Taylor & Francis
Taylor & Francis Group

Journal of Business & Economic Statistics

ISSN: 0735-0015 (Print) 1537-2707 (Online) Journal homepage: http://amstat.tandfonline.com/loi/ubes20

Multivariate Seasonal Adjustment, Economic Identities, and Seasonal Taxonomy

Tucker McElroy

To cite this article: Tucker McElroy (2017) Multivariate Seasonal Adjustment, Economic Identities, and Seasonal Taxonomy, Journal of Business & Economic Statistics, 35:4, 511-525, DOI: 10.1080/07350015.2015.1123159

To link to this article: http://dx.doi.org/10.1080/07350015.2015.1123159

Multivariate Seasonal Adjustment, Economic Identities, and Seasonal Taxonomy

Tucker McElroy
Center for Statistical Research and Methodology, U.S. Census Bureau, Washington, DC 20233-9100 (*tucker.s.mcelroy@census.gov*)

This article extends the methodology for multivariate seasonal adjustment by exploring the statistical modeling of seasonality jointly across multiple time series, using latent dynamic factor models fitted using maximum likelihood estimation. Signal extraction methods for the series then allow us to calculate a model-based seasonal adjustment. We emphasize several facets of our analysis: (i) we quantify the efficiency gain in multivariate signal extraction versus univariate approaches; (ii) we address the problem of the preservation of economic identities; (iii) we describe a foray into seasonal taxonomy via the device of seasonal co-integration rank. These contributions are developed through two empirical studies of aggregate U.S. retail trade series and U.S. regional housing starts. Our analysis identifies different seasonal subcomponents that are able to capture the transition from prerecession to postrecession seasonal patterns. We also address the topic of indirect seasonal adjustment by analyzing the regional aggregate series. Supplementary materials for this article are available online.

KEY WORDS: Co-integration; Dynamic factor models; Seasonality; Time series; Trends; VAR.

1. INTRODUCTION

The notion that an observed time series is composed of several unobserved components has a long history in economics, going back to mid nineteenth century England, according to Nerlove, Grether, and Carvalho (1979). The study of seasonality goes back to this epoch as well, in the field of meteorology (Buys Ballot 1847); changing seasonality was noted by Gilbart (1852). The view of an economic time series being composed of trend, seasonal, cycle, and residual effects was exposited by Persons (1919), which substantially impacted later work in the 20th century. The idea that the latent effects, especially the seasonal, could be common to multiple series (or at least, highly correlated) can be dated at least to Nerlove (1964, p. 263): "Indeed, seasonality does not occur in isolated economic series, but seasonal and other changes in one series are related to those in another. Hence, ideally one should formulate a complete econometric model in which the causes of seasonality are incorporated directly in the equations." Bell and Hillmer (1984) provided a slightly dated, but still pertinent, overview of the literature and topics of interest in the field of seasonal adjustment. Nerlove's proposal is the subject of this article.

This article studies the modeling and extraction of seasonality in multiple time series, using an unobserved components framework where each latent component is described via a dynamic factor model. Our proposed model for m-variate time series data $\{y_t\}$ takes the form

$$\begin{aligned} y_t &= \mu_t + \xi_t + \iota_t \\ (1-B)^d \, \mu_t &\sim \mathrm{WN}(0, \Sigma^{\mu}) \\ (1 + B + B^2 + \cdots + B^{p-1})\, \xi_t &\sim \mathrm{WN}(0, \Sigma^{\xi}) \\ \iota_t &\sim \mathrm{WN}(0, \Sigma^{\iota}), \end{aligned}$$

where B is the backshift operator, p is the seasonal period, and the latent processes are trend $\{\mu_t\}$, seasonal $\{\xi_t\}$, and irregular $\{\iota_t\}$. Each of these latent processes is driven by a vector white-noise (WN) process of covariance Σ. As will be shown later, the Cholesky factorization of each Σ matrix entails an interpretation of each latent component as a dynamic factor process, potentially of reduced dimension.

While there is a tremendous amount of economic literature on dynamic factor models (Stock and Watson 2011 provided an overview), there is somewhat less material on structural dynamic factor models, wherein the identification of the factor ranks for each latent component is considerably more challenging. There are even fewer serious attempts at multivariate modeling and seasonal adjustment using an unobserved components framework. However, there exists persistent (and publicly expressed) interest among economists in multivariate seasonal adjustment, as has been highlighted recently by public concerns about the performance of univariate seasonal adjustment procedures (such as X-12-ARIMA) under the extreme economic conditions engendered by the Great Recession (GR); see Zentner, Amemiya, and Greenberg (2011), Alexander and Greenberg (2012), Alexander, Zentner, and Greenberg (2012), and Feroli (2012).

In the last 20 years or so, development of the concept of co-integration—along with its extension to component co-integrating rank—and better techniques of computation (such as state-space methods) have facilitated the multivariate signal extraction project, although only a few articles have addressed seasonal adjustment, such as Bartelsman and Cleveland (1993), Krane and Wascher (1999), Birrell, Steel, and Lin (2011), Koopman, Ooms, and Hindrayanto (2012), and Greenaway-McGrevy (2013). Facilitating the stochastic approach proposed by Geweke (1978) is the software STAMP (Koopman et al. 2009), which can be used to produce multivariate seasonal adjustments.

In the Public Domain
Journal of Business & Economic Statistics
October 2017, Vol. 35, No. 4
DOI: 10.1080/07350015.2015.1123159
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/jbes.

This article develops tools to model and seasonally adjust multiple economic time series, using the concept of component co-integrating rank, and moreover explores the topics of: (i) efficiency gain in multivariate signal extraction versus univariate approaches, (ii) the preservation of economic identities, and (iii) seasonal taxonomy via the device of seasonal co-integrating rank. In Section 2, we introduce these three facets. Section 3 provides the modeling and seasonal adjustment methodology, and Section 4 discusses two empirical illustrations. Section 5 concludes, and Section 6 reviews the supplementary materials—Appendices, Data, and R code (the *sigex* program).

2. RANK, IDENTITIES, AND TAXONOMY

2.1 Dimension Reduction and Latent Rank

Dynamic factor models, which constitute our basic modeling framework in this article, have been found to be quite empirically useful; see Sargent and Sims (1977), Giannone, Reichlin, and Sala (2005), and Stock and Watson (2011, 2012). We next provide the statistical motivation for dynamic factor models, from the standpoint of dimension reduction and model parsimony.

Consider an m-variate time series $\{y_t\}$. When m is large, it is important to consider a dimension reduction approach to the data analysis, because the number of parameters will otherwise be too large for meaningful results. Bozik and Bell (1987) appeared to be one of the first articles advocating principal components analysis (PCA) to estimate a dynamic factor model, which can then be used to reduce $\{y_t\}$ to a lower-dimensional factor time series $\{x_t\}$:

$$y_t = \Lambda\, x_t + \iota_t. \tag{1}$$

Here $\{x_t\}$ is a latent vector seasonal process of dimension k and $\{\iota_t\}$ is the error process in this decomposition. Λ is an $m \times k$ matrix, called the loading matrix. The hope of such a factor analysis is that we can take k much smaller than m, while the error $\{\iota_t\}$ contributes little to the overall variation, and moreover is not serially uncorrelated. If this is the case, one can estimate Λ and the factor series using PCA, and proceed with modeling the estimated factors.

However, when working with monthly or quarterly raw economic time series such a decomposition often fails to result in dimension reduction, as much of the interesting trend and seasonal structure is shifted into the estimated errors, unless k is taken quite large. Notwithstanding, some authors (Stock and Watson 2012) have found this PCA approach to be useful on seasonally adjusted data, and have drawn useful interpretations from the resulting factor series. Unfortunately, when seasonality is present in the data, it happens that the estimated factors have trend, seasonality, and cyclicality all mingled together in each factor series (i.e., each component of the vector process $\{x_t\}$ will contain *all* the dynamics of the data process), and so this type of decomposition is not useful for separating out different dynamics. (Another approach to estimating such a decomposition is independent components analysis (ICA), but the same empirical behavior has been observed by the author.)

Because the single factor model (1) is ineffective at dimension reduction in seasonal time series, and cannot cleanly separate disparate dynamics, it is natural to look for a more complex decomposition that achieves these objectives. The key is to have several factor components, each one associated with particular dynamics that are present in the series. This takes the form

$$y_t = \mu_t + \xi_t + \iota_t. \tag{2}$$

Equation (2) decomposes the observed series into trend $\{\mu_t\}$, seasonal $\{\xi_t\}$, and irregular $\{\iota_t\}$ processes. The trend and seasonal in turn could be written as a dynamic factor model $\Lambda\, x_t$, where x_t is a latent trend or seasonal process, respectively, effectively generalizing (1). One could also introduce a latent cyclical process to model moderate 2–10 year swings about the long-term trend, but we avoid such devices in this article.

The framework of (2), when appropriately extended, has proven useful in our empirical work for separating out dynamics of different types (it is also used in Koopman, Ooms, and Hindrayanto 2012), while also allowing for some dimension reduction. This article adopts the latent dynamic factor framework, and seeks to answer the following questions: how do we fit models to economic time series, such that we allow for dimension reduction? Given the models, how can we estimate and remove seasonality? What are the advantages over a univariate approach? Although there are some recent publications treating multivariate modeling and seasonal adjustment, there seems to be no systematic treatment of latent rank, and its ramifications on seasonal adjustment. In fact, our research has been partially motivated by strong public criticism of univariate seasonal adjustment procedures, such as X-12-ARIMA, and the challenges implicit in the GR. See Wright (2013) for an overview of these concerns. One might hope that modeling multiple series would facilitate superior estimation of latent seasonal effects, thereby ensuring that quickly evolutive seasonality does not pollute trend and business cycle extractions, such as can occur during epochs of great change.

2.2 Maintaining Economic Identities

Although of little interest to the theoretician, economic accounting rules are extremely important to the publishers and consumers of economic data. These accounting rules may be aggregation relations across stratifications (e.g., male unemployment plus female unemployment equals total unemployment), regions (e.g., housing starts series for South, West, Northeast, and Midwest, must sum to the Total series), or epochs (e.g., three monthly figures must sum to the corresponding quarterly figure, for a flow time series). A discussion of accounting constraints is provided in Quenneville and Fortier (2012). The key challenge that arises is that data arise from diverse sources (e.g., surveys and/or administrative records) of varying quality, and are typically subject to sampling error—see Tiller (2012) for an overview—and are revised over time. Various recipes, adapted to the protocols and culture of each particular agency, are used to balance accounting constraints and ensure economic identities hold—these recipes include raking and using controls (to a more reliable data source).

Here, we are concerned with the disturbance of accounting rules resulting from seasonal adjustment. Any signal extraction, even when linear in the data, will disturb the raw data's economic identities. The naïve solution is to declare the aggregate variable's seasonal adjustment to be the appropriate aggregation of the component variables' seasonal adjustments—a procedure known as indirect adjustment. However, in many cases this results in inadequate seasonal adjustment of the aggregate, as seasonality yet remains. Where does it come from? One explanation is that the component variables are cross-correlated in their seasonal dynamics, and this is unaccounted for by a univariate seasonal adjustment methodology (see the excellent discussion in Geweke 1978), resulting in an aggregate seasonal adjustment that still has seasonality.

Suppose that the seasonal coherency in the raw disaggregate time series is modeled and measured, and accounted for in the multivariate seasonal adjustment. Because the multivariate seasonal adjustment is produced via a filter that acts upon *all* the input series, we do not expect the cross-spectra to have seasonal peaks. This heuristic argument says that indirect seasonal adjustment is safer when using a multivariate approach; see Planas and Campolongo (2001) for support of this point, and a study of the precision gains in the indirect approach. We pursue this idea further in Section 4.3, through the analysis of regional housing starts.

2.3 Seasonal Taxonomy

The pattern of seasonality varies greatly by industry and series type (e.g., retail vs. construction), but certain facets are common to batches of coherent time series. It is of interest to group and classify series by these features, to understand which series are driven by a common latent seasonal process. This classification, or taxonomy, can assist in detecting new patterns (e.g., departure of one series from its cluster); it can help in understanding redundancies and coherence (e.g., do some series lead or lag others with respect to their seasonal movements?); and it can provide a general portrait of the economic variable (Granger 1966). The applications of this taxonomy are at this stage speculative, but may include the following: identification of batches of series suitable for joint multivariate analysis and adjustment (or forecasting); identification of structural changes to the economy, when series that were formerly classified as belonging to the same species no longer do so; identification of data inaccuracies, when co-integrating relations are violated at particular sample points.

Given that taxonomy is of interest, tools are needed to provide measures of clustering. Coherence is the analog of correlation for time series, but here we focus on seasonal coherence, that is, high spectral coherence at seasonal frequencies (those of the form $2\pi j/p$ for integer j, and p the seasonal period). The latent component model (2) implies that the spectral density (of the differenced series) evaluated at seasonal frequencies is equal to the spectrum driving the seasonal $\{\xi_t\}$, so that there is an immediate connection (further discussed in Definition 1). As is shown in Section 3, reduced rank in the covariance matrix of the white-noise driving $\{\xi_t\}$ is known by the term "reduced seasonal rank," and implies there is seasonal co-integration in the observed series; it corresponds to maximal possible coherence at the seasonal frequencies. We propose to use seasonal rank as a measure of taxonomic proximity. In particular, if m series have seasonal rank equal to one, they belong to the same species, all of them being driven by the same latent one-dimensional seasonal process.

3. MODELING METHODOLOGY

This section provides a discussion of a multivariate time series model involving latent components for trend, seasonality, cycle, and irregular. We discuss co-integration and latent rank, and review signal extraction methodology. For further background, see Harvey (1989), Durbin and Koopman (2001), and Koopman, Ooms, and Hindrayanto (2012).

3.1 Latent Component Models

Let us consider (2) in more detail; we will further decompose the seasonal into its "atomic" components corresponding to the complete factorization of the seasonal unit root differencing polynomial. The latent processes are related to the observed m-dimensional time series $\{y_t\}$ via (2), and each latent process in turn is a difference stationary vector time series driven by potentially collinear white noise. When the white noise is collinear, the latent process is said to be "common," and there is a reduced dimension representation. When the white noise is not collinear, the latent process is said to be "related," and there is no dimension reduction possible; if the white-noise covariance matrix is also diagonal, then the latent process is said to be "unrelated."

We assume that the latent trend process $\{\mu_t\}$ is differenced to stationarity by application of $(1-B)^d$ for $d = 0, 1, 2$, where this scalar polynomial is applied to each component of the process, that is,

$$(1-B)^d \mu_t = \epsilon_t^{\mu}$$

for an m-variate white-noise process $\{\epsilon_t^{\mu}\}$ of covariance matrix Σ^{μ}. Likewise, the latent seasonal process $\{\xi_t\}$ is reduced to stationarity by application of $U(B) = 1 + B + \cdots + B^{p-1}$, where p is the seasonal period. For some applications, it is possible to use the model $U(B)\,\xi_t \sim \mathrm{WN}(0, \Sigma^{\xi})$, referred to as the *aggregate seasonal model*, but for some data a more flexible structure is needed, which we next describe. The $U(B)$ operator is factorized into

$$U(B) = \prod_{j=1}^{p/2-1} (1 - 2\cos(2\pi j/p)B + B^2)(1+B)$$

when p is even, and when p is odd we obtain

$$U(B) = \prod_{j=1}^{(p-1)/2} (1 - 2\cos(2\pi j/p)B + B^2).$$

We focus on the p even case below, as modifications for p odd are trivial to make. We next introduce *atomic seasonal processes* $\{\xi_t^{(j)}\}$ for $1 \leq j \leq p/2$, which are defined such that a single factor of $U(B)$ reduces each atomic seasonal to stationarity. By definition, the jth atomic seasonal $\{\xi_t^{(j)}\}$ for $1 \leq j \leq p/2$ satisfies

$$\begin{aligned}(1 - 2\cos(2\pi j/p)B + B^2)\,\xi_t^{(j)} &= \epsilon_t^{(j)} \qquad 1 \leq j \leq p/2-1 \\ (1+B)\,\xi_t^{(p/2)} &= \epsilon_t^{(p/2)},\end{aligned}$$

where each $\{\epsilon_t\}$ is m-variate white noise with covariance matrix Σ (i.e., $\Sigma^{(j)}$ for the jth atomic seasonal, and Σ^{ι} for the irregular). The latent seasonal process $\{\xi_t\}$ is related to the atomic seasonal components via aggregation

$$\xi_t = \sum_{j=1}^{p/2} \xi_t^{(j)}. \tag{3}$$

It is easy to check that its minimal differencing polynomial is $U(B)$ when all the atomic seasonal processes are nonzero.

Although using atomic seasonals (as opposed to the simpler aggregate seasonal) introduces additional parameters, it affords the model an additional flexibility that is crucial for modeling retail and construction data over the GR. As elaborated in Section 4, this more nuanced seasonal structure is able to capture the transition of seasonal behavior from pre- to post-GR. Essentially, the atomic seasonals allow each of the six spectral peaks in the spectral density to have separate parameters controlling width and height. If instead only one parameter (or matrix) controls all the spectral peaks' features, then some peaks may be modeled with heights and/or widths that are inappropriate. This can gravely impact seasonal adjustment: if a peak is modeled too narrowly, then the resulting seasonal adjustment filter's frequency response function will have seasonal troughs that are too narrow, and seasonality at that particular seasonal frequency may well remain. This problem is of lesser concern in times of economic regularity, but when transitioning between economic regimes (e.g., from pre-GR to post-GR), the spectral seasonal peaks grow wider, reflecting the more highly evolutive nature of seasonality.

The difference polynomial for the aggregate $\{y_t\}$ is $(1-B)^d U(B)$, and by its application we obtain

$$\begin{aligned}
(1-B)^d U(B) y_t &= g_\mu(B)\,\epsilon_t^\mu + \sum_{j=1}^{p/2} g_j(B)\,\epsilon_t^{(j)} + g_\iota(B)\,\iota_t \\
g_\mu(B) &= U(B) \\
g_j(B) &= (1-B)^d \prod_{k\neq j} \delta^{(k)}(B) \\
g_\iota(B) &= (1-B)^d U(B) \\
\delta^{(j)}(B) &= 1 - 2\cos(2\pi j/p)B + B^2 \quad 1 \le j \le p/2-1 \\
\delta^{(p/2)}(B) &= 1 + B.
\end{aligned} \tag{4}$$

The differenced observed process on the left-hand side of (4) will then be denoted by $\{\partial y_t\}$. For all of our applications (model fitting and signal extraction), it is necessary to compute the autocovariance function of each summand process in (4), so we now discuss how these functions can be easily computed. First note that the spectral density of $\{\partial y_t\}$ is real-valued and is given by

$$f(\lambda) = |g_\mu(z)|^2\,\Sigma^\mu + \sum_{j=1}^{p/2} |g_j(z)|^2\,\Sigma^{(j)} + |g_\iota(z)|^2\,\Sigma^\iota \tag{5}$$

with $z = e^{-i\lambda}$, under the assumption that all the latent process' white noises are uncorrelated with one another, and are uncorrelated with $\{\iota_t\}$. Each summand of (5) is a known scalar function times a covariance matrix Σ, and hence corresponds to the spectral density of a simple vector moving average—the autocovariance is then extremely easy to compute, and we can simply sum up these autocovariances to obtain the sequence for $\{\partial y_t\}$. These autocovariances, together with the multivariate Durbin–Levinson algorithm (Brockwell and Davis 1991), provide a stable and efficient method for computing the Gaussian likelihood. (For computation of the Gaussian likelihood, our R code *sigex* adopts the paradigm of Bell (1984), wherein initial values are assumed to be uncorrelated with the various $\{\epsilon_t\}$ white noises. This differs from the diffuse likelihood of state-space approaches, as implemented in STAMP. See McElroy and Trimbur (2015) for discussion.)

There are models implied for each of the individual series, which of course can differ quite a bit from a univariate model fitted to the particular series. Due to the extremely simple structure of the unobserved component models, these implied models are simple to derive. Let e_ℓ denote the ℓth unit vector of $\mathbb{R}^m$, then the ℓth series has trend innovation variance $e_\ell' \Sigma^\mu e_\ell = \Sigma^\mu_{\ell,\ell}$, and so forth. If we filter the ℓth series $\{y_t^{(\ell)}\}$ with the univariate signal extraction filter (for details, see below) corresponding to this implied univariate model, we obtain $\mathbb{E}[s_t^{(\ell)} | \{y_t^{(\ell)}\}]$, which can be quite different from $\mathbb{E}[s_t^{(\ell)} | \{y_t\}]$. Moreover, for a Gaussian process the mean square estimate (MSE) that is generated from the same methodology will correspond to $\mathrm{var}[s^{(\ell)} | \{y_t^{(\ell)}\}]$ rather than the smaller $\mathrm{var}[s^{(\ell)} | \{y_t\}]$.

3.2 Collinearity, Orthogonality, and Co-Integration

The case of collinear latent innovations can now be discussed. Each latent process (trend, irregular, etc.) in the model is driven by a white-noise innovation process $\{\epsilon_t\}$ of covariance matrix Σ. The various matrices Σ^μ, Σ^ι, etc., contain all the parameters of the model (excepting an m-vector of means). Any or all of these latent processes can have collinear innovations—this happens if the corresponding Σ has reduced rank. Each such covariance matrix Σ has a unit lower triangular generalized Cholesky decomposition:

$$\Sigma = L\,D\,L',$$

where L is unit lower triangular and D is diagonal with nonnegative entries. In such a decomposition, the diagonal entries of D are interpretable as partial variances (see Appendix A for discussion). If the rank is $k \le m$, then $m-k$ of these partial variances will be zero; let J denote the subindices of $\{1, 2, \ldots, m\}$ such that $d_j > 0$ for $j \in J$. Then with $L_{\cdot j}$ denoting the jth column of L, we can write

$$\Sigma = \sum_{j\in J} d_j\, L_{\cdot j}\, L_{\cdot j}'. \tag{6}$$

Note that the partial variances need not be ordered, so that zero values of the diagonals can occur at any index (however, a value of $d_1 = 0$ will typically not occur in practice, as it means that the first variable of that latent component is deterministic). If estimating Σ through a parametric model—say via maximum likelihood estimation (MLE)—we can proceed as described in Pinheiro and Bates (1996): all lower triangular values of L are unconstrained real numbers, whereas the nonzero values of D can be described as the exponentials of real numbers. Clearly, the number and format of such parameters depends on knowing the rank k, and also the sub-indices composing J; each choice of J constitutes a different model, requiring separate estimation.

So each configuration J of restrictions on the rank of D constitutes a nested model within the nesting model, which is the fully unconstrained case wherein all covariance matrices have full rank. To obtain a more parsimonious model, it is of interest to determine whether the innovations are collinear. Because co-integration tests in the econometric literature are focused on the case of common trends (see Nyblom and Harvey 2000, 2001), we take a fresh approach to the problem of seasonal co-integration that is based on the MLEs. Our method here lacks a distribution theory, although given the asymptotic distribution of the MLEs it seems plausible that such a theory could be developed. (Although Busetti (2006) made extensions of Nyblom–Harvey to the seasonal case, his procedures conduct inference on the rank k and not on the configuration J.) We are looking for any partial variances that are suitably close to zero—this is further developed in the next subsection.

Given one or more small values for the partial variances, we consider the nested model given by setting these quantities to zero, thereby obtaining the index set J of size k. The corresponding columns of L are also eliminated—these are the $L_{\cdot j}$ for $j \notin J$ in (6). Labeling the resulting rectangular lower triangular matrix by Λ and the diagonal matrix of corresponding nonzero partial variances by Δ, it is seen that this Λ exactly corresponds to the factor loading matrix described in (1). The corresponding factor time series corresponds to a (nonstationary) latent stochastic process of dimension k, which has no cross-correlation; indeed, the innovations driving the factor latent stochastic process (whether trend or seasonal) will have covariance matrix Δ. This is the latent dynamic factor model interpretation of the model described herein.

Determining which partial variances to replace with zero requires some care (next subsection), but note that any such zeroes result in a model that is nested on the boundary of the parameter space (since each $d_j \geq 0$). The likelihoods for the nested and nesting models can then be directly compared. Because the distribution—under the null hypothesis that the nested model is correct—of the log-likelihood ratio is not χ^2, and the true mixture distribution is unknown, we recommend instead that one should do an Akaike information criterion (AIC) comparison. An important caution is that such zeroes should not be placed in Σ^{ι}, because then the spectral density f will be noninvertible. Actually, f is allowed to be noninvertible at a finite number of frequencies (see the discussion in McElroy and Trimbur 2015). For example, it is permissible to have collinear trend innovations and/or collinear seasonal innovations, so long as the irregular has full rank.

The entries of L have a statistical interpretation—in particular, the ijth entry of L is proportional to the partial covariance between the ith and jth variables (i.e., components $\epsilon_t^{(i)}$ and $\epsilon_t^{(j)}$ of the vector ϵ_t), conditional on variables one through $j-1$ (see Appendix A for more discussion). Hence, a zero value in L corresponds to conditional uncorrelation, and hence if all the off-diagonal entries are constrained to be zero there will be no cross-correlation between components of the corresponding random vector. This restriction is tantamount to fitting a univariate model to each series, with model fit determined in an aggregate sense across all m series. We refer to this particular submodel as the orthogonality restriction; in a sense it is the opposite of the collinear innovations case, and also involves a reduction in the number of parameters, from $\binom{m+1}{2}$ down to m.

The decision to replace a small entry of L with a zero can be made on the basis of the statistical uncertainty of the parameter. Because we use a Gaussian likelihood, the inverse of the Hessian of the log-likelihood provides an estimate of the parameter error covariance matrix, due to the efficiency of MLEs. The resulting nested model can be checked against the nesting competitor via the generalized likelihood ratio (GLR) test, using χ^2 quantiles, because the parameter restriction of zero does not lie on the boundary of the parameter space. Because collinearity can eliminate entire columns of the L matrix, one should determine collinearity first, and then pursue orthogonality.

We next discuss the relationship to co-integration, which is also discussed in McElroy and Trimbur (2015). Generalizing the basic concepts presented in Engle and Granger (1987) and Stock and Watson (1988), we say that when an m-vector α exists such that $\{\alpha' y_t\}$ has reduced nonstationarity, then α is a co-integrating vector. By reduced nonstationarity, we mean that the minimal differencing polynomial required to reduce $\{\alpha' y_t\}$ to stationarity (up to fixed effects) has lower degree than the polynomial required for the original $\{y_t\}$. If $\{\alpha' y_t\}$ has only trend nonstationarity, α is said to be a seasonal co-integrating vector, whereas if there only seasonal nonstationarity remains, α is said to be a trend co-integrating vector.

Given our particular latent factor model, α is a jth atomic seasonal co-integrating vector if and only if α is a left null-vector of $\Sigma^{(j)}$, whereas α is a trend co-integrating vector if and only if α is a left null-vector of Σ^{μ}. This follows from the form of our model—see Equation (5). A basis for the co-integrating vectors' space can be computed from the rows of L^{-1}, using the rows that correspond to zero d_j values. For example, a rank of one implies that there exists a basis of $m-1$ co-integrating vectors. From the standpoint of taxonomy, we say that all such time series belong to the same latent species, where the type of species is defined by the particular latent frequency. For example, if $\Sigma^{(j)}$ has rank one we say that all the series belong to the same jth-atomic seasonal species.

Observe that Σ (excepting the case of the irregular covariance matrix) is the value of the spectral density f of the differenced time series at the latent process' corresponding frequency, that is, $\Sigma^{\mu} = f(0)$ and $\Sigma^{(j)} = f(\pi j/p)$. This discussion leads to the following definition.

Definition 1. Two time series following a latent dynamic factor model are j-equivalent if and only if the bivariate spectrum evaluated at the jth seasonal frequency ($0 \leq j \leq p/2$) has rank one. We denote this equivalency with the notation $\sim_j$.

This is well-defined, in the sense that permuting the series' order does not change the rank. Moreover, we have the following result:

Proposition 1. $\sim_j$ is an equivalence relation, and therefore partitions the set of difference stationary time series.

We make a few comments about this result. First, when estimating the covariance matrix of multiple time series, we can only make probabilistic assertions about the rank, and therefore statistical errors can arise; also, changing samples can alter the species classification. Second, there is a classification pertaining to each frequency $2\pi j/p$, so that up to $p/2$ different partitions exist for these types of time series. Two series might be equivalent according to frequency zero, but might not be according to frequency $j = 1$. This concept is applied in Section 4.1, with $j = 6$.

3.3 Taxonomic Identification

Having proposed a theoretical definition for taxonomy, it is crucial to have an empirical procedure for its application on data. Formal tests for whether particular d_j's are zero are the ideal target for the statistician, but using the Hessian—as with the partial covariances in L—will lead to incorrect inference because the parameter space for the partial variances has boundary. Our recommendation is to examine the d_j parameter estimates in the context of other entries in Σ, and thereby obtain nested models—by reducing the rank of each Σ matrix—and finally to use AIC comparisons to evaluate the models. To that end, we describe how condition numbers can be defined from the partial variances of each correlation matrix.

First consider the $m = 2$ case, which contains the main features of the general case as well. The partial correlation between the first and second variables (it is the same as the unconditional correlation) is

$$\kappa_{21} = \frac{\sqrt{d_1}\, L_{21}}{\sqrt{d_1\, L_{21}^2 + d_2}}.$$

This formula (see Lemma 1 of Appendix A) shows that the absolute correlation approaches unity as $L_{21} \to \pm\infty$, or if $d_2 \to 0$. Defining $\tau_2 = \log(d_2) - \log(L_{21}^2\, d_1 + d_2)$, we have $\kappa_{21}^2 = 1 - e^{\tau_2}$, and values of τ_2 tending to $-\infty$ correspond to the rank of Σ going from two to one. Observe that the determinant of Σ is equal to the product of the diagonal entries times $1 - \kappa_{21}^2$, which represents a scale-free quantity. We can define a condition number by computing the scale-free determinant of Σ in log-scale, that is,

$$\log(1 - \kappa_{21}^2) = \tau_2.$$

For example, if $|\kappa_{21}| = 0.9$ then $\tau_2 = -1.66$, whereas $|\kappa_{21}| = 0.99$ corresponds to $\tau_2 = -3.92$, and $|\kappa_{21}| = 0.999$ implies $\tau_2 = -6.22$. The advantage of examining τ_2 and κ_{21} to determine approximate singularity is that we have removed the scale of the series from the analysis.

Now let us generalize to $m > 2$: defining for $i > j$ the partial correlation $\kappa_{ij} = \mathrm{corr}(\epsilon_t^{(i)}, \epsilon_t^{(j)} | \epsilon_t^{(1)}, \ldots, \epsilon_t^{(j-1)})$, and with R the m-dimensional correlation matrix corresponding to Σ, it follows from Theorem 1 of Appendix A that the jth partial variance of R is equal to

$$\prod_{j=1}^{i-1} (1 - \kappa_{ij}^2).$$

Letting τ_i denote the log of this quantity, we have the proposed condition number

$$\log \det R = \sum_{i=1}^{m} \tau_i.$$

The τ_i can be conveniently calculated as the logged ith partial variance of R. So for a given i, any values of κ_{ij}^2 close to unity (or equivalently, large negative values of τ_i)—for any j—indicate that the ith variable can potentially be eliminated. We apply these measures to a trivariate analysis in Section 4.1; the goal of such elimination of variables is primarily parsimony—and secondarily, taxonomy.

3.4 Multivariate Filtering

Here, we consider the main application of the preceding modeling methodology to signal extraction by describing the minimum MSE linear filters corresponding to the fitted structural model. We rely upon the formulas derived in Theorem 2 of McElroy and Trimbur (2015). To implement the smoothing formula, we write the data vector as collected by time and listed over vector components. Then the covariance matrices for the differenced latent components can be computed quite easily. Let s and n denote signal and noise, where s consists of the sum of any components given in (2) that are of interest, and n consists of the remaining components. For example, n could consist of the sum of all $p/2$ atomic seasonals, and s consist of the sum of trend and irregular; then the signal extraction corresponds to seasonal adjustment.

Identification of the signal components of interest in turn implies a signal differencing operator $\delta^s(B)$, and a spectrum f_s for the differenced signal; similarly, we will have a noise differencing operator $\delta^n(B)$ and noise spectrum f_n. The signal and noise spectra will actually correspond to various summands of (5), in the following sense: the squared gain of the noise differencing operator will multiply the signal spectrum in (5), whereas the squared gain of the signal differencing operator will multiply the noise spectrum. In the seasonal adjustment example, the noise differencing operator is $U(B)$ and the signal spectrum is

$$f_s(\lambda) = \Sigma^{\mu} + |1 - z|^{2d}\, \Sigma^{\iota}.$$

On the other hand, the signal differencing operator is $(1 - B)^d$ and the noise spectrum is

$$f_n(\lambda) = \sum_{j=1}^{p/2} \Big| \prod_{k \neq j} \delta^{(k)}(z) \Big|^2\, \Sigma^{(j)}.$$

This is just one example; we might be interested in various atomic seasonals as signals, or combinations of such, and in each case f_s and f_n can be defined. The signal extraction filter for a bi-infinite sample has frf given by

$$f_s(\lambda)\, f(\lambda)^{-1}\, |\delta^n(z)|^2,$$

as proved by McElroy and Trimbur (2015) for cases including co-integration. For samples of finite length we instead use a matrix filter F, whose formula is also computed from the signal and noise spectra, as well as the differencing operators. For a sample of size T, there are two alternative ways of stacking the data into a matrix. First, we have

$$Y = [y_1\, y_2 \ldots y_T],$$

which is $m \times T$ dimensional. We call this *series-by-time*, and is conventional in many textbooks. The other representation is

$$Y' = [y^{(1)}\, y^{(2)} \ldots y^{(m)}],$$

where $y^{(j)}$ is a T-vector consisting of all observations for the jth series. Thus, Y' is $T \times m$ dimensional, and is referred to as the *time-by-series* representation. The description of F in McElroy and Trimbur (2015) presumes the time-by-series representation, so that $F\, \mathrm{vec}[Y']$ yields vecced signal extraction estimates written in time-by-series format.

Now the application of F to the vectorization of the time-by-series data matrix Y' is appropriate when the mean is zero. When the mean of the differenced series is nonzero, say given by an m-vector $\mathbf{m}$, then instead we apply F to the mean-corrected time-by-series data, where the mean correction involves subtracting $\mathbf{m} \otimes \tau$ from $\mathrm{vec} Y'$, where τ is a column vector $\tau = [1^d, 2^d, \ldots, T^d]'$ for $d \geq 0$. For the seasonal, cycle, and irregular components, we compute $F\, [\mathrm{vec} Y' - \mathbf{m} \otimes \tau]$, but for the trend we compute $\mathbf{m} \otimes \tau + F\, [\mathrm{vec} Y' - \mathbf{m} \otimes \tau]$. This procedure is justified in Appendix A.

Moreover, the error covariance matrix—whose diagonal entries are the signal extraction MSEs, or conditional variances—is given by a matrix V, which is expressed (McElroy and Trimbur 2015) as the difference of two positive definite matrices. Essentially, the first matrix corresponds to univariate signal extraction error, and the second matrix brings cross-series information into play, to increase precision when warranted. Clearly, when series' latent components are orthogonal, the multivariate approach offers no additional signal extraction precision over the univariate approach, while in the case of a common component the precision gain is maximal. In between these extremes is the case of related components, and there may be substantial gains to precision in this case (see discussion in Geweke 1978).

Section 4.2 further explores the precision increases due to multivariate signal extraction. As discussed in Section 3.1, we propose to measure the ratios

$$\frac{\mathrm{var}[s_t^{(j)} | \{y_t\}]}{\mathrm{var}[s_t^{(j)} | \{y_t^{(j)}\}]} \tag{7}$$

for each $1 \leq j \leq m$; here the numerator is given by the appropriate entries of the matrix V, and the denominator is computed from V under the assumption that all series are uncorrelated with one another. This ratio of MSEs will give an idea of how much reduction in MSE is attributable to the multivariate filtering. (Birrell, Steel, and Lin (2011) employed the same measure in their work.)

3.5 Indirect Seasonal Adjustment

The problem of indirect seasonal adjustment is that the total of the seasonal adjustments of several disaggregate series (e.g., corresponding to regions) might not equal the seasonal adjustment of the total, if this latter adjustment is done separately; see Ghysels (1997) for background. This problem is actually a repercussion of conditional expectation calculations, and has nothing to do per se with nonlinearity of filtering. The direct seasonal adjustment of the aggregate would be $\mathbb{E}[\sum_j n_t^{(j)} | \{\sum_j y_t^{(j)}\}]$, where $n_t^{(j)}$ is the jth series' nonseasonal component; indirect multivariate adjustment proceeds by summing the disaggregate multivariate adjustments, namely, $\sum_j \mathbb{E}[n_t^{(j)} | \{y_t\}]$. Note that the latter information set includes the information set of the direct case, so we should favor indirect adjustment. We should also prefer this indirect multivariate adjustment to univariate adjustment, which proceeds by summing the univariate disaggregate adjustments: $\sum_j \mathbb{E}[n_t^{(j)} | \{y_t^{(j)}\}]$.

We propose adopting the indirect multivariate seasonal adjustment, as its expectation conditions on the most amount of information. The economic identity—which, for example, states that the total shall equal the sum of the disaggregate series—is preserved automatically, and both the total seasonal adjustment and the individual seasonal adjustments are coherent, as they are computed from the same information set. Part of the problem with direct seasonal adjustment is that it proceeds from the information set $\{\sum_j y_t^{(j)}\}$, which can be quite different from the univariate information sets.

The signal extraction MSE for the total can be determined from the error covariance matrices for the individual series; if V is the error covariance matrix for the multivariate seasonal adjustment, then

$$[[1, 1, \ldots, 1] \otimes 1_T] \;\; V \;\; [[1, 1, \ldots, 1] \otimes 1_T]' \tag{8}$$

is the error covariance matrix for the aggregate, whose diagonal entries provide the time-varying MSEs (here 1_T is the $T \times T$ identity matrix). This concept is further explored in Section 4.3.

This approach becomes problematic if the data have been log transformed. The application of a log transformation implies a multiplicative decomposition in the original scale of the data, so that $y_t = s_t \cdot n_t$ for signal and noise vectors s_t and n_t, and $\cdot$ denoting Hadamard product. The total is defined as $z_t = [1, 1, \ldots, 1]' y_t$, and it is unclear how to define the corresponding signal and noise decomposition for z_t, since it equals

$$z_t = \sum_{j=1}^{m} s_t^{(j)} n_t^{(j)}. \tag{9}$$

If s_t is the seasonal, and n_t is the nonseasonal component, then one possibility is to arbitrarily define $\sum_{j=1}^m n_t^{(j)}$ as the nonseasonal component of the aggregate z_t, and their quotient as the seasonal. To apply classical signal extraction methodology, a log transformation is applied to each component of y_t, which transforms the multiplicative decomposition into an additive decomposition. After applying the signal extraction methodology in the log domain, one can exponentiate all the estimates to translate results into the original scale. (Some statistical agencies (e.g., the Bureau of Economic Analysis) prefer using a log transform, so that the final seasonal adjustment results can be interpreted as percentage adjustments, and use benchmarking algorithms to enforce accounting rules. This approach is alien to the methodology of this article, which employs model-based methods so as to carefully quantify signal extraction uncertainty.) Unfortunately, these transformations will interfere with quantifying the MSE of the indirect multivariate adjustment. More fundamentally, such an approach implicitly views $\sum_{j=1}^m n_t^{(j)}$ as the target nonseasonal component of z_t in (9), and this is an ad hoc choice corresponding to setting $s_t^{(j)} = 1$.

An alternative to this nonlinear procedure is to shorten the series' length to a degree such that a log transformation is no longer necessary. (Many series manifest their seasonal amplitude as proportional to the trend level, and thus warrant a multiplicative decomposition; by shortening the series, the impact of trend growth becomes linearized.) For monthly data, 10 years of data are typically sufficient to fit common univariate models, and longer spans actually seem to warrant more complicated models (e.g., time-varying coefficient models, or regime-switching models). When modeling multivariately, we have additional sample cross-sectionally, and hence there is less danger in restricting the series length. We henceforth proceed to work with such shortened spans, finding that our simple models work quite well without a log transformation.

3.6 Implementation and Computation

We here provide an overview of the methods discussed above, which we have implemented in the R code *sigex*. We first suppose the data have been cleaned of outliers and moving holiday effects (see Appendix C for details). The analyst makes choices about whether to log transform the series, and which latent components to include in the initial nesting model. In *sigex*, each desired component is specified through the associated differencing polynomial $\delta(B)$ and the rank configuration J. For example, an $I(1)$ trend for trivariate data, where the second partial variance is set to zero, would have $\delta(B) = 1 - B$ and $J = \{1, 3\}$. For the nesting model used for the retail and construction examples of this article, there are eight latent components with $J = \{1, 2, \ldots, m\}$ in each case, with differencing operators $(1 - B)^2$ for trend and $\delta^{(j)}(B)$ for the jth atomic seasonal $(1 \leq j \leq 6)$, given in (4).

For fast estimation, *sigex* uses a method of moments (MOM) based on Equation (5), discussed in Appendix B. One can check the residuals from this fit, and if adequate (assessed via examination of the autocorrelation plot) the analyst can proceed to dimension reduction. If more time is available, MLEs for the nesting model can be determined via nonlinear optimization of the Gaussian likelihood. Model inadequacy at this stage may indicate the merit of including additional latent components, such as a stochastic cycle or vector autoregression—but such variants are rarely needed for seasonal adjustment of retail, manufacturing, import, or construction data.

Supposing the model is adequate, the condition number $\log\det R$ can be computed for each latent component, and decomposed in terms of the log partial variances τ_i $(1 \leq i \leq m)$. Any such τ_i that is less than our threshold is flagged, and the corresponding index i is deleted from the rank configuration set J. The default threshold is -6.22, corresponding to a partial correlation of 0.999. In this way J is determined for each component, and a proposed nested model is formed. Note that these calculations can be done for both MOM and MLE estimates, but in the latter case we can refit the data to the nested model. Finally, if the nested model has adequate diagnostics, we can do an AIC comparison to judge whether there is a substantial loss indicated by the collinearity restrictions.

When the threshold for τ_i is set to a small number such as -6.22, the actual entries of Σ change extremely little (relative to the matrix' scale) when the matrix is altered to the reduced rank version, because the partial variance is very close to zero in such a case. (This follows from the representation (6).) We could now consider orthogonality restrictions, or replacing certain entries of each latent component's L matrix with a zero, if the MLE is small relative to its standard error.

Next, the analyst can proceed to signal extraction. In *sigex*, one specifies any desired aggregation of latent components as a signal. For example, in the monthly atomic seasonal model with trend and irregular, specifying components 2 through 7 yields $\xi_t = \sum_{j=1}^{6} \xi_t^{(j)}$ as the target signal, whereas the first component is the trend, and the eighth component is the irregular. The software computes matrices F and V for each such desired signal, as described in Section 3.4. Applying F to the data vector (accounting for the mean, if the signal is composed of trend) then yields the signal extraction, and confidence intervals can be constructed using V. If some aggregation over series' signals is desired (e.g., a total trend), *sigex* can be used to take linear combination of signal extractions.

We close with a few comments on computation. Everything described above (except the MOM) could be implemented in state space, though *sigex* does not do this, using recursive algorithms instead. Likelihood evaluation via the Kalman filter and the Durbin–Levinson algorithm is comparable in terms of speed and stability, though the values can differ slightly (see discussion in McElroy and Trimbur 2015). Signal extraction estimates are more quickly computed using a state-space smoother, as opposed to direct matrix formulas, but if the full matrix V is needed than the latter approach is preferable. Slightly different assumptions about the data-generating process are used in each approach, but the resulting MLEs and signal extractions differ very little in our experience.

4. EMPIRICAL ILLUSTRATIONS

We first model three retail series, and give an application of taxonomy; then, we consider multivariate seasonal adjustment of four regional construction series, with application to indirect adjustment.

4.1 Retail Series and Seasonal Taxonomy

In the course of modeling retail and construction series, we found that inclusion of a business cycle component gives little improvement to the overall models, and moreover had an obnoxious impact on seasonal adjustment: either one allows the cycle period to freely vary—in which case it can become coincident with seasonal frequencies and lead to misidentification—or one constrains the cycle period arbitrarily to some band. This latter choice produced period estimates on the boundary, and moreover there was little evidence in the estimated spectra to indicate a cycle's existence in the first place. Thus, we employ the atomic seasonal model given by (2) and (3), without a cycle.

For our first example, we study three series from the Advance Monthly Sales for Retail and Food Services data (representing a preliminary estimate of each series, featuring the largest retailers), which are published each month. We consider the following fairly highly aggregated series: 448 (Clothing and Clothing Accessories Stores); 451 (Sporting Goods, Hobby, Book, and Music Stores); and 452 (General Merchandise Stores). The sampling period was 1992 through 2012, and values pertain to the entire U.S. geography. Each series was first cleaned of fixed

Table 1. Log partial variances by component for trivariate retail analysis

Component	Series 448	Series 451	Series 452	Ranks J
Trend	0.00	−0.09	−2.23	{1, 2, 3}
First seasonal	0.00	−0.04	−1.03	{1, 2, 3}
Second seasonal	0.00	−0.08	−0.48	{1, 2, 3}
Third seasonal	0.00	−0.07	−1.59	{1, 2, 3}
Fourth seasonal	0.00	−0.44	−1.65	{1, 2, 3}
Fifth seasonal	0.00	−0.05	−1.32	{1, 2, 3}
Sixth seasonal	0.00	−0.02	−14.60	{1, 2}
Irregular	0.00	−0.10	−0.24	{1, 2, 3}

effects such as outliers, trading day effects, and moving Easter effect (details in Appendix C), and then modeled with an unrestricted model including $I(2)$ trend, irregular, and six atomic seasonals.

There are 6 parameters governing each latent component, and 3 mean parameters, for a total of 51 in the initial nesting model. Focusing on MLEs in our discussion, all estimates (with t statistic significance) are given in Appendix C. (MOM results are also discussed in Appendix C.) The AIC was −1006.64, and residuals indicate an adequate model. The τ_i are directly computed from the MLEs, and are given in Table 1; using the threshold −6.22, the given rank configurations J are readily deduced.

The initial trivariate model can be modified to allow for co-integration at the sixth frequency. The partial correlations for the sixth atomic seasonal are $\kappa_{21} = 0.147$, $\kappa_{31} = 0.99994$, and $\kappa_{32} = 0.99999$. The first two are actually interpretable as straight correlations, whereas the third (κ_{32}) is the correlation for series 451 and 452 conditional on series 448. The straight correlation between 451 and 452 is 0.158. Thus, at the sixth seasonal frequency series 448 and 451 have little relationship, whereas 448 and 452 are highly linked. The co-integration measures are $\tau_2 = -0.02$ and $\tau_3 = -14.6$; the latter value indicates that d_3 should be set to zero to get a nested model.

The nested model with 50 parameters has AIC of −1008.64, which is definitely superior—in fact, the likelihoods (up to numerical precision) are the same for both models, so that the nested model is preferred. Signal extraction estimates for trend as well as the six seasonals (and their aggregate) for the three series are given in Figure 1. In this case, the signal extraction uncertainty was quite low, but there is a subtle shading around each estimate corresponding to a two standard error width confidence interval.

It is possible to consider three sets of bivariate analyses—the pairings (448, 451), (451, 452), and (448, 452)—and the results should be compatible with the trivariate analysis. This is indeed the case, but we only provide summaries here. The first two analyses were run with unconstrained models, and no co-integration (for any of the components) was identified. The third analysis—as expected given the structure of the trivariate covariance matrix for the sixth atomic seasonal—yields a high degree of correlation (0.9999996) between the two series, arguing that we can reduce to a rank one nested model. This is confirmed by the second log partial variance: $\tau_2 = -14.09$. Refitting, the AIC drops from −370.76 to −372.75 with the loss of one parameter; the likelihoods are virtually the same.

In terms of taxonomy, we have 448 $\sim_6$ 452, whereas 448 $\not\sim_6$ 451 and 451 $\not\sim_6$ 452. The co-integrating vector β for the three series, that is, that vector such that its application reduces the nonstationarity by the factor $1 + B$, is given by taking the nesting trivariate model and computing the bottom row of L^{-1}, or in other words

$$\beta' = [L_{32} L_{21} - L_{31}, -L_{32}, 1] = [-0.947, -0.058, 1].$$

If we apply β' to the three series, we should obtain a series that has reduced order of integration. That is, excepting possible deterministic terms in the null space of $1 + B$ (i.e., sequences proportional to $(-1)^t$), the application of $(1 - B)^2(1 + B^2 + \cdots + B^{10})$ (observe that $U(B)$ divided by $1 + B$ equals $1 + B^2 + \cdots + B^{10}$) should reduce the series to stationarity. To the extent that the signal extraction estimates share the co-integration properties of the underlying signals, we can expect that application of the co-integrating vector to the sixth seasonal extraction will be a stationary time series, plus a deterministic function of period 2. This is exactly the case: we computed the application of β to the sixth seasonal extraction, and find the result to be purely deterministic sine wave of period 2; similarly, the co-integrating vector for the third bivariate analysis is $\beta' = [-0.947, 1]$, and again its application to the sixth seasonal extraction is purely deterministic with frequency π.

This confirms that 448 and 452 are in the same species (according to $\sim_6$), and as a result their sixth seasonal extractions are approximately—up to a deterministic function—scale multiples of one another, across all time points. While this has some bearing on the seasonal adjustment (Figure 1), of main interest is the taxonomy of the series. A full taxonomic classification of the retail database would provide insight into how different variables are related, and also indicate which batches of series would be amenable to multivariate seasonal adjustment. This could have possible ramifications to missing data problems, changing sampling frequency, and anticipation of data revisions—these speculations are left unto future work for refinement, but are only mentioned here to provoke interest.

4.2 Modeling Housing Starts

For a second illustration, we consider housing starts data that are published by the U.S. Census Bureau on a monthly basis, for the regions corresponding to South, West, Northeast (NE), and Midwest (MW). We study "New Residential Construction 1964–2012, Housing Units Started, Single Family Units" from the Survey of Construction of the U.S. Census Bureau, available at *http://www.census.gov/construction/nrc/how_the_data_are_collected/soc.html*. We use the co-integration techniques described in the previous subsection rather freely, and will report only the ultimate refined models.

The four series were first cleaned of fixed effects, such as trading day and outliers (Appendix C), and then multivariate structural models involving $I(2)$ trend, irregular, and six atomic seasonals were applied. As mentioned before, results involving a cycle and/or an aggregate seasonal gave poor results, and were

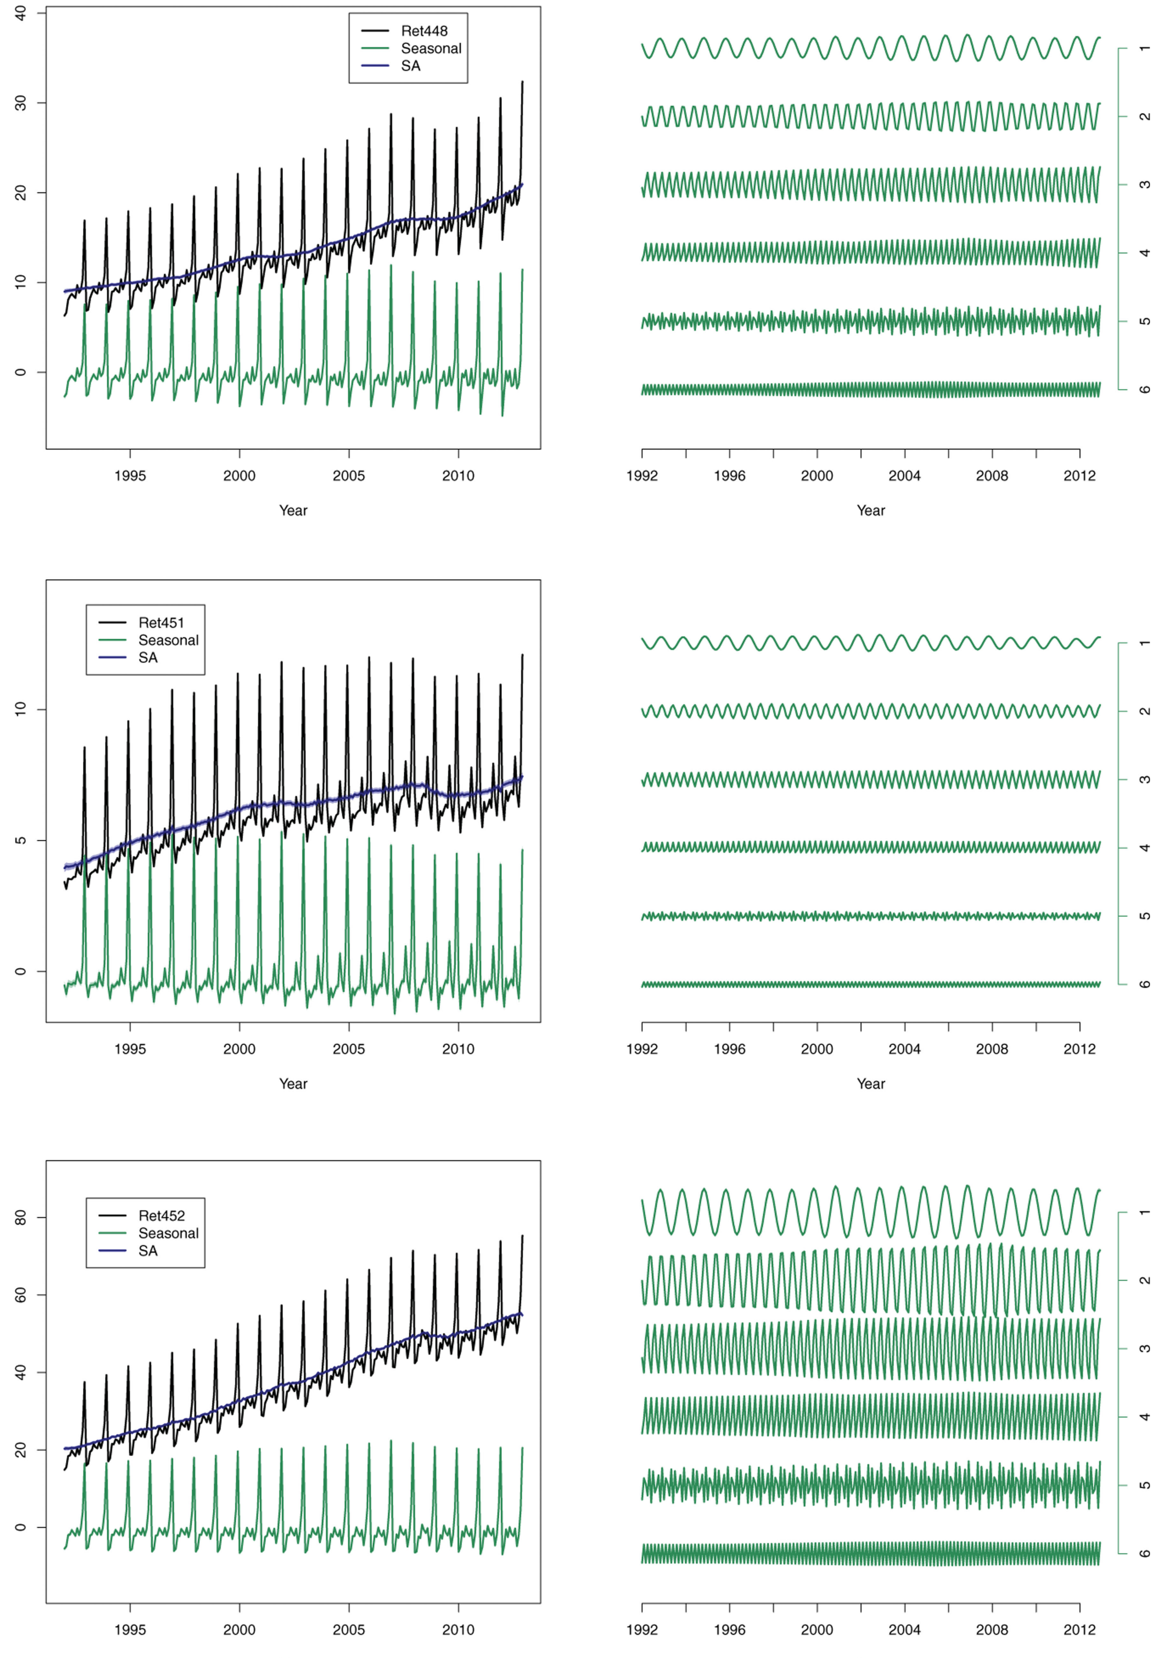

Figure 1. Left panel: retail series (black) for 1992–2012 period, for 448, 451, and 452, with seasonal adjustment (blue) and seasonal (green) extractions based on the best (MLE) fitted trivariate model. Right panel: seasonal extractions for first through sixth seasonal components. Shaded bands correspond to confidence intervals of width given by two standard errors.

Table 2. Log partial variances by component for quadvariate starts analysis

Component	Series				Ranks
	South	West	NE	MW	J
Trend	0.00	−2.76	−2.65	−15.29	{1, 2, 3}
First seasonal	0.00	−1.82	−1.40	−11.86	{1, 2, 3}
Second seasonal	0.00	−0.01	−2.68	−5.21	{1, 2, 3, 4}
Third seasonal	0.00	−10.23	−0.01	−11.09	{1, 3}
Fourth seasonal	0.00	−0.33	−10.18	−9.27	{1, 2}
Fifth seasonal	0.00	−0.08	−0.48	−11.57	{1, 2, 3}
Sixth seasonal	0.00	−0.50	−9.26	−10.38	{1, 2}
Irregular	0.00	−0.15	−0.04	−5.86	{1, 2, 3, 4}

abandoned. The atomic seasonal model is necessary, because the data span covers the GR; the data reflect a change *to trend as well as seasonal patterns*. We have run analyses for the entire span of data, but here focus our discussion on the span 2004–2012, 9 years that include both pre- and post-GR time periods. The sample size is $T = 108$ with $m = 4$ series; the unconstrained nesting model has AIC of 1009.11, with 84 parameters (in each latent component, there are six parameters in L and four parameters in D, plus four mean parameters), whereas the restricted nested model has AIC of 984.11 with 71 parameters. This substantial increase in parsimony is achieved by using a nested model determined from the τ_i computed from the unrestricted MLEs (see Table 2). (The likelihoods for both models were the same, up to numerical error, so the AIC discrepancy is purely attributable to the reduction of 13 parameters.)

As usual, the irregular is enforced to have full rank. In terms of taxonomy, the four series belong to different species, although it is likely that certain pairs can be classified in the same species. We observe that our selection of rank threshold (-6.22) results in quite conservative refinements of the nesting model—the resulting discrepancy in each Σ matrix is so minute, that likelihoods tend to be numerically the same at both the unconstrained and constrained optima. Relaxing this threshold to -3.92 or -1.66 can result in the likelihoods being different, and then the AIC comparison is less trivial (and there is potentially much greater gains in terms of parsimony).

The final residuals from both this nested model, as well as the nesting model, are both adequate, neither set of plots indicating any substantive residual serial correlation. All MLEs for the nested model are given in Appendix C. Signal extraction results, including the aggregate seasonal and the seasonal adjustment along with shaded uncertainty, are given in Figure 2. Multivariate signal extraction has increased precision over univariate approaches, due to the information we can glean from other related series; this precision gain is measured via the relative precision of (7). MSE can be reduced substantially in the case of starts—more than 30% in some cases. (See Appendix C for details.)

We make several observations about the series and the results. The data are initially rising, but by the end of 2006 the decline has begun. Although the low frequency behavior of the series is of substantial interest, we draw attention to the rapidly evolving seasonal pattern. We enunciate here a stylized fact that is well known to the seasonal adjustment community, but is yet to be absorbed by the broader enclave of economists—the seasonal pattern can change rapidly, making the antiquated use of seasonal regressors a dire mistake. The seasonal pattern is fairly stable when focus is restricted to the pre-GR years, but the transition to mid- and post-GR behavior involves a gradual and yet substantial change to the seasonal pattern. The change is on both amplitude and yearly pattern. Standard checks—spectral plots of seasonal extractions and seasonal adjustments, as well as autocorrelation plots—indicate that the resulting adjustment is adequate.

The integrated random walk trend model is evidently flexible enough to accommodate rapid changes in level, and moreover—at least for this sample period—the trend disturbances are so highly cross-correlated as to support a related trend specification (there are three stochastic trends driving the four series). Although Wright (2013) suggested that failure to adequately capture changing trend dynamics will lead to distortions in seasonal adjustment (Greenaway-McGrevy 2013 emphasized multivariate trend modeling to improve the estimation of the seasonal component), in our experience it is more vital to have a sufficiently nuanced seasonal model (such as the atomic seasonal). The aggregate seasonal model (not displayed), in contrast, failed to capture the post-GR shift in seasonal pattern, essentially passing the pre-GR seasonal pattern forward to the GR years; this resulted in salient residual seasonal swings in the irregular component seasonal adjustment in years 2010, 2011, and 2012. Essentially, pre-GR seasonal patterns were imposed on post-GR years. This error can be avoided by using seasonal adjustment filters constructed from a sufficiently flexible seasonal model that captures the highly evolutive seasonality. In short, the extracted trend and seasonal dynamics seem to be suitably separated by our proposed model.

A potential criticism of the results in Figure 2 is that the resulting new seasonal is much more dynamic and swiftly changing than many seasonal adjusters would be comfortable with. However, this directly follows from the model that we used, applied to the highly evolutive period of the GR. This model accommodates regime change in seasonal patterns, and thus offers a gradual and gentle alternative to modeling the GR with ramp regressors or other intervention effects, as some statisticians have pursued (see Buszuwski and Scott (1988) for background, and Ciammola, Cicconi, and Marini (2010) and Maravall and Pérez (2012) for recent efforts).

We proceed to analyze seasonal adjustment sensitivity to model span. We look at eight prior spans of the construction series, each beginning in 1996 and extending through 2005, 2006, 2007, 2008, 2009, 2010, 2011, and 2012, respectively. The same atomic seasonal nesting model was fitted to each, but different nested models were derived in each case—less collinear restrictions tend to be identified when the sample size grows longer. Appendix C contains the resulting seasonal adjustments; overall the pattern of revisions are dramatically reduced by the second full year of new data (i.e., the first annual revision can have some substantial changes, but the later revisions are relatively small). Of course, altering the collinearity thresholds τ_i will alter the results, potentially making the results more sensitive to revision—our own choices have been conservative (less

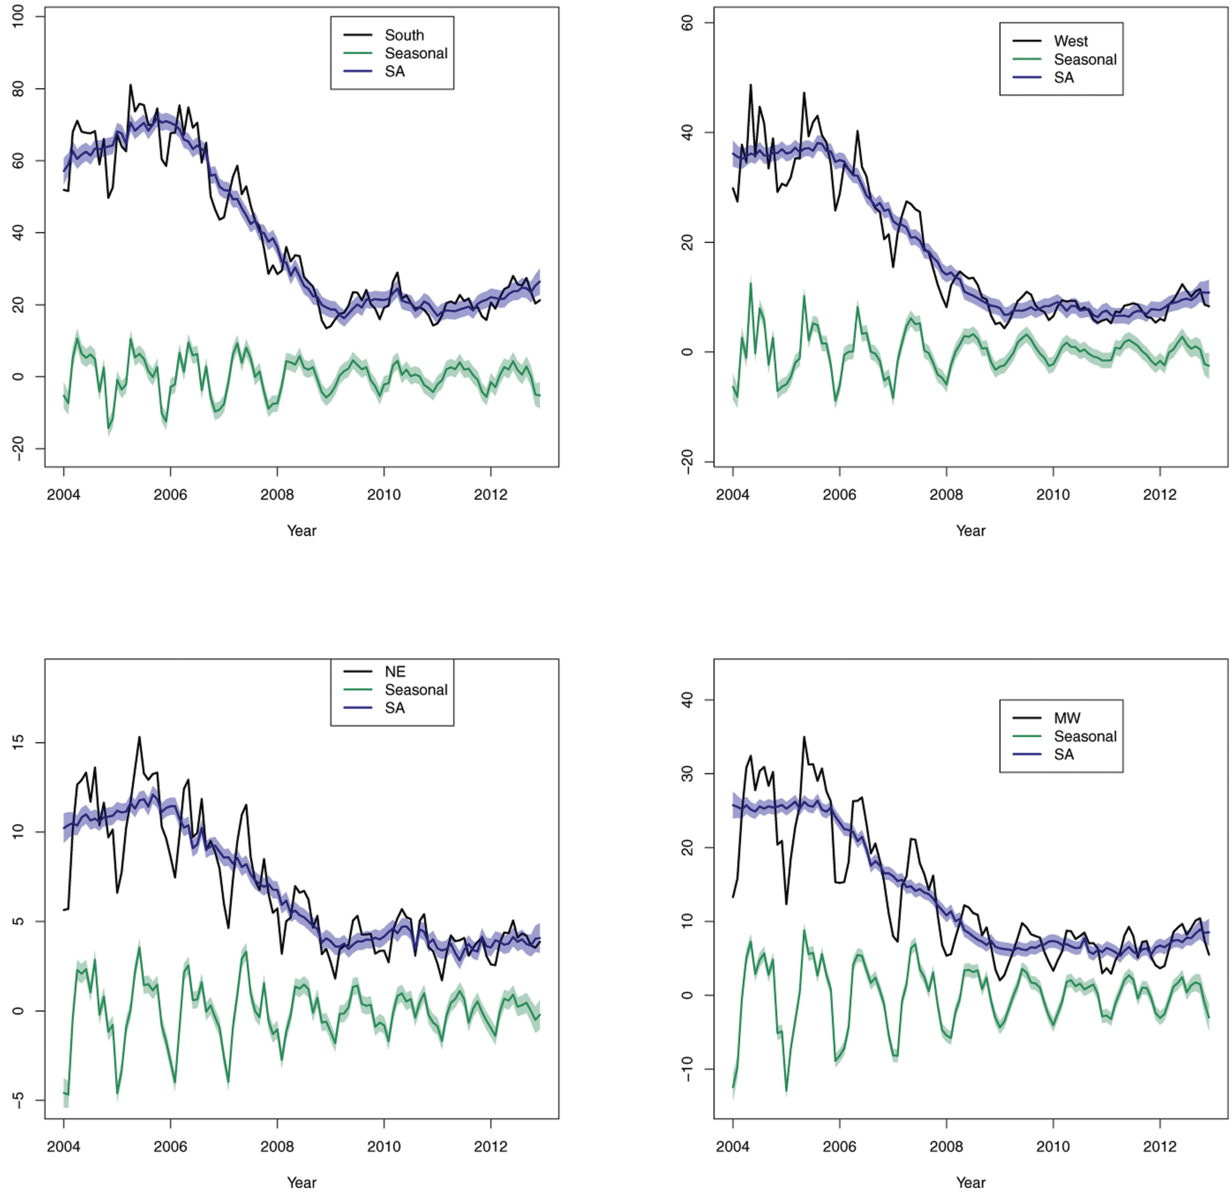

Figure 2. Housing starts (black) for 2004–2012 period, for four regions of the United States (in thousands of housing units), with seasonal adjustment (blue) and seasonal (green) extractions based on the best (MLE) fitted multivariate model. Shaded bands correspond to confidence intervals of width given by two standard errors.

parsimonious than could be attempted empirically), resulting in smaller revisions.

4.3 Regional Aggregation of Housing Starts

It is also of interest to seasonally adjust the Total series, defined to be the straight sum of the four regional series; we wish this to be true of the signal extraction estimates as well. We proceed to illustrate the indirect multivariate seasonal adjustment procedure on the housing starts data. We begin with fitting univariate models (these consist of $I(2)$ trend, irregular, and six atomic seasonals, for nine parameters total) to each of the four series, and display the signal extraction results in Appendix C. Comparing with Figure 2, there is indeed a discrepancy between signal extraction results, although ultimately both seasonal adjustments—multivariate and univariate—are adequate, both having the ability to adapt their extraction filters to the GR regime change. We may then very well ask: what are the advantages of the multivariate method?

One can simply add the univariate extractions to get an indirect extraction for totals, but there will be no quantification of uncertainty—one must either model the regional series jointly, or model the totals directly to get signal extraction uncertainty. This type of "univariate" indirect extraction yields results that are ultimately similar, in this case, to the multivariate extraction for totals shown in the left panel of Figure 3. Uncertainty for the total is obtained by (8). Alternatively, one can model totals with a univariate model and produce a direct adjustment, with signal extraction error quantified, but there is no longer any guarantee that aggregation constraints are respected (they are not); this estimate is displayed in the right panel of Figure 3. Comparing both panels of this figure, we see that the multivariate indirect

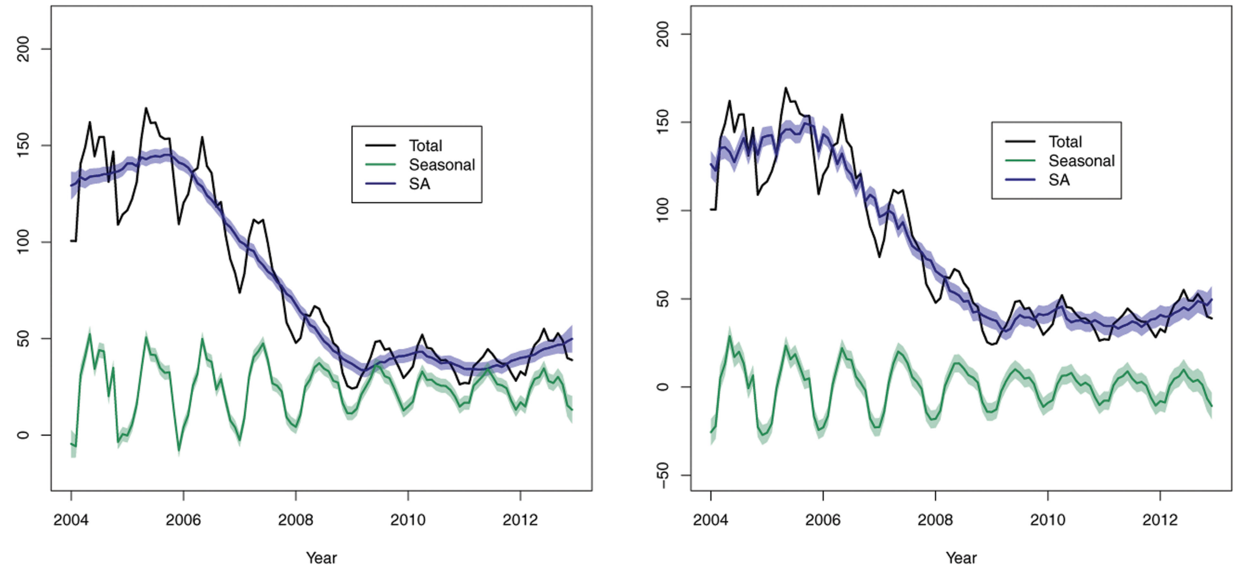

Figure 3. Housing starts (black) for 2004–2012 period, for total starts of the United States (in thousands of housing units), with seasonal adjustment (blue) and seasonal (green) estimates. Shaded bands correspond to confidence intervals of width given by two standard errors. The indirect adjustment (left panel) is constructed from the multivariate model, whereas the direct adjustment (right panel) is constructed from the univariate model fitted to totals.

extractions and the univariate direct extractions are broadly similar, and both seasonal adjustments are indeed adequate, but the MSE for the former (left panel) is less than for the latter (right panel).

These results agree with previous literature (e.g., Geweke 1978). To summarize, the multivariate indirect method seems to be superior to both the univariate indirect method and the univariate direct method: against the former methodology, the advantage is quantification of uncertainty; against the latter methodology, the advantage is the respecting of accounting rules. This discussion has omitted the possibility of raking or other ad hoc reconciliation measures (Quenneville and Fortier 2012), because these nonparametric techniques destroy all possibility of quantifying signal extraction uncertainty.

We mention that some modelers may prefer to take a log transformation of the data, although there are some drawbacks (Section 3.5). Model-fitting results are not reported here, but were similar or slightly simpler than the results given above for no transformation. The signal extraction estimates are extremely similar, although the uncertainty for the four regions is quite a bit lower—this benefit can be weighed against the lack of an uncertainty quantification for the Total.

4.4 Discussion

To summarize these analyses, we have successfully modeled and extracted both trend and seasonal dynamics in retail and construction data. We emphasize that the signals are cleanly separated by the atomic seasonal model, which is vital for the highly evolutive seasonality occurring in housing starts data over the GR. However, such features of adequacy are not the ultimate motivation of this multivariate methodology—because many univariate methods are also capable of generating an adequate separation of trend and seasonal dynamics. The preceding analyses actually demonstrate the method's capacity to address taxonomy, increased precision, and accounting rules.

Sections 3.2 and 3.2 lay out the tools (the log partial variance τ_i) for identifying co-integration, and the trivariate retail analysis of Section 4.1 shows how this is applied to taxonomy. This application is repeated on the construction data (Section 4.2), where we also quantify precision gains over univariate methods (details in Appendix C). Section 4.3 emphasizes how the multivariate indirect method—while yielding an adequate separation of signals—can preserve accounting rules, and still quantify uncertainty.

5. CONCLUSION

This article addresses an important and long-standing question in seasonal adjustment and signal extraction, namely, *is there a benefit to multivariate techniques*? Our proposals herein rely on available tools of multivariate time series analysis (encoded in *sigex*). We have motivated these models as latent dynamic factor models that expand the basic dynamic factor model (1) in a manner that takes account of time series structure, associating additional latent dynamic factors with frequencies of interest in the process' spectral density. The factor loadings of each latent factor are then naturally associated with the lower Cholesky factors of the respective innovations' covariance matrix. Each entry of these lower Cholesky factors is interpretable as a scaled partial covariance, and therefore gives some information about how the respective time series are related to one another at trend or seasonal frequencies.

A pleasing facet of these models is their ready interpretability. The reduced rank in a latent process' innovation covariance matrix, corresponding to collinear innovations, is easily modeled, and moreover can be interpreted as frequencies of noninvertibility for the differenced process' spectral density. This in turn implies a co-integration interpretation for the undifferenced series, in a generalized sense; the basis of the co-integrating spaces are obtained at once by taking the appropriate rows of the inverse of the lower Cholesky factors. It has been shown that application of these co-integrating vectors reduces the order of nonstationarity of the original process, by exactly eliminating the need for the differencing operator corresponding to that particular latent process. We elaborate this interpretation with a trivariate retail analysis.

Having identified the ranks of each latent component, we can then contemplate taxonomy of economic data, because co-integration for such processes is the same thing as full spectral coherency among series at the respective frequency. We propose that fully reduced rank, of unity, be used as the definition of species, and establish some preliminary results for taxonomic classification. A key empirical facet is the ability to determine the actual rank of each latent process' innovation covariance matrix, and we describe some new tools involving partial correlations to define condition numbers, which enable one to tease out viable reduced rank models. We illustrate this procedure on the three retail series, showing how these partial correlation measures do indeed indicate which series are redundant for a particular latent component (the sixth atomic seasonal in our illustration). Once the full and restricted models have been fitted, an AIC comparison can be used to decide between competitors; other parameters (the entries of the lower Cholesky factors, or factor loading matrices) can also be zeroed out if warranted by a likelihood ratio test. A central aspect of this methodology is the ability to compute likelihoods and signal extraction results with relative ease.

One benefit of the multivariate signal extraction methodology is increased precision, as demonstrated through the precision comparisons on the housing starts data. Another benefit is the improvement of the indirect method of seasonal adjustment for preservation of economic identities. The direct method—running a univariate methodology on the totals—fails to preserve aggregation relations, while univariate indirect methods (summing the individual adjustments) will not take into account cross-series correlation, and will not allow for quantification of the aggregate series' signal extraction uncertainty. The multivariate indirect method addresses both of these latter issues, while preserving economic identities.

Given the benefits in terms of interpretability, taxonomic classification, and preservation of economic identities, what are the demerits of the multivariate methodology? We explored signal extraction revisions for construction series, demonstrating that the models are able to adapt to pre- and post-GR phenomena, indeed having the flexibility to accommodate rapidly changing seasonal patterns. This accommodation resulted in moderate revisions, which is no surprise given that the new information radically altered prior understandings of trend and seasonal patterns. Overall, the revisions behavior seemed satisfactory, although we noted that the actual models identified (the particular co-integrating ranks) can change dramatically as the data span is altered.

In our own opinion, the chief criticism is in the additional time required of the analyst to perform the modeling task; second, and related, is the huge number of parameters involved when $m = 4$ or higher. Third, the use of log transformations interferes with our proposed method of handling the preservation of economic identities—but only if quantification of uncertainty is requisite. Regarding the first two points, to achieve parsimony and a feasible computation time, one is naturally led to seeking co-integrating relationships and other reductions of the parameter space. The chief bottle-neck is not likelihood evaluation (even for $m = 10$, this is rapid), but optimization over high-dimensional parameter manifolds.

Two promising avenues for further research are the method-of-moments estimator proposed herein (Appendix B), and potential further parameter dimension reduction achieved through informative metadata—for example, if series are geographically connected, a spatial covariance function could possibly be used to govern the values of each Σ.

SUPPLEMENTARY MATERIALS

Appendices: Appendix A (Proofs and Derivations), Appendix B (Method of Moments), Appendix C (Data Analysis). (pdf)

Data: Processed retail and construction data. (zip file multiSAdata)

R-code: Script and code for analyses used in the article. (zip files Scripts and sigex)

ACKNOWLEDGMENTS

The author has benefitted from discussions with Bill Bell, Tom Evans, David Findley, Andrew Harvey, Siem Jan Koopman, Kathleen McDonald-Johnson, Brian Monsell, Anindya Roy, Thomas Trimbur, and Jonathan Wright. The work has been presented at the Bureau of Labor Statistics and the Reserve Bank of India. The primary impetus for this work was a comment of Dale Jorgensen, who encouraged the Census Bureau to devote research resources to the topic of multivariate seasonal adjustment. This report is released to inform interested parties of research and to encourage discussion. The views expressed on statistical issues are those of the author and not necessarily those of the U.S. Census Bureau.

[Received February 2015. Revised September 2015.]

REFERENCES

Alexander, L., and Greenberg, J. (2012), "Echo of Financial Crisis Heard in Recent Jobless Claims Drop," available at *www.nomura.com/research*. [611]

Alexander, L., Zentner, E., and Greenberg, J. (2012), "Seasonal Bias in Sentiment Indicators: Reduced, But Not Eliminated," available at *www.nomura.com/research*. [611]

Bartelsman, E., and Cleveland, W. (1993), *Joint Seasonal Adjustment of Economic Time Series (Finance and Economics Discussion Series)*, Washington, DC: Federal Reserve Board. [611]

Bell, W. (1984), "Signal Extraction for Nonstationary Time Series," *The Annals of Statistics*, 12, 646–664. [614]

Bell, W., and Hillmer, S. (1984), "Issues Involved With the Seasonal Adjustment of Economic Time Series," *Journal of Business and Economics Statistics*, 2, 291–320. [611]

Birrell, C. L., Steel, D. G., and Lin, Y. (2011), "Seasonal Adjustment of An Aggregate Series Using Univariate and Multivariate Basic Structural Models," *Journal of Statistical Theory and Practice*, 5, 179–205. [611,617]

Bozik, J., and Bell, W. (1987), "Forecasting Age Specific Fertility Using Principal Components," Research Report RR-87/19, Bureau of the Census, available at *http://www.census.gov/srd/papers/pdf/rr87-19.pdf*. [612]

Brockwell, P., and Davis, R. (1991), *Time Series: Theory and Methods*, New York: Springer. [614]

Busetti, F. (2006), "Tests of Seasonal Integration and Cointegration in Multivariate Unobserved Component Models," *Journal of Applied Econometrics*, 21, 419–438. [615]

Buszuwski, J. A., and Scott, S. (1988), "On the Use of Intervention Analysis in Seasonal Adjustment," in *Proceedings of the ASA Section on Business and Economic Statistics*, pp. 337–342. [621]

Buys Ballot, C. (1847), *Les Changements Pérodiques de Tempèrature*, Utrecht: Kemink et Fils. [611]

Ciammola, A., Cicconi, C., and Marini, M. (2010), *Seasonal Adjustment and the Statistics Treatment of the Economic Crisis: An Application to Some Italian Time Series*, Italy: National Institute of Statistics. [621]

Durbin, J., and Koopman, S. (2001), *Time Series Analysis by State Space Methods*, Oxford: Oxford University Press. [613]

Engle, R., and Granger, C. (1987), "Cointegration and Error Correction: Representation, Estimation, and Testing," *Econometrica*, 55, 251–276. [615]

Feroli, M. (2012), "Another Springtime Rope-a-Dope?" J.P. Morgan. North America Economic Research. [611]

Geweke, J. (1978), "The Temporal and Sectoral Aggregation of Seasonally Adjusted Time Series," in *Seasonal Analysis of Economic Time Series*, ed. A. Zellner, Department of Commerce, Bureau of the Census, pp. 411–427. [611,613,617,623]

Ghysels, E. (1997), "Seasonal Adjustment and Other Data Transformations," *Journal of Business and Economic Statistics*, 15, 410–418. [617]

Giannone, D., Reichlin, L., and Sala, L. (2005), "Monetary policy in real time," in *NBER Macroeconomics Annual 2004, Volume 19*, Cambridge, MA: MIT Press, pp. 161–224. [612]

Gilbart, J. (1852), "On the Laws of the Currency of Ireland, as Exemplified in the Changes That Have Taken Place in the Amount of Bank Notes in Circulation in Ireland, Since the Passing of the Acts of 1845," *Journal of the Statistical Society of London*, 15, 307–325. [611]

Granger, C. (1966), "The Typical Spectral Shape of an Economic Variable," *Econometrica*, 34, 150–161. [613]

Greenaway-McGrevy, R. (2013), *A Multivariate Approach to Seasonal Adjustment*, Washington DC: Bureau of Economic Analysis. [611,621]

Harvey, A. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge: Cambridge University Press. [613]

Koopman, S. J., Harvey, A. C., Doornik, J. A., and Shephard, N. (2009), *STAMP 8 Structural Time Series Analyser, Modeler, and Predictor*, London: Timberlake Consultants. [611]

Koopman, S. J., Ooms, M., and Hindrayanto, I. (2012), "A Multivariate Periodic Unobserved Components Time Series Analysis for Sectoral U.S. Employment," in *Economic Time Series: Modeling and Seasonality*, eds. W. Bell, S. Holan, and T. McElroy, New York: Chapman and Hall, pp. 3–35. [611,612,613]

Krane, S., and Wascher, W. (1999), "The Cyclical Sensitivity of Seasonality in U.S. Unemployment," *Journal of Monetary Economics*, 44, 523–553. [611]

Maravall, A., and Pérez, D. (2012), "Applying and Interpreting Model-Based Seasonal Adjustment The Euro-Area Industrial Production Series," in *Economic Time Series: Modeling and Seasonality*, eds. William R. Bell, Scott H. Holan, and Tucker S. McElroy, Boca Raton, FL: CRC Press-Chapman and Hall, pp. 281–313. [621]

McElroy, T., and Trimbur, T. (2015), "Signal Extraction for Non-Stationary Multivariate Time Series With Illustrations for Trend Inflation," *Journal of Time Series Analysis*, 36, 209–227. Also: Finance and Economics Discussion Series, Working Paper No. 2012-45, Federal Reserve Board. Available at *http://www.federalreserve.gov/pubs/feds/2012/201245/201245abs.html*. [614,615,616,617,618]

Nerlove, M. (1964), "Spectral Analysis of Seasonal Adjustment Procedures," *Econometrica*, 32, 241–286. [611]

Nerlove, M., Grether, D., and Carvalho, J. (1979), *Analysis of Economic Time Series: A Synthesis*, New York: Academic Press. [611]

Nyblom, J., and Harvey, A. (2000), "Tests of Common Stochastic Trends," *Econometric Theory*, 16, 176–199. [615]

——— (2001), "Testing Against Smooth Stochastic Trends," *Journal of Applied Econometrics*, 16, 415–429. [615]

Persons, W. (1919), "Indices of Business Conditions," *Review of Economics and Statistics*, 1, 5–107. [611]

Pinheiro, J., and Bates, D. (1996), "Unconstrained Parametrizations for Variance-Covariance Matrices," *Statistics and Computing*, 6, 289–296. [614]

Planas, C., and Campolongo, F. (2001), "The Seasonal Adjustment of Contempo-Raneously Aggregated Series," Working Paper 8221011/1-Lot3, EUROSTAT, Luxembourg. [613]

Quenneville, B., and Fortier, S. (2012), "Restoring Accounting Constraints in Time Series–Methods and Software for a Statistical Agency," in *Economic Time Series: Modeling and Seasonality*, eds. W. Bell, S. Holan, and T. McElroy, New York: Chapman and Hall, pp. 231–253. [612,623]

Sargent, T. J., and Sims, C. A. (1977), "Business Cycle Modeling Without Pretending to Have Too Much A Priori Economic Theory," *New Methods in Business Cycle Research*, 1, 145–168. [612]

Stock, J., and Watson, M. (1988), "Testing for Common Trends," *Journal of the American Statistical Association*, 83, 1097–1107. [615]

——— (2011), "Dynamic Factor Models," in *Oxford Handbook on Economic Forecasting*, eds. M. J. Clements and D. F. Hendry, Oxford: Oxford University Press, pp. 35–59. [611,612]

——— (2012), "Disentangling the Channels of the 2007–2009 Recession," Prepared for the Brookings Panel on Economic Activity, March 22–23. [612]

Tiller, R. (2012), "Frequency Domain Analysis of Seasonal Adjustment Filters Applied to Periodic Labor Force Survey Series," in *Economic Time Series: Modeling and Seasonality*, eds. W. Bell, S. Holan, and T. McElroy, New York: Chapman and Hall, pp. 135–158. [612]

Wright, J. (2013), "Unseasonal Seasonals?" Fall 2013 Brookings Panel on Economic Activity, September 19–20. [612,621]

Zentner, E., Amemiya, A., and Greenberg, J. (2011), "Stronger Data Ahead: Explanation and Implication," available at *www.nomura.com/research*. [611]