Impact. Provided the storm does not move away from the Gulf, notice from Figs. 5–7 that the demand distributions in stage 2 will be at least as large as those in the first stage. For more intense first stage storms, the demands are therefore most likely to be high in both stages and the retailer can allocate more inventory in the first stage, which in turn yields a lower recourse value for the same reasons we discussed in Section 5.2. This explains why the VOR decreases as Current Intensity increases in Fig. 4.

From Fig. 4 we see that everything else equal, an increase in intensity can change the VOR by less than 10 percentage points, with only a couple of exceptions. In contrast, Forecasted Region of Impact can change VOR by more than 30 percentage points.

5.4. Transportation costs

In the case of a threatening hurricane, the retailer emphasizes shortage costs over transportation costs because of the humanitarian concern of supplying important commodities for the community. However, since it is clear that transportation costs could affect the allocation decisions and VOR, we considered a range of transportation costs in our numerical analysis. The results in Fig. 8 depict the changes in VOR as the transportation cost varies from 0.05 to 0.50 for different shortage costs under each first stage scenario with an initial inventory level at the DC of 200,000 (as Fig. 4 shows, this inventory level does not impose a constraint on the VOR).

Fig. 8 illustrates that transportation costs moderate the VOR under all storm scenarios. If the shortage costs are high, then the value dampens less quickly than when the shortage costs are relatively smaller. The retailer has emphasized that humanitarian concerns (embedded in shortage costs) are the driving factor for hurricane preparedness which supports the conclusion that recourse is still valuable under most storm scenarios.

6. Managerial insights

Our analysis provides a number of insights for both the demand estimation and inventory management aspects of this problem.

Weather forecasts available to consumers through news media contain a significant amount of predictive information on hurricane demand. We demonstrate this for two products with different hurricane demand characteristics.

More uncertain hurricane forecasts do not always lead to delaying inventory allocations. For example, when a five day forecast cone intersects both regions there is a great deal of uncertainty about the landfall location of the hurricane, yet there is relatively less value for a recourse action. When customers are faced with an uncertain

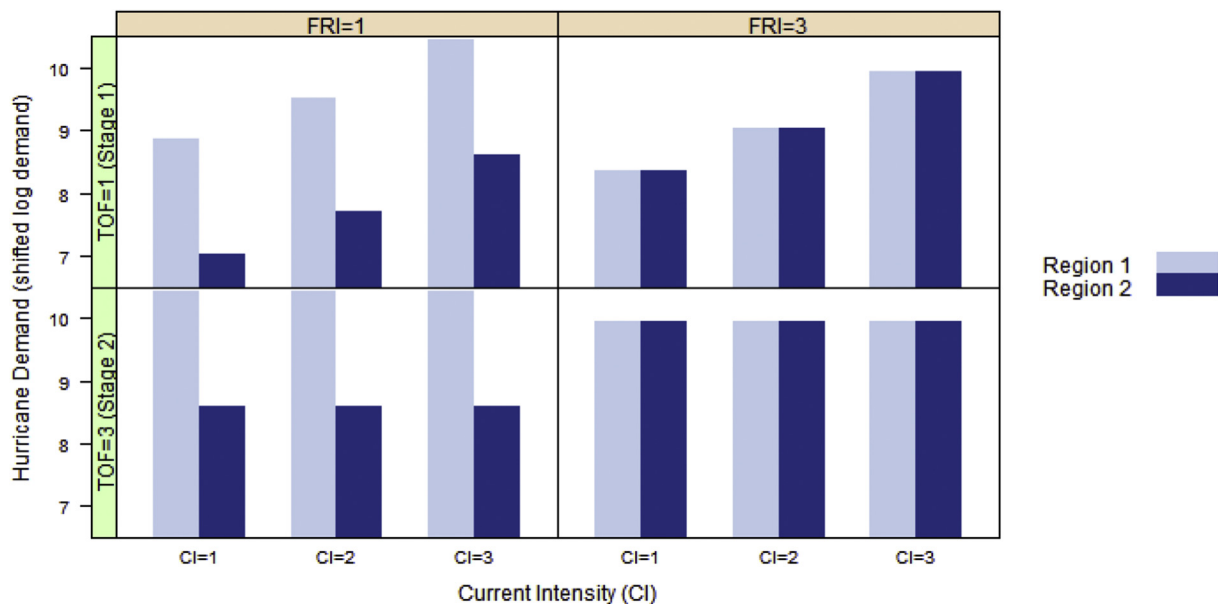


Fig. 7. Mean hurricane demand by state of storm and region for $FRI = 1$ and 3 for both stages.

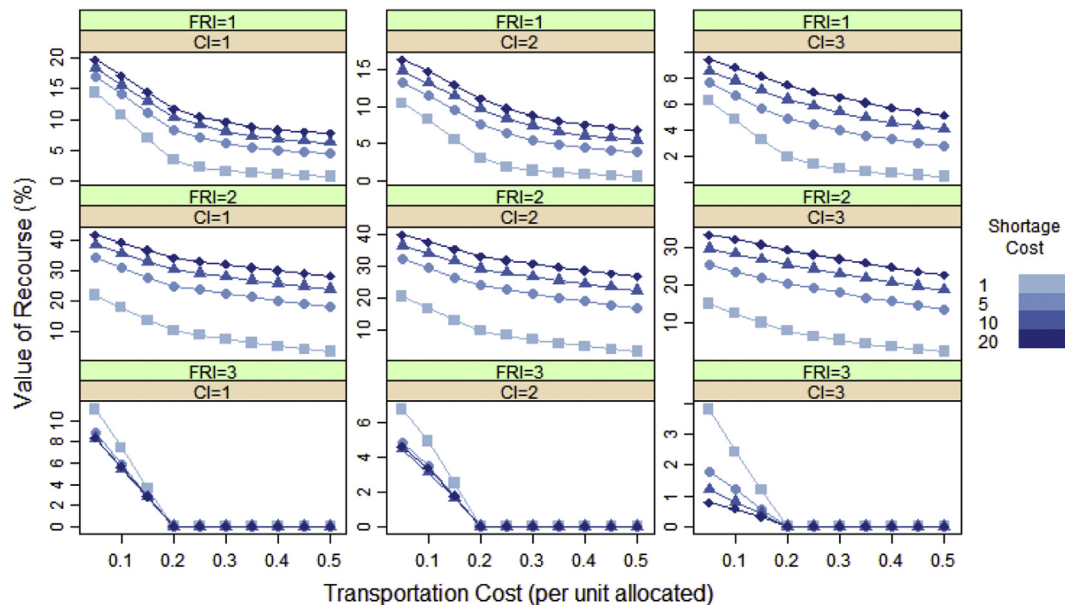


Fig. 8. Product 1 value of recourse vs. transportation cost with DC initial inventory level (I_1) = 200,000.

forecast, demand spikes in the second stage reduce risk for the retailer and lower the value of recourse.

The association between hurricane storm intensity and hurricane sales is comparatively weaker than expected landfall location. Since consumers appear to prepare in roughly the same manner for any size storm (tropical storm and above) in all but a 5-day forecast projected to strike only Region 1, this simplifies the resultant demand estimation and inventory management problem for the retailer.

Higher initial inventories at the DC motivate the managers to make a second inventory allocation to take advantage of an opportunity for recourse. When the initial inventory level at the DC is low, there is little value of recourse. As the initial inventories at the DC increase, it becomes preferable to make a partial allocation of the initial inventories and then wait for the 3-day forecast to revise the inventory position at the retail locations but eventually the value of recourse levels off.

When making a recourse decision, the impact of shortage costs should be considered in conjunction with the inventories at the DC and the initial forecast of the hurricane. As discussed in Section 5, with very low initial inventory levels, the retailer does not have the flexibility to delay inventory decisions and thus minimizes shortage risk by pushing out large portions of inventory holdings in the first stage. Higher shortage costs result in even larger first stage allocations to mitigate the shortage penalty in the first stage which reduces the recourse value. With sufficiently large initial inventory levels, the retailer benefits from recourse. However, as we have already demonstrated in Section 5, the shortage penalty's impact on the value of recourse at these higher inventory levels depends on the initial hurricane forecast in a non-trivial way. Since higher shortage costs reflect the retailer's humanitarian concerns, their impact on the inventory management approach used during a hurricane event need to be well understood.

Transportation costs moderate the value of recourse in a fairly predictable manner. While higher transportation costs dampen the VOR, the dampening tends to happen less rapidly for relatively higher shortage costs which is more representative of retailer's views when a hurricane threatens.

Initial forecast of the hurricane's landfall location is a key factor in

recourse decisions. The 5-day Forecasted Region of Impact has a significant impact on the allocation policy; e.g., when it includes only Region 2 (FRI = 2), then the managers should be more cautious about inventory allocation as the hurricane's impact on retail demand is highly unpredictable. In this case waiting and observing more information about the hurricane and partially deferring the inventory allocation after observing the 3-day forecast is more valuable. Recourse can create up to 40% more value compared to no recourse. In contrast, when the 5-day Forecasted Region of Impact includes both Regions 1 and 2, the value of recourse is much more modest (less than 7%).

7. Impact on practice

Community preparation for a hurricane involves the efforts of non-profit organizations, government agencies, and retailers that supply necessary goods to prepare for a storm. Our focus on this paper has been on how retailers can optimally supply necessary commodities for people in affected areas by considering the tradeoffs among lost sales, excess inventories, and transportation costs in a two-stage recourse model. Our work provided direct benefit to the retailer in this regard. Historically, the retailer used a no recourse model, but it was interested in investigating the value of recourse. We were able to show the retailer circumstances under which there can be significant benefit from using an optimal recourse policy. However, we were also able to demonstrate circumstances when recourse would not provide significant value even in the face of uncertain storm forecasts, because customers absorb the risk by purchasing large amounts of supplies.

Additionally, the retailer had been experimenting with simple recourse options of holding back a fixed fraction of hurricane depot inventory in stage one as a hedge against uncertainty in the second stage. However, since the retailer lacked a formal model for demand based on state of the storm forecasts, the fraction held back in the recourse policy was determined in an ad hoc manner. We showed the retailer that it could do better using a more formal state-space demand model and following an optimal state dependent recourse policy, including situations where no inventory is held back for the second stage.

As a result of our work and other initiatives, the retailer improved its hurricane operations management which enabled the retailer to play a strategic role in the state of Texas as a member of a special hurricane disaster preparedness task force (Texas Emergency Management Digest, 2009). Based on this public/private partnership, the State of Texas experienced marked improvement in its response to hurricane events. Thus, these improvements not only benefitted the retailer, but also population residing along the Texas Gulf Coast.

The retailer shared with us a number of direct learnings from their collaboration with the public sector during a hurricane event. Of particular relevance for this paper, the retailer described how the public and private sectors work together to leverage each other's strengths. Subject to safety requirements, the retailer focuses on keeping stores open and well-stocked so consumers can get the supplies they need, and the public sector officials focus on ensuring access to the stores (e.g., directing traffic, clearing roadways, maintaining public safety). As a corollary, both parties collaborate to do everything they can to enable people to take care of themselves as much as possible. The retailer informed us that the collaboration with the public sector began with the establishment of temporary distribution centers during every hurricane event. These were found to be less effective than leveraging the private sector's existing infrastructure. Consequently, the retailer found our research on better inventory allocation policies particularly beneficial to the strategy of keeping their stores open and stocked during a hurricane event. As a further learning, public sector officials rely on the retailer's understanding of consumer demographics to provide the right supplies to people in each part of the state. Our data analysis provided the retailer with improved insights on consumer purchasing patterns during hurricane events.

8. Conclusions

Inventory management under the threat of a hurricane poses challenges for managers in the retail industry. Uncertainty about the hurricane's possible path and its impact on demand needs to be carefully assessed to make sound inventory decisions to help communities prepare for a natural disaster. In this paper, we consider the inventory allocation decisions of a retailer from a DC with fixed capacity to two geographical regions, while anticipating the demand from a hurricane threat. In particular, an early allocation decision runs the risk of the hurricane moving to an unanticipated location leading to excess inventory in one location while leaving shortages in other locations. A wait-and-see approach provides more information about the possible path of the hurricane but runs the risk of unmet demand and lost sales. In practice, it is also possible to make sequential allocations as the hurricane information is gathered over time, but this recourse ability involves setup and transportation costs. Hence, the value of recourse is of significant managerial importance and determines when to pull the trigger and allocate inventories. Our analysis allows us to identify situations when it is better to make a large initial allocation vs. making a series of sequential allocations. The empirical approach developed in this paper can be applied to any retailer whose sales are impacted by an approaching storm or other predictable supply chain disruption.

We consider a retailer located in a developed country, which has access to good meteorological information during a storm event and a well-developed transportation infrastructure for shipping product to stores. Hence, our insights are likely more directly transferrable to parts of the world with similar characteristics. While the modeling framework is general enough to apply to dissimilar regions of the world, the model parameters (e.g., the costs and forecasts) might differ enough to result in different

recommendations from the model. For example, in a country with a less developed information and transportation infrastructure firms may benefit less from recourse due to their limited capabilities to exploit their recourse option. In addition, future advances in hurricane trajectory prediction may reduce the uncertainty faced by the retailers and the value of recourse.

While our empirical estimates were based on 15 storms, the state-space framework is flexible enough to conduct sensitivity analysis on the transition probabilities or to be updated easily as more storm histories become available. For example, through the end of 2012, seven additional storms threatened the retailer's regions resulting in five more 5-day forecasts and four more 3-day forecasts. Combined with more retailer sales data, it could be used to update both the regression models and the state transition probabilities of Section 3.2 and 3.3, respectively.

In this work, we developed an inventory management model that could be solved for each first-period storm state scenario. It may be interesting to consider planning for the hurricane season where the initial inventory position at the depot is a decision variable. Additional work could revolve around the optimal product mix for the retailer to provide to consumers stocking up for a hurricane.

References

- Assimon, S., 2009. H-E-B Gets Jump on Hurricane Preparedness. Last accessed January 18, 2016. <http://supermarketnews.com/latest-news/h-e-b-gets-jump-hurricane-preparedness>.
- Banjo, S., 2012. Big box retailers spring into action. Wall Str. J. October 31, 2012, Business Section.
- Birge, J.R., Louveaux, F., 1997. Introduction to Stochastic Programming. Springer-Verlag, New York.
- Eppen, G., 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Manag. Sci.* 25, 498–501.
- Eppen, G., Schrage, L., 1981. Centralized ordering policies in multi-echelon systems with lead-times and random demand. In: Schwarz, L.B. (Ed.), *Multi-Level Inventory/Production Control Systems: Theory and Practice*. North Holland, Amsterdam, pp. 51–67.
- Federgruen, A., Zipkin, P., 1984. Approximation of dynamic, multi-location production and inventory problems. *Manag. Sci.* 30 (1), 69–84.
- GAMS, 2016. General Algebraic Modeling System. Last accessed January 18, 2016. <http://www.gams.com>.
- Garcia-Donato, G., Forte, A., 2014. BayesVarSel: Bayesian Variable Selection in Linear Models. Last accessed December 31, 2014. <http://cran.r-project.org/web/packages/BayesVarSel/index.html>.
- Horwitz, S., 2009. Wal-Mart to the rescue: private enterprise's response to Hurricane Katrina. *Indep. Rev.* 13 (4), 511–528.
- Lodree, E.J., Taskin, S., 2009. Supply chain planning for hurricane response with wind-speed information updates. *Comput. Oper. Res. Spec. Issue Disaster Recovery Plan.* 36, 2–16.
- Lodree, E.J., Ballard, K.N., Song, C.H., 2012. Pre-positioning hurricane supplies in a commercial supply chain. *Soc.-Econ. Plan. Sci.* 46 (4), 291–305.
- McGavin, E.J., Ward, J.E., Schwarz, L.B., 1997. Balancing retailer inventories. *Oper. Res.* 45 (6), 820–830.
- NHC Forecast Verification, 2016. National Hurricane Center Forecast Verification. Last accessed January 18, 2016. <http://www.nhc.noaa.gov/verification/verify2.shtml>.
- NHC Graphics Archive, 2016. National Hurricane Center Tropical Cyclone Advisory Archive. Last accessed January 18, 2016. www.nhc.noaa.gov/archive.
- NHC Saffir-Simpson, 2016. National Hurricane Center about Saffir-Simpson Hurricane Wind Scale. Last accessed January 18, 2016. <http://www.nhc.noaa.gov/aboutshws.php>.
- Raths, D., 2010. Working together: private-sector organizations earn a seat in the emergency operations center. *Emerg. Manag.* 28–34. May/June 2010.
- Rawls, C.G., Turnquist, M.A., 2010. Pre-positioning of emergency supplies for disaster response. *Transp. Res. Part B* 44 (4), 521–534.
- Regnier, E., 2008. Public evacuation decisions and hurricane track uncertainty. *Manag. Sci.* 54 (1), 16–28.
- Regnier, E., Harr, P.A., 2006. A dynamic decision model applied to hurricane landfall. *Weather Forecast.* 21 (5), 764–780.
- Salmeron, J., Apte, A., 2010. Stochastic optimization for natural disaster asset prepositioning. *Prod. Oper. Manag.* 19 (5), 561–574.
- SAS, 2016. SAS Statistical Software. Last accessed January 16, 2016. <http://www.sas.com/>.
- Song, J., Zipkin, P., 1993. Inventory control in a fluctuating demand environment. *Oper. Res.* 41 (2), 351–370.
- Song, J., Zipkin, P., 1996. Managing inventory with the prospect of obsolescence.

- Oper. Res. 44 (1), 215–222.
- Target, 2016. Disaster Preparedness and Response. Last accessed January 18, 2016. <https://corporate.target.com/corporate-responsibility/safety-preparedness/disaster-preparedness-response>.
- Taskin, S., Lodree, E.J., 2011. A Bayesian decision model with hurricane forecast updates for emergency supplies inventory planning. J. Oper. Res. Soc. 62 (6), 1098–1108.
- Texas Emergency Management Digest, 2009. Texas team of Teams Ready to Roll. http://www.txdps.state.tx.us/dem/documents/digest/digest_2009_summer.pdf. Last accessed January 18, 2016.