The Power of Atlas

Why In-Store Shopping Behavior Matters

JACOB SUHER
The Wharton School
suherja@gmail.com

HERB SORENSEN
TNS Sorensen
herb.sorensen@
tns-global.com

This report introduces Atlas, a model that accurately describes shopper behavior across a series of supermarkets in terms of traffic across the store, share of shopping at various locations, resulting purchases, and the amount of time shoppers invest in these activities. The value of the tool to measure the in-store "audience" for advertisers is assessed as is its utility for allowing store management to evaluate a wide range of "what if" movements of categories to alternate locations, with alternate display sizes. This report focuses on center-of-store aisles but points the way to deployment for the full store, including end-caps and perimeter displays.

BACKGROUND

Shopper Behavior

Shoppers have been speaking to retailers and manufacturers with an increasingly clear voice since the 1970s. This voice has been heard predominantly through the votes they make with their dollars at the checkout counters. These are the ultimate "votes" based on actual behavior, not opinion or speculation on anyone's part.

The voting begins well before the checkout counter as the shopper selects first a particular retail establishment where purchases will be made and then their pathway through the store that will expose them to the varied wares and communication available. Each area the shopper visits, pauses for shopping or media intake, selects this or that items, and spends time in these activities constitute mini-votes leading to the placement of "coin-on-the-counter," the final expression of their consumer voice. This earlier behavior leading up to the purchase is very much a part of the shopper's voice. It is to this voice that we look for understanding of how and why they are speaking so clearly at the checkout.

Listening to the shopping voice—not just the buying voice—is what shopper behavior is all about. If this voice is to be anything but cacophony, we must learn to distinguish all the notes, tones, and chords. It is necessary not just to make measurements but to have a framework for understanding those measurements. The framework outlined here is based not only on common sense observations in stores, but on patterns observed in a score of stores subjected to detailed descriptive analysis.

In-Store Media

We will address issues here from the in-store advertiser's point of view and store, category, and brand management merchandising's point of view. Some of the parameters used have exactly the same numerical values for advertisers and management, but with potentially different meanings. For example, an advertiser is likely looking to *reach* the shopper whereas the self-service retailer is looking for the shopper to *visit* the merchandise. The advertiser is likely to expect to measure reach as exposures. For Atlas purposes, we will deal with reach and visits as similar concepts, both meaning that the shopper and the merchandise have occupied the same space at some point.

Further distinction between *physical* reach and *visual* reach is necessary. It is a truism that what is not seen does not exist for the shopper. Anything that appears in the field of vision has been seen. Given these considerations, we can tighten the first condition of a sale to *visual reach*. This distinguishes the physical presence from the visual

DOI: 10.2501/S0021849910091142

presence, *reach* from *visual reach*, or *exposures*. Visual reach is the proper conception of exposures. Only when this occurs can we fairly say that the shopper has been *exposed* to either merchandise or point-of-purchase media (P-O-P).

Media metrics provide both scientific understanding of the relation of media to the audience and a commercial basis for the sale of media exposures of the shopper "audience." As noted, everything physically reached is not visually reached. *The fact is that any accurate measurement of exposures in the store must be a fraction of physical visits to the media*. This is because the visual cone is only the field of vision in *front* of the shoppers' eyes, leading to true exposures of less than one-fourth of visits—a significant reduction in exposures available to be sold.

A major function of shopping is filtering out most of what is seen: that is, to discard most of what is visually reached to focus on, or engage with, very specific items or features of those items. That is, exposure may lead to some form of engagement, make an impression, but exposures themselves are *not* impressions. Impressions are tied to the actual *point* upon which the shopper's eyes focus, a tiny fraction—well less than 1 percent of the visual field (based on the size of the foveal vision, the actual point of focus area).

There are two reasons for going into this much detail about the shopping experience. The first is because the Atlas tool is based on accurate measurements of shoppers' physical movements around stores. The second reason (as detailed in "Long Tail Media in the Store," *Journal of Advertising Research* 48, 3), is that the number of exposures, with even fewer impressions, is quite small for this or that P-O-P, product or category. Think of this decrescendo:

Population > Stores > Visits/Reach
> Exposures > Impressions > Sales

Shoppers have been speaking to retailers and manufacturers with an increasingly clear voice since the 1970s.

We begin with very large numbers on the left and deliver very small numbers (for individual items) on the right. Media folk are accustomed to dealing with very large numbers on the left, and experience with mass media such as television leaves them unprepared to deal with the radical fall-off of numbers across this decrescendo. However, a holistic view of shopping requires that we start with the tiny individual sales—on the right—and deconstruct the events that lead to the purchase.

Visits Versus Purchases

A good deal can be learned about shopper dynamics in the store by considering shoppers' relationship to even a single category. For illustration, consider shoppers *visiting* the cookie-and-cracker aisle (Figure 1), and making *purchases*, across a series of nine representative national supermarkets.

The first thing to notice here is the relative constancy of category purchases across the stores. The average is 12 percent of baskets with a category purchase (±4 percent; 2 standard deviations). The reason for this constancy is the *relative* constancy of the human race: Any thousand shoppers across the United States will buy about the same amount of the cookie-and-cracker category. This is not to minimize the importance of the difference between

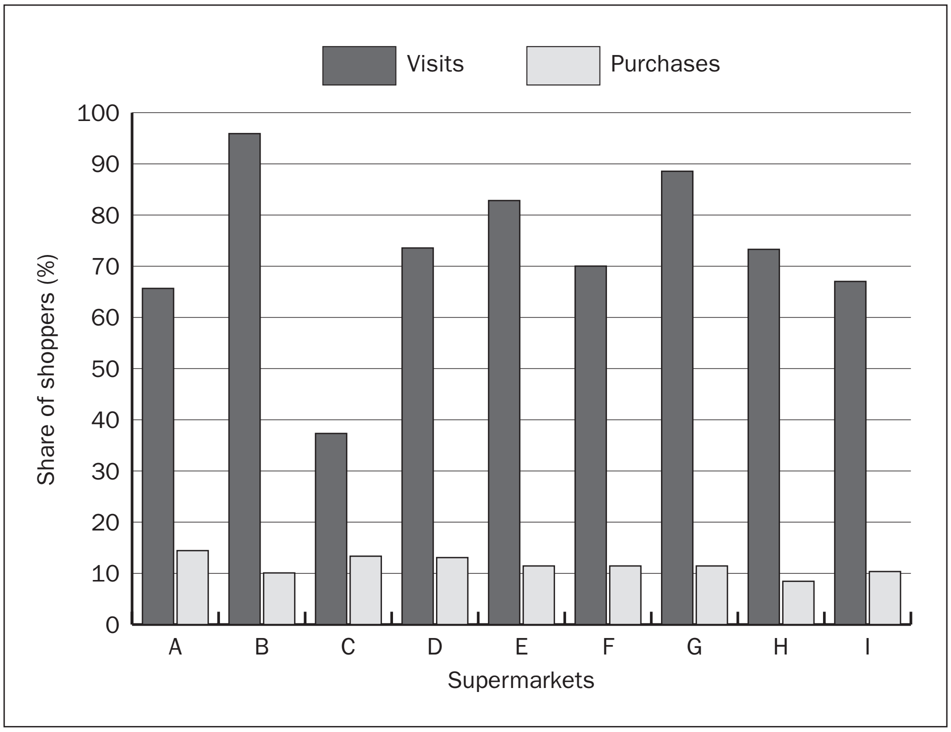

Figure 1 Cookie and Cracker Shopping Behavior

9 percent (the lowest share of purchasers) and 15 percent (the highest share of purchasers), a major difference in category sales performance.

The variation in purchase percentages is nothing in comparison to the variation in the number of shoppers visiting/reaching the cookie-and-cracker aisle across stores. This varies from a low of 37 percent to a high of 96 percent of shoppers visiting the category. Even more significant is that the lowest percentage of visitors (37 percent) delivered the second highest percentage of sales (14 percent), whereas the highest percentage of visitors (96 percent) delivered the second lowest percentage of sales (10 percent). The correlation between visits and purchases here is insignificant at the $p = 0.05$ level. This is a direct consequence of the fact that 60 to 80 percent of the typical shopping trip is not spent purchasing but rather is wasted with ineffective wandering (Hui, Bradlow, and Fader, 2007). Most aisles see a lot of traffic that is simply using the aisle for what aisles are meant to be: a way to get from point A to point B.

Although purchases are the primary focus of both retailer and shopper, the cookie-and-cracker example shows that understanding of the process that leads to the purchase is not a simple extrapolation from the purchase itself. There is a complex interaction among the categories, their geographic locations within the store, what point in the trip the shopper actually arrives at the category (early or near the end), and other factors. It is a systematic understanding of these factors that is the foundation of Atlas.

The overall objective is to develop a mathematical model of shopping behavior that will predict aggregate shopper behavior on a category-by-category basis, given the store design, merchandising, and inventory.

OBJECTIVE

The overall objective is to develop a mathematical model of shopping behavior that will predict aggregate shopper behavior on a category-by-category basis, given the store design, merchandising, and inventory. This is an ambitious goal that becomes achievable when broken into suitable component parts. Selecting the parts creates the framework to allow a methodical accretion of learning that stands a reasonable chance of delivering the objective. It should be noted that Atlas is not a speculative tool per se but systematically reports exactly how categories actually performed in various physical locations, with reasonable interpolations and extrapolations from similar category configurations, measured across stores.

METHODOLOGY

V-S-P/T

The ultimate goal is to predict shopper behavior. The model's dependent variables are the following:

- **V**isits: the percentage of trips that come within the vicinity of the designated category; also described as reach or, more specifically, physical reach
- **S**hops: the percentage of trips wherein the shopper slows up or stops within the vicinity
- **P**urchases: the percentage of trips wherein a purchase of an item/category occurs as evidenced by scan data at the checkout
- **T**ime: the elapsed time for one or more of the foregoing variables.

These constitute the foundational quartet of PathTracker measurements, which may be summarized as V-S-P/T (Sorensen, 2009.) A number of measures can be derived from these to allow focus on individual aspects of the shopping behavior. BuyTime (BT) is used to refer to the total amount of time that the shopper spends in the vicinity of an item that is ultimately purchased.

Domain

Shoppers exhibit varied distinctive behaviors based on their current display environment, or "domain." For example, a shopper instinctively behaves differently when faced with a wall of merchandise as compared to a series of end-caps or open spaces with lower display tables and fixtures. This study deals only with the long-constrained aisles in the center-of-store "domain." We select the center-of-store (CoS) aisles for this article primarily because strong, consistent statistical patterns have been observed in these aisles from store to store.

The CoS aisles are bounded on either side by a series of product displays, about 4 feet wide and 6 feet high. Each of these displays constitutes one product point (PP). The location of each PP is assigned both a continuous and categorical location in the store. Cartesian coordinates provide a continuous location for all PPs, and a nine-section "location" grid divides the CoS into discrete pieces.

The Data Set

The data set is a selection of six congruent stores from the PathTracker database. Store congruency is necessary to simplify the modeling process. All stores have similar sizes, counter-clockwise traffic flows, and contiguous CoS aisle domains with

	Product point level data	
Description	**Explanatory variables**	**Dependent variables**
6 stores	Observable store characteristics	Visits/reach %
All congruent aisles	Normalized aisle location	Shops %
All with-in aisle locations	Normalized with-in aisle location	Purchases %
All product points	Product point number	BuyTime
	Category (1, 2, ..., N)	

Figure 2 The Atlas Data Set

few transverse (cut-through) aisles. The data set is described in the chart above (Figure 2).

The model allows the input of the first five explanatory measures and delivers the last four dependent measures as output. The model output of shopper metrics (Visits/Reach, Shops, Purchases, and BuyTime) is useful in two formats. Performance predictions at the PP level will provide specific results for a 4-foot section of the CoS aisle. On its own, this information is not the final goal of a predictive engine. Category-level (CAT-level) performance needs to be delivered. The last four dependent measures need to be available at the CAT-level with their performance statistics correctly aggregated across product points.

Model Specification

The proposed model is a two-step process. Step 1 is to deliver PP-level performance. Step 2 is to compute CAT-level Visits/Reach percent, Shops percent, Purchases percent, and BuyTime (Figure 3). The sequential organization of the model will allow the output of Step 1 (PP level) to become the input of Step 2 (CAT-level).

SAS data analysis software uses the method of least squares to fit general linear models to the data. After testing and evaluating the in-sample fit and prediction of several analyses, multiple linear regression models were chosen for their ease of interpretation and parsimonious use of explanatory variables.

MODEL RESULTS

Both Steps 1 and 2 of the modeling process returned successful results. As a whole, the equations were very significant, with model level *p*-values of less than .0001 and strong adjusted *r*-squared values. The model statistics explain that above a 99.999 percent level of confidence, the equations have at least one significant parameter. The *r*-squared values tell us that Step 1 explains between 40 and 50 percent of the variation in the display-level data and Step 2 explains between 75 percent and 90 percent of the variation in the category-level data. These are outstanding model fit statistics. A quick examination of some

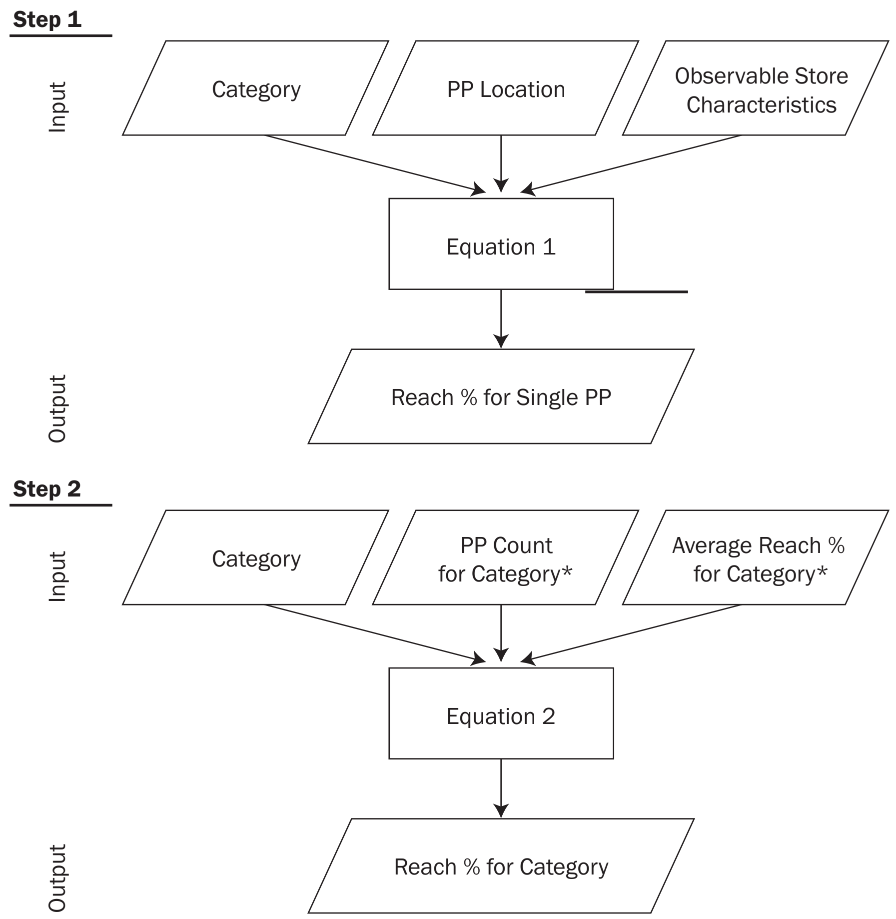

* Intermediate calculations compute PP Count for Category and Average Reach % for Category

Figure 3 Modeling Shopper Reach in Center-of-Store Aisles of the parameters' significance values demonstrates the power of our analytical framework (Tables 1 and 2).

The majority of the parameters of interest across models returned very significant results. A display's performance is predicted by its location within the CoS and which category is shelved there. An entire category's performance is predicted using its individual display statistics and the number of displays assigned to the category.

Model Validation

In keeping with the two-step modeling process, two questions must be asked to validate the performance of the models specified earlier. First, does Step 1 provide accurate predictions of display level performance when predicting a real store outside of the dataset? Then, are the CAT-level predictions from Step 2 supported by historical category performance across supermarkets?

A "hold-out" test was performed to validate the success of Step 1. The sixth and final store in the original data set was initially omitted from model development. This limited-data model was used to generate performance predictions for the omitted store. The mean absolute error across all CoS displays in the validation sample was 12.13 percent for Visits/Reach percent, 1.68 percent for Shops percent, 1.09 percent for Purchase percent, and 0.60 seconds for BuyTime. The hold-out test demonstrated that the model does an adequate job of capturing the immense variability of human shopping behavior.

The CAT-level predictions from Step 2 reasonably meet expectations created by years of tracking grocery store performance. The completed model was used to simulate an existing grocery store's layout. This exercise demonstrated the ability of the model to predict CAT-level performance within a single standard deviation of the actual data.

Model validation tests are an important step in modeling CoS data. Beyond providing objective feedback on prediction accuracy, data used for validation tests quickly become incorporated into the model. The omitted store used to validate Step 1 had unique store-level characteristics. Though the model did a reasonable job during the hold-out test, the predictions were the result of extrapolation. The hold-out store is now incorporated in the dataset. As the underlying data continues to expand, so will the capabilities of the prediction engine.

DISCUSSION

It will be helpful to look at a simple Atlas application to better understand its implications for both in-store media, and the management of categories and merchandising (Figure 4).

Here we see the effect of moving the 10 carbonated soft drink displays (40 ft.) from the back of the first aisle on the right (typically the first aisle shopped in a right-entry store) to the front of the first aisle on the left (typically the last aisle shopped in such a store). The inputs are variables supplied by the Atlas user, including the number and locations of displays. Prices

TABLE 1
Step 1 – Display-level Parameter Statistics

	Visits/Reach %		Shops %		Purchases %		BuyTime	
Parameters	***F* value**	**Pr > *F***	***F* value**	**Pr > *F***	***F* value**	**Pr > *F***	***F* value**	**Pr > *F***
Aisle Location	9.05	<0.0001	30.21	<0.0001	9.59	<0.0001	2.61	0.0502
Within Aisle Location	40.4	<0.0001	11.11	<0.0001	4.4	0.0043	41.58	<0.0001
Category	19.02	<0.0001	19.89	<0.0001	14.21	<0.0001	3.9	<0.0001

TABLE 2
Step 2 – Category-level Parameter Statistics

	Visits/Reach %		Shops %		Purchases %		BuyTime	
Parameters	***F* value**	**Pr > *F***	***F* value**	**Pr > *F***	***F* value**	**Pr > *F***	***F* value**	**Pr > *F***
Display Performance	994.99	<0.0001	570.22	<0.0001	19.42	<0.0001	1578.04	<0.0001
# of Displays	14.11	<0.0001	105.53	<0.0001	15.45	<0.0001	1.35	0.2456
Category	12.98	<0.0001	11.58	<0.0001	9.4	<0.0001	3.67	<0.0001

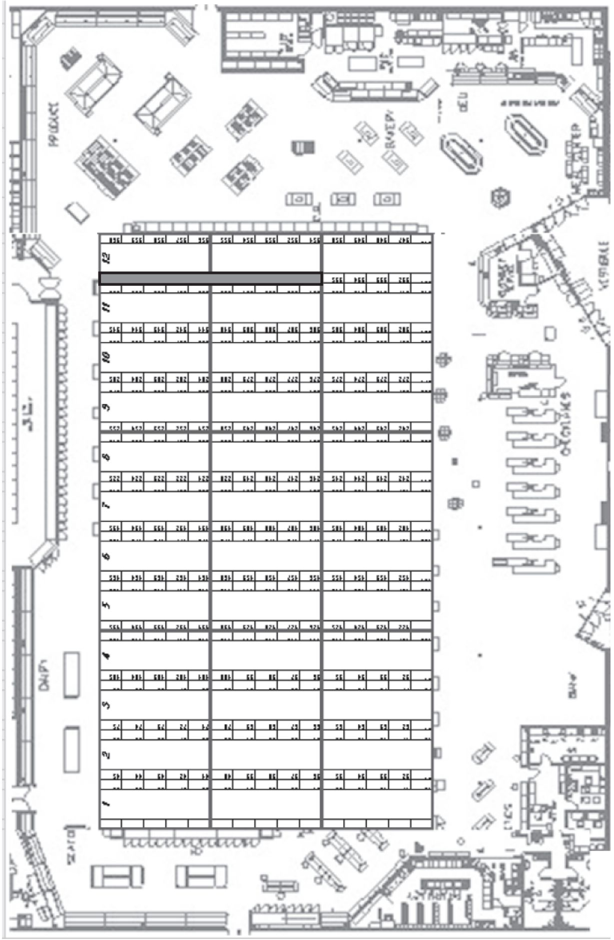

Soft Drinks – early in the trip

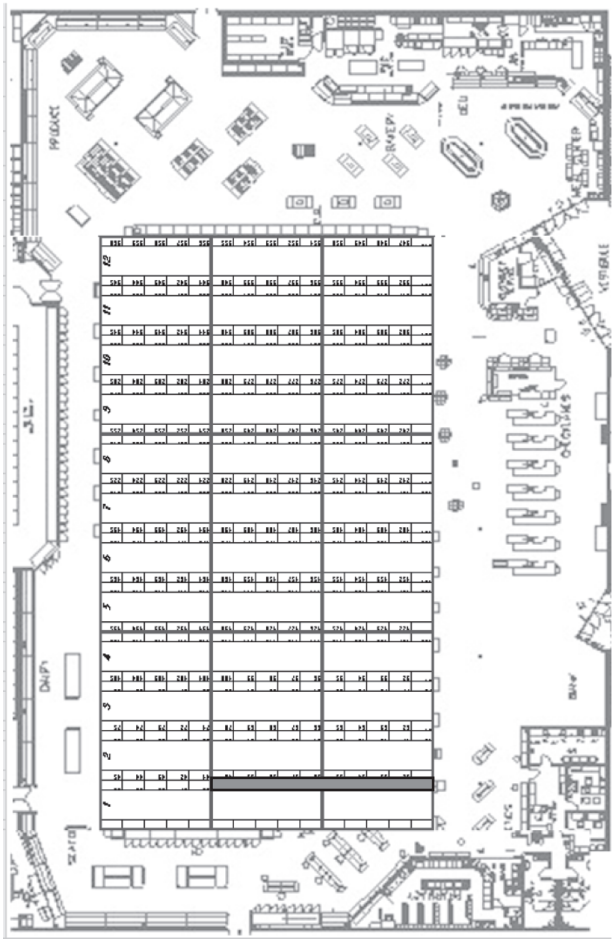

Soft Drinks – late in the trip

Carbonated Soft Drinks: Early-in-Trip vs. Late-in-Trip Comparison

	Atlas Input			Atlas Shopper Metrics – Output							
Position	Displays (× 4ft)	$ Price average	Margin % average	Reach (Visits)	Stopping V to S	Shops	Closing S to P	Purchase	Holding BuyTime	$ Sales expected	Margin expected
Early	10	$2.85	3.1%	36.5%	62.9%	22.9%	27.0%	6.2%	11.6	$3,529.87	$109.75
Late	10	$2.85	3.1%	30.4%	71.8%	21.8%	26.7%	5.8%	12.9	$3,329.07	$103.50
Change				–6.1%	8.9%	–1.1%	–0.2%	–0.4%	1.3	–$200.81	–$6.24

Figure 4 A Simple Atlas Application

are average category spends, not individual items, and margins are category margins.

Notice that the reach drops substantially in going from early to late, and it takes a bit longer to close the sale (BuyTime). However, the impact on sales themselves is quite modest. The stopping power at the later location is substantially higher, which leads to nearly the same amount of shopping and purchasing.

This example illustrates that an advertiser will *reach* more shoppers in the first aisle, but that won't necessarily result in more sales. This is one reason we have pointed out (*Journal of Advertising Research* 48, 3) that the small numbers of exposures that actually close sales is of genuine value compared to broadcasting to large numbers of disinterested shoppers, which may simply contribute to much of the "noise" in the store.

The example also illustrates how store management can have confidence in moving categories, increasing/decreasing display size, and seeing what the impact will be on sales and margins. It should be noted that this example focuses only on the CoS. For categories such as carbonated soft drinks, as much as 50 percent of sales may come from end-caps and other secondary displays.

In-Store Media

Just as the Atlas model provides an accurate measure of the shopper's actual behavior, foot by foot and second by second, down the aisle, the Atlas *visits* data are an accurate measure of the *physical reach* of the P-O-P in terms of the number of shoppers who come within range of this media. This of course presupposes knowledge of the distribution of the categories and placement of the media in relation to those locations. In its initial implementation, Atlas will not provide EyeShare measures for CoS aisles but will provide these for end-caps at the ends of the aisles. EyeShare is a computed measure of exposure of displays based on the accumulated shopper traffic (Sorensen, 2006).

Managerial Implications

The mathematical models developed here challenge the prevailing belief that product categories will perform the same regardless of geographic location within the stores aisles. No longer can managers trust the retail idea that "if you build it, they will come."

Managers will be able to immediately utilize this study's results. The parameters for the independent variables of product categories and store locations can lead to actionable decisions. Product categories can be segmented by their ability to drive performance metrics. Store locations can be leveraged to influence the success (or failure) of products. The combination of these two sets of parameters allows both the retail manager and product brand manager to make educated decisions on the placement of their products across the CoS aisles.

The hold-out test explores a further level of managerial implications. With a successful prediction model, retail decision makers have the ability to manage their stores with unprecedented ease. Rather than physically rearranging a store and waiting to collect new receipt data, the hold-out model can immediately predict PP performance.

Retail Performance Modeling Implications

These modeling efforts neatly fit between two large pools of store management data. On one side, numerous papers explore the optimal placement of products on the vertical shelf space. Commercial applications such as Nielsen's Spaceman have brought this research into retail managers' hands. On the other side, research has been completed studying a retail store's physical location in relation to neighborhood demographics. Another Nielsen product, Spectra, clusters stores and their customers to increase return on investment and store efficiency. The importance of physical location driving the Atlas model bridges the gap between these two research pools. The placement of product categories within the physical space of a store is an original contribution to the retail performance-modeling field. In the future, the most advanced researchers will learn to combine and leverage several pools of research to understand the complete shopping experience.

OBJECTIVE RECAP

The overall objective of the Atlas: Modeling Shopping Behavior project was to develop a mathematical model of shopping behavior that will predict stores' sales on an item-by-item, category-by-category basis, given the store design, merchandising, and inventory. Much of this has been achieved. The CoS model presented in this study predicts shopping behavior on a category basis given store design and merchandising.

LOOKING FORWARD

The immediate future of developing mathematical models of shopping behavior is to expand outside of the CoS domain. With enough data, an entire store could be modeled with the same methods developed in this article. We are now adapting the methodology for end-cap displays and perimeter domains. The promotional displays and aisle end-caps are of particular interest to managers and academics alike. These dynamic locations change most frequently and are often used for in-store promotions, a controversial retailing practice. Though these new territories provide unique challenges of their own, they will benefit from a standardized method of prediction developed from the CoS data. Initial data sampling and statistical analysis provide evidence that the new domains will be just as successful as the CoS.

The use of equations to predict future performance, based on historical performance, reveals the underlying characteristics of each square foot of store real estate. Which areas of the store drive purchases? What is the best location to accelerate the performance of a slow-moving product? Atlas answers these types of questions with actionable solutions.

Further, Atlas provides highly detailed and accurate in-store media metrics, associated with categories and their locations. However, it also provides detailed shopping characteristics in the environment wherein media planners may intend to position in-store communication. This includes, in addition to physical reach, stopping power, holding power, and closing power: stopping and closing representing the necessary conversions from reach to effect purchases. JAR

Jacob Suher is a visiting researcher for the TNS Global Retail & Shopper Practice. He graduated in the spring of 2009 from The Wharton School at the University of Pennsylvania with a bachelors of science in economics. His research interests mainly focus on holistic shopper behavior and service operations management.

Herb Sorensen is the Global Scientific Director, Retail & Shopper Insights at TNS Sorensen, where he presides over the development of new services and research methods. He is the author of *Inside the Mind of the Shopper: The Science of Retailing* (Wharton School Publishing, 2009).

REFERENCES

Hui, S. K., P. S. Fader, and E. T. Bradlow. "The Traveling Salesman Goes Shopping: The Systematic Deviations of Grocery Paths From TSP-optimality." *Marketing Science*, forthcoming.

Sorensen, H. "The In-store 'Audience'." *Journal of Advertising Research*, 49, 2, (2009): 170–175.

Sorensen, H. "Long Tail Media in the Store." *Journal of Advertising Research*, 48, 3, (2008): 329–338.

Sorensen, H. "Purchase Selection Behavior Analysis System and Method Utilizing a Visibility Measure." U.S. Patent No. 7,006,982 B2. February 28, 2006.

Sorensen, H. *Inside the Mind of the Shopper*. Philadelphia, PA: Wharton School Press, 2009a.

Sorensen H. "Shopping Environment Analysis System and Method with Normalization." U.S. Patent No. 7,606,728 B2, 2009b

APPENDIX

The Appendix contains a summary of the statistical tests and robustness analysis results that were carried out (1–3). It also contains material relevant to retail store geography (4).

1. Variable Transformations

The set of explanatory variables were inspected for normality and other underlying regression assumptions using graphical methods. Graphical displays (e.g., histograms, stem-and-leaf plots, and normal quantile plots) revealed that some variables violated the normality assumption. Transformations of these variables were assessed for improved normality and their ability to predict the dependent variables in univariate regression.

2. Multicollinearity Tests

Both steps of the final model were checked for collinearity using the variance inflation factor test. If an explanatory variable exhibited collinearity, it was removed from the model, and the regression was re-fitted with that variable excluded. This process was repeated until all collinearity was resolved. The final models proposed in this study are absent of collinear variables.

3. Excursions

The final Atlas model is the result of many data and modeling considerations. This section provides a brief description of the various model excursions that were conducted. Our model excursions fall into three categories: the *location* hypothesis, observable store characteristic specification (Step 1), and Step 2 development.

The premise of this study is based on the null hypothesis, H_0, that *location* is not a significant predictor of store performance (i.e., visits, shops, purchases, and buy-time). To explore and ultimately reject H_0, we specified models that isolated either location or category performance. The result of this was two additional nested models. The *location*-only model was significant overall and contained several statistically significant location parameters giving us evidence to reject the null hypothesis.

Step 1 of the final model utilizes observable store characteristics as explanatory variables in replacement of store specific parameters. This allows for the prediction of supermarket performance of stores outside the dataset. Observable store characteristics were considered through the use of scatter plots, correlation measures, and univariate regressions. We used these visual and quantitative measures to select a parsimonious set of store characteristics that adequately replaced the individual store dummy variables. The variables were selected on the basis of their ability to control for performance variation across stores while remaining independent of each other. We were able to replace the store covariates with three observable store characteristics that minimally reduced the final model's explanatory power.

The dependent measure of the final model's second step is CAT-level performance. Several key pieces of information and alternative models helped inform the specification of Step 2. The high correlation between average category display level performance, display size, and category level performance was evidence that the prediction of category level performance was feasible. Several model iterations determined the set of independent variables that best explained category performance. From this, Step 2 was specified to include the interaction between display size and average display performance among other explanatory variables.

4. Center-of-Store Domain

Not all CoS aisles are created equal. We make some preliminary sub-classifications by calling out the following configurations:

- 50+ × 7 walled aisles (visual isolation from adjacent aisles)
- 50+ × 7 broken/walled aisles (transverse aisle intersects series of aisles)
- Shorter versions of the first type
- "Stub" walls (no visual isolation on right, left or both)–freezer coffins
- aisle widths greater than 7 feet
  - Freezer doors
  - "Bazaar" type widened aisle
  - Other extra-wide aisles (similar for extra-narrow aisles?)
- Others (e.g., shorter versions of stub walls, wide aisles)

In general, all of the CoS aisles are bounded on either end by some portion of the perimeter racetrack, although a bazaar or other domain may terminate an aisle. However, there are two circumstances wherein the racetrack may create abnormal CoS behavior. Without detailing all aspects of the racetrack here (reserved for later modeling), suffice it to say that the racetrack has four main sub-classifications:

- Ascension: whereby the shopper moves from the front of the store to the back
- Rear transverse: across the back of the store, typically between the rear end-caps and the perimeter wall
- Descension: typically the final or dominant path by which a shopper may return to the front of the store
- Front transverse: across the front of the store, typically between the front end-caps and the checkouts or service areas

Ascension and descension seriously affect CoS aisle behavior. At a minimum, they create an extraordinary amount of traffic in the affected aisles, a good deal of which may be simply passing through. Also, the direction and/or magnitude of flow may be altered as shoppers may be more set on getting to the back (or front) of the store than is typical for an aisle in that particular location in the store.

Now that we have looked at the selection and description of the domain, we will turn our attention to the internal structure of a typical CoS aisle, which we may refer to as the subdomains constituting the domain. The CoS domain then consists of a series of aisles, which are in turn compartmentalized into subdomains.

Each aisle is bounded on either side by a series of product displays, typically about 4 feet wide and 6 feet high. Each of these displays constitutes one PP. For computational purposes, each PP is defined by x, y coordinates of the point on the floor at the center of the display. At a minimum, the location of every SKU in the store is defined by its PP(s). If detailed planograms are available, the exact x, y, z coordinates of the product (including its height from the floor) can be used. The aisle is divided longitudinally in half to give a left and right half (defined as looking down the aisle from the front of the store.) It is also divided from front to back into six equal-length areas. This effectively subdivides the aisle into 12 segments. A typical 60-foot aisle will have approximately three PPs per aisle segment.