

## Demand and distribution relationships in the ready-to-drink iced tea market: A graphical approach

Robert E. Krider · Tieshan Li · Yong Liu ·  
Charles B. Weinberg

Received: 19 November 2006 / Accepted: 6 August 2007 /  
Published online: 13 September 2007  
© Springer Science + Business Media, LLC 2007

**Abstract** The relationship between distribution coverage and market share for an incumbent brand and for new entrants in the ready-to-drink iced tea market during its growth phase is explored using state space diagrams. This graphical visualization method was originally developed to disentangle lead–lag relationships between short nonstationary time series, a situation in which standard econometric methods have difficulty. In this research we show the usefulness of this method for long time series in offering complementary insights to econometric models, in providing a simple and managerially useful tool, and in conducting exploratory data analysis to guide subsequent modeling decisions. In the ready-to-drink iced tea market, usage of this method shows that during introduction of new brands and growth of the category, demand leads distribution, but that as the market matures, the dominant incumbent can defend with a distribution-leading-demand dynamic. Further, distribution coverage eventually becomes relatively stable so that short term fluctuations in demand (probably due to responses to promotion) have minimal impact on distribution.

**Keywords** Distribution · Market evolution · Time series ·  
Measurement and inference · Retailing · Competitive analysis

---

R. E. Krider  
Simon Fraser University, 8888 University Drive, Burnaby, BC V5A 1S6, Canada  
e-mail: rkrider@sfu.ca

T. Li  
John Molson School of Business, Concordia University, Montreal, QC H3G 1M8, Canada  
e-mail: tieshli@jmsb.concordia.ca

Y. Liu (✉)  
Eller College of Management, University of Arizona, Tucson, AZ 85721, USA  
e-mail: yoliu@eller.arizona.edu

C. B. Weinberg  
Sauder School of Business, University of British Columbia, Vancouver, BC V6T 1Z2, Canada  
e-mail: weinberg@sauder.ubc.ca

## 1 Introduction

Understanding causal relationships among marketing variables over time is a central concern of researchers and managers. Over the past decade substantial progress has been made in this field (Pauwels et al. 2004). Most notably, vector autoregression (VAR) models have been used to disentangle complex time-dependent relationships among advertising, pricing, distribution, sales, profits, and other key marketing variables. In many cases, the initial step in such modeling exercises is to explore the data with simple statistical and graphical tools; such explorations are often presented in research articles to build intuition for the overall analysis. For example, Dekimpe and Hanssens (1995) show plots of the univariate time series that is used in their subsequent VAR model. They note that the visible trends suggest evolving, non-stationary series. Unit root tests are then used to confirm such observation and to help develop the subsequent VAR models (e.g., Bronnenberg et al. 2000; Pauwels 2001). Bronnenberg et al. (2000) show the market share time series of four brands of ready-to-drink tea, and note that the shares drift for more than two years, although they eventually seem to settle down to a stable level.

In all cases, the time series plots presented are univariate. In this research, we show how a recently introduced simple graphical method (Krider et al. 2005) can be used for initial explorations of two coupled time series. The method is based on state-space diagrams. Rather than plotting the values of two time series, say  $A(t)$  and  $B(t)$ , on the  $y$ -axis and time  $t$  on the  $x$ -axis, we define the *state* of the system at any time  $t$  as  $(A(t), B(t))$  and the *state-space* as the space of values  $(x, y)$  that this pair can take on. At each time  $t$ , the pair  $(A(t), B(t))$  is plotted as a single point in this state-space coordinate system. The trajectory of the system in the state-space is typically shown by joining the plotted states with a line in the order of increasing time. By plotting state-space diagrams for two related time series and observing the patterns of the trajectory, Krider et al. (2005) discuss how this graphical method takes advantage of the eye's ability to see dynamic patterns that algebra based formulations cannot easily recognize. The state-space plots show not only evolving trends, but also lead–lag relationships between the variables. Combined with institutional knowledge of the industry, these plots provide useful information to modelers during exploratory investigations and generate distinctive insights of their own. Furthermore, because state-space diagrams are easy to generate, managers can use them to quickly make inferences about the relationships among variables.

The use of state-space diagrams in Krider et al. (2005) was motivated by the desire to untangle the causal relationship between demand and distribution for very short nonstationary time series, with which econometric methods may have difficulty. In an application to motion pictures, they show that the method provides more robust results than a VAR model with the Granger causality test. In this study, we apply this method to long time series and demonstrate how it generates valuable insights in a straight-forward manner. While the VAR models in this case provide good, detailed multivariate information, the simplicity of the bivariate plots makes them useful for the exploratory phase of guiding modeling decisions, or in situations where the econometric machinery is not available. While state-space diagrams are limited to revealing relationships between two variables, understanding the dynamics of such a relationship is a key managerial concern. Moreover, as we illustrate in this paper, the

graphical approach can provide insights that may not be apparent with an econometric approach.

There is growing interest in the use of graphical approaches to illustrate and understand dynamic relationships. For example, Pauwels and Hanssens (2007) estimate the trend parameter after decomposing a single time series into three linear components: trend as captured by time  $t$ , a seasonal dummy, and an error term. Once estimated for various time windows, these trend parameters (referred to as “performance barometers”) are used to indicate if there is a positive, negative, or neutral trend and if this trend is increasing or decreasing over the time windows. Once plotted along time, the parameters indicate the trend of firm performance. Marketing mix variables are then used to explain these trends. Our graphical method adds to this growing interest by focusing on the lead–lag relationships between time series. It differs from most of the plots in traditional time series analysis in that it uses state-space diagrams jointly constructed with two time series, and that no parameter estimation is needed to construct these plots.

In the next section, we review the mechanisms of the graphical method. Simulated examples are used to illustrate how the lead–lag relationship in long time series can be identified by the state-space diagrams and to build intuition for the interpretation of these diagrams. We then apply this method to the demand and distribution time series in the ready-to-drink iced tea market. These data were analyzed by Bronnenberg et al. (2000). Since we find consistent results, we rely on that analysis to provide additional support for the graphical approach. Moreover, the graphical approach enables us to provide additional insights and to extend Bronnenberg et al.’s results. For example, it indicates that the demand-leading-distribution dynamic dominates for new products, whereas the opposite is true for mature products in this category.


## 2 Interpreting lead–lag relations with bivariate state-space trajectories

We first present a simple example to illustrate the basic mechanism of the graphical method.<sup>1</sup> Assume advertising is periodic (e.g., “pulsed”) and that sales respond to advertising with a short delay. Let the deviations from the mean advertising be  $X_t = A \sin(2\pi f \cdot t)$ . Measuring time in weeks, this series has frequency  $f$  in cycles per week, cycle period  $1/f$  in weeks, and maximum advertising spending of  $A$ . Let  $Y_t$  be the lagging sales at time  $t$ ,  $S$  the mean sales over time, and

$$Y_t = S + kX_{t-\Delta t} = S + kA \sin(2\pi f \cdot t - \theta).$$

Sales lag advertising by  $\Delta t$  weeks. The phase difference between the two series is  $\theta = 2\pi f \cdot \Delta t$ , which appears in the argument of the sine function. To visualize this series, set the frequency  $f = 1/(2\pi)$  and the phase difference  $\theta = \pi/2$ . The corresponding period of the series is  $1/f$ , which equals 6.2 weeks, and the lag  $\Delta t$  is  $\theta/2\pi f$ , which equals 1.6 weeks. With constants  $k$  and  $A$  both set to 1 and  $S$  to zero,


<sup>1</sup> This example is taken from Krider et al. (2005) for illustration purpose. Detailed model development and a discussion of limitations can be found there. Due to the reliance on state-space plots for visual inspection, the method is useful for two time series—the limitation to bivariate data is further discussed in the summary section of this paper.



**Fig. 1** Stylized example of sales lagging advertising expenditure

the two series are plotted in Fig. 1 as functions of time  $t$ . The period can be read directly from the graph as the time between successive peaks of either curve, and the lag as the zero crossing time differences between the two curves.

Figure 2 shows the same two series but as a state-space diagram by plotting the pairs  $(X_t, Y_t)$  for  $t=(1, \dots, 6)$ . With advertising leading sales in this case, and advertising plotted on the horizontal axis, the curvature of the trajectory is counter-clockwise. Reversing the axes reverse the curvature to be clockwise, so it is important to pay attention to the axes. Intuitively, this unique pattern can be seen by following a sequence of points to infer lead-lag behavior in Fig. 2, which provides a simple example to illustrate the underlying process. From time 1 to time 2, both advertising and sales are decreasing. Between times 2 and 3, advertising increases slightly, but sales continue to decline. At time 4 sales decrease by a small amount and begin to increase in period 5. The increase in sales lags the advertising increase.



**Fig. 2** A state-space diagram for the sales-lagging-advertising example

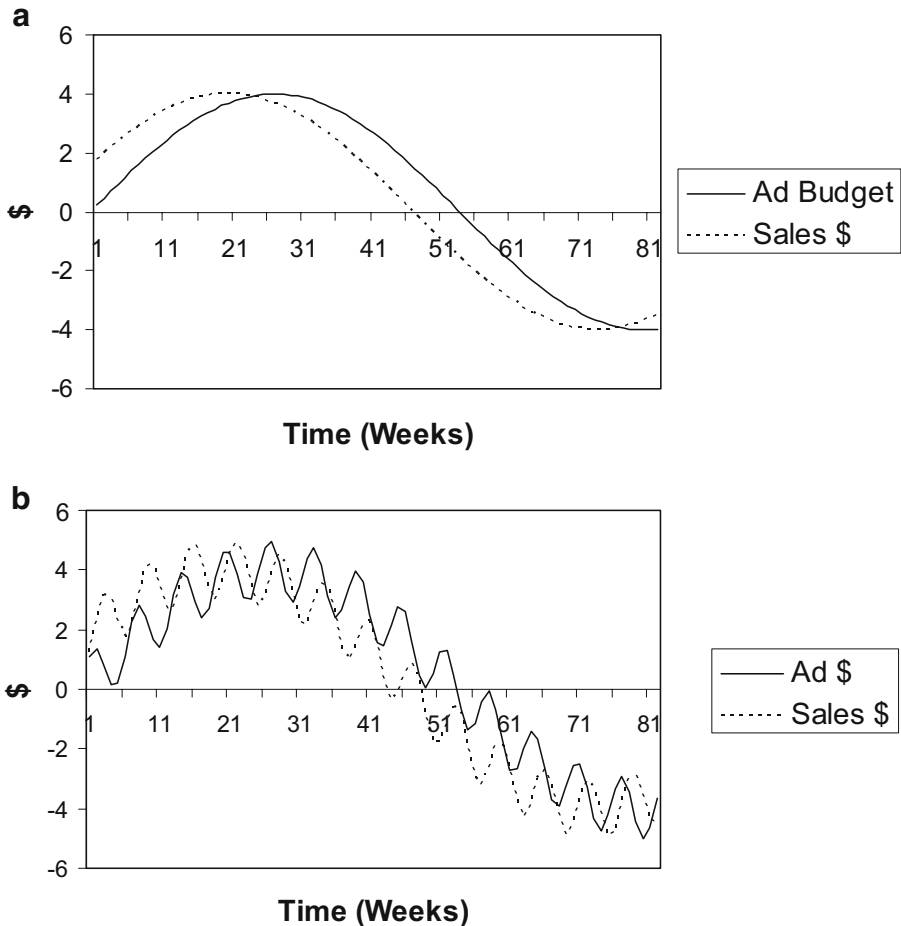
Shortly after time 5, advertising peaks and reverses, but sales continue to increase, again lagging the advertising reversal. Note that if the two series are perfectly in phase, corresponding to positively correlated contemporaneous effects, the plot will collapse to a straight line with a positive slope. The two series will move up and down the straight line in tandem, reaching their maximum and minimum at the same time. If they are perfectly out of phase (i.e., phase difference is  $180^\circ$ ), corresponding to negatively correlated contemporaneous effects, the plot will collapse to a straight line with a negative slope. As discussed in Krider et al. (2005), state-space diagrams offer the useful feature that curvature suggests a lead–lag relation and which series leads. Their main advantages are that if there is a lead–lag relationship, the curvature can be quickly detected by eye, and that changes in the direction as the market evolves can be easily determined.<sup>2</sup>

To further build intuition for the interpretation of the lead–lag relationship between two time series, and to illustrate how the state-space diagrams can be useful in detecting other subtleties in these dynamic relations, we now present two additional simulated scenarios. The first is a case where one series both leads and lags another series, but on different time scales. The second case shows a path-dependent effect, where the difference in lag time depends on whether the leading series is increasing or decreasing.

We illustrate the first case by returning to our advertising-sales example. Suppose that, in addition to the lagged sales response to advertising pulsing, the dynamics are also affected by an advertising budget baseline—the level around which the expenditures pulse—that is given by corporate headquarters on the basis of past sales. Specifically, suppose the advertising budget is set monthly on the basis of total sales in the previous four months, and that there is an exogenous business cycle of approximately two years in this industry that is reflected in long-run sales. The four month sales window used by corporate headquarters smoothes out the effects of the shorter term pulses, and effectively delays changes in the advertising budget to changes in the smoothed sales by 2 months, the midpoint of the window. In our lead–lag language, the long-period ad budget lags the long-period sales by two months. As before, the marketing department advertises in the six-week periodic pulsing pattern and sales pulses follow with a 1.6-week lag. The smooth long period sales cycle and delayed advertising budget are illustrated for 80 weeks in Fig. 3a. Short period fluctuations in advertising expenditures (for clarity of exposition taken to be identical in amplitude to those in the example in Fig. 1) and the resulting lagged short period sales results are superimposed in Fig. 3b.

The long period state-space trajectory corresponding to Fig. 3a is shown in Fig. 4a, with the *sales-leading* clockwise curvature. Figure 4b shows this long period trajectory, but now including the short period advertising expenditure and the *sales-following* pulses. The large-scale (long period) clockwise curvature remains, with the small-scale (short period) *counterclockwise* curvature added. An analyst presented with Fig. 4b could infer that sales leads advertising on a long time scale, and advertising leads sales on a short time scale. As a further modification, to illustrate

<sup>2</sup> Krider et al. (2005) conduct a test with three independent observers. In the test, the observers were asked to identify various curvature patterns provided to them. The test shows that the subjects were able to detect the nature of the dynamic relationship between the two variables in a highly consistent manner.

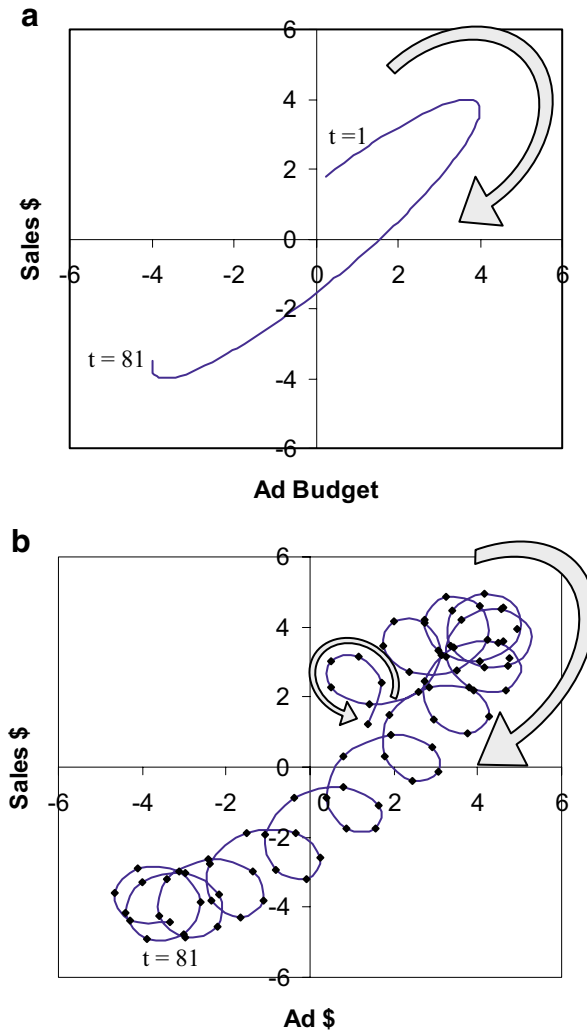


**Fig. 3** **a** Simulated smoothed sales *lead* advertising budgets by 8 weeks. Sales have an exogenous 2-year cycle. The y-axis is scaled to deviations from the mean as in Fig. 1. **b** Short advertising expenditure oscillations *lead* short period sales oscillations, but the long period advertising *lags* the long period sales

the effect of cycle amplitude, Fig. 5a and b show the time series and state-space trajectories with the business cycle reduced to the same amplitude as the pulsing cycles. While the shapes of the small versus large loops change accordingly, all the curvatures remain the same as before.


The second case we present to help build intuition illustrates that the sharpness of the curvature reflects the speed of the response: a sharper curve means a shorter lag time.<sup>3</sup> It also shows how a lag time that changes as rapidly as the fluctuations in the time series can be observed, as well as how a path-dependent lag time appears. As an

<sup>3</sup> In a static context, sharpness of curvature would be measured by the inverse of the radius of curvature. In a dynamic situation, such as the state-space trajectories in time, it would be measured as the second derivative with respect to time of a vector anchored at the origin that traces the trajectory over time. See the technical appendix in Krider et al. (2005) for the analytic definition of curvature.



**Fig. 4** **a** The trajectory corresponding to Fig. 3a, with the curvature resulting from a budget increase (or decrease) occurring after a sales increase (or decrease). **b** The trajectory corresponding to Fig. 3b, with sales leading advertising expenditures on a long cycle (*large clockwise arrow*), and advertising leading sales on a short cycle (*small counterclockwise arrow*)

example to conceptualize this, consider the transmission of price fluctuations through a distribution channel, for instance, the transmission of wholesale (refinery) oil price changes to retail gasoline (pump) price changes. We set up a simulated scenario of several cycles of wholesale prices with an increasing trend, which are followed by cycles of lagged retail price changes. In the simulated market, the retail pump price responds much quicker to wholesale price increases than to wholesale price decreases. Using smooth sinusoidal functions again, this scenario is shown as two time series in Fig. 6a, and as a state-space trajectory in Fig. 6b. The sharper curvature at the bottom left of each of the cycles in Fig. 6b, and the broader curvature at the top right, are the result of the asymmetric path-dependent response to price



**Fig. 5** **a** A simulation similar to Fig. 3b, with the same advertising pulsing amplitude, but with a lower-amplitude business cycle. **b** The trajectory corresponding to **a**. The advertising-leading counter clockwise pulses give the same size loops as in Fig. 4b, but the smaller amplitude business cycle spreads the loops out less

increases and decreases. The reader is encouraged to confirm this by following a cycle through as described for Fig. 2.

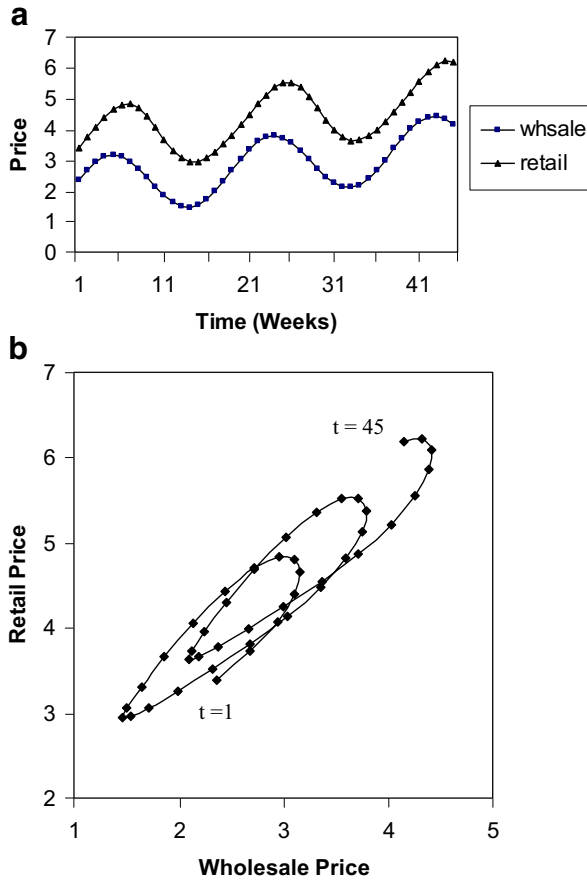
In summary, these stylized examples illustrate that our eye, with the aid of the state-space diagrams and some practice, is capable of quickly sorting out fairly complex lead–lag relations among other interesting dynamics in time series data.

### 3 Application to ready-to-drink iced tea

In the context of demand and distribution, we investigate weekly market share and distribution coverage for four brands of ready-to-drink iced tea during the period May 26, 1991 to April 21, 1996.<sup>4</sup> A brief history of this product category is as follows. Lipton introduced the category in 1972 and monopolized it until 1987, when

<sup>4</sup> In this dataset, distribution coverage is all-commodity value weighted, so that larger stores are given more weight (Bronnenberg et al. 2000).





**Fig. 6** **a** Path-dependent lags: a simulation of a market where retail price changes lag wholesale price changes, but the lag is greater for price decreases than price increases. **b** The state space trajectory of **a**. The quicker response of retail prices to wholesale price increases appears as a sharper curve at the bottom of the price cycles

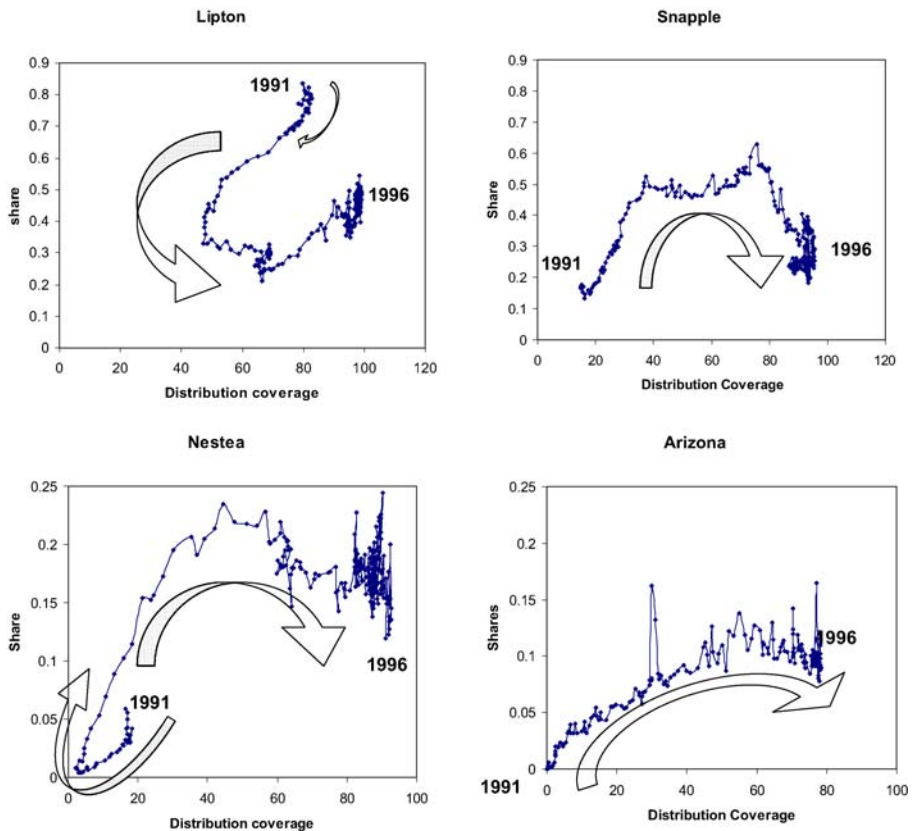
Snapple entered. In 1991, Lipton partnered with PepsiCo, and Nestea entered early in 1992 in a joint venture with Coca-Cola. The category began to grow rapidly, but Lipton initially lost sales and distribution to both Snapple and Nestea. Arizona entered in May 1992. Over the 5-year period covered by the data, the category volume grew seven-fold, with the strongest growth occurring in 1992 and 1993, and slowing through 1995 and 1996.

Bronnenberg et al. (2000) provide a thorough econometric analysis of this data set, using VAR models and impulse response function methodology. Their primary objective is to explore feedback relations between distribution and market share, and to investigate the resulting possibility of momentum (a shock to a variable has an effect on the system for some time after the shock, although it may eventually decay) and persistence (the effect of a shock does not decay to the original level, and the system stabilizes at a new level).

Figure 7 plots the state-space trajectories of market share and distribution coverage for the four brands. The arrows show the direction of increasing time and the associated curvature. In the following discussion, we shall illustrate where the plots provide inferences that are supported by Bronnenberg et al. (2000), where they offer additional and clarifying results, and highlight the types of insights that are generally useful to guide econometric modeling decisions.

The first noteworthy point is the obvious nonstationarity of both series for all four brands. The market shares and distribution coverage evolve or “accumulate” until near the end of the observed interval. At that point, as indicated by the clustering of the data points, market shares and distribution stop evolving and settle to a small region of the state space. Distribution coverage is particularly stable, while share continues to vary by about five or ten percent, consistent with promotional activities. Bronnenberg et al. (2000) use unit root tests to demonstrate this accumulation, and note that the evidence is weaker towards the end of the period. The state-space plots are consistent with this, and visually show the change from an evolving to a stable system.

To the modeler who is using the plots for initial exploratory work, the curvatures in the plots, as described in more detail below, also suggest that there is likely to be delayed feedback effects of changes in share or distribution coverage (i.e., positive



**Fig. 7** Lead-lag relationships in the ready-to-drink iced tea market, 1991–1996

feedback and momentum), and that these effects may be weaker or nonexistent at the end of the observed interval. Therefore an econometric model should allow for time dependent weakening of momentum.

The three recent entrants (i.e., Snapple, Arizona and Nestea) show strong clockwise curvature during the category growth phase. We can infer that distribution coverage was following market share gains with some lag. Retailers who had not yet carried them observed that sales of these brands were doing well and began to stock them as a result. The pioneer, Lipton, initially shows a clockwise curvature like the other brands, although in this case it is with a declining market share. Market share began to fall and distribution coverage followed, with some delay, as retailers stopped stocking the brand, presumably in favor of the new entrants. When market share dropped to 50%, distribution coverage stabilized. This suggests that Lipton put some intense push effort into stopping the decline of coverage, and slowing the loss of share (recall that the category is growing). Continued efforts resulted in recovering some distribution, followed by increasing market share, but with a considerable lag. In this stage, the direction of the effects reverses as indicated by the change in curvature direction—distribution leads market share. Bronnenberg et al. (2000) find a general result for all the brands that lagged effects of both market share on distribution and distribution on market share exist. They use interaction with time to show that both effects decrease with time. The plots provide additional insight on which brands each of the two effects arise from, and at what stages in the product lifecycle the different effects dominate.

All four series eventually settle into a small region of their demand-distribution state-space. The curvature becomes indeterminate as the category matures and stabilizes. This coincides with the loss of detectable positive feedback effects noted by Bronnenberg et al. (2000). Snapple and Lipton end with near complete distribution coverage, and market shares of about 25 and 40% respectively. Nestea and Arizona have distribution coverage of 90 and 80% respectively, and market shares of around 15 and 10%.

The state-space diagrams cannot provide all the details that an econometric approach can for long time series. However, they can quickly and very simply suggest modeling choices and provide consistent and, importantly, complementary information. An example of such complementary information is the dramatic difference between the established brand, Lipton, and the three entrants. The graphical method encourages studying the brands separately and thus we find these brand-specific results. It also suggests the sorts of other data, such as category sales disaggregated by new and mature brands, and pricing, that might be useful to help explain what underlies the patterns. While Lipton dominated an apparently mature category for many years, the new entrants redefined the category—they essentially created a new product category that entered a strong growth phase. In the early stages of the product lifecycle, gaining some significant market share with limited distribution was essential to trigger the share-leading-distribution feedback process and grow market share. Lipton had extensive distribution, and as its share fell, also showed the share-leading-distribution behavior. But in order to stop the slide and recover share, the reverse distribution-leading-share dynamic dominated. It is tempting to speculate that for new products, the share-leading dynamic should be exploited, for example by introducing a new product in a restricted geographical area and allocating resources to pull

strategies. Similarly, for mature products the distribution-leading dynamic should be exploited by using push strategies.

#### 4 Summary and discussion

Both the VAR model and state-space plots lead to the conclusion that, during the time period studied, the time series are evolving, or have momentum; that the market share and distribution can affect each other with a delay; and that the momentum declines over time. The state-space plots cannot distinguish between momentum that decays versus that which is persistent. On the other hand, the interesting difference between the well-established incumbent and the new entrants was striking in the plots, but not addressed in econometric models.

In addition to the limitations of the graphical method discussed in Krider et al. (2005), we note that the state-space plots cannot detect autoregressive effects of the variables on themselves, as the VAR methodology can, nor quantify the specific lag times of delayed effects as cleanly as VAR models. The bivariate plots are an improvement over univariate time series plots for exploratory work. However, they can only address multivariate problems by examining two variables at a time. Although a third or additional time series cannot be included in the state-space plots, it could be useful to examine such data on marketing mix variables (such as promotion) to see how they relate to dynamic patterns observed in the bivariate plots. Due to data availability, we could not do so for the four brands of iced tea in this paper.

State-space diagrams provide the analyst with insights into dynamic relationships by visually examining the trajectory and curvature of the state-space trace. Like other graphical and visualizations techniques, this method can provide useful direction in selecting among the myriad model-building choices faced by econometric modelers, as well as quickly and easily provide useful insights on its own for researchers and managers.

**Acknowledgement** The authors would like to thank Bart Bronnenberg for allowing us to use the ready-to-drink iced tea data. The financial support of the Social Sciences and Humanities Research Council of Canada is gratefully acknowledged.

#### References

- Bronnenberg, B., Mahajan, V., & Vanhoner, W. (2000). The emergence of market structure in new repeat-purchase categories: A dynamic approach and an empirical application. *Journal of Marketing Research*, 37(1), 16–31.
- Dekimpe, M., & Hanssens, D. (1995). The persistence of marketing effects on sales. *Marketing Science*, 14, 1–21.
- Krider, R. E., Li, T., Liu, Y., & Weinberg, C. B. (2005). The lead-lag puzzle of demand and distribution: A graphical method applied to movies. *Marketing Science*, 24(4), 635–645.
- Pauwels, K. (2001). The drivers of online relationship: A decomposition approach. Social Science Research Network, Working Paper 01–08.
- Pauwels, K., Currium, I., Dekimpe, M., Ghysels, E., Hanssens, D., Mizik, N., et al. (2004). Modeling marketing dynamics by time series econometrics. *Marketing Letters*, 15(4), 167–183.
- Pauwels, K., & Hanssens, D. M. (2007). Performance regimes and marketing policy shifts. *Marketing Science*, 26(3), 293–311.