REX YUXING DU and WAGNER A. KAMAKURA*

Trendspotting has become an important marketing intelligence tool for identifying and tracking general tendencies in consumer interest and behavior. Currently, trendspotting is done either qualitatively by trend hunters, who comb through everyday life in search of signs indicating major shifts in consumer needs and wants, or quantitatively by analysts, who monitor individual indicators, such as how many times a keyword has been searched, blogged, or tweeted online. In this study, the authors demonstrate how the latter can be improved by uncovering common trajectories hidden behind the coevolution of a large array of indicators. The authors propose a structural dynamic factor-analytic model that can be applied for simultaneously analyzing tens or even hundreds of time series, distilling them into a few key latent dynamic factors that isolate seasonal cyclic movements from nonseasonal, nonstationary trend lines. The authors demonstrate this novel multivariate approach to quantitative trendspotting in one application involving a promising new source of marketing intelligence—online keyword search data from Google Insights for Search—in which they analyze search volume patterns across 38 major makes of light vehicles over an 81-month period to uncover key common trends in consumer vehicle shopping interest.

Keywords: marketing intelligence, market sensing, quantitative trendspotting, online searches, factor analysis, multivariate time-series analysis, common trends

Quantitative Trendspotting

To keep fingers on the pulse of consumers, “market sensing” has traditionally involved tracking and analyzing a few vital marketplace indicators. As more consumers use search engines to gather product-related information, conduct transactions and reviews online, and communicate with others through social media such as blogs, Twitter, and Facebook, an abundance of real-time indicators of their interests, opinions, and behaviors becomes available as a promising new source of information, allowing marketers to conduct market sensing in ways previously unattainable, leading to more timely and potentially more insightful marketing intelligence.

Indeed, many online services already gather and track these individual indicators, which are often charted over time and labeled as “trends.” The better known of such online consumer interest trending services is Google Trends (google.com/trends)/Google Insights for Search (google.com/insights/search), which provides search volume indexes for queries people have been entering into the Google search engine since 2004. Another well-known online consumer interest trending service is BlogPulse (blogpulse.com) by Nielsen, which scours the blogosphere on a daily basis to monitor the mention of individual keywords. Similarly, researchers can track real-time trends on any individual term in the Twitter space using services such as Trendistic (trendistic.com). To cover multiple platforms simultaneously, researchers can use more integrated services such as Trendrr (trendrr.com), which allows users to monitor public conversations about a product, service, or brand in real-time across a large variety of digital and social media.

Increasingly, marketers believe they can not only monitor but also get an early read on real-world trends by studying tracking metrics from online sources (e.g., trendly.com, recordedfuture.com, icerocket.com). Anecdotally, an analysis of Twitter activity reported by the *Wall Street Journal* showed a “gradual decline in the number of people tweeting about being unemployed and looking for work over the six months ending Feb. 16 [2010]” (Merrit 2010), presaging an upturn in employment. Similarly, Yellowbook.com tracked the types of businesses consumers search for and concluded

*Rex Yuxing Du is Hurley Associate Professor of Marketing, Bauer College of Business, University of Houston (e-mail: rexdu@bauer.uh.edu). Wagner A. Kamakura is Ford Motor Co. Global Marketing Professor, Fuqua School of Business, Duke University (e-mail: kamakura@duke.edu). Vithalo Rao served as associate editor for this article.

ISSN: 0022-2437 (print), 1547-7193 (electronic)

that an "upturn in searches related to home improvement and remodeling began in October [2009].… In late January, Home Depot and Lowe's reported increased fourth-quarter sales" (Merrit 2010).

While online trending services such as the aforementioned are potentially useful, they provide tracking data by charting the time courses of individual terms. For example, when a topic, brand, or celebrity is deemed to be trending up or down, it is usually based on the trajectory of a single indicator. Unfortunately, individual indicators can be noisy and unreliable by themselves, containing idiosyncrasies (e.g., seasonal fluctuations, measurement errors, pure random shocks) that can mask or confound the real, nonseasonal trend lines hidden underneath. Moreover, individual indicators are likely to show strong co-movements. For example, increasing searches for one product may be accompanied by decreasing searches for its substitutes and increasing searches for its complements. Thus, to increase the signal-to-noise ratio and avoid reading too much into any individual indicator, it seems reasonable to simultaneously examine the time courses of many related (or even seemingly unrelated) indicators, from which a few key common trajectories may be distilled, leading to a more reliable read on the underlying trends that have manifested through the co-movements of a large array of individual indicators.

To help marketers take a more holistic view of the multitude of tracking metrics available through online (and offline) sources, we propose a modeling framework through which analysts can methodically identify, interpret, and project common trend lines hidden behind a large number of time series. Hereinafter, we refer to such an endeavor as "quantitative trendspotting," the ultimate goal of which is to help marketers uncover hidden gems of insight buried underneath the rubble of high-dimensional tracking data.

QUANTITATIVE VERSUS QUALITATIVE TRENDSPOTTING

Before we delve into the conceptual framework and statistical model for quantitative trendspotting, it is important to compare and contrast it with its qualitative counterpart. Long before the emergence of online consumer interest trending services, marketers have carried out trend analyses as an integral part of marketing intelligence, attempting to identify and track general tendencies in consumer interest and behavior. However, until the widespread use of search engines, blogs, Twitter, Facebook, and other digital and social media, trendspotting was mostly a qualitative endeavor, involving highly sensitive "trend hunters" adept at detecting far-reaching new developments by combing through everyday life and interacting with recognized trendsetters or opinion leaders. We refer to this traditional form of trendspotting as "qualitative trendspotting" because it relies heavily on the acuity, judgment, and foresight of a few people in correctly recognizing and interpreting the manifested signs of a new trend.

In our view, the main distinction between quantitative and qualitative trendspotting is that they are about different types of trends. Quantitative trendspotting is about trends in the tradition of time-series analysis—that is, the trajectories or trend lines of a collection of longitudinal measures. In contrast, qualitative trendspotting is about foretelling the onset of a path-breaking shift in consumer interest or behavior, such as the emergence of a new paradigm in fashion or lifestyle. In other words, quantitative trendspotting is more evolutionary, focusing on uncovering hidden trend lines that already exist, which may then be extrapolated into the near future, whereas qualitative trendspotting is more revolutionary, attempting to seek out radical departures from the past that may potentially reshape the marketplace for years to come.

A CONCEPTUAL FRAMEWORK FOR QUANTITATIVE TRENDSPOTTING

Distinctive purposes aside, quantitative and qualitative trendspotting are both exploratory endeavors, aimed at making better sense of the observed marketplace indicators, as opposed to prescribing exact courses of action. The main goal of quantitative trendspotting is to help managers distill massive amounts of historical indicators into a few key trend lines that have jointly shaped the evolution of these indicators. In short, as a tool for market sensing, we envision quantitative trendspotting as an exploratory process consisting of the following five steps:

- *Step 1: Gathering individual indicators*: Assemble potential indicators from various sources. The challenge lies in the proverbial embarrassment of riches: massive amounts of data increasingly available, leading to tens or even hundreds of seemingly relevant indicators. The way to deal with this challenge is to examine all of them holistically and systematically, instead of arbitrarily creating some sort of composite index or subjectively preselecting some indicators while ignoring others, which could be inefficient or even misleading.
- *Step 2: Extracting key common trends*: Distill all the individual indicators into a few common factors. Simultaneous analysis of all the indicators is necessary because otherwise the co-movement patterns they share may be indistinguishable. Furthermore, individual indicators often include substantial inaccuracies and noise that must be filtered out to uncover the more reliable underlying trend lines shared by multiple indicators, thus lowering the risk of overreacting to idiosyncrasies of any individual indicator.
- *Step 3: Interpreting the identified common trends*: Make sense of the trend lines identified previously. As we explain subsequently when we illustrate our proposed model, this interpretation is aided by parameters measuring the correlation between the observed indicators and the latent common trends. The resulting insights should shed light on the underlying forces that have shaped the bulk of the observed dynamics. Furthermore, cyclical movements must be separated from noncyclical trends. The former are typically driven by seasonal factors (e.g., weather, annual events). The latter are less predictable but are often more significant, reflecting genuine shifts in consumer interest or behavior over time.
- *Step 4: Generating insights by relating the identified trends to other variables*: After the latent trends have been identified and interpreted, to further establish their validity and generate insights, subsequent analyses of the trends themselves can be carried out. For example, the identified trends can be related to potential causal variables to understand what may have shaped them in the past. Similarly, the identified trends can be related to outcome measures to better understand their business impacts and managerial relevance.
- *Step 5: Projecting the identified trends*: The trend lines identified from historical tracking data can be extrapolated into the near future. Equipped with such projections, managers should be better positioned in deciding what to do proactively in an

attempt to reverse certain unfavorable trends while leveraging the favorable ones.

The rest of the article proceeds as follows: First, we provide an overview of dynamic factor analysis (DFA), through which latent common trends that have jointly shaped the observed individual time series can be identified. Then, we present our proposed model, which extends the standard DFA model to allow for a more flexible and more interpretable dynamic process governing the evolution of the latent common trends. We illustrate the application of our model using data from Google Insights for Search (GIFS), in which search volume indexes across 38 major makes of light vehicles are analyzed simultaneously over an 81-month period. We demonstrate how our modeling framework may allow analysts and managers to better leverage the multitude of marketplace indicators available to them, uncovering trends and insights that would not be visible by examining these individual indicators separately.

OVERVIEW OF DYNAMIC FACTOR ANALYSIS

Researchers are often presented with high-dimensional longitudinal data (e.g., a large number of macroeconomic indicators) and want to uncover a few key factors (e.g., the health of the economy) that have shaped the observed multivariate time series. Instead of analyzing each series separately or focusing on a small subset of them by arbitrarily aggregating or preselecting certain series, DFA offers a formal statistical approach for analyzing large panels of time series simultaneously, allowing the researcher to extract a small number of common factors from the entire data set by uncovering the predominant co-movement patterns.

Researchers in many different fields have long used DFA to identify common trend lines hidden underneath multiple time series. Early applications can be found in econometrics (e.g., Engle and Watson 1981; Geweke 1977; Harvey 1989; Lutkepolh 1991), psychometrics (e.g., Molenaar 1985; Molenaar, Gooijer, and Schmitz 1992), and statistics (e.g., Shumway and Stoffer 1982). Geweke (1977), in an analysis of multiple macroeconomic input and output time series, is often credited with the first DFA model in econometrics.[1] In psychometrics, Molenaar (1985) first proposed a DFA model to analyze multivariate physiological measures taken repeatedly from a single subject, where the set of measurement occasions gives rise to an ordered sample of indicators of the subject's latent behavioral states.

More recently, Zuur et al. (2003) introduced DFA to environmetrics as a tool for uncovering common trajectories shared by a large number of biological and environmental time series, while Ludvigson and Ng (2007), in a study of risk and return in the U.S. stock market, effectively summarized the information contained in 209 series of macroeconomic indicators and 172 series of financial indicators using two eight-factor DFA models. Aruoba, Diebold, and Scotti (2009) demonstrate how DFA can be applied for real-time measurement of aggregate business conditions, which they argue should be treated as a latent variable that is not tied to any single observed indicator but rather is manifested through the shared dynamics of many observed indicators. In the same spirit, Doz and Lenglart (2001) apply DFA to extract one common factor from the co-movements of 30 time series tracked by European industrial business surveys. They show that the uncovered common factor can be used as a composite business cycle index because the past evolution of this factor, and in particular its turning points, can be easily interpreted in terms of major economic events in the Euro area.

At its core, DFA is a technique for dimension reduction; each of the n time series under study is modeled as the sum of a residual term representing idiosyncratic movements, and a linear combination of p ($<<$ n) unobserved variables representing the underlying drivers of the observed co-movements (i.e., the latent common trends). Just as in standard factor analysis (FA), the basic principle of DFA is to use as few latent factors as possible to retain as much variation observed in the data as possible. The argument against applying FA to multivariate time-series data (Anderson 1963) is that FA assumes that the factor scores are independent across observations, which is plausible when each observation represents an element from a cross-section but inappropriate when the observations are taken over time. If we apply FA to multivariate time series, the model will produce identical factor scores and loadings even if we rearrange the order of the observations. In other words, FA completely ignores temporal relationships inherent to longitudinal observations, which is inefficient and can be erroneous.

It is expedient to view DFA as a merger between standard FA (for dimension reduction) and univariate time-series analysis (for temporal dynamics). In its state-space form, a DFA model consists of one set of equations for the observed variables and another set of state equations for the latent factors. The observation equations describe the relationship between the n observed variables and the p latent factors. The state equations describe the law of motion for the p latent factors. Broadly speaking, DFA models fall into two main types. In the first type (e.g., Engle and Watson 1981; Zuur et al. 2003), the observation equations specify that an observed variable at time t is a function of the latent factors at the same instant t. These observation equations are identical to those in standard FA. What makes these models distinct from standard FA is the formulation of the state equations, which specify the latent factors at time t as a function of their lagged values. As a result, the state equations introduce dynamics into the system, capturing common temporal patterns through the estimated serial relationship between the latent factors and their lagged values.

In contrast, the second type of DFA models (e.g., Forni et al. 2000; Sargent and Sims 1977; Stock and Watson 2002) does not incorporate dynamics in the state equations, assuming that a latent factor at time t is independent of its lagged values. What makes these models distinct from standard FA is the formulation of the observation equations, which specify each observed variable at time t as a function of not only the current but also the lagged latent factor scores. In such a formulation, common temporal patterns

[1]Under the condition that the number of measurement occasions is small and the observed indicators have a stationary lagged covariance structure, a DFA model can be reformulated into a structural equation model (e.g., Jöreskog 1979), which can then be calibrated with LISREL. However, the assumption of stationary lagged covariance structure is too restrictive for quantitative trendspotting exercises, which must handle a large number of time series that can follow many types of seasonal and nonstationary trajectories over extended time periods. Thus, here, we focus on DFA models that stem from the econometrics literature (i.e., Geweke 1977).

embedded in the observed multivariate time series will all be contained in the observation equations.

To develop our new DFA model, we choose the first type because the factor scores at time t can be interpreted unambiguously as the "state" of the system, summarizing all that has happened in the system up to time t; the observed variables are manifestations of this state. In other words, the factor scores directly represent the latent trends. In contrast, the interpretation of the factor scores from the second type of formulation is less straightforward, because they have an impact on the observed variables over multiple time periods. Another (secondary) reason for choosing the first type of formulation is that it requires substantially fewer parameters because the dynamic structure is on the p state equations rather than the n (>> p) observation equations. Admittedly, our overview of DFA is by no means exhaustive, and researchers have continued to introduce new applications and new methods of inference. For recent advances in this area, see Croux, Renault, and Werker (2004) and Molenaar and Ram (2009).

OTHER TIME-SERIES METHODS

Although DFA has been well established in many different fields, to the best of our knowledge, it has not appeared in the marketing literature, which is a bit surprising given that marketing researchers have widely adopted standard FA and how often they must deal with high-dimensional time-series data (e.g., the multitude of tracking measures collected through repeated surveys). In contrast, we argue that current time-series methods familiar to marketing researchers are not well suited for the task of quantitative trendspotting. For example, it is impossible to conduct quantitative trendspotting as we have envisioned using univariate time-series techniques such as spectral analysis (Priestley 1981), wavelet analysis (Shumway and Stoffer 2000), and Box-Jenkins/ARIMA (Ljung 1987), because these methods can only be applied to one indicator series at a time, thus making them incapable of examining multiple indicators simultaneously to uncover their shared trend lines.

As for methods that can be applied simultaneously to multiple time series, vector autoregression (VAR) and vector autoregressive moving average models (VARMA) are probably the most popular among marketing researchers. Unfortunately, standard VAR/VARMA models tend to break down under the curse of dimensionality as the number of time series reaches double digits. More recently, Bayesian VAR (BVAR) has been advocated as an alternative for dealing with high-dimensional time series (e.g., De Mol, Giannone, and Reichlin 2008): It overcomes the curse of dimensionality by restricting the parameter space a priori through highly informative priors, and to avoid overfitting, the tightness of the prior increases as the number of series increases (Carriero, Kapetanios, and Marcellino 2009).

The effectiveness of BVAR inevitably depends on the quality of the priors, which can pose a challenge when the dimensionality is high and there is little prior knowledge about the data under study. Furthermore, imposing tight priors may not help when presumptions about the underlying dynamic structure are ultimately wrong. In contrast, DFA, similar in spirit to other exploratory factor-analytic methods, lets the data speak for themselves, revealing, post hoc, the underlying joint dynamics through the structure of the estimated factor loadings. Finally, and more important, neither standard nor Bayesian VAR produces outputs that can be interpreted as common trend lines hidden beneath the observed time series, thereby precluding their use for quantitative trendspotting, the main purpose of our study. Table 1 summarizes the comparisons between DFA, our extension of it (which is introduced next), and other time-series methods.

STRUCTURAL DFA MODEL

In this section, we present our novel DFA model for quantitative trendspotting, which we named structural DFA (SDFA). It is structural in the sense that, instead of assuming that each latent factor follows a random walk or a simple autoregressive process (as DFA models typically do), we impose a directly interpretable and more flexible structure on the law of motion, decomposing each latent factor into a seasonal component and a nonseasonal component that follows a locally quadratic trend line. Such a statistical decomposition of the underlying dynamic process falls into what is known in the econometrics literature as structural time-series analysis (Harvey 1989; Harvey and Shephard 1993). A major distinction is that our SDFA model uses the structure in the latent factor space and therefore can be applied to uncover trends from high-dimensional time series, while existing structural time-series models cannot.

Table 1
SUMMARY OF METHOD COMPARISONS

	ARIMA	*VAR/VARMA*	*Bayesian VAR*	*Standard FA*	*DFA*	*Structural Time-Series Analysis*	*SDFA*
Examines multivariate time series simultaneously to account for co-movements	No	Yes	Yes	No	Yes	No	Yes
Accounts for temporal interdependence	Yes	Yes	Yes	No	Yes	Yes	Yes
Applies to large panels of time series without breaking down under the curse of dimensionality	Yes	No	Yes	Yes	Yes	Yes	Yes
Makes flexible projections into the future	Yes	Yes	Yes	No	No	Yes	Yes
Helps the analyst see the big picture by identifying common trend lines hidden behind all the available indicator series	No	No	No	No	Yes	No	Yes
Decomposes law of motion into directly interpretable stochastic, nonstationary seasonal and nonseasonal movements	No	No	No	No	No	Yes	Yes

Suppose that an analyst observes a vector y_t containing n indicators over periods t = 1, 2, …, T and wants to identify, track, and extrapolate the key common trend lines behind these n time series. Our proposed model is formulated as the following system of vector equations:

(1)	$y_t = B + Lf_t + u_t$	$u_t \sim N(0, \Sigma_u)$	observed indicator
(2)	$f_t = \alpha_t + \gamma_t$		unobserved dynamic factor
(2.1)	$\alpha_t = \alpha_{t-1} + \beta_{t-1} + \varepsilon_t$	$\varepsilon_t \sim N(0, \Sigma_\varepsilon)$	nonseasonal trend component
(2.2)	$\beta_t = \beta_{t-1} + \delta_{t-1} + \eta_t$	$\eta_t \sim N(0, \Sigma_\eta)$	slope of nonseasonal trend
(2.3)	$\delta_t = \delta_{t-1} + \zeta_t$	$\eta_t \sim N(0, \Sigma_\zeta)$	change rate of slope
(2.4)	$\gamma_t = -\Sigma_{j=1}^{s-1}\gamma_{t-j} + \xi_t$	$\xi_t \sim N(0, \Sigma_\xi)$	seasonal component

Equation 1 models the n-dimensional indicators (y_t) as the sum of (1) intercepts, B; (2) a linear combination of p unobserved dynamic factors, f_t; and (3) an irregular component u_t that is assumed to be independent over time and normally distributed with mean 0 and variance Σ_u. Matrix L (n × p) reflects how the p unobserved dynamic factors f_t load onto the n observed indicators y_t, with $p \ll n$. All else being equal, the larger the absolute value of L_{gh}, the stronger is the correlation between the hth unobserved dynamic factor and the gth observed indicator. Indicators with strong co-movements will have high loadings on the same factors.

The set of state equations (Equations 2.1–2.4) indicates that f_t, the dynamic factors, can be further decomposed into a nonseasonal trend component, α_t, and a seasonal component, γ_t. The dynamics of the nonseasonal trend α_t is governed by Equations 2.1–2.3, while the dynamics of the seasonal component γ_t is governed by Equation 2.4. Equation 2.1 states that the nonseasonal trend α_t is updated from its lagged value α_{t-1} as a result of two additive terms, ε_t and β_{t-1}. We assume the first term, ε_t, (a random shock received at t), to be a priori independent over time and normally distributed with mean 0 and diagonal variance Σ_ε. The second term, β_{t-1}, represents a time-varying slope in the latent trend. β_{t-1} evolves over time according to Equation 2.2, where η_t (a random shock received at t) is assumed a priori to be independent over time and normally distributed with mean 0 and diagonal variance Σ_η. In Equation 2.2, δ_{t-1} represents the expected slope change rate at t, which evolves over time following a random walk defined by Equation 2.3, where diagonal Σ_ζ denotes the variance of ζ_t (a random shock received at t).

To better explain the dynamics that Equations 2.1–2.3 can generate and how our SDFA extends DFA, it helps to set aside the seasonal component γ_t for the time being and focus on several more restricted formulations of the law of motion for the nonseasonal trend component α_t. Beginning with the most basic, remove α_{t-1} and β_{t-1} from Equation 2.1, such that $\alpha_t = \varepsilon_t$, and the system degenerates into standard FA. Relaxing this assumption so that $\alpha_t = \alpha_{t-1} + \varepsilon_t$ produces the standard DFA. This extension of FA to DFA is deceptively simple because with $\alpha_t = \alpha_{t-1} + \varepsilon_t$, as opposed to $\alpha_t = \varepsilon_t$, the likelihood function becomes much more complex, requiring joint calculation of the stochastic term ε_t across all time periods.

A key limitation of standard DFA lies in that it is not well suited for projecting trends because, given $\alpha_t = \alpha_{t-1} + \varepsilon_t$, the model would extend a flat trend line into the future at the end of the observation window. In quantitative trendspotting, analysts and managers are often interested in distinguishing trend lines that are headed up from those that are headed down, and trend lines that are accelerating from those that are losing momentum, which becomes the main motivation for our extension of DFA into SDFA by relaxing some of its restrictions.

First, we can relax the standard DFA formulation ($\alpha_t = \alpha_{t-1} + \varepsilon_t$) by making $\alpha_t = \alpha_{t-1} + \beta_0 + \varepsilon_t$, which can be rewritten as

$$\text{(2.1.a)} \qquad \alpha_t = \alpha_0 + \beta_0 t + \sum_{k=1}^{t} \varepsilon_k,$$

which can be viewed as a random walk in the level (ε_k) superimposed on a linear trend line that has a constant slope of β_0.

Furthermore, instead of using a constant slope, we allow the slope to evolve over time by making $\alpha_t = \alpha_{t-1} + \beta_{t-1} + \varepsilon_t$ and $\beta_t = \beta_{t-1} + \delta_0$, leading to a nonseasonal trend α_t that can be rewritten as

$$\text{(2.1.b)} \qquad \alpha_t = \alpha_0 + \beta_0 t + \delta_0 \frac{t(t-1)}{2} + \sum_{k=1}^{t} \varepsilon_k,$$

which can be viewed as a random walk in the level (ε_k) superimposed on a quadratic trend line that has an initial slope of $\beta_0 - \delta_0/2$ and a constant slope change rate of δ_0.

Moreover, instead of a constant slope change rate, we allow the change rate to evolve over time by making $\beta_t = \beta_{t-1} + \delta_0 + \eta_t$. This leads to a nonseasonal trend α_t that can be rewritten as

$$\text{(2.1.c)} \quad \alpha_t = \alpha_0 + \beta_0 t + \delta_0 \frac{t(t-1)}{2} + \sum_{k=1}^{t} \varepsilon_k + \sum_{k=1}^{t-1} \eta_k (t-k),$$

which can be viewed as a random walk in the level (ε_k) superimposed to a quadratic trend line that has an initial slope of $\beta_0 - \delta_0/2$, which then changes over time following a random walk (η_k) with a constant drift of δ_0.

Finally, our complete SDFA model relaxes the constant drift restriction by adding a random walk to the slope change rate (i.e., Equation 2.3, $\delta_t = \delta_{t-1} + \zeta_t$). Together, Equations 2.1, 2.2, and 2.3 imply the following law of motion for the nonseasonal trend component α_t:

$$\text{(2.1.d)} \qquad \alpha_t = \alpha_0 + \beta_0 t + \delta_0 \frac{t(t-1)}{2} + \sum_{k=1}^{t} \varepsilon_k + \sum_{k=1}^{t-1} \eta_k (t-k) + \sum_{k=1}^{t-2} \zeta_k \frac{(t-k)(t-k-1)}{2}.$$

The upshot of Equation 2.1.d is that our SDFA model implies a *locally* quadratic trend line for the latent nonseasonal component α_t. It is important to distinguish a locally quadratic trend line from a globally quadratic one. With a globally quadratic trend line, the level and linear and quadratic terms are all constant over time. With a locally quadratic trend line, as in Equation 2.1.d, the level and linear and quadratic terms are all dynamically adjusted depending on the stochastic shocks (i.e., ε_k, η_k, and ζ_k) received over time, leading to a highly flexible trajectory. It is also important to note that in our SDFA model, the diagonal variances of the stochastic shocks (i.e., Σ_ε, Σ_η, and Σ_ζ) are all estimated from data. If they turn out to be empirically indistinguishable from zero, the nonseasonal component α_t degenerates into following a globally quadratic trend line shaped by α_0, β_0, and δ_0, which are also empirically determined. If the diagonals of Σ_ε, Σ_η, and Σ_ζ prove to be empirically greater than zero, the trend line will be less driven by α_0, β_0, and δ_0 and more by the shocks received over time (i.e., ε_k to the level, η_k to the slope, and ζ_k to the change rate of the slope). For a recent application of the locally quadratic trend model to a single time series, see Harvey (2010).

Equation 2.4 describes the law of motion for the seasonal component γ_t , which implies that each of the s seasons has a distinct effect and the sum of any s consecutive seasonal effects will have an expectation of zero and a diagonal variance of Σ_ξ (to illustrate this, we can rewrite Equation 2.4 as $\Sigma_{j=0}^{s-1}\gamma_{t-j} = \xi_t$). The choice of s is exogenous and typically straightforward (e.g., 12 for monthly, 4 for quarterly, 52 for weekly data).

Empirically, if the diagonal of Σ_ξ approaches zero, Equation 2.4 degenerates into s fixed seasonal dummy effects that are constrained to sum to zero. As the diagonal of Σ_ξ increases, the seasonal effects will gradually evolve over time. Modeling seasonal fluctuations separately enables us to hone in on the genuine nonseasonal movements, which are, in general, more significant and less predictable. This also constitutes an important extension to standard DFA, which does not distinguish seasonal from nonseasonal movements in the dynamic latent factors.

In summary, each component of our SDFA model as shown in Equations 1, 2, and 2.1–2.4 has an easy-to-understand interpretation. Tracing out the α_t and γ_t over time would generate, respectively, p-dimensional nonseasonal and seasonal trend lines. In addition, tracing out the β_t and δ_t would show, respectively, how the slopes of the nonseasonal trend lines and the change rates of those slopes have evolved over time. Appendixes A and B provide details about (1) how the latent state variables (i.e., $z_t \overset{\text{def}}{=} [\alpha_t\ \beta_t\ \delta_t\ \gamma_t\ \ldots\ \gamma_{t-s-2}]'$) can be inferred on the basis of y_t observed from t = 1 through T and extrapolated h periods into the future for t = T + 1 through T + h and (2) how our SDFA model can be reformulated in a state-space form and calibrated with a highly efficient and robust expectation-maximization algorithm.

To recap, we propose a novel model for uncovering and projecting common trend lines hidden behind a large number of time series, the SDFA model. First, it is *factor analytic* because, similar to standard FA, it extracts a small number of common latent factors from a large number of observed variables. By reducing the dimensionality of the observed data, the latent factors help identify key patterns of comovement. Second, our model is *dynamic* because it takes into account that indicators observed over time are temporally dependent. Finally, our model is *structural* because, unlike standard DFA, which extends a flat trend line, our model allows much more flexibility. The particular dynamic structure we impose on the latent factors consists of seasonal and nonseasonal components, both directly interpretable. The nonseasonal component is modeled as following a locally quadratic trend line governed by three separate stochastic processes—one in the level, one in the slope, and one in the slope change rate.

QUANTITATIVE TRENDSPOTTING WITH SDFA

We apply our SDFA model in a scenario that is becoming increasingly common. Consider a business analyst who has been charged with gathering marketing intelligence on the U.S automotive industry. The analyst is particularly interested in finding out how consumer interests in various makes of vehicles have evolved over the years and whether there have been any predominant common trend lines. Moreover, the analyst is interested in learning more about what may have shaped these trend lines, how they may be correlated with performance measures such as sales, and where these trend lines may be headed in the near future.

Because the Internet has become such a dominant source of information for vehicle shopping, the analyst plans to use GIFS as the primary data source for this task, which offers many advantages. First, GIFS provides volume indexes for any queries people have been entering into the Google search engine, dating back to January 2004. Second, realtime data are readily accessible to the public for free. Third, GIFS allows filtering by categories and subcategories (e.g., searches for “Toyota” can be limited to the “Automotive” category, which can be narrowed down to subcategories such as “Vehicle Shopping” or “Vehicle Maintenance”).

The analyst gathered the data needed as follows: First, a list of the top 38 makes of light vehicles in the U.S. was obtained from *Automotive News*. Second, using the Keyword Tool from Google AdWords, the analyst obtained the search terms that are commonly used for each vehicle make, including popular abbreviations and misspellings (e.g., Volkswagen + VW + Volkswagon, Mercedes + Benz + Mercedez, Chevrolet + Cheverolet + Chevorlet + Chev). Third, the search terms were entered into the query box on the home page of GIFS, with the filters set to “Web Search” in the “United States” for “2004–present” within the “Vehicle Shopping” subcategory. Monthly search volume indexes (measured on a 100-point scale, with the maximum set to 100) were obtained from the displayed charts, from January 2004 to September 2010. Finally, all the raw search indexes were log-transformed to reduce skewness and then standardized to have zero mean and unit standard deviation (for more discussion on the issue of standardization, see Web Appendix B at www.marketingpower.com/jmr_webappendix). Figure 1, Panels A and B, presents the resulting 38 time series.[2]

A cursory observation of Figure 1, Panels A and B, reveals a few notable temporal patterns regarding consumer interest in vehicle shopping. First, many series show strong seasonality, with peaks and valleys at regular intervals, but not necessarily sharing the same cycle. Second, seasonal

[2]The standardized data are used in all subsequent analyses. Beyond the existing data period, the analyst can return to the nonstandardized scale by simply multiplying the standard deviation and adding the mean, both calculated according to raw data from the existing data period.

Figure 1

SEARCH VOLUME PATTERNS FOR THE TOP 38 VEHICLE MAKES IN THE UNITED STATES FROM JANUARY 2004 TO SEPTEMBER 2010

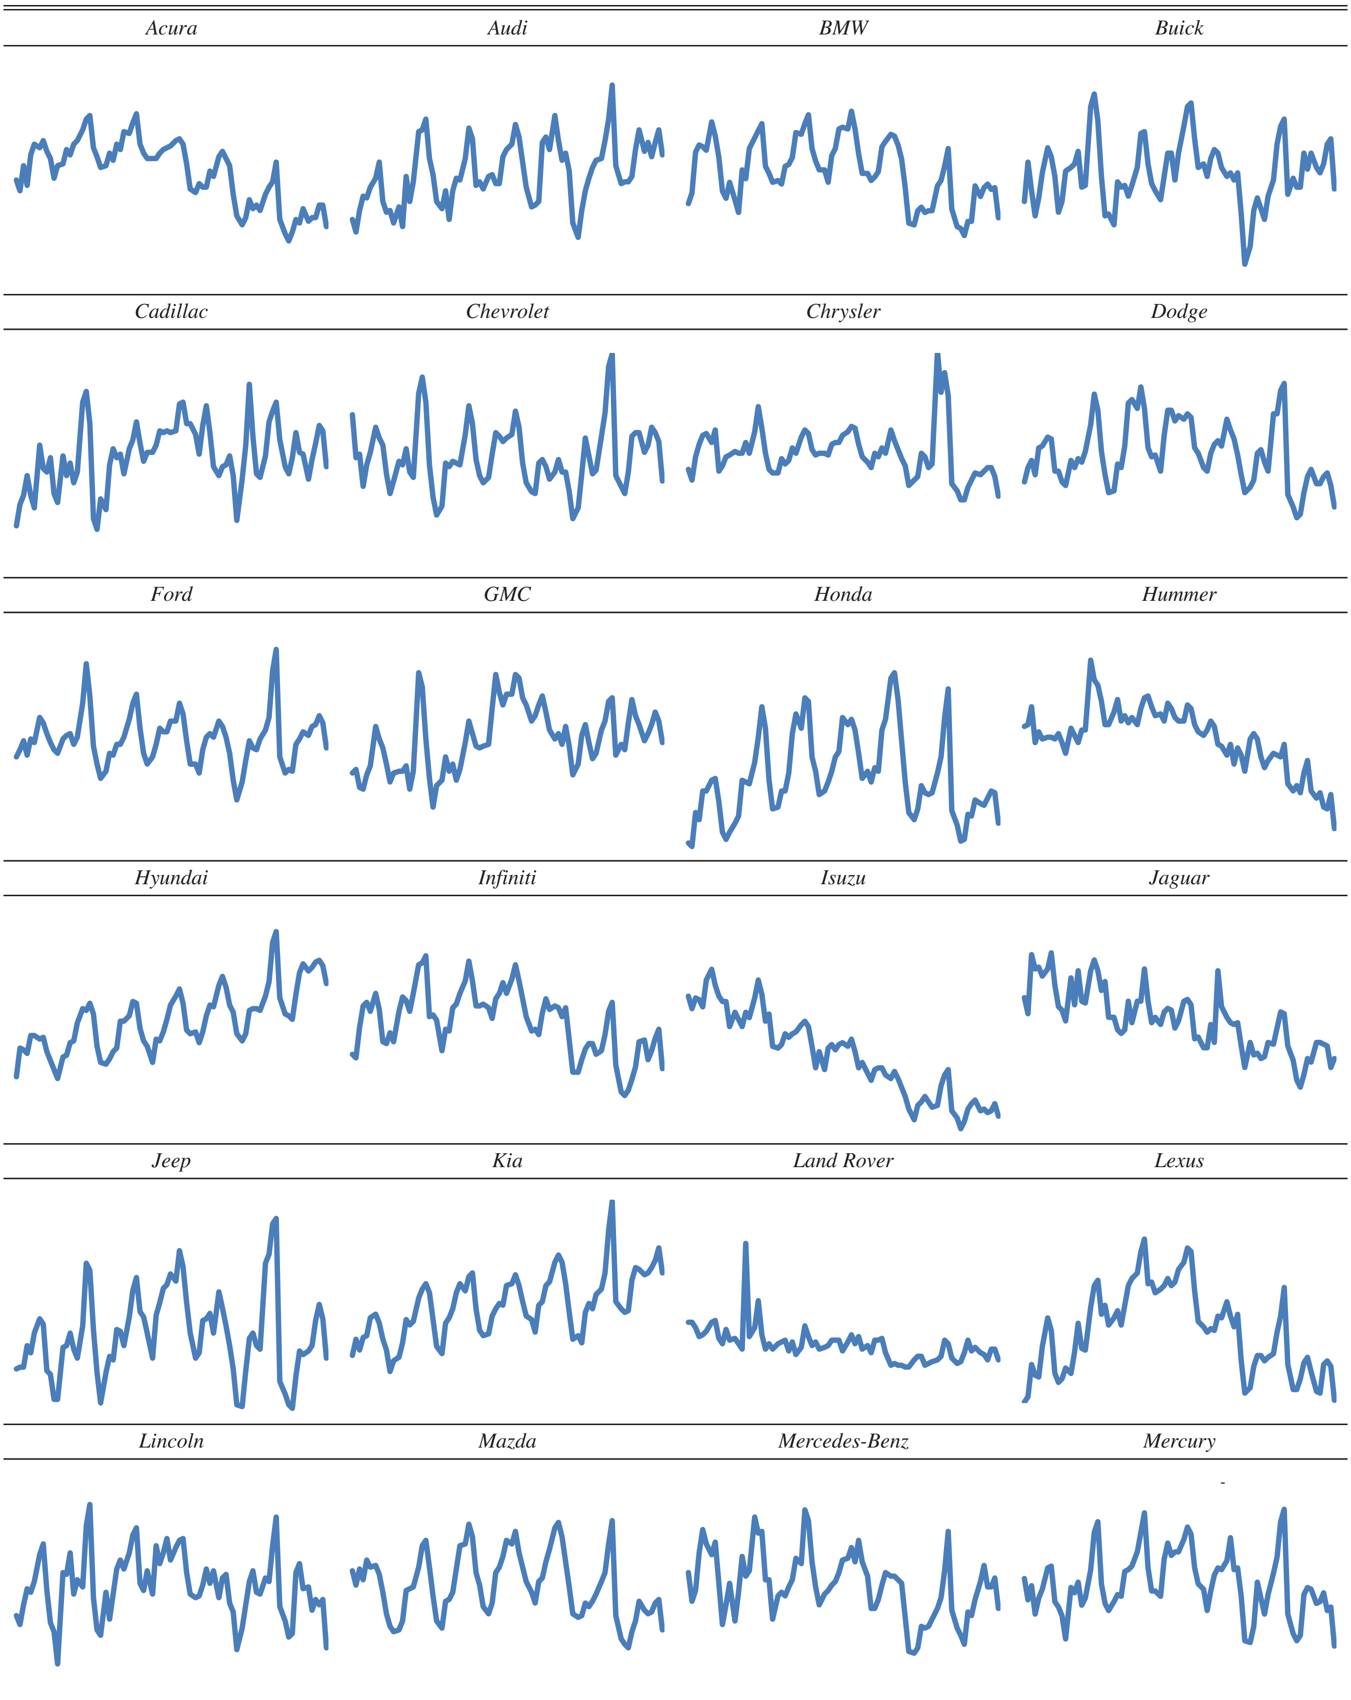

Figure 1
CONTINUED

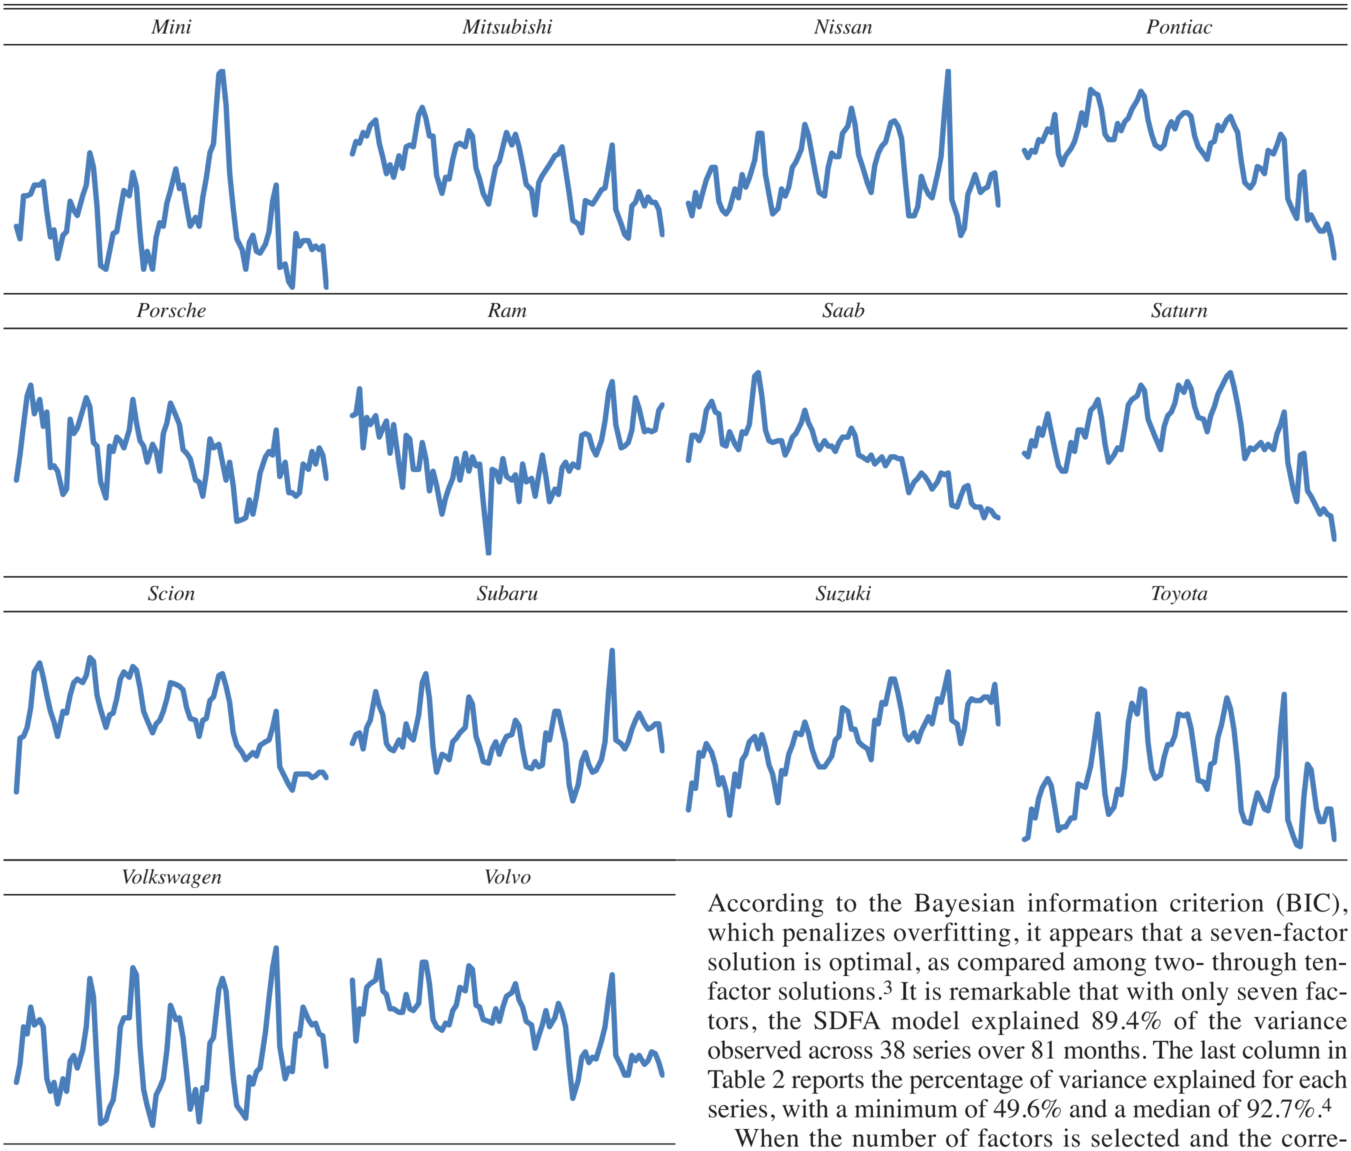

fluctuations aside, most series seem nonstationary, with a few showing substantial and consistent growth (e.g., Hyundai, Kia, Suzuki) or decline (e.g., Hummer, Isuzu, Saab) during the 81-month observation window. To see through the 38 individual series and grasp the “big picture,” the analyst needs a device that can help identify the common trajectories hidden underneath all the observed data. As we demonstrate subsequently, when uncovered using our SDFA model, such common trajectories can stimulate an open, exploratory process, leading to insights that would otherwise be unattainable.

Model Outputs and Interpretations

Our SDFA model offers a dynamic dimension-reduction tool that allows the analyst to capture as much of the comovement patterns shared by the 38 series with as few latent factors as possible. Just as in standard FA, to determine the optimal number of latent factors, the analyst needs to balance goodness of fit, parsimony, and interpretability. According to the Bayesian information criterion (BIC), which penalizes overfitting, it appears that a seven-factor solution is optimal, as compared among two- through ten-factor solutions.[3] It is remarkable that with only seven factors, the SDFA model explained 89.4% of the variance observed across 38 series over 81 months. The last column in Table 2 reports the percentage of variance explained for each series, with a minimum of 49.6% and a median of 92.7%.[4]

When the number of factors is selected and the corresponding model is estimated, the analyst can focus on two sets of outputs: (1) the time-invariant latent factor loadings (i.e., the $L_{n \times p}$ matrix in Equation 1), which maps p latent dynamic factors (f_t) onto the n observed time series (y_t), and (2) the time-varying seasonal (γ_t in Equation 2.4) and nonseasonal components (α_t, β_t, and δ_t in Equations 2.1–2.3) of the latent dynamic factors f_t.

Table 2 presents the estimated factor loadings, obtained at the end of the M-step of the estimator when the expectation-maximization algorithm converges and a Varimax rotation is applied.[5] These loadings estimates can be interpreted in

[3]For the various factor solutions, the BICs are 5071 (two-factor), 4229 (three-factor), 3769 (four-factor), 3548 (five-factor), 3446 (six-factor), 3370 (seven-factor, the smallest), 3415 (eight-factor), 3425 (nine-factor), and 3471 (ten-factor).

[4]The percentage of variance explained measure is similar to the measure of communality in standard FA.

[5]For more details on the expectation-maximization estimator and Varimax rotation, see Web Appendix B at www.marketingpower.com/jmr_webappendix.

Table 2

LATENT FACTOR LOADINGS AND PERCENTAGES OF VARIANCE EXPLAINED (R^2)

id	Manufacturer (Country of Origin)	Make	Foreign Mass	U.S. Mass	Euro Lux	GM Survived	Lexus/ (Ram)	Subaru/ (Mazda and Saturn)	GM Cancelled and Isuzu/ (Hyundai Kia and Suzuki)	R^2
1	Honda (Japan)	Acura	.31	.21	.36	.08	.32	.08	.42	96.7%
2	Volkswagen (Germany)	Audi	.56	.36	.21	.13	.13	.24	–.54	86.3%
3	BMW (Germany)	BMW	.32	.22	.73	.16	.16	–.12	.02	95.3%
4	GM (United States)	Buick	.31	.47	.35	.63	–.07	.28	.13	87.3%
5	GM (United States)	Cadillac	.14	.44	.02	.63	.20	.01	.01	77.8%
6	GM (United States)	Chevrolet	.16	.90	.24	.21	–.10	.07	.11	98.5%
7	Chrysler (United States)	Chrysler	.27	.62	.20	–.02	.14	–.21	.18	64.8%
8	Chrysler (United States)	Dodge	.37	.61	.27	.04	.33	–.24	.11	92.1%
9	Ford (United States)	Ford	.38	.72	.31	.15	–.01	.20	–.02	92.9%
10	GM (United States)	GMC	.10	.34	.18	.84	.01	–.10	.00	98.3%
11	Honda (Japan)	Honda	.78	.21	.09	.00	.39	–.06	–.16	99.1%
12	GM (United States)	Hummer	.03	.23	.10	.21	.36	–.19	.71	87.1%
13	Hyundai_Kia (Korea)	Hyundai	.46	.32	.06	–.05	–.02	.16	–.77	96.7%
14	Nissan (Japan)	Infiniti	.31	.26	.53	.22	.29	.00	.18	89.1%
15	Isuzu (Japan)	Isuzu	.01	.19	.44	.08	–.19	.09	.77	97.7%
16	Jaguar Land Rover (United Kingdom)	Jaguar	.17	.23	.59	.03	–.19	.05	.58	83.4%
17	Chrysler (United States)	Jeep	.26	.66	.41	.15	.24	–.18	–.20	87.3%
18	Hyundai_Kia (Korea)	Kia	.57	.36	.08	–.04	.14	.17	–.71	96.3%
19	Jaguar Land Rover (United Kingdom)	Land Rover	–.09	.24	.35	.23	–.23	.16	.56	49.5%
20	Toyota (Japan)	Lexus	.29	.24	.26	.19	.73	–.07	–.09	95.5%
21	Ford (United States)	Lincoln	.34	.57	.20	.24	.26	–.02	.13	74.1%
22	Mazda (Japan)	Mazda	.48	.42	.48	–.03	.06	–.35	.02	98.2%
23	Daimler (Germany)	Mercedes-Benz	.22	.37	.78	.13	.01	.09	–.04	84.0%
24	Ford (United States)	Mercury	.42	.61	.25	.17	.25	–.15	.14	92.7%
25	BMW (Germany)	Mini	.92	–.03	.07	.13	–.27	–.04	.41	96.0%
26	Mitsubishi (Japan)	Mitsubishi	.34	.33	.50	.04	–.13	.01	.51	94.5%
27	Nissan (Japan)	Nissan	.64	.43	.19	.21	.19	–.08	–.04	94.9%
28	GM (United States)	Pontiac	.45	.19	–.01	.10	.30	–.11	.70	94.5%
29	Porsche (Germany)	Porsche	.01	.33	.94	.03	–.02	.03	–.11	81.4%
30	Chrysler (United States)	Ram	–.19	.49	.15	–.20	–.61	.06	–.08	74.9%
31	GM (United States)	Saab	.20	.13	.25	.24	.06	.07	.82	92.8%
32	GM (United States)	Saturn	.59	.12	–.03	.20	.22	–.34	.65	95.7%
33	Toyota (Japan)	Scion	.60	–.02	.20	.08	.24	.11	.43	95.0%
34	Subaru (Japan)	Subaru	.57	.48	.24	.07	–.20	.51	.05	89.1%
35	Suzuki (Japan)	Suzuki	.56	.14	–.07	–.01	.00	.03	–.69	92.3%
36	Toyota (Japan)	Toyota	.63	.30	.15	.09	.46	–.08	–.09	94.3%
37	Volkswagen (Germany)	Volkswagen	.61	.46	.32	–.26	–.14	.14	–.37	91.7%
38	Ford (United States)	Volvo	.29	.37	.42	.10	.05	.07	.51	90.9%

the same way as those from standard FA, labeling each factor according to which observed time series have the largest corresponding loadings. The larger the absolute value of the loading L_{gh} (i.e., the gth row and hth column in Table 2), the stronger is the correlation between the gth time series and the hth factor. Series with strong co-movements will result in large loadings on the same subset of factors, revealing the trends they share.

Because the observed time series y_t are standardized, each entry in Table 2 indicates that for every one unit increase in the hth factor, the gth time series would change by L_{gh} standard deviations. Thus, if the absolute value of L_{gh} is greater than that of L_{kh}, we can conclude that factor h is more correlated with time series g than with series k. To label each of the seven columns in Table 2, the analyst needs to identify the rows with the largest absolute values and a theme that can somehow tie the corresponding vehicle makes together. Admittedly, such a task is inherently a subjective exercise: Different analysts observing the same pattern of loadings could potentially reach different conclusions. However, the subjective exercise of labeling each factor could potentially stimulate creative thinking, especially when there are patterns that may appear counterintuitive or unexpected at first blush but ultimately make sense post hoc. The loadings reported in Table 2 led us to the following interpretations and labels for the seven latent dynamic factors identified by our model:

- Factor 1:"Foreign Mass," with largest loadings on Honda, Nissan, Toyota, Volkswagen and, somewhat unexpectedly, Mini and Scion, indicating that although conventional wisdom would suggest that the latter two may be considered niche players, empirical results show they experience the same ebbs and flows in consumer interest as the most popular mass-market foreign makes.
- Factor 2: "U.S. Mass," with largest loadings on mass-market domestic makes: Chevrolet, Ford, Mercury, Chrysler, Dodge, and Jeep. Lincoln, the luxury make of Ford, also has a relatively large loading.
- Factor 3: "Euro Lux," with largest loadings on luxury European makes: Porsche, Mercedes-Benz, BMW, and Jaguar. Infiniti, Nissan's luxury make, which is known to have a European style, also has a relatively large loading.
- Factor 4: "GM Survived," with largest loadings on Buick, Cadillac, and GMC, three of the four makes that survived

GM's recent portfolio restructuring. The fourth survivor, Chevrolet, has a much larger loading on Factor 2, "U.S. Mass."
- Factor 5: "Lexus/(Ram)," with large positive loading on Lexus, but large negative loading on Ram, suggesting, somewhat unexpectedly, that consumer interest in these two makes tend to move in opposite directions.
- Factor 6: "Subaru/(Mazda and Saturn)," with relatively large positive loading on Subaru but large negative loadings on Mazda and Saturn, suggesting that consumer interest in Subaru moves in the opposite direction of Mazda and Saturn.
- Factor 7: "GM Canceled and Isuzu/(Hyundai, Kia, and Suzuki)," with large positive loadings on Saab, Hummer, Pontiac, and Saturn (the four makes dropped in GM's recent portfolio restructuring) and Isuzu but large negative loadings on Hyundai, Kia, and Suzuki, suggesting that the decline in consumer interest for the four lackluster makes of GM and Isuzu has been accompanied most closely by increase in interest for relatively low-priced Asian makes.

Taken together, these results suggest that our SDFA model has done what a typical exploratory factor analysis is designed to achieve: uncovering and quantifying intuitive as well as somewhat unexpected patterns of covariation. For example, the emergence of three unambiguous main factors—Foreign Mass, U.S. Mass, and Euro Lux, each anchored by makes with similar, well-established positions—provides face validity for our model. Moreover, although it may not be obvious a priori, our model's ability to separate GM Survived from GM Canceled and band together Hyundai and Kia, two makes from the same parent company with a lingering image of low-priced Korean imports, provides further face validity for our model. Factor 6, Subaru/(Mazda and Saturn), may be unexpected a priori but not surprising in hindsight, as it is conceivable that these three makes may indeed attract the same niches of consumers. It is a bit puzzling to see Factor 5, with Lexus and Ram anchoring opposite ends, indicating that gains in consumer interest in the luxury Japanese make are often accompanied by losses for the U.S. truck make.

Despite the distinctions discussed previously, it is important to note that for any given row in Table 2, it is rarely the case that each observed series is tied exclusively to a single factor. The number of loadings per row with absolute value greater than .2 has a minimum of 2 and a median of 4, indicating that the movements of many series are jointly shaped by multiple underlying factors. For example, although Infiniti loads relatively heavily on Euro Lux (.53), it also has a lighter but nevertheless significant loading on Foreign Mass (.31).

After interpreting and labeling the latent factors based on the estimated loadings, the next step is to focus on the estimated dynamic factor scores (f_t). As we detail in Web Appendix A (www.marketingpower.com/jmr_webappendix), various components of f_t ($= \alpha_t + \gamma_t$) can be derived through a Kalman filter and smoother, given the observed data (y_t from t = 1 through T) and model parameters (B, L, Σ_u, Σ_ε, Σ_η, Σ_ζ, and Σ_ξ in Equations 1, 2, and 2.1–2.4). The best way to convey the information contained in f_t is to chart its seasonal (γ_t) and nonseasonal components (α_t, β_t, and δ_t) over time, as in Figures 2–4.

Figure 2 plots the seasonal component (γ_t) for each latent factor. Given that the data are monthly, the number of seasons, s, is set to 12. Although our SDFA model allows seasonal patterns to evolve over time, our results show that this component is remarkably stable during the observation window. The first three latent factors (i.e., Foreign Mass, U.S. Mass, and Euro Lux) show considerably stronger seasonality than the other four. In addition, the seasonal patterns for the first three factors are not only strong but also distinctive from one another. While Foreign Mass shows consumer interest peaking in May–July and bottoming in December, U.S. Mass shows two peaks, one in January and the other in July, and one valley in November. Euro Lux shows a similar seasonal pattern to Foreign Mass, except that the bottoming occurs in January instead of December.

After partialing out the seasonal component (γ_t), the nonseasonal component (α_t) reveals the genuine trend line that is often of more interest, which we plot in Figure 3 with a solid line. Given Equation 2.1, which we can rewrite as $\alpha_t = \alpha_0 + \Sigma_{k=0}^{t-1}\beta_k + \Sigma_{k=1}^{t}\varepsilon_k$, the nonseasonal trend can be viewed as an initial level (α_0) plus the cumulative of slopes ($\Sigma_{k=0}^{t-1}\beta_k$) and the cumulative of random shocks ($\Sigma_{k=1}^{t}\varepsilon_k$). By removing the random shocks, which are more volatile than the slopes, we produce a smoother and easier-to-read trend line ($\alpha_0 + \Sigma_{k=0}^{t-1}\beta_k$), which appears as a dashed line in Figure 3.

According to the deseasonalized and smoothed trend lines in Figure 3, we can tell that (1) consumer interest in Foreign Mass peaked in approximately May 2008 and declined fast ever since, despite a small bump in the second half of 2009, and (2) U.S. Mass peaked in mid-2009 and trended downward ever since. The direction of these two Mass trend lines did not bode well for a quick turnaround of the U.S. economy in general and the automobile market in particular. Coincidentally, according to the August 2011 report from Thomson Reuters/University of Michigan Surveys of Consumers, consumer confidence dropped by 19.2% from August 2010 to August 2011, accompanied by similar declines in buying plans for vehicles and other household durables. Furthermore, we note that the trend line of Foreign Mass is clearly distinct from that of U.S. Mass, suggesting that declines of interest in one are not necessarily accompanied by increases in the other, contrary to what conventional wisdom might take for granted (e.g., Datamonitor 2005).

Figure 3 also shows that consumer interest in Euro Lux bottomed in the second half of 2009 and trended upward afterward. Notably, although there is a distinct trend line for Euro Lux, our model did not indicate a distinct trend line for either American Lux or Japanese Lux, suggesting that it can be overly simplistic to assume that there exists a single overarching trend governing U.S. consumers' interest in luxury vehicles (e.g., *The Wall Street Journal* 2008).

Finally, Figure 3 indicates that (1) consumer interest in GM Survived has remained more or less stable in the past few years, after experiencing increases from 2004 through 2007; (2) Lexus/(Ram) peaked in 2006 and has sustained continued drops since that time; (3) Subaru/(Mazda and Saturn) went through a large swing and has trended up in the past few years; and (4) GM Cancelled and Isuzu was on a declining path during the entire observation window, coinciding with the ascendance of Hyundai, Kia, and Suzuki. Note that although the rise of Hyundai and Kia in the United States is a well-recognized trend in recent years (e.g., Buss 2011), our results indicate that the growth of consumer interest in these two Korean makes has been accompanied

Figure 2
SEASONAL COMPONENT (γ_t) OF THE LATENT DYNAMIC FACTOR

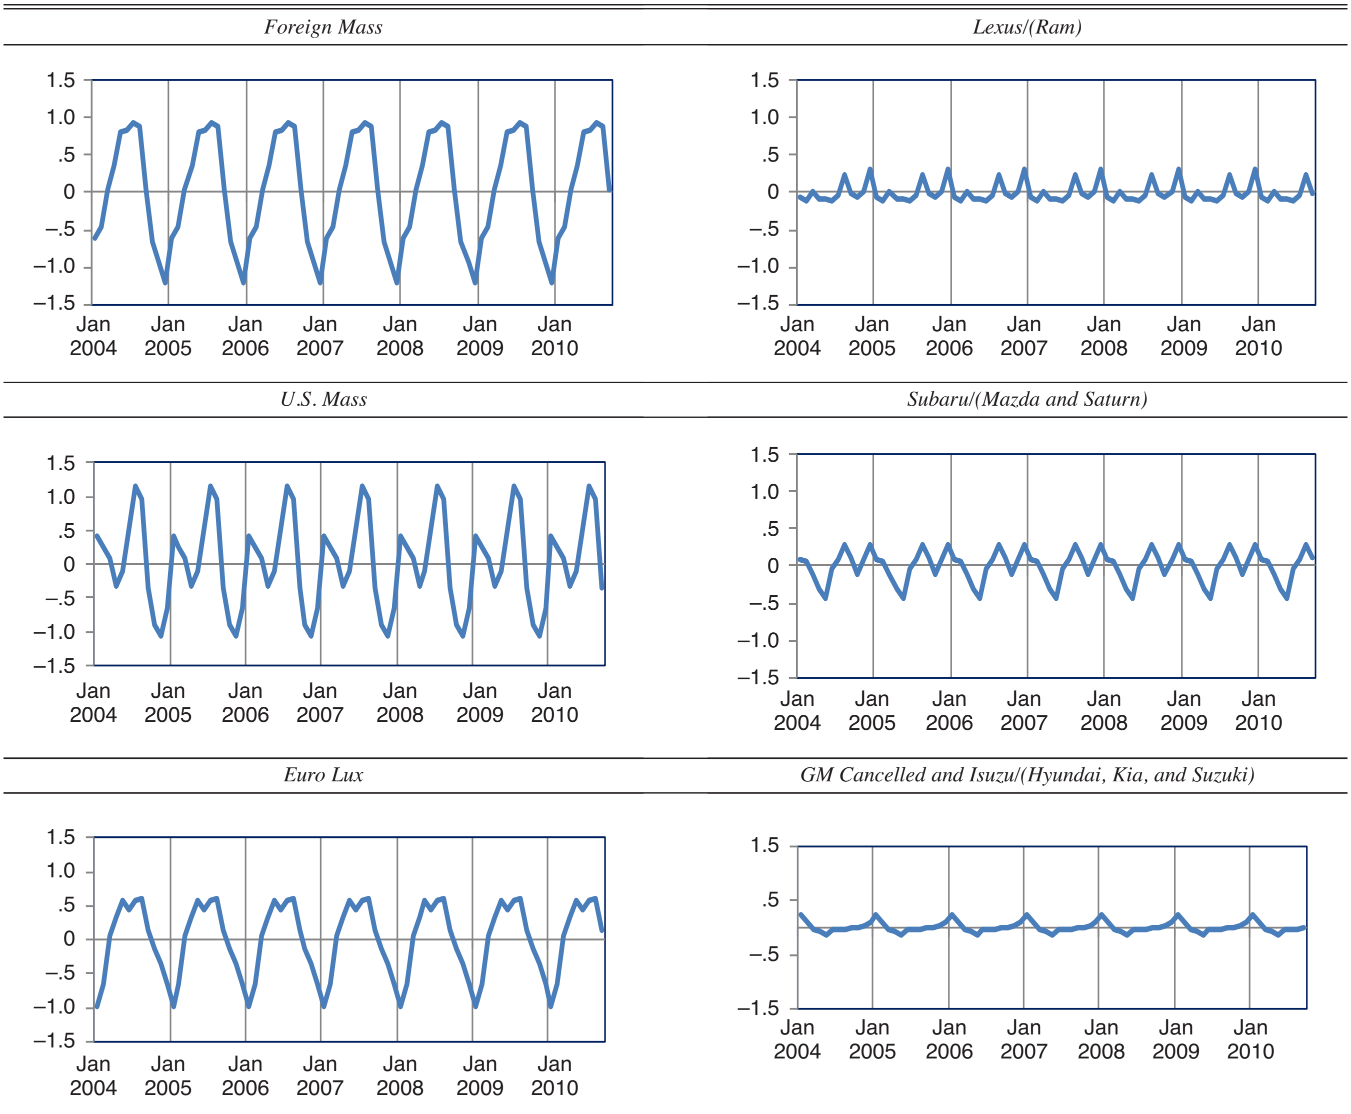

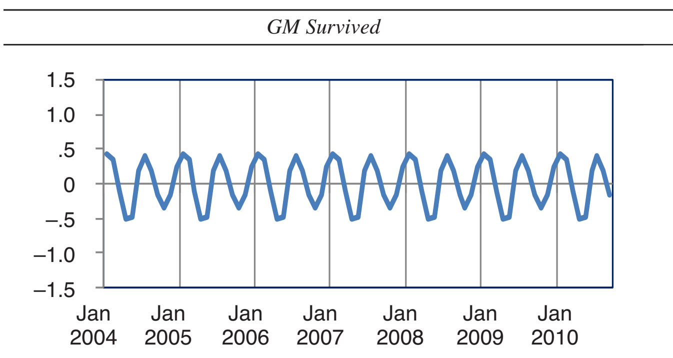

most closely by decreasing interest in GM Cancelled and Isuzu.

Figure 4, Panel A, plots the slope (β_t) over time, which can be viewed as the first derivative of the deseasonalized and smoothed trend (the dashed line in Figure 3). Through β_t, we can examine the direction (sign of β_t) and speed (size of β_t) of each trend line on a finer scale, obtaining a clearer view of the periods of growth ($\beta_t > 0$) and decline ($\beta_t < 0$). For example, Figure 4, Panel A, indicates that Euro Lux went through two periods of decline, from early 2004 to early 2005 and from early 2007 to mid-2009. Figure 4, Panel B, charts the slope change rate (δ_t), which can be viewed as the second derivative of each deseasonalized and smoothed trend line. When β_t and δ_t are of the same sign, the trend is accelerating; otherwise, it is decelerating.

Two pieces of information from Figure 4, Panels A and B—the slope (β_t) and its change rate (δ_t) at the last period of the observation window—should be of particular interest to the analyst, because they can be used to extrapolate the trend lines h-period ahead according to Equation A4.1 in the Web Appendix (www.marketingpower.com/jmr_webappendix) (i.e., $\alpha_{T+H} = \alpha_T + \beta_T h + \delta_T[h(h-1)/2]$). Take Foreign Mass as an example: Toward the end of the observation window, the reading of β_T from Figure 4, Panel A, is approximately –.2, and δ_T from Figure 4, Panel B, is approximately –.005, suggesting that Foreign Mass, as of September 2010, is trending downward (a negative slope) at an accelerating speed (a negative change rate for a negative slope). In con-

Figure 3

NONSEASONAL TREND COMPONENT α_t (SMOOTHED AND UNSMOOTHED) OF THE LATENT DYNAMIC FACTOR

Figure 3
CONTINUED

GM Survived

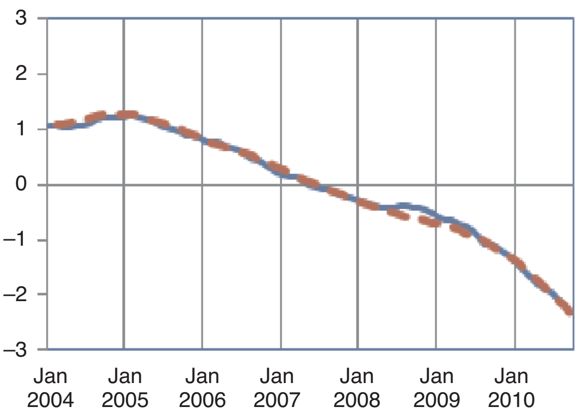

trast, Figure 4, Panels A and B, suggests that, at the end of the observation window, Euro Lux has an upward trend (β_T of approximately .1) that is gaining momentum (δ_T of approximately .005).

Testing Trend Projections

To ascertain how well the uncovered trend lines can be extrapolated into the near future, we (1) refit the SDFA model using data from January 2004 to May 2007, (2) infer the latent trend lines with the Kalman filter and smoother detailed in Web Appendix A (www.marketingpower.com/jmr_webappendix), and (3) predict the nonseasonal trend component α_t in November 2007 (i.e., a six-step-ahead forecast) using Equation A4.1 (see www.marketingpower.com/jmr_webappendix). From that point onward, with each additional month of data, we repeat the preceding steps, extending the target date of the six-step-ahead forecast by one month. This enables us to evaluate the ability of the SDFA model to extrapolate the latent trend lines into the near future. Figure 5 compares the smoothed trend lines of the nonseasonal component α_t based on complete data from January 2004 to September 2010 (i.e., $\alpha_{t|T}$) against projections of α_t using only data available six months earlier (i.e., $\alpha_{t|t-6}$). Figure 5 shows that, overall, the two matched well, suggesting that the SDFA model can not only uncover latent common trend lines from high-dimensional time series but also extrapolate them into the near future with reasonable reliability.

While the tests displayed in Figure 5 attest to the reliability of our SDFA model's trend projections, they do not ascertain the reliability and validity of the trends themselves. Next, we present two tests of predictive and face validity of these trends. Web Appendix C (www.marketingpower.com/jmr_webappendix) presents a third test verifying the split-half reliability of the dynamic factor structure.

Testing Predictive Validity

Although our seven-factor SDFA model explained 89.4% of variance observed in the 38 time series, it could nevertheless be a result of overfitting. If that is the case, the model should do less well in out-of-sample forecasting of the observed time series. To test such predictive validity, we recalibrated the model with data from January 2004 to September 2009. We used the recalibrated model to forecast each latent factor f_t, which we then combined through the loadings matrix L to arrive at forecasts for each individual time series y_t in the next 12 months (i.e., October 2009–September 2010). As benchmarks, we use three models. The first is series-specific ARIMA, identified and fit to each time series using the first 69-month data, which was then used to forecast for the next 12 months. The second is VAR(1) with first differencing to remove nonstationarity in the original series (i.e., $\Delta y_t = A\Delta y_{t-1} + e_t$). The third is Bayesian VAR(1), for which, following Carriero, Kapetanios, and Marcellino (2009), we impose a driftless random walk prior and set the shrinkage parameter θ to .0001, so that cross-series dynamics enter into the picture only if there is overwhelming evidence of them in the data.

We evaluate out-of-sample predictive performances in terms of mean absolute error (MAE) and mean absolute percentage error (MAPE). Across the 38 series, our proposed model, SDFA, produced an overall MAE (MAPE) of .16 (41%), compared with .26 (67%) for ARIMA, .34 (88%) for VAR(1), and .27 (71%) for Bayesian VAR(1). Comparing out-of-sample fit for each individual series, SDFA outperformed ARIMA, VAR(1), and Bayesian VAR(1) in, respectively, 29, 37, and 33 of 38 cases. To provide a more concrete sense of in-sample as well as out-of-sample fit, Figure 6 plots eight selected series, consisting of actual data denoted by smoothed lines and SDFA's predictions by solid squares. These eight series ranked, respectively, 1st, 2nd, 13th, 14th, 25th, 26th, 37th, and 38th in terms of out-of-sample MAE, covering the full range of fit.

Taken together, we view the preceding out-of-sample performances as a strong sign of predictive validity, suggesting that SDFA is not an overfitted model for our application. Its remarkable improvement in predictive fit over the benchmarks comes mainly from two sources: (1) information sharing across series such that forecasts for any individual series are informed by other series according to their historical co-movement patterns and (2) the parsimonious factor-analytic structure and thus less overfitting of noise in the data. That said, we caution against treating SDFA as a device designed for multivariate time series *forecasting*, because it is not what we intended with our framework for quantitative trendspotting. When the goal is purely to achieve the best out-of-sample predictive fit, ARIMA, standard or Bayesian VARMA/VAR, and other parametric or nonparametric methods may well be more suitable forecasting tools than dynamic factor models such as ours. However, what SDFA can and other methods cannot do is to uncover common trend lines hidden behind high-dimensional time series and decompose those trend lines into directly interpretable seasonal and nonseasonal components.

Testing Face Validity

Given that the latent trends uncovered by our SDFA model are supposed to represent the predominant trajectories in consumer interest for all the major automobile makes in the United States, we would expect these trends to have some explanatory power on vehicle sales over time. In other words, relating these trends to sales should provide a test for both face validity and managerial relevance. To conduct

Figure 4
NONSEASONAL TREND COMPONENTS

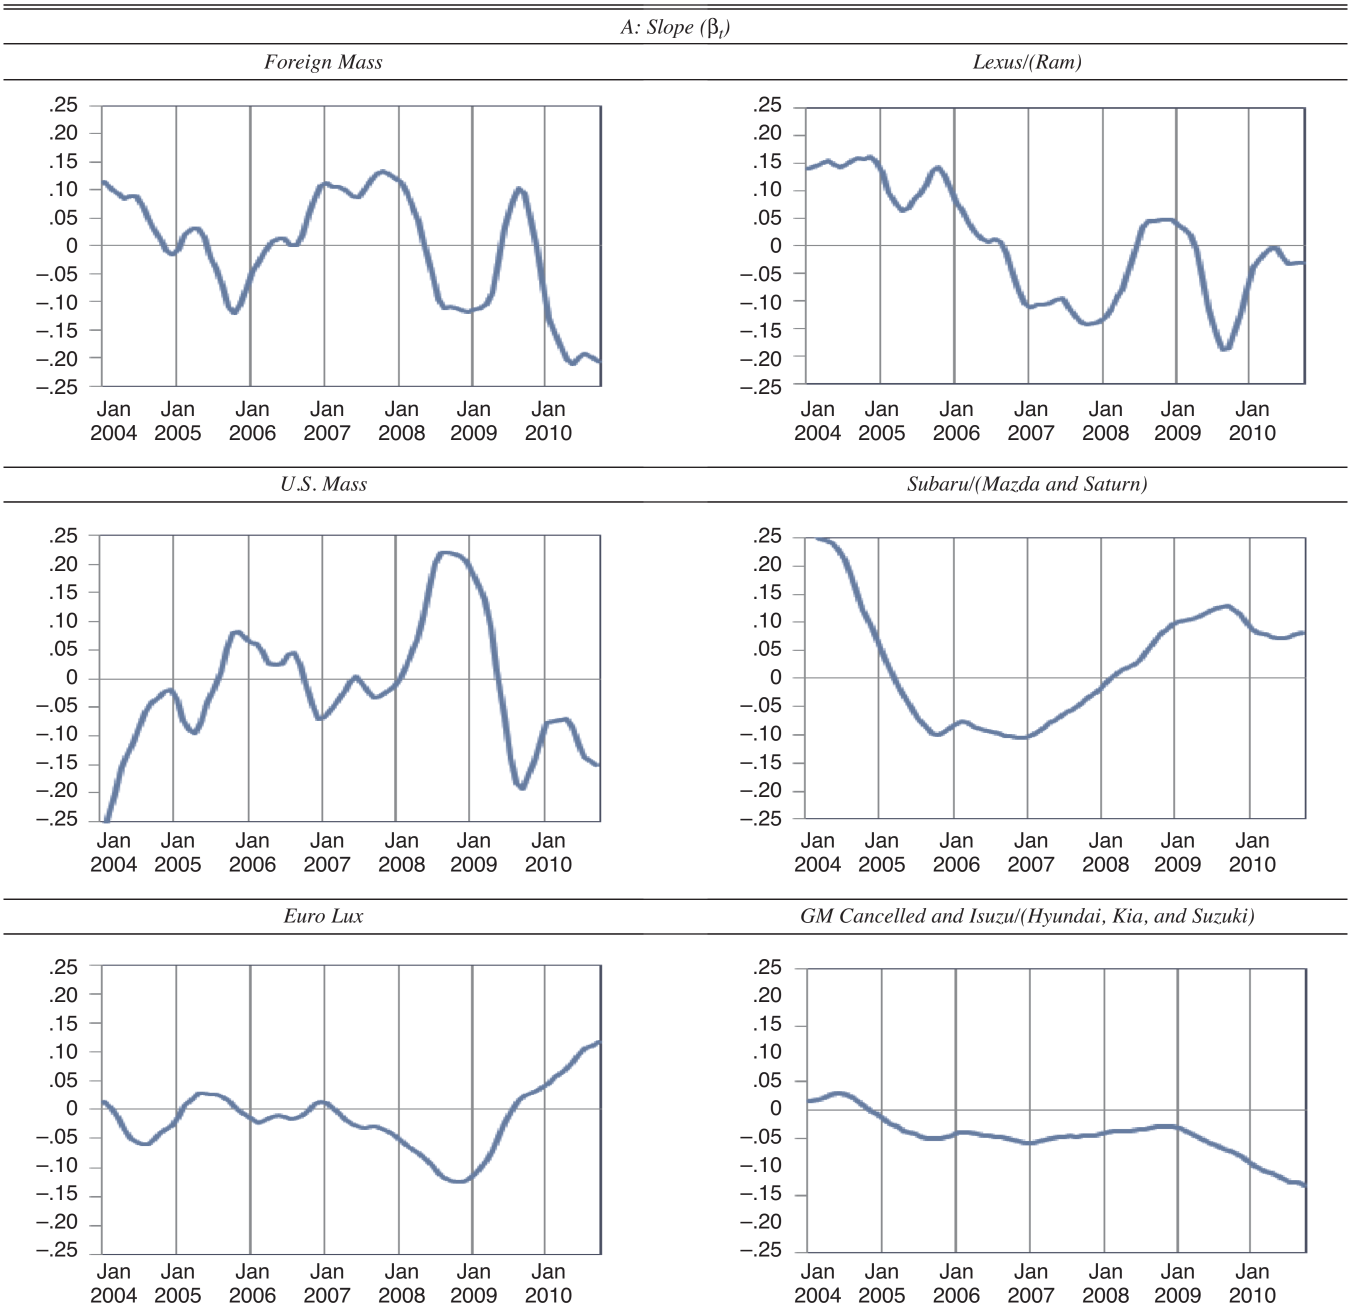

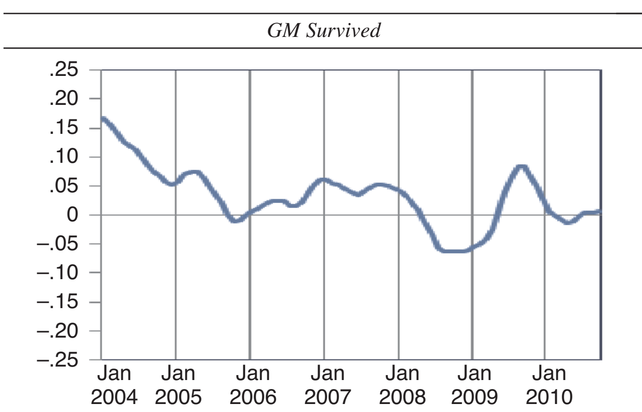

such a test, we gathered from *Automotive News* monthly make-level new vehicle unit sales. Then, we regressed the log-transformed monthly sales for each of the 38 makes against the seven latent trends and compared the R-squares with an alternative in which a given make's log-transformed monthly sales is regressed against its log-transformed search index (i.e., only one explanatory variable). Table 3 reports the results, which make two things clear.

First, Google search indexes are clearly but *imperfectly* related to new vehicle sales. (The R-squares from the regressions of each make's sales against its own search index have a mean of 31.6%, median of 27.1%, minimum

Figure 4
CONTINUED

B: Change Rate (δ_t)

Foreign Mass

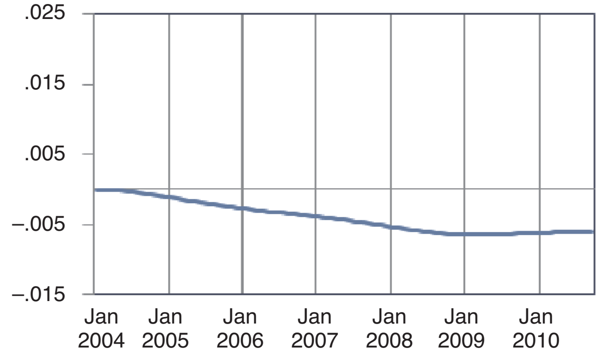

Lexus/(Ram)

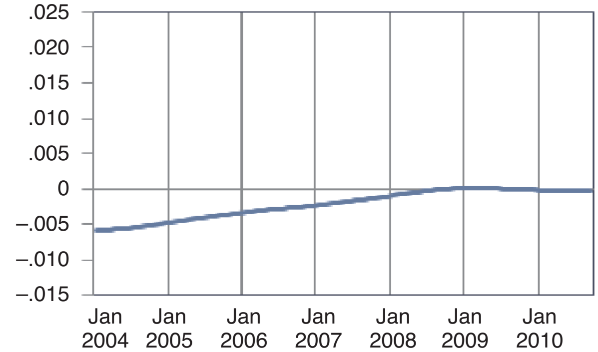

U.S. Mass

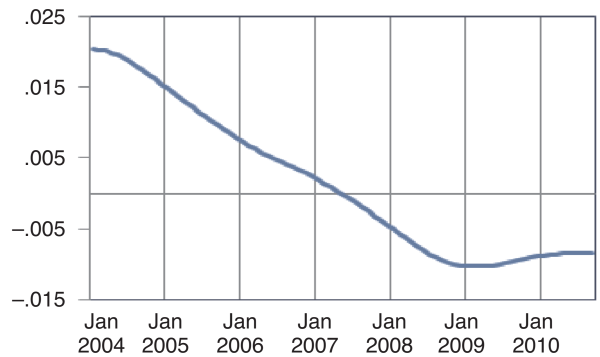

Subaru/(Mazda and Saturn)

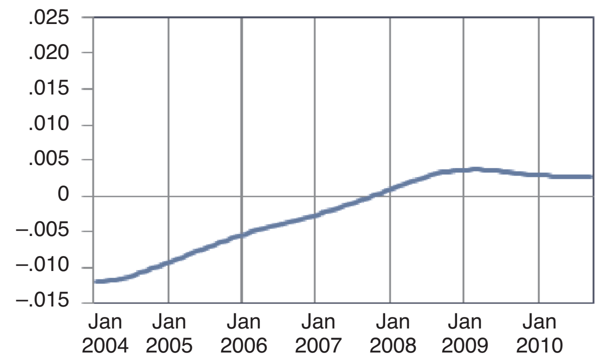

Euro Lux

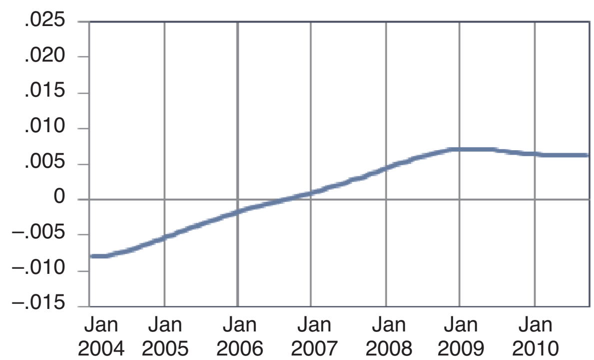

GM Cancelled and Isuzu/(Hyundai, Kia, and Suzuki)

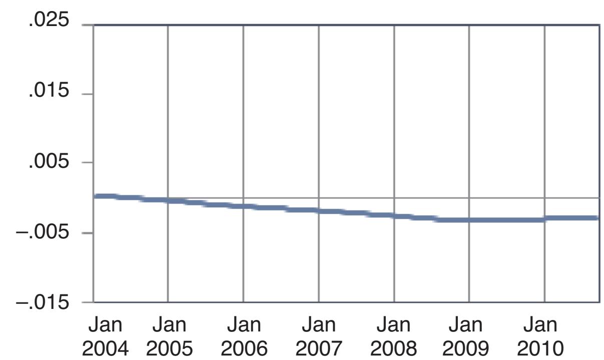

GM Survived

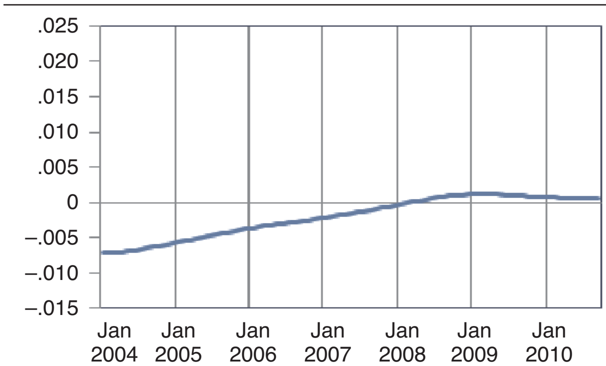

of .4%, and maximum of 81.9%.) The finding that the relationship between search and sales can be tenuous for many makes is not surprising; consumer online searches should only reflect interest/consideration at the early stages of the purchase funnel, while a myriad of other factors can influence final decisions.

The second and more important takeaway from Table 3 is that by synthesizing information contained across all 38 search indexes, through the seven latent trends uncovered by our model, the collective explanatory power of the search indexes increased dramatically: The R-squares from the regressions of each make's sales against the seven latent

Figure 5
LATENT TRENDS ($\alpha_{t|T}$) VERSUS PROJECTIONS OBTAINED SIX MONTHS EARLIER ($\alpha_{t|t-6}$)

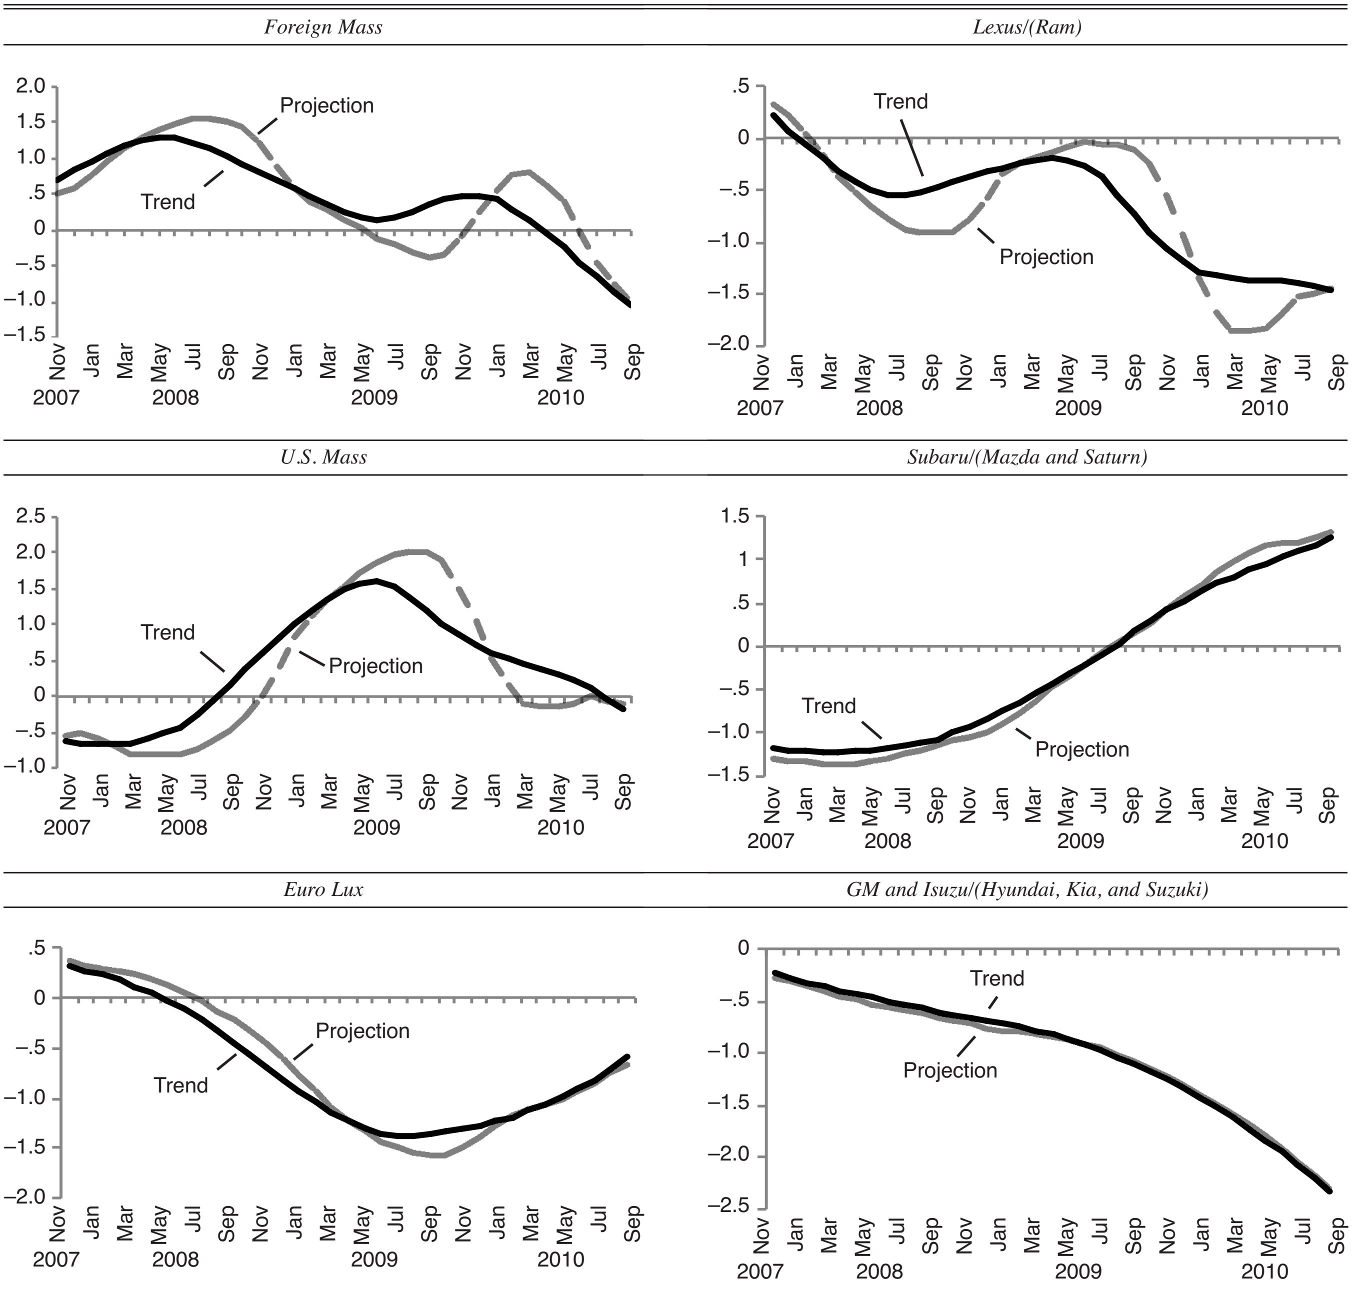

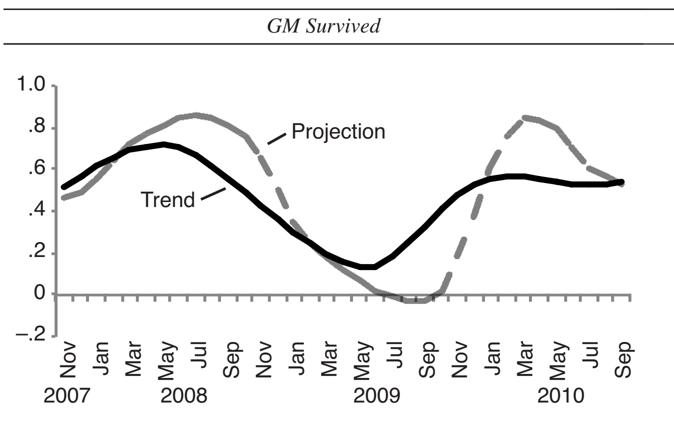

trends have a mean of 73.9%, median of 77.0%, minimum of 44.9%, and maximum of 93.1%. The high R-squares aside, it is also reassuring, from a face validity standpoint, to note that the Euro Lux search trend was the strongest predictor of sales for BMW, Jaguar, Mercedes-Benz, and Porsche and that the GM Survived trend was closely tied to the sales of Buick, Cadillac, and GMC.

We must admit that the remarkable explanatory power of the seven latent trends on sales exceeded our original expectation; its generalizability in contexts other than vehicle shopping deserves more systematic investigation. Nevertheless, our basic conclusion should hold in other contexts; that is, although individual online search indexes can be inform-

Figure 6

IN-SAMPLE (FIRST 69 MONTHS) AND OUT-OF-SAMPLE (LAST 12 MONTHS) FIT

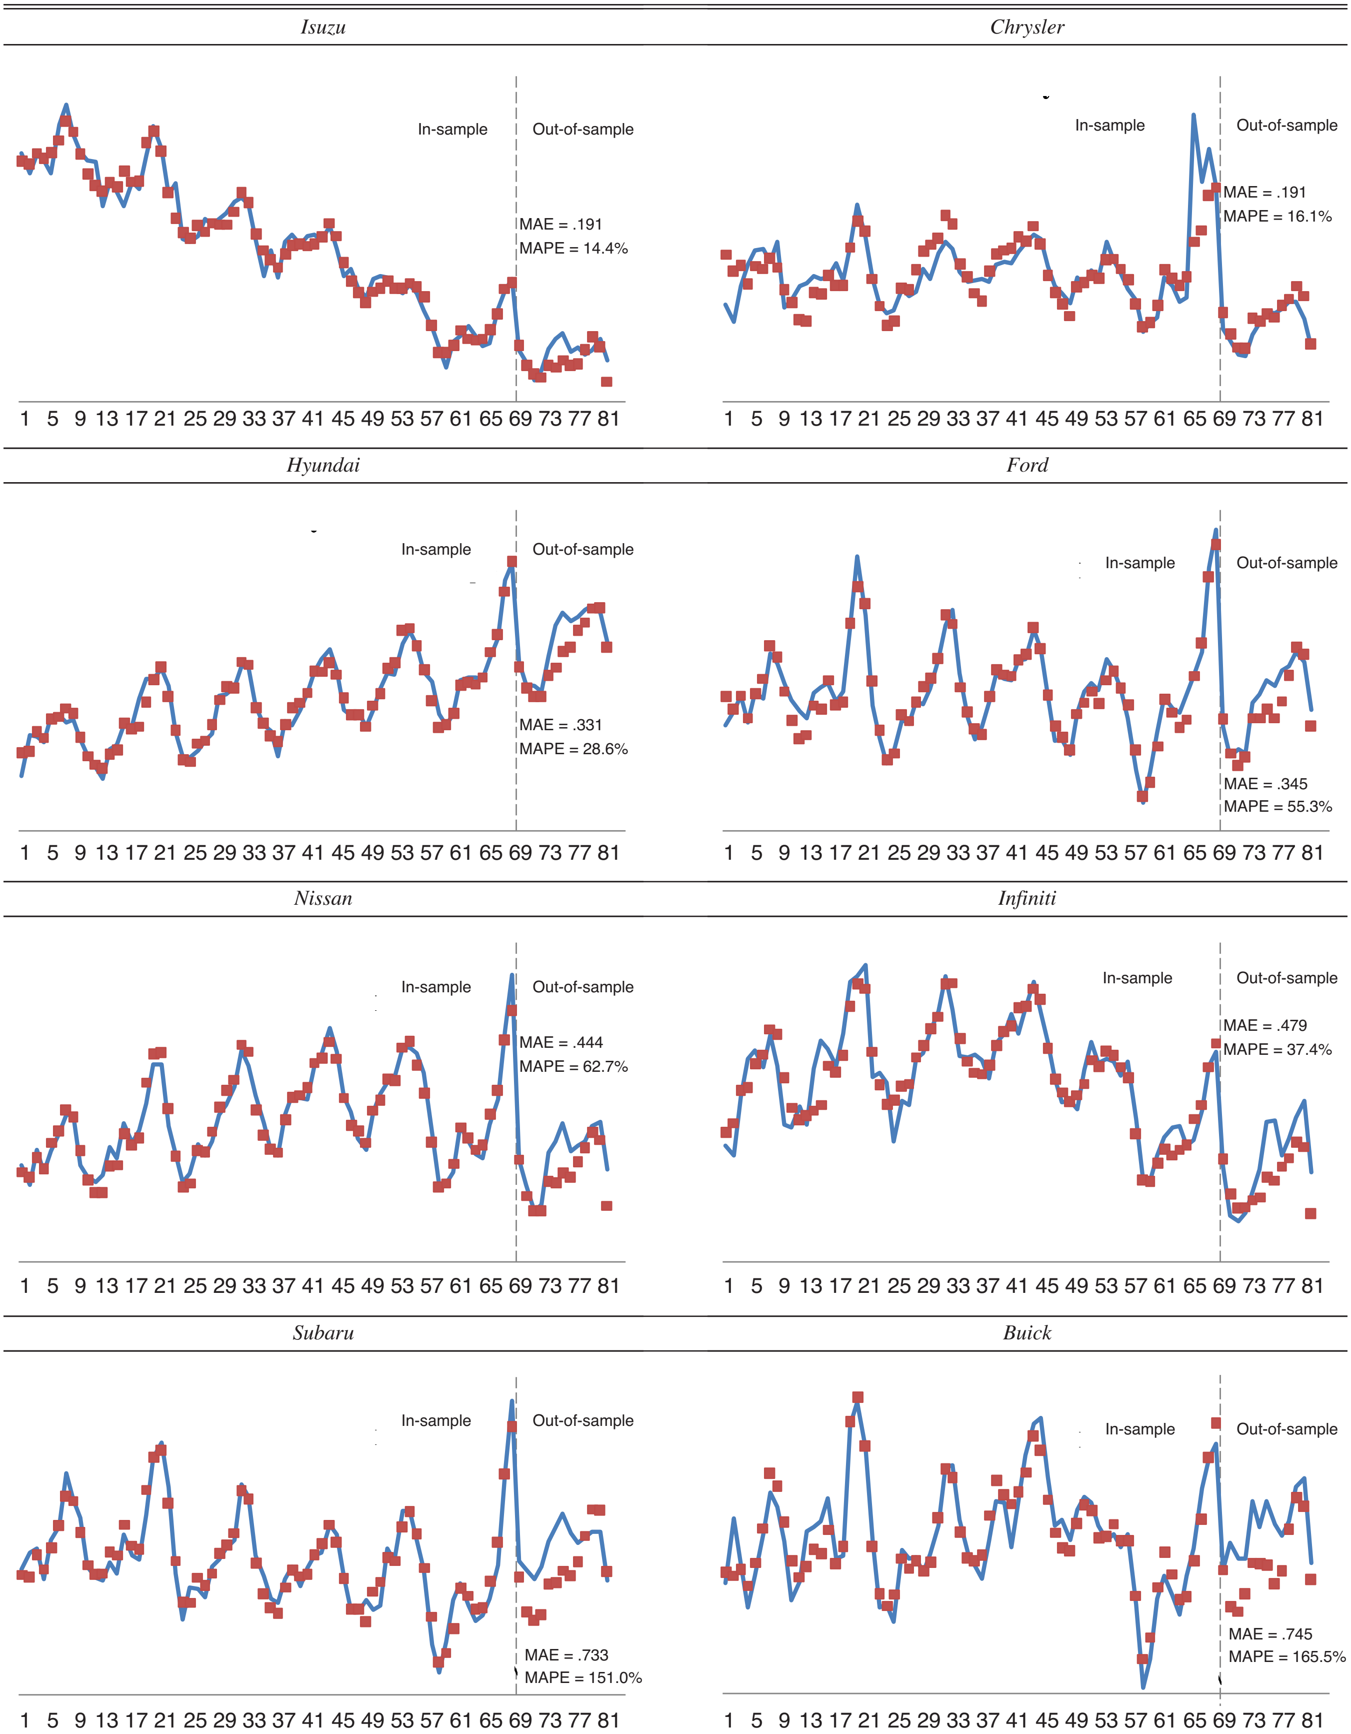

Table 3
R^2 OF INDIVIDUAL SEARCH INDEX VERSUS LATENT TRENDS IN EXPLAINING SALES

Dependent Variable (Sales)	*R^2 with Different Explanatory Variables* Individual Search	Latent Trends	Improvement
Ram	5.5%	93.1%	87.6%
Chrysler	2.8%	88.3%	85.5%
Cadillac	.4%	83.2%	82.8%
Suzuki	5.6%	87.3%	81.7%
Chevrolet	3.6%	82.6%	79.0%
Buick	8.3%	86.9%	78.6%
GMC	.6%	79.1%	78.5%
Jeep	7.1%	80.2%	73.1%
Ford	12.1%	81.2%	69.1%
Land Rover	3.3%	67.8%	64.5%
Mercury	12.5%	76.8%	64.2%
Lincoln	1.5%	60.8%	59.3%
Dodge	7.4%	65.8%	58.3%
Hyundai	25.7%	69.4%	43.7%
Audi	10.3%	52.8%	42.5%
Mercedes-Benz	12.3%	50.8%	38.5%
Jaguar	51.1%	87.4%	36.2%
Porsche	28.5%	62.6%	34.1%
Saturn	48.1%	80.8%	32.7%
Lexus	28.0%	59.7%	31.7%
Infiniti	36.5%	67.5%	31.0%
Mini	15.7%	44.9%	29.2%
Nissan	26.2%	54.8%	28.6%
Subaru	25.8%	54.0%	28.2%
Mitsubishi	58.4%	85.2%	26.7%
Volkswagen	35.8%	60.9%	25.1%
BMW	45.4%	70.5%	25.0%
Volvo	62.2%	86.1%	23.9%
Acura	64.4%	86.2%	21.7%
Honda	54.7%	75.2%	20.5%
Toyota	46.8%	67.1%	20.2%
Mazda	38.9%	58.2%	19.3%
Pontiac	58.1%	77.2%	19.1%
Kia	48.5%	65.8%	17.3%
Saab	73.7%	89.8%	16.1%
Hummer	74.8%	89.3%	14.5%
Isuzu	78.3%	89.5%	11.3%
Scion	81.9%	90.3%	8.4%

ative of sales by themselves (an average R-square of 31.6% in our test), these indexes can be far more powerful when key common underlying trends are distilled from them and used in subsequent analyses (an average R-square of 73.9% in our test). The challenge of practicing such a "look at everything together" approach lies in the curse of dimensionality, an issue that our SDFA model can potentially turn into an advantage.

Understanding What May Have Shaped the Uncovered Latent Trends

Next, we try to better understand the nonseasonal trend lines (i.e., α_t) in terms of major events and variables that may have shaped them. Figure 7 displays the "shocks" (ε_t) to the level of the deseasonalized trend lines. A naked-eye inspection would reveal a small number of visible major shocks. Searching through the archives of Google News from approximately the time of these shocks, we identified two major events that coincided with several of them. In June 2005, General Motors offered its generous employee discounts to any buyer during that month. During approximately the same period, there is a large positive shock in consumer interest for U.S. Mass and GM Survived. The subsequent negative shock, especially in GM Survived, indicates that the increase in consumer interest was only transient, disappearing as soon as the promotion was over, a possible sign of forward buying (Busse, Simester, and Zettelmeyer 2010). Similarly, on July 1, 2009, the U.S. government initiated a Cash for Clunkers program, which began to process claims on July 24 and ended on August 24. As might be expected, there are large positive shocks during July and August 2009 in consumer interest for Foreign Mass and U.S. Mass. These large positive shocks were followed by large negative shocks, again pointing to the short-lived impacts of cash incentives on consumer interest. Although there might be other events related to other shocks shown in Figure 7, these two examples illustrate how the outputs from our SDFA model can be used to reverse-engineer the events that may have caused large shocks in the uncovered trend lines.

It is evident that major interventions by a company or the government are not the only forces shaping trends in consumer interest. It would be expected that intention to buy a vehicle and search for a particular make would also be driven by other economic and market conditions. To address this issue, we must relate the nonseasonal trends (α_t) to exogenous variables, which requires a better model than the standard linear regression because the impacts of exogenous variables on the trends can extend through multiple periods (i.e., carryover effects), and exogenous variables may lead the trends by one or more periods (i.e., time delays). For these reasons, we identified and estimated the best transfer function (Box, Jenkins, and Reinsel 2008) for each deseasonalized trend (α_t) using the following exogenous variables: (1) unemployment rate (%), (2) Consumer Sentiment Index (Michigan Survey Research Center), (3) regular gasoline price (in U.S. dollars per gallon), (4) GM Employee Discount for all buyers (June 2005), and (5) Cash for Clunkers program (July–August 2009).

Table 4 displays the parameter estimates of the best transfer functions identified for each deseasonalized trend. The estimates for Foreign Mass indicate that increases in unemployment lead to decreases in this trend one month later (–.389), while increases in consumer sentiment translate into relatively small concurrent decreases in this trend (–.021). The effect of gas prices on Foreign Mass is more complex, with an increase in prices leading to a concurrent positive (.904) shift on this trend and a negative (–.618) shift one month later. The Cash for Clunkers program produced a positive shift of .977 in the Foreign Mass trend, followed by an opposite shift at the end of the program. Unfortunately, for the other trends (except GM Survived), it would be difficult to interpret the estimated transfer functions directly. Fortunately, the dynamics captured by these functions are much easier to interpret in the form of impulse response functions (Box, Jenkins, and Reinsel 2008), which represent the expected change in the deseasonalized trend in each time period, in response to a unit impulse in the exogenous variable at time zero. The impulse response functions displayed in Figure 8 suggest that the trends most affected by the exogenous variables are Foreign Mass, U.S. Mass, GM Survived, and, to a lesser extent, Euro Lux. As expected, the Cash for Clunkers program had a positive concurrent effect on Foreign Mass and U.S. Mass (see the

Figure 7
SHOCKS IN NONSEASONAL TREND (ε_t)

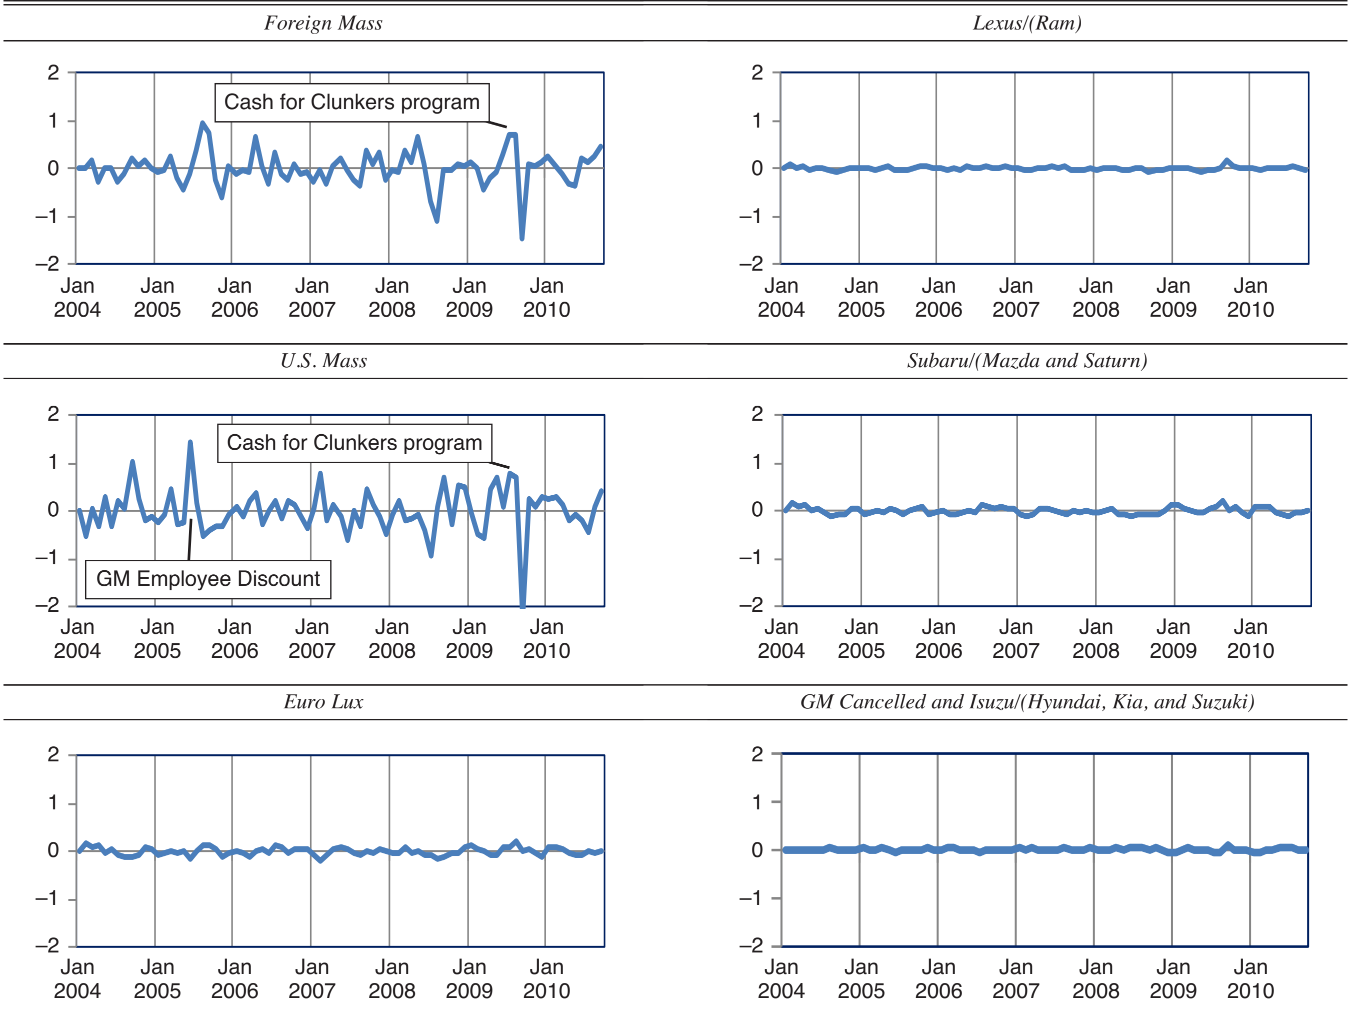

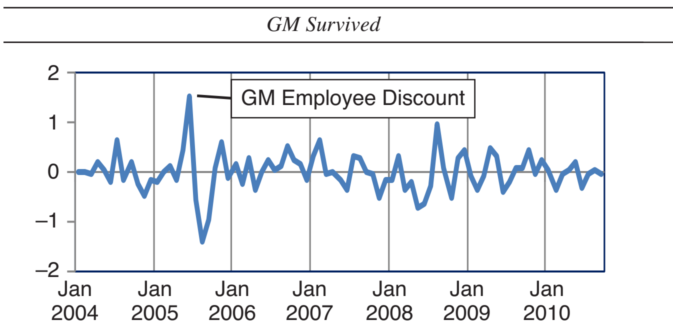

solid circles at time 0), while GM's decision to extend employee discounts to all buyers had a strong concurrent positive effect on GM Survived and U.S. Mass (see the solid squares at time 0). Unemployment had a negative effect on Foreign Mass, lagged by one month, and a complex distributed-lag effect on Euro Lux (see the solid diamonds at time 1 and afterward). Increases in gas prices raised concurrent interest in Foreign Mass, and decreased interest in U.S. Mass (see the solid triangles at time 0), and in both cases, the impulse response function shows a trough in the second month (see the solid triangles at time 2).

Identifying Impactful Events Affecting Individual Series (but Not the Common Trends)

In addition to helping identify the latent common trends and understand what may have shaped them over time, our SDFA model can also help tease out idiosyncratic shocks in each individual series, after partialing out shifts already accounted for by movements in the common factors. In our application, these idiosyncratic shocks (u_t) can be investigated to identify events that might affect consumer interest in a particular vehicle make, but not the common trends.

Figure 9 displays the standardized idiosyncratic shocks (μ_t) for 4 of the 38 individual series in our study. To illustrate, we focused on μ_t values that are at least three standard deviations in size and searched the archives of Google News around the time of these major idiosyncratic shocks to identify possible events behind them, which are indicated on top of the respective charts in Figure 9. This illustration shows how an analyst could use SDFA outputs to distinguish movements that are common to multiple indicators from those that are unique to an individual indicator and reverse-engineer the major events behind them accordingly. The opposite of the preceding analysis, in which the events

Table 4

TRANSFER FUNCTION MODELS RELATING DESEASONALIZED TRENDS TO EXOGENOUS VARIABLES

Estimated Model	R^2
$\Delta\text{ForeignMass}_t = -.389\Delta\text{Unemp}_{t-1} - .021\Delta\text{Csent}_t + (.904 - .618B)\Delta\text{GasPrice}_t + .977\Delta\text{CashClunk}_t + \varepsilon_t$	.503
$(1.933 - .821B)\text{USMass}_t = -.803\Delta\text{GasPrice}_{t-1} - 1.718\Delta\text{CashClunk}_t + .599\Delta\text{GMEmployee}_t + \varepsilon_t$	.791
$(1 - .558B)\Delta\text{EuroLux}_t = [-.157/(1 - 1.031B + .665B^2)]\Delta\text{Unemp}_{t-1} + \varepsilon_t$	.434
$\Delta\text{GM}_t = -.031\Delta\text{Csent}_t - .914\Delta\text{GMEmployee}_t + \varepsilon_t$	.283
$\Delta^2\text{LexusRam}_t = -.094\Delta^2\text{CashClunk}_t - .060\Delta^2\text{GasPrice}_t + \varepsilon_t$	.277
$\Delta^2\text{SubaMazSat}_t = (.143 + .117B^{10})\Delta^2\text{Unemp}_{t-5} + \varepsilon_t$	.254
$\Delta^2\text{GMIsuzu}_t = [-.058/(1 + .261B + .752B^2)]\Delta^2\text{Unemp}_{t-5} - .054\Delta^2\text{CashClunk}_t + (1 - .479B)\varepsilon_t$	.513

Notes: All parameters are statistically significant at the .01 level, and portmanteau (Q) tests could not reject (at the .01 level) the hypotheses that the residuals are white noise.

are known a priori and the μ_t are used to verify their unique impact on an individual series, would also be useful (if not more so) for learning whether a specific indicator has shifted in response to the events under study.

CONCLUDING REMARKS

In this article, we propose an SDFA model that can simultaneously analyze large panels of time series, distilling them into a few latent dynamic factors that isolate seasonal cyclic movements from nonseasonal nonstationary movements. Through its application in the context of quantitative trendspotting, we demonstrate how such a model can be used to uncover a small number of key common trend lines hidden behind the co-evolution of a large array of marketplace indicators, turning the curse of dimensionality into an advantage. Equipped with such a market-sensing tool, instead of relying too much on any individual indicator (and thus risking overreacting to its idiosyncratic ups and downs), managers can systematically synthesize information across multiple indicators, generating more reliable market intelligence by zeroing in on the predominant trajectories shared by these indicators.

It must be kept in mind that, as with virtually any data used by marketing researchers, the indicators used for quantitative trendspotting may contain two types of errors: random and systematic. Our method can filter random errors out because, by definition, they are idiosyncratic to individual indicators and will not manifest as co-movements across series, the source of identification for common trends. As for errors that may systematically affect multiple indicators, they will inevitably hinder the correct identification of trends. To guard against being misled by such systematic data errors, we have demonstrated how to validate the identified trend lines through various tests. For example, the estimated factor loadings could be examined to determine whether they show strong face validity. The outlier shocks to the common trends and individual series could be examined to determine whether they coincide with any major events. More formally, the identified trend lines can be related to other variables that are measured with accuracy (e.g., in our application, gas prices, vehicle sales) to determine whether there is any meaningful relationship.

In addition to offering a conceptual framework for quantitative trendspotting and a novel statistical model for implementing it, our study introduces a powerful dimension-reduction method for marketing researchers whose work requires the analysis of high-dimensional (tens or even hundreds) time series. With a DFA model such as ours, analysts no longer need to arbitrarily aggregate or preselect subsets of time series and can objectively assess common trends embedded in the entire data set. Although our application used data from a promising online consumer interest trending service—GIFS—it would be straightforward to apply our framework and SDFA model to time-series data from more traditional sources of marketing intelligence (e.g., tracking measures collected through repeated surveys, transaction data gathered by syndicated services or loyalty programs).

As directions for further research, we suggest several promising avenues. First, for firms operating in multiple markets, data for market sensing could come in a three-way mode—that is, multiple indicators tracked over multiple time periods across multiple markets. In such a case, it would be helpful if the model could tease apart trends that are common across markets from those that are market-specific, as a dynamic version of three-mode factor analysis.

In our empirical illustration, we used indicators available at the monthly level, and there was no missing information throughout the observation window. In practice, firms may track indicators at different frequencies (e.g., daily, weekly, monthly, quarterly, yearly), and not all indicators are available for the whole duration under study. In addition, in our application, we used indicators from a single source that are measured on the same scale. In practice, firms may monitor multiple sources (e.g., searches, blog posts, tweets) for indicators that are measured on different scales. Our model, as it is currently specified, can be directly applied across measures of different scales after the standardization of each time series. The real intriguing issue lies in blending indicators from different sources that reflect fundamentally different constructs (e.g., consumer searches indicating mere consideration vs. purchases indicating eventual choices).

Our current model formulation assumes stability in the latent factor structure, which a split-half reliability test showed to be valid in our application. However, major events might indeed change the entire co-movement pattern, in which case, the whole model would need to be recalibrated. The possible existence of these situations, especially in the long run, implies that robustness checks such as split-half reliability tests are needed. When correlation patterns (and therefore estimated factor loadings structure) vary sub-

Figure 8
IMPULSE RESPONSE FUNCTIONS RELATING DESEASONALIZED TRENDS TO EXOGENOUS VARIABLES

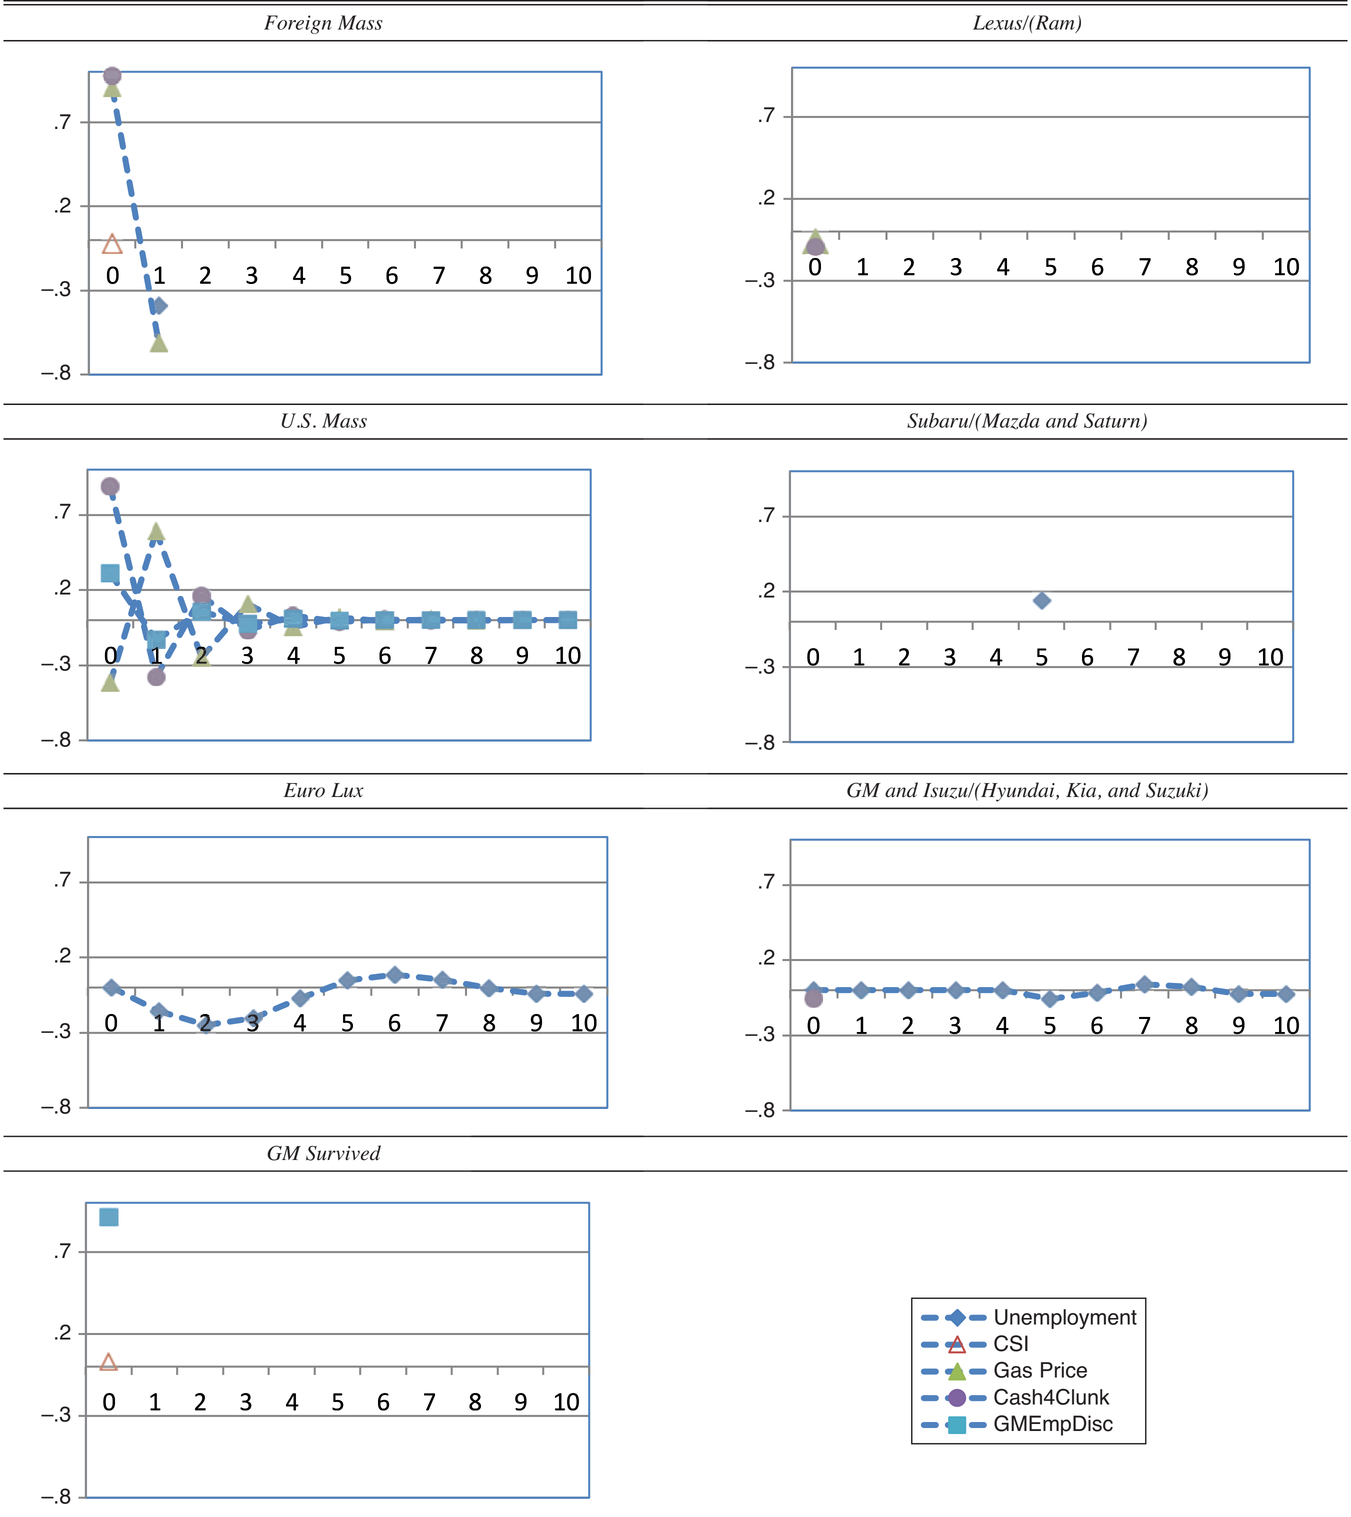

stantially, it is usually a sign that certain past common trends may have been replaced by new ones. It would be useful, albeit challenging, to develop a formal statistical model that can automatically detect such regime shifts.

Another promising area for extension would be new product adoption and diffusion. For example, an innovation is introduced to the marketplace, and multiple postlaunch tracking measures are gathered. Our model can potentially be modified to include a dynamic latent factor that follows an S-shaped time path. The co-movement pattern shared by the postlaunch tracking measures should help pin down the underlying diffusion curve. Such a DFA approach should

Figure 9
STANDARDIZED RESIDUALS (μ_t) FOR FOUR INDIVIDUAL TIME SERIES

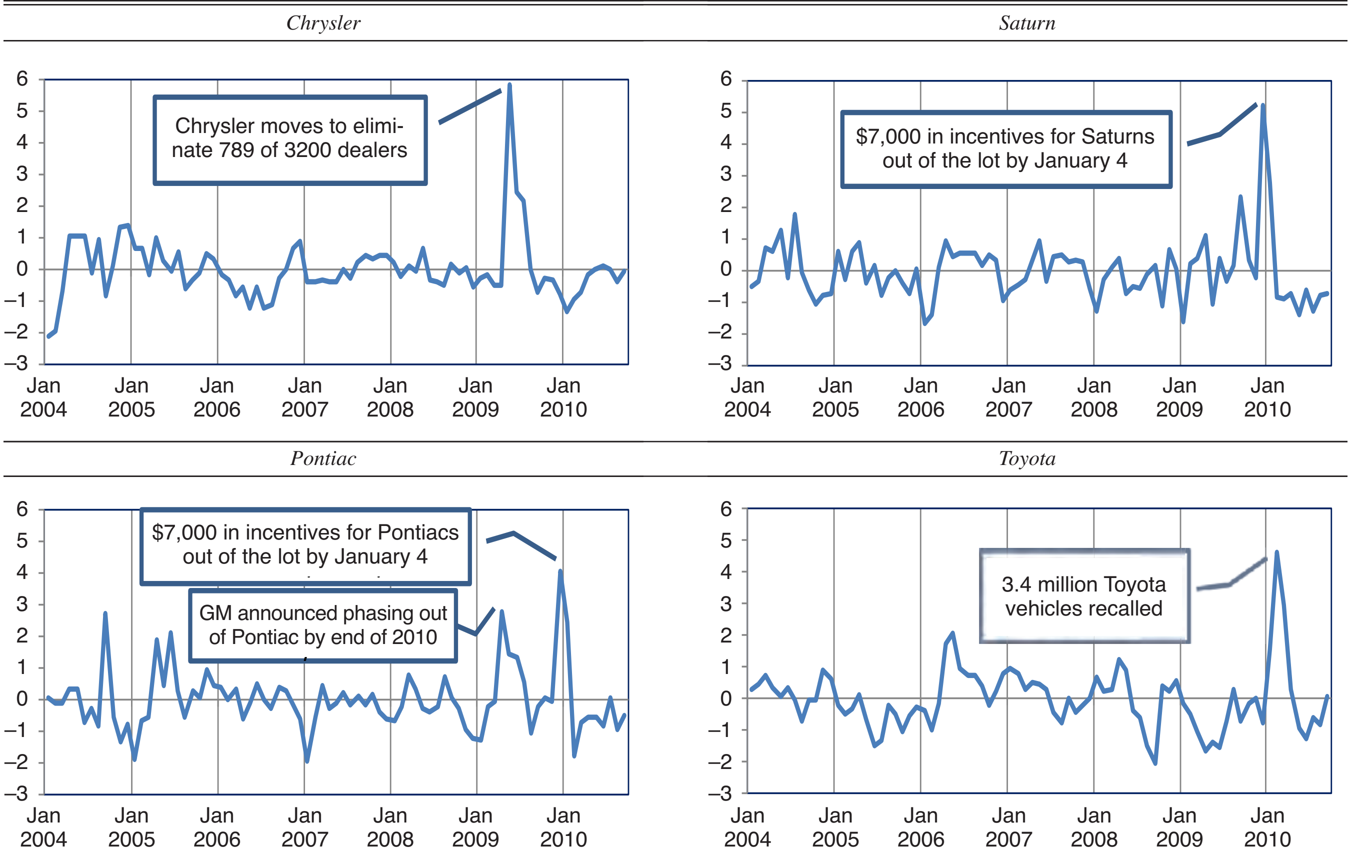

produce a more reliable read on the dynamics of the diffusion process than approaches that rely on any single indicator series.

Finally, as we have shown in our test of face validity, the seven latent trends distilled from 38 Google search indexes demonstrated an extraordinary amount of explanatory power on new vehicle sales. Two directions for more systematic exploration include (1) the relationship between online consumer interest trending metrics and sales (and how this relationship may differ across product categories) and (2) a sales forecasting system fully leveraging the multitudes of online consumer interest trending metrics that are often available in real time.

REFERENCES

Anderson, Ted W. (1963), "The Use of Factor Analysis in the Statistical Analysis of Multiple Time Series," *Psychometrika*, 28 (1), 1–25.

Aruoba, S. Boragan, Francis X. Diebold, and Chiara Scotti (2009), "Real-Time Measurement of Business Conditions," *Journal of Business & Economic Statistics*, 7 (4), 417–27.

Box, George E.P., Gwilym M. Jenkins, and Gregory Reinsel (2008), *Time Series Analysis: Forecasting & Control*, 3d ed. Hoboken, NJ: John Wiley & Sons.

Buss, Dale (2011), "Hyundai and Kia Prosper Now but Look to Future," *AutoObserver*, (July 25), (accessed February 24, 2012, [available at http://www.autoobserver.com/2011/07/hyundai-and-kia-prosper-now-but-look-to-future.html].

Busse, Meghan R., Duncan Simester, and Florian Zettelmeyer (2010), "'The Best Price You'll Ever Get': The 2005 Employee Discount Pricing Promotions in the U.S. Automobile Industry," *Marketing Science*, 29 (2), 268–90.

Carriero, Andrea, George Kapetanios, and Massimiliano Marcellino (2009), "Forecasting Exchange Rates with a Large Bayesian VAR," *International Journal of Forecasting*, 25 (2), 400–417.

Croux, Christophe, Eric Renault, and Bas J. M. Werker (2004), "Dynamic Factor Models," *Journal of Econometrics*, 119 (2), 223–30.

Datamonitor (2005), "Global Round-up: U.S. Vehicle Sales: Asian Inroads," *MarketWatch*, 4 (6), 13.

De Mol, Christine, Domenico Giannone, and Lucrezia Reichlin (2008), "Forecasting Using a Large Numer of Predictors: Is Bayesian Shrinkage a Valid Alternative to Principal Components?" *Journal of Econometrics*, 146 (2), 318–28.

Doz, Catherine and Fabrice Lenglart (2001), "Dynamic Factor Analysis: Estimation and Test with an Application to European Business Surveys," working paper, CEPR/Banca d'Italia.

Engle, Robert and Mark Watson (1981), "A One-Factor Multivariate Time Series Model of Metropolitan Wage Rates," *Journal of the American Statistical Association*, 76 (376), 774–81.

Forni, Mario, Marc Hallin, Marco Lippi, and Lucrezia Reichlin (2000), "The Generalized Dynamic-Factor Model: Identification and Estimation," *Review of Economics and Statistics*, 82 (4), 540–54.

Geweke, John (1977), "The Dynamic Factor Analysis of Economic Time Series," in *Latent Variables in Socio-Economic Models*, Dennis J. Aigner and Arthur S. Goldberger, eds. Amsterdam: North-Holland, 365–83.

Harvey, Andrew C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge, UK: Cambridge University Press, 449–56.

——— and Neil Shephard (1993), "Structural Time Series Models," in *Handbook of Statistics*, Vol. 11, G.S. Maddala, C.R. Rao and H.D. Vinod, eds. Amsterdam: Elsevier Science, 261–302.

——— (2010), "The Local Quadratic Trend Model," *Journal of Forecasting*, 29 (1/2), 94–108.

Jöreskog, Karl G. (1979), "Statistical Estimation of Structural Models in Longitudinal-Developmental Investigations," in *Longitudinal Research in the Study of Behavior and Development*, J.R. Nesselroade and P.B. Baltes, eds. New York: Academic Press, 303–351.

Ljung, Lennart (1987), *System Identification: Theory for the User*. Englewood Cliffs, NJ: Prentice Hall.

Ludvigson, Sydney C. and Serena Ng (2007), "The Empirical Risk-Return Relation: A Factor Analysis Approach," *Journal of Financial Economics*, 83 (1), 171–222.

Lutkepohl, Helmut (1991), *Introduction to Multiple Time Series Analysis*. Berlin: Springer-Verlag, 421–24.

Merrit, Jennifer (2010), "Studies of Web Searches Offer Preview of Hiring Trends," *The Wall Street Journal*, (March 5), (accessed February 24, 2012), [available at http://online.wsj.com/article/SB10001424052748704541304575099450389844456.html].

Molenaar, Peter C.M. (1985), "A Dynamic Factor Model for the Analysis of Multivariate Time Series," *Psychometrika*, 50 (2), 181–202.

———, Jan Gooijer, and Bernhard Schmitz (1992), "Dynamic Factor Analysis of Non-stationary Multivariate Time Series," *Psychometrika*, 57 (3), 333–49.

——— and Nilam Ram (2009), "Advances in Dynamic Factor Analysis of Psychological Process," in *Dynamic Process Methodology in the Social and Developmental Sciences*, J. Valsiner et al., eds. Heidelberg, Germany: Springer, 255–68.

Priestley, Maurice B. (1981), *Spectral Analysis and Time Series*. New York: Academic Press.

Sargent, Thomas J. and Christopher A. Sims (1977), "Business Cycle Modeling Without Pretending to Have Too Much a Priori Economic Theory," in *New Methods in Business Research*, Christopher A. Sims, ed. Minneapolis, MN: Federal Research Bank of Minneapolis, 45–109.

Shumway, Robert H. and David S. Stoffer (1982), "An Approach to Time Series Smoothing and Forecasting Using the EM Algorithm," *Journal of Time Series Analysis*, 3 (4), 253–64.

——— and ——— (2000), *Time Series Analysis and Its Applications*. New York: Springer-Verlag.

Stock, James H. and Mark W. Watson (2002), "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business & Economic Statistics*, 20 (2), 147–62.

The Wall Street Journal (2008), "Eyes on the Road: Luxury Auto Makers Gird for Sales Slump," (October 20), (accessed February 24, 2012), [available at http://online.wsj.com/article/SB122443910753248171.html].

Zuur, Alain F., R.J. Fryer, I.T. Jolliffe, R. Dekker, and J.J. Beukema (2003), "Estimating Common Trends in Multivariate Time Series Using Dynamic Factor Analysis," *Environmetrics,* 14 (7), 665–85.

Quantitative Trendspotting

Rex Yuxing Du and Wagner A. Kamakura

Web Appendix A – Inferring and Projecting the Latent Dynamic Factors

The procedure for inferring the latent state variables (i.e., $z_t \stackrel{\text{def}}{=} [\alpha_t \quad \beta_t \quad \delta_t \quad \gamma_t \quad \ldots \quad \gamma_{t-s-2}]'$), from the observed indicators (i.e., y_t from t = 1 through T), given the model parameters (i.e., B, L, Σ_u, Σ_ε, Σ_η, Σ_ζ, and Σ_ξ in Equations 1, 2, and 2.1 through 2.4), consists of two passes through the data: a forward pass (periods 1 through T) referred to in the literature as Kalman filter, and a backward pass (periods T through 1) often referred to as Kalman smoother (cf. Shumway and Stoffer 2000). Below we provide details on this inference procedure, which is important for understanding the mapping from the observed indicators to the latent state variables, which in turn determine the latent dynamic factors (f_t).

Kalman filter (forward pass)[1]

Step 1 – initialize $z_{0|0}$ and $var(z_{0|0})$, i.e., prior of the initial state at period 0

- $z_{0|0} = \mu_0$
- $var(z_{0|0}) = \Sigma_0$

Repeat Steps 2 and 3 for t = 1, …, T

Step 2 – calculate $z_{t|t-1}$ and $var(z_{t|t-1})$ i.e., expectation and uncertainty about the state in t given data observed up until t-1

- $z_{t|t-1} = Cz_{t-1|t-1}$

[1] The parametric definitions of matrices μ_0, Σ_0, A, B, C, Σ_u and Σ_v that appear in the procedure described below are given in Equations B1 and B2 in Web Appendix B.

- $var(z_{t|t-1}) = Cvar(z_{t-1|t-1})C' + \Sigma_v$

Step 3 – calculate $z_{t|t}$ and $var(z_{t|t})$ i.e., expectation and uncertainty about the state in t given data observed up until t

- $K_t = var(z_{t|t-1})A'[Avar(z_{t|t-1})A' + \Sigma_u]^{-1}$ (a.k.a. Kalman gain from t)
- $z_{t|t} = z_{t|t-1} + K_t[y_t - Az_{t|t-1} - B]$ (additional data from t used to update expectation)
- $var(z_{t|t}) = [I - K_tA]var(z_{t|t-1})$ (additional data from t used to reduce uncertainty)

Kalman smoother (backward pass)

Repeat Step 4 for t = T, T-1, …, 1

Step 4 – calculate $z_{t-1|T}$ and $var(z_{t-1|T})$, i.e., expectation and uncertainty about the state in period t-1, given all the data available through T

- $J_{t-1} = var(z_{t-1|t-1})C'[var(z_{t|t-1})]^{-1}$
- $z_{t-1|T} = z_{t-1|t-1} + J_{t-1}[z_{t|T} - z_{t|t-1}]$ (all data observed through T to update expectation)
- $var(z_{t-1|T}) = var(z_{t-1|t-1}) + J_{t-1}[var(z_{t|T}) - var(z_{t|t-1})]J_{t-1}'$ (all data observed through T to reduce uncertainty).

Step 5 – initialize $cov(z_{T|T}, z_{T-1|T})$, i.e., covariance of uncertainties about states at T and T-1

- $cov(z_{T|T}, z_{T-1|T}) = [I - K_TA]var(z_{t-1|t-1})$

Repeat Step 6 for t = T, T-1, …, 2

Step 6 – calculate $cov(z_{t-1|T}, z_{t-2|T})$, i.e., covariance of uncertainties about states at t-1 and t-2, given all the data available through T

- $cov(z_{t-1|T}, z_{t-2|T}) = var(z_{t-1|t-1})J_{t-2}' + J_{t-1}[cov(z_{t|T}, z_{t-1|T}) - var(z_{t-1|t-1})]J_{t-2}'$

The ultimate goal of applying the above Kalman filter and smoother is to infer, for t = 1 through

T, $z_{t|T}$, $var(z_{t|T})$ and $cov(z_{t|T}, z_{t-1|T})$, i.e., expectations, variances and lagged covariances of the state variables in any period during the observation window. Charting $z_{t|T}$'s over time would show the trend lines of the state variables, and $var(z_{t|T})$'s would determine the confidence band surrounding the trend lines.

Given the model parameters, the inferred state variables and the corresponding dynamic factor scores are determined by the complete history of the observed data, i.e., y_t for t = 1 through T. When data from additional periods are observed, the above procedure needs to be applied to update the entire course of the state variables. At the end of the last observation period (T), h-step ahead forecasts can be carried out as follows, given the model parameters and $z_{T|T}$:

(A1) $\mathrm{y_{T+h|T} = A^h z_{T|T} + B}$

(A2) $\mathrm{z_{T+h|T} = C^h z_{T|T}}$.

Or, more intuitively (but equivalently),

(A3) $\mathrm{y_{T+h|T} = L f_{T+h|T} + B}$ (predicted indicator)

(A4) $\mathrm{f_{T+h|T} = \alpha_{T+h|T} + \gamma_{T+h|T}}$ (predicted dynamic factor score)

(A4.1) $\mathrm{\alpha_{T+h|T} = \alpha_{T|T} + \beta_{T|T} h + \delta_{T|T} \frac{h(h-1)}{2}}$ (predicted non-seasonal trend component)

(A4.2) $\mathrm{\beta_{T+h|T} = \beta_{T|T} + \delta_{T|T} h}$ (predicted slope of the trend)

(A4.3) $\mathrm{\delta_{T+h|T} = \delta_{T|T}}$ (predicted change rate of the slope)

(A4.4) $\mathrm{\gamma_{T+h|T} = \gamma_{T|T}}$, if $remainder(h/s) = 0$ (predicted seasonal component)

$\mathrm{\gamma_{T+h|T} = -\sum_{j=0}^{s-2} \gamma_{T-j|T}}$, if $remainder(h/s) = 1$

$\gamma_{T+h|T} = \gamma_{T-s+\text{remainder}(h/s)|T}$, if $remainder(h/s) > 1$.

From the standpoint of trend projection, Equation A4.1 is of particular interest. It shows that a quadratic non-seasonal trend line can be extrapolated into the future, which has $\alpha_{T|T}$ at the origin, an initial slope of $\beta_{T|T}$ that will change at a rate of $\delta_{T|T}$ (Equation A4.2). A positive (negative) $\beta_{T|T}$ would indicate a trend line heading up (down). When $\beta_{T|T}$ and $\delta_{T|T}$ are of the same (opposite) sign, it indicates the trend will accelerate (decelerate). Given $\alpha_{T+h|T}$ and $\gamma_{T+h|T}$, we have $f_{T+h|T}$ (Equation A4), which in turn leads to predicted indicators, $y_{T+h|T}$ (Equation A3).

Web Appendix B – Model Calibration

Although Equations 1, 2 and 2.1 through 2.4 facilitate interpretation, in order to calibrate our SDFA model it is actually more convenient to rewrite it in a state-space form:

(B1) $\mathrm{y_t = Az_t + B + u_t}$ $\quad\quad$ $\mathrm{u_t \sim N(0, \Sigma_u)}$ $\quad\quad$ observation equation

(B2) $\mathrm{z_t = Cz_{t-1} + v_t}$ $\quad\quad$ $\mathrm{v_t \sim N(0, \Sigma_v)}$ $\quad\quad$ state equation

In Equation B1, hereafter referred to as the *observation equation*, y_t $(n \times 1)$, B $(n \times 1)$, and u_t $(n \times 1)$ are the same as in Equation 1, representing, respectively, the observed indicators from period t, and the intercepts and the irregularities in these indicators. The irregularities are assumed to have mean zero and variance Σ_u. The difference between Equation B1 and Equation 1 lies in z_t, an $m \times 1$ vector of latent state variables, $z_t \stackrel{\text{def}}{=} [\alpha_t \quad \beta_t \quad \delta_t \quad \gamma_t \quad \dots \quad \gamma_{t-s-2}]'$, and $A \stackrel{\text{def}}{=} [L_{n\times p} \quad 0_{n\times 2p} \quad L_{n\times p} \quad 0_{n\times(s-2)p}]$, the matrix determining how each state variable affects each observed indicator. The dimensionality of the state variables is $m = (3 + s - 1)p$, where s denotes the number of periods within a seasonal cycle, and p the number of dynamic factors.

Equation B2, hereafter referred to as the *state equation*, indicates that the unobserved state variables are updated from z_{t-1} to z_t subject to transition matrix

$$C \stackrel{\text{def}}{=} \begin{pmatrix} I_p & I_p & 0 & 0 & 0 & 0 \\ 0 & I_p & I_p & 0 & 0 & 0 \\ 0 & 0 & I_p & 0 & 0 & 0 \\ 0 & 0 & 0 & -I_p & \cdots & -I_p \\ 0 & 0 & 0 & I_p & 0 & 0 \\ 0 & 0 & 0 & 0 & I_{(s-3)p} & 0 \end{pmatrix}$$, and shocks occurred during period t, i.e., $v_t \stackrel{\text{def}}{=} \begin{bmatrix} \varepsilon_t \\ \eta_t \\ \zeta_t \\ \xi_t \\ 0 \end{bmatrix}$, which is assumed a priori to be normally distributed with mean zero and variance

$$\Sigma_v \stackrel{\text{def}}{=} \begin{bmatrix} \Sigma_\varepsilon & 0 & 0 & 0 & 0 \\ 0 & \Sigma_\eta & 0 & 0 & 0 \\ 0 & 0 & \Sigma_\zeta & 0 & 0 \\ 0 & 0 & 0 & \Sigma_\xi & 0 \\ 0 & 0 & 0 & 0 & 0_{(s-2)p} \end{bmatrix}.$$

In Equations B1 and B2, i.e., the state-space form of our SDFA model, matrices A, B, Σ_u, and Σ_v contain parameters to be estimated. The expected value (μ_0) and uncertainty (Σ_0) of the initial state (z_0) also need to be estimated. We describe next the Expectation-Maximization (EM) algorithm we use to calibrate these parameters. For this description, denote the collection of all the model parameters as $\{\Theta\}$, which includes:

1. L, $n \times p$ matrix of loadings in Equation 1, mapping $p \times 1$ latent factors $\mathrm{f_t}$ to $n \times 1$ observed time series y_t,
2. B, $n \times 1$ vector of intercepts in Equation 1,
3. Σ_u, $n \times n$ variance-and-covariance matrix of the idiosyncratic component (u_t) in Equation 1,
4. Σ_ε, Σ_η, Σ_ζ and Σ_ξ, $p \times p$ variance-and-covariance matrices in Equations 2.1 through 2.4, indicating, respectively, the sizes of stochasticity in the level (ε_t), slope (η_t) and slope change rate (ζ_t) of the non-seasonal trend line of the latent factor, and the sizes of stochasticity in the seasonal component of the latent factor (ξ_t),

5. μ_0 and Σ_0, mean and variance of $z_0 \stackrel{\text{def}}{=} [\alpha_0 \quad \beta_0 \quad \delta_0 \quad \gamma_0 \quad \ldots \quad \gamma_{0-s-2}]'$, state of various components of the latent factors before the first observation of the manifest variables.

For parsimony, Σ_u is assumed to be diagonal, resulting in n variance terms to be estimated, as opposed to $n \times (n+1)/2$ terms. The gain in parsimony can be substantial when n is large (e.g., 38 in our application). By imposing this assumption it also means that we assume all the co-movements among the observed time series are caused by the latent common factors. We also assume Σ_ε, Σ_η, Σ_ζ, Σ_ξ and Σ_0 to be diagonal because we intend to extract latent factors that are independent of one another, each capturing a distinct temporal pattern.

The model parameters $\{\Theta\}$ are to be estimated given the observed time series $y_t \in Y$ for t = 1 through T. For a given number (p) of latent factors, the following EM algorithm can be used for model calibration, which consists of five basic steps:

- Step 1: **Initialize the parameters** $\{\Theta^0\}$.
- Step 2: **The E-step**. Given $\{\Theta^0\}$ and Y, we use the Kalman filter (the forward pass) and Kalman smoother (the backward pass) described in Web Appendix A to derive $z_{t|T}$ and $var(z_{t|T})$. That is, given the observed data and the model parameters, we can predict the expected mean and variance of the latent factors that could have generated the observed time series.
- Step 3: **The M-step**. Given Y, by maximizing the likelihood function given $z_{t|T}$ and $var(z_{t|T})$ for t = 0 through T from Step 2, $\{\Theta^0\}$ can be updated to $\{\Theta^1\}$ as follows,

  <u>For parameters related to the initial state</u>

  $$\mu_0 = z_{0|T}$$

  $$\Sigma_0 = \text{diag}[\text{var}(z_{0|T})]$$

  <u>For parameters related to the observation equation</u>

  $$A_{\alpha\gamma} \stackrel{\text{def}}{=} [I_{p\times p} \quad 0_{p\times 2p} \quad I_{p\times p} \quad 0_{p\times(s-2)p}]$$

  $$z^*_{t|T} \stackrel{\text{def}}{=} A_{\alpha\gamma} z_{t|T}$$

$$[L^h \quad B^h] = \left[\sum_{t=1}^{T} y_t^h z_{t|T}^{*}{}' \quad \sum_{t=1}^{T} y_t^h\right] \begin{bmatrix} \sum_{t=1}^{T} z_{t|T}^{*} z_{t|T}^{*}{}' + A_{\alpha\gamma} var(z_{t|T}) A_{\alpha\gamma}{}' & \sum_{t=1}^{T} z_{t|T}^{*} \\ \sum_{t=1}^{T} z_{t|T}^{*}{}' & T \end{bmatrix}^{-1}$$

$\Sigma_u^h = \frac{1}{T}\sum_{t=1}^{\mathrm{T}}\left[(y_t^h - L^h z_{t|T}^* - B^h)(y_t^h - L^h z_{t|T}^* - B^h)' + (L^h A_{\alpha\gamma}) var(z_{t|T})(L^h A_{\alpha\gamma})'\right]$, where the superscript h indicates the h-th row of the corresponding matrix.

<u>For parameters related to the state equation</u>

$\Sigma_\varepsilon^{\alpha_j} = \frac{1}{T}\sum_{t=1}^{\mathrm{T}}\left[\left(z_{t|T}^{\alpha_j} - C^{\alpha_j} z_{t-1|T}\right)\left(z_{t|T}^{\alpha_j} - C^{\alpha_j} z_{t-1|T}\right)' + I^{\alpha_j} var(z_{t|T}) I^{\alpha_j}{}' + C^{\alpha_j} var(z_{t-1|T}) C^{\alpha_j}{}' - I^{\alpha_j} cov(z_{t|T}, z_{t-1|T}) C^{\alpha_j}{}' - C^{\alpha_j} cov(z_{t|T}, z_{t-1|T}) I^{\alpha_j}{}'\right]$, where the superscript α_j indicates the row corresponding to the α component of the j-th factor.

Formula of the same form can be applied to calculate $\Sigma_\eta^{\beta_j}$, $\Sigma_\zeta^{\delta_j}$ and $\Sigma_\xi^{\gamma_j}$

- Step 4: **Go back to Step 2** with updated $\{\Theta^1\}$ and iterate until convergence.
- Step 5: **Varimax Rotation**. Like standard factor analysis or any other factor-analytic model, our SDFA model is invariant to orthogonal rotation. In other words, any orthogonal rotation (i.e., Q s.t. $Q \times Q' = I$) to the loadings matrix (L) along with an inverse rotation to the factor scores (f) would produce the same fit to the observed time series (i.e., $L \times f = (L \times Q) \times (Q' \times f) \stackrel{\text{def}}{=} L^* \times f^*$). To obtain a unique solution for the factor loadings and scores, we apply a Varimax rotation to L and $z_{t|T}$, seeking a distinctive factor structure such that, for each factor, large loadings will result for a few variables and the rest will be small[2].

Lastly, in our empirical implementation we standardized all the observed time series before applying the EM algorithm given above. It is important to note that standardization will not lead to different trends. To see this, let $y_t^* = \Sigma^{-1} y_t$, where y_t^* is the standardized series and Σ is a diagonal

[2] Varimax, the most common factor rotation routine in marketing research, should be available in most popular software packages (e.g., Proc Factor in SAS). For more on Varimax and other common factor rotation routines (e.g., Quartimax and Eqimax), see Basilevsky (1994), *Statistical Factor Analysis and Related Method: Theory and Applications*. John Wiley & Sons; and Kaiser (1958), “The Varimax Criterion for Analytic Rotation in Factor Analysis,” *Psychometrika*, 23 (3).

matrix containing standard deviations. As a result, the observation equation (Equation 1, $y_t = Lf_t + B + u_t$) is equivalent to $y_t^* = L^* f_t + B^* + u_t^*$, where $L^* = \Sigma^{-1}L$, $B^* = \Sigma^{-1}B$, $u_t^* = \Sigma^{-1}u_t$. This shows that f_t, the latent factor scores, will not be affected at all and the effect of standardization is equivalent to dividing the factor loadings, intercept and noise term by the corresponding time series' standard deviation. After standardizing each time series, we can interpret the factor loading estimates ($\hat{L}$) as "for every unit change in the latent factor, one would expect to see a change of $\hat{L}$ standard deviations in the corresponding time series." Without standardization, differences in factor loading estimates would confound differences in the scale/variance of the times series with differences in the correlation between the latent factor and the series. In analyses such as ours, one care about correlation between the factors and the observed indicators, as opposed to their covariance.

Web Appendix C – Split-Half Reliability Test

Our SDFA model assumes that the latent factor loadings (L) will be time-invariant during the observation window, and any observed temporal variations in the time series are caused by variations in the latent factor scores. One might question such an assumption and wonder whether the structure of co-movements among the time series may change over time as well. To address this concern, we conducted a split-half reliability test, where we recalibrated our model twice, once with data from the first half of the observation window and another with data from the second half. The table below reports three sets of correlations between the estimated loadings: 1) the full sample and the first half, 2) the full sample and the second half, and 3) the first half and the second half. Strong correlation coefficients in all cases indicate that the factor loadings structure has remained largely stable across the samples (full vs. first vs. second half).

Factor \ Correlation	Full sample & first half	Full sample & second half	First half & second half
Foreign Mass	0.92	0.93	0.95
U.S. Mass	0.83	0.92	0.89
Euro Lux	0.92	0.89	0.92
GM Survived	0.92	0.93	0.86
Lexus / (RAM)	0.88	0.74	0.76
Subaru / (Mazda & Saturn)	0.66	0.82	0.69
GM Cancelled & Isuzu / (Hyundai Kia & Suzuki)	0.74	0.87	0.81

The above results aside, we stress the need to distinguish between non-stationarity in the trend lines and non-stationarity in the loadings. Major shifts can take place in the trend lines with no significant changes in the loadings structure, which is determined by co-movements among the observed time series. Regardless of the values of the observed indicators, as long as they follow the same co-movement patterns, the resulting loadings will remain unchanged. To see this, suppose there was a sudden increase in gas prices due to a crisis in the Middle-East, and as a result the trend line representing consumer interest in fuel-efficiency spiked. Despite this sudden change, as long as vehicles of similar fuel efficiency are affected in a similar way, the correlation between them would remain the same, which would result in the same loadings structure.