

Leveraging Trends in Online Searches for Product Features in Market Response Modeling

Evolving tastes can change the relative importance of product features in shaping consumers' purchase decisions, which in turn can shift the relative attractiveness of products with different feature levels. The challenge lies in finding a reliable yet cost-effective way to monitor the weights consumers place on various product features. In the context of the U.S. automotive market, the authors explore the potential of using trends in online searches for feature-related keywords as indicators of trends in the relative importance of the corresponding features (e.g., fuel economy, acceleration, cost to buy, cost to operate, body type). By augmenting marketing-mix data with feature search data in a market response model, they show substantial improvements in goodness-of-fit both in and out of sample. The authors find empirical support for the hypothesis that feature search trends are positively correlated with feature importance trends. They discuss how managers may make better decisions by monitoring feature search trends and leveraging those trends strategically.

Keywords: Google trends, market response model, product feature, conjoint analysis, big data

At a fundamental level, the attractiveness of a product can be viewed as a function of two sets of factors: product features and marketing efforts (e.g., advertising, incentives). Because the feature levels of existing products typically do not undergo major changes often, marketers have focused on marketing efforts as the main force in shaping the dynamics of product sales (Hanssens, Parsons, and Schultz 2001). As a result, most existing market response models treat the baseline attractiveness of a product—the part of product attractiveness that is not tied to marketing efforts—as a nuisance parameter (e.g., Ailawadi, Lehmann, and Neslin 2001; Montgomery and Rossi 1999; Neslin 1990).

This study was motivated by the notion that many factors beyond marketers' control can lead to shifts in the relative importance of various product features in shaping consumers' purchase decisions. Such shifts can change the relative attractiveness of products with different feature levels. For example, increasing concerns about climate change may cause fuel economy to weigh more heavily in consumers' vehicle purchase decisions. In this case, all else

being equal, vehicles offering higher (lower) gas mileage would become more (less) attractive.

The key challenge in practice lies in finding a reliable yet cost-effective way to track how the weights that consumers place on various product features evolve over time. To do so, one option is to run repeated conjoint studies over time, which would enable researchers to generate trend lines that can reveal the evolution of each feature's weight in determining consumer utility trade-offs. By incorporating such trends in a market response model, researchers could capture how the importance of various features has shifted over time and how such shifts have led to the observed dynamics in sales beyond what marketing efforts can account for.

However, monitoring trends in feature importance weights through repeated conjoint studies can be cost prohibitive and time consuming, especially when the analyses need to be carried out at a high frequency (e.g., monthly) and the number of respondents needed for a representative sample is large. Furthermore, in light of declining response rates, it has become increasingly challenging to track the evolution of consumer tastes through surveys. In this study, instead of monitoring consumers' stated preferences with methods such as repeated conjoint or tracking surveys, we explore a new, unobtrusive data-gathering alternative that has become viable with the advent of numerous online consumer interest tracking services (e.g., Attensity.com, ConverseOn.com, Conversation.com, DataSift.com, Lithium.com, Synthesio.com, NetworkedInsights.com, Sysomos.com, VisibleTechnologies.com). Among such online consumer interest tracking services, Google Trends (www.google.com/trends, previously known as Google Insights for Search) is probably the best known and most widely used.

Rex Yuxing Du is Hurlley Associate Professor of Marketing (e-mail: rexdu@bauer.uh.edu), and Ye Hu is Associate Professor of Marketing (e-mail: yehu@uh.edu), Bauer College of Business, University of Houston. Sina Damangir is Assistant Professor of Marketing, College of Business, San Francisco State University (e-mail: damangir@sfsu.edu). The authors thank Betsy Gelb, Jackie Kacen, and participants at the 2012 Marketing Science Conference and the 2014 American Marketing Association Advanced Research Techniques Forum for helpful comments and suggestions. They also thank the *JM* review team for their numerous inputs in improving the original manuscript. Peter Danaher served as area editor for this article.

As consumers become more dependent on the Internet for product information, their reliance on search engines as a gateway increases. This has opened a promising new avenue for tracking shifts in consumer interest by monitoring changes in the intensity of searches for various product-related keywords (e.g., brand names). Indeed, recognizing the potential value of such tracking data to marketers, Google introduced Google Trends in 2008; although it is meant for marketers, it is accessible to any user, free of charge.

As a potential source for marketing intelligence, Google Trends presents several appealing characteristics. First, it allows for the tracking of any queries that consumers have typed into Google's search box, generating volume indexes going as far back as January 2004. Second, data from Google Trends are updated in near real time and aggregated on a weekly basis (or daily, for the most popular queries), enabling users to track consumer interest with little time delay. Third, search volume indexes from Google Trends are highly customizable. For example, search terms can be combined or excluded to formulate composite queries, and searches can be filtered by geography (e.g., countries, states, cities), time range (e.g., July 2010 through May 2013), and category (e.g., Autos & Vehicles, Computers & Electronics). Fourth, Google is by far the most popular search engine. According to the Pew Research Center's Internet & American Life Project Poll (Purcell, Brenner, and Rainie 2012), Google is the search engine most often used by 83% of U.S. Internet users, followed by Yahoo (6%) in distant second. Given the ubiquity of consumer online searches and Google's dominance in this space, the volume of Google searches can plausibly be viewed as a reflection of the collective intentions of Internet users (Batelle 2005).

In this study, we tap into Google Trends and extract search volume indexes for product feature-related terms ("feature search trends" hereinafter). As we demonstrate in our empirical analysis, feature search trends can potentially serve as reflective indicators of trends in the importance weights consumers place on the corresponding product features ("feature importance trends" hereinafter). As motivating examples, Panels A–D of Figure 1 illustrate four sets of feature search trends among U.S. consumers between 2004 and 2011.

Figure 1, Panel A, plots the search trends for four food product-related features: calories, carbohydrates ("carbs"), cholesterol, and fiber. Among queries that Google classifies into the "Nutrition" subcategory (under the "Health" category), we observe that search interest for calories increased steadily and substantially over the years (more than 80% higher at the end of 2011 than 2004). In contrast, search interest for cholesterol declined by approximately 40% during the same period. As for carbs, the intensity of consumer searches followed a U-shaped trend line, bottoming in 2005 and rising since that time. In contrast with the large movements in search interests for calories, cholesterol, and carbs, the intensity of consumer searches for fiber remained largely stable.

Figure 1, Panel B, plots the search trends for four laptop computer-related features. Among queries that Google classifies into the "Laptops & Notebooks" subcategory, we

observe that search interests for memory and central processing unit (CPU) have declined or remained largely flat, whereas search interests for screen size and battery life have increased substantially over the years (increasing by more than 250% for screen size and nearly 200% for battery life). Similarly, Figure 1, Panel C, shows a diverging pattern for two features associated with digital cameras: whereas search interest for resolution remained largely flat, search interest for weight increased by more than 60%. Finally, Figure 1, Panel D, plots search trends associated with five vehicle features, showing that (1) searches for terms related to fuel economy experienced major fluctuations, (2) searches for terms related to cost to buy and cost to operate increased substantially, (3) searches for terms related to acceleration declined substantially, and (4) searches for sport utility vehicles (SUVs) bottomed in 2008 and bounced back afterward.

In summary, the trend lines presented in Figure 1, Panels A–D, illustrate that consumer online searches for terms related to various product features can vary substantially over time and follow very different patterns. This raises an important empirical question: Are feature search trends positively correlated with feature importance trends? If so, it would mean that (1) feature importance can also vary substantially over time and (2) marketers can potentially treat feature search trends, which have become readily available through online consumer interest tracking services such as Google Trends, as indicators of feature importance trends, which can otherwise be difficult to monitor over time.


To address this empirical question, we set out to investigate the extent to which feature search trends can help predict product sales, beyond the effects of marketing efforts, and whether feature search trends relate to product sales in a manner that is consistent with the hypothesis that feature search trends are positively correlated with feature importance trends. Our aim with this research is to contribute to the marketing literature in two main ways:

1. *We aim to demonstrate a creative use of Google Trends, a promising new source of consumer intelligence.* So far, the most common uses of Google Trends data, by marketing academics and practitioners alike, have focused on search trends related to product brand names, which are treated as leading indicators that can help forecast demand. However, consumers search online for much more than product brand names. We show that marketers should go beyond the "brand focus" and tap into search trends related to nonbrand keywords such as product features, which can potentially reveal more fundamental changes in the consumer's underlying preference structure. Such a broadened perspective can open the door to a whole new set of possibilities in consumer interest tracking.¹
2. *We aim to present a novel, readily implementable market response model that enables managers to better leverage trends in evolving consumer tastes.* Most market response


¹For example, a major U.S. food manufacturer has adopted our proposed approach. Managers at the company have begun to systematically track search trends for various food ingredients, nutritional facts and dietary benefits and concerns, which has shed more light on consumer tastes than merely tracking searches for food product brand names.

FIGURE 1
Google Trends for Feature Searches

A: Category Filter: Nutrition




B: Category Filter: Laptops & Notebooks



models in the extant literature have treated product baseline attractiveness as a nuisance that is either time invariant or purely stochastic. We propose a novel market response model in which the baseline attractiveness is modeled as a function of feature levels and the importance weights consumers place on each feature. We allow the feature importance weights to evolve over time as a function of the corresponding feature search trends. Like conjoint analysis, our model produces estimates of feature importance weights, which managers can use to quantify how sales would shift in response to changing feature levels. However, unlike

conjoint analysis, which requires individual preference data based on designed experiments, our model is more readily implementable in the sense that, in addition to whatever managers have been using for market response modeling, it only requires data about product feature levels and feature search trends, both of which are readily available. Equipped with our model and projected feature search and, thus, feature importance trends, managers can make not only better sales forecasts but also better decisions in product design, budget allocation, ad copy development, and so on.

FIGURE 1
Continued



Notes: To generate these time plots, we gathered Google Trends data for the keywords related to the corresponding features. For Panels A, B, and C, the queries are the same as the labels. For Panel D, we used the queries listed in Table 1. We used a separate category filter for each panel: “Nutrition” for Panel A, “Laptop & Notebook” for Panel B, “Camera & Photo Equipment” for Panel C, and “Autos & Vehicles” for Panel D. We calculated the 12-month moving averages to make the long-term trend lines easier to discern. For each plot, the first point corresponds to December 2004, which is based on the average of 2004. The plots show the percentage change in the moving average, as compared with 2004.

The remainder of the article proceeds as follows. The next section provides a quick overview of two literature streams: one focusing on how existing market response models have dealt with the dynamics of product baseline attractiveness and the other on how Google Trends data have been used as tracking measures of consumer interests.

We then present our proposed model. To illustrate, we use sales and marketing-mix data for 80 major passenger vehicles sold in the United States between January 2004 and April 2011, augmented with Google Trends data for keywords that are commonly associated with fuel economy, acceleration, cost to buy, cost to operate, and vehicle body

type. The empirical results show that our proposed model substantially outperforms benchmark models that do not leverage feature search trends, both in and out of sample. Furthermore, the estimates of how feature search trends moderate the relationship between sales and feature levels are all significant and of the expected signs: an upward trend in searches for desirable features (features for which, all else being equal, the higher the feature level, the better for consumers; e.g., fuel economy, acceleration) makes sales more positively elastic to levels of these features. In contrast, an upward trend in searches for undesirable features (features for which, all else being equal, the lower the feature level, the better for consumers; e.g., cost to buy, cost to operate) makes sales more negatively elastic to levels of these features.

After establishing predictive and face validity, we explore the managerial implications of our research: beyond improving sales forecasts, managers can improve the bottom line by monitoring feature search trends and leveraging those trends strategically. To conclude, we discuss the limitations of our research and potential extensions to make better use of Google Trends and other online consumer interest tracking services.

Literature

Products can often be viewed as bundles of features, with their attractiveness determined as a compensatory function of feature levels (Fishbein and Ajzen 1975; Lancaster 1966). Such a basic view of product attractiveness lies at the foundation of conjoint analyses (e.g., Bradlow, Hu, and Ho 2004; Ding, Park, and Bradlow 2009; Green and Rao 1971). Similarly, in both market response modeling and individual choice modeling, researchers have often treated the baseline attractiveness of a product as a function of its feature levels (e.g., Boatwright and Nunes 2001; Fader and Hardie 1996). However, because the feature levels of existing products do not undergo major changes frequently, product baseline attractiveness is typically treated as time invariant, leaving mainly time-varying marketing efforts to explain the dynamics in product sales (e.g., Ailawadi, Lehmann, and Neslin 2001; Danaher, Bonfrer, and Dhar 2008; Neslin 1990; Neslin, Henderson, and Quelch 1985).

Treating product baseline attractiveness as time invariant is plausible when the window of observation is short and the weights consumers place on different product features remain largely constant. However, consumer tastes can change substantially over time, often manifesting in shifts in the relative importance of various product features. These shifts in turn can lead to changes in the relative attractiveness and, thus, sales of products with different feature levels.

To allow for such possibilities, a few previous studies have considered time-varying baseline attractiveness. For example, Blattberg and Levin (1987) include both seasonality and a time trend in baseline sales. Kopalle, Mela, and Marsh (1999) allow the baseline attractiveness of a product to evolve over time stochastically. However, although these approaches can accommodate time-varying baseline attractiveness, they are limiting in that they cannot help marketers understand—beyond what marketing efforts can account for—why certain products become more or less attractive over time.

In this study, we propose a market response model in which (1) a product's baseline attractiveness is treated as a compensatory function of its feature levels and (2) the importance weight of each feature is allowed to vary over time. By tapping into Google Trends and treating feature search trends as potential indicators of feature importance trends, we investigate (1) the extent to which feature search trends can help predict product sales, beyond the effects of marketing efforts, and (2) whether feature search trends relate to product sales in a manner that is consistent with the hypothesis that feature search and importance trends are positively correlated.

In theory, as consumers rely increasingly on Internet search engines such as Google in their acquisition of product information (J.D. Power and Associates 2008, 2012), the search intensity for terms related to a particular feature should evolve over time with the weight that consumers place on that feature (Bronnenberg, Kim, and Mela 2014; Ratchford, Lee, and Talukdar 2003; Ratchford, Talukdar, and Lee 2007). In other words, all else being equal, when a feature gains in importance, one should expect, among other things, that consumers will seek more information about it online as well as offline. Indeed, our literature review has led us to a burgeoning area of research that leverages search trends extracted from Google Trends as signs of shifting consumer interests.

For example, in epidemiology, Ginsberg et al. (2009) and Pelat et al. (2009) show that search trends for disease-related terms can serve as real-time indicators of disease incidence rates and are cheaper and faster than tracking measures collected through conventional epidemic surveillance methods. In macroeconomics, research has shown that search trends can improve forecasts of housing market price and sales (Wu and Brynjolfsson 2009), unemployment rates (Askatas and Zimmermann 2009; Choi and Varian 2009a), and household expenditures (Vosen and Schmidt 2011). In finance, Da, Engelberg, and Gao (2011) show that search trends for stock tickers can help predict stock prices.

More relevant to marketing, Choi and Varian (2009b) demonstrate that search trends can help predict demand in various industries (e.g., retailing, automotive, housing, tourism). Du and Kamakura (2012) show that seven common trends extracted from Google search data for 38 major vehicle brands can explain 74% of brand-level new car sales in the United States. Hu, Du, and Damangir (2014) demonstrate how search trends can be combined with sales data to decompose advertising's overall impact into its impacts on generating consumer interest in prepurchase information search and converting that interest into sales. Taken together, these studies suggest that there can be strong ties between what consumers search online and what they purchase.

A common aspect of the aforementioned studies is that they have all focused on search terms that are directly tied to the subject of study. For example, in relating online searches to vehicle sales, Choi and Varian (2009b), Du and Kamakura (2012), and Hu, Du, and Damangir (2014) focus on the linkage between searches for a vehicle brand name (e.g., "Prius") and the sales of that vehicle (hereinafter, we refer to this type of search as "brand search"). In this study,

we extend beyond brand search by also including feature search (e.g., “fuel economy” in addition to “Prius”). We argue that such an extension constitutes a substantial contribution in the following ways.

Conceptually, our extension is consistent with the notion that consumers often engage in both brand search and feature search when they gather product-related information, especially for expensive durables such as automobiles (Bronnenberg, Kim, and Mela 2014; Ratchford, Lee, and Talukdar 2003; Ratchford, Talukdar, and Lee 2007). It is highly conceivable that searches for brand names send different signals about consumer interests than searches for product features. To the best of our knowledge, although Google Trends has been freely available since 2008, no existing academic studies have systematically examined feature search trends and how they relate to sales.

More importantly, although brand search trends may reveal which brands are gaining or losing consumer interest, they do not indicate why brand popularity is shifting. In contrast, feature search trends can reveal at a more fundamental level how the underlying preference structure may have evolved. This is akin to conjoint studies in that although it is important to get a good handle on the part-worths of brand names, it is even more useful to quantify the relative importance of different features because the resulting insights into consumer utility trade-offs are often more actionable than brand preferences alone. In the marketing literature, many studies have made important contributions by extending models of consumer demand from brands only to brands plus features (e.g., Boatwright and Nunes 2001; Fader and Hardie 1996).

That said, a common challenge in utilizing both brand and feature search trends is that uncertainties and/or ambiguities exist in the causes of consumer searches for a particular brand or feature. Indeed, this challenge is intrinsic to any study that attempts to leverage Google Trends data because this so-called database of intentions is bound to be noisy, making it difficult to establish a direct link between search trends and trends in consumers’ collective interests (Battelle 2005). Consequently, one of this study’s intended contributions is to construct a market response model that would enable us to test whether feature search trends are positively correlated with feature importance trends—a task made more challenging by the issue that feature importance trends are not readily observable. We argue that such market response model-based hypothesis testing adds rigor to emerging literature that hinges on whether online consumer tracking can provide reliable indicators of genuine shifts in consumer interest.

Finally, we explicitly control for the impacts of own and competitive marketing efforts. We show that, after controlling for marketing efforts, including only brand search trends in a market response model does not improve model fit (only marginal improvement in sample and no improvement out of sample). In contrast, including feature search trends improves model fit substantially both in and out of sample. This contrast highlights the importance of both extending beyond brand searches and controlling for marketing efforts, which has often been missing in studies that

have used search data to improve sales forecasts (e.g., Choi and Varian 2009b; Du and Kamakura 2012).

Model

Consider N competing products, each of which is characterized with a brand name and a set of features. Among the features, some are time invariant (e.g., vehicle body type), whereas the others are time varying (e.g., vehicle fuel economy). Given the empirical context in which we apply our model (i.e., the U.S. automotive market) and for ease of exposition, we assume that the time-invariant features are nominal and can be coded as L dummy variables (e.g., 1 for SUV and 0 for non-SUV) and that the K time-varying features are continuous and can be log-transformed (e.g., miles per gallon). If needed, the model we present next can be readily extended to accommodate time-invariant features that cannot be dummy-coded or time-varying features that cannot be log-transformed.

We treat the sales of product i ($i = 1, \dots, N$) in month t ($t = 1, \dots, T$), y_{it} , as a function of baseline attractiveness α_{it} (specified in Equation 2), own and competitive marketing efforts x_{ijt} ($j = 1, \dots, J$), and lagged sales $y_{i,t-1}$ and $y_{i,t-12}$. We adopt a log-log formulation that has proved to be robust in modeling sales responses (e.g., Christen et al. 1997; Danaher, Bonfrer, and Dhar 2008; Kopalle, Mela, and Marsh 1999; Wittink 1977; for a review of market response models, see Hanssens, Parsons, and Schultz 2001):

$$(1) \quad \ln(y_{it}) = \alpha_{it} + \sum_{j=1}^J \beta_{ij}^x \ln(x_{ijt}) + \rho_{i1} \ln(y_{i,t-1}) + \rho_{i2} \ln(y_{i,t-12}) + \varepsilon_{it},$$

where β_{ij}^x captures the impact of marketing effort j ; ρ_{i1} and ρ_{i2} capture, respectively, carryover and seasonal effects; and ε_{it} captures unobserved factors affecting product i ’s sales, which is assumed to be distributed i.i.d. normal with mean 0 and standard deviation σ_{ε_i} .

As we discussed previously, most market response models treat the baseline attractiveness α_{it} as a nuisance, assumed to be either time invariant or purely stochastic. The focus has been on relating observed sales dynamics to own and competitive marketing efforts. In contrast, in this study we focus on the dynamics of α_{it} and treat it as a compensatory function of product brand name searches b_{it} and feature levels $d_{i\ell}$ and f_{ikt} (Fader and Hardie 1996; Fishbein and Ajzen 1975; Lancaster 1966):

$$(2) \quad \alpha_{it} = \beta_i^0 + \beta_i^b \ln(b_{it}) + \sum_{\ell=1}^L \beta_{i\ell}^d d_{i\ell} + \sum_{k=1}^K \beta_{ikt}^f \ln(f_{ikt}),$$

where β_i^0 is the intercept; β_i^b captures the extent to which the search trend for product i ’s brand name b_{it} is tied to product i ’s attractiveness over time; $\beta_{i\ell}^d$ captures the weight consumers place on the ℓ th dummy-coded time-invariant feature $d_{i\ell}$ in determining product i ’s attractiveness in period t ; and, similarly, β_{ikt}^f captures the weight placed on $\ln(f_{ikt})$, the k th log-transformed continuous feature of product i in period t . Substituting Equation 2 into Equation 1, we have the following:

$$(3) \quad \ln(y_{it}) = \beta_i^0 + \beta_i^b \ln(b_{it}) + \sum_{\ell=1}^L \beta_{i\ell t}^d d_{i\ell} + \sum_{k=1}^K \beta_{ikt}^f \ln(f_{ikt}) \\ + \sum_{j=1}^J \beta_{ij}^x \ln(x_{ijt}) + \rho_{i1} \ln(y_{i,t-1}) + \rho_{i2} \ln(y_{i,t-12}) + \varepsilon_{it}, \\ \varepsilon_{it} \sim N(0, \sigma_{\varepsilon_i}^2).$$

Intuitively, we expect β_i^b to be positive because an upward trend in product brand name search should, on balance, indicate an upward trend in product attractiveness. To complete our model formulation, we allow $\beta_{i\ell t}^d$ and β_{ikt}^f , the feature importance weights, to vary as a function of $z_{\ell t}^d$ and z_{kt}^f , the intensity of consumer searches for the corresponding features in period t :

$$(4a) \quad \beta_{i\ell t}^d = \gamma_{\ell}^d \ln(z_{\ell t}^d) + \zeta_{i\ell t}^d, \quad \ell = 1, \dots, L, \text{ and}$$

$$(4b) \quad \beta_{ikt}^f = \gamma_k^f \ln(z_{kt}^f) + \zeta_{ikt}^f, \quad k = 1, \dots, K,$$

where γ_{ℓ}^d and γ_k^f capture the extent to which feature importance trends ($\beta_{i\ell t}^d$ and β_{ikt}^f) are correlated with the corresponding feature search trends ($z_{\ell t}^d$ and z_{kt}^f) and $\zeta_{i\ell t}^d$ and ζ_{ikt}^f are assumed to be distributed i.i.d. normal with, respectively, means $\bar{\zeta}_{\ell}^d$ and $\bar{\zeta}_k^f$ and standard deviations $\sigma_{\zeta_{\ell}^d}$ and $\sigma_{\zeta_k^f}$.

Finally, to enable information pooling across products in model estimation, we assume the product-specific effects of brand search trends (β_i^b), marketing efforts (β_{ij}^x), and sales carryover and seasonality (ρ_{i1} , ρ_{i2}) to follow normal distributions across products: $\beta_i^b \sim N(\bar{\beta}^b, \sigma_{\beta^b}^2)$; $\beta_{ij}^x \sim N(\bar{\beta}_{ij}^x, \sigma_{\beta_{ij}^x}^2)$, for $j = 1, \dots, J$; $\rho_{i1} \sim N(\bar{\rho}_1, \sigma_{\rho_1}^2)$; $\rho_{i2} \sim N(\bar{\rho}_2, \sigma_{\rho_2}^2)$. The parameters $\bar{\beta}^b$, $\bar{\beta}_{ij}^x$, $\bar{\rho}_1$, and $\bar{\rho}_2$ are the means and $\sigma_{\beta^b}^2$, $\sigma_{\beta_{ij}^x}^2$, $\sigma_{\rho_1}^2$, and $\sigma_{\rho_2}^2$ are the variances of the hierarchical distributions.

Testing Whether Feature Search and Importance Trends Are Positively Correlated

Analogous to conjoint analysis, β_{ikt}^f in Equation 3 can be viewed as the importance weight of feature k during period t in determining product i 's attractiveness. Alternatively, in the nomenclature of market response models, β_{ikt}^f can be interpreted as product i 's sales elasticity to feature k level (i.e., $\partial \ln(y_{it}) / \partial \ln(f_{ikt}) = \beta_{ikt}^f$). That is, for a 1% increase in feature k level, one would expect a $\beta_{ikt}^f\%$ shift in product i 's sales, which should be positive for desirable features and negative for undesirable features. Furthermore, β_{ikt}^f is allowed to vary over time as a function of z_{kt}^f , the intensity of consumer searches for feature k at time t . If an upward trend in z_{kt}^f is indeed a manifestation of feature k becoming more important, we expect γ_k^f in Equation 4b to be positive for a desirable feature and negative for an undesirable feature.

To observe the logic more intuitively, consider fuel economy, a desirable vehicle feature, as an example. All else being equal, when fuel economy gains importance in shaping vehicle purchase decisions, two upward trends emerge: (1) consumers conduct more searches containing fuel economy-related keywords (i.e., an upward trend in z_{kt}^f) and (2) a vehicle's sales become more positively elastic to its fuel economy level (i.e., an upward trend in β_{ikt}^f). Taken together, as the importance of fuel economy evolves over time, one should expect to observe a positive correla-

tion between β_{ikt}^f and z_{kt}^f , which would be captured through a positive γ_k^f .

Similarly, this logic also applies to undesirable features. For example, all else being equal, when cost to operate becomes more of a concern to consumers, one would expect people to conduct more searches containing keywords related to cost to operate (i.e., an upward trend in z_{kt}^f). In the meantime, a vehicle's sales would be expected to become more negatively elastic to its cost to operate (i.e., a downward trend in β_{ikt}^f). Taken together, as the importance of cost to operate evolves over time, a negative correlation between β_{ikt}^f and z_{kt}^f is expected, which would be captured through a negative γ_k^f .

In summary, on the basis of one first principle (i.e., when a desirable [undesirable] feature becomes more important to consumers, sales should become more positively [negatively] elastic to the level of that feature) and one hypothesis (i.e., when a feature becomes more important, one of its manifestations will be an upward trend in searches for keywords related to that feature), we conjecture that the γ_k^f estimates in our model should be positive for desirable features (e.g., fuel economy, acceleration) and negative for undesirable features (e.g., cost to buy, cost to operate). If this conjecture is borne out empirically, it should lend strong face validity to our proposed model and the treatment of feature search trends as reflective indicators of feature importance trends, which are latent constructs that would be difficult to track otherwise.

Projecting Search Trends into the Future

Although feature search data are readily available through Google Trends in near real time, monitoring historical feature search trends can only inform managers about how feature importance weights have evolved in the past (up to the point when the latest data are gathered). To make our model more useful in helping managers become proactive in their leveraging of feature search data, we present a model that projects feature search trends into the near future (e.g., the next one to two years). Equipped with forecasted feature search trends and our proposed market response model, managers can improve not only sales forecasts but also planning for the future.

We use a basic structural time-series model with a local linear trend and a seasonal component to capture the dynamics in feature search trends (Commandeur and Koopman 2007). Specifically, we model the search intensity for feature k during period t , z_{kt} , as follows:

$$(5) \quad \ln(z_{kt}) = \phi_{kt}^z + \psi_{kt}^z + v_{kt}^z, \quad v_{kt}^z \sim \text{i.i.d.}N(0, \sigma_{v_k^z}^2) \\ \phi_{kt}^z = \phi_{k,t-1}^z + \Delta_{k,t-1}^z + w_{kt}^z, \quad w_{kt}^z \sim \text{i.i.d.}N(0, \sigma_{w_k^z}^2) \\ \Delta_{kt}^z = \Delta_{k,t-1}^z + h_{kt}^z, \quad h_{kt}^z \sim \text{i.i.d.}N(0, \sigma_{h_k^z}^2) \\ \psi_{kt}^z = - \sum_{m=1}^{11} \psi_{k,t-m}^z + u_{kt}^z, \quad u_{kt}^z \sim \text{i.i.d.}N(0, \sigma_{u_k^z}^2),$$

where ϕ_{kt}^z represents the local linear trend and ψ_{kt}^z is the seasonal component.

We use a standard procedure in SAS (Proc UCM) to calibrate this model and to make projections in the given forecasting window (e.g., 24 months in holdout validation). We apply the same model setup and estimation procedure to project brand search trends (b_{it}).

Benchmark Models

We have proposed a novel market response model in which feature search trends enter into the system as predictors of feature importance weights, which means that z_{kt}^f affect sales only through interactions with feature levels. We have argued that the signs and significance of γ_k^f will provide a strong test of our hypothesis that feature search trends are positively correlated with feature importance trends. However, another important empirical question remains: What is the incremental predictive power of feature search trends, beyond what has typically been included in existing market response models? To address this issue, we benchmark our model against two alternatives and evaluate their performances in in- and out-of-sample goodness-of-fit.

Equation 6 presents the first benchmark, a seasonal autoregressive model with marketing efforts (x_{ijt}) as additional covariates. This model assumes that the baseline attractiveness of product i (β_i^0) remains constant over time, which leaves marketing efforts as the sole driver of the observed dynamics in sales:

$$(6) \quad \ln(y_{it}) = \beta_i^0 + \sum_{j=1}^J \beta_{ij}^x \ln(x_{ijt}) + \rho_{i1} \ln(y_{i,t-1}) + \rho_{i2} \ln(y_{i,t-12}) + \varepsilon_{it}.$$

Equation 7 presents the second benchmark, which extends Equation 6 by including product brand name search b_{it} as an additional covariate:

$$(7) \quad \ln(y_{it}) = \beta_i^0 + \beta_i^b \ln(b_{it}) + \sum_{j=1}^J \beta_{ij}^x \ln(x_{ijt}) + \rho_{i1} \ln(y_{i,t-1}) + \rho_{i2} \ln(y_{i,t-12}) + \varepsilon_{it}.$$

The second benchmark is consistent with the notion that trends in product brand name searches can be indicative of trends in product sales, which has been the predominant way Google Trends data have been leveraged (e.g., using online search volumes for the keyword “Prius” as a predictor of current or future Prius sales). This is essentially what Choi and Varian (2009b) and Du and Kamakura (2012) find (even though, unlike Equations 6 and 7, neither study controlled for the impacts of marketing efforts).

In summary, by comparing our proposed model against the two benchmarks, we can determine the incremental value of brand search trends over marketing efforts and lagged sales (benchmark 2 [Equation 7] vs. benchmark 1 [Equation 6]). We can also determine the incremental value of feature search trends over brand search trends, marketing

efforts, and lagged sales (our proposed model [Equations 3 and 4] vs. benchmark 2 [Equation 7]).

Endogeneity and Model Calibration

We aim to achieve two goals through our empirical analyses. We hope to (1) establish face validity by examining the signs and significance of model parameter estimates to determine whether they are mostly consistent with our expectations and (2) establish predictive validity by comparing in- and out-of-sample fit between the alternative models. To achieve the first goal, we recognize that, for both our proposed model and the benchmarks, endogeneity could bias the parameter estimates because there might be (1) unobserved factors influencing both the dependent variable (sales) and the covariates and (2) reverse causality between sales and marketing efforts or between sales and search trends (Joo et al. 2013).

Following Lancaster (2004) and Van Heerde et al. (2013), we adopt a two-stage least squares approach in dealing with potential endogeneity in marketing efforts and search trends. Following Van Heerde et al., we use the following instrumental variables (IVs) for a focal vehicle’s marketing efforts and brand search trend: average lagged cash incentive, average lagged advertising spend, and average lagged brand search, with all the averaging taken over vehicles *outside* the class to which the focal vehicle belongs. For example, when Honda Civic is the focal vehicle, we use as IVs the averages of lagged cash incentives, advertising spend, and brand searches from vehicles that are not in the compact sedan class (e.g., SUVs, full-size sedans). Lamey et al. (2012) and Ma et al. (2011) adopted similar IVs (lagged marketing activities outside the focal product class). To address potential endogeneity in feature search trends, which are not specific to any particular product, we use as IVs the average current feature levels (of all products) and lagged industry-wide cash incentives and advertising spend.² All the IVs we have adopted passed the Angrist–Pischke (2009) test for strength ($p < .05$) and Sargan test for overidentification ($p > .10$).

Although it is important to adjust for potential endogeneity in determining the face validity of our model (through the signs and significance of parameter estimates), endogeneity poses less of a threat in determining the predictive validity of our model. Indeed, as Ebbes, Papies, and Van Heerde (2011) demonstrate, one should not use models that are adjusted for endogeneity in holdout sample validation. Accordingly, when we attempt to establish predictive validity, we recalibrate our proposed model and the bench-

²It is important to note that feature searches, unlike brand searches, are not tied to any specific product and thus have much larger volumes in general. For example, in the United States the search volume for keywords related to vehicle fuel economy is more than five times larger than the search volume for Prius. Consequently, compared with brand searches, feature searches are less susceptible to idiosyncratic/brand-specific forces that affect only a few products and do not reflect genuine shifts in the underlying preference structure. For example, a product recall might trigger more searches for Prius, with little impact on searches for vehicle fuel economy.

marks without using any IVs and focus on comparing the models' out-of-sample fit.

We estimate all the models, endogeneity-adjusted or not, using the Gibbs sampler programmed in WinBUGS (Lunn et al. 2000) and use diffuse priors on all the parameters, with normal (0, 1,000) for the means and inverse-gamma (.01, .01) for the variances. Of the 120,000 draws from the Gibbs sampler, the first 100,000 serve as burn-ins, and the posterior estimates are based on the last 20,000 draws. We determine convergence by using the Brooks–Gelman–Rubin diagnosis in WinBUGS (Brooks and Gelman 1998). We deem an estimate statistically significant when the posterior credible interval between the 2.5th and 97.5th percentile does not contain 0.

Data

We apply our model in the context of the U.S. automotive market. In buying expensive durables such as automobiles, consumers are highly motivated to conduct information search (Beatty and Smith 1987; Moorthy, Ratchford, and Talukdar 1997; Punj and Staelin 1983; Srinivasan and Ratchford 1991), and increasingly those searches are carried out over the Internet (Bronnenberg, Kim, and Mela 2014; J.D. Power and Associates 2008, 2012; Ratchford, Lee, and Talukdar 2003; Ratchford, Talukdar, and Lee 2007; Zettelmeyer, Morton, and Silva-Risso 2006). Furthermore, the automotive market in the United States is highly differentiated, with a large number of established vehicle models that can be characterized by a set of well-defined and readily searchable features. We consider this market a suitable context for testing our proposed market response model.

We gathered monthly new vehicle sales data (y_{it}) from Automotive News for the period between January 2004 and April 2011. During the 88-month period, 80 nonluxury passenger vehicles had been continuously available in the U.S. market, all of which are included in our study. Such a selection criterion means that we had to exclude newly launched or discontinued vehicles. That said, the 80 vehicles that we do include account for more than 71% of total U.S. nonluxury passenger vehicle sales during the observation window.

In terms of each vehicle's own marketing efforts, we assembled monthly average consumer cash back per vehicle (x_{11t}) from Automotive News ("own incentives" hereinafter) and monthly total advertising spend in the United States (x_{12t}) from Kantar Media's AdSpender ("own ad spend" hereinafter). For each focal vehicle, competitive marketing efforts include monthly average consumer cash back per vehicle (x_{13t} , "competitive incentives" hereinafter) and total ad spend (x_{14t} , "competitive ad spend" hereinafter) by all the competitors in the focal vehicle's class (e.g., compact sedan, compact SUV). We identified each focal vehicle's class and all the within-class competitors using Automotive News's market classification system.

From Google Trends, we extracted 80 brand search trends (b_{it}) for the 80 vehicles in our sample, one search trend (z_{1t}^d) for vehicle body type (SUV or not, a time-invariant feature), and four search trends for time-varying features: fuel economy (z_{1t}^f), acceleration (z_{2t}^f), cost to buy (z_{3t}^f), and

cost to operate (z_{4t}^f). For fuel economy, acceleration, and cost to operate, we gathered the actual levels (f_{ikt}) of each vehicle during each model-year from Edmunds.com. For cost to buy, we use the manufacturer's suggested retail price (MSRP) gathered from MSN Autos. Just as prices are often treated as a product feature in conjoint analyses, we include cost to buy (MSRP) as a product feature rather than as a marketing effort variable. We chose cost to buy as the label (instead of sticker price) to better contrast with cost to operate, which includes costs associated with maintenance, repair, fuel, and so on. Finally, for vehicles that offer multiple trims, we used the median feature level.

We chose to focus on vehicle body type, fuel economy, acceleration, cost to buy, and cost to operate because these features are considered the most relevant to an average U.S. consumer, according to major automotive information websites such as Edmunds.com and JDPower.com. In general, for nonluxury passenger vehicles, fuel economy and acceleration are considered desirable features, whereas cost to buy and cost to operate are considered undesirable features. As we argued previously, if indeed feature search trends are positively correlated with feature importance trends, we expect a vehicle's sales to become more positively elastic to its fuel economy and acceleration levels when searches for these desirable features trend upward (i.e., γ_1^f and γ_2^f are expected to be positive and significant). Conversely, we expect a vehicle's sales to become more negatively elastic to its cost-to-buy and cost-to-operate levels when searches for these undesirable features trend upward (i.e., γ_3^f and γ_4^f are expected to be negative and significant). Finally, the body type of a vehicle (SUV or not) does not change over time and is a nominal feature that is neither universally desirable nor universally undesirable. Nevertheless, intuition suggests that, all else being equal, when searches for SUVs intensify, SUV sales should be expected to increase (i.e., γ_1^d should be positive).

A key challenge in extracting the search trend for a particular feature lies in the identification of relevant search terms. When consumers seek information related to the same underlying product feature, they may use a wide variety of terms (e.g., fuel economy vs. fuel efficiency vs. gas mileage vs. miles per gallon), including abbreviations (e.g., "MPG" for miles per gallon), minor variations, singulars/plurals, and misspellings. Fortunately, Google Trends allows users to construct composite queries by joining multiple terms with plus signs. For example, the composite query we used to extract the search trend for fuel economy includes ten terms (i.e., city mileage + fuel consumption + fuel economy + fuel efficiency + fuel efficient + gas mileage + highway mileage + hybrid + mile per gallon + mpg). Table 1 presents the composite queries we used to extract feature search trends from Google Trends.

To compile the terms used to form the composite queries, we went through a careful multistep procedure. First, we generated a comprehensive list of candidate terms so as not to miss any popular ones used by consumers. We began by scanning consumer reviews on Edmunds.com and selected terms that seemed relevant for each feature. Each term that resulted from this step was then entered into

TABLE 1
Google Trends Queries Used for Extracting Feature Search Trends

	Feature	Query
k = 1	Fuel economy	city mileage + fuel consumption + fuel economy + fuel efficiency + fuel efficient + gas mileage + highway mileage + hybrid + mile per gallon + mpg
k = 2	Acceleration	acceleration time + "acceleration times" + "0-60 time" + "0 to 60 time" + "0-60 times" + "0 to 60 times" + "quarter mile time" + "quarter mile times"
k = 3	Cost to buy	price + msrp + discount + rebate + "finance rate" + "cash back" + cashback
k = 4	Cost to operate	powertrain-warranty + reliability + reliable + "cost to own" + "maintenance cost" + "maintenance costs"
$\ell = 1$	Body type: SUV	suv

Notes: A single-term query will match all the searches containing that term. If multiple terms are used in a query, Google Trends uses a set of simple logical operations in matching searches: terms inside quotation marks will match searches having the exact phrase, terms separated by a space will match searches with all the terms, minus signs before a term/phrase will exclude searches with the corresponding term/phrase, and plus signs between terms/phrases cause the query to match searches with either of the terms/phrases.

Google AdWords for suggestions of additional keywords to further expand the list of candidate terms. Subsequently, we pared down this list by excluding terms that can be intended for things other than the focal feature (e.g., we do not use the term "acceleration" by itself because the result can be confounded with searches related to the accelerator pedal). We relied on two independent judges, and whenever disagreement arose, we used a third to make the decision regarding whether to keep or remove a term. Finally, we removed terms that have much lower search volumes than the popular ones (both Google AdWords and Google Trends can be used to determine the relative popularity for different terms). We followed the same procedure in constructing composite queries for vehicle brand names, which mainly consist of vehicle make and model, along with common variations (e.g., volkswagen beetle + vw beetle + volkswagon beetle + volkswagen beatle + vw beatle + volkswagon beatle).

To ensure that the feature search trends are indeed related to automobiles (as opposed to other product categories), we relied on the "Autos & Vehicles" category filter in Google Trends. In constructing the feature search trends used in our model, we divided each raw feature search index from Google Trends by the search index for the entire "Autos & Vehicles" category. This gives us a normalized trend line that captures the search intensity for a particular feature relative to the search intensity for the entire "Autos & Vehicles" category. Such normalization provides us with a more reliable tracking measure of how feature importance has evolved over time because it removes variations in feature searches that are due to variations in consumers' overall category search interest. For example, seasonal increases in vehicle shopping activities may increase all vehicle-related searches, but that does not mean a particular vehicle feature has become more important. By dividing raw feature search indexes by the search index for the entire "Autos & Vehicles" category, such category-wide fluctuations can be removed. Like feature search trends, brand search trends are also normalized by dividing the raw Google Trends indexes by the search index for the entire "Autos & Vehicles" category. Table 2 presents summary statistics for all the variables used in our empirical analyses, including sales, own and competitive marketing efforts, brand and feature search trends, and the actual feature levels.

Results

In- and Out-of-Sample Fit

We compare our proposed model against the two benchmarks introduced previously. The goal is to determine the incremental value of feature search trends in predicting sales, beyond what can be accomplished with marketing efforts and brand search trends. To do so, we compare the models in terms of both in- and out-of-sample fit. For in-sample comparisons, we calibrate all the models using data from the first 64 months of our observation window (i.e., January 2004 through April 2009) and compare the mean absolute error (MAE) between the actual sales $\ln(y_{it})$ and the predicted sales. The first row of Table 3 reports the overall in-sample MAEs (across 80 vehicles and 64 months).

We observe that Model 1, a seasonal autoregressive model that uses lagged sales and own and competitive marketing efforts as predictors, produces an MAE of .190. By adding brand search trends to Model 1, Model 2 improves in-sample MAE little (.190 vs. .188). In contrast, by adding feature levels and feature search trends to Model 2, our proposed Model 3 reduces MAE by 54.3% (from .188 to .086). Such a remarkable improvement in fit indicates that feature levels and feature search trends capture a substantial amount of dynamics in sales, beyond what has already been captured by lagged sales, marketing efforts, and brand search trends.³

Although a 54.3% reduction in MAE is impressive, superior in-sample fit can be misleading due to the risk of overfitting. To address this issue, we use data from the last 24 months of our observation window (i.e., May 2009 through April 2011) to conduct a holdout test. As with the in-sample comparisons, we calibrated the parameters of all

³We explored two alternative benchmarks, which extend, respectively, Models 1 and 2 by including time-varying feature levels (i.e., fuel economy, acceleration, cost to buy, and cost to operate) as additional covariates. Time-invariant features such as vehicle body type cannot be included as covariates because the models already have product-specific time-invariant intercepts. We chose not to include these two alternative benchmarks because they underperformed their simpler counterparts in out-of-sample fit.

TABLE 2
Summary Statistics

Variable	M	SD	Min	Mdn	Max
Sales	8.568	1.129	4.615	8.668	11.751
Feature Search Trends					
Fuel economy	3.314	.365	2.750	3.206	4.470
Acceleration	3.841	.237	3.309	3.825	4.284
Cost to buy	4.108	.118	3.868	4.102	4.340
Cost to operate	4.205	.110	3.920	4.206	4.439
Body type: SUV	4.042	.157	3.631	4.028	4.389
Feature Levels					
Fuel economy	3.176	.217	2.680	3.179	3.882
Acceleration	5.344	.321	4.631	5.347	6.080
Cost to buy	10.002	.342	9.138	9.976	11.122
Cost to operate	10.094	.548	8.375	9.984	11.151
Body type: SUV	.537	—	.000	—	1.000
Brand search trends	3.712	.639	.000	3.868	4.567
Own incentive	2.871	3.641	.000	.000	9.521
Own ad spend	6.745	2.065	.000	7.066	10.958
Competitive incentive	5.395	2.682	.000	6.459	8.604
Competitive ad spend	10.283	1.062	6.585	10.592	12.088

Notes: All variables are log-transformed except Body type: SUV, which is dummy-coded with 1 for SUV and 0 otherwise.

TABLE 3
Model Comparison

Overall Goodness-of-Fit	Benchmark Models		Proposed Model	% Change
	Model 1	Model 2	Model 3	Model 3 Versus Model 2
MAE (in-sample) ^a	.190	.188	.086	-54.3%
MAE (out-of-sample) ^b	.370	.370	.331	-10.3%

^aBased on the first 64 months of data.

^bBased on the last 24 months of data.

the models using data from the first 64 months. In the out-of-sample comparisons, instead of using the actual brand and feature search trends during the 24-month holdout period, we use their forecasts.

More specifically, for each brand search trend and feature search trend, we calibrate a separate forecasting model (Equation 5) using data from the first 64 months, which is then projected into the 24-month holdout period. These forecasted search trends are then plugged into the market response models calibrated using data from the first 64 months, producing sales forecasts for the 24-month holdout period. Note that, for marketing efforts and feature levels, we use their actual values in the holdout period. We do so for three reasons. First, they are decision variables that are under managerial control. Second, using the actual marketing efforts and feature levels and the forecasted brand and feature search trends in the holdout period allows us to focus on the incremental value of search trends in predicting sales, which is a main goal of this study. Third, the way we carry out the holdout test mimics the common practice in sales forecasting, in which the main source of uncertainty, from the manager's perspective, lies in future consumer preferences, the evolution of which manifests in search trends for not only product brand names but also product features.

The second row of Table 3 reports the overall MAEs (across 80 vehicles and 24 months) based on the holdout

test. First, we note that the out-of-sample MAEs are larger than their in-sample counterparts. This is not surprising given that the models are calibrated using data from the first 64 months, which can contain idiosyncrasies that do not exist in the 24-month holdout period. For Models 2 and 3, the increases in MAEs from in sample to out of sample also reflect errors in forecasted brand and feature search trends.

Second, we note that the out-of-sample MAEs are the same between Models 1 and 2, which suggests that adding brand search trends to Model 1 failed to improve predictive performance. This finding runs counter to Choi and Varian (2009b) and Du and Kamakura (2012), who find that trends in online searches for vehicle brand names could improve forecasts of vehicle sales. A key difference between our study and theirs is that we controlled for own and competitive marketing efforts, which raises the possibility that, after controlling for marketing efforts, including only brand search trends in a market response model does not improve model fit (in our case, marginal in-sample improvement and no out-of-sample improvement).

Third, and most importantly, by comparing Model 3 with Model 2, we note that by adding feature search trends, our proposed model reduces the out-of-sample MAE by 10.3% (from .370 to .331), an empirical result that should be considered impressive. Taken together, such improvements in fit (54.3% in sample and 10.3% out of sample) indicate that vehicle feature search trends can explain a

large portion of dynamics in vehicle sales, beyond what can be accounted for by lagged sales, marketing efforts, and brand search trends.

Furthermore, the incremental predictive power of our proposed model does not depend on the amount of lead time between feature searches and product purchases. Rather, the incremental predictive power results from the notion that feature importance—and thus, the elasticity of sales to feature levels—can change substantially over time; as a result, as long as feature search trends and feature importance trends are correlated over the long run, our proposed model would be able to leverage the signals about consumer preferences contained in feature search trends and capture sales dynamics that cannot be foreseen otherwise.

Estimates of γ_k^f

Although the improvements in in- and out-of-sample fit suggest strong predictive validity, to further investigate whether the improvements in fit are a statistical coincidence, we examine the estimates of the model parameters that link feature search trends to sales dynamics (i.e., the γ_k^f). Recall that product i 's sales elasticity to time-varying feature k level is captured through $\beta_{ikt}^f = \partial \ln(y_{it}) / \partial \ln(f_{ikt})$ (all else being equal, for a 1% increase in the feature's level, one would expect a $\beta_{ikt}^f\%$ change in product i 's sales). Our proposed model allows β_{ikt}^f to vary as a function of feature search trend z_{kt}^f ($\partial \beta_{ikt}^f / \partial \ln(z_{kt}^f) = \gamma_k^f$). Simply put, γ_k^f captures how the search trend for feature k moderates the sales elasticity of feature k .

If our model's superior performance in capturing sales dynamics is not a statistical fluke but rather is due to the finding that feature search trends contain genuine information about feature importance trends, an intuitive pattern in the γ_k^f estimates would be expected. Table 4 indicates that the expected pattern has been borne out: the γ_k^f estimates for the two desirable features, fuel economy and acceleration, are both positive and significant ($\gamma_1^f = .019$ and $\gamma_2^f = .106$, $p < .05$), while the γ_k^f estimates for the two undesirable fea-

tures, cost to buy and cost to operate, are both negative and significant ($\gamma_3^f = -.050$ and $\gamma_4^f = .103$, $p < .05$). These results mean that (1) all else being equal, when consumer searches for keywords related to fuel economy and acceleration intensify, vehicle sales become more positively elastic to those features and (2) when consumer searches for keywords related to cost to buy and cost to operate intensify, vehicle sales become more negatively elastic to those features.

Taken together, the finding that all the γ_k^f estimates for the desirable and undesirable features are significant and of the expected signs suggests that our proposed model's superior performance in capturing sales dynamics is by no means a statistical coincidence. It also lends strong empirical support to the notion that the evolution of feature search intensity contains genuine information about shifting consumer preferences and that feature search trends can be tracked as a way to monitor feature importance trends (at least in our empirical context).

Estimates of Other Model Parameters

In addition to the γ_k^f estimates, Table 4 also reports the estimates for the mean effects of brand search trends ($\bar{\beta}^b$), carryover and seasonality ($\bar{\rho}_1$ and $\bar{\rho}_2$), and own and competitive marketing efforts ($\bar{\beta}_1^x$, $\bar{\beta}_2^x$, $\bar{\beta}_3^x$, and $\bar{\beta}_4^x$). According to Model 3, vehicle brand search trends are positively linked with vehicle sales ($\bar{\beta}^b = .138$, $p < .05$). The carryover and seasonality effects are both positive and significant ($\bar{\rho}_1 = .480$ and $\bar{\rho}_2 = .243$, $p < .05$). The effects of own and competitive incentives are both insignificant ($\bar{\beta}_1^x = -.002$ and $\bar{\beta}_3^x = .012$, $p > .05$), while the effects of own and competitive ad spend are both positive and significant ($\bar{\beta}_2^x = .035$ and $\bar{\beta}_4^x = .015$, $p < .05$).

Concluding Remarks

A basic idea behind conjoint analysis is that the attractiveness of a product can be modeled as a function of feature levels and the importance weights that consumers place on

TABLE 4
Parameter Estimates

Variable	Parameter	Benchmark Models		Proposed Model
		Model 1	Model 2	Model 3
Feature Search Trends				
Fuel economy	γ_1^f			.019
Acceleration	γ_2^f			.106
Cost to buy	γ_3^f			-.050
Cost to operate	γ_4^f			-.103
Body type: SUV	γ_1^d			.002
Brand search trends	$\bar{\beta}^b$ ($\sigma_{\beta^b}^2$)		.122 (.021)	.138 (.007)
Sales (1 month lagged)	$\bar{\rho}_1$ ($\sigma_{\rho_1}^2$)	.558 (.013)	.529 (.015)	.480 (.005)
Sales (12 months lagged)	$\bar{\rho}_2$ ($\sigma_{\rho_2}^2$)	.244 (.013)	.239 (.013)	.243 (.006)
Own incentive	$\bar{\beta}_1^x$ ($\sigma_{\beta_1^x}^2$)	-.014 (.004)	-.017 (.004)	-.002 (.001)
Own ad spend	$\bar{\beta}_2^x$ ($\sigma_{\beta_2^x}^2$)	.041 (.006)	.039 (.006)	.035 (.002)
Competitive incentive	$\bar{\beta}_3^x$ ($\sigma_{\beta_3^x}^2$)	.010 (.007)	-.0005 (.007)	.012 (.002)
Competitive ad spend	$\bar{\beta}_4^x$ ($\sigma_{\beta_4^x}^2$)	.037 (.005)	.036 (.005)	.015 (.001)

Notes: Boldface indicates that the 95% posterior credible interval excludes zero.

them. In this study, we extend this basic idea to market response modeling, allowing product baseline attractiveness to vary as a function of feature levels and feature importance weights, which we allow to vary over time as a function of feature search trends. By applying our proposed model to the U.S. automotive market, we find empirical support for the notion that feature search trends, which are readily available in near real time, may be treated as reflective indicators of feature importance trends, which are latent constructs that would otherwise be difficult to track. Furthermore, we show that, even after accounting for marketing efforts and brand search trends, the predictive power of a market response model can be substantially improved by allowing product baseline attractiveness to vary over time as a function of feature search trends.

We note that the managerial relevance of our proposed approach goes beyond sales forecasting, although that can be a worthy goal in and of itself. First and foremost, managers should recognize that, as our empirical results have shown, feature search and importance trends tend to be positively correlated. When increasing (decreasing) consumer searches for certain product features have been spotted, managers should infer with reasonable confidence that those features are likely gaining (losing) importance in shaping consumers' purchase decisions. Second, given those trends in feature searches (and thus, feature importance), managers should reevaluate the baseline attractiveness of each product that is under their purview, which should reveal that some products will become more attractive, while others will become less so. Finally, equipped with those foresights, we identify the following avenues through which managers may improve their decision making by leveraging feature search trends strategically.

- *Product design.* Managers can continuously monitor consumer tastes through online search trends and enhance their product features to meet changing needs. In particular, managers can put more effort into developing products that are superior in features that attract increasingly more consumer searches. When considering adjusting certain product features, managers can combine our proposed market response model with projected feature search trends to conduct what-if analyses and quantify the impacts of feature level changes on future sales.
- *Budget allocation.* Managers can adjust their budget allocation across a product portfolio on the basis of how each product's baseline attractiveness is projected to change as a result of shifting feature search trends. More specifically, to maximize profits for a given budget, managers should consider increasing marketing spend (e.g., advertising, incentives) on products that are becoming more attractive because they will generate more "bang for the buck" (i.e., the same sales lift is applied to an increased baseline). Or, if the goal is to meet sales targets while minimizing total marketing spend, managers should consider decreasing marketing spend on products that are becoming more attractive because they are more likely to "sell themselves." In short, without paying close attention to feature search trends, managers risk over- or underspending on products whose baseline attractiveness is shifting.
- *Advertising.* Managers may dynamically adjust their ad copies and messaging by increasing emphasis on features for which their products are considered superior and consumer searches are trending upward.

- *Production and inventory planning.* Faced with increasing or decreasing consumer searches for certain features, managers need to know how demands for different products would change as a result. They may consider expanding production and/or inventory for products that are superior in features whose searches are trending upward.

To implement our proposed approach in practice, we note the following challenges and limitations that must be put into proper perspective:

- *Ambiguity of consumer search intentions.* Although we have argued and provided empirical evidence in support of the notion that feature search and importance trends are positively correlated, we acknowledge that it would be difficult to prove that increasing feature search intensity can always be viewed as a sign of increasing feature importance. Such ambiguity exists because there will always be uncertainties about what might have caused consumers to conduct more or fewer online searches on a particular feature (e.g., it could be due to non-shopping-related reasons or to a change in information availability from offline sources). Indeed, the challenge we face here is intrinsic to any study that attempts to leverage Google Trends data, which is bound to be more ambiguous than transaction or survey data. We took an empiricist approach in this study and relied on predictive and face validity tests to gain confidence in our results. In contexts in which such tests may not be viable, researchers must be cautious in interpreting what a particular search trend stands for.
- *Model stability.* Model stability is related to the previous challenge. For example, the now-famous Google Flu Trends tool must also contend with the fact that not all searches for flu-related terms are tied to flu incidences and that trends in flu searches are not perfectly correlated with trends in flu incidences. Although such a lack of perfect correlation did not prevent Google Flu Trends from becoming a useful tool in helping detect influenza epidemics (Ginsberg et al. 2009), recent studies (e.g., Butler 2013; Lazer et al. 2014) have found that predictions can go seriously wrong if the forecasting model is not dynamically recalibrated. We share the concern and emphasize that one should treat all models that leverage Google Trends data with caution because (1) the data-generating process could change over time and (2) the dynamic relationship between search trends and the target variables could also change over time. In other words, what works now might not work as well in the future. Consequently, it is important to periodically recalibrate models (e.g., update the keyword list and coefficient estimates) and pay close attention when systematic biases begin to emerge in model predictions.
- *Subjectivity in the construction of feature search indexes.* A challenge in extracting the search trend for a particular feature lies in the identification of relevant search terms. In our study, we first generated a comprehensive list of candidate terms so as not to miss any popular ones used by consumers. We then pared down the list by removing terms that could be intended for things other than the focal feature or terms that have much lower search volumes. Although we adopted a systematic procedure in constructing the composite query for each search trend, subjectivity in such a process is inevitable. In future studies, one possibility is to consider employing multiple independent teams in the construction of feature search indexes, resolving any between-index discrepancies before proceeding to subsequent analyses.

These challenges aside, we see many promising avenues for extending our work, including, but not limited to, the following:

- *Studying other product categories.* As the examples in Figure 1 show, the application of our approach can go beyond the automotive market. It should be readily extendable to categories in which products can be characterized by a set of well-defined features and for which a significant number of consumers search for feature-related information online.
- *Extending the approach to explain spatial-temporal variation in consumer demands.* It is conceivable that relative feature importance can differ across geographic regions and its evolution can follow different paths. Because Google Trends provides search indexes at both national and regional levels, marketers can potentially monitor feature search trends by geographic region and tailor their product offerings and marketing efforts to the distinct and evolving tastes of each local market.
- *Investigating drivers of feature importance trends.* In this research, we treat feature search trends as reflective indicators of feature importance trends. However, we do not address what may cause the relative importance of various product features to evolve over time. It would be worthwhile to investigate the extent to which feature importance trends are shaped by factors that are under marketers' control (vs. environmental factors that are beyond marketers' control). Given that feature importance trends are difficult to track, future studies can treat feature search trends as their manifestations and examine the latter's drivers.

- *Using data from other online consumer interest-tracking services.* Although Google Trends is the best known and most readily available service, it captures only online consumer searches. It is conceivable that platforms with user-generated content such as Facebook, Twitter, blogs, and product review sites may share some of our findings and yield new insights. With these other social media platforms, researchers could monitor not only the frequency at which a particular feature is mentioned over time but also the context in which it is mentioned (e.g., the valence, the brands most often associated with it, co-mentions of different features).

In conclusion, we intend this research to serve as a call to action for marketers to better account for the impacts of evolving consumer tastes in market response modeling. Thanks to the emergence of big data sources such as Google Trends, it will become increasingly cost effective for marketers to monitor the evolution of consumer tastes. We hope this article serves as a prototype empirical example, illustrating what could be achieved by leveraging online consumer interest tracking data in analyses of market performance dynamics, an area of research that is bound to grow as marketers continue to explore this promising new source of marketing intelligence.

REFERENCES

- Ailawadi, Kusum L., Donald R. Lehmann, and Scott A. Neslin (2001), "Market Response to a Major Policy Change in the Marketing Mix: Learning from Procter & Gamble's Value Pricing Strategy," *Journal of Marketing*, 65 (January), 44–61.
- Angrist, Joshua D. and Jörn-Steffen Pischke (2009), *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.
- Askitas, Nikolaos and Klaus F. Zimmermann (2009), "Google Econometrics and Unemployment Forecasting," *Applied Economics Quarterly*, 55 (2), 107–120.
- Battelle, John (2005), *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*. New York: Portfolio.
- Beatty, Sharon E. and Scott M. Smith (1987), "External Search Effort: An Investigation Across Several Product Categories," *Journal of Consumer Research*, 14 (1), 83–95.
- Blattberg, Robert C. and Alan Levin (1987), "Modelling the Effectiveness and Profitability of Trade Promotions," *Marketing Science*, 6 (2), 124–46.
- Boatwright, Peter and Joseph C. Nunes (2001), "Reducing Assortment: An Attribute-Based Approach," *Journal of Marketing*, 65 (July), 50–63.
- Bradlow, Eric T., Ye Hu, and Teck-Hua Ho (2004), "A Learning-Based Model for Imputing Missing Levels in Partial Conjoint Profiles," *Journal of Marketing Research*, 41 (November), 369–81.
- Bronnenberg, Bart J., Jun B. Kim, and Carl F. Mela (2014), "Zooming In on Choice: How Do Consumers Search for Cameras Online?" working paper, Tilburg University.
- Brooks, Stephen P. and Andrew Gelman (1998), "General Methods for Monitoring Convergence of Iterative Simulations," *Journal of Computational and Graphical Statistics*, 7 (4), 434–55.
- Butler, Declan (2013), "When Google Got Flu Wrong," *Nature*, 494 (7436), 155–56.
- Choi, Hyunyoung and Hal Varian (2009a), "Predicting Initial Claims for Unemployment Benefits," research report, Google Inc., (July 5), (accessed October 31, 2014), [available at <http://static.googleusercontent.com/media/research.google.com/en/us/archive/papers/initialclaimsUS.pdf>].
- and ——— (2009b), "Predicting the Present with Google Trends," research report, Google Inc., (April 10), (accessed October 31, 2014), [available at http://static.googleusercontent.com/media/www.google.com/en/us/googleblogs/pdfs/google_predicting_the_present.pdf].
- Christen, Markus, Sachin Gupta, John C. Porter, Richard Staelin, and Dick R. Wittink (1997), "Using Market-Level Data to Understand Promotion Effects in a Nonlinear Model," *Journal of Marketing Research*, 34 (August), 322–34.
- Commandeur, Jacques J.F. and Siem Jan Koopman (2007), *An Introduction to State Space Time Series Analysis*. New York: Oxford University Press.
- Da, Zhi, Joseph Engelberg, and Pengjie Gao (2011), "In Search of Attention," *Journal of Finance*, 66 (5), 1461–99.
- Danaher, Peter J., André Bonfrer, and Sanjay Dhar (2008), "The Effect of Competitive Advertising Interference on Sales for Packaged Goods," *Journal of Marketing Research*, 45 (April), 211–25.
- Ding, Min, Young-Hoon Park, and Eric T. Bradlow (2009), "Barter Markets for Conjoint Analysis," *Management Science*, 55 (6), 1003–1017.
- Du, Rex Yuxing and Wagner A. Kamakura (2012), "Quantitative Trendspotting," *Journal of Marketing Research*, 49 (August), 514–36.
- Ebbes, Peter, Dominik Papies, and Harald J. van Heerde (2011), "The Sense and Non-Sense of Holdout Sample Validation in the Presence of Endogeneity," *Marketing Science*, 30 (6), 1115–22.
- Fader, Peter S. and Bruce G.S. Hardie (1996), "Modeling Consumer Choice Among SKUs," *Journal of Marketing Research*, 33 (November), 442–52.
- Fishbein, Martin and Icek Ajzen (1975), *Belief, Attitude, Intention and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley.
- Ginsberg, Jeremy, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, and Larry Brilliant (2009), "Detecting Influenza Epidemics Using Search Engine Query Data," *Nature*, 457, 1012–14.

- Green, Paul E. and Vithala R. Rao (1971), "Conjoint Measurement for Quantifying Judgmental Data," *Journal of Marketing Research*, 8 (August), 355–63.
- Hanssens, Dominique M., Leonard J. Parsons, and Randall L. Schultz (2001), *Market Response Models: Econometric and Time Series Analysis*, 2nd ed. Norwell, MA: Kluwer Academic Publishers.
- Hu, Ye, Rex Yuxing Du, and Sina Damangir (2014), "Decomposing the Impact of Advertising: Augmenting Sales with Online Search Data," *Journal of Marketing Research*, 51 (June), 300–319.
- J.D. Power and Associates (2008), *New Autosshopper.com Study*. Agoura Hills, CA: J.D. Power and Associates.
- (2012), *New Autosshopper.com Study*. Agoura Hills, CA: J.D. Power and Associates.
- Joo, Mingyu, Kenneth C. Wilbur, Bo Cowgill, and Yi Zhu (2013), "Television Advertising and Online Search," *Management Science*, 60 (1), 56–73.
- Kopalle, Praveen K., Carl F. Mela, and Lawrence Marsh (1999), "The Dynamic Effect of Discounting on Sales: Empirical Analysis and Normative Pricing Implications," *Marketing Science*, 18 (3), 317–32.
- Lamey, Lien, Barbara Deleersnyder, Jan-Benedict E.M. Steenkamp, and Marnik G. Dekimpe (2012), "The Effect of Business Cycle Fluctuations on Private-Label Share: What Has Marketing Conduct Got to Do with It?" *Journal of Marketing*, 76 (January), 1–19.
- Lancaster, Kelvin J. (1966), "A New Approach to Consumer Theory," *Journal of Political Economy*, 74 (2), 132–57.
- Lancaster, Tony (2004), *An Introduction to Modern Bayesian Econometrics*. Oxford, UK: Blackwell Publishing.
- Lazer, David, Ryan Kennedy, Gary King, and Alessandro Vespignani (2014), "The Parable of Google Flu: Traps in Big Data Analysis," *Science*, 343 (6176), 1203–1205.
- Lunn, David J., Andrew Thomas, Nicky Best, and David Spiegelhalter (2000), "WinBUGS—A Bayesian Modelling Framework: Concepts, Structure, and Extensibility," *Statistics and Computing*, 10 (4), 325–37.
- Ma, Yu, Kusum L. Ailawadi, Dinesh K. Gauri, and Dhruv Grewal (2011), "An Empirical Investigation of the Impact of Gasoline Prices on Grocery Shopping Behavior," *Journal of Marketing*, 75 (March), 18–35.
- Montgomery, Alan L. and Peter E. Rossi (1999), "Estimating Price Elasticities with Theory-Based Priors," *Journal of Marketing Research*, 36 (November), 413–23.
- Moorthy, Sridhar, Brian T. Ratchford, and Debabrata Talukdar (1997), "Consumer Information Search Revisited: Theory and Empirical Analysis," *Journal of Consumer Research*, 23 (4), 263–77.
- Neslin, Scott A. (1990), "A Market Response Model for Coupon Promotions," *Marketing Science*, 9 (2), 125–45.
- , Caroline Henderson, and John Quelch (1985), "Consumer Promotions and the Acceleration of Product Purchases," *Marketing Science*, 4 (2), 147–65.
- Pelat, Camille, Clement Turbelin, Avner Bar-Hen, Antoine Flahault, and Alain-Jacques Valleron (2009), "More Diseases Tracked by Using Google Trends," *Emerging Infectious Diseases*, 15 (8), 1327–28.
- Punj, Girish N. and Richard Staelin (1983), "A Model of Consumer Information Search Behavior for New Automobiles," *Journal of Consumer Research*, 9 (4), 366–80.
- Purcell, Kristen, Joanna Brenner, and Lee Rainie (2012), "Main Findings," Pew Research Internet Project. (March 9), (accessed November 10, 2014), [available at <http://www.pewinternet.org/2012/03/09/main-findings-11/>].
- Ratchford, Brian T., Myung-Soo Lee, and Debabrata Talukdar (2003), "The Impact of the Internet on Information Search for Automobiles," *Journal of Marketing Research*, 40 (2), 193–209.
- , Debabrata Talukdar, and Myung-Soo Lee (2007), "The Impact of the Internet on Consumers' Use of Information Sources for Automobiles: A Re-Inquiry," *Journal of Consumer Research*, 34 (1), 111–19.
- Srinivasan, Narasimhan and Brian T. Ratchford (1991), "An Empirical Test of a Model of External Search for Automobiles," *Journal of Consumer Research*, 18 (2), 233–42.
- Van Heerde, Harald J., Maarten J. Gijsenberg, Marnik G. Dekimpe, and Jan-Benedict E.M. Steenkamp (2013), "Price and Advertising Effectiveness over the Business Cycle," *Journal of Marketing Research*, 50 (April), 177–93.
- Vosen, Simeon and Torsten Schmidt (2011), "Forecasting Private Consumption: Survey-Based Indicators vs. Google Trends," *Journal of Forecasting*, 30 (6), 565–78.
- Wittink, Dick R. (1977), "Exploring Territorial Differences in the Relationship Between Marketing Variables," *Journal of Marketing Research*, 14 (May), 145–55.
- Wu, Lynn and Erik Brynjolfsson (2009), "The Future of Prediction: How Google Searches Foreshadow Housing Prices and Quantities," in *ICIS 2009 Proceedings*, Paper 147, (accessed November 10, 2014), [available at <http://aisel.aisnet.org/icis2009/147>].
- Zettelmeyer, Florian, Fiona Scott Morton, and Jorge Silva-Risso (2006), "How the Internet Lowers Prices: Evidence from Matched Survey and Automobile Transaction Data," *Journal of Marketing*, 43 (April), 168–81.