Forecasting the Global Electronics Cycle with Leading Indicators: A VAR Approach

by

Hwee Kwan Chow* and Keen Meng Choy†

*School of Economics and Social Sciences, Singapore Management University, Singapore 259756.

Corresponding author. Tel.: +65-6-8220868; fax: +65-6-8220833; e-mail: hkchow@smu.edu.sg

†Department of Economics, National University of Singapore, Singapore 117570. Tel.: +65-6-8744874; fax: +65-6-7752646; e-mail: ecsckm@nus.edu.sg

Abstract

Developments in the global electronics industry are typically monitored by tracking indicators that span a whole spectrum of activities in the sector. However, these indicators invariably give mixed signals at each point in time, thereby hampering efforts at prediction. In this paper, we propose a unified framework for forecasting the global electronics cycle by constructing a VAR model that captures the economic interactions between leading indicators representing expectations, investments, orders, inventories and prices. The ability of the indicators to presage world semiconductor sales is assessed by Granger causality tests. The VAR model is also used to derive the dynamic paths of adjustment of global chip sales in response to shocks in each of the leading variables. These impulse response functions confirm the leading qualities of the selected indicators. Finally, out-of-sample forecasts of global chip sales are generated from the VAR model and compared with predictions from a univariate model as well as a model which uses a composite index of the leading indicators. An evaluation of their relative accuracy suggests that the VAR model's forecasting performance is superior to that of the univariate model and comparable to that of the composite index model.

Key Words and Phrases: Leading indicators; Global electronics cycle; VAR; Forecasting

1 Introduction

The semiconductor industry sets the pace of global economic growth, more so than any other single sector, and its vitality is a leading indicator of the world's economic health. As fundamental building blocks of final electronic products, semiconductors (also known as chips) are used as inputs in a wide variety of sectors such as information and communication technology, consumer electronics, as well as the industrial and transportation sectors. Thus, chips serve as a cornerstone to the global electronics industry. A key characteristic of the semiconductor industry is the acceleration of technology which renders each new generation of semiconductors obsolete fairly quickly.[2] Consequently, product cycles are short and this, in turn, results in a compression of the overall global electronics cycle. At the same time, the commoditization of semiconductors—whereby an innovation initially generating high profits plunges in value as the technology for producing it becomes widespread and standardized—brings on wide fluctuations in the electronics industry.

The inherent volatility of the global electronics cycle is perhaps most vividly illustrated by the information technology boom during the 1990s, followed by the bursting of the technology bubble in late 2000. It is evident that worldwide economic growth, particularly the domestic business cycles of economies that are heavily reliant

[2] The semiconductor industry is driven by Moore's Law which says that the number of transistors on a chip doubles every 18 to 24 months, resulting in ever faster and cheaper semiconductors.

on electronics exports, is severely impacted by such swings in electronics demand. It follows that close monitoring of the electronics industry is essential for assessing the health of the world economy, which means that timely and accurate forecasts of the global electronics cycle are indispensable.

Developments in the electronics industry have typically been monitored by tracking a host of diverse indicators, such as those measuring expectations, investments, orders, inventories, production, shipments, prices and profits. As these indicators span a whole spectrum of activities, they invariably give mixed signals at each point in time, thereby hampering efforts to predict world electronics activity. Apart from product cycles, global electronics demand can also be affected by other factors and the predictive value of each indicator might vary depending on which causal factors are pre-eminent in a particular cyclical episode. There is, therefore, a need for a systematic examination of the predictive potential of each indicator. Yet, the approach that has been adopted to circumvent the problem of mixed signals in electronics indicators—and for that matter, in leading indicators of the economy—is to aggregate them to form a composite index. For instance, the *Monetary Authority of Singapore* has developed an electronics composite leading index comprising five indicators to forecast Singapore's domestic electronics output and exports (Ng et al., 2004), while *Gartner Research* has a composite index of semiconductor market leading indicators for predicting growth in the world semiconductor industry.

In this paper, we propose a unified framework for forecasting the global electronics cycle by constructing a vector autoregressive (VAR) model which incorporates a set of leading indicators identified from a longer list of electronics series. To the best of our knowledge, this has hitherto not been done in the literature. Given the endogeneity of and dynamic interactions between the economic variables influencing the world electronics cycle, forecasting within a VAR framework may confer advantages. Firstly, it frees us from the implicit assumption made in the index approach of a single common factor underlying the movements in electronics indicators, possibly associated with the product cycle. Secondly, the flexibility of the VAR model means that it can potentially accommodate the different lead times of indicators, which might partly account for the conflicting signals received.

We initially use the VAR model to perform Granger causality tests that assess the ability of the selected leading indicators to presage world semiconductor sales. Following this, their leading qualities are examined through an impulse response analysis by tracing out the dynamic adjustment paths of global chip sales in response to orthogonalized shocks in each of the indicators. The VAR model is next employed to generate out-of-sample forecasts of global semiconductor sales. Finally, we evaluate the relative predictive accuracy of the VAR model against a benchmark univariate model and an alternative model which uses a composite index constructed from the leading indicators.

2 Leading Indicator Selection

The first task in forecasting the global electronics cycle is to search for plausible leading indicators. We began with a list of indicators that covers, *inter alia*, US time series on electronics new orders, inventories and shipments. Also included in the list are producer prices for dynamic random access memory (DRAM), the Institute of Supply Management's (ISM) manufacturing Purchasing Managers' Index (PMI), the book-to-bill ratio of semiconductor equipment and Nasdaq stock prices, all of which are widely used as de facto leading indicators of the global electronics cycle by private sector analysts. In addition, US corporate profits and private fixed investment in information processing equipment and in computers and peripherals were also considered as possible proxies of the final end-user demand for electronics.

The selection of leading indicators from the pool of economic variables at our disposal could be a potentially daunting exercise. Assuming that 5 indicators are to be picked from 15 series, there are over 3000 combinations of indicators to choose from. We resolved the conundrum by appealing to the classical criteria used by researchers at the National Bureau of Economic Research (NBER) to select leading indicators for the macroeconomy. These include 'economic significance', 'currency' and 'conformity' (Zarnowitz, 1992, pp. 317–319). We ensured that the first criterion is satisfied i.e., there should be an economic reason for why an indicator leads. Accordingly, US shipments of electronics was dropped as it appears by definition to be more nearly coincident with the global electronics cycle. The PMI also did not qualify as a leading indicator because the share of electronics production in US manufacturing output is fairly small. The currency criterion, interpreted as a timeliness constraint, meant that quarterly time series should be eschewed in favour of monthly ones, thereby precluding the selection of the profits and investment series as leading indicators.

As a measure of an indicator's conformity, we calculated its cross correlation coefficients at various lead times with the coincident indicator of the electronics cycle used in our study—global semiconductor sales. This indicator represents world billings or shipments of semiconductor products, as reported by the Semiconductor Industry Association (SIA) at its website (we have seasonally adjusted the raw data using the Census X-12 multiplicative method). We chose to use global chip sales as the coincident series because it is commonly viewed as the best available indicator of the unobserved state of the world electronics sector.[3] The conformity criterion, taken together with the need to ensure timeliness, further eliminated electronics series that exhibited statistically insignificant correlations or very short leads of less than three months, resulting in the eventual selection of five variables as putative leading indicators of the global electronics cycle.[4]

[3] Some might argue that the use of a coincident index of world electronics activity, analogous to the one developed for the US technology cycle by Hobijn et al. (2003), is preferable to relying on a single indicator. However, the construction of such an index is beyond the scope of this paper.

[4] The cross correlation results are available upon request from the authors.

The identified variables are the Nasdaq composite index (NASDAQ), the North American book-to-bill ratio for semiconductor equipment (BTB), US new orders of electronics (NO), the inverted change in US electronics inventories (INVENT)[5], and the US producer price index for DRAM (PPI). The Nasdaq index was downloaded from Datastream, the book-to-bill ratio from the Semiconductor Equipment and Materials International (SEMI) website, the seasonally adjusted new orders and inventories series from the Census Bureau website (series codes are A34SNO and A34STI respectively), and the PPI from the Bureau of Labour Statistics website (the series code is PCU3344133344131A101). The overlapping sample period of these monthly datasets is 1992:2–2004:1, which is therefore the time period used in the paper.

We end this section with a discussion of the economic rationales behind our chosen set of leading indicators that draws on ideas in Zarnowitz (1992) and de Leeuw (1991). The Nasdaq stock price index is a good proxy for firms' expectations about future global electronics activity. At the root of the leading relationship is the market's sensitivity to the discounted future earnings of technology firms that supply to world markets, which are ultimately dependent on the final demand for electronics products. A drawback of stock prices is that they tend to be affected by other factors, including speculation, thus occasionally giving rise to false signals. Like the Nasdaq

[5] In its latest revisions to the historical data, the Census Bureau has excluded semiconductors from the new orders series but included them in the inventory series. We would have preferred to use indicators with a consistent coverage had they been available.

index, the book-to-bill ratio responds to feedback from the end-user demand for semiconductors, as well as from chip prices. Being the three-month moving average ratio of new orders to sales received by North American-based manufacturers of semiconductor equipment, however, the ratio has a tendency to lead global chip sales because investment decisions by equipment makers temporally precede other processes.

New orders of electronics is synonymous with demand and serve as an indicator of the early stage in the production process. This indicator might be expected to lead electronics activity because it usually takes time to translate an order into actual production and sales; it works especially well as a leading indicator if firms adopt 'just-in-time' manufacturing technologies. However, given that firms do try to anticipate future sales, only unexpected changes in orders will presage global chip sales.

The level of electronics inventories has a propensity to lag the electronics cycle. But when its inverted changes are taken or it is considered in relation to sales as in the shipment-to-inventory ratio, the series becomes a leading indicator.[6] Changes in inventories help firms smooth production by acting as a buffer to unexpected fluctuations in demand. For example, an increase in orders could be met by a temporary drawdown in inventories before prices are adjusted. Indeed, anecdotal evidence suggests that the elimination of excess inventory in a downturn is a pre-requisite for

[6] We do not consider the shipment-to-inventory ratio in order to avoid duplicating the coverage of the inverted inventory change series.

future increases in prices and sales. DRAM prices respond in turn to both anticipated and unforeseen imbalances in demand and supply, making them a leading indicator in much the same way as the prices of sensitive materials.

3 A VAR Analysis of Electronics Leading Indicators

In this section, we carry out empirical analyses to demonstrate the leading qualities of the identified electronics indicators. These latter were converted into natural logarithms to stabilize their variances and mitigate departures from normality. We investigated the integration status of the transformed series by applying the DF-GLS unit root test developed by Elliot, Rothenberg and Stock (1996), in conjunction with the modified AIC for selecting the lag length proposed by Ng and Perron (2001). The DF-GLS test is an asymptotically more powerful variant of the augmented Dickey-Fuller (ADF) test obtained via generalized least squares detrending.

The results are shown in Table 1. Except for the inverted change in electronics inventories and the semiconductor book-to-bill ratio, which is apparently stationary, the indicator series were found to be integrated of order one. Given this, we checked for cointegration between them using Johansen's trace test with six lags and an unrestricted constant. The trace statistic for the null hypothesis that there is at most two cointegrating relations in the data is 36.92, thus making it impossible to reject the hypothesis even at the 10% significance level.

Table 1: Unit Root Tests

Variable	Lag Length	τ^{GLS}	5% Critical Value
NASDAQ	1	-1.260	-2.977
BTB	3	-2.823	-2.058
NO	2	-0.850	-2.965
INVENT	5	-2.236	-2.042
PPI	1	-2.632	-2.977
CHIP	5	-1.755	-2.924

Notes: The tests are for the logarithms of variables. A trend was included except in the cases of BTB and INVENT. Critical values are from Cheung and Lai (1995).

In the light of these findings, the empirical analyses are performed in the framework of a vector autoregression (VAR) in levels given by:

$$\mathbf{y}_t = \boldsymbol{\tau} + \mathbf{\Pi}_1\mathbf{y}_{t-1} + \cdots + \mathbf{\Pi}_k\mathbf{y}_{t-k} + \boldsymbol{\varepsilon}_t, \qquad t = 1, ..., T \tag{1}$$

where $\mathbf{y}_t = (NASDAQ,\ BTB,\ NO,\ INVENT,\ PPI,\ CHIP)'$, the $\mathbf{\Pi}_i$ are fixed (6×6) matrices of parameters, $\boldsymbol{\tau}$ is a (6×1) vector of constants and $\boldsymbol{\varepsilon}_t \sim MN(0, \mathbf{\Sigma})$ is multivariate normal white noise with zero mean. The optimal lag length k selected by minimizing information criteria such as the AIC and the Hannan-Quinn criterion was 3. However, the residuals that resulted from including only 3 lags in the VAR model exhibited autocorrelation and were also not normally distributed, with attendant complications for post-estimation inferences. As a remedy to both problems, we decided to use 6 lags in the analyses that follow.

3.1 Causality Tests

The standard Granger causality test entails specifying the VAR in (1) and testing to see if the subset of coefficients associated with a given leading indicator is jointly and significantly different from zero in the equation for global chip sales. Under the null hypothesis of no Granger causality, the test statistic follows a χ^2 distribution with m degrees of freedom in large samples, m being the number of zero restrictions imposed.

A summary of the empirical results from the Granger causality tests is presented in Table 2. The null hypothesis of non-causality can be rejected at the 5% significance level for three out of the five electronics indicators—inverted inventory change, the DRAM chip price and the Nasdaq stock index. The other two indicators do not Granger-cause global chip sales at the usual significance levels. Although these results may seem a little disappointing, it should be borne in mind that the use of a multivariate VAR for causality testing imposes relatively stringent requirements on the information content of electronics indicators. When pairwise causality tests were performed instead, the book-to-bill ratio was found to Granger-cause world semiconductor sales at the 10% significance level even though non-causality still cannot be rejected for the new orders series.[7]

[7] This finding might be explained by the fact that the shipments of electronics industries which

Table 2: Causality Tests

	Granger		Toda-Yamamoto	
Variable	χ_6^2	p-value	χ_6^2	p-value
NASDAQ	15.633	0.016	11.892	0.064
BTB	7.525	0.275	5.817	0.444
NO	4.237	0.645	2.980	0.811
INVENT	16.549	0.011	9.375	0.154
PPI	16.643	0.011	9.088	0.169

Notes: The VAR is estimated with six lags in the Granger tests and seven lags in the Toda-Yamamoto tests. The χ_6^2 values are the test statistics for the null hypothesis that a variable does not Granger-cause global chip sales.

The Granger causality tests carried out above based on levels estimation are asymptotically valid because the VAR is consistently estimated in the presence of cointegration (Sims, Stock and Watson, 1990). However, the results are conditional on the prior outcomes of the tests for unit roots and cointegration. To avoid possible pretest bias and at the same time provide a robustness check, we also implemented the causality test proposed by Toda and Yamamoto (1995). The advantage of the Toda-Yamamoto test is its invariance with respect to the integration and cointegration

do not produce to order are counted as part of new orders.

status of variables. It is similar to the standard Granger test in that an augmented VAR with $p = k + d_{\max}$ lags is estimated in place of (1), where $d_{\max}$ is the maximal order of integration suspected in the time series under consideration. If we take it that $d_{\max} = 1$, the estimated VAR is:

$$\mathbf{y}_t = \hat{\boldsymbol{\tau}} + \hat{\mathbf{\Pi}}_1 \mathbf{y}_{t-1} + \cdots + \hat{\mathbf{\Pi}}_k \mathbf{y}_{t-k} + \hat{\mathbf{\Pi}}_{k+1} \mathbf{y}_{t-k-1} + \hat{\boldsymbol{\varepsilon}}_t \tag{2}$$

The above equation can be re-written in more compact matrix notation as

$$\mathbf{Y}_0' = \hat{\boldsymbol{\tau}}\mathbf{i}' + \hat{\mathbf{\Phi}}\mathbf{X}' + \hat{\mathbf{\Pi}}_{k+1}\mathbf{Y}_{k+1}' + \hat{\boldsymbol{\xi}} \tag{3}$$

where $\mathbf{Y}_i = (\mathbf{y}_{1-i}, ..., \mathbf{y}_{T-i})$, $\mathbf{X} = (\mathbf{Y}_1', ..., \mathbf{Y}_k')'$, $\hat{\boldsymbol{\xi}} = (\hat{\boldsymbol{\varepsilon}}_1, ..., \hat{\boldsymbol{\varepsilon}}_T)$, $\hat{\mathbf{\Phi}} = (\hat{\mathbf{\Pi}}_1, ..., \hat{\mathbf{\Pi}}_k)$ and $\mathbf{i}$ is a $(T \times 1)$ vector of ones. The Wald statistic to test for Granger non-causality is

$$\mathcal{W} = f(\hat{\phi})'[F(\hat{\phi}) \left\{ \hat{\mathbf{\Sigma}}_\varepsilon \otimes (\mathbf{X}'\mathbf{Q}\mathbf{X})^{-1} \right\} F(\hat{\phi})']^{-1} f(\hat{\phi}) \tag{4}$$

where $\hat{\phi} = vec(\hat{\mathbf{\Phi}})$, $F(\hat{\phi}) = \partial f(\hat{\phi})/\partial \hat{\phi}'$, $\mathbf{Q} = \mathbf{Q}_\tau - \mathbf{Q}_\tau \mathbf{Y}_{k+1}(\mathbf{Y}_{k+1}'\mathbf{Q}_\tau \mathbf{Y}_{k+1})^{-1}\mathbf{Y}_{k+1}'\mathbf{Q}_\tau$, $\mathbf{Q}_\tau = \mathbf{I}_T - \mathbf{i}(\mathbf{i}'\mathbf{i})^{-1}\mathbf{i}'$, and $\hat{\mathbf{\Sigma}}_\varepsilon = T^{-1}\hat{\boldsymbol{\xi}}'\hat{\boldsymbol{\xi}}$. Notice that the parameter restrictions in $f(\hat{\phi}) = 0$ do not involve the coefficients of $\hat{\mathbf{\Pi}}_{k+1}$, since the latter are all zero under the assumption that the true lag length is k. Toda and Yamamoto (1995) prove that $\mathcal{W}$ converges in distribution to a χ_m^2 random variable irrespective of whether the $\mathbf{y}_t$ process is stationary, integrated or cointegrated.[8]

[8] The asymptotic results will hold as long as $p \geqslant k + d_{\max}$. This implies that the test is still valid if the true lag length is smaller than six.

The Toda-Yamamoto tests produced results that are similar to the Granger tests in levels, as can be seen from Table 2. However, the p-values for inverted inventories and the DRAM chip price now exceed 10% while the χ_6^2 statistics for the book-to-bill ratio and new electronics orders continue to suggest that these two variables are not causally prior with respect to global chip sales. We present further evidence in the next sub-section that overturns these ambiguous findings and demonstrates the leading abilities of all the electronics indicators.

3.2 Impulse Response Analysis

The second use to which we put the VAR model is the derivation of impulse response functions, which show the dynamic paths taken by global chip sales in response to innovations in the leading series. Traditionally, impulse response analysis in leading indicator research has been conducted using the bivariate methodology of transfer function models (Koch and Rasche, 1988; Veloce, 1996). We prefer to adopt a VAR approach because it accounts for the endogeneity of the electronics variables and also captures the dynamic economic relationships between the leading and coincident indicators.

The impulse response functions generated by the VAR model will only be meaningful if innovations to the variables in the system are serially and mutually uncorrelated. Granted this, the innovations can then be interpreted as unanticipated shocks to the leading indicators. Justifying the causal ordering with the economic rationales of the leading indicators discussed in the previous section, we orthogonalize these shocks by resorting to a Choleski decomposition of the estimated variance-covariance matrix of the residuals. In theory, if the individual series have distinct lead times over global chip sales, the contemporaneous correlations between their residuals in the VAR will be small and alternative causal orderings will yield impulse responses that look alike. This is in fact true for the majority of the empirical correlations. In any event, we tried putting the Nasdaq index after the book-to-bill ratio and new orders on the grounds that the share prices of technology firms might very well react to the release of new data on electronics indicators, but this makes virtually no difference to the results. Similarly, switching the positions of inverted inventories and chip prices in the system leave the impulse response functions qualitatively unchanged.

The estimated impulse response functions are depicted in Figures 1–5. Following the advice of Sims and Zha (1999), we have included 68% asymptotic confidence intervals to gauge the statistical significance of the responses.[9] In every case, unanticipated shocks to the leading indicators produce significant movements in world semiconductor sales. The time horizon over which the dynamic adjustment paths of chip sales are plotted following the innovations to each of the leading series extends to 24 months, by which time the responses are all insignificantly different from zero.

[9]Bootstrapped standard errors did not materially alter the widths of the intervals.

Figure 1: Impulse Response of Global Chip Sales to Nasdaq Shock

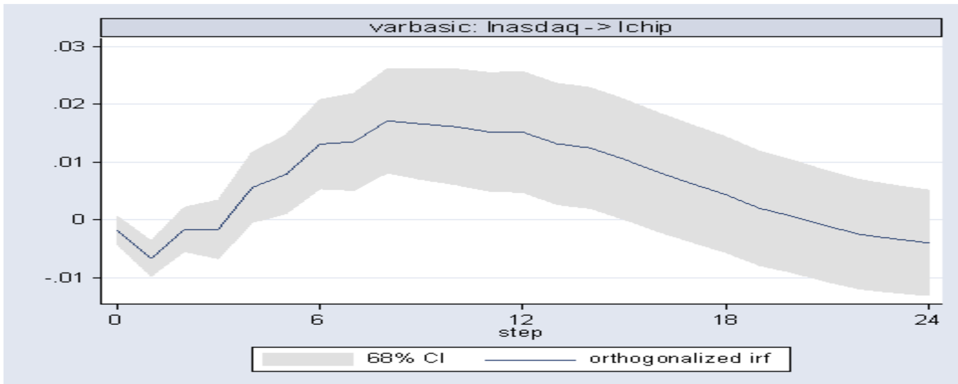

Figure 2: Impulse Response of Global Chip Sales to Book-to-Bill Shock

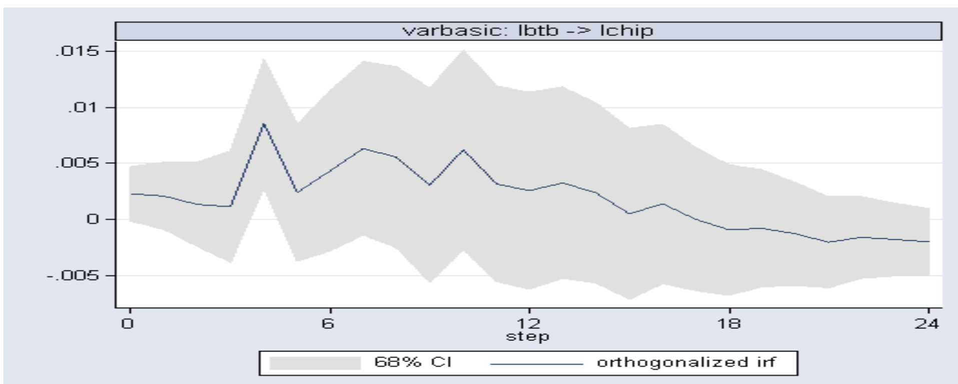

Figure 3: Impulse Response of Global Chip Sales to New Orders Shock

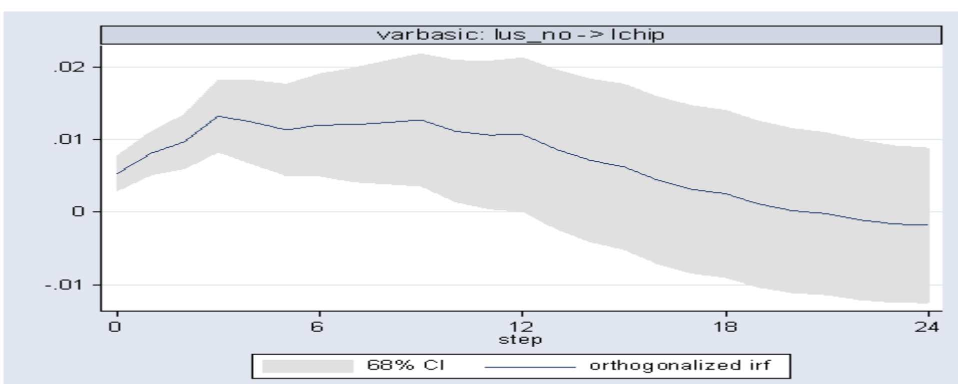

Figure 4: Impulse Response of Global Chip Sales to Inventory Shock

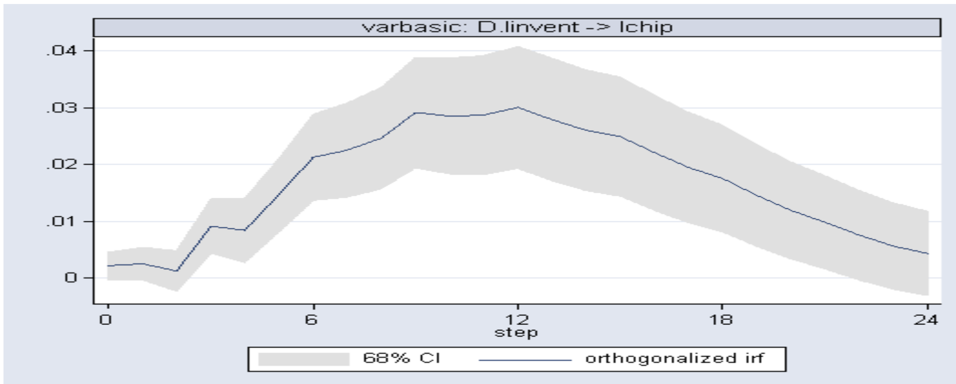

Figure 5: Impulse Response of Global Chip Sales to DRAM Price Shock

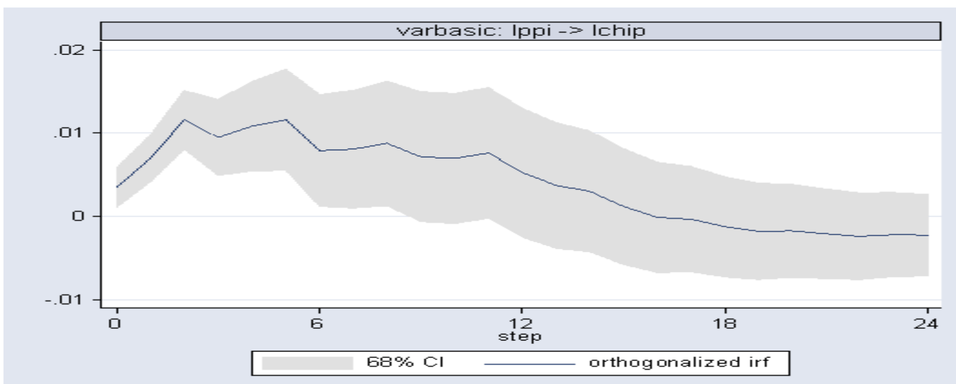

Apart from the response for the book-to-bill ratio, the graphs share the same hump-shaped feature so often observed in the impulse responses reported in business cycle studies. In our context, this characteristic demonstrates the leading qualities of the electronics indicators, including the US new orders series. The indicators differ, however, on the number of months it takes for the dynamic response of global chip sales to reach a peak, which gives us an idea of the average lead in a series. The impulse responses indicate that the average lead for the inverted change in inventories, at 10–12 months, is the longest, followed by the Nasdaq index at 9 months. The lead times for new orders and DRAM prices coincide at about 3–4 months. And though the response of chip sales to a shock in the book-to-bill ratio has an irregular shape and is insignificant for the most part, there is a prominent spike corresponding to an average lead time of 5 months. To sum up, the impulse response functions from the VAR model confirm that all the selected indicators presage world electronics activity, albeit with different lead times.

4 Forecast Performance of VAR Model

The VAR model in (1) incorporating our five leading indicators is next used to generate *ex ante* forecasts of global chip sales. Since our primary concern is in predicting the state of the electronics cycle rather than its growth rates, the forecasts are generated from the model specified in levels.[10] We will compare the predictive performance of the VAR with two alternative models of chip sales. The first is the univariate autoregressive (AR) process, which is a frequently used benchmark model. The presence of a unit root in the sales series suggests modelling in logarithm first differences, and the following AR model of order 5 was found to fit the data well:

$$\triangle y_t = \tau + \sum_{k=1}^{5} \phi_k \triangle y_{t-k} + \varepsilon_t \tag{5}$$

[10] We eschew the VECM specification for simplicity and to avoid misspecification errors that might arise from estimated cointegrating relationships, albeit at the expense of some loss in efficiency.

The forecasts of chip sales from this model are converted into levels for comparison with the VAR model.

The second forecasting model we consider is a bivariate specification involving a composite index derived from the leading indicators. As mentioned at the beginning, it is customary to combine leading series into a composite index to give a summary measure of their movements. Using the methodology employed by The Conference Board for compiling the US Leading Index, we constructed a similar index for the global electronics cycle.[11] This leading index (z_t) was found to be cointegrated with global chip sales (y_t), motivating us to build a bivariate VAR model in the logarithm levels of these two series. Modelling in levels instead of differences facilitates comparison with the multivariate VAR model. Both the AIC and the Hannan-Quinn criterion selected an optimal lag length of 6 for the leading index model, hence we estimate these two equations:

$$\begin{aligned} y_t &= \tau_1 + \sum_{k=1}^{6} \phi_{1k} y_{t-k} + \sum_{k=1}^{6} \alpha_{1k} z_{t-k} + \varepsilon_{1t} \\ z_t &= \tau_2 + \sum_{k=1}^{6} \phi_{2k} y_{t-k} + \sum_{k=1}^{6} \alpha_{2k} z_{t-k} + \varepsilon_{2t} \end{aligned} \tag{6}$$

[11] This entails the computation of symmetrical month-to-month percentage changes in each indicator, followed by a standardisation process to prevent the more volatile series from dominating the rest. These are then summed to yield the monthly percentage changes in the composite index, thus effectively assigning equal weights to each component. Finally, the index levels are derived recursively after setting the first month's value of the index to 100.

It is fortuitous to have a common lag length for all three models in terms of predicting the level of global chip sales, as this enhances forecast comparability. For the purpose of evaluating each model's predictive performance, we divided our data set into two parts. The first spans the period from 1992:3 to 2003:1 and was used only for estimation; the remaining 12 data points, spanning 2003:2 to 2004:1, were used for post-sample prediction. We do not use a longer post-sample prediction period in view of the shortness of the data series as well as the size of the VAR model. Forecast horizons of one, three and six months are considered. Reflecting what a forecaster would be able to do in practice, we estimated each model recursively so that the forecast for time $t + h$ is computed with data up to time t.

As is conventional in the literature, we use the root mean square prediction error (RMSE) and the mean absolute prediction error (MAE) as measures of forecast accuracy. The results from the univariate AR model serve as a yardstick against which we measure the predictive abilities of the other two models; that is, we compute the ratio of the latter's RMSE or MAE to that of the AR model. Whenever the relative RMSE or MAE of the VAR or leading index model is smaller (larger) than one, its forecasting performance is better (worse) than the benchmark model. Table 3 reports the relative RMSE and MAE associated with the out-of-sample forecasts of global chip sales generated from the VAR and index models.

Table 3: Forecast Performance of VAR and Index Models

	Relative RMSE		Relative MAE	
Forecast Horizon	VAR	Index	VAR	Index
1 month	1.098	0.925	1.103	0.939
3 months	0.861	0.906	0.833	0.845
6 months	0.882	0.812	0.682	0.721

Note: Relative RMSE or MAE is expressed as a ratio to the univariate AR model.

The inclusion of information from the leading indicators in the VAR and index forecasting models clearly leads to an improvement in predictive accuracy over the benchmark AR model, especially at the 3 and 6 months forecast horizons. However, it does not come as a surprise that the 1-step ahead forecasts from these models only improved moderately or even worsened *vis-à-vis* the univariate model, since ARIMA models are known to produce very accurate forecasts in the short term. As for the relative predictive performances of the VAR and index models, we do not observe any one model being consistently superior to the other: at the 3 months horizon, the VAR model fares unambiguously better but at the 6 months horizon, it outperforms the index model only in terms of the MAE criterion.

To ascertain if the differences in predictive accuracy found between the models are statistically significant, we conduct formal tests of forecast performance. In particular, we employ the following Diebold-Mariano (1995) test statistics (DM) and its small sample version ($DM^{\dagger}$) proposed by Harvey, Leybourne and Newbold (1997):

$$DM = \frac{\bar{d}}{\sqrt{V(\bar{d})}} \tag{7}$$

$$V(\bar{d}) = \frac{1}{T}\left(\hat{\gamma}_0 + 2\sum_{k+1}^{h-1}\hat{\gamma}_k\right)$$

$$DM^{\dagger} = \sqrt{\frac{T+1-2h+h(h-1)/T}{T}}\, DM \tag{8}$$

where T is the number of forecasts made, h is the forecast horizon, $\bar{d}$ is the sample mean of the differences between the squared or absolute forecast errors from any two competing models, $V(\bar{d})$ is the approximate asymptotic variance of $\bar{d}$, and $\hat{\gamma}_k$ is the estimated kth order autocovariance of the forecast error differences. These test statistics are shown in Table 4 and compared with critical values from the t-distribution with $T-1$ degrees of freedom.

It is evident from the table that where the 1-month ahead forecasts are concerned, there is no appreciable difference in forecast performance between the three competing models as all the test statistics turned out to be insignificant. At the 3 months forecast horizon, however, the DM tests indicate that the VAR and index models deliver significantly more accurate predictions than the univariate AR model. The corresponding $DM^{\dagger}$ statistics are marginally insignificant for squared forecast errors and significant only in the case of the absolute forecast errors generated by the index model. Interestingly, the hypothesis of equal predictive ability between the VAR and index models cannot be rejected at the 10% significance level at the same horizon.

Table 4: Predictive Accuracy Tests

	DM		$DM^{\dagger}$	
	Sq. Errors	Abs. Errors	Sq. Errors	Abs. Errors
$h = 1$				
VAR vs AR	0.511	0.395	0.490	0.378
Index vs AR	−1.170	−0.721	−1.072	−0.691
VAR vs Index	0.834	0.879	0.764	0.841
$h = 3$				
VAR vs AR	-1.641^{*}	-1.385^{*}	−1.228	−1.037
Index vs AR	-1.703^{*}	-2.426^{*}	−1.274	-1.815^{*}
VAR vs Index	0.004	0.095	0.003	0.071

Notes: * denotes significance at the 10% level. The one-tailed critical values for the 1 and 3 months forecast horizons are −1.363 and −1.383 respectively.

The Diebold-Mariano test statistics are undefined for the 6-steps ahead forecasts because $V(\bar{d})$ took on a negative value in every case, requiring the evaluation of the square root of a negative number in equation (7). In such pathological situations, Diebold and Mariano (1995) suggest that the null hypothesis of equal forecast accuracy be rejected. To help us infer which is the outperforming model, we rely on the summary measures in Table 3 and visual inspection of the 6-months ahead forecasts generated by the three models, which are plotted alongside global chip sales in Figure 6. While all the models missed the strong upturn in semiconductor sales during the post-sample sub-period 2003:7–2003:10, the VAR arguably gives the best visual forecasts. The plot, and the RMSE and MAE measures, indicate that the benchmark AR model is inferior to both the VAR and index models for forecasting at the 6 months horizon. But as we saw earlier, the relative ranking of the predictive ability of these two models is ambiguous, being dependent on the metric used for evaluation in Table 3. Figure 6 shows that the forecasts from the leading index model appear to be slightly worse than those of the VAR model.

A priori, it is difficult to predict which model will do better. On the one hand, the VAR model extracts information from a diversified set of leading indicators, thereby obviating the need to form a composite index and avoiding the problems associated with index construction, such as the weights to be assigned to the component indicators. The flexibility of the VAR model means that the different lead times of indicators are accommodated and it also frees us from the implicit assumption of a single common factor underlying the movements in the indicators. On the other hand, forecasting with an index results in a much more parsimonious model—the number of autoregressive parameters drops from 216 to 24—hence averting the overfitting problem and yielding more efficient estimates of the parameters in the index model. Moreover, when the mixed signals provided by the individual electronics indicators are caused by measurement errors and random disturbances, the use of a single composite index leads to noise reduction. That said, our results show that the VAR generates forecasts with an accuracy that can rival the predictions from the index model. We surmise that the greater flexibility of the VAR model offsets its less parsimonious structure, thereby resulting in gains to forecasting in practice.

Figure 6: Forecasts Comparisons for 6 Months Forecast Horizon

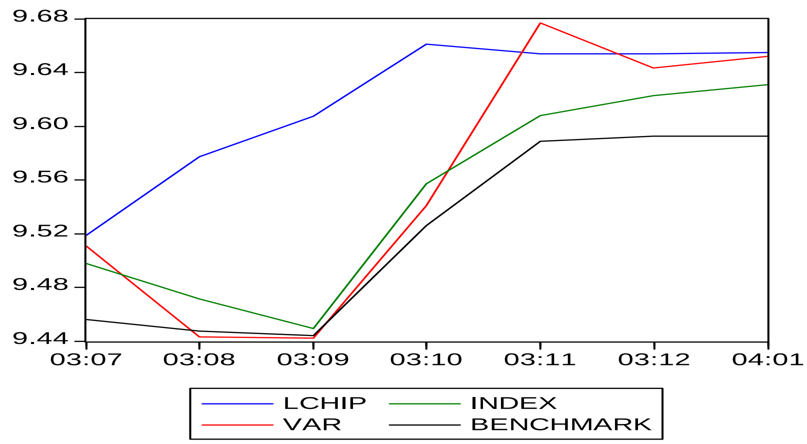

5 Conclusion

In this study, we identified from a list of frequently monitored electronics indicators five monthly leading series that are economically significant and show the potential to presage global semiconductor sales. These are the Nasdaq composite index, the semiconductor industry book-to-bill ratio, US new orders of electronics, inverted changes in US electronics inventories, and DRAM chip prices. We then construct for this set of leading indicators and our chosen coincident indicator of the global electronics cycle a VAR model that reflects the dynamic interactions in the electronics market.

Besides providing a natural framework for performing Granger causality tests which establish the leading qualities of most of the selected indicators, the VAR system is also used to characterize the dynamic paths of adjustment of global chip sales in response to orthogonalized shocks in each of the leading series. These impulse response functions with their hump-shaped features confirm that our chosen set of electronics indicators presage the world electronics cycle by distinct lead times.

From a methodological point of view, the principal objective of adopting a VAR approach is to provide a unified framework for forecasting the global electronics cycle with leading indicators, without having to make the restrictive assumption of a single common factor underlying the movements in the indicators. To this end, post-sample predictions of global chip sales were generated from the VAR model and their accuracy compared with forecasts from two alternative models—a benchmark AR model and a model which uses a composite index constructed from the same set of leading indicators. An evaluation based on standard measures of forecast accuracy and formal tests of predictive ability suggests that the VAR model's forecasting performance is superior to that of the benchmark model, and is comparable to that of the composite index model. Our results are therefore in contrast to recent studies that compare the relative forecasting efficacy of index and VAR models, and find that index models generally predict better (Camba-Mendez et al., 2002; Bodo et al., 2000).

Although we conclude that the proposed VAR model incorporating our set of identified leading indicators is useful for forecasting the global electronics cycle, there is scope for further work. For one thing, one might want to consider the ability of the model to anticipate turning points in the global electronics cycle. Forecasters in the electronics industry might be more interested to predict the timing of peaks and troughs rather than in the type of quantitative forecasts that we focused on in this paper. We did not address this issue partly because of the paucity of turning points in our relatively short sample period, but also due to the inherent difficulty of defining cyclical turning points. Nonetheless, future research along these lines is warranted. Another possible extension of this study is to explore the forecasting power of Bayesian VAR (BVAR) models based on our set of leading indicators. The use of the BVAR as a more parsimonious alternative to the VAR might just strike the right balance between the objectives of flexibility and noise reduction.

Acknowledgements

This paper has its origins in a joint study with the Monetary Authority of Singapore (MAS). We are grateful to the MAS for providing us with data. However, the views expressed in this paper are those of the authors and should not be attributed to the MAS.

References

Bodo, G., Golinelli, G., & Parigi, G. (2000). Forecasting Industrial Production in the Euro Area. *Empirical Economics 25*, 541—561.

Camba-Mendez, G., Kapetanios, G., Weale, M.R., & Smith, R.J. (2002). The Forecasting Performance of the OECD Composite Leading Indicators for France, Germany, Italy, and the U.K. In Clements, M.P., & Hendry, D.F. (Eds.), *A Companion to Economic Forecasting*, Blackwell Publishers, Oxford, pp. 386—408.

Cheung, Y., & Lai, K.S. (1995). Lag Order and Critical Values of a Modified Dickey-Fuller Test. *Oxford Bulletin of Economics and Statistics 57*, 411–419.

de Leeuw, F. (1991). Toward a Theory of Leading Indicators. In Lahiri, K., & Moore, G.H. (Eds.), *Leading Economic Indicators: New Approaches and Forecasting Records*, Cambridge University Press, Cambridge, pp. 15–56.

Diebold, F.X., & Mariano, R.S. (1995). Comparing Predictive Accuracy. *Journal of Business and Economic Statistics 13*, 253–263.

Elliot, G., Rothenberg, T.J., & Stock, J.H. (1996). Efficient Tests for an Autoregressive Unit Root. *Econometrica 64*, 813–36.

Harvey, D., Leybourne, S., & Newbold, P. (1997). Testing the Equality of Prediction Mean Squared Errors. *International Journal of Forecasting 13*, 281–291.

Hobijn, B., Stiroh, K.J., & Antoniades, A. (2003). Taking the Pulse of the Tech Sector: A Coincident Index of High-Tech Activity. Federal Reserve Bank of New York, *Current Issues in Economics and Finance 9*, 1–7.

Koch, P.D. & Rasche, R.H. (1988). An Examination of the Commerce Department Leading-Indicator Approach. *Journal of Business and Economic Statistics 6*, 167–187.

Ng, S., & Perron, P. (2001). Lag Length Selection and the Construction of Unit Root Tests with Good Size and Power. *Econometrica 69*, 1519–1554.

Ng, Y.P., Tu, S.P., Robinson, E., & Choy, K.M. (2004). Using Leading Indicators to Forecast the Singapore Electronics Industry. *MAS Staff Paper No. 30.* Available at http://www.mas.gov.sg.

Sims, C.A., & Zha, T. (1999). Error Bands for Impulse Responses. *Econometrica 67*, 1113–1155.

Sims, C.A., Stock, J.H., & Watson, M.W. (1990). Inference in Linear Time Series Models with Some Unit Roots. *Econometrica 58*, 113–144.

Toda, H.Y., & Yamamoto, T. (1995). Statistical Inference in Vector Autoregressions with Possibly Integrated Processes. *Journal of Econometrics 66*, 225–250.

Veloce, W. (1996). An Evaluation of the Leading Indicators for the Canadian Economy Using Time Series Analysis. *International Journal of Forecasting 12*, 403–416.

Zarnowitz, V. (1992). *Business Cycles: Theory, History, Indicators, and Forecasting*, NBER Studies in Business Cycles Volume 27, The University of Chicago Press, Chicago.