

Temporal Big Data for Tire Industry Tactical Sales Forecasting

Yves R. Sagaert^{a,c,*}, El-Houssaine Aghezzaf^a, Nikolaos Kourentzes^b, Bram Desmet^c

^a*Department of Industrial Management, Ghent University, Technologiepark 903, Zwijnaarde 9052, Belgium*

^b*Department of Management Science, Lancaster University Management School, Lancaster, LA1 4YX, UK*

^c*Solventure NV, Kranepoelpad 4, Aalter 9880, Belgium*

Abstract

We propose a forecasting method to improve accuracy for tactical sales predictions at a major supplier to the tire industry. This level of forecasting serves as direct input for the demand planning, steering the global supply chain and is typically up to a year ahead. The case company has a product portfolio that is strongly sensitive to external events. Univariate statistical methods, which are common in practice, are unable to anticipate and forecast changes in the market, while human expert forecasts are known to be biased and inconsistent. The proposed method is able to automatically identify key leading indicators that drive sales from a massive set of macro-economic indicators, across different regions and markets and produce accurate forecasts. Our method is able to handle the additional complexity of the short and long term dynamics from the product sales and the external indicators. We find that accuracy is improved by 16.1% over current practice with proportional benefits for the supply chain. Furthermore, our method provides transparency to the market dynamics, allowing managers to better understand the events and economic variables that affect the sales of their products.

Keywords: Forecasting, time series, regression, temporal big data, supply chain planning

Top management decision making requires long term tactical-strategic forecasts. These forecasts are very difficult to formulate and thus are often inaccurate (Fildes et al., 2008). Naturally, improved long-term forecast accuracy is of paramount importance to management. Current practice is to rely on expert judgement for such predictions. However, this poses two challenges: experts are well known to be biased with detrimental effects to accuracy; and experts are often not able to consider a large number of external indicators simultaneously and weight them appropriately, focusing on the ones that are subjectively considered as most critical (Fildes and Goodwin, 2007).

Supply chain managers have access to an increasing amount of external data, holding great potential for describing and quantifying the dynamics of their markets. Based on insights coming from external data, for example macro-economic indicators, managers could assess future evolution of the markets and the impact on their supply chain. Yet, it is very difficult to do this without analytical support on how to use it appropriately. Nonetheless, experts believe this information is important. Thirty percent of the respondents in an in-

*Yves R. Sagaert

Email address: yves.sagaert@ugent.be (Yves R. Sagaert)

dustry survey by Russom et al. (2011) identified that capturing the market sentiment trend is one of the unexploited benefits from big data.

In this paper we propose a solution to this problem. We describe a forecasting framework that automatically identifies key leading indicators to predict sales from a massive set of macro-economic indicators and thus tracking changes in the strategic business environment and producing superior forecasts. We apply our temporal big data solution to a supplier to the tire industry and demonstrate how this improves forecast accuracy for longer tactical horizons.

To illustrate the potential of macroeconomic leading indicators, we give following example in the tire market, which is highly sensitive to external effects. An economic upswing causes increased road haulage, and more trucks transporting goods makes truck tires to wear faster. As a consequence, replacement tire consumption goes up. However, the tire industry can be divided in two main end market segments: replacement tires and tires for new vehicles. If the economic upswing is lasting, companies are planning for extra capacity, which results in new investments and increasing tire sales for new trucks. The combined effect has an impact on both sales volume and the timing of the sales. Truck tires are replaced after the old ones get worn and therefore lags behind the increased the economic activity. This in turn has the potential to be a leading indicator of tire replacement sales. Analogous complex dynamics are observed with other commercial and passenger car tires.

In this, the gross domestic product is often used as a proxy for economic activity. However such an aggregate variable does not detail changes in the different sectors and activities in the economic system. On the other hand, using more detailed macro-economic indicators that capture activities in different areas, such as education, unemployment and healthcare, can highlight information that can play a role when families are replacing their car tires or buying a new car. Therefore, the number of potentially useful leading indicators increases substantially.

The number of interesting variables can increase even further, as the case company has a global supply chain, supplying several major tire manufactures in different markets and countries. Different indicators will describe these various aspects separately. The size of the problem of identifying and using the relevant information from so many variables, as well as the predictive nature of it makes it challenging, particularly if experts have to rely solely on their judgement; yet, these indicators could improve long-term forecasts.

The tactical planning horizon of the case company is 12 months ahead, which requires forecasting sales for the same horizon. Good sales forecasts are essential as this enables the management to optimise the supply chain. The orders are placed 1 to 2 months in advance, but these are sometimes altered one week prior to delivery. However, on the supply side, the procurement department requires raw material orders 4-6 months in advance to be able to conduct price negotiations. When the final order is placed, it takes 3-4 months to deliver the raw materials. This makes the total lead time of raw materials to vary between 7 and 10 months. Moreover, in order to optimise capacity utilisation, the future demand should be known 9-12 months in advance. This makes it possible to free the necessary additional capacity among the plants worldwide. Based on the tactical sales forecasting, the supply chain management decides on production planning, capacity, inventory and procurement on a tactical level.

The long lead times make estimating future demand challenging. From a statistical

forecasting perspective extrapolative forecasting methods, such as exponential smoothing, are widely used to produce forecasts (Weller and SF, 2012). Such methods extract time series features, like seasonality and trend, and use them to predict future sales. Obviously without additional information these cannot anticipate external events and they cannot foresee an economic downturn or upswing. Therefore these methods are unfit for tactical sales forecasting.

Although expert can adjust the baseline statistical predictions with additional market insights, this is not always feasible or accurate. Long term expert forecasts are notoriously difficult because of the complexity of the economic system and the biased and inconsistent nature of human judgement (Fildes and Goodwin, 2007). The Delphi method can be used to reduce bias and increase the accuracy of such forecasts (Goodwin and Wright, 2010). However, using multiple experts can be very expensive and since it cannot be automated it is not scalable to global portfolios that include several different product and market combinations. Alternatively, supply chain management could rely on external market growth forecasts for different countries by analysts. Irrespective of the quality of analyst predictions, these are often not available on the detail that is required to feed to the tactical forecasts of organisations. Furthermore, research has shown that field knowledge is not always incorporated in the most effective way through human judgement (Petropoulos et al., 2016) and that the forecasting accuracy gain of these managerial adjustments is inconsistent (Trapero et al., 2013). As a result, our motivation for this work is to provide a statistical method, that is not limited in the same ways, to incorporate external data consistently.

This need was identified by the company where this research is implemented. They asked our team to help improve the existing statistical sales forecasts by incorporating macro-economic leading indicators in a fully automated way. To achieve this we developed a statistical forecasting framework that is able to deal with massive datasets of external indicators, and connected it to a fully automated data stream of exogenous macro-economic information.

A crucial issue in resolving this is to identify the most relevant leading indicator variables for the sales of the particular company. The available exogenous macro-economic indicators are numerous and a lot of them are publicly available by various statistics agencies and central banks around the world. However, due to limited monthly sales history, which is common in practice, we face the so called ‘short and fat data problem’, with $p \gg n$, where p is the potential predictors and n is the available data points. There are insufficient degrees of freedom to estimate the full model. Furthermore, Verleysen and François (2005) note that traditional algorithms do not scale to the exogenous data due to the curse of dimensionality. This means that calculation time and memory rapidly increase with the increase of number of dimensions of the data.

Here we are not only trying to identify exogenous macro-economic variables that are causal to the company sales, but also require that such information is leading, i.e. predict the changing behaviour of sales multiple months in advance. Even if the optimal set of indicators was known, the leading order of is unknown. Therefore, the macro-economic indicators should be shifted in time, and tested for relevant leading effects. Adding one leading effect to all the available indicators doubles the size of the database. Since the amount of potential shifts in time is unknown, the number of different leads (1) makes the original big set of indicators to grow substantially. We refer to this as temporal big data,

incurring a much bigger selection problem in the order of $pl \gg p \gg n$. Figure 1 gives insight how the temporal big data is build and inserted into the model.

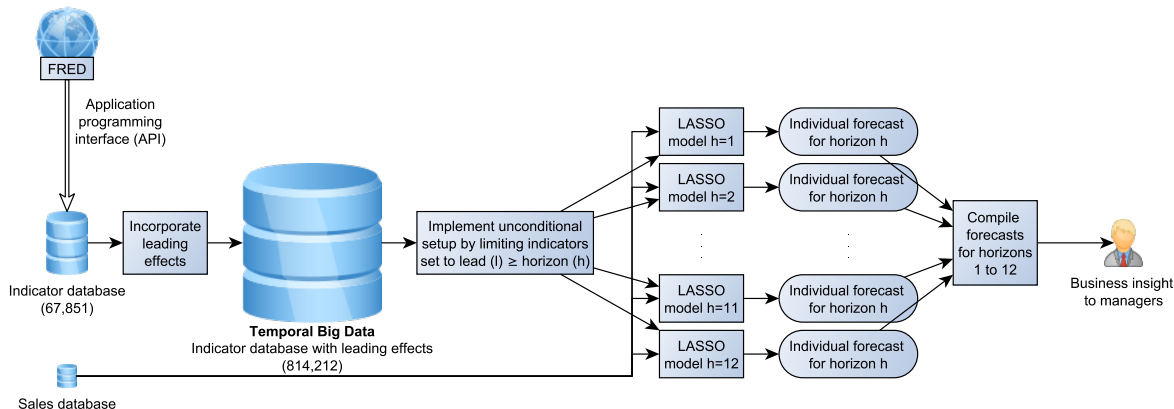


Figure 1: Algorithm flowchart: data trajectory of Temporal Big Data.

Since our forecasts are conditional on the known values of the explanatory indicators the set of available variables changes with the forecast horizon, requiring a new forecasting model for each (Figure 1). To understand this better, suppose we need to forecast only one period ahead. All p explanatory variables with lead orders from 1 to l periods can be used, allowing to capture both short and long dynamics. This results in a model that has to select from pl number of potential inputs. On the other hand, when forecasting six months ahead we cannot use short lead orders, from 1 to 5 months, as such information would require knowledge of the future. In turn that means that the forecasting model has available only $(l - 5)p$ variables to choose from, capturing medium to long term effects. Therefore, to produce tactical sales forecasts from 1 to 12 months ahead requires 12 individual models, each choosing the appropriate leading indicators from a different set of variables, raising the complexity of producing automatic statistical forecasts twelvefold. Naturally if the selected predictors are important we expect that these will be chosen by several models, albeit with different lead orders.

In the case company implementation we limit the analysis to monthly macro-economic leading indicators, as the demand forecasts are also reviewed in monthly periods. The available monthly historical data contain 66 monthly data points, based on which we want to select the relevant leading macro-economic indicators. We collected 67,851 macro-economic indicators from 25 sections of the Federal Reserve Economic Data (FRED) database of the federal reserve bank of St. Louis. The FRED database is a well-established source of information, which addresses quickly changing economic trends with continuous publications. Since we are interested for predictions up to 12 months ahead, we construct variables with leads from 1 to 12 months. In this way, the resulting $pl \gg n$ problem contains 814,212 $(67,851 \times 12)$ variables and only 66 data points, which is impossible to solve with conventional causal regression predictive modelling.

In order to address this we use Least Absolute Shrinkage Selection Operator (LASSO) regression (Tibshirani, 1996, Efron et al., 2004). This has been widely used in several applications. However when modelling time series data, often the univariate autoregressive

information is not considered and models focus solely on the exogenous information. This is a severe limitation, as the univariate autoregressive information contains potential sales dynamics.

Therefore the methodology we develop explicitly allows to capture relevant seasonality, autoregressions, macro-economic indicators and their leading effect. As we mentioned the $pl \gg n$ problem cannot be solved by conventional regression. Even if the number of variables was much smaller, including different macro-economic indicators via multiple linear regression suffers from multicollinearity, i.e. multiple variables contain similar information. This does not permit producing reliable forecasts, as there is limited confidence in the estimated model parameters. Introducing multiple leading effects of the same indicator increases the multicollinearity even further. This has been a main motivation for us to use LASSO.

LASSO can effectively deal with a large number of variables by performing variable selection in an alternative way to conventional regression modelling. LASSO uses an altered optimisation cost function to the Ordinary Least Squares regression for the estimation of the model coefficients. The new cost function balances two competing objectives: the mean squared fitting error and the sum of absolute values of the coefficients multiplied by a regularisation term, which is a scalar. From the one hand, it attempts to get the best fit to the data, by minimising the fitting error, like conventional regression. From the other hand, by increasing the cost for coefficients with large values, it effectively shrinks them towards zero. Therefore, only very informative variables will end up having non-zero coefficients, and be included in the model. All other coefficients are forced to zero as otherwise they would increase the cost. Furthermore, by shrinking the coefficients LASSO avoids over-fitting and tackles multicollinearity, as both will occur only when the coefficients can take large values in an unrestricted fashion. Due to way that variables are included in the model, LASSO makes it possible to identify informative variables even when $pl \gg n$, and due to its conservative nature, the final model contains only a few but relevant variables. The regularisation term controls the strength of the shrinkage and balances the two competing objectives of the cost function. This is identified automatically by cross-validation. As various unit measures can cause different sized variable coefficients, which will affect the penalty part of the cost function, all the input variables are standardised, as well as the target series (Friedman et al., 2001). A final advantage of LASSO is that the final model is similar in nature to conventional regression and can be interpreted in the same way, thus being transparent to the users.

We coded the statistical forecasting method using the R platform and set up a separate SQL database that retrieves all required data from FRED via an application programming interface (API). This interface enables full automatic synchronisation between the public database and our private SQL database. The monthly updates of the database include updating the new values of the macro-economic indicators, as well as separately updating adjusted historical data points. The statistical forecasts are provided to the supply chain managers, together with the list of selected relevant macro-economic indicators and their regression coefficients.

We benchmark our solution against the existing company univariate forecasting method: Holt-Winters, which allows for trend and multiplicative seasonality. We also use as a benchmark a conventional predictive causal regression that combines any relevant univariate information, such as seasonality, and exogenous information. To avoid multicollinearity, we

only select the maximally correlated indicator with the sales series in this regression. Table 1 summarises the results over all time series, for 7 consecutive forecasting cycles. The Mean Absolute Percent Error (MAPE) is shown as forecasting accuracy for each forecast horizon individually, and overall. The best performing forecast is highlighted in boldface. Figure 2 provides a visual overview of the forecast accuracy.

Table 1: Summarised MAPE results.

MAPE for each horizon	1	2	3	4	5	6	7	8	9	10	11	12	Overall
Company benchmark	11.0	13.8	13.7	16.1	16.4	17.3	18.5	19.7	21.8	25.2	25.8	24.0	18.6
Linear regression	15.7	17.1	17.2	17.9	15.4	16.5	18.1	16.2	19.4	29.0	30.7	30.3	20.3
LASSO	13.3	13.6	13.5	13.8	13.7	13.9	14.2	16.5	17.5	17.9	20.7	18.2	15.6




Figure 2: Mean Absolute Percent Error (MAPE) evolution over forecast horizon.

We observe that in the majority of cases LASSO outperforms the other forecasts with a substantial difference. The differences increase for longer forecast horizons, demonstrating the benefits of using leading indicators in the calculation of the tactical level sales.

Figure 3 shows a detailed analysis the forecast errors, for each forecast horizon. The boxplots show the error distribution for each case, showing the minimum, first quartile, median, third quartile and maximum errors. Overall, this figure is in line with the MAPE results. Here, we can also see that the errors of LASSO have a smaller dispersion, particularly for longer horizons that are of interest.

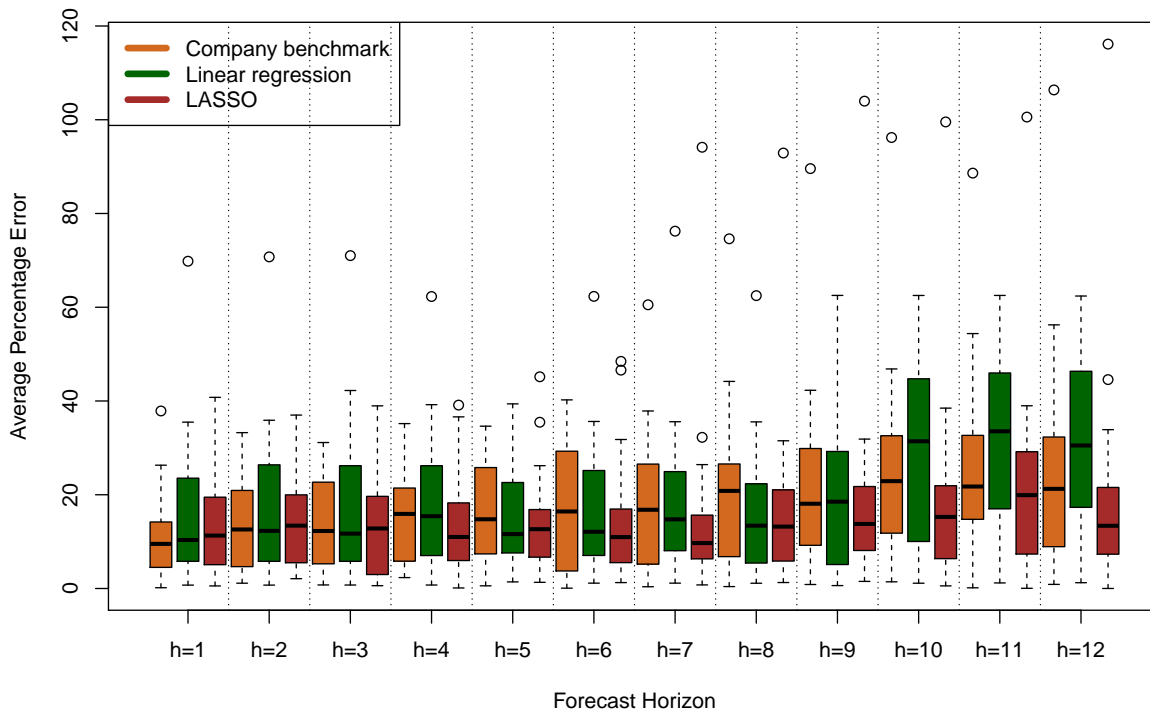


Figure 3: Analysis of average percentage errors over rolling origins.

The results of this implemented research are twofold: first, the tactical sales forecasts improve over the company benchmark by 16.1%. The linear regression, combining seasonality with one leading indicator on average does not improve on the company benchmark. Second, the forecasting model provides insights in the relevant macro-economic leading indicators and their individual contribution to the forecast. The tactical sales time series consist of product groups for passenger car tires and trucks tires. The forecasts are made for the US and Europe markets.

Table 2 exhibits the type of indicators that were chosen over all series, over several horizons and several origins. The selection percentage that indicators from these groups were selected are shown for each type. Likewise, the most occurring indicator units are displayed for each group. We notice that the selected indicators originate most from the groups of labour, consumer price index and government services. The most frequent type of indicators are trade, private services and manufacturing.

In figure 4, we review the selected leading order of the selected indicators. As outlined in figure 1, leading effects are only available when forecasting a horizon that is smaller or equal than this leading effect. Due to this unconditional setup, the leading effect 12 can be selected for all horizons, while the leading effect 1 can only be selected for one horizon. In figure 4, we adjust for this effect. We observe that the leading order of 12 months is selected frequently. The group of second most selected leading effects is 1, 3, 6, 8, 9, 10 and 11 months ahead. Leading effects that are multiples of 3 can be explained by leading quarterly

Table 2: Most selected types of indicators.

Type of indicator	Percentage selected	Most frequent indicator units
Labour	18.5	Persons, USD/Week, %, % change, Hours
Consumer price index	12.4	Index, Growth rate, %, % change
Government services	9.6	Persons, % change, Index
Retail trade	9.3	Persons, Index, Hours
Financial Activities	8.5	Persons, Index, % change, USD
Private services	8.3	Persons, %, % change
Transport and Automotive	6.4	Persons, USD, Index, % change
Manufacturing	5.8	Hours, USD, Net %, Index
Food Manufacturing	5.0	Persons, Rate, % change
Education & Health services	4.4	Persons, % change
Wholesale Trade	3.6	Persons, USD, %, USD/Week
Tourism	2.6	Persons
Recession Indicator	1.5	Index
Construction	1.1	Hours, Net %, Index, USD/Week
Mining	1.0	Persons, Hours, % change
International trade	1.0	Growth rate, USD, National currency, Index
Exchange rate	0.6	Index, %, National currency units/USD
Housing	0.2	Persons, % change, Units
Investment goods	0.1	%, USD

dynamics of the indicators. The effect of 1 month ahead could indicate a short-term close related effect.

The resulting LASSO forecast is transparent, providing insight in the selected leading indicators. Experts can benefit from gaining additional market understanding, This can help experts refine their understanding of market dynamics and interactions. Furthermore, by gaining new insights in the relevant leading indicators, sales people are given a tool to better manage and react to markets. This can motivate the sales department to increased involvement in the forecasting process. The supply chain director, of the case company, expects on a longer term further forecast accuracy improvements through closer interaction between departments. Moreover, it allows experts to understand and validate the produced forecasts by investigating the selected indicators. Note that the selection is done both for the nature and the lead order of the leading indicators fully automatically, addressing the temporal big data challenge of this particular predictive problem.

Finally, as the forecasts are model based we can produce prediction intervals. In contrast, judgemental managerial forecasts do not have this property. Prediction intervals are important as they permit gauging the uncertainty of the forecasts. This is paramount for inventory management and safety stock calculations. Prediction intervals can be used in the decision process for robust planning, as they allow assessing the risk of future scenarios. The resulting framework allows supply chain managers to better manage inventory, production scheduling and work-in-capital in their Sales & Operations Planning process.

The supply chain director has expressed his expectation that the inventory costs will linearly decrease with increase of forecast accuracy. This corresponds to prior research by Hosoda et al. (2008) and Trapero et al. (2012), who incorporated external data from an Electronic Point Of Sales (EPOS) system. He also expects supply chain savings due to




Figure 4: Most selected leading effects among the indicators.

more stability in the planning of manufacturing and logistics, that will allow for a more cost efficient approach. For example, logistics costs can be decreased by avoiding emergency shipping.

Further research for refining this tool aims at including weather effects in these forecasting models. Whereas winter tires are produced in the summer period, a low sales volume of winter tires will result in a lower need for raw materials supply in the next summer. Although one would normally expect these patterns to be captured in the seasonality of sales, extreme or changing weather patterns are expected to cause additional variability in sales. Quantifying and adjusting for global sudden external events is another next step for this tool. We expect these global sudden events to even further improve the model. For example, the change in import tariffs on Chinese tires in July 2015 caused a reduction in import of Chinese tires in the US. As a consequence, US production rose, causing raw materials demand to increase. Although the implementation of this solution was customised to the case company, it has potential for other firms and sectors.

Acknowledgements

This research is funded by IWT-Vlaanderen (project IWT 130843) and Solventure NV.

Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. *The Annals of statistics*, 32(2):407–499.

- Fildes, R. and Goodwin, P. (2007). Against your better judgment? How organizations can improve their use of management judgment in forecasting. *Interfaces*, 37(6):570–576.
- Fildes, R., Nikolopoulos, K., SF, C., and AA, S. (2008). Forecasting and operational research: a review. *Journal of the Operational Research Society*, 59(9):1150–1172.
- Friedman, J., Hastie, T., and Tibshirani, R. (2001). *The elements of statistical learning*. Springer series in statistics Springer, Berlin.
- Goodwin, P. and Wright, G. (2010). The limits of forecasting methods in anticipating rare events. *Technological Forecasting and Social Change*, 77(3):355 – 368.
- Hosoda, T., Naim, M. M., Disney, S. M., and Potter, A. (2008). Is there a benefit to sharing market sales information? Linking theory and practice. *Computers & Industrial Engineering*, 54(2):315 – 326.
- Petropoulos, F., Fildes, R., and Goodwin, P. (2016). Do 'big losses' in judgmental adjustments to statistical forecasts affect experts' behaviour? *European Journal of Operational Research*, 249(3):842–852.
- Russom, P. et al. (2011). Big data analytics. *TDWI Best Practices Report, Fourth Quarter*.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58(1):267–288.
- Trapero, J. R., Kourentzes, N., and Fildes, R. (2012). Impact of information exchange on supplier forecasting performance. *Omega*, 40(6):738–747.
- Trapero, J. R., Pedregal, D. J., Fildes, R., and Kourentzes, N. (2013). Analysis of judgmental adjustments in the presence of promotions. *International Journal of Forecasting*, 29(2):234–243.
- Verleysen, M. and François, D. (2005). The curse of dimensionality in data mining and time series prediction. In *Computational Intelligence and Bioinspired Systems*, pages 758–770. Springer.
- Weller, M. and SF, C. (2012). Supply chain forecasting: Best practices & benchmarking study.