Sensing endogenous seasonality in the case of a coffee supply chain

Abstract

Rogue seasonality, or endogenously generated cyclicality (in variables), is common in supply chains and known to adversely affect performance. This paper explores a technique for sensing rogue seasonality at a supply chain echelon level. A signature and index based on cluster profiles of variables, which are meant to sense echelon-level generation and intensity of rogue seasonality, respectively, are proposed. Their validity is then established on echelons of a downstream coffee supply chain for five stock keeping units (SKUs) with contrasting rogue seasonality generation behaviour. The appropriateness of spectra as the domain for representing variables, data for which is daily sampled, is highlighted. Time-batching cycles which could corrupt the sensing are observed in variables, and the need to therefore filter them out in advance is also highlighted. The knowledge gained about the echelon location, intensity and time of generation of rogue seasonality could enable timely deployment of specific mitigation actions.

Keywords: Supply chain risk, endogenous seasonality, bullwhip, sense and respond, clustering

1 Introduction

Supply chains are prone to disturbances, including those that arise endogenously from the use of inappropriate control systems and/or information to match supply with demand (Chopra and Sodhi, 2004). One such disturbance, characterised by endogenous generation of cyclicality in the profiles of system variables, is endogenous or rogue seasonality (Forrester 1961; McCullen and Towill 2002), a key but under-researched component of the bullwhip effect (Lee, Padmanabhan and Wang 1997; Fransoo and Wouters 2000; Holland and Sodhi 2004; Hussain and Saber 2012). Evident in several studies such as Kaipia, Korhonen, and Hartiala (2006) and Thornhill and Naim (2006), rogue seasonality tends to be misinterpreted as a cyclicality/seasonality in market demand and unnecessary managed (through ramping up and down of production and/or increase in stock levels) thereby increasing operating costs. The (adverse) cost implications are particularly severe for multi-echelon contexts where rogue seasonality could propagate to other echelons and cause a similar impact. Metters (1997) showed that costs to the extent of 10–20% could be reduced by the elimination of such seasonal variations.

One way of managing rogue seasonality, as suggested by Shukla, Naim, and Thornhill (2012) and Shukla and Naim (2015) and discussed in resilience terms by Purvis et al (2016), could be through 'sense and respond' (Haeckel 1999), where 'sense' means assessing a problem context through system information and appropriate models/filters/data-mining techniques, and 'respond', to the subsequent initiation of suitable corrective action/s. Their suggested approach for sensing rogue seasonality involves a signature and an index, where the signature is to indicate if rogue seasonality is present/generated in a supply chain, and the index, to provide a measure of its intensity. Both the signature and the index are defined on the basis of cluster profiles of variables such as order, inventory and work in process that are represented as spectra, or amplitudes of sinusoids at various frequencies, which was identified to be the most appropriate domain for representation. Different supply chain contexts, both simulated and empirical, though only upstream ones were used for their validation. However, these studies give rise to further questions such as:

1) Instead of a supply chain level of sensing, is it possible to sense rogue seasonality at an individual echelon (in a supply chain) level through an appropriately defined signature and index? Also, whether such a sensing could be done over time in a monitoring mode? Rogue seasonality could then be more precisely determined in terms of generation location, intensity and time of generation, thereby enabling more specific and timely mitigative actions to be deployed.

2) Whether the signature and index could work with downstream or distribution/production-distribution supply chain contexts? These differ from the upstream contexts considered in previous studies in the following ways: i) There are (additional) time-batching related cyclicality/ies in order and dispatch variables (Cachon 1999; Potter and Disney 2006) that could corrupt the signature and index's logic; ii) Their operative level is more disaggregated, both product and time-wise; data, such as at an SKU level and daily sampled that is appropriate for such contexts tends to be noisy, and could affect the development of cluster profiles (Keogh and Kasetty (2003), and thereby the logic/validity of the signature and index.

This study is aimed at addressing these questions. The downstream end of a multi-echelon coffee supply chain is used for the investigations, with the dynamics of variables for five SKUs being separately analysed. The multi-echelon structure enables the validity of the echelon-level signature and index (that are defined) to be assessed at different echelons. Daily data is considered for the variables, and which is chosen for a length of time that enables the assessments to be made at multiple/different time intervals. While the primary contribution of this study is towards sensing of rogue seasonality, it also contributes towards empricial rogue seasonality investigations and supply chain dynamics assessments with low-level data, where there is a paucity of studies in the literature.

The rest of the paper is structured as follows. Previous studies on rogue seasonality as well as its sensing are discussed in the next section, while in Section 3, we explore the development of a signature and an index for sensing it at an echelon (in a supply chain) level. Investigation of the coffee supply chain from the perspective of rogue seasonality generation and its sensing is discussed in

Section 4. In Section 5 the practical aspects of applying the signature and index are discussed with Section 6 being the concluding section.

2 Rogue seasonality and its sensing

2.1 About rogue seasonality

It is a seasonality/cyclic pattern (of the order of months/weeks/days) in order and other supply chain variables that does not originate from consumer/exogenous demand but instead (endogenously) from the inventory and production control system used (Forrester 1961; McCullen and Towill 2002); hence the term endogenous or rogue seasonality. Though commonly observed (Kaipia, Korhonen and Hartiala 2006; Torres and Moran 2006; Thornhill and Naim 2006) academic interest on rogue seasonality has been limited. Few studies have exclusively focussed on it with most considering it together with the Bullwhip effect which involves amplification of orders from downstream to upstream echelons (Forrester 1961; Miragliotta 2006; McCullen and Towill 2002). This is true for the works of Dejonckheere et al (2003) and Jaksic and Rusjan (2008) as well; the focus there again is on the Bullwhip effect, but because the investigations are done in the frequency domain, rogue/endogenously generated cycles have indirectly gotten analysed. The analyses though is analytical (rather than empirical), and is for one variable (orders) only rather than multiple ones typically associated with rogue seasonality; the range of cycles considered too are appropriate for Bullwhip rather than rogue seasonality (investigation). Other more rogue seasonality focussed studies incl. those of Kim and Springer (2008, 2010) and Spiegler and Naim (2017). Their focus is on studying rogue seasonality generation characteristics and proposing design formulations such as optimum ordering policies to be used to minimise the strength of the generation. Analytical system dynamics approaches are used, which therefore meant tractable/simplifying assumptions such as dyadic structures and rationality in ordering having to be made. Though useful, policies developed under such simplistic assumptions may have limited relevance for real world contexts characterised by dynamic uncertainties, differing member objectives/constraints and behavioural biases in decision making. Top down design based approaches for rogue seasonality therefore need to be complemented with the bottoms-up sense/detect and respond based approach (Haeckel, 1999).

2.2 Rogue seasonality sensing approaches

Three studies, Thornhill and Naim (2006), Shukla, Naim, and Thornhill (2012) and Shukla and Naim (2015), have explored sensing of rogue seasonality, with 'sense and respond' (Haeckel 1999) as the underlying approach. While Thornhill and Naim's (2006) sensing technique has subjective elements, Shukla, Naim, and Thornhill's (2012) approach, that is based on a signature to assess presence/generation and an index to indicate the intensity of rogue seasonality in a supply chain is more objective, versatile and automation-friendly. The signature is defined as follows: *Rogue seasonality is considered generated/present in a supply chain when its system variables are clustered* *away from exogenous demand, and not otherwise*, where clustering is used to record the similarity/dissimilarity relationships. The index is based on the logic of the signature, but with numerical values used to represent variable profiles and profile relationships. The specific logic is: a greater intensity of rogue seasonality in a supply chain would mean system variables across the supply chain being more tightly aligned to each other, causing their mutual dissimilarities (the denominator in the index) to be lower, while increasing their dissimilarities with exogenous or consumer demand (the numerator in the index), to ultimately cause the index to have a larger value. Two versions of the index, identified as index definition 1 and 2 are proposed, where the difference is in terms of whether the dissimilarity in the numerator is "minimum" or "average" as can be seen below. The rationale here is to assess if the logic of the index is robust to profiles of some system variables being different to the rest.

$$\textit{Rogue Seasonality Index for a supply chain (Definition 1)} = \frac{\textit{Minimum dissimilarity between all system variables and exogenous demand}}{\textit{Average dissimilarity between all system variables excluding exogenous demand}}$$

$$\textit{Rogue Seasonality Index for a supply chain (Definition 2)} = \frac{\textit{Average dissimilarity between all system variables and exogenous demand}}{\textit{Average dissimilarity between all system variables excluding exogenous demand}}$$

On testing, both the index definitions were found to be similarly effective. Spectra or amplitudes of sinusoids at different frequencies (after Fourier transform) was identified to be the most appropriate domain for representing variables when deriving the signature and the index. While Shukla, Naim, and Thornhill (2012) used different simulated contexts and an empirical steel context, Shukla and Naim (2015) established the signature's and index's validity with a more complex supply chain system that incorporated backlogs and quantity batching in ordering and shipping and considered more realistic ordering policies. While the index was seen to be valid as such, the signature required tweaking in the form of a threshold for proportion of variables that could be clustered with exogenous demand (and the context to still be classified as having rogue seasonality present/generated) needing to be specified.

2.3 Gaps in previous sensing approaches

While Shukla, Naim, and Thornhill's (2012) signature and index propositions for sensing rogue seasonality are novel, they are defined for a supply chain as a whole; also, the rogue seasonality is assumed to be present in all/most variables across the supply chain except consumer/exogenous demand (the reference variable). Discrimination and identification of problematic supply chains (from a rogue seasonality perspective) is therefore possible. But what about individual echelons within supply chains? They are where rogue seasonality generation actually takes place (not in supply chains), and where countermeasures need to be applied in practice. By sensing rogue seasonality at an echelon (in a supply chain) level and identifying the problematic one/s, specific (echelon level) corrective actions would be possible to further minimise the economic consequences of rogue seasonality. Such a level of sensing makes sense from a technical perspective as well: in many supply chain contexts, all echelons may not show rogue seasonality that is assumed in the existing (supply chain level) signature and index definitions. This could happen, for example, where a) the customer-facing echelon and other echelons at that end operate on a make-to-order basis that is associated with no rogue seasonality generation (Shukla, Naim, and Thornhill 2012); or where b) rogue seasonality is generated in one of the middle echelons in a supply chain and while it may propagate upstream, echelons downstream do not show it. Echelon level sensing of rogue seasonality is therefore relevant. However, the question that arises then is whether we need a (separate) signature and index for that level; or the existing supply chain level definitions could be applied as such given that an echelon is also technically a supply chain. Prima facie, the latter would work for the most downstream echelon but not for others. For the most downstream echelon, "all system variables" in the definition would mean those for that echelon, while the reference (for sensing rogue seasonality) would be exogenous demand; the supply chain level signature and index (refer to the definitions discussed earlier) can therefore be seen to apply for this echelon. For being able to sense rogue seasonality at other echelons however, the reference would need to be not just exogenous demand but also all echelon variables downstream of the echelon under consideration; only then would it be possible to isolate the rogue seasonality generation characteristics for that echelon. The signature and index definitions for sensing rogue seasonality at an individual echelon level would therefore be different from those for the supply chain as a whole, and which need to be developed and then validated.

Another deficiency of the previous work is that the signature's and index's effectiveness is not established for sensing rogue seasonality at different time intervals (in a monitoring mode); rather, the related evaluations in Shukla, Naim, and Thornhill (2012) and Shukla and Naim (2015) are done at one time with all available data. Success with the former could enable earlier detection of rogue seasonality, and thereby in quicker initiation of corrective response/s to ultimately lower the economic consequences of rogue seasonality.

Finally, the signature and index in Shukla, Naim, and Thornhill (2012) and Shukla and Naim (2015) are based on upstream supply chain contexts only; downstream or distribution/production-distribution contexts, whose dynamic characteristics are more complex are not considered. Even Shukla and Naim's (2015) work, which is based on the Beer game supply chain (Sterman 1989) that is technically a production-distribution or downstream context, does not incorporate these complexities, which are:

a) Time-batching related cyclicality in ordering/despatch. This cyclicality is because of transportation efficiency and ease of management considerations (Cachon 1999; Potter and Disney 2006), and could differ across echelons in a supply chain due to differences in their cost structures and volume of goods handled (Burbidge 1994). For a production-distribution system, this difference could stretch upto the production echelon, with production batching cycles differing from those pertaining to dispatches to downstream echelons (Hejazi and Hilmola 2006; Taylor and Fearne 2009). Presence of these (additional) endogenously generated time-batching cycles could interfere with the signature and index's logic/validity.

b) Disaggregated nature of the operations, both product and time-wise. This is on account of downstream contexts' proximity to end customers (that have specific product demands) and their decision making requirements being more operational in nature respectively. The granular data, such as at an SKU level and which is daily sampled that is relevant to such contexts tends to be noisy and irregular; the cluster profiles for such a data may not conform to the logic of the signature and index (Keogh and Kasetty, 2003; Liao, 2005), and this needs to be assessed. Previous rogue seasonality sensing investigations have considered only monthly and product-wise (on a weight basis) aggregated data. In general also, few studies, including those on the bullwhip effect, have assessed supply chain dynamics with SKU level daily data. This, despite the rationale for doing so being supported by the literature (Fransoo and Wouters 2000; Chen and Lee 2012), and the availability of granular data (through RFID technology) and associated analytical techniques (de Kok et al. 2005; Lee and Ozalp 2007) increasing the practical relevance of such studies.

This study seeks to fill the above discussed gaps through investigations on a downstream coffee supply chain; the dynamics of variables (from a rogue seasonality sensing perspective) for five SKUs are analysed, both across echelons and over time, with daily data being used for the variables. However, echelon-level signature and index definitions need to be developed first.

3 Rogue seasonality signature and index for an echelon (in a supply chain) level

Given that supply chain-level signature and index definitions already exist, we explore if these could be adapted to work at an echelon level, and what, if any, changes are needed. In these (supply chain level) definitions, the supply chain as a whole is considered to be one entity with profiles of all the variables across the supply chain (i.e. covering all constituent echelons) together being therefore compared against market or exogenous demand's, which is considered as the reference variable. Applying this logic to an echelon level (of sensing) would mean that variables whose profiles are considered for comparison as well as the variables to whose profiles they are compared with, or the reference variable/s, would change with the echelon under assessment. For example, for sensing at the most downstream echelon, profiles of variables for that echelon would need to be compared against that for exogenous or market demand. For the next upstream echelon, it would mean comparing the profiles of variables for that echelon with those for all downstream echelon variables as well as exogenous or market demand. Generalising this pattern and using the existing (supply chain level) signature definition as the basis, the signature for an echelon in a supply chain can be defined as follows: *Rogue seasonality is considered generated at an echelon when system variables for that echelon are clustered away from all downstream echelon variables and exogenous demand, and not* *otherwise*. The index to indicate the intensity of rogue seasonality generated at an echelon level can be similarly adapted from its supply-level version as follows:

$$\begin{array}{l}\textit{Rogue Seasonality Index}\\ \textit{for an echelon in a supply chain}\\ \textit{(Definition 1)}\end{array} = \frac{\begin{array}{l}\textit{Minimum dissimilarity between system variables at that echelon and all downstream}\\ \textit{echelon variables incl. exogenous demand}\end{array}}{\textit{Average dissimilarity between all system variables at that echelon}}$$

$$\begin{array}{l}\textit{Rogue Seasonality Index}\\ \textit{for an echelon in a supply chain}\\ \textit{(Definition 2)}\end{array} = \frac{\begin{array}{l}\textit{Average dissimilarity between system variables at that echelon and all downstream}\\ \textit{echelon variables incl. exogenous demand}\end{array}}{\textit{Average dissimilarity between all system variables at that echelon}}$$

These echelon-level index definitions are expected to work in the same way as the supply chain level ones discussed earlier in Section 2.2. Though logical, they need to be validated with a realistic context, as is discussed in the next section.

4 Coffee supply chain case study

The supply chain structure and flows as well as the nature of SKU's, system variables and data considered are explained/justified first. Profiles of the system variables are then used to support the assumptions made about downstream supply chain contexts as well as to investigate rogue seasonality generation characteristics. Finally, the signature and index are derived for individual SKU's at different echelons, and their validity is established. The signature and index are also derived at different time intervals to assess their (rogue seasonality) monitoring effectiveness.

4.1 Structure and flows

A downstream supply chain for coffee is considered for the investigations (see Figure 1 below), whose focal company, referred to as company M, is a part of a global food corporation.

Add Figure 1 here

Based on information gathered through interviews with company M's planning and distribution managers, the material flow across the supply chain is understood to take place as follows: Coffee is produced in batches at company M's factory and then transferred to two distribution centres (DCs), both company managed, with almost no stock held at the factory. Though company M actually has three factories, each product/SKU is produced in one specific factory only, and hence the representation of a single factory in the figure. From the DCs, the products/SKUs are shipped to company M's primary customers, which are large Retailers/Wholesalers, as well as to a Distributor (referred to as WD) that services company M's small customers. Some inter-DC transfer of products/SKUs also takes place due to errors both in forecasting individual DC's requirements and in allocation. Supplies to WD are delivered to its nine depots, each uniquely assigned to the DCs, from where they are further supplied to the respective small customers. These depots are used by WD to service its own customers as well, given the (additional) Wholesaler role that it plays. As regards ordering by the small customers, they place their orders with company M, which then coordinates with WD for their servicing; in the case of a shortfall being anticipated at one of the depots, the required quantity of goods is dispatched there in advance. Inventory management at the DCs is done at an SKU level and is based on a periodic (weekly) review policy; in any week T, the status for week T+4 is assessed, and in the case of a shortfall being anticipated, a production request is initiated. This request is usually for a quantity which is greater of economic production quantity (EPQ) or one week's requirement for that SKU.

4.2 Nature of data considered

4.2.1 SKUs considered

With the focus being on investigations with disaggregated data, SKU-level data is considered. As will be seen later, this is an appropriate (product-wise) disaggregation level for the context under investigation here. A sample of five SKUs with varying sales volumes (in cases) and produced in different factories are considered, which are as follows: SKU1: 1% of total sales volume, produced in Factory 1; SKU2: 1% of total sales volume, produced in Factory 2; SKU3: 3% of total sales volume, produced in Factory 3; SKU4: 9% of total sales volume, produced in Factory 3; and SKU5: 1% of total sales volume, produced in Factory 3. Different sales volumes could mean different regularity in variable profiles that could make for a more robust assessment of the signature and index, as will be seen later.

4.2.2 Variables considered

Sensing, including of rogue seasonality, requires time series data on system variables. For the coffee supply chain, these variables (in units of cases), are '*A*', '*B*', '*C*', '*D*', '*E*', '*F*', '*G1 to G5*', '*H1 to H4*', '*I1 to I9*' and '*J1 to J9*', with their descriptions given in figure 1. No data was available on WD's sales to its own customers at the depots, and in order to ensure that this does not influence the findings significantly, only SKUs where these sales were estimated to be small were chosen. For the chosen SKUs 1, 2, 5, 4 and 3, only 2%, 9%, 15%, 35% and 37%, respectively, of the products going into WD depots make their way to WD's own customers. Transaction times between supply chain entities, as per inputs of company M's managers, being of the order of days, daily data on variables was considered. As will be seen later, this is an appropriate (time-wise) disaggregation level for the context being investigated here. Another issue was the length of data to be collected, which needed to be such that signature and index assessments could be done at multiple/different time intervals, but structural shifts in data, if any, are avoided. One year's worth of data was considered appropriate, and therefore daily data for one year for the above variables for each of the five SKUs was used for the rogue seasonality sensing investigations. This data was extracted from company M's and WD's IT systems.

4.2.3 Profiles of the variables considered

One of the justifications of this study is that downstream supply chain contexts show certain distinguishing characteristics which could affect the validity of the signature and the index for such contexts. Ascertaining if these characteristics, such as time-batching, operative dynamics being at a disaggregated level which necessitates use of such a level of data, and this data being noisy and irregular, are actually present in the downstream (coffee supply chain) context being investigated here is therefore important. To do so, time series and spectra profiles (refer to Appendix I for its description as well as for other terms covered later) of the variables are plotted; Excel© and Matlab© are used for the computations and plots. Spectra or amplitude of sinusoids at different frequency channels (derived through Fourier transformation) is considered, as it is known to be effective for data with cyclical characteristics, and which was also observed in Shukla, Naim, and Thornhill (2012) and in Shukla and Naim (2015). These profiles for one of the SKUs, SKU1, are presented in Figure 2 below.

Add Figure 2 here

The profiles are based on de-trended and normalised data. This enables time series profiles to be compared at the same scale, while unnecessary features near the zero frequency channel that could complicate assessments are avoided for spectra profiles. For the spectra plot in Figure 2, the horizontal axis is a normalised frequency axis such that 1.0 represents the sampling frequency. This means that for the daily data that is used here, a spectral peak on the frequency axis, at say 0.2/day, would correspond to a cycle in the time series of a period of 1/0.2 or 5 days. The spectra plot is seen to stop at 0.5/day on the frequency axis because of the Nyquist sampling theorem, which requires a sinusoidal signal to be sampled at least twice per cycle (Chatfield 2004). The spectra are scaled to the same maximum peak height to enable better visualisation of their characteristics. The following can be observed from the plots in Figure 2:

- The time series plots clearly highlight the irregular and discontinuous nature of the data. However, apart from this, no other insights including on time-batching are gained from these plots.
- The spectra profiles, on the other hand (and as expected) are seen to be more effective: peaks, representative of a structured repeating pattern, are evident in the spectra for many variables. For example, for variables '*A*' and '*B*', spectral peaks are seen at a frequency of 0.15/day (~ 7 day cycle), while for others, such as '*E*', '*F*', '*G1 to G5*', '*H1 to H4*' and '*J1 to J9*', these are seen at some/all of the following frequencies: 0.15/day (~ 7 day cycle), 0.3/day (~ 3 day cycle) and 0.43/day (~ 2 day cycle). These peaks are all essentially reflecting the time-batching cycles present. These cycles also

appear to be different for different echelons, as per variables for those echelons. For example, while time-batching cycles of frequencies 0.15/day (~7 day cycle), 0.3/day (~3 day cycle) and 0.43/day (~2 day cycle) are seen in variables '*J1 to J9*' that correspond to the WD depot echelon, those of only 0.15/day (~7 day cycle) are seen in variables '*A*' and '*B*', corresponding to the factory echelon. These time-batching cycles, though of endogenous origin, appear to have a rational basis: for example, those in variables '*J1 to J9*' appear to be on account of customers with different purchase volumes and who are locationally dispersed being serviced at different frequencies to ensure cost effectiveness, while that in variables '*A*' and '*B*' seems to be due to the weekly nature of the planning. The (order of days) cyclicality seen in variables also supports the choice of daily data for the investigations, which would have been masked if weekly/monthly data was used. Finally, features at low frequencies of between 0.03/day and 0.005/day (~ 30 day to 200 day cycles) are seen in the spectra for variables '*A*', 'B', '*C*', '*D*' and the majority of the '*G*', '*H*' and '*I*' variables, which could be due to the end of month/Easter/Christmas and other such factors.

While the above discussion concerned SKU1, time series and spectra profiles of a similar nature were observed for the other SKUs as well, though these are not included here due to a paucity of space. The downstream coffee supply chain context here therefore shows the characteristics expected of a generic downstream context except for (product-wise) disaggregated data, such as the SKU level being an appropriate level; this will be done in the next section. We now focus on utilising the variable profiles to investigate echelon-level sensing of rogue seasonality. The preference will be for spectra profiles that were seen to be effective in uncovering cyclical features in the data. Though this effectiveness was demonstrated by Shukla, Naim, and Thornhill (2012) and Shukla and Naim (2015), these researchers utilised continuous data as opposed to the irregular and discontinuous kind seen here.

4.3 Rogue seasonality investigation

4.3.1 Aggregation of entities and variables

The downstream coffee supply chain under investigation has both a vertical and a horizontal structure (refer to Figure 1): while Factory, Distribution Centre (DC), Large Retailers/Wholesalers/WD and end customers are a part of the former, parallel entities at these stages, such as two DCs, nine WD depots and two sets of end customers, i.e. WD's own customers and company M's customers, constitute the latter. As various studies have shown rogue seasonality to propagate upstream, our focus here is essentially on the vertical structure; the horizontal structure could therefore be compressed to create a simplified and potentially more insight-yielding structure. This would mean combining parallel entities at the relevant stages as long as those entities, as assessed by comparing the profiles of their variables, are similar. This is to ensure integrity in the aggregation.

Combining the parallel entities and associated variables (data) results in the creation of the following aggregated variables: total despatches from factory to DCs represented by variable '*A plus B*'; total stock at DCs represented by variable '*C plus D*'; total despatches from DCs to large Retailers/Wholesalers represented by variable '*E plus F*'; total despatches from DCs to WD depots represented by variable '*G plus H total*'; total stock at WD depots represented by variable '*I total*' and total customer M orders serviced at WD depots represented by variable '*J total*'. Though information on WD's sales to its own customers at the depots was unavailable, and which therefore could not be added to '*J total*', its influence on the dynamics is nonetheless curtailed through choice of SKUs, as discussed earlier. Also, '*E plus F*' and '*G plus H total*' are kept separate rather than being combined together, as their magnitudes were noted to be significantly different, which was expected to cause their dynamic characteristics to differ, as will be corroborated later. The de-trended and normalised time series and spectra profiles of the aggregated variables for each of the five SKUs are given in figure 3 below.

Add Figure 3 here

It is important to assess if the aggregation done was valid, and therefore spectra profiles of both the variables that are aggregated and of the resulting aggregated variable were all compared to each other. The three were found to be similar in all cases. For example, in the case of SKU1, profiles for variables '*A'* and '*B'* (refer to Figure 2, right-hand panel) and of the aggregated variable '*A plus B'* (refer to Figure 3, bottom half, first panel) are all similar with a spectral peak at frequency 0.15/day (~ 7 day cycle). Similarly, spectra profiles for each of the '*J1*' *to* '*J9*' variables and the aggregated variable '*J total*' are also similar with spectral peaks at frequencies of 0.15/day (~7 day cycle), 0.3/day (~3 day cycle) and 0.43/day (~ 2 day cycle). The aggregated variables therefore represent the dynamics accurately and will be used for the subsequent investigations. One final point about the profiles for variables '*E plus F*' and '*G plus H total*' is that these are seen to be different for most SKUs (refer to Figure 3, bottom half), thereby supporting the original decision of not combining them together.

4.3.2 Rogue seasonality presentation

We now try to understand the nature of rogue seasonality present/generated, and therefore, profiles of the aggregated variables in Figure 3 are examined.

The time series profiles (refer to Figure 3, top half), as with unaggregated variables earlier, do not appear to be useful with no cyclic features identifiable. What is highlighted is that the data is irregular, that this irregularity is significantly more in order and dispatch variables than in stocks variables and that it varies across SKUs; this last point increases our confidence that the sensing assessments are robust. No other patterns are evident except in the case of SKU 5, where periods when no supplies were made from factory to DC'S (refer to variable "A plus B"), which caused a reduction in total stocks at DC's (refer to variable "C plus D"), that in turn led to a reduction in despatches from DC's to Large Retailers/Wholesalers (refer to variable "E plus F") and to WD depots (refer to variable "G plus H total"), the last potentially affecting WD's ability to service M's customers from its depots (refer to variable "J total") are seen.

The spectra profiles, on the other hand (refer to Figure 3, bottom half), as with unaggregated variables earlier, are seen to be more effective, with the following being evident: 1) time-batching cycles (reflected in the form of spectral peaks), which are different in some/all variables for an SKU; and 2) elevated spectra values at low frequencies of between 0.03/day and 0.005/day (~ 30 day to 200 day cycles) for most variables. However, variables are also seen to show different time-batching cycles in different SKUs. For example, variable '*J total*' shows cycles of 2, 3 and 7 days for SKUs 1, 3 and 4, but only shows cycles of 7 and 3 and 7 days in the case of SKUs 2 and 5, respectively. While this can be explained as being due to the nature of factors that cause time-batching and their inputs being different, it also demonstrates the unique nature of each SKU's dynamics. The choice of such a (product-wise) disaggregated level, i.e. the SKU level, for the investigation is therefore justified.

We now use the spectra profiles further to assess rogue seasonality or endogenously generated cyclicality in variables. It is useful to clarify that unlike in previous rogue seasonality studies, where rogue seasonality was the only endogenously generated cyclicality present, here, another cyclicality of the same kind (of time-batching), and which is perfectly legitimate, is also present. In such cases, the different endogenously generated cyclicalities present require further discrimination, with only the illegitimate one, appearing as an unexplained cyclicality, being classified as rogue seasonality. The spectra profiles are examined again closely, but now from this perspective. No unexplained spectral peaks are apparent in any variable in the case of SKUs 1 to 4, with only those pertaining to time-batching and at low frequencies discussed earlier being seen. This is true for variables in the case of SKU 5 as well except for variable '*A plus B*', where, in addition to the spectral peak at low frequencies of between 0.03/day and 0.005/day (~ 30 day and 200-day cycles), and another time-batching-related one at frequency 0.15/day (~14 day cycle), a (smaller) unexplained peak at frequency 0.07/day (~14 day cycle) is also evident. Interviews were conducted with company M's personnel to clarify this unexplained cycle, as per whom manufacturing-related problems had caused production stoppages at certain times during the year (corresponding to the data). This is also supported by the time series profile for the '*A plus B*' variable for SKU 5, where several nil values, suggestive of no despatches being made from the factory, are evident. The non-production in certain periods appears to have caused a change in the factory production schedule and consequently a change in the schedule of despatches made from there to the DCs, or in the '*A plus B*' variable, to create the 14 day cycle therein that can be classified as rogue seasonality. Manufacturing problems, such as equipment breakdowns and other supply disruptions, are known to affect the dynamics of variables (Taylor 1999; Paik and Bagchi 2007), but their causing rogue seasonality generation as seen here has not been discussed before. The rogue seasonality here is seen in only a single echelon, the factory echelon (as per the variable for this echelon, '*A plus B*'), rather than echelons across the supply chain. An echelon, rather than a supply chain level of sensing, is therefore more relevant here, as was discussed earlier in Section 2.2. This factory-generated rogue seasonality could be transmitted to echelons upstream and cause economic consequences there. Timely sensing and subsequent speedy management of this rogue seasonality is therefore essential.

4.3.3 Rogue seasonality sensing at an echelon level

After understanding the nature of rogue seasonality generated/present, we now explore its sensing at an echelon level. This would mean trying out the echelon-level signature and index propositions discussed in Section 3 and assessing their effectiveness. The assessments can be made at each of the two echelons in the (aggregated) coffee supply chain of factory and WD depots; DCs are not considered as a separate echelon, as their operations are jointly managed with the factory's.

4.3.3.1 Sensing based on the signature

Referring to the echelon-level signature definition (refer to Section 3), and applying it at the factory and WD depots echelons (for the five SKUs) would mean the following: for the factory echelon, the nature of the clustering of '*A plus B*' (the only variable for this echelon) vis-a-vis the rest of the variables (the downstream variables in this case; refer to Figure 1) will need to be compared; for the WD depots echelon, this comparison will need to be between the two variables for this echelon of '*G plus H total*' and '*I Total*', vis-a-vis the lone downstream variable, '*J total*'. Exogenous demand in the definition is not relevant here as respective proxies, '*E plus F*' and '*J total*', are used. For clustering, representing variables as spectra, as in Figure 3, would not be enough; these spectra would need to be pre-processed as well.

Features at unrelated frequencies in the spectra could interfere with the logic and computation of the signature and index. The aim of pre-processing here is therefore to filter these frequencies out in advance. Some frequencies of this kind with elevated spectra values are seen in the low frequency range (refer to Figure 3, bottom half), similar to that in Thornhill and Naim (2006) and Shukla, Naim, and Thornhill (2012); if not filtered out, they would tend to mask the rogue seasonality present. Given that transaction times between supply chain entities is of the order of days and a similarly sampled data is used, frequencies lower than 0.03/day (or cycles of time periods greater than a month) could be considered as low and filtered out from the spectra of all the variables. Another set of frequencies that need to be filtered out here in order to effectively uncover rogue seasonality are those pertaining to time-batching. For a particular echelon under evaluation, time-batching related frequencies present in each of the relevant variables would need to be filtered out from all of those variables. For example, in the case of SKU 2, for the factory echelon where all six variables are relevant, the time-batching related frequencies seen in these variables of 0.15/day (~7 day cycle), 0.3/day (~3 day cycle) and 0.43/day (~ 2 day cycle) (refer to Figure 3, bottom half, second panel) would need to be filtered out from all six variables. On the other hand, for the evaluation of the WD depots echelon for the same SKU, the time-batching related frequency of 0.15/day (~7 day cycle) seen in the relevant variables of 'G plus H total', 'I Total' and 'J total' would need to be filtered out from all (these) three variables. From the spectra plots in Figure 3, it is clear that for all the SKUs except SKU2, the filtered out time-batching related frequencies for factory and WD depots echelon evaluations would be the same. It is important to point out that while identification of the time-batching related frequencies, including those which need to be filtered out, is done visually here, for a new/unknown context, this would have to be done on the basis of historical data/patterns. Spectra of variables are filtered for the factory and WD depots echelon evaluations as above for each SKU. This is done by replacing the spectra values at the relevant frequencies with zeros. Variables are then clustered using Ward's algorithm (and XLTAT© software), with Euclidean distance as the dissimilarity measure. These cluster profiles are presented in Figure 4 below. To conserve space, cluster profiles pertaining to the evaluation of the factory echelon are only presented, with those pertaining to WD depots echelon being interpreted from the same.

Add Figure 4 here

Evaluating the signature at the factory echelon, with unrelated/interfering frequencies being filtered away from all the relevant variables during pre-processing, rogue seasonality, if generated at this echelon, would cause '*A plus B*' to be dissimilar to others. Therefore, it would also be clustered separately. Such a cluster profile can be expected for the factory echelon in the case of SKU 5, where rogue seasonality generation was evident (refer to Section 4.3.2). However, for the other four SKUs with no rogue seasonality generation at the factory echelon, there is no reason for the '*A plus B*' variable to be clustered separately from the other variables. Examination of the cluster profiles in Figure 4 shows this to be true in both cases: while '*A plus B*' is clustered separately from other variables in the case of SKU 5, it is clustered with them for SKUs 1 to 4. The signature is therefore seen to work in detecting generation/non-generation of rogue seasonality at the factory echelon.

We now assess the effectiveness of the echelon-level signature for the WD depots echelon. As per the discussion in Section 4.3.2, no rogue seasonality generation was apparent in either of the two variables, '*G plus H total*' and '*I Total*', for this echelon for any SKU. Therefore, if the logic of the signature were to hold and non-generation of rogue seasonality for this echelon correctly indicated, variables '*G plus H total*' and '*I Total*' and the lone downstream variable '*J total*' would need to be not clustered separately. On examining the cluster profiles, this appears to be only partially true: while '*G plus H total*' and '*J total*' are seen to be not separately clustered, this is not the case for '*I Total*' and '*J total*'. As discussed in Shukla and Naim (2015), specification of a threshold for proportion of variables that do not cluster as expected but for which the signature is still valid therefore becomes relevant.

Given the visual nature of the signature, which makes it subjective and operationally difficult to apply in contexts involving a large number of supply chains/SKUs, a quantified version of the signature in the form of an index of rogue seasonality becomes relevant.

4.3.3.2 Sensing based on the index

The echelon-level index definitions (definitions 1 and 2) and logic are discussed earlier in Section 3. Though index values as per definitions 1 and 2 were seen to be strongly correlated for the supply chain version of the index, this may not be true for the echelon-level index and needs to be assessed.

We now seek to derive index values for the factory and WD depots echelons for each SKU. For the WD depots echelon, no problems are anticipated: dissimilarity among the echelon's variables, i.e. '*G plus H total*' and '*I Total*', can be calculated, as can the dissimilarity between them and the downstream variable, '*J total*'. However, for the factory echelon with only one variable, '*A plus B*', average dissimilarity among echelon variables, or the denominator in the index, cannot be calculated. One option could be to use average dissimilarity among all (its) downstream echelon variables (which means the rest of the variables in the present context) instead. However, this could result in underestimation of the index value, especially for the case where that echelon, i.e. the factory echelon, is generating rogue seasonality. This is because the average dissimilarity between the echelon's, i.e. the factory echelon's, variables in such a situation would actually be lower (on account of the shared rogue cyclicality) than the one between (its) downstream echelon variables that is assumed. However, in the case of no rogue seasonality generation at that echelon, i.e. the factory echelon, echelon variables and downstream echelon variables profiles may not significantly differ, and therefore use of the latter may not bias the index values significantly. Unlike the denominator, computing the numerator in this case, which involves dissimilarities between '*A plus B*' and the rest of the variables, is relatively easy. Index values were computed as per above for the factory and WD depots echelons for each SKU, and based on both index definitions. These are presented in Table 1 below. Pre-processed spectra were used for the computations, as in the case of the signature, with dissimilarities between them measured in terms of the Euclidean distance.

Add Table 1 here

On examining the index values in Table 1 and correlating them with generation/non-generation of rogue seasonality in different contexts, we find a good correspondence. Rogue seasonality, as previously seen in Section 4.3.2, is only generated at the factory echelon in case of SKU 5, with all other contexts, i.e. factory and WD depots echelons in the case of SKUs 1 to 4 and WD depots echelon in the case of SKU 5, showing no rogue seasonality generation. The index values in Table 1 accurately reflect the same: index value for the factory echelon for SKU5 (shaded) is greater than that for each of the other contexts, and for both index definitions. However, to incorporate rigour into the assessment, the index value for the former is compared with the pooled/average index value for the latter, which includes all the contexts with no rogue seasonality generation; the difference between the two of 2 and 3.5 times standard deviation of index values for index definitions 1 and 2, respectively, is significant. This (significant) difference is despite the index value for the former, i.e. for the context with rogue seasonality generation, being an underestimation due to the nature of assumptions made in computation, as discussed earlier. The index (both definitions) therefore appears to be effective in discriminating echelons with different intensities of rogue seasonality generation, which for the present case meant generation/non-generation of rogue seasonality. Correlation between index values based on the two index definitions was observed to be high at 0.83, suggesting both to be similarly representative of the index's logic. Either could therefore be used, though thresholds for classifying whether an echelon is generating/not generating rogue seasonality would be different, and which would need to be specified, for each case.

4.3.4 Rogue seasonality sensing over time

After establishing the feasibility of sensing rogue seasonality at an echelon level, we now explore if this sensing could be done across time. This would require investigations of the previous section that were done in a one-time mode, i.e. at the end of a year using the previous 365 days' data, to be repeated at different (shorter) time intervals. With shorter time intervals (or shorter time series), frequency channels in the spectra representation would be fewer, and resolution of the spectral peaks would therefore be poorer; the signature's and the index's effectiveness for sensing rogue seasonality may, as a result, be compromised. This is even more relevant in the present context, where there are a significant number of zero entries in variable data (refer to Figure 3, top half) due to the investigations being at a disaggregated level. However, a shorter time series has its benefits, such as rogue seasonality generation being sensed earlier. We consider an assessment done every 3 months, but based on the previous 6 months' data, to be an appropriate compromise: while the former ensures sensing of rogue seasonality to be reasonably timely, the latter ensures that there are enough data points for the signature and the index to be effective in sensing. The echelon-level signature's and index's effectiveness were therefore assessed as follows: at the end of month 6 using the previous 6 months' data, at the end of month 9 using data from months 4 to 9 and at the end of month 12 using data from months 7 to 12. All the steps followed earlier, i.e. representing variables as spectra, pre-processing them, clustering them for the signature and computing index values for each of the factory and WD depots echelons for each SKU, are now repeated for each of the above time intervals.

We first examine the spectra profiles of variables at different time intervals. The spectra profiles for SKU5, where rogue seasonality generation is evident (at the factory echelon), and for SKU2, where this is not so, are plotted in Figure 5 below. Rogue seasonality index values for the factory echelon at different time intervals for each case are also presented. These were computed using filtered spectra as earlier.

Add Figure 5 here

Examining the profile of variables for SKU5 (refer to Figure 5, top half panels), no significant difference is seen in the profile of a variable across the three time intervals except in the case of variable '*A plus B*'. For '*A plus B*', while the 14 day rogue seasonality seen earlier in its full-year data-based profile is evident in its profiles based on data from months 1 to 6 and months 4 to 9 (refer to the first two panels), this is not so for its profile based on months 7 to 12 (refer to the last panel). From this, it can easily be deduced that the rogue seasonality generation (at the factory echelon) in this case was happening between months 4 to 6, and perhaps between months 1 and 3 as well, although the latter cannot be confirmed. A shorter time interval of sensing, i.e. quarterly instead of yearly, has therefore enabled the rogue seasonality generation determination to be earlier, i.e. the end of month 3 or 6 rather than end of the year. The echelon-level signature based on cluster profiles (not presented here to conserve space) was also seen to accurately indicate generation of rogue seasonality for the first two time intervals and non-generation for the last. The index's effectiveness is also evident, with high index values being seen for the first two time intervals and a low value for the last. In contrast to SKU5, for SKU2, no rogue seasonality generation (at the factory echelon) is apparent for any of the three time intervals (refer to Figure 5, bottom half panels), as was the case with its full-year data. This is also accurately reflected in the index values; index values for all three time intervals for this SKU are much lower than that for the time interval corresponding to rogue seasonality generation in the case of SKU5.

The above discussion pertained to only the factory echelon and for only two SKUs. We now seek to include the rest of the echelon contexts, i.e. WD depots echelon for all five SKUs and factory echelon for SKUs 1, 3 and 4, in the assessment. None of these contexts showed rogue seasonality generation as per the year-long data-based assessment discussed earlier. Index values for these as well as the earlier discussed contexts are provided in Table 2 below. These are computed from filtered spectra of relevant variables as was discussed in Section 4.3.3.2 earlier.

Add Table 2 here

Examination of Table 2 reveals the following:

1) Index values for all three time intervals for the echelon contexts where no rogue seasonality generation was apparent earlier are seen to be low (refer to the unshaded values in Table 2); the average index value for these contexts across time intervals is 0.86 and 0.95 (as per index definition 1 and 2, respectively), with 0.047 being the standard deviation of index value in both cases.

2) The index value for each context and time interval associated with rogue seasonality generation (shaded values in Table 2) is seen to be greater than where that is not the case (unshaded values in Table 2), and which is true for both index definitions. Moreover, the difference between the two is significant: the average index value for the former at 1.03 and 1.17 (as per index definitions 1 and 2) is greater than those for the latter by 3.6 and 4.5 standard deviation of the index value, respectively. The echelon-level index therefore appears to be effective in discriminating rogue seasonality generation intensity across time.

3) The correlation between index values derived on the basis of the two index definitions is seen to be around 0.8. While marginally lower than that seen for the full-year data, it still enables either of the two index definitions to be used.

5 Managerial Implications

The echelon-level signature's and index's effectiveness was established (as above) with an empirical context; their application in practice therefore looks promising. Their use would be particularly relevant to multiple supply chain analysis contexts, including those in the retail sector, where a quick, consistent and efficient identification of problematic echelon/s would be useful. This would require the index to be computed for all echelons of each investigated supply chain at specified time intervals, with an echelon at a certain time interval being identified as problematic, in the case that the index value for it exceeded a specified threshold. Figure 6 summarises the process of practically applying the index for a single supply chain.

Add Figure 6 here

Some of the key application related considerations are:

Supply chains and products/SKU's (within each supply chain) to be considered: While it would be beneficial to include all supply chains and all products/SKU's within each, a selective focus may be appropriate where there are resource constraints; only critical supply chains and key products/SKU's (within each) from a cost and customer service perspective could be considered in such cases. It is important to recognise that given the many echelons in a supply chain, the volume of sensing assessments associated with echelon level sensing would be significantly more than that with the supply chain level ones; the greater complexity would mean more technology and related resource requirements, and therefore, resource availability/constraints would be an important consideration during implementation. Blackhurst et al. (2011) provide a good basis for deciding the criticality of supply chains and products from a disturbance management perspective.

Time series data on system variables: Echelon-wise data on system variables such as orders, receipts, despatches and inventory for the concerned supply chains is needed for computing the signature and index. Given the increasing emphasis on information exchange and collaboration, further boosted by the discourse on big data and its benefits, as well as the availability of enabling technologies such as RFID (Herrmann et al. 2015), this is not expected to be a significant hindrance. Another requirement is for the data to be available at an appropriately sampled frequency, as dictated by the Nyquist sampling theorem. While daily data, if available, would be appropriate for most cases as it would be able to handle the cyclicality of the order of 2 days or more, in the case that only slower sampled data such as weekly data is available, its appropriateness vis-a-vis the context would need to be assessed.

Index computation related: There are several aspects to this:

Filtering out unrelated/irrelevant frequencies from the spectra of variables before index computation: This is important so that these (unrelated/irrelevant) frequencies, typically encountered in the low frequency range on account of factors such as exogenous seasonality/long term trends, do not interfere with the index's logic/validity. The cut-off for low frequency (so as to filter out all frequencies below it from the spectra of all the variables) could be based on the sampling frequency: for e.g., for daily data, frequencies lower than 0.03/day (or cycles of time periods greater than a month) could be considered low. Further refinement could be done on the basis of previous and expert knowledge about operating dynamics. Another set of (unrelated/irrelevant) frequencies needing to be filtered out are those pertaining to time-batching; these could be determined on the basis of historical data/patterns. Such frequencies present in the spectra of individual variables would need to be filtered out from all the variables before index computation. Here all variables includes those which are relevant to index computation for a particular supply chain echelon (as per the index definition).

Index as per which definition, definition 1 or 2

Given that the index as per both definitions was found to be similarly effective, either could be used. However, it is important that whichever is chosen is used consistently across SKU's and over time. Also, the fact that the index can do the job of the signature also, i.e. discriminate presence and absence of rogue seasonality, and do so more objectively, means that it alone is sufficient for sensing. The signature can however be used as a complement in select cases, especially where, detailed insights such as on the nature of the endogenous cycle generated and its causal origins are needed.

Threshold index value

It is important to specify a threshold index value to classify whether an echelon is generating or not generating rogue seasonality and/or for triggering of mitigation options. This can be a fixed value decided on the basis of historical index values for known rogue and non-rogue seasonality generation cases as was done in the case of the coffee supply chain. Alternatively, it can be dynamic such as three standard deviations above the average index value for all echelons and SKU's during each evaluation.

Frequency/time interval used

This could be determined based on the on the basis of the trade-off between the increased effort/cost associated with more frequent sensing evaluations and the benefits from earlier detection. Another factor that needs to be considered here is sensing accuracy; more frequent (or shorter time intervals) sensing could mean lower accuracy/more false alarms in sensing. Here the extent of previous data or the length of the rolling window is also important as seen earlier in Section 4.3.4; the goal here should be to keep as short a length as gives requisite sensing accuracy.

Mitigation options after sensing: The usefulness of echelon-level sensing of rogue seasonality is dependent on how quickly and specifically countermeasures to reduce its intensity or eliminate it altogether are deployed. Actions such as those proposed by Kim and Springer (2008), including reducing the supply lead time or adjusting the ordering/replenishment parameters, would need to be implemented. For the latter, it is important that the adjusted parameters do not have a negative effect on other parts of the supply chain. In the case of supply disruption being the cause of generation, goods/material flow would need to be enabled and steps taken to avoid its recurrence in the future. It is important to note here that it is the echelon (rather than a supply chain) which is more practically relevant both from a rogue seasonality generation as well as its mitigation perspective.

6 Conclusions

In this paper we proposed a signature and index for sensing rogue seasonality generation and its intensity, respectively, at a supply chain echelon level. While the signature is based on cluster profiles of variables, the index involves quantification of the same into a numerical value. Both the signature and the index were found to be effective, though for the former, an additional threshold for proportion of variables clustered as per Shukla and Naim (2015) may need to be specified. Their effectiveness was seen over time as well; in different time interval assessments, both could discriminate intervals that involved (and did not involve) rogue seasonality generation at an echelon. As a result, rogue seasonality can be sensed more precisely in terms of the echelon where it is generated, its intensity and its time of generation (through choice of an appropriately short time interval for assessments). This could enable specific and timely mitigative actions to be applied, and the economic consequences of rogue seasonality, minimised. The signature and index are versatile vis-a-vis input data and were seen to be effective with SKU level and daily sampled data. Spectra was found to be an appropriate domain for representing variables when deriving the signature and index. Finally, the presence of additional endogenously generated cyclicality in variables such as time-batching (besides rogue seasonality), and the need to identify and remove them in advance of the signature and index assessment was highlighted.

While this study has advanced the case for sensing rogue seasonality and for the sense and respond approach in general, some gaps have also emerged. One of them pertains to rogue seasonality generation from supply disruptions, which was observed in this study, but still requires a more detailed understanding through, e.g. the use of simulation approaches. Similarly, horizontal propagation of rogue seasonality, i.e. between entities at the same echelon, can be explored, since only vertical propagation is examined in most studies (including this one). Another stream of research could focus on time-batching and its dynamic implications, which is an infrequently discussed topic in the literature. Future research would also benefit from testing the echelon-level signature and index with other empirical contexts that have more echelons, and where data on multiple variables per echelon is available.

References

Blackhurst, J., K. S. Dunn, and C. W. Craighead. 2011. "An Empirically Derived Framework of Global Supply Resiliency." *Journal of Business Logistics*, 32(4): 374-391.

Burbidge, J. L. 1994. "The use of period batch control (PBC) in the implosive industries." *Production Planning & Control: The Management of Operations* 5(1): 97-102.

Cachon, G.P. 1999. "Managing Supply Chain Demand Variability with Scheduled Ordering Policies." *Management Science* 45(6): 843-856.

Chatfield, C. 2004. *The Analysis of Time Series: An Introduction.* CRC Press.

Chen, L., and H. L. Lee. 2012. "Bullwhip Effect Measurement and Its Implications." *Operations Research* 60(4): 771-784.

Chopra, S., and M. S. Sodhi. 2004. "Managing Risk To Avoid Supply-Chain Breakdown." *MIT Sloan Management Review* 46(1): 53-61.

Dejonckheere, J., S. Disney, M. Lambrecht, and D. R. Towill. 2003. "Measuring and avoiding the bullwhip effect: A control theoretic approach." *European Journal of Operational Research* 147(3): 567-590.

de Kok, T., F. Janssen, J. V. Doremalen, E. V. Wachem, M. Clerkx, and W. Peeters. 2005. "Philips Electronics Synchronizes Its Supply Chain to End the Bullwhip Effect." *Interfaces* 35(1): 37-48.

Forrester, J. W. 1961. *Industrial dynamics.* M.I.T. Press.

Fransoo, J. C., and M. J. F. Wouters. 2000. "Measuring the bullwhip effect in the supply chain." *Supply Chain Management: An International Journal* 5(2): 78-89.

Haeckel, S. H. 1999. *Adaptive Enterprise: Creating and Leading Sense-And-Respond Organizations.* HBS Press.

Hejazi, A., and O. P. Hilmola. 2006. "Manufacturing lot sizing as a source of the Bullwhip effect: a case study of electronics and furniture supply chains." *International Journal of Services and Operations Management* 2(3): 237-255.

Herrmann, S., H. Rogers, M. Gebhard, and E. Hartmann. 2015. "Impact of information integration on decision-making in a supply chain network." *Production Planning & Control* 26(12): 983-993.

Holland, W., and M. S. Sodhi. 2004. “Quantifying the Effect of Batch Size and Order Errors on the Bullwhip Effect Using Simulation.” *International Journal of Logistics Research and Applications* 7(3): 251-261.
Hussain, M., and H. Saber. 2012. “Exploring the bullwhip effect using simulation and Taguchi experimental design.” *International Journal of Logistics Research and Applications* 15(4): 231-249.
Jaksic, M., and B. Rusjan. 2008. "The effect of replenishment policies on the bullwhip effect: A transfer function approach." *European Journal Of Operational Research* 184(3): 946-961.
Kaipia, R., H. Korhonen, and H. Hartiala. 2006. “Planning nervousness in a demand supply network: an empirical study.” *International Journal Of Logistics Management* 17(1): 95-113.
Keogh, E., and S. Kasetty. 2003. “On the need for time series data mining benchmarks: A survey and empirical demonstration.” *Data Mining and Knowledge Discovery* 7 (4): 349-371.
Kim, I., and M. Springer. 2008. “Measuring endogenous supply chain volatility: beyond the bullwhip effect.” *European Journal of Operational Research* 189(1): 172-193.
Springer, M., and I. Kim. 2010. “Managing the order pipeline to reduce supply chain volatility.” *European Journal of Operational Research* 203(2): 380-392.
Lee, H. L., V. Padmanabhan, and S. Wang. 1997. “The Bullwhip effect in supply chains.” *Sloan Management Review* 38(3): 93-102.
Lee, H., and O. Ozalp. 2007. “Unlocking the Value of RFID.” *Production and Operations Management* 16(1): 40-64.
Liao, T. W. 2005. “Clustering of time series data-a survey.” *Pattern Recognition* 38(11): 1857-1874.
McCullen, P., and D. R. Towill. 2002. "Diagnosis and reduction of bullwip in supply chains." *Supply Chain Management: An International Journal* 7(3): 164-180.
Metters, R. 1997. “Quantifying the bullwhip effect in supply chains.” *Journal of Operations Management* 15: 89-100.
Miragliotta, G. 2006. "Layers and mechanisms: A new taxonomy for the Bullwhip Effect." *International Journal of Production Economics* 104(2): 365-381.
Paik, S. K., and P. K. Bagchi. 2007. "Understanding the causes of the bullwhip effect in a supply chain." *International Journal of Retail & Distribution Management* 35(4): 308-324.
Potter, A., and S. M. Disney. 2006. “Bullwhip and batching: An exploration.” *International Journal of Production Economics* 104 (2): 408-418.
Purvis, L., S. Spall, M.M.Naim, and V. Spiegler. 2016. “Developing a resilient supply chain strategy during ‘boom’ and ‘bust’.” *Production Planning and Control* 27 (7-8): 579-590.
Shukla, V., M. M. Naim, and N. F. Thornhill. 2012. “Rogue seasonality detection in supply chains.” *International Journal of Production Economics* 138 (2): 254-272.
Shukla, V., and M. M. Naim. 2015. "Rogue seasonality in supply chains: an investigation and a measurement approach." *Journal of Manufacturing Technology Management* 26(3): 364 – 389.
Spiegler, V.L.M., and M.M. Naim. 2017. “Investigating sustained oscillations in nonlinear production and inventory control models.” *European Journal of Operational Research* 261(2): 572-583.
Sterman, J. 1989. “Modelling managerial behaviour: Misperceptions of feedback in a dynamic decision making experiment.” *Management Science* 35 (3): 321–339.
Taylor, D. H. 1999. "Measurement and analysis of demand amplification across the supply chain." *International Journal of Logistics Management* 10(2): 55-70.
Taylor, D.H., and A. Fearne. 2009. “Demand management in fresh food value chains: a framework for analysis and improvement.” *Supply Chain Management: An International Journal* 14(5): 379-392.
Thornhill, N. F., and M. M. Naim. 2006. “An exploratory study to identify rogue seasonality in a steel company’s supply network using spectral principal component analysis.” *European Journal of Operational Research* 172(1): 146-162.
Torres, O. A. C., and F. A. V. Moran. 2006. "Managing the Bullwhip effect through Benchmarking." *The Bullwhip effect in supply chains: A review of methods, components and cases*, O. A. C. Torres and F. A. V. Moran, eds., Palgrave Macmillan, New York, 215-242.

Appendix I

Key terms and their descriptions

Term	Description
De-trending (of data)	Refers to removing/subtracting the trend (or the increasing/decreasing) component from the time series data; this component is typically determined by fitting a line (in the least-squares sense) to the data. De-trending is required to be done for all relevant variables under consideration so that their profiles (with different trends) become comparable.
Normalisation (of data)	Refers to mean centering the time series data (of a variable), i.e. subtracting the (arithmetic) mean from all values, followed by scaling or dividing all values by standard deviation of the data. Normalisation is required to be done for all relevant variables under consideration so that their profiles (with different mean levels and of different scales) become comparable.
Frequency	Number of cycles per unit time; if the length of a cycle is 2 days, frequency = 1 cycle/2 days = 0.5 cycles/day
Fourier Transformation	Technique that decomposes a time series into its constituent sinusoids (or sine waves) at different frequencies
Spectra	Amplitudes of (constituent) sinusoids at different frequencies (determined post Fourier transformation of the time series data); peak amplitude/s at specific frequency/ies reflect corresponding periodicity in the data
Filtered spectra	Spectra where some cycles (or frequencies) have been filtered out or removed; this is done by replacing the amplitude/s for those frequency/ies with 0 values
Clustering	Technique that partitions data sets (variable profiles in time/frequency domain in this case) into a small number of homogenous groups or clusters
Euclidean distance (for dissimilarity)	The square root of the sum of squared differences between corresponding elements of the two relevant vectors $d_E = \sqrt{\sum_{k=1}^{p}\left(x_{ik} - y_{jk}\right)^2}$ d_E = Euclidean distance x_i and y_j are p dimensional vectors *For time series*: x_i and y_j would be the values at different time instants, and p, the total number of time instants *For spectra*: x_i and y_j would be amplitudes of the sinusoids at different frequencies, and p, the total number of frequencies If x_i and y_j are exactly similar, d_E computes to 0 indicating no dissimilarity
Ward's algorithm for clustering	A popular hierarchical clustering technique where each point (time series/spectra represented as a point in multidimensional space) is merged into clusters based on their relative closeness or similarity relationships; the clusters formed are merged again on the same basis and this is repeated until there is one all encompassing cluster. Merger of clusters at every stage is done by computing the sum-of-squares variance (i.e. squares of the Euclidean distances) for all possible two cluster mergers and then going with the one with the smallest value.

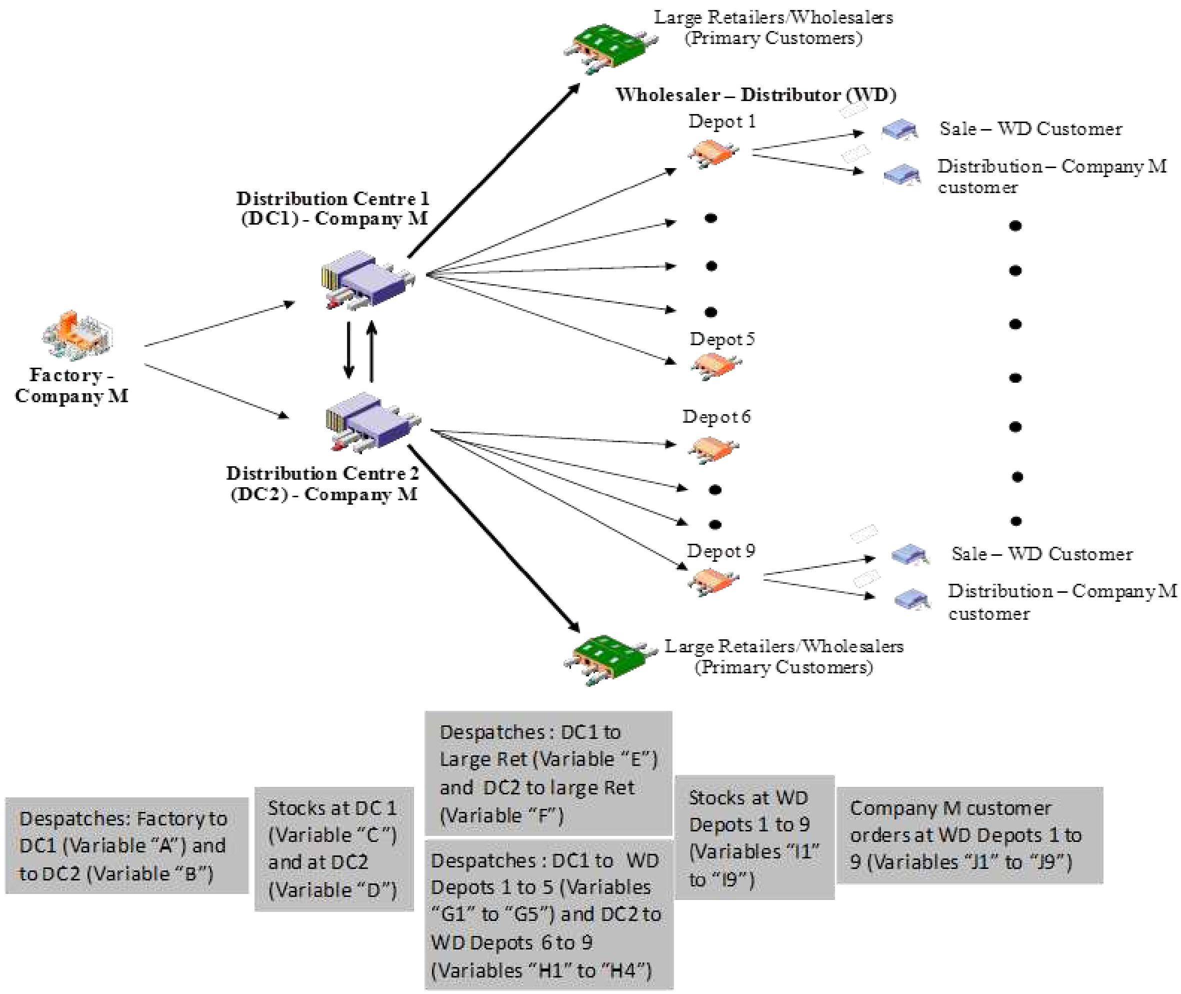

Figure 1. Downstream coffee supply chain under study and the variables considered

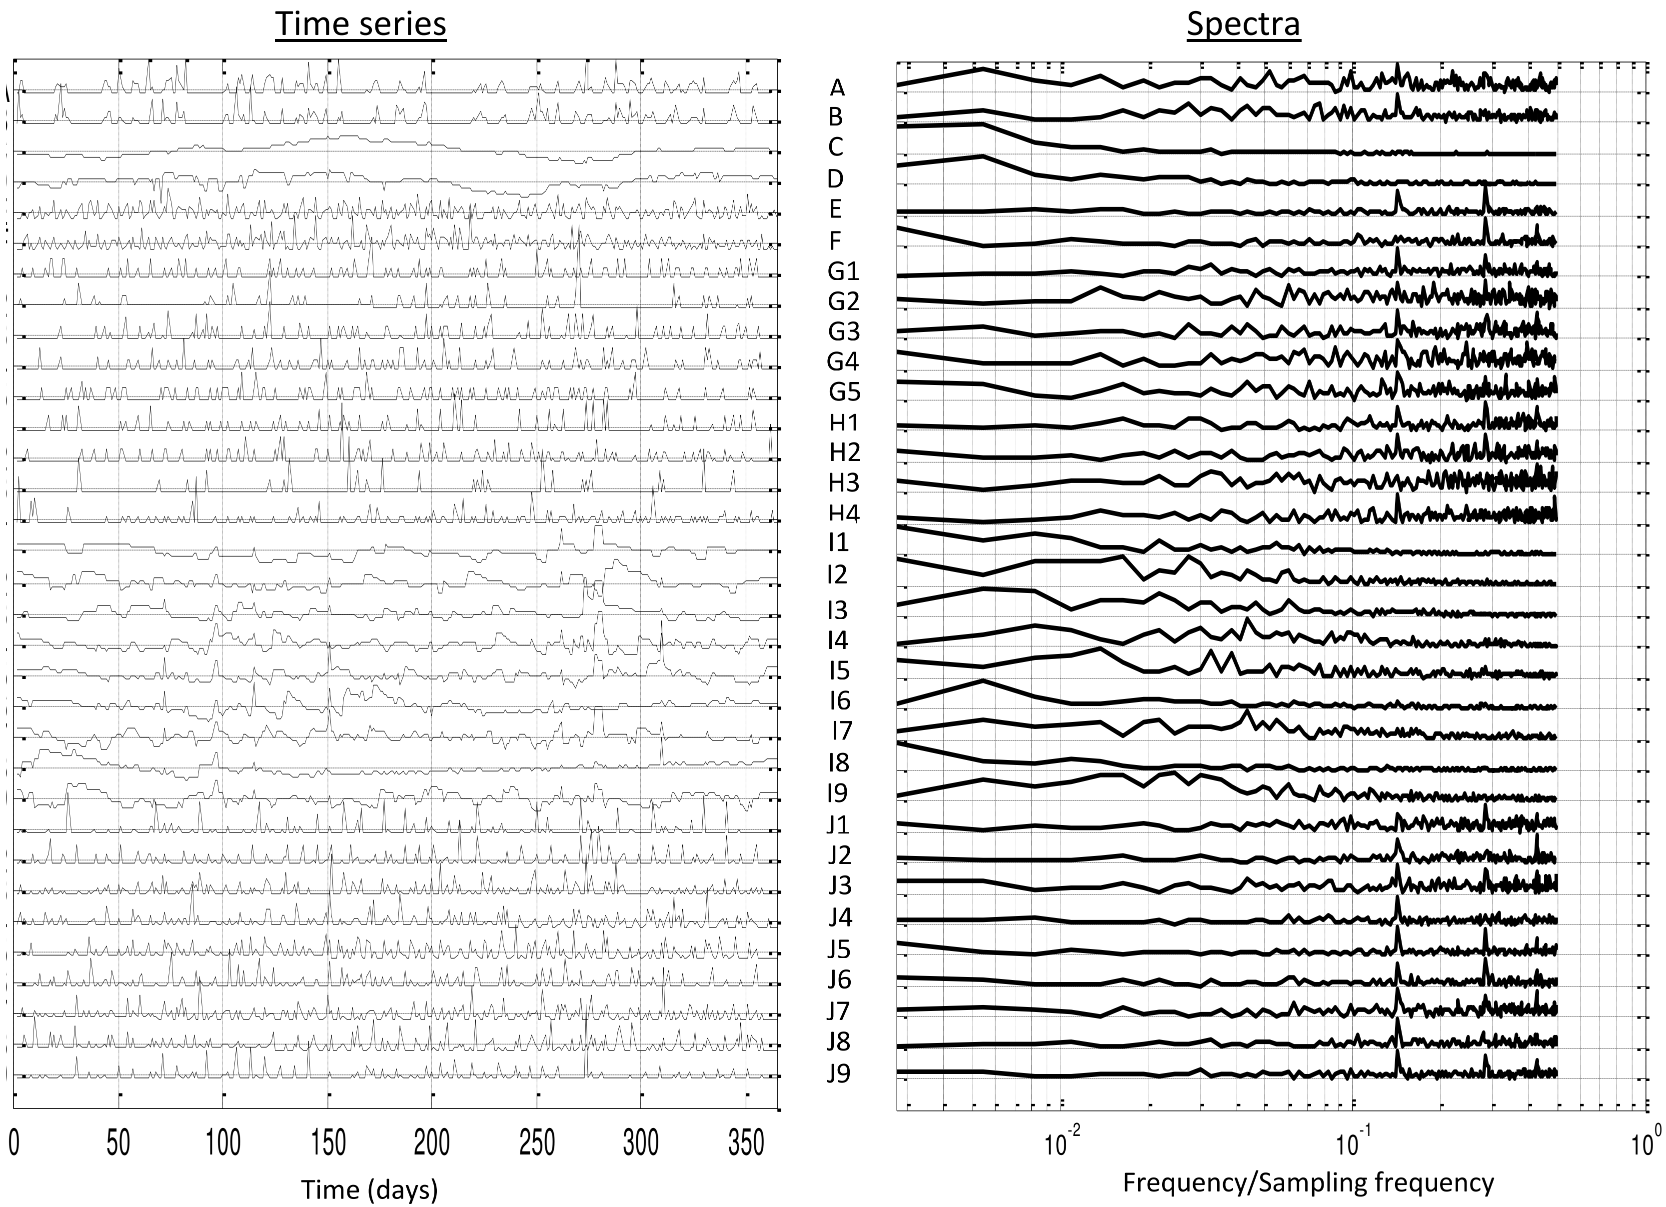

Figure 2: Time series and spectra profiles of variables for SKU1. Spectra are scaled to the same maximum peak height.

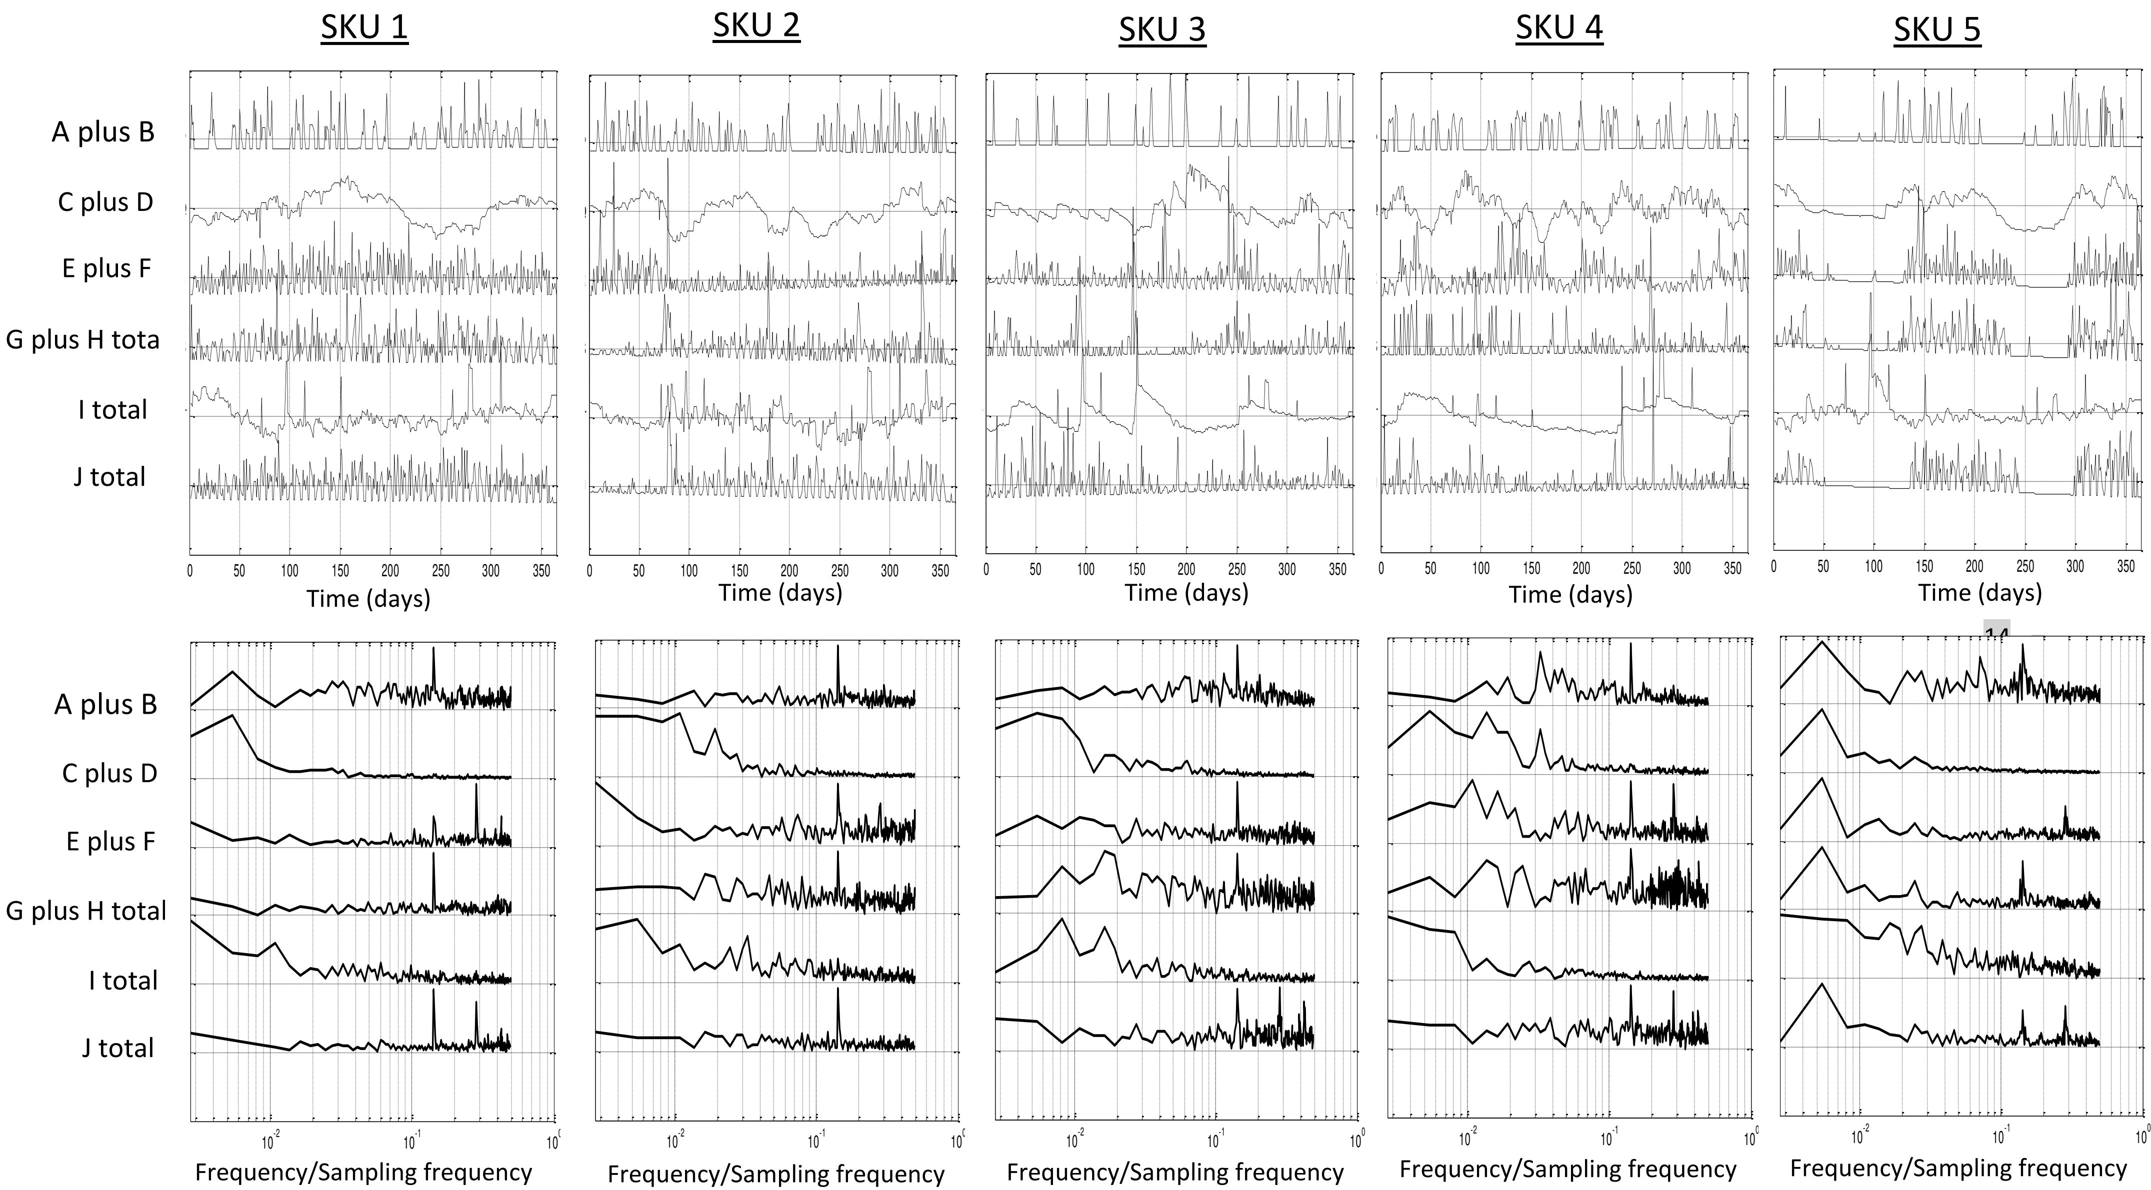

A plus B: Total despatches from factory to DCs; *C plus D*: Total stocks at DC's; *E plus F*: Total despatches from DCs to large Retailers/Wholesalers; *G plus H total*: Total despatches from DC's to WD depots; *I Total*: Total stocks at WD depots; *J total*: Total company M customer orders serviced at WD depots; X: Cycle time period (days)

Figure 3. Time series and spectra profiles of aggregated variables for different SKUs. Spectra are scaled to the same maximum peak height.

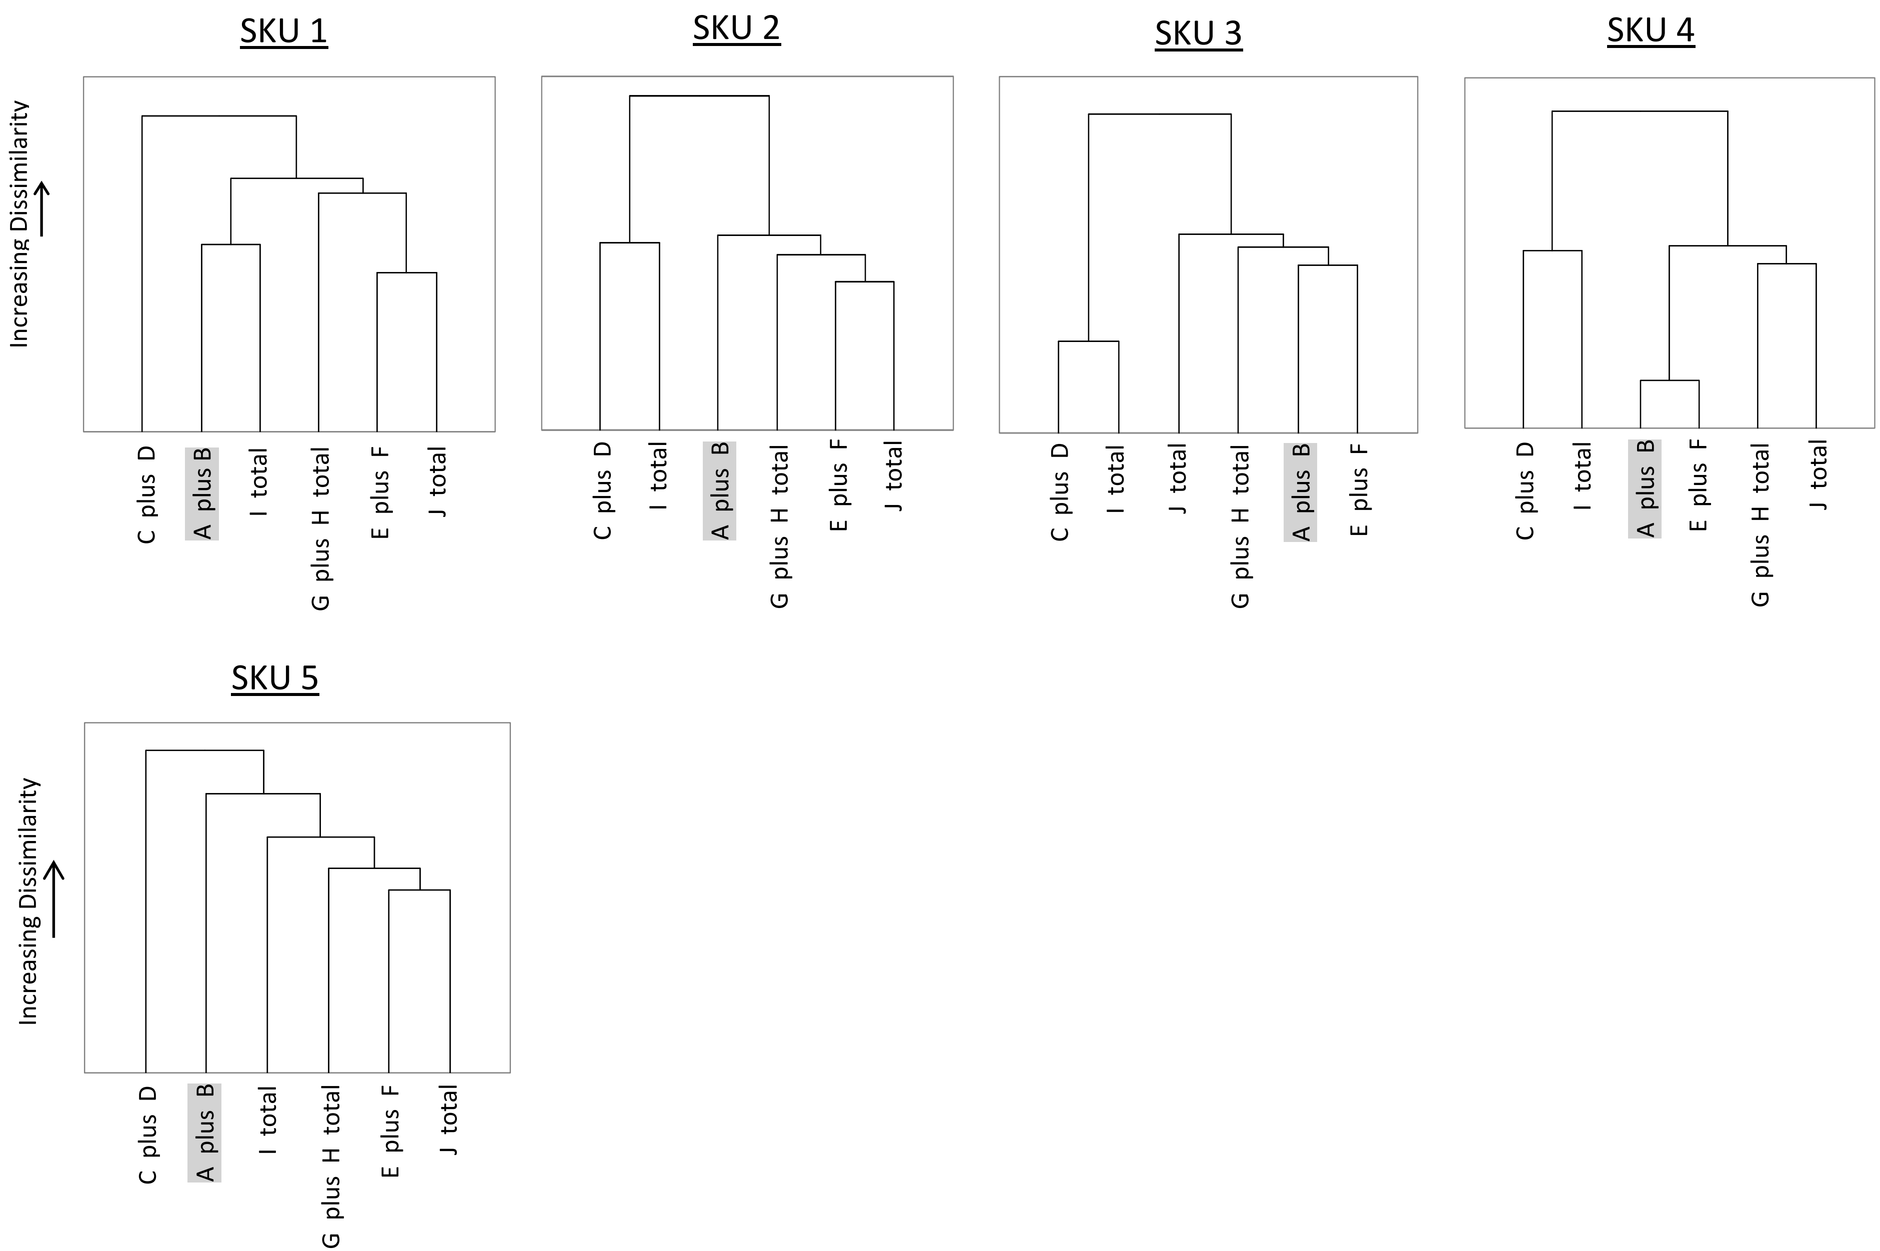

Figure 4. Cluster profiles of filtered spectra of variables for different SKUs.

Table 1

Rogue seasonality index values* at different echelons of the (aggregated) coffee supply chain for individual SKUs

At the Factory echelon		
SKU	**Index value (as per index definition 1)**	**Index value (as per index definition 2)**
SKU1	0.91	0.99
SKU2	0.93	0.99
SKU3	0.94	0.99
SKU4	0.87	0.97
SKU5	1.00	1.09
At the WD depots echelon		
SKU	**Index value (as per index definition 1)**	**Index value (as per index definition 2)**
SKU1	0.98	1.01
SKU2	0.96	1.00
SKU3	0.87	0.91
SKU4	0.90	0.98
SKU5	0.91	0.93

**Computed using filtered spectra (or amplitudes of sinusoids at various frequencies) of variables*

Avg index value (as per index definition 1) for all except shaded = 0.92

Std deviation of index values (as per index definition 1) for all except shaded = 0.04

Avg index value (as per index definition 2) for all except shaded = 0.97

Std deviation of index values (as per index definition 2) for all except shaded = 0.03

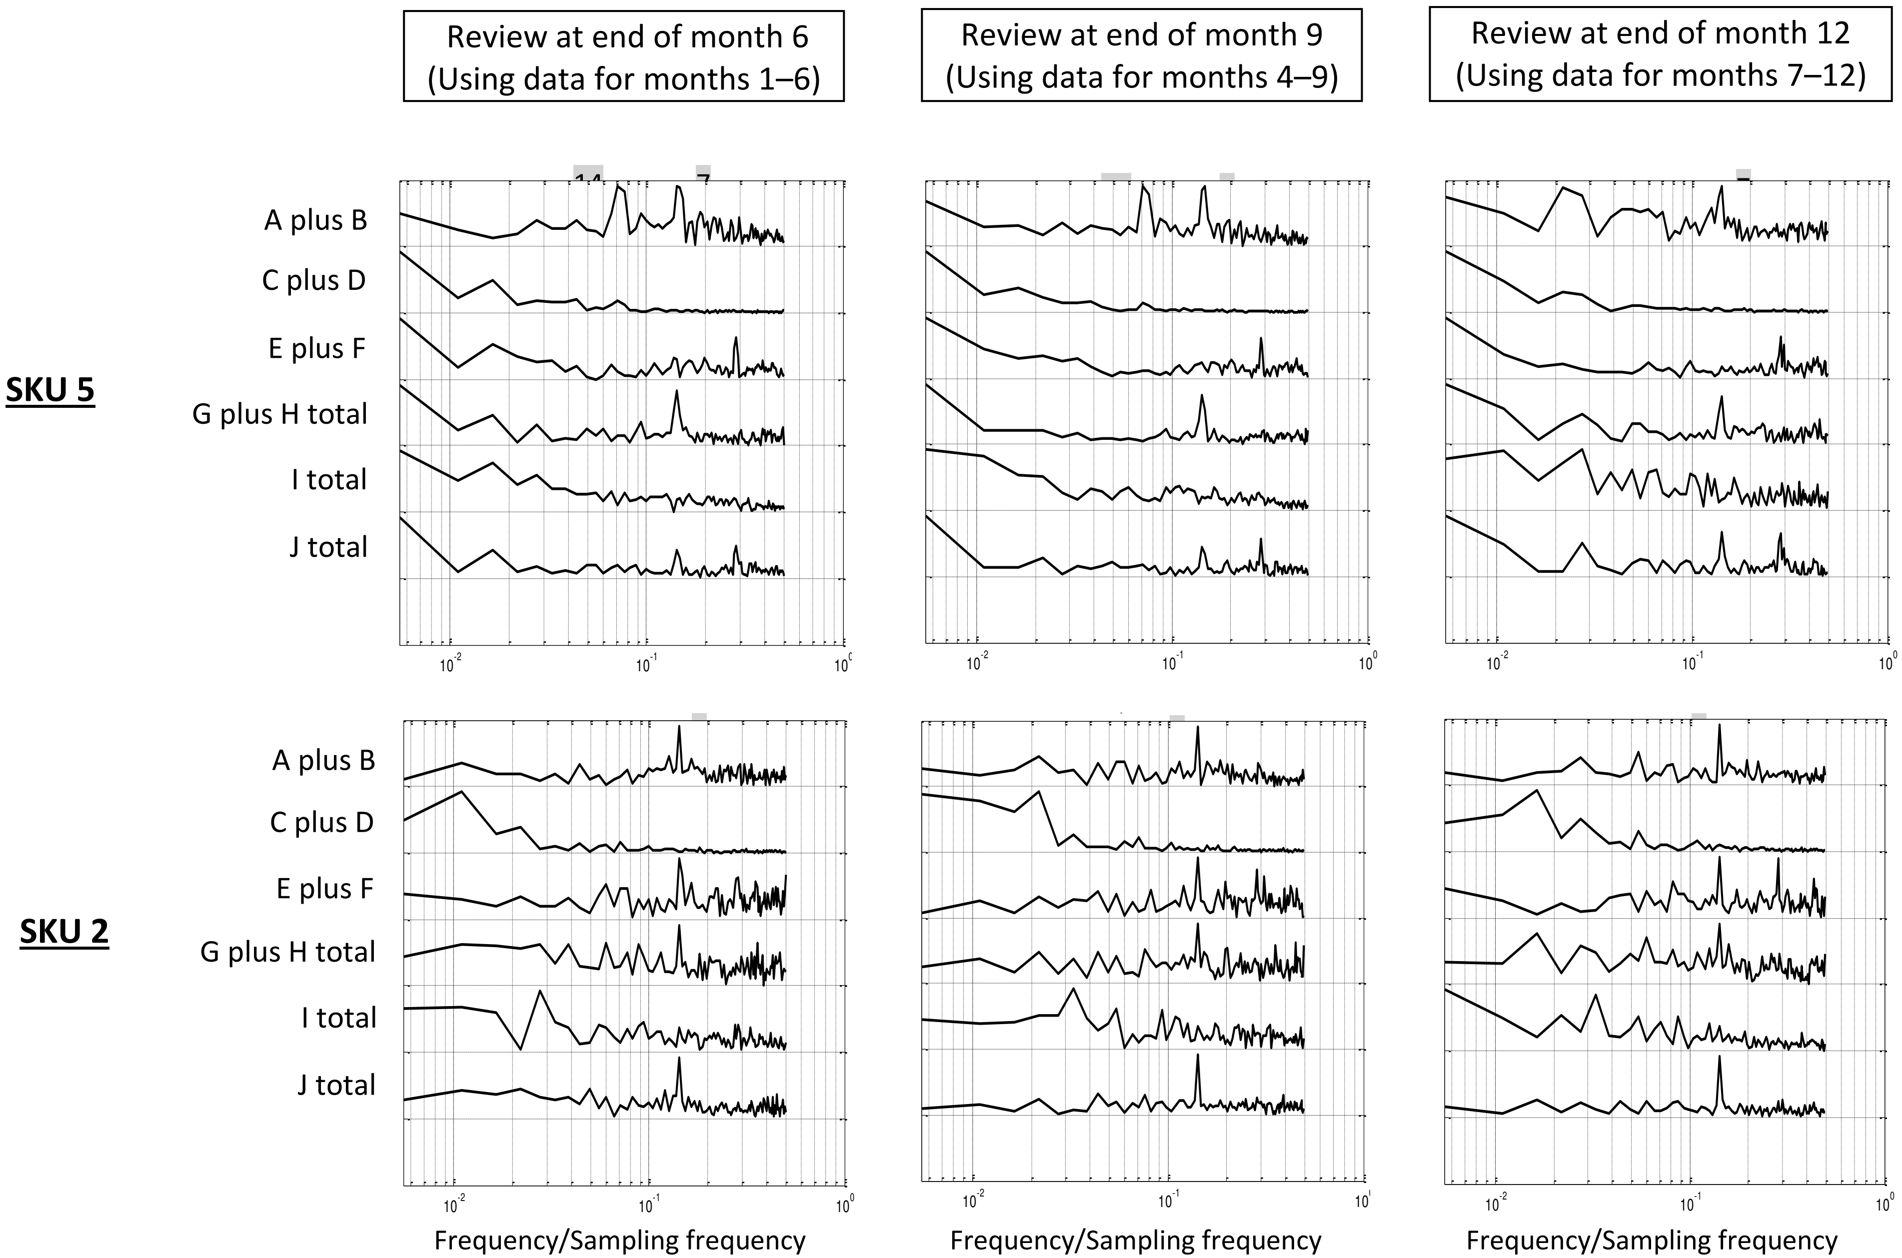

Figure 5. Spectra profiles of variables and rogue seasonality index values for the factory echelon at different time intervals.
Index value A (B) implies that A is based on index definition 1 and B on index definition 2.

Table 2

Rogue seasonality index values* at different echelons of the (aggregated) coffee supply chain at different time intervals

Index value A (B) implies that A is based on index definition 1 and B on index definition 2

SKU	At the end of month 6 (Based on data for months 1–6)		At the end of month 9 (Based on data for months 4–9)		At the end of month 12 (Based on data for months 7–12)	
	At the WD depots echelon	At the Factory echelon	At the WD depots echelon	At the Factory echelon	At the WD depots echelon	At the Factory echelon
SKU1	0.91 (0.94)	0.84 (0.97)	0.80 (0.90)	0.82 (0.98)	0.78 (0.83)	0.84 (0.97)
SKU2	0.83 (0.92)	0.94 (1.03)	0.93 (0.99)	0.91 (1.00)	0.94 (0.97)	0.94 (0.98)
SKU3	0.91 (0.98)	0.90 (0.97)	0.85 (0.93)	0.81 (0.92)	0.84 (0.89)	0.83 (0.96)
SKU4	0.92 (0.93)	0.86 (1.00)	0.83 (0.93)	0.86 (1.03)	0.90 (0.95)	0.84 (0.99)
SKU5	0.88 (0.93)	1.00 (1.17)	0.82 (0.87)	1.05 (1.18)	0.87 (0.93)	0.81 (0.95)

**Computed using filtered spectra (or amplitudes of sinusoids at various frequencies) of variables*

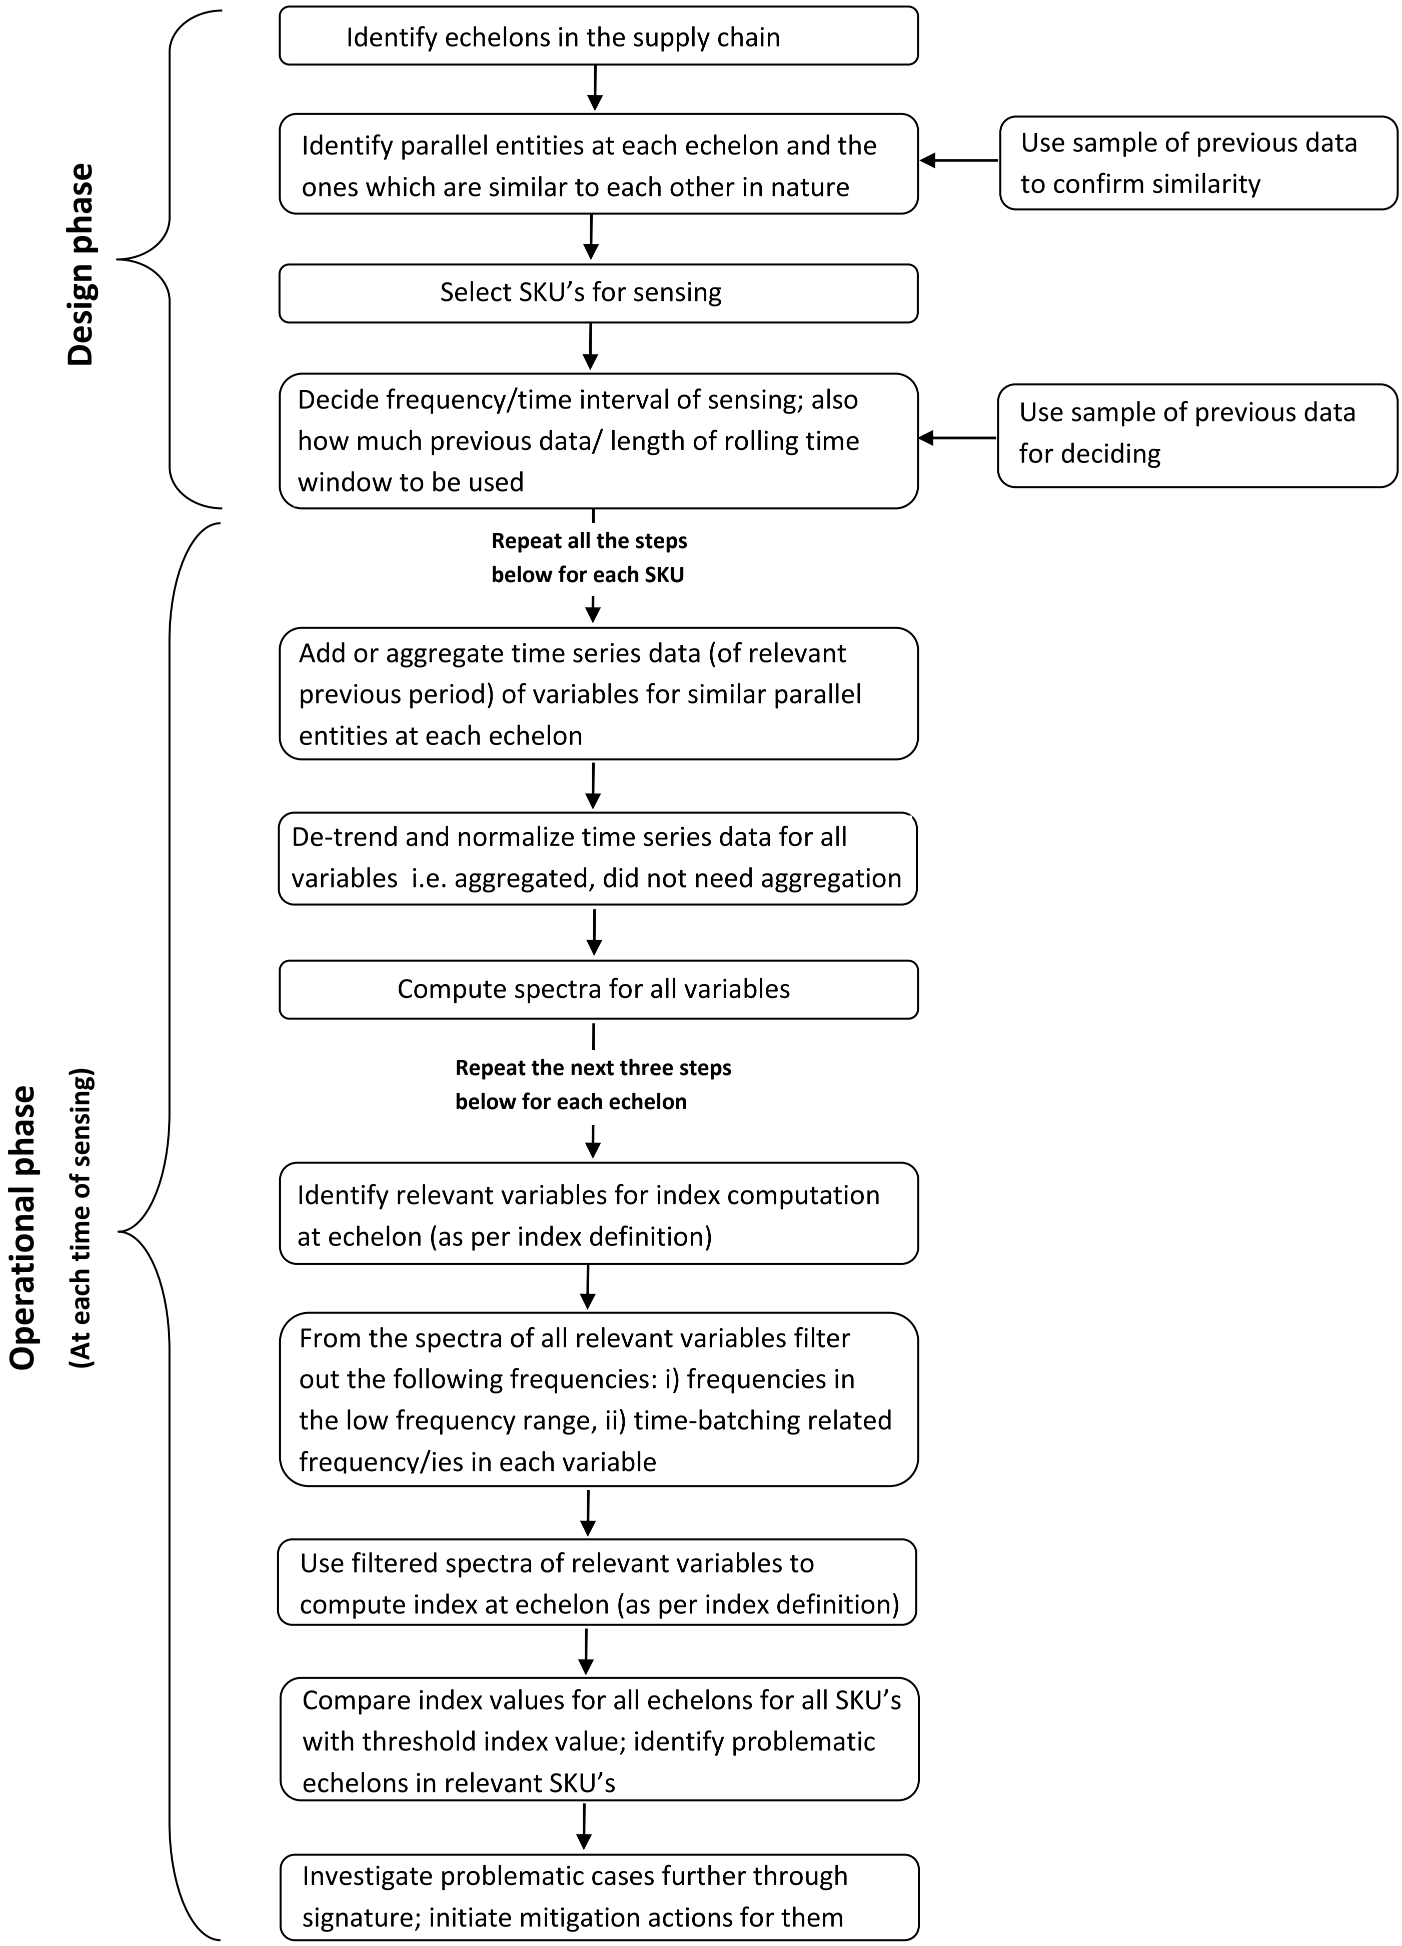

Figure 6. Flow chart for applying rogue seasonality index at a supply chain echelon level