BUREAU OF THE CENSUS
STATISTICAL RESEARCH DIVISION
Statistical Research Report Series
No. 92/14

Non-Gaussian Season Adjustment: X-12 ARIMA Versus Robust Structural Models

by
Andrew G. Bruce & Simon R. Jurke
Statistical Sciences, Inc.
Seattle, WA 98109

Report Issued: 12/10/92

Non-Gaussian Seasonal Adjustment: X-12-ARIMA Versus Robust Structural Models

Andrew G. Bruce
and
Simon R. Jurke

November 16, 1992

This work was started while Dr. Bruce was a visiting research statistician at the U.S. Bureau of the Census. It was completed through the support of a Joint Statistical Agreement between the U.S. Bureau of the Census and Victoria University of Wellington.

The authors are indebted to David Findley (U.S. Bureau of the Census), who inspired, supported and guided this project. Peter Thomson (Victoria University) also deserves special thanks for his very substantive contributions. William Bell, Brian Monsell, and Mark Otto (U.S. Bureau of the Census) gave invaluable support, providing the data, X-12-ARIMA software, REGCMPNT software, and the choice of ARIMA models for X-12-ARIMA. In addition, Genshiro Kitagawa provided software upon which the MING program is based. Finally, Alistair Gray (New Zealand Department of Statistics), Magdalena Cordera, and Jim Durbin contributed many insightful comments.

Non-Gaussian Seasonal Adjustment: X-12-ARIMA Versus Robust Structural Models

Andrew G. Bruce
and
Simon R. Jurke

November 16, 1992

Statistical Sciences, Inc.,
1700 Westlake Ave. N,
Seattle, WA 98109

Non-Gaussian Seasonal Adjustment: X-12-ARIMA Versus Robust Structural Models

Andrew G. Bruce and Simon R. Jurke

November 15, 1992

Statistical Sciences, Inc., 1700 Westlake Ave. N, Seattle, WA 98109

Abstract

This study compares X-12-ARIMA and MING, two new seasonal adjustment methods designed to handle outliers and structural changes in a time series. X-12-ARIMA is a successor to the X-11-ARIMA seasonal adjustment method, and is being developed at the U.S. Bureau of the Census (Findley et al. (1988)). MING is a "MIxture based Non-Gaussian" method for seasonal adjustment using time series structural models. It was developed for this study based on methodology proposed by Kitagawa (1990).

The procedures are compared using 29 macroeconomic time series from the U.S. Bureau of the Census. These series have both outliers and structural changes, providing a good testbed for comparing non-Gaussian methods. For the 29 series, the X- 12-ARIMA decomposition consistently leads to smoother seasonal factors which are as or more "flexible" than the MING seasonal component. On the other hand, MING is more stable, particularly in the way it handles outliers and level shifts.

This study relied heavily on graphical tools for comparing seasonal adjustment methods. Use of graphics is critical in forming the conclusions of this paper.

This work was started while Dr. Bruce was a visiting research statistician at the U.S. Bureau of the Census. It was completed through the support of a Joint Statistical Agreement between the U.S. Bureau of the Census and Victoria University of Wellington.

The authors are indebted to David Findley (U.S. Bureau of the Census), who inspired, supported and guided this project. Peter Thomson (Victoria University) also deserves special thanks for his very substantive contributions. William Bell, Brian Monsell, and Mark Otto (U.S. Bureau of the Census) gave invaluable support, providing the data, X-12-ARIMA software, REGCMPNT software, and the choice of ARIMA models for X-12-ARIMA. In addition, Genshiro Kitagawa provided software upon which the MING program is based. Finally, Alistair Gray (New Zealand Department of Statistics), Magdalena Cordera, and Jim Durbin contributed many insightful comments.

Abstract

This study compares X-12-ARIMA and MING, two new seasonal adjustment methods designed to handle outliers and structural changes in a time series. X-12-ARIMA is a successor to the X-11-ARIMA seasonal adjustment method, and is being developed at the U.S. Bureau of the Census (Findley et al. (1988)). MING is a "MIxture based Non-Gaussian" method for seasonal adjustment using time series structural models. It was developed for this study based on methodology proposed by Kitagawa (1990).

The procedures are compared using 29 macroeconomic time series from the U.S. Bureau of the Census. These series have both outliers and structural changes, providing a good testbed for comparing non-Gaussian methods. For the 29 series, the X- 12-ARIMA decomposition consistently leads to smoother seasonal factors which are as or more "flexible" than the MING seasonal component. On the other hand, MING is more stable, particularly in the way it handles outliers and level shifts.

This study relied heavily on graphical tools for comparing seasonal adjustment methods. Use of graphics is critical in forming the conclusions of this paper.

Keywords: Seasonal adjustment, X-12-ARIMA, non-gaussian time series, structural models.

Non-Gaussian Seasonal Adjustment: X-12-ARIMA Versus Robust Structural Models

Andrew G. Bruce and Simon R. Jurke

November 15, 1992

Statistical Sciences, Inc., 1700 Westlake Ave. N, Seattle, WA 98109

Abstract

This study compares X-12-ARIMA and MING, two new seasonal adjustment methods designed to handle outliers and structural changes in a time series. X-12-ARIMA is a successor to the X-11-ARIMA seasonal adjustment method, and is being developed at the U.S. Bureau of the Census (Findley et al. (1988)). MING is a "MIxture based Non-Gaussian" method for seasonal adjustment using time series structural models. It was developed for this study based on methodology proposed by Kitagawa (1990).

The procedures are compared using 29 macroeconomic time series from the U.S. Bureau of the Census. These series have both outliers and structural changes, providing a good testbed for comparing non-Gaussian methods. For the 29 series, the X- 12-ARIMA decomposition consistently leads to smoother seasonal factors which are as or more "flexible" than the MING seasonal component. On the other hand, MING is more stable, particularly in the way it handles outliers and level shifts.

This study relied heavily on graphical tools for comparing seasonal adjustment methods. Use of graphics is critical in forming the conclusions of this paper.

This work was started while Dr. Bruce was a visiting research statistician at the U.S. Bureau of the Census. It was completed through the support of a Joint Statistical Agreement between the U.S. Bureau of the Census and Victoria University of Wellington.

The authors are indebted to David Findley (U.S. Bureau of the Census), who inspired, supported and guided this project. Peter Thomson (Victoria University) also deserves special thanks for his very substantive contributions. William Bell, Brian Monsell, and Mark Otto (U.S. Bureau of the Census) gave invaluable support, providing the data, X-12-ARIMA software, REGCMPNT software, and the choice of ARIMA models for X-12-ARIMA. In addition, Genshiro Kitagawa provided software upon which the MING program is based. Finally, Alistair Gray (New Zealand Department of Statistics), Magdalena Cordera, and Jim Durbin contributed many insightful comments.

1 Introduction

Seasonal adjustment at most statistical agencies is currently done using a procedure based on the X-11-ARIMA method for seasonal adjustment. While X-11-ARIMA has proven to be reliable and effective, it is primarily an *ad hoc* method. Researchers have explored *model* based alternatives to X-11-ARIMA. Gaussian "time series structural models," introduced by Gersch and Kitagawa (1983) and Harvey and Todd (1983), are a class of models currently enjoying a surge of interest. Time series structural models are based on using simple intuitive component models for the trend, seasonal and irregular. In a comparison of X-11 and structural models, den Butter and Mourik (1990) and Jain (1992) conclude that structural models are a competitive method.

Outliers and "structural" changes (e.g., level shifts or ramps) cause problems with both X-11-ARIMA and the Gaussian structural models. While X-11-ARIMA provides some protection against outliers, it is not fully robust and cannot handle level shifts or other structural changes. The X-12-ARIMA procedure, a successor to X-11-ARIMA, is being developed at the U.S. Bureau of the Census to handle additive outliers and level shifts (Findley et al. (1988), Monsell (1990)).

Time series structural models can be adapted to non-Gaussian situations by assuming that the innovations of the component models are non-Gaussian (Kitagawa (1990)). For this study, we have developed "MING", which extends the structural model based seasonal adjustment to handle outliers and structural changes. MING is derived from a computer program by Kitagawa (1991).

X-12-ARIMA and MING differ in very significant ways. X-12-ARIMA is a non-parametric method while MING is model based. The X-12-ARIMA seasonal filters are manually selected on the basis of diagnostic plots. By contrast, the seasonal decomposition of MING is automatically obtained by maximizing the likelihood. Finally, the two procedures adopt very different methods for handling outliers and level shifts.

The X-12-ARIMA procedure is compared with MING using 29 macroeconomic time series from the U.S. Bureau of the Census. These series were chosen as a basis for comparison since each series contains both outliers and structural changes. The three main conclusions of our study are:

1. For the 29 series, the X-12-ARIMA decomposition consistently leads to seasonal factors which "smoother" and more "flexible" than the MING seasonal component. This leads us to conclude that the X-12-ARIMA seasonal adjustments are generally more appealing for these series.

2. According to sliding span statistics Findley et al. (1990), the MING procedure for handling outliers and level shifts leads to more stable seasonal adjustments

than X-12-ARIMA. The lack of stability in X-12-ARIMA can be partly attributed to the discontinuous nature of its outlier/level shift detection scheme.

3. Simple diagnostics are not adequate for comparing seasonal adjustment methods. Graphical tools are essential for effective assessment.

Section 2 discusses the outlier handling scheme for X-12-ARIMA. The MING procedure is described in section 3. Section 4 describes the data and the associated filters and models. The heart of the paper lies in sections 5, 6, and 7 which discuss the three main conclusions listed above. Other conclusions are given in section 8 and the MING method is explored further in section 9. Finally, directions for future research are discussed in section 10.

This paper is based on a fuller report by Bruce and Jurke (1992a), which we will refer to as [BJ92a]. Associated with the report is a "book" of plots (Bruce and Jurke (1992b)).

2 X-12-ARIMA

The X-11 method for seasonal adjustment was developed at the U.S. Bureau of the Census by Shiskin et al. (1967). X-11-ARIMA is an extension of the X-11 method, developed at Statistics Canada by Dagum (1980). X-11-ARIMA eliminates the asymmetric filters of X-11 by using ARIMA models to forecast beyond the ends of the series. Both X-11 and X-11-ARIMA are nonparametric procedures, with a design based on practical considerations.

Numerous empirical studies have examined the X-11-ARIMA method: see, for example, Dagum (1978), Dagum and Morry (1984), den Butter et al. (1985), and Jain (1989). X-12-ARIMA offers several new features, including a new "language oriented" interface and the "sliding spans" diagnostics (Findley et al. (1990)). The primary new feature of interest in this study is the procedure for automatic detection of additive outliers and level shifts. This procedure is discussed in more detail below.

2.1 X-12-ARIMA outlier/level shift identification procedure

To avoid problems caused by additive outliers (AO's) and level shifts (LS's), X-12-ARIMA does a prior adjustment. AO's and LS's are identified using hypothesis tests based on the appropriate parametric intervention and ARIMA model. The series is adjusted using the estimated interventions.

The idea of doing hypothesis tests to identify the type of outlier was first introduced by Fox (1972). Suppose X_t is a time series which behaves according to the multiplicative Gaussian ARIMA $(p,d,q)\times(P,D,Q)_S$ model. Let Y_t be the observed series, which is related to X_t by

$$Y_t = X_t + \sum_{j=1}^{k} \zeta_j Z_t^{(j)}$$

Y_t might contain outliers, level shifts, etc., which are modeled by the 0-1 processes $Z_t^{(j)}$ and the parameters ζ_j. To model an "additive outlier" (AO) at time T, we set

$$Z_t^{(j)} = \begin{cases} 1 & t = T \\ 0 & t \neq T \end{cases}$$

A "level shift" (LS) at time T is given by

$$Z_t^{(j)} = \begin{cases} 1 & t \geq T \\ 0 & t < T \end{cases}$$

A hypothesis test for the presence of an AO (or LS) at time T takes the form

$$\begin{aligned} \mathrm{H}_0: &\quad \zeta_j = 0 \\ \mathrm{H}_1: &\quad \zeta_j \neq 0 \end{aligned}$$

A large test statistic is indicative of an AO (or LS). These ideas generalize to other types of interventions, such as innovations outliers, ramps, or variance changes.

X-12-ARIMA incorporates tests for AO's and LS's in an iterative method for estimating parameters in a multiplicative seasonal ARIMA model. Suppose we have an initial estimate of the ARIMA parameters $\hat{\alpha}_0$. The algorithm proceeds as follows:

Step 0: $j = 0$.

Forward Addition

Step 1: Given $\hat{\alpha}_j$, compute the t-statistics $\hat{\tau}_t^{\mathrm{AO}}$ and $\hat{\tau}_t^{\mathrm{LS}}$ corresponding to the hypothesis tests for AO's and LS's at times $t = 1,2,\ldots,N$.

Step 2: If

$$\max_t \left\{ \left|\hat{\tau}_t^{\mathrm{AO}}\right|, \left|\hat{\tau}_t^{\mathrm{LS}}\right| \right\} < C,$$

then go to step 5. Otherwise, flag the observation which is the most significant AO or LS according to the t-statistics.

Step 3: Subtract the least squares estimate $\hat{\zeta}_j$ of the flagged intervention from the series Y_t. Re-estimate the parameters $\hat{\alpha}_{j+1}$ with the adjusted data.

Step 4: $j = j + 1$. Go to step 1.

Backward Elimination

Step 5: Let Ω be the set of indices corresponding to the identified AO's and LS's. Re-estimate the t-statistics for all identified AO's and LS's.

Step 6: If

$$\min_{t \in \Omega} \{|\hat{\tau}_t|\} > C,$$

then we are done. Otherwise, drop the least significant estimated intervention from the index set Ω and go to step 5.

For step 1, an efficient algorithm is available for ARIMA models, reducing the computational burden of computing the test statistic for all observations simultaneously. In step 2, the cutoff C is used to determine if there are any more significant AO's or LS's remaining in the series. In this study, a cutoff of $C = 3.1$ is used.

Note that only the most significant AO or LS is identified on each iteration. This "one-at-a-time" approach is computationally slower than identifying all significant AO's and LS's on each pass. However, for several of the series examined in this study, the multiple identification procedure is unstable and leads to poor decompositions. Hence, the multiple identification option is not recommended for general use.

Other types of interventions could be incorporated into the procedure. However, for economic time series, the most important and natural situations to attempt to model in this manner seem to be additive outliers and level shifts. This iterative identification procedure was first developed by Chang and Tiao (1983). Hillmer et al. (1983) applied this iterative estimation method in the context of ARIMA model based seasonal adjustment. See Bell (1986), Chang et al. (1988), and Tsay (1988) for further development of the method.

3 MING

Many model based approaches to seasonal adjustment have been proposed. Advantages of model based seasonal adjustment are articulated by Bell and Hillmer (1984). Models provide an interpretable decomposition whose characteristics adapt to the nature of each series. One approach towards model based seasonal adjustment is based on fitting ARIMA models: see, for example, Box et al. (1978). The ARIMA model is decomposed into trend, seasonal and irregular components, maximizing the variance of the irregular. This is often called the canonical decomposition.

In this study, we work with an approach based on time series structural models. Gersch and Kitagawa (1983) and Harvey (1984) have explored the use of structural models for seasonal adjustment (see also Kitagawa and Gersch (1984), Harvey

(1989), den Butter and Mourik (1990) and Jain (1992)). There are several advantages of structural model-based seasonal adjustment. Structural models are constructed by using simple component models for the trend, seasonal, and irregular. Harvey (1989) and Harvey and Valls Pereira (1989) argue that structural models are more interpretable yield superior seasonal decompositions and adjustments than the canonical decomposition of ARIMA models. The fitting process is simpler for structural models, with only one or two basic model forms needed for a broad range of series (of the three models considered in this study, one model was consistently superior). Finally, structural models easily and naturally incorporate simple structural changes, such as level shifts and ramps.

3.1 Gaussian Time Series Structural Models

A time series structural model is based on forming models directly for each of the components in the decomposition

$$Y_t = T_t + S_t + I_t. \tag{1}$$

Y_t is some suitable transformation of the original observed series (in this study, Y_t is the log-transformed data). T_t, S_t, and I_t are the trend, seasonal and irregular. The irregular is usually considered to be Gaussian white noise with zero mean and variance σ_I^2. This is denoted by $I_t \sim \mathrm{GWN}(0, \sigma_I^2)$.

A typical model for the trend is given by

$$T_t = T_{t-1} + b_{t-1} + \eta_t \tag{2}$$

where $\eta_t \sim \mathrm{GWN}(0, \sigma_\eta^2)$. The term b_t acts as a "slope", and is permitted to evolve according to a random walk

$$b_t = b_{t-1} + \xi_t$$

where $\xi_t \sim \mathrm{GWN}(0, \sigma_\xi^2)$.

Following Harvey (1984), we consider two different models for the seasonal component. Let s be the seasonal period (for monthly data $s = 12$). The first seasonal model, which makes up part of Harvey's "Basic Structural Model" (BSM), is defined by

$$S_t = -\sum_{j=1}^{s-1} S_{t-j} + \omega_t \tag{3}$$

where $\omega_t \sim \mathrm{GWN}(0, \sigma_\omega^2)$. An alternative and more flexible seasonal model is given by

$$S_t = \sum_{j=1}^{[s/2]} \gamma_{j,t} \tag{4}$$

with

$$\begin{aligned}\gamma_{j,t} &= \gamma_{j,t-1} cos\lambda_j + \gamma^*_{j,t-1} sin\lambda_j + \omega_{j,t} \\ \gamma^*_{j,t} &= -\gamma_{j,t-1} sin\lambda_j + \gamma^*_{j,t-1} cos\lambda_j + \omega^*_{j,t}\end{aligned}$$

where $\lambda_j = 2\pi j/s$. The innovations $\omega_{j,t}$ and $\omega^*_{j,t}$ are uncorrelated with $\omega_{j,t} \sim \mathrm{GWN}(0, \sigma^2_{\omega,j})$ and $\omega^*_{j,t} \sim \mathrm{GWN}(0, \sigma^2_{\omega,j})$ for $j = 1, \ldots, 6$.

Modeling Calendar Effects

Structural models are easily extended to handled such things as calendar and holiday effects. Instead of (1), we might use the model

$$Y_t = T_t + C_t + S_t + I_t. \tag{5}$$

where C_t represents the calendar effect. C_t is estimated as a fixed effect by construction of the appropriate regression variables: see Bell and Hillmer (1983). Dynamic models for trading days also fit nicely within the structural model framework: see Monsell (1983) and Dagum et al. (1988).

3.2 Robustness through Gaussian Mixtures Models

One of the strengths of the structural model is the simplicity and interpretability of the component models. This is illustrated when we consider non-Gaussian extensions to the trend model. It can readily be seen that outliers in I_t, η_t, and ξ_t translate directly into additive outliers, local level shifts, and ramps respectively. Hence, these types of events can be accommodated for in the model by assuming that I_t, η_t, and ξ_t are generated from an appropriate outlier producing distribution.

Gaussian mixture distributions are one way to model outliers. For example, to generate additive outliers, we assume that

$$I_t \sim \begin{cases} \mathrm{N}(0, \sigma^2_{\mathrm{I}}) & \text{with probability } 1 - \epsilon_{\mathrm{I}} \\ \mathrm{N}(0, \tilde{\sigma}^2_{\mathrm{I}}) & \text{with probability } \epsilon_{\mathrm{I}} \end{cases} \tag{6}$$

where $\tilde{\sigma}^2_{\mathrm{I}} \gg \sigma^2_{\mathrm{I}}$. The factor ϵ_{I} represents the "prior" probability of an additive outlier.

To avoid too many mixture terms with small probabilities, it can be assumed that only one type of structural change can occur at a given time point. In other words, we shall assume that either a level shift or a ramp can occur, but not both.

Hence, the joint distribution of η_t and ξ_t is given by

$$\begin{pmatrix} \eta_t \\ \xi_t \end{pmatrix} \sim \begin{cases} \mathrm{N}\left(\mathbf{0}, \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\xi^2 \end{pmatrix}\right) & \text{with probability } 1-\epsilon_\eta-\epsilon_\xi \\ \mathrm{N}\left(\mathbf{0}, \begin{pmatrix} \tilde{\sigma}_\eta^2 & 0 \\ 0 & \sigma_\xi^2 \end{pmatrix}\right) & \text{with probability } \epsilon_\eta \\ \mathrm{N}\left(\mathbf{0}, \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & \tilde{\sigma}_\xi^2 \end{pmatrix}\right) & \text{with probability } \epsilon_\xi \end{cases} \tag{7}$$

where $\tilde{\sigma}_\eta^2 \gg \sigma_\eta^2$ and $\tilde{\sigma}_\xi^2 \gg \sigma_\xi^2$. The factors ϵ_η and ϵ_ξ represent the prior probability of a level shift and ramp.

The Gaussian mixture model has formed the basis for much of the research into non-Gaussian time series modeling. In the statistics literature, the model defined by (6) and (7) was introduced by Harrison and Stevens (1976), who called it the "Multiprocess Model" (see also Harrison and Stevens (1971)). This model has been successfully used in a Bayesian setting for several applications: see, for example, Smith and West (1983) and Gordon and Smith (1990). Kitagawa (1990) uses a similar model for robust seasonal adjustment, except that σ_η^2 and $\tilde{\sigma}_\eta^2$ are constrained to be zero (so ramps but not level shifts are modeled).

Seasonal breaks are not considered in this study, but (7) could easily be generalized to do so. However, from the perspective of generating appropriate non-Gaussian disturbances, the seasonal models given by (3) and (4) may not be so useful.

3.3 Technical Issues

Evaluation of the Likelihood

In structural time series models, the likelihood function is often decomposed in the form

$$L(Y_1, Y_2, \ldots, Y_N) = p(Y_1)p(Y_2|Y_1)\ldots p(y_N|Y_1, Y_2, \ldots, Y_{N-1}) \tag{8}$$

In the purely Gaussian case, (8) is readily computed by casting the model in state space form and applying the Kalman filter (see, for example, Harvey (1989)). For the Gaussian mixture model of (6) and (7), exact computation of (8) involves an algorithm with complexity of 6^N! This is because the one-step ahead predictive distribution for Y_t is a Gaussian mixture of 6^t components.

Different approaches have been adopted to circumvent this difficulty. Alspach and Sorenson (1972) develop a "Gaussian sums" method in which low probability components of the mixture are pruned. Pruning low probability components is known to cause problems (Kitagawa (1990)). Harrison and Stevens (1976) invoke a collapsing procedure in which a number of terms in the mixture at each time are replaced with a Gaussian distribution. The replacement is done through moment matching, and minimizes the Kulback-Leibler distance. Kitagawa (1990) adopts a collapsing approach similar to Harrison and Stevens. The main difference is that while Harrison and Stevens collapse the same densities at each time, Kitagawa successively collapses the pair of densities which are "closest" in terms of Kulback-Leibler distance. Bruce and Martin (1992) combine a variety of pruning and collapsing methods in an adaptive tree growing algorithm.

Obtaining the seasonal decomposition

To obtain a seasonal decomposition, we need an estimate of the "smoothed" trend and seasonal. Most procedures yield the expected trend and seasonal: $E(T_t|Y_1, Y_2, \ldots, Y_N)$ and $E(S_t|Y_1, Y_2, \ldots, Y_N)$. Using a "Gaussian sum" smoother, Kitagawa (1990) shows how we can actually get an estimate of the densities $p(T_t|Y_1, Y_2, \ldots, Y_N)$ and $p(S_t|Y_1, Y_2, \ldots, Y_N)$. From the densities, we could obtain a point estimate using the expected value. However, in the non-Gaussian setting, a more natural point estimate is given by the *median.*

A particularly nice feature about this approach is that we can readily obtain confidence intervals. In addition, Bruce and Cordera (1992) show how to obtain estimates of the posterior probabilities of outliers, level shifts, and ramps.

Initialization of the Filters

To compute the likelihood (8) and to use the two-filter smoother of Kitagawa (1990), we need to handle densities such as $p(Y_t|Y_1, Y_2, \ldots, Y_{t-1})$. For $t < 14$ with monthly data and the models considered above, certain assumptions are needed about initial conditions. In the purely Gaussian case, the most natural approach is to assume a "diffuse prior": see Ansley and Kohn (1985), Bell and Hillmer (1987), and De Jong (1991). In the Gaussian mixture case, a "diffuse prior" could be used as well. The exact distribution, though, is a Gaussian mixture involving an intolerable number of terms for t bigger than 5 or 6. The various schemes for reducing the number of components do not work: the distributions are partially diffuse and it is not possible to determine which observations are likely outliers, level shifts, etc.. As a result, the current implementation of MING does not properly handle the initialization. This causes problems in the seasonal adjustments with three series (BTAPRI, BTNDRI, and ITVRUO). See section 6.5 of [BJ92a] for details.

3.4 The MING program

A non-Gaussian seasonal adjustment method based on time series structural models, called "MING", was developed for this study. It is based on a computer program developed by Kitagawa (1991). It permits trend models of the form (2) and a choice of either (3) or (4) for the seasonal model. The irregular is assumed to be white noise. Gaussian mixture distributions of the form (6) and (7) can be specified.

The mixture distribution for structural changes, given by (7), can be extended to accommodate seasonal breaks as well. The method of Kitagawa (1990) is used for reducing the number of mixture terms in computing the likelihood and smoothed estimates. MING is implemented in Fortran-77 as a function in the S-PLUS language (S-PLUS (1991)).

Ease of Use

Potentially, MING is easier for the naive seasonal adjuster to apply than X-12-ARIMA. In this study, only the transformation choice was not automatic (but it could easily be done so). By contrast, both an ARIMA model and the seasonal filters had to be specified for X-12-ARIMA.

Offsetting these advantages is the relative computational inefficiency of MING. In general, MING is an order of magnitude slower than X-12-ARIMA. However, much can be done to improve the speed of MING. With the increasing computing power available to a broad spectrum of users, this should not be a major factor in the near future.

Continuity property of MING

An important property of the MING method is that it employs a continuous scheme for handling outliers and level shifts. MING can adapt to different magnitudes of AO's or LS's using appropriate posterior probabilities. Large AO's or LS's are given probabilities close to 1 while small AO's or LS's are given probabilities close to 0. By contrast, the X-12-ARIMA outlier prior adjustment procedure declares observations as either AO's or LS's or neither. Essentially, X-12-ARIMA assigns a posterior probability of either 1 or 0. As a result, we can expect the X-12-ARIMA seasonal adjustment to change discontinuously as an observation passes the threshold and is declared as an AO or LS. While this discontinuity is mitigated by the outlier treatment intrinsic to X-11, we shall see that the MING outlier method leads to more stable seasonal adjustments.

4 The Data, Filters, and Models

4.1 The Data

The empirical study involves 29 monthly U.S. macroeconomic time series, selected by time series staff at the Statistical Research Division, U.S. Bureau of the Census. These series have both outliers and structural changes (such as level shifts). Table 1 lists the series along with their abbreviations. Of the 29 series, 13 are retail trade series, 7 are housing starts series, and 9 are inventory series. The series exhibit a range of seasonal behavior. The retail trade series usually display very strong seasonal patterns. The construction series are often very erratic and quite difficult to adjust. The inventory series tend to have large level shifts but relatively stable adjustments. On the whole, this collection of economic time series gives a broad range of problems with which to assess and compare seasonal adjustment methods.

	Abbrev-iation	Series Description	SABL transform
B1	BAUTRS	Retail Sales of Automobiles	0
B2	BFRNRS	Retail Sales of Furniture	0
B3	BGMRRI	Retail Sales of General Merchandise	-0.5
B4	BGRCRS	Retail Sales of Groceries	0
B5	BHDWWS	Wholesale Sales of Hardware	0
B6	BLQRRS	Retail Sales of Liquor	0
B7	BMNCRS	Retail Sales of Men's Apparel	-0.25
B8	BSHORS	Retail Sales of Shoes	0
B9	BSPGWS	Wholesale Sales of Sporting Goods	0.25
B10	BTAPRI	Total Retail Sales of Apparel	-0.25
B11	BTNDRI	Retail Sales of Nondurables	-0.25
B12	BVARRS	Variety Store Retail Sales	0.25
B13	BWAPRS	Retail Sales of Women's Apparel	0
C14	CMW1HS	One Family Housing Starts in the Midwest	-0.25
C15	CMWTHS	Total Housing Starts in the Midwest	0
C16	CNE1HS	One Family Housing Starts in the Northeast	0.25
C17	CNETHS	Total Housing Starts in the Northeast	0.25
C18	CSOTHS	Total Housing Starts in the South	0
C19	CWETHS	Total Housing Starts in the West	0.25
C20	C24THS	Total Housing Starts - 2 to 4	0.25
I21	IBEVTI	Total Inventories of Beverages	-1
I22	ICMETI	Total Inventories of Communications Equipment	-1
I23	IFATTI	Total Inventories of Fats and Oils	-1
I24	IFMETI	Total Inventories of Farm Machinery and Equipment	0
I25	IGLCTI	Total Inventories of Glass Containers	-0.25
I26	IHAPTI	Total Inventories of Household Appliances	-0.25
I27	INEWUO	Unfilled Orders for Newspapers and Magazines	0.25
I28	ITVRTI	Total Inventories of TV's and Radios	0.25
I29	ITVRUO	Unfilled Orders for TV's and Radios	0.5

Table 1: List of abbreviations for the 29 series studied and power transformations used by the SABL seasonal adjustment procedure. The choice of powers is roughly consistent with a logarithmic transform (power = 0).

4.2 Log Additive Seasonal Decomposition

Time series staff at the Statistical Research Division identified a multiplicative decomposition in all of the series for the X-12-ARIMA seasonal adjustments. In this study, we transform the data by taking logarithms and multiplying by 1000 and then apply an additive seasonal adjustment. The log-additive seasonal adjustment is used so that MING can be fairly compared with X-12-ARIMA. While multiplicative and log-additive seasonal adjustments should be similar, there is a consistent downward bias in the trend from the log-additive procedure (see Ozaki and Thomson (1992)).

For simplicity and conceptual clarity, we have chosen not to consider other possible transformations (e.g., square root). Table 1 displays the transformation powers for each of the series chosen by the robust seasonal adjustment procedure SABL (Cleveland and Devlin (1980)). The powers range from -1 to 0.5, with the majority being ± 0.25. While the SABL transforms may not be "optimal" (see Shulman and McKenzie (1984)), these results support the use of a logarithmic transformation.

4.3 X-12-ARIMA Filters

For the 29 series in this study, the options required by X-12-ARIMA were provided by the time series staff of the U.S. Bureau of the Census. This includes the choice of ARIMA models, filters, and trading day and Easter effects. For the retail trade and inventory series, the default filters are used (see Dagum (1980)). For the construction series, a 3×9 moving average is used, yielding smoother seasonal factors than the default filter. Trading day prior adjustment is done for all of the retail trade series and three of the construction series. Prior adjustment for the timing of Easter is done for five of the retail trade series.

4.4 MING Models

Three types of seasonal models are fit to all 29 series:

BSM: the "BSM" seasonal model given by (3),

TRIG-1: the trigonometric seasonal given by (4) with the assumption that all of the noise terms $\omega_{j,t}$ have a common variance σ^2_ω,

TRIG-6: and the trigonometric seasonal given by (4) allowing different variances $\sigma^2_{\omega,j}$.

We will use the acronyms BSM, TRIG-1, and TRIG-6 to refer to these models.

For each model, we fit the variance parameters σ^2_{I}, σ^2_η, σ^2_ξ, and σ^2_ω (or $\sigma^2_{\omega,j}$ in the case of TRIG-6). We also fit the prior probabilities of an additive outlier and a level shift (ϵ_{I} and ϵ_η).

We did not optimize over the variances of the outlier and level shift processes, $\tilde{\sigma}^2_{\mathrm{I}}$ and $\tilde{\sigma}^2_\eta$. Instead these variances are set to:

$$\tilde{\sigma}^2_{\mathrm{I}} = \tilde{\sigma}^2_\eta = 100\sigma^2_{t|t-1}$$

where $\sigma^2_{t|t-1}$ is the steady state one-step ahead prediction variance. The value of 100 was chosen based on practical considerations. It seemed big enough to ensure that the model would handle all conceivable outliers and level shifts but not so big as to distort the likelihood. Our experience agrees with that of Smith and West (1983) who suggest that the likelihood is relatively invariant with respect to $\tilde{\sigma}_{\rm I}^2$ and $\tilde{\sigma}_\eta^2$.

A simplified version of the level shift and ramp model (7) is used in this study. The prior probability of a ramp ϵ_ξ is set to zero, reducing the Gaussian mixture in (7) to just two terms. This restricted model provides for most of what is desired in terms of modeling structural changes while significantly reducing the computational burden. Section 9.4 explores a more general ramp model.

4.5 Model Parameters

Table 2 gives the maximum likelihood estimates of the parameters for TRIG-6 for each of the series. See appendix A for details on the fitting procedure. The table

	Variances				Probabilities		Type of convergence
	mean σ_η^2	slope σ_ξ^2	irregular σ_ϵ^2	seasonal $\bar{\sigma}_\omega^2$	outlier $\epsilon_{\rm I}$	level shift ϵ_η	
BAUTRS	1488	1.228e-08	0.004631	1.657	0.008848	0.0001354	X
BFRNRS	285.9	9.559e-08	67.15	0.3738	5.943e-05	0.0005969	R
BGMRRI	81.22	1.514	0.005892	0.4867	0.0003415	0.0284	R
BGRCRS	4.624	0.5419	100	9.696e-14	9.294e-05	0.0004777	R
BHDWWS	471.2	7.394e-05	354.3	0.6132	5.643e-05	0.0001864	R
BLQRRS	137.4	0.05058	121.7	0.2775	0.006969	0.004569	F
BMNCRS	123.3	0.09107	335.6	4.963	8.679e-05	8.67e-05	R
BSHORS	408	0.001657	317.2	1.616	7.074e-05	0.008381	F
BSPGWS	1934	9.571e-05	843.3	4.77	0.006177	0.0002729	B
BTAPRI	132.6	2.174e-05	6.783e-06	0.1578	6.098e-05	0.002837	R
BTNDRI	41.84	0.002471	0.05045	0.2091	0.0001406	0.01906	F
BVARRS	134.1	0.0261	342.5	0.8838	0.000809	0.01346	F
BWAPRS	189.1	0.0149	217.6	2.153	0.000111	0.0001306	B
C24THS	8946	8.627e-05	11830	3.336	0.0007437	0.0006874	R
CMW1HS	5728	0.0008224	9063	4.575	0.028	0.0095	R
CMWTHS	11940	0.0001437	9377	9.969	0.008784	0.00821	
CNE1HS	4596	5.12e-05	11230	34.2	0.00503	0.000751	R
CNETHS	7683	0.001412	18370	17.77	0.01361	0.001678	X
CSOTHS	5014	0.001275	3083	0.4219	0.0004325	0.001477	R
CUSTHS	4276	3.132e-06	816.3	2.903	4.054e-05	4.344e-05	F
CWETHS	8078	0.000324	3979	2.895	0.0005571	0.0005949	R
IBEVTI	152.6	0.2058	5.47e-08	0.521	9.929e-05	0.001752	R
ICMETI	20.16	17.83	23.91	0.05469	1.052e-08	0.006383	R
IFATTI	4044	0.002189	2.959e-05	4.974	5.976e-05	0.002057	B
IFMETI	284.6	6.747	26.04	1.503e-09	0.0009216	0.01885	R
IGLCTI	517.3	0.1818	0.001409	0.1159	0.0001653	0.01075	R
IHAPTI	443.3	0.02705	1.685e-05	0.6787	0.0001042	0.005982	X
INEWUO	2221	0.9796	0.000118	4.472	7.64e-05	0.002101	R
ITVRTI	751.4	18.72	0.006499	0.5567	0.0008571	1.561e-05	R
ITVRUO	6574	0.01101	906.3	85.11	0.02277	0.003414	F

Table 2: Structural model parameters as estimated by GAUSUM-STM for the TRIG-6 model. Note that $\bar{\sigma}_\omega^2 = \sum_{j=1}^{6} \sigma_{\omega,j}^2/6$. The individual seasonal variances are given in Table 3.

displays the variances σ_η^2, σ_ξ^2, σ_I^2, the mean of the variances of the seasonal component $\bar{\sigma}_\omega^2 = \sum_{j=1}^{6} \sigma_{\omega,j}^2/6$. and the prior probabilities ϵ_η and ϵ_I.

The type of convergence achieved by the optimizer is also given in Table 2. Five types of convergence are possible: (R)elative function convergence, (X)-convergence, (B)oth X- and relative function convergence, and (F)alse convergence (see appendix A for details). These are denoted in the table by the letters in the parentheses. A convergence code of R, X, or B indicates that the optimizer successfully found a local maximum. A convergence code of F means that the optimizer may be stuck at a non-critical value.

Note that a few series have false convergence for TRIG-6. By contrast, almost all of the fits for the BSM and TRIG-1 achieved successful convergence (Bruce and Jurke (1992b)). This is probably due to the number of seasonal parameters and the relative flatness of the likelihood. Since the TRIG-6 uses the maximum likelihood estimates of TRIG-1 as starting values, we expect that reasonably good estimates are obtained in all cases. Furthermore, our experience with the optimizer indicates that many of the false convergences are at a local maximum.

Comparison of Seasonal Parameters

Table 3 compares the variances of the seasonal components for TRIG-1 and TRIG-6. The TRIG-6 variances are given relative to the TRIG-1 variance. Let $\sigma_{\omega,j}^2$ for $j = 1, \ldots, 6$ be the variances for the TRIG-6 model and let σ_ω^2 be the variance for the TRIG-1 model. Table 3 gives σ_ω^2 and $\tilde{\sigma}_{\omega,j}^2 \equiv \sigma_{\omega,j}^2/\sigma_\omega^2$.

For each series, the largest variance $\sigma_{\omega,j}^2$ is highlighted by a surrounding box. The dominant variability is in either at the seasonal frequency or the first harmonic, as one would expect. There are a few interesting exceptions: for the C24THS series, the dominant variance is the last harmonic!

Table 3 also gives Akaike's Information Criterion (AIC) for the BSM, TRIG-1, and TRIG-6 models. AIC is defined by

$$\text{AIC} = -2 \times \log \text{L} + 2p$$

where $\log \text{L}$ is the log-likelihood at the maximum and p are the number of parameters fit. AIC gives a guide towards selecting the "best" model, and models with lower AIC values are preferable. The model with the minimum AIC value is marked by a box in table 3.

In regards to goodness of fit, Harvey (1989) (p. 43) states that it is usually not necessary to optimize over all six variances of the trigonometric seasonal model. This is important since optimizing over six variances is far more laborious than optimizing over just one variance. However, the AIC's of Table 3 favor TRIG-6 in several instances, sometimes by a substantial amount. This suggests that we should not blindly optimize over just one variance without good reason. In particular, it would be useful to have a diagnostic that indicates when further optimization would not lead to a significantly better fit.

5 Flexibility and Smoothness in the Seasonal Factors

In this study, we focus on the *flexibility* and *smoothness* of the seasonal factors. The most important result of this study is that many of the 29 series compared, the

	TRIG-1	TRIG-6 variances						AIC		
	σ^2_ω	$\tilde{\sigma}^2_{\omega,1}$	$\tilde{\sigma}^2_{\omega,2}$	$\tilde{\sigma}^2_{\omega,3}$	$\tilde{\sigma}^2_{\omega,4}$	$\tilde{\sigma}^2_{\omega,5}$	$\tilde{\sigma}^2_{\omega,6}$	BSM	TRIG-1	TRIG-6
BAUTRS	0.84	0.0020	**8.84**	2.44	0.54	0.0025	0.0027	2748	2742	**2736**
BFRNRS	0.40	0.0011	0.86	**2.37**	1.51	0.48	0.37	2351	**2350**	2355
BGMRRI	0.19	**12.09**	2.90	0.88	0.041	0.43	0.015	1764	1751	**1729**
BGRCRS	7.19e-15	27.87	0.37	0.069	17.69	0.046	**34.9**	**2049**	**2049**	2059
BHDWWS	0.31	1.14	1.54	0.045	**8.03**	0.76	0.48	2566	**2565**	2572
BLQRRS	0.068	**18.58**	0.066	4.59	0.36	0.12	0.95	2290	**2289**	2296
BMNCRS	2.71	**5.96**	3.00	0.94	0.85	0.15	0.089	2577	2559	**2544**
BSHORS	1.61	**2.93**	0.0037	0.26	2.37	0.39	0.071	2607	2604	**2603**
BSPGWS	2.59	**6.16**	0.64	2.74	0.75	0.26	0.51	2888	**2884**	2888
BTAPRI	0.059	**11.7**	1.25	2.01	1.21	3.12e-05	7.92e-06	1667	1664	**1663**
BTNDRI	0.073	**13.69**	2.08	1.39	0.00041	0.025	0.0029	1519	1500	**1484**
BVARRS	0.76	**3.81**	1.65	3.27e-06	0.97	0.53	0.016	2506	2493	**2489**
BWAPRS	0.99	1.10	**5.77**	4.94	1.07	0.094	0.081	2441	2435	**2434**
C24THS	3.33	0.010	0.23	0.10	1.29	0.39	**3.99**	3855	**3853**	3860
CMW1HS	1.78	**11.18**	7.55e-07	3.98	3.12e-05	3.11e-05	0.2555	3880	**3879**	3883
CMWTHS	0.0067	**8352**	535.7	82.7	0.80	0.25	0.19	**3858**	3889	3865
CNE1HS	34.01	2.41	**2.51**	0.56	0.034	0.48	0.043	3877	**3866**	3868
CNETHS	5.54	**10.86**	5.81	1.22	0.0048	1.34	0.0036	3970	**3968**	3974
CSOTHS	3.25e-08	9437	42650	18040000	**59730000**	141800	250.6	**3564**	**3564**	3570
CWETHS	0.86	**15.87**	0.00042	0.037	0.00092	2.48	1.87	3686	**3684**	3691
IBEVTI	0.072	**33.44**	6.86	2.73	0.31	0.28	6.12e-05	2460	2455	**2435**
ICMETI	0.055	0.00036	**3.28**	1.07	1.34	0.26	0.016	1959	**1958**	1963
IFATTI	1.22	**15.94**	3.64	4.47	2.32e-07	0.34	0.052	3327	3323	**3315**
IFMETI	4.66e-10	1.91	**10.19**	0.0087	1.49	2.34	3.43	**2879**	**2879**	2889
IGLCTI	0.040	1.93	0.83	0.22	**13.84**	0.30	0.36	**3003**	**3003**	3011
IHAPTI	0.019	**207.6**	6.32	0.00031	0.00018	1.78	1.19	2913	2913	**2911**
INEWUO	2.0	**7.09**	2.76	2.95	0.34	0.225	0.071	3186	3177	**3169**
ITVRTI	0.25	5.97	**6.52**	0.44	0.24	0.024	1.84e-05	2895	2882	**2877**
ITVRUO	49.98	1.22	**6.24**	0.50	0.91	0.066	1.28	3818	3798	**3782**

Table 3: Seasonal parameters as estimated by GAUSUM-STM for the TRIG-1 and TRIG-6 models. The TRIG-6 variances are given as ratios to the TRIG-1 varaince: $\tilde{\sigma}^2_{\omega,j} \equiv \sigma^2_{\omega,j}/\sigma^2_\omega$. The AIC values are given for the BSM, TRIG-1, and TRIG-6. The model with the minimum AIC value is marked with a "*".

X-12-ARIMA decomposition leads to smoother and more flexible seasonal factors as compared with the structural models. As we argue below, the "roughness" of the structural models is a serious problem which needs to be addressed in future research.

This study also casts doubt about the usefulness of the "BSM" seasonal model, and indicates that optimizing over six variances in the trigonometric models often makes a significant difference. These results are discussed in more detail below.

5.1 X-12-ARIMA Smoother and More Flexible

Flexibility measures the amount that the seasonal effect for a given month is allowed to "bend" or change from year to year. In many series, the seasonal factors evolve over time. A certain amount of flexibility in a seasonal adjustment procedure is desirable to allow adaptation to this evolution. A procedure which is too rigid will not remove enough of the seasonality from the series.

Seasonal factors of a given flexibility can either be slowly varying or rapidly changing. This corresponds to a monthly effect which evolves smoothly or roughly from year to year. All things being equal, smoother seasonal factors are preferable. Hannan (1964) states

> ...there seems little point in allowing for anything more than the very slowest change in seasonal variation. It would seem wrong here to concern oneself too much with faithfully representing a possibly rapidly changing seasonal because of the consequent risk of seriously distorting the series.

In other words, we do not want to allow seasonal patterns to evolve into very different shapes in a short time frame. Rapid changing seasonal patterns are not what "most users" think of as a seasonal effect. We are probably better off putting "excess" local variation of a seasonal pattern into the trend or irregular component.

In summary, *long term* flexibility is generally desirable but *short term* variability is not. We will show below that the *long term* flexibility of X-12-ARIMA is as great or greater than that of MING. However, the *short term* evolution of seasonal patterns for MING are much "rougher". The X-12-ARIMA seasonal adjustments are generally more appealing for the business and inventory series.

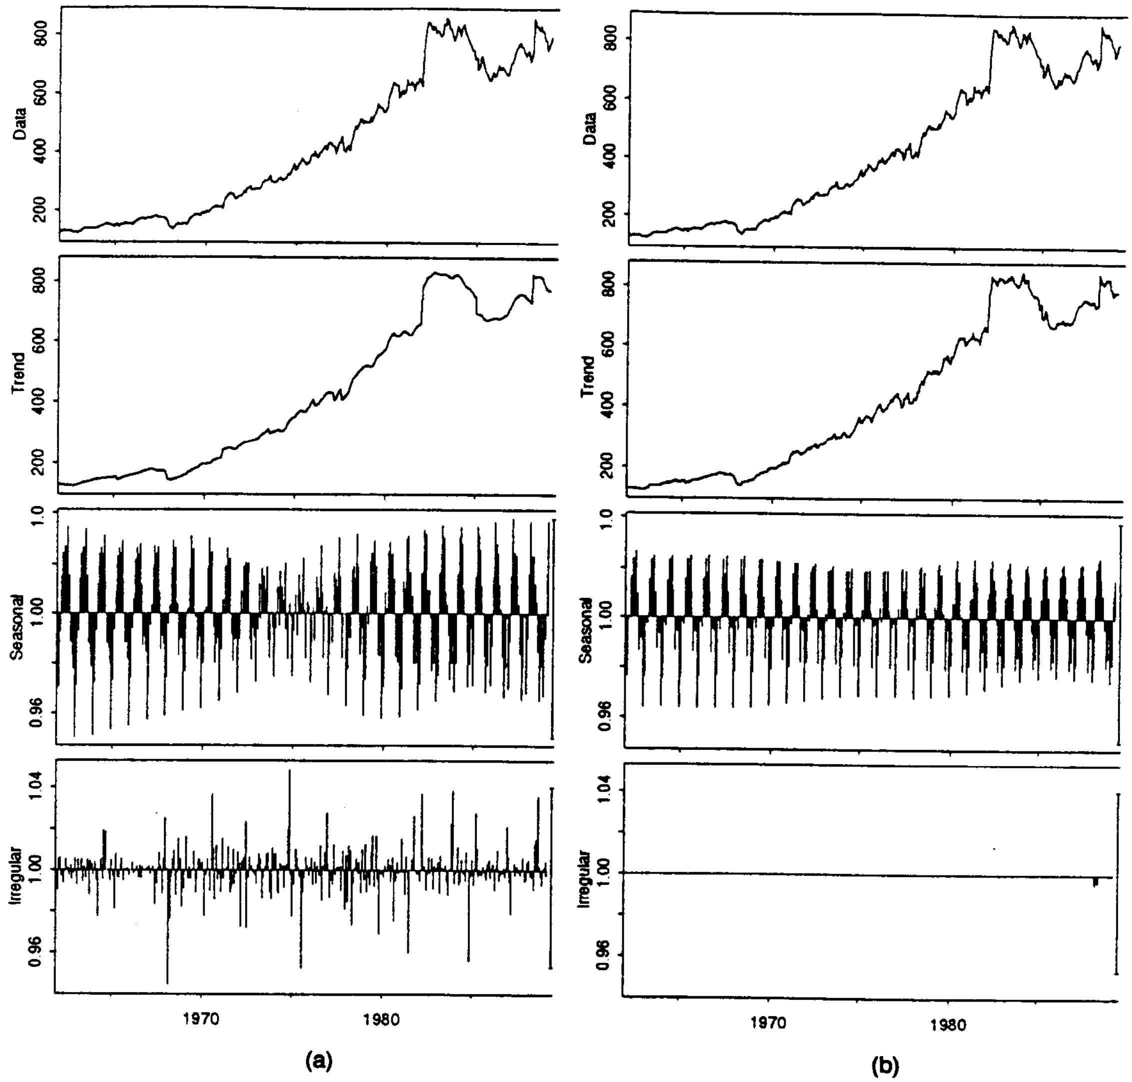

Figure 1: Decomposition for IGLCTI based on (a) X-12-ARIMA and (b) the TRIG-6 structural model. The top plot gives the untransformed data. The second, third and fourth plots show the trend, seasonal, and irregular. The vertical bar to the right of bottom two plots are the same length in real coordinates. The bars compare the relative strengths of the seasonal and irregular components.

An Example: X-12-ARIMA More Flexible

Figure 1 gives the seasonal decompositions obtained by the X-12-ARIMA and TRIG-6 methods for the IGLCTI series. The data, trend, seasonal and irregular are displayed in the four plots (from top to bottom). Since a log-additive decomposition is used, the trend, seasonal and irregular are exponentiated:

$$Y_t = \exp^{T_t} \exp^{S_t} \exp^{I_t}$$

Focusing on the seasonal component, the most obvious feature in this plot is the vastly greater flexibility of the X-12-ARIMA seasonal.

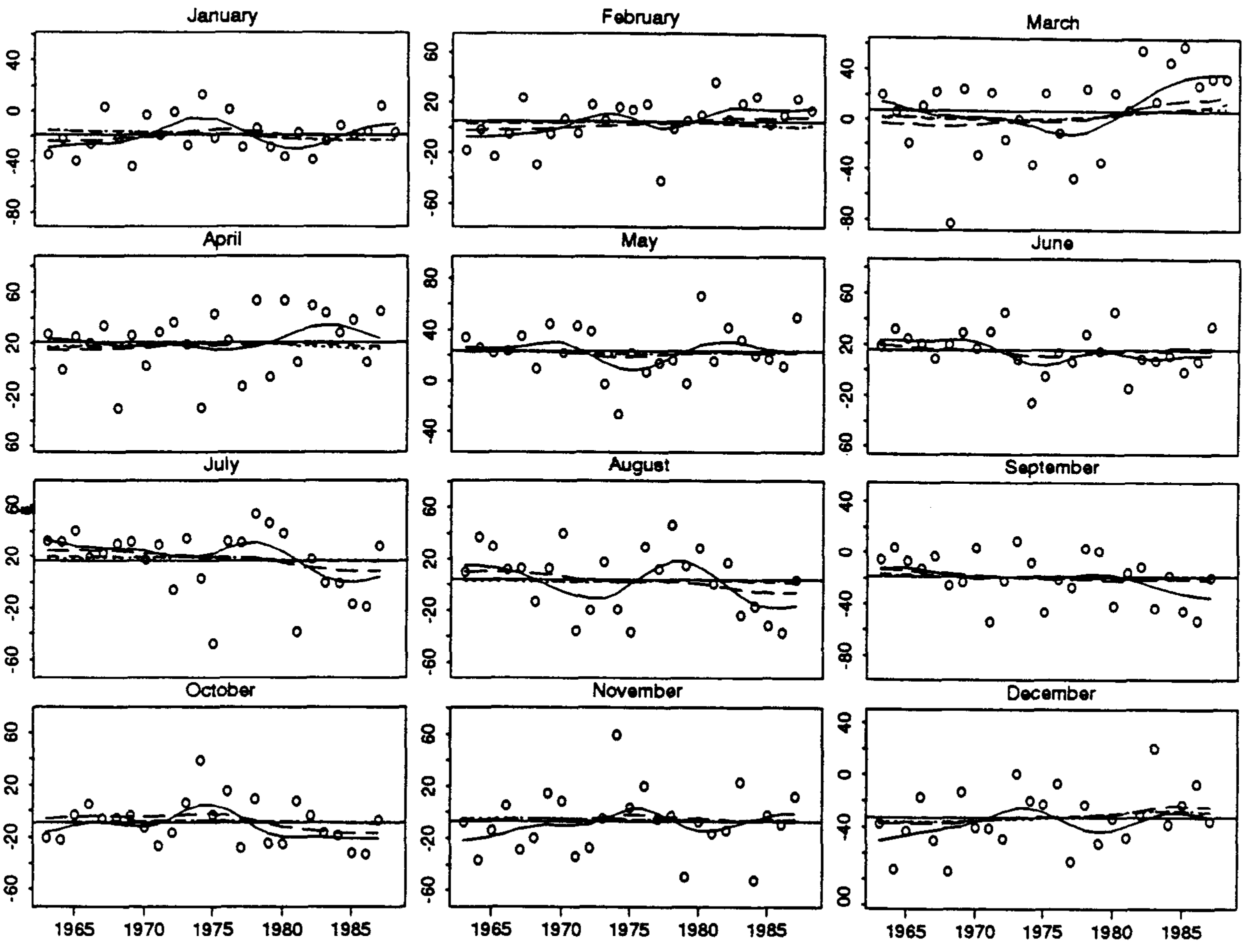

Figure 2: SSI-plot for IGLCTI. The points are the log-transformed detrended series plotted by month. The estimated seasonal factors are compared for the BSM (short dashed line), TRIG-1 (medium dashed line), TRIG-6 (long dashed line) and X-12-ARIMA (solid line). The horizontal solid line corresponds to the mean.

To compare the different natures of the seasonal factors, we turn to the "SSI-Plot" of Cleveland and Terpenning (1982). Figure 2 gives a version of the SSI-Plot for the IGLCTI series, comparing X-12-ARIMA with the structural models. For each month, figure 2 displays the detrended transformed data as points (see appendix B for a description of the detrending procedure). The horizontal solid line corresponds to the mean. The seasonal factors are displayed for X-12-ARIMA (bending solid line), the BSM (short dashed line), TRIG-1 (medium dashed line), and TRIG-6 (long dashed line).

The structural models show very little deviation from the mean effect, as we would expect from figure 1. However, X-12-ARIMA smoothly adapts to the data, capturing such features as the slow variation of the August effect.

Is this adaptivity is desirable? In other words, are we removing "true" seasonality from the series, or are we distorting the series by putting too much flexibility into the seasonal component. While this is partly a subjective issue, one way to check this is to see how much power has been removed around the seasonal frequency (recall that changing seasonal patterns over time will result in smearing of power in the periodogram).

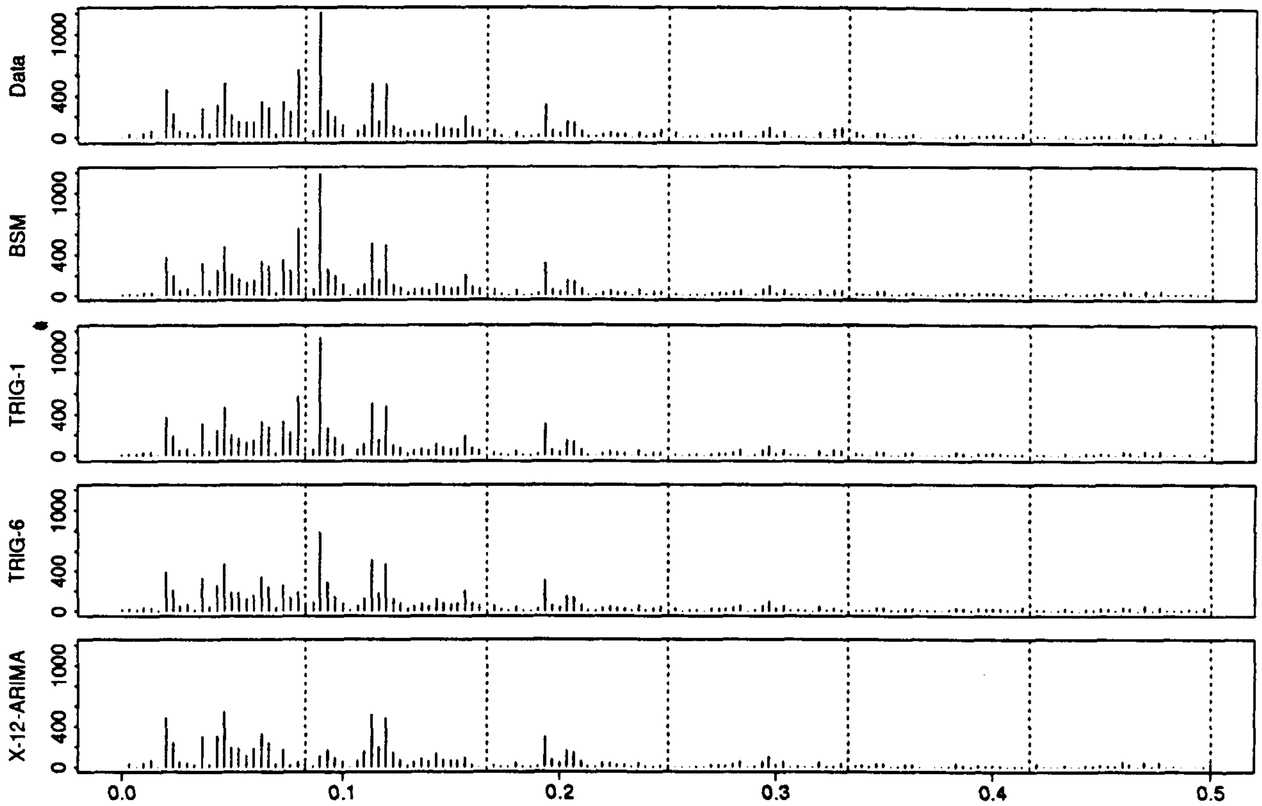

Figure 3: The top plot is the periodogram of the detrended data for IGLCTI. The seasonal and its harmonics, marked by dashed lines, are suppressed in this periodogram. The other plots give the periodograms for the detrended seasonally adjusted data for BSM, TRIG-1, TRIG-6, and X-12-ARIMA.

Figure 3 compares the periodograms of the detrended seasonal adjusted data (see appendix B for details). The top plot in figure 3 is the periodogram for the original data with the values at the seasonal frequency and its harmonics suppressed. This corresponds roughly to what one would obtain by fitting a fixed seasonal pattern. The subsequent plots show the periodograms for the BSM, TRIG-1, TRIG-6 and X-12-ARIMA procedures. We can see that X-12-ARIMA removes considerably more power around the seasonal frequency.

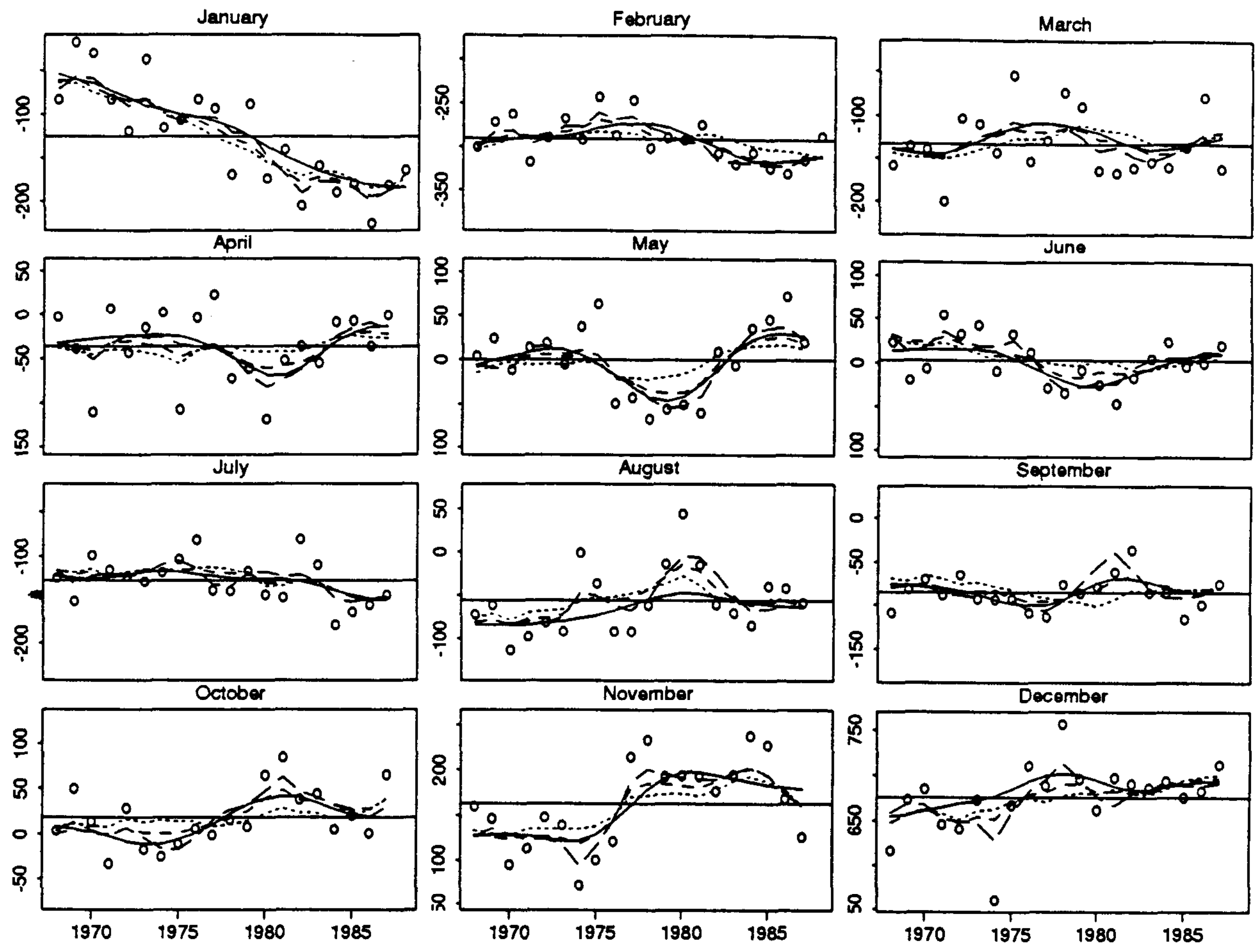

Figure 4: SSI-plots for BMNCRS. The points are the log-transformed detrended series plotted by month. The estimated seasonal factors are compared for the BSM (short dashed line), TRIG-1 (medium dashed line), TRIG-6 (long dashed line) and X-12-ARIMA (solid line). The horizontal solid line corresponds to the mean.

Another Example: X-12-ARIMA Smoother

An equally revealing example is given by the BMNCRS series. Figure 4 gives the SSI-Plot for the BMNCRS series (this is analogous to figure 2). In this case, the structural models are as or even more flexible then X-12-ARIMA. In fact, examination of the periodogram shows that TRIG-6 removes slightly more power around the seasonal frequency and its harmonics than X-12-ARIMA (see figure 15 of [BJ92a]). However, this flexibility is achieved at a significant increase in roughness. The seasonal factors for structural models exhibit a great deal of seemingly undesirable local variation.

Some Conclusions

The help quantify these results, we have constructed some diagnostics based on the notions of flexibility and smoothness. We measure long-term seasonal flexibility by the year to year change in the smoothed seasonal factors. Denote the seasonal factors of the log-transformed data by S_t. Let $\tilde{S}_t$ be as smoother version of S_t obtained by

$$\tilde{S}_t = \sum_{j=-1}^{1} \left(\frac{2-|j|}{4} \right) S_{t+12j}$$

The diagnostic is given by

$$\text{SEAS FLEX} \equiv \frac{100}{N} \sum_{t=t_0+12}^{t_N-12} |\Delta^{12} \tilde{S}_t|. \tag{9}$$

SEAS FLEX corresponds roughly to the mean annual percentage change in the smoothed seasonal:

$$\frac{\exp \tilde{S}_t - \exp \tilde{S}_{t-12}}{\exp \tilde{S}_{t-12}} \approx \Delta^{12} \tilde{S}_t.$$

Seasonal roughness is measured by simply looking at the mean absolute residuals from the smooth:

$$\text{SEAS ROUGH} \equiv \frac{100}{N} \sum_{t=t_0}^{t_N} |S_t - \tilde{S}_t|. \tag{10}$$

Extra flexibility in the seasonal factors is only desirable if it removes seasonality from the data. To measure this, we look at the remaining seasonality by the seasonally as reflected in the periodogram near the seasonal frequency $s = \pi/6$. Specifically, we define the statistic SEAS REMAIN by

$$\text{SEAS REMAIN} \equiv \sum_{j=-2}^{2} \left(\frac{3-|j|}{9} \right) I_{s+j} \tag{11}$$

where I_{s+j} is the magnitude of the periodogram at frequency $s + 2\pi j/N$.

Figure 5 displays the diagnostics SEAS ROUGH, SEAS FLEX and SEAS REMAIN. Two outlying series were removed from these plots for clarity (the structural models performed quite poorly for these series). The diagnostics have been "median corrected:" the median of the diagnostic for the four procedures is subtracted out for each series.

For the moment, we concentrate on the comparison of X-12-ARIMA with TRIG-6. For the business series, according to SEAS ROUGH, X-12-ARIMA produces smoother seasonal factors than the TRIG-6 models. At the same time, the X-12-ARIMA seasonal is as or more flexible than TRIG-6. For the inventory series, the

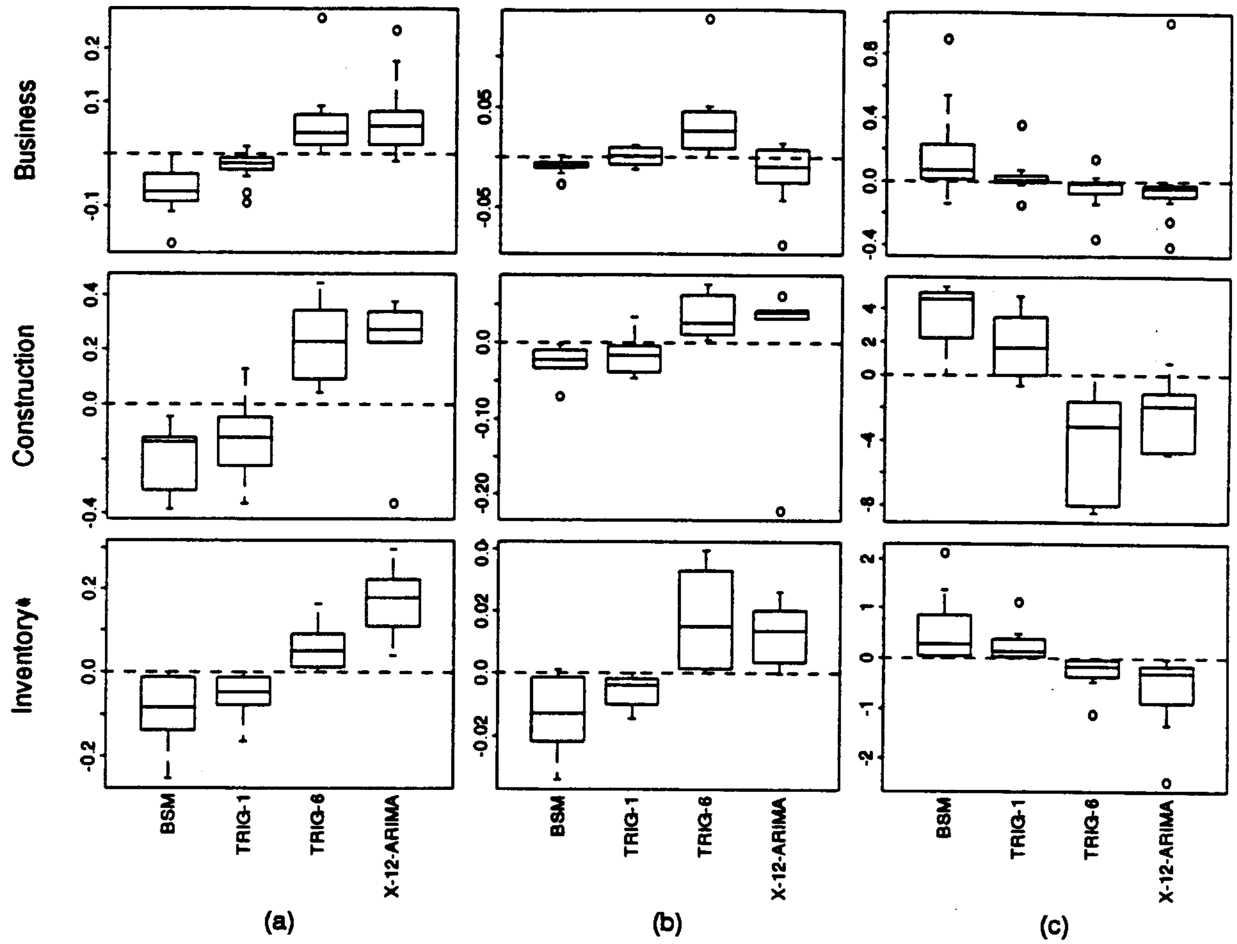

Figure 5: (a) The diagnostic SEAS FLEX, given by (9), (b) the diagnostic SEAS ROUGH, given by (10), and (c) the diagnostic SEAS REMAIN, given by (10). To remove the "series effect", the diagnostics for each series are median corrected (i.e., the median of the diagnostic for the four procedures is subtracted for each series).

X-12-ARIMA seasonal factors are more flexible and remove more seasonality than TRIG-6 without being any rougher.

These results are not surprising. The smoothness of X-12-ARIMA seasonal factors has been explicitly incorporated into the procedure based on practical considerations. By contrast, the seasonal factors for structural models are based on the maximum likelihood estimates. This does not guarantee seemingly desirable features such smoothness and long term flexibility.

Only for the construction series are the results ambiguous. The construction series tend to be difficult to adjust and perhaps are less useful for evaluating procedures (since often no procedure is really adequate).

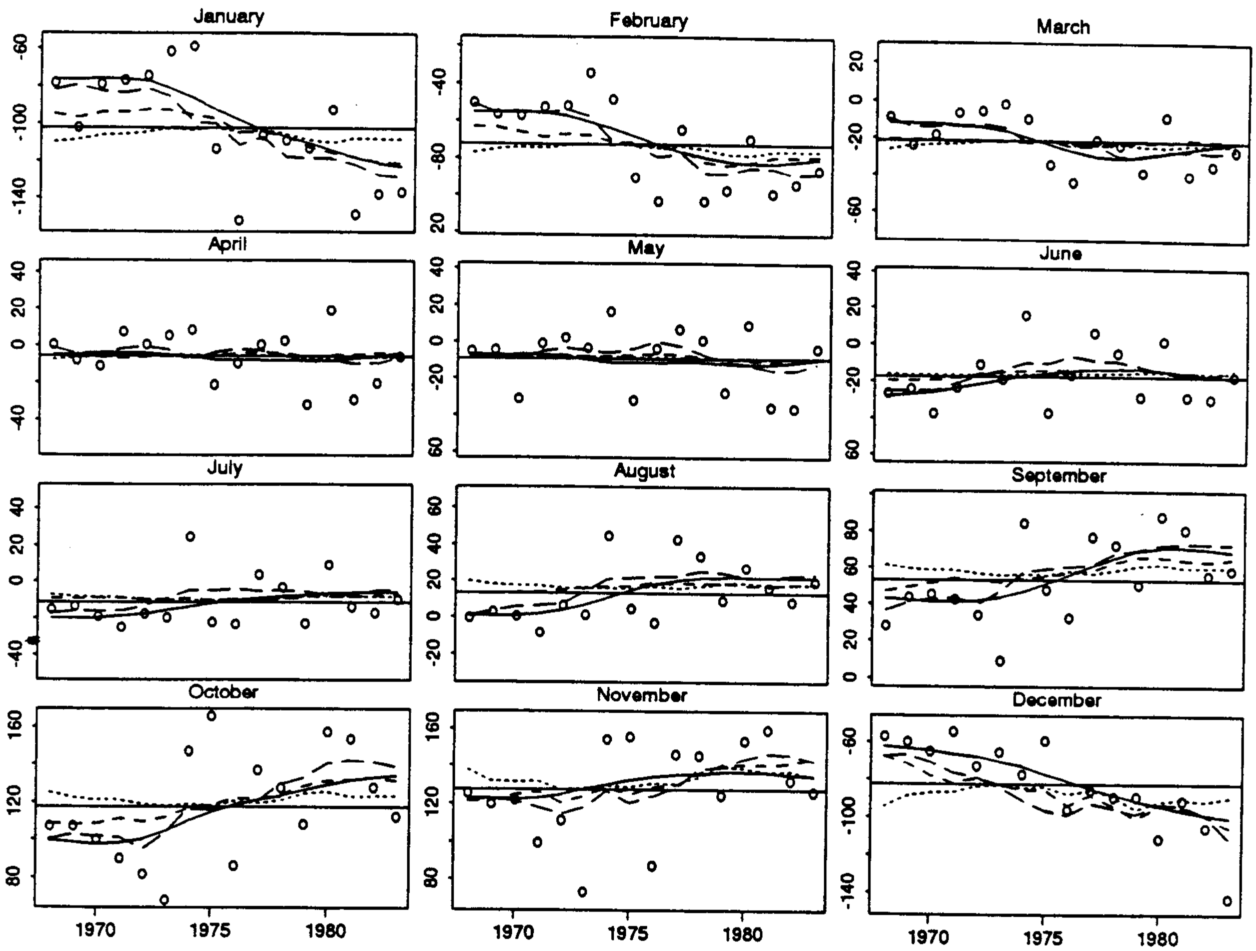

Figure 6: SSI-plots for BGMRRI. The points are the log-transformed detrended series plotted by month. The estimated seasonal factors are compared for the BSM (short dashed line), TRIG-1 (medium dashed line), TRIG-6 (long dashed line) and X-12-ARIMA (solid line). The horizontal solid line corresponds to the mean.

5.2 Rigidity of the BSM

The BSM seasonal model is not as flexible as the trigonometric seasonal model. This is evidenced by figure 5(c), which shows that the BSM consistently removes leaves more power in the spectrum around the seasonal frequency. According to the diagnostic plots, TRIG-1 adapts significantly better to changing seasonal patterns than the BSM in 8 series (B3, B7, B8, B10, B12, I21, I27, I28). Recall that TRIG-1 and BSM have the same number of parameters. If only one model is to be considered for seasonal adjustment, on the basis of these results, we would prefer TRIG-1.

The BGMRRI series provides a dramatic example of the inadequacy of the BSM seasonal model. Figures 6 and 7 give the SSI-Plot and periodogram plot for the

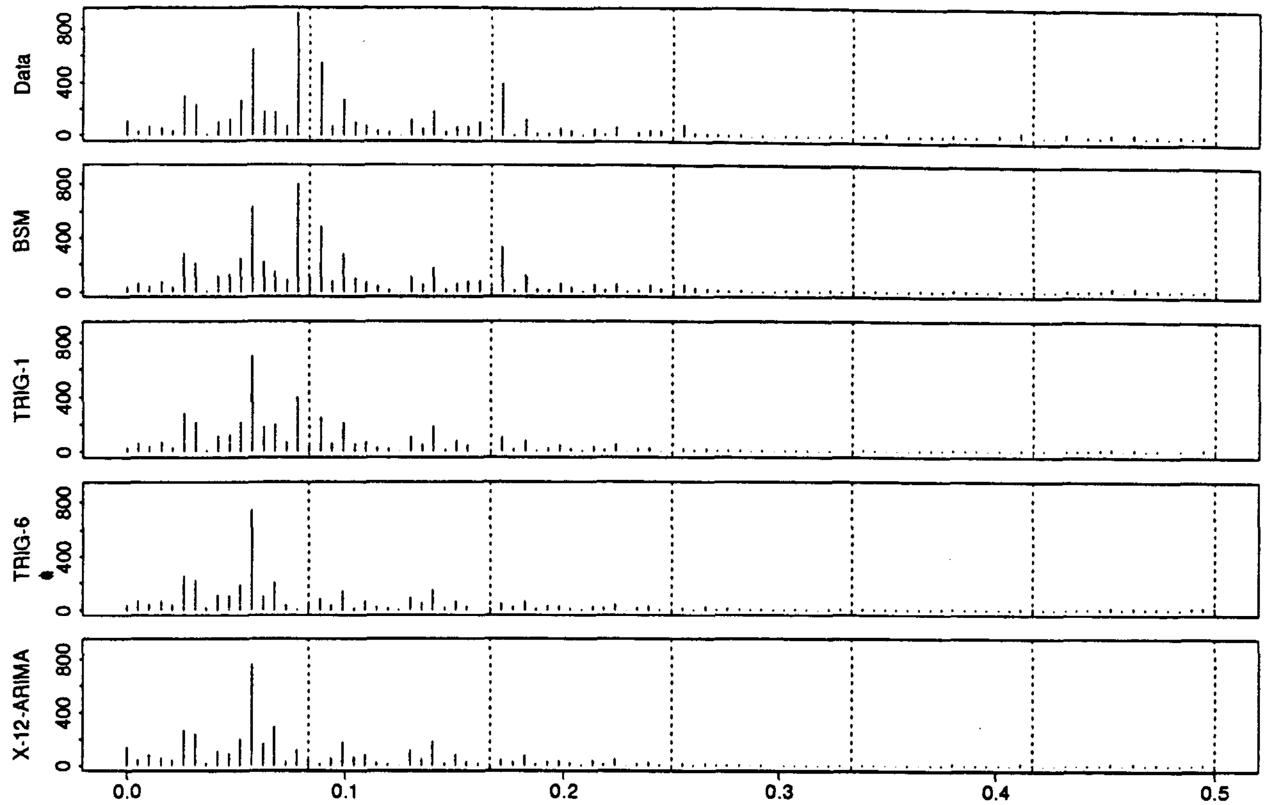

Figure 7: The top plot is the periodogram of the detrended data for BGMRRI. The seasonal and its harmonics, marked by dashed line, are suppressed in this periodogram. The other plots give the periodograms for the detrended seasonally adjusted data for BSM, TRIG-1, TRIG-6, and X-12-ARIMA.

BGMRRI series. The movements in the BSM seasonal factors bear little resemblance to the data! Considerable power is left in the periodogram both at and near the seasonal frequency.

5.3 TRIG-6 More Flexible Than TRIG-1

Figure 5(a) and 5(c) show that TRIG-6 is much more flexible and removes more seasonality than either TRIG-1 or the BSM. This flexibility, however, is obtained at a significant increase in the roughness of the seasonal factors: see figure 5(b). Hence, in regards to seasonal adjustment, optimizing over additional parameters makes a significant difference.

A good example of the flexibility gained (but smoothness lost) by optimizing over all six variances is given by the BGMRRI series: see Figures 6 and 7. Often the seasonal factors progress in flexibility from BSM to TRIG-1 to TRIG-6. The

X-12-ARIMA seasonal looks like a smoothed version of the TRIG-6 one.

5.4 Which Would We Choose?

Choosing between the four seasonal adjustment procedures, we prefer X-12-ARIMA in 18 of the 29 series. In another 7 series, we have little or no preference between X-12-ARIMA and TRIG-6. The structural models are clearly preferred in only 4 series. These choices are based on a collection of diagnostic statistics and plots (see [BJ92a] and Bruce and Jurke (1992b)) and ignore stability issues.

We can classify our preferences of X-12-ARIMA into two types: those resembling the BMNCRS example and those resembling the IGLCTI example. With the BMNCRS series, X-12-ARIMA has as much "long term" flexibility as TRIG-6, but the seasonal factors are smoother. This holds true in 12 of the 29 of the series (B1, B2, B3, B5, B7, B10, B11, B13, I21, I23, I27, I29). With the IGLCTI series, X-12-ARIMA is distinctly more flexible and removes more seasonality but is as smooth as TRIG-6. This holds true in 6 series (C14, C20, I22, I24, I25, I28).

In another 7 series, we do not have a strong preference for either X-12-ARIMA or TRIG-6 (B6, B8, B9, B12, C18, C19, I26). In most of these series, though, the differences between X-12-ARIMA and TRIG-1 or BSM are significant.

In the remaining 4 series, X-12-ARIMA is either "too flexible" with no apparent advantage (B4), removes less seasonality (C15, C16), or has problems in its outlier treatment method (C17).

In general, X-12-ARIMA dominates for the business and inventory series. In series with little structure or linear changes in seasonal patterns, the differences between the methods is small. Both procedures can capture basically linear evolution in a seasonal cycle. The differences are the greatest in series which exhibit strong, non-linear changing seasonal patterns.

6 Stability in the Seasonal Adjustments

Stability in the seasonal adjustments is a crucial measure. In this study, we have chosen to study stability using the sliding spans statistics of Findley et al. (1990). Our results indicate that the structural models lead to more stable seasonal adjustments than X-12-ARIMA. This is to be expected since stability and flexibility are conflicting goals. However, part of the instability of X-12-ARIMA can be attributed to the discontinuous nature of its outlier/level shift detection scheme. We believe that the MING procedure for handling outliers and level shifts leads to more stable seasonal adjustments than X-12-ARIMA.

6.1 Sliding Spans Statistics

The sliding spans of Findley et al. (1990) are one measure of volatility involving adjustment of the data using four overlapping spans of approximately eight years in length. We focus on the stability of the month-to-month percentage change in the seasonal adjustments. For time t and span k, these are defined as

$$MM_t(k) = \frac{\exp A_t(k) - \exp A_{t-1}(k)}{\exp A_t(k)}$$

where $A_t(k)$ is the log transformed seasonally adjusted value at time t in span k. Using the notation of Findley et al. (1990), the sliding spans statistic for time t is

$$MM_t^{\max} = \max_k MM_t(k) - \min_k MM_t(k) \tag{12}$$

where k varies over those spans that contain both months t and $t-1$. Seasonal adjustments with more than 35% of the months with $MM_t^{\max} > 0.03$ are almost never acceptable. Good seasonal adjustments seem to have less than 15% greater than the cutoff. This criteria is designed for the X-11-ARIMA procedure and may not be suitable for the structural models.

6.2 X-12-ARIMA Generally Less Stable

We return to the IGLCTI series to illustrate the relative stability of X-12-ARIMA and the BSM fit. Note that the BSM fit is used for comparison in this case rather than TRIG-6. Figure 8 display the sliding span statistics $MM_t^{\max}$ for the IGLCTI series. The statistics for the X-12-ARIMA and BSM adjustments are displayed in the top and bottom plots respectively. The plots display $MM_t^{\max}$ over the period of sliding spans. As a benchmark to judge stability, a dashed horizontal line is drawn at 0.03. The histogram at the right of this plot shows the distribution of the $MM_t^{\max}$.

As figure 8 shows, the X-12-ARIMA adjustment is slightly less stable than the BSM adjustment. Some tradeoff between stability and flexibility is inevitable. Recall that the BSM seasonal factors are much less flexible than those of X-12-ARIMA for this series: see figures 2 and 3. In this case, the slight decrease in stability is probably a price worth paying for the significant increase in sensitivity. In [BJ92a], a similar example is given by the IFMETI series.

According to sliding spans statistics, the X-12-ARIMA seasonally adjusted data is less stable than MING with the BSM fit in 16 series (B1, B2, B7, B8, B9, B12, C14, C15, C18, C19, C20, I24, I25, I27, I28). X-12-ARIMA is significantly more stable for only two series, and only then because of problems with the initialization of the MING method (B10 and B11; see section 3.3).

Figure 9(a) gives scatterplots of the sliding span statistics $MM_t^{\max}$ for the seasonal adjustments of X-12-ARIMA (horizontal axis) and the BSM (vertical axis).

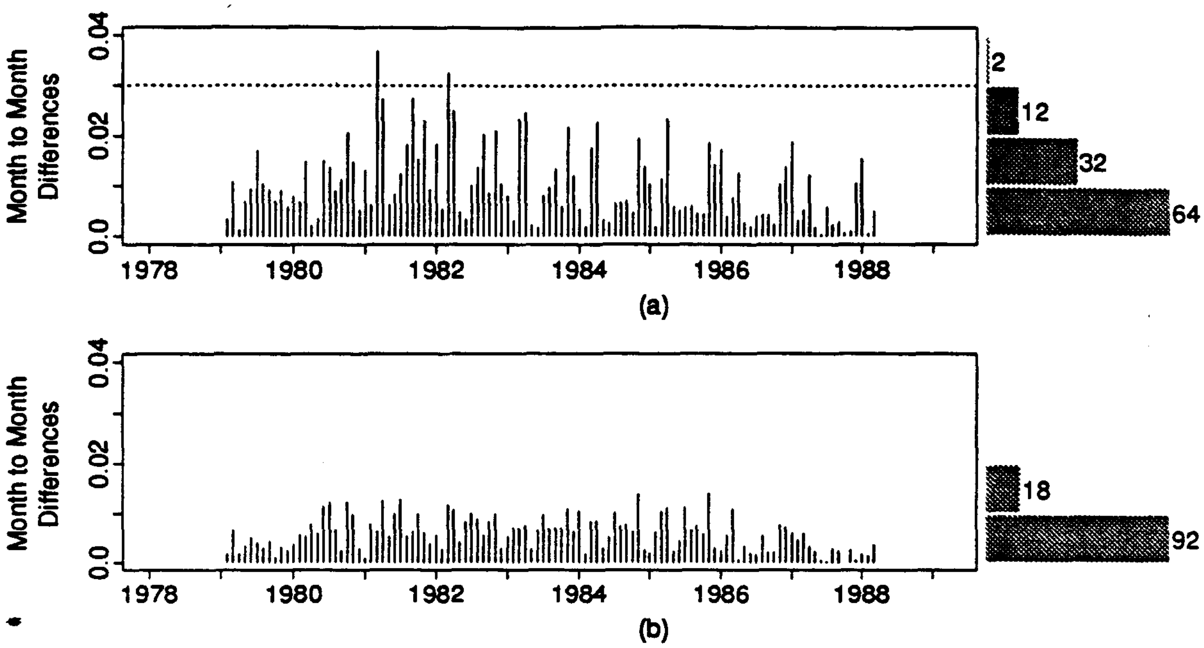

Figure 8: The sliding span statistics MM_t^{max} for the IGLCTI series for the (a) X-12-ARIMA and (b) BSM seasonal adjustments. As a benchmark to judge stability, a dashed horizontal line is drawn at 0.03. The histogram at the right of this plot shows the distribution of the MM_t^{max}.

The .03 significance lines are marked by dased lines. Clearly more points lie below the line, indicating that the BSM produces more stable adjustments. A density estimate of the difference

$$MM_t^{\Delta} \equiv (MM_t^{max} \text{ for X-12-ARIMA}) - (MM_t^{max} \text{ for BSM})$$

is given in the upper right corner. The 45° line indicates where $MM_t^{\Delta} = 0$. The median difference, marked by a line through the density estimate, is greater than zero.

We emphasize that difference in stability is, in large part, to be expected. It is not a major concern given the increased flexibility of X-12-ARIMA. Figure 9(b) gives scatterplots of MM_t^{max} for the seasonal adjustments of X-12-ARIMA (horizontal axis) and X-11-ARIMA (vertical axis). As indicated by the median lines, there is not much difference in the stabilities of X-12-ARIMA and X-11-ARIMA. Hence, most of the difference in stability we observe in figure 9(a) can be attributed to the difference in the flexibility of the seasonal adjustments.

Figure 9(b) does hint towards a problem: while the median stability is not much different, X-12-ARIMA tends to have far more large values of MM_t^{max}. We suspect this is due to the X-12-ARIMA outlier procedure. This is explored below.

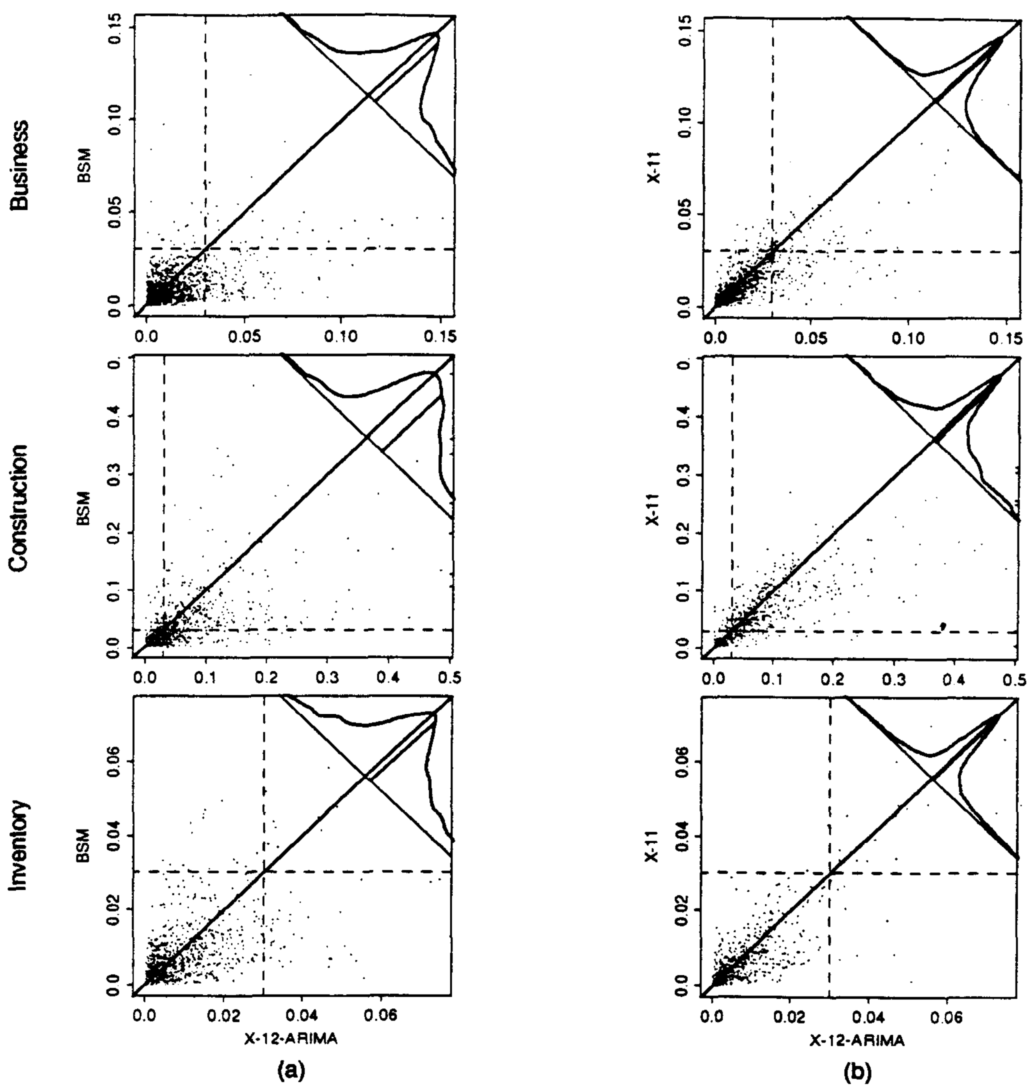

Figure 9: Scatterplots of the sliding span statistics $MM_t^{\max}$ for the seasonal adjustments of X-12-ARIMA (horizontal axis) and the BSM (vertical axis). The .03 significance lines are marked by dased lines. The density estimate in the upper right corner is of MM_t^{Δ}. The 45° line indicates where $MM_t^{\Delta} = 0$. The median difference, marked by a line through the density estimate, is greater than zero.

6.3 Instability of X-12-ARIMA Outlier Procedure

A more worrisome cause of instability in X-12-ARIMA seasonal adjustments is instability in its outlier identification procedure. The root of the problem lies in the discontinuous nature of the X-12-ARIMA method for handling outliers and level shifts (see section 2). An observation is either declared as an outlier, a level shift, or neither. There is no mechanism in the current procedure for smoothly transitioning between these different states. This naturally leads to instability in the seasonal adjustments.

By contrast, the MING procedure uses a continuous scheme: an observation is assigned a posterior probability of being an outlier, level shift, or neither. The probabilities are estimated from the data and range from 0 and 1. As a result, the MING's outlier procedure is less likely to cause instability in the seasonal adjustments.

An Example

To illustrate this we look at the BVARRS series. Figure 10 displays the sliding span statistics for the X-12-ARIMA adjustment BVARRS. The first plot shows the untransformed seasonally adjusted data obtained when the procedure is applied to each of four 8 year spans from 1978 to 1989. The second plot displays $MM_t^{\max}$ of (12) over this time period, as in figure 8. As a benchmark to judge stability, dashed horizontal lines are drawn at $k \times 0.01$ for $k = 3, 4, \ldots$. To pick up patterns of seasonal instability, boxplots of the $MM_t^{\max}$ by month are shown in the fourth plot. In contrast to the IGLCTI series, the X-12-ARIMA adjustment is borderline in terms of stability: about 15% of $MM_t^{\max} > .03$ with the maximum value at .08.

The third plot of figure 10 shows the outliers and level shifts detected by X-12-ARIMA in each span. One outlier was detected in each of the last three spans, as indicated by the points connected to the level by vertical lines. A small level shift was also detected in the second span, as indicated indicated by the "step" in the level. The size of the step or vertical line indicates the relative magnitude of the level shift or outlier. The source of the instability as reflected by $MM_t^{\max}$ is clear: different outlier/level shift combinations are detected in each of the four spans! Large values of $MM_t^{\max}$ can be directly associated with the different identifications of outliers and level shifts.

The sliding spans plot for the BSM fit is given in figure 11. The first, second, and fourth plots are the same as in figure 10. By contrast with X-12-ARIMA, the BSM adjustment is still extremely stable: no values of $MM_t^{\max}$ exceed .03.

The third plot of figure 11 shows the posterior probabilities of the outliers and level shifts for each span. The probability of a level shift at each point in the series is plotted as a vertical line extending downwards from the horizontal line. Similarly, outliers would be indicated by vertical lines extending upwards (none were detected). The length of the lines correspond the posterior probability of (with the dashed line being one). The identification of outliers and level shifts is very stable. The posterior probability of the level shifts ranges from fairly high in the second span to moderate in the fourth span. Evidently, the ability to assign a range of posterior probabilities avoids the problems encountered by X-12-ARIMA.

The BSPGWS series gives another example of the problems caused by the discontinuous nature of the X-12-ARIMA outlier identification scheme (see [BJ92a]).

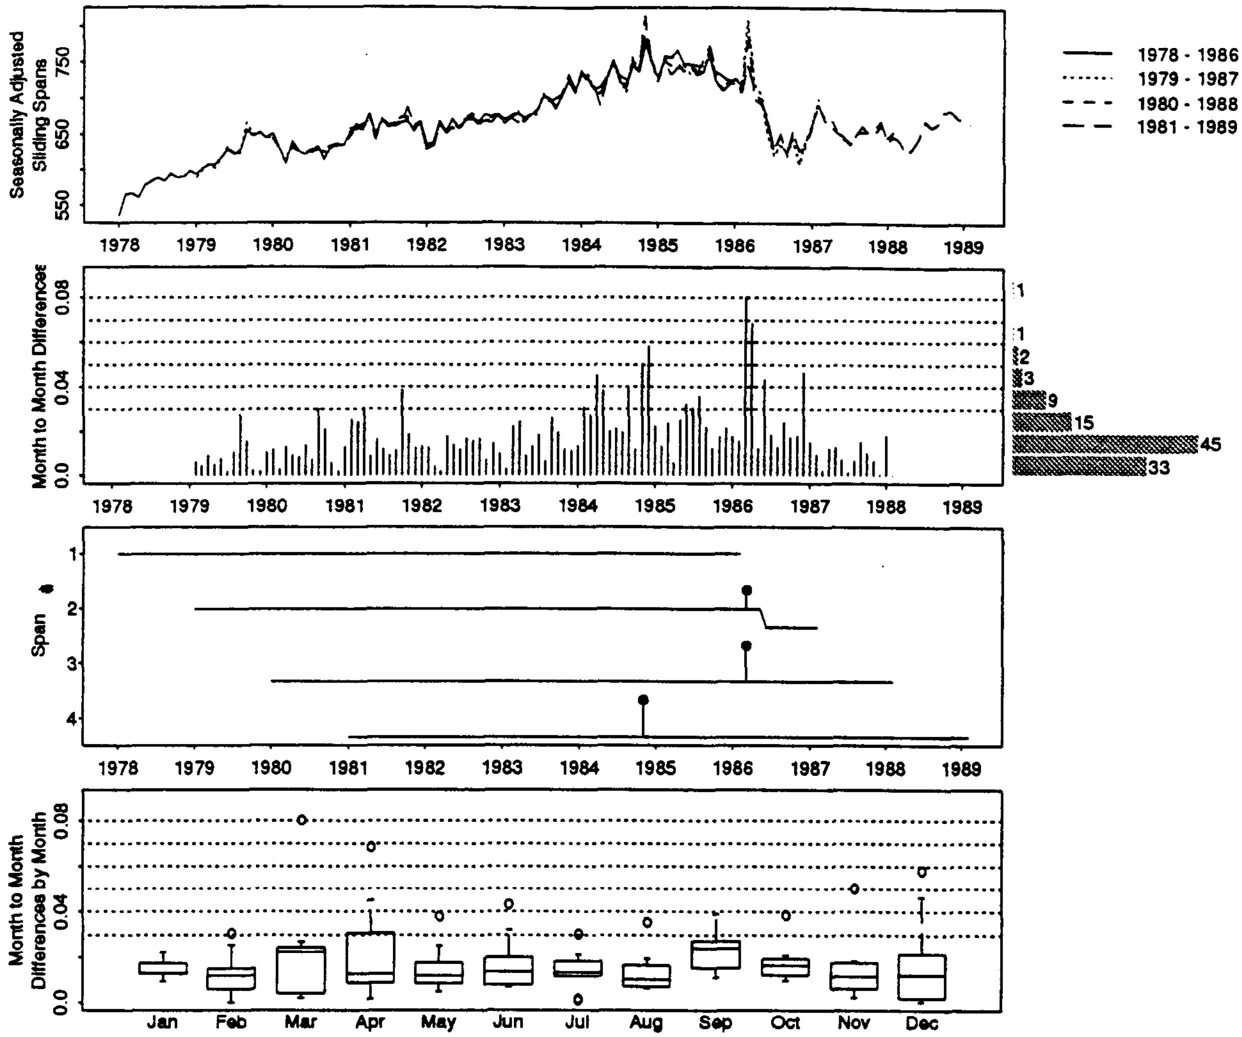

Figure 10: X-12-ARIMA sliding spans for BVARRS. The top plot shows the seasonally adjusted data for the four spans The second and fourth plots display the statistic $MM_t^{\max}$ over time and by months. The dashed horizontal line indicates the .03 cutoff: too many statistics bigger than .03 indicate an unstable adjustment. The third plot compares the outlier treatments over the spans. Outliers are indicated by points and level shifts by steps.

Some General Conclusions

X-12-ARIMA outlier identification procedure is less stable than the MING method for 13 series (B1, B2, B8, B9, B12, B13, C15, C17, C18, C19, C20, I25, and I27). In all but one of these series, this instability appears to lead to significantly less stable seasonal adjustments (as measured by the diagnostic $MM_t^{\max}$).

The outlier procedure of MING is less stable for 6 series (B6, B10, B11, C16, I24, I29). Three of these (B10, B11, and I29) are due to a problem with the initial-

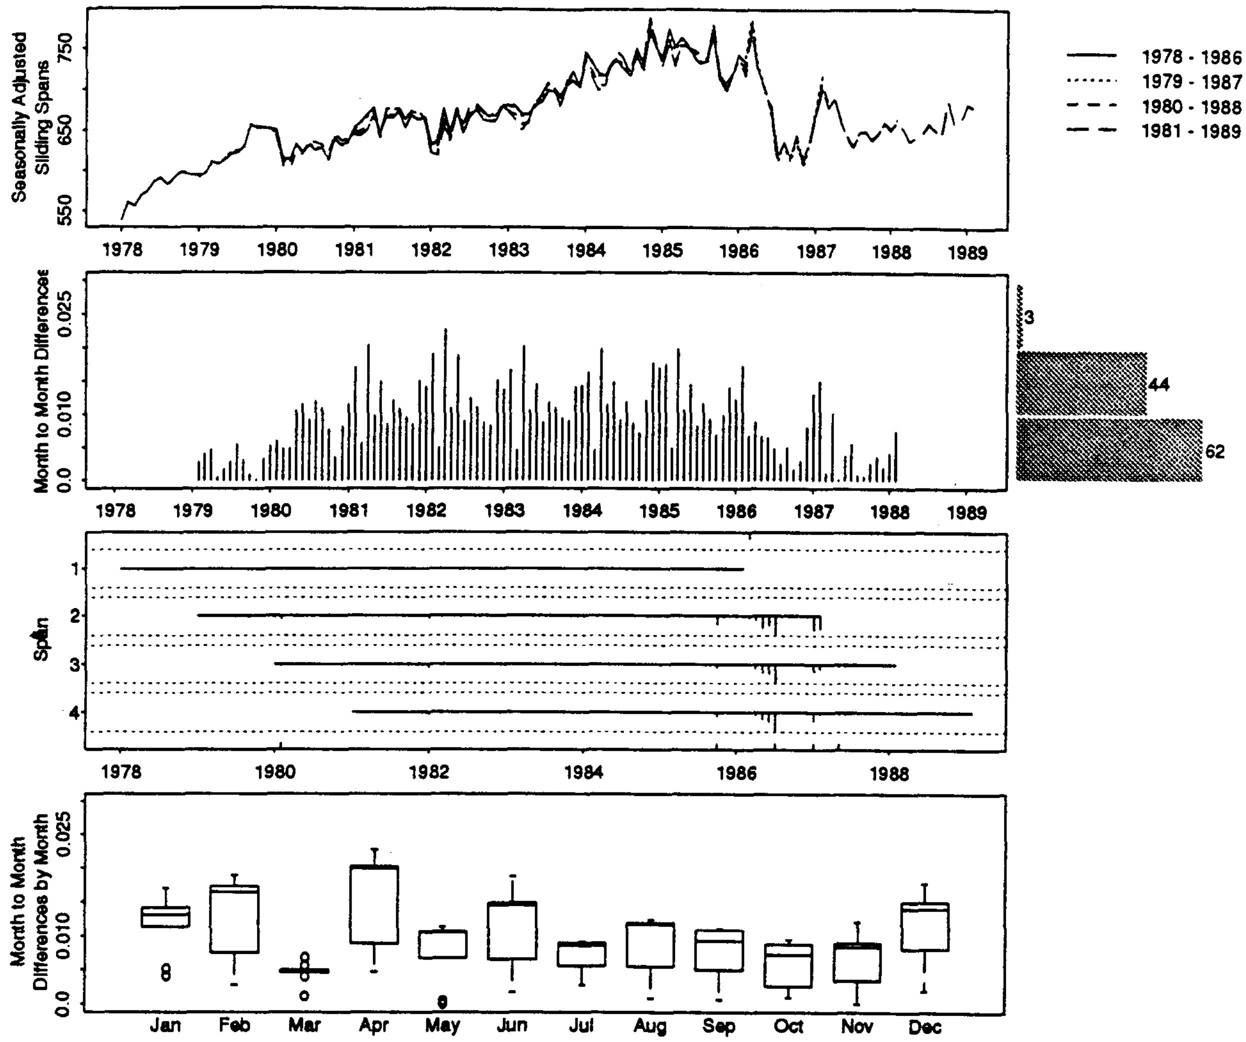

Figure 11: BSM sliding spans for BVARRS. The top plot shows the seasonally adjusted data for the four spans The second and fourth plots display the statistic MM_t^{max} over time and by months. The dashed horizontal lines indicate the .03 cutoff: too many statistics bigger than .03 indicate an unstable adjustment. The third plot compares the outlier treatments over the spans. Outliers are indicated by points and level shifts by steps.

ization procedure (see below). In two others (B6 and I24), MING estimates several additional level shifts with moderate or low probability in one span. This has very little effect on the seasonal adjustments. For CNE1HS, an outlier is only identified in two spans by MING, leading to quite a large value of MM_t^{max} for that month (this it the type of instability which is more typical of X-12-ARIMA).

It seems reasonably safe to conclude that the X-12-ARIMA outlier identification procedure is less stable. Furthermore, the instability of X-12-ARIMA has a greater effect on the instability of the seasonal adjustments. These results all point to problems with the discontinuous outlier detection method of X-12-ARIMA. Similar results should hold in adding new observations or revising observations in a time series.

7 The Need for Graphics

We started this study with the aim of proving the superiority of the MING procedure for dealing with series laden with outliers and level shifts. We had to change course as we were confronted with what we perceived as problems in MLE/structural model based seasonal adjustment. In hindsight, these problems may seem obvious for the series studied. However, they only became apparent after perusing through literally hundreds of plots (see Bruce and Jurke (1992b)). Without these plots, it is quite likely that we would have reached very different conclusions.

A fixed set of numerical diagnostics can not be relied upon to uncover all potential problems in a seasonal adjustment. Looking at plots is vital for uncovering unusual features. With the advent of the modern computer workstation, perusal of many plots is no longer the burden it once was. With the right computing environment, it is now feasible to rapidly validate seasonal adjustments using both numerical and graphical diagnostics.

7.1 Plot Driven Diagnostics

Initially, our comparison was based on the diagnostics similar to those proposed by den Butter and Mourik (1990). The conclusions we reached in section 5 are not supported by these diagnostics! It turns that these diagnostics are not sensitive enough to distinguish flexibility from roughness in seasonal factors, to show how well the seasonal factors fit the data, or to accurately determine the amount of seasonality remaining in the data.

The diagnostics SEAS FLEX, SEAS ROUGH, and SEAS REMAIN were all constructed after we had formed preliminary conclusions through the plots. In particular, we found the SSI-Plot (as in figure 2) and a plot of the periodograms (as in figure 3) especially valuable.

While we looked at many other diagnostics (see [BJ92a]), we did not find these particularly revealing or enlightening in comparing X-12-ARIMA with the structural models. Hence, we have have not reported these here.

7.2 Inadequacy of Simple Diagnostics

The simple one-way ANOVA test for residual seasonality, while widespread in use, has little power and can be misleading. For the IFMETI series, the test has a "*p*-

value" of essentially 1 for all structural models, indicating no seasonality remaining in the residuals. On the other hand, the p-value for X-12-ARIMA is 0.761. On the surface, this would indicate that the structural models are quite adequate, and do a better job of removing seasonality from the data. Examination of the periodograms tells otherwise (see figures 19 and 20 of [BJ92a]): X-12-ARIMA seasonal factors adapt to the data in an appealing manner and considerable power is reduced in the periodogram around the seasonal frequency.

The results also indicate that a better fitting model, according to AIC, doesn't mean a more appealing seasonal factor. For example, compared with TRIG-1 or BSM, TRIG-6 has a much lower AIC value for the ITVRUO series. However, examination of the seasonal factors with the SSI Plot does not show a strong preference for TRIG-6. The ITVRUO series is difficult to fit, and in that sense is atypical (the series undergoes a variance shift in the latter portion). A more typical example is given by BLQRRS, for which TRIG-6 has a slightly higher AIC value but significantly more flexible seasonal factors.

8 Other Results

The MING procedure is intrinsically more stable: see section 6. Beyond the question of stability, though, the different ways in which MING and X-12-ARIMA handle outliers and level shifts does not have a big impact on the seasonal adjustments. Whatever impact it does have is swamped by the fundamental differences in the decompositions (see section 5). Nonetheless, we did uncover several interesting results which are discussed below.

8.1 Advantages of the MING outlier method

For many of the series, MING detects fewer outliers/level shifts or the same number with lower probability (B1, B2, B4, B5, B7, B9, B13, C16, C17, C19, I21, I23, I26, I27, I28). For example, for the IFATTI series, X-12-ARIMA detects eight outliers and ten level shifts By contrast, TRIG-6 detects only 3 moderate probability level shifts and 5 low probability level shifts. This is a reflection of the rather arbitrary level at which the outlier threshold is set for the X-12-ARIMA procedure. Setting it to a higher level would obviously reduce the number of series for which X-12-ARIMA detects more outliers/level shifts.

In some series, MING detects numerous low probability level shifts or outliers not identified by X-12-ARIMA (B3, B6, B8, B12, C14, C15, C20, I22, I24, I25, I29). For example, MING picks up several small level shifts for ICMETI not identified by X-12-ARIMA. These level shifts are barely visible in a plot of the seasonally adjusted data. For the series BLQRRS, MING models a "ramp" using two successive level shifts of moderate to high probability while X-12-ARIMA uses one large level shift.

The MING procedure has two apparent advantages. First, it has an automatic way to adapt the "cutoff" level based on the likelihood. Second, it can incorporate small level shifts or outliers by giving them low probability. However, for all of the series listed above, the two procedures handle major outliers and level shifts in a similar manner. The only difference for these series is the way in which the methods handle small outliers and level shifts. The advantages advantages offered by MING are more theoretical than practical.

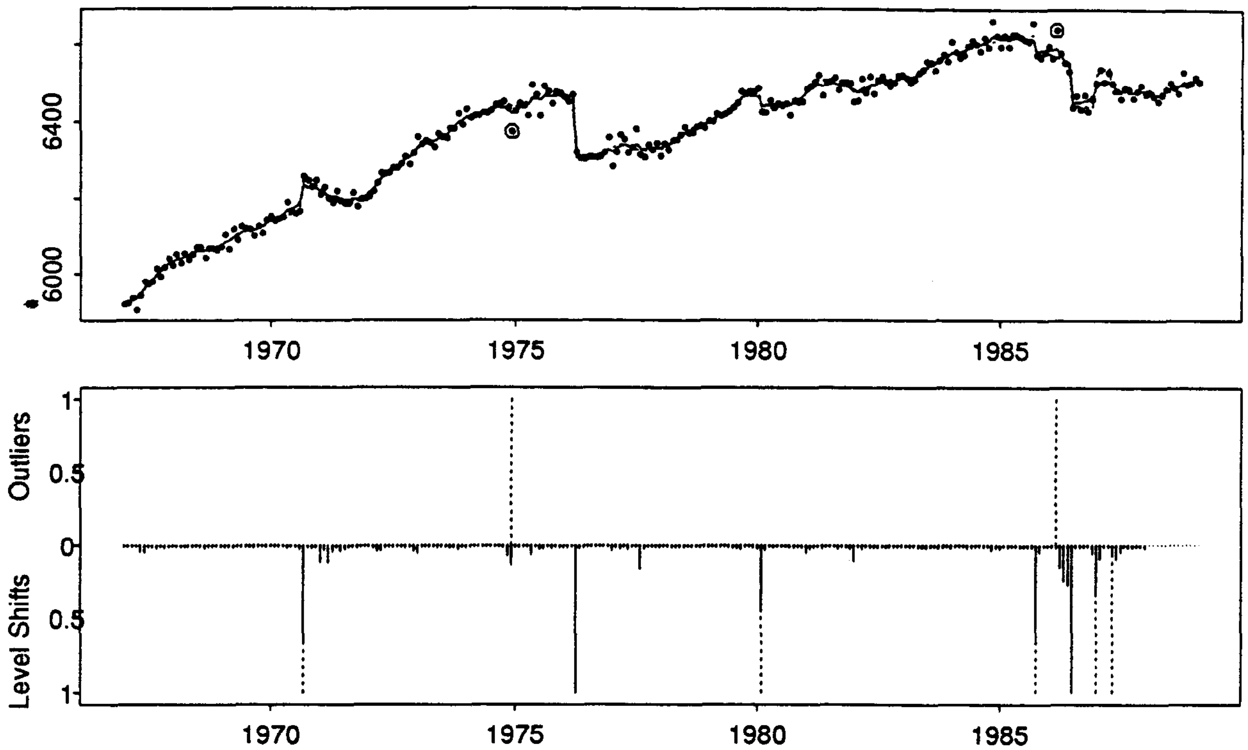

Figure 12: The top plot shows the TRIG-6 seasonally adjusted data for BVARRS (points) with the TRIG-6 trend (solid line) and X-12-ARIMA trend (dashed line). Two outliers identified by X-12-ARIMA are circled. The bottom plot shows the posterior probability of an observation being an outlier (solid upward lines) or a level shift (solid downward lines) as estimated by TRIG-6. Also shown are outliers (dashed upward lines) and level shifts (dashed downward lines) as identified by X-12-ARIMA.

8.2 Problems with MING

Figure 12 illustrates a common problem with the MING procedure. The top plot shows the TRIG-6 seasonally adjusted data for BVARRS (points) with the TRIG-6 trend (solid line) and X-12-ARIMA trend (dashed line). Two outliers identified by X-12-ARIMA are circled. The bottom plot shows the posterior probability of an observation being an outlier (solid upward lines) or a level shift (solid downward lines) as estimated by TRIG-6. Also shown are outliers (dashed upward lines) and level shifts (dashed downward lines) as identified by X-12-ARIMA.

Both X-12-ARIMA and MING pick up several major level shifts, but only the X-12-ARIMA procedure identifies outliers at 12/74 and 3/86. The plot of the seasonal adjusted data strongly supports this: 12/74 and 3/86 stand out as outliers. MING fails to detect the "obvious" in this case.

This an illustration a general problem with the MING procedure. In several series, MING has a problem detecting moderate outliers if level shifts are present (B3, B12, I21, I23, I24, I26). When the series has mostly level shifts with a couple of outliers, the estimated posterior probability of the occurrence of an outlier tends to be very small. The converse holds as well: when a series has many outliers and a single level shift, the estimated posterior probability of a level shift is small (B1).

In the BVARRS series, where there are several obvious level shifts and two moderate outliers, MING estimates a relatively high prior probability to level shifts ($\epsilon_\eta > .01$). These give a small degree of protection against the two outliers in this series (by increasing the variance). By giving a nearly zero probability to outliers ($\epsilon_I \approx 0$) the likelihood is optimized since it does not incur the "penalty" of modeling outliers when there aren't any (as is the case for all but two observations).

The root of the problem is the decision to optimize of the prior probabilities ϵ_I and ϵ_η. Alternative approaches, such as constraining the parameters, may produce better behavior in this regard.

8.3 Local level shift or an outlier patch?

X-12-ARIMA and MING occasionally use very different approaches to modeling non-Gaussian behavior in a number of series (B3, B10, B11, C18, I24, I29). Whereas the MING tends to treat these series using one or more level shifts, X-12-ARIMA identifies an outlier "patch". An example of this is given by the BGMRRI series (see [BJ92a]). In general, when X-12-ARIMA and MING differ in this way, neither method is demonstrably superior.

8.4 Difference in outlier treatments not as important

While MING and X-12-ARIMA lead to quite different seasonal adjustments, the difference in the outlier procedures is *usually* only a second order contributing effect to the difference in seasonal adjustments. The difference in the seasonally adjusted data is primarily due to the different estimation methods of the seasonal factors (see section 5). See [BJ92a] for an example based on the IFMETI series illustrating that the outlier treatments can be substantially different without much effect on the seasonal factors.

In a few cases, the outlier treatments lead to quite notable differences in the seasonal adjustments (B1, B3, B6, B10, B12, C16, C17, I23). For CNETHS, the X-12-ARIMA procedure identifies 5 outliers in January and 4 in February. TRIG-6 also identifies most of these outliers, but generally with probability less than one. As a result of the outlier identifications, the X-12-ARIMA seasonal factors for January are elevated. According to the periodogram, there is considerable seasonality remaining in the X-12-ARIMA decomposition. For this series, it would seem that exceptionally low January values are part of the seasonal effect. The outlier procedure of X-12-ARIMA is perhaps adjusting too much for these values.

Another example is given by the series BVARRS. Recall that MING has difficulty in picking up a fairly major outlier in this series at time 3/86 (see the above discussion and figure 12). This outlier seems to have leaked into the seasonal pattern for the MING decomposition (see figure 30 of [BJ92a]). A less dramatic, and more typical, example is given by the BAUTRS series. Some of the largest differences between the seasonal adjustment are at times in which the outlier treatments are different. However, the difference in the seasonal adjustments for BAUTRS are still mainly due to the different seasonal factors.

For seasonal adjustment, the crucial thing is to deal with large or moderately large outliers and level shifts in an adequate manner. It is not crucial to handle small outliers or level shifts. The way in which the procedure deals with non-Gaussian behavior is not especially important: e.g., either local level shifts or an outlier patch may suffice.

8.5 X-12-ARIMA trends are smoother

X-12-ARIMA has smoother trends in all but 3 series (the exceptions are B4, B7, I22). The decompositions of IGLCTI, given in figure 1, give a typical example of this. The X-12-ARIMA trends are visually more appealing. This is a well known feature of the time series structural model seasonal decomposition. Harvey and Valls Pereira (1989) defend the rough trends yielded by structural models.

Smoother trends for structural models can be obtained simply by constraining the variances (see section 9.1). Including a local AR trend in the STM's may also produce smoother trends, although the fitting would become more difficult and perhaps unstable.

9 More on MING

We also pursued several other extensions and research issues regarding the MING procedure for seasonal adjustment. These include constraining the variances to obtain an adjustment closer to X-12-ARIMA, estimation of standard errors, modeling calendar effects, and incorporation of different outlier models.

9.1 Constraining the Variances

The results of section 5 indicate that maximum likelihood estimation of structural models does not necessarily produce very good seasonal adjustments. One approach to improving the seasonal adjustments for the structural models is to fit a *constrained* version of the model. The basic idea is to find a model which closely mimics the X-11 filters.

Let Y_t denote a time series for which the default filters of X-11 are "optimal" in a mean square error sense. Maravall (1985) derives constraints on the parameters of the basic structural model (BSM) so that its autocorrelation function (*acf*) closely matches the *acf* of $\Delta\Delta^{12}Y_t$. These parameters are given by (Maravall (1985); Harvey and Valls Pereira (1989)):

$$\sigma_\epsilon^2 = \sigma^2 \quad \sigma_\eta^2 = 0.133\sigma^2 \quad \sigma_\xi^2 = 0.167\sigma^2 \quad \sigma_\omega^2 = 0.067\sigma^2 \tag{13}$$

To fit this model, a single variance σ^2 is optimized. We call this constrained fit BSM-CONS.

We fit BSM-CONS to the 29 series to see if we could obtain seasonal adjustments which mimic those of X-12-ARIMA. The result was a resounding no: for the business and inventory series, the seasonal factors are only slightly more flexible and rougher than the origianl BSM fits For the construction series, non-default filters were used for X-12-ARIMA. Hence, the BSM-CONS adjustments are very poor and not comparable for the construction series. See [BJ92a] for details.

Maravall (1985) acknowledges that equivalent *acf*s do not translate into equivalent decompositions. Indeed, the unobserved components ARIMA model which yields a decomposition similar to that of X-11 involves more MA terms than BSM-CONS for both the seasonal and trend components. As a result, X-12-ARIMA produces smoother seasonal and trend components than BSM-CONS.

9.2 Estimation of standard errors

A big advantage of a model based procedure over X-12-ARIMA is in the availability of standard errors for the seasonally adjusted data. MING is especially good in this regard, since it incorporates outliers and structural changes within the model. MING actually produces an estimate of the posterior density, not just standard errors. Kitagawa (1987) gives several nice examples of the advantages of non-Gaussian confidence intervals. See Kitagawa (1988) for examples in the context of seasonal adjustment.

Figure 13 gives the width of the 99% confidence intervals for the seasonal adjustments for the BVARRS series. The intervals widen towards the ends of the series, reflecting "end effects". The intervals also tend to get wider near level shifts and outliers. However, the greatest uncertainty is *not* near the large level shift in 1976,

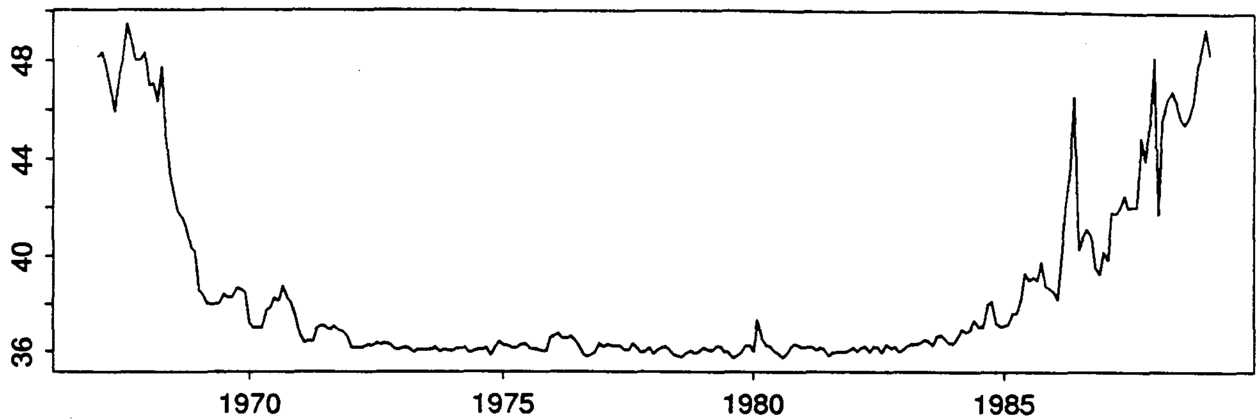

Figure 13: Width of the 99% confidence intervals for the BSM seasonally adjusted BVARRS data. The intervals widen towards the ends of the series and near level shifts and outliers. However, there is not much uncertainty near the large level shift in 1976, which is easy to identify and model (see figure 12).

which is easy to identify and model (see figure 12). Rather, it is near the series of smaller level shifts in 1986.

One could not expect the parametric outlier identification procedure of X-12-ARIMA to perform as well in this regard. This is because the standard errors are calculated under the assumption that the location of the outlier or level shift is *known*. This can make a big difference, since the timing of a local level shift or outlier patch is often in doubt, especially towards the ends of the series. In addition, the intervals produced by the ARIMA outlier identification procedure are purely Gaussian, and cannot capture the long tailed nature of the densities.

9.3 Modeling calendar effects

In the fits done for MING, trading day and Easter effects were handled by prior adjustment based on X-12-ARIMA. Optimizing over trading day and Easter regression variables is not very critical, and is unlikely to lead to significantly different results. This is mainly because both MING and the ARIMA model underlying X-12-ARIMA give reasonable fits to the data and adequately deal with outliers and level shifts. Hence, fitting a fixed effects regression variable such as for trading day should be roughly equivalent with either procedure. See [BJ92a] for an example.

9.4 Building in Ramps and Other Outlier Models

For the sake of parsimony, simplicity, and computational efficiency, the model used to fit the series only accommodated level shifts. A slight generalization of this model can be obtained by inflating both the variances of η_t and ξ_t in the second component of the Gaussian mixture model (7):

$$\begin{pmatrix} \eta_t \\ \xi_t \end{pmatrix} \sim \begin{cases} \mathrm{N}\left(\mathbf{0}, \begin{pmatrix} \sigma_\eta^2 & 0 \\ 0 & \sigma_\xi^2 \end{pmatrix}\right) & \text{with probability } 1 - \epsilon_{\eta,\xi} \\ \mathrm{N}\left(\mathbf{0}, \begin{pmatrix} \tilde{\sigma}_\eta^2 & 0 \\ 0 & \tilde{\sigma}_\xi^2 \end{pmatrix}\right) & \text{with probability } \epsilon_{\eta,\xi} \end{cases} .$$

This more general "ramp" yields very similar estimates for the trend or seasonal in a representative subset of eight series. The series experiencing the largest change in the estimated trend is IFMETI, for which the maximum difference is only ±0.5%.

The main benefit from including ramps may be to improve the overall fit from the model: see [BJ92a].

9.5 Handling Doublets

Often economic data has two or three adjacent aberrant values, which we call a "doublet" or "triplet". These are due to strikes, weather, or any condition which has a temporary effect on the economy. For example, the CMW1HS series has a doublet at 1/79 and 2/79, presumably caused by unusually cold weather. MING handles this patch using a combination of outliers and level shifts. This results in an "unnatural" trend which chases after the peaks and valleys.

This behavior of MING stems from shortcomings in the model: the occurrence of an outlier is assumed to be independent of whether an outlier occurred at the previous observation. This is counter to what we know about economic (and many other) time series: outliers often come in patches. Indeed, the number of outliers in a patch often depends on the sampling interval. A sensible generalization of the outlier model is to allow Markov behavior in the outlier generating process.

Let Z_t be a 0-1 process which indicates whether an outlier has occurred at time t. We assumed in section 3.2 that $p(Z_t = 1|Z_1, \ldots, Z_{t-1}) = p(Z_t = 1) = \epsilon_\mathrm{I}$. A more natural assumption is

$$p(Z_t = 1|Z_1, \ldots, Z_{t-1}) = p(Z_t = 1|Z_{t-1}) = \begin{cases} \epsilon_\mathrm{I} & \text{if } Z_t = 0 \\ \epsilon_\mathrm{I}^0 & \text{if } Z_t = 1 \end{cases} \qquad (14)$$

where $\epsilon_\mathrm{I}^0 \gg \epsilon_\mathrm{I}$. Hence, if an outlier occurs at time $t-1$, then an outlier is much more likely to occur at time t.

Figure 14 gives a plot of the seasonally adjusted data for CMW1HS. The trends are compared for the "doublet" outlier model (14) and the standard outlier model (6). The trend for the standard model (dashed line) chases after the outlier pair in 1979. The trend for the doublet model (solid line) completely excludes the outlier pair. In addition, although masked in figure 14, the trend for the doublet model smooths several sharp peaks prominent in the original trend.

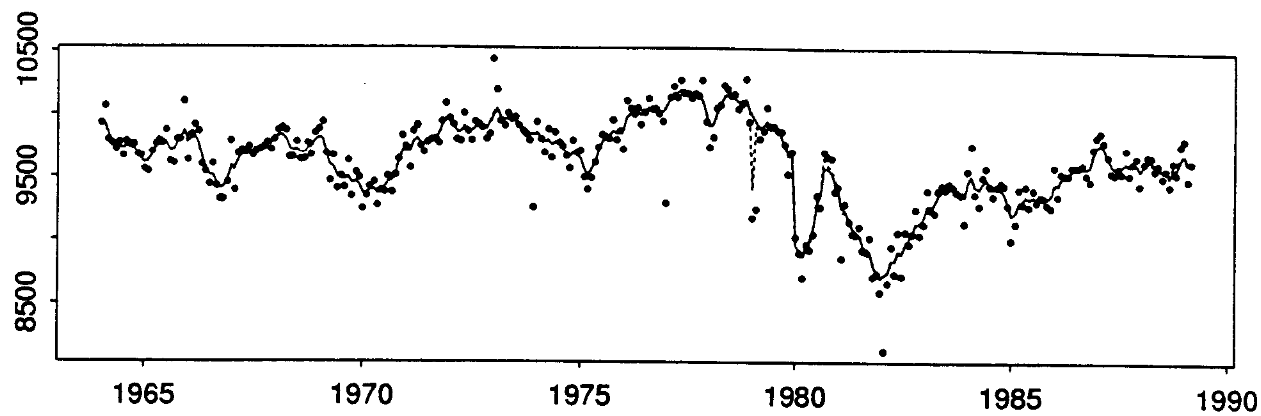

Figure 14: A plot of the seasonally adjusted data for CMW1HS. The trends are compared for the "doublet" outlier model (solid line) and the standard outlier model (dashed line). The doublet outlier model interpolates through the outlier patch while the standard model produces a very unnatural looking trend.

For this example, ϵ_I^0 is set to 0.25. This prior was choosen based on empirical experimentation. Smaller values ($\epsilon_I^0 < 0.1$) tended to leave a small spike in the trend. Much larger prior values would be likely to drive the estimate ϵ_I to zero.

For the CMW1HS series, X-12-ARIMA also has a very unappealing looking trend which chases after the doublet. The problem is that X-12-ARIMA models the outlier pair using two adjacent level shifts. This is an artifact of the "one-at-a-time" outlier and level shift fitting approach. Fortunately, an *ad hoc* solution exists to make the X-12-ARIMA trends more appealing: search for adjacent or near adjacent level shifts of opposite sign, and replace them with an outlier patch.

10 Conclusions

On one level, this study can be viewed as an endorsement of X-12-ARIMA. The procedure adequately handles most of the series with both outliers and structural changes. The decompositions would appear to be more appealing than those generated by a structural model based method.

However, X-12-ARIMA has some significant shortcomings, such as the discontinuous nature of the outlier identification procedure. In addition, a procedure such as MING offers several potential advantages, including estimates of standard errors, generalization to multivariate seasonal adjustment, and an appealing underlying methodology. Hence, we should not give up on alternatives to X-12-ARIMA.

10.1 Current Research

We are currently extending and refining the MING program. These include the following developments:

1. MING is relatively slow. Several changes could be made to speed up the likelihood evaluations. In particular, use of a different criteria to determining which densities should be collapsed could lead to substantial improvements. More dramatic computation savings could be achieved if an adaptive tree approach, as in Bruce and Martin (1992), is adopted.

2. The initialization method used for MING is clearly inadequate, as exhibited by several examples. One possible solution is to use an "EM approach", estimating the posterior probabilities of outliers and level shifts in the beginning of the series using a backwards filter. This method is relatively time consuming to program. A simpler, but computationally expensive, approach is to estimate the initial conditions (see De Jong (1988) for the Gaussian case).

3. Only one aspect of stability was examined in this paper. We are now examining the stability of the methods in regards to adding new observations on to the ends of the series.

We plan to have an initial public domain release of the MING program by the end of January 1993. It is based on S-PLUS (1991) and will be available through the Statlib software library.

10.2 Open Problems

We feel this study has brought to attention a number of interesting research areas. Some of these include:

Seasonal Models: The seasonal model used in the BSM is not sufficiently flexible for the purposes of seasonal adjustment. While the seasonal trigonometric model TRIG-6 is flexible enough, it does not produce sufficiently smooth seasonal factors when fit by maximum likelihood. The assumptions underlying seasonal models need to be investigated and alternative models should be explored: see, for example, Hannon et al. (1970); Harvey and Valls Pereira (1989); Young (1988); Young and Ng (1990).

Alternatives to MLE: The seasonal decompositions for structural models in this paper are based on the maximum likelihood estimates of the models discussed in section 3.1. While MLE's have good statistical properties, they did not yield particularly good seasonal adjustments for the series studied. Since time series structural models tend to have flat likelihood surfaces, it may be possible to obtain more appealing seasonal factors by constraining the model parameters without sacrificing much in terms of goodness of fit.

In section 9.1, we constrained the parameters so that the adjustment for the BSM would more closely mimic that of X-12-ARIMA. However, the results were not very good, perhaps due to the inadequacy of the BSM seasonal model. A more promising approach, adopted by Young (1988); Young and Ng (1990), is to use a rich trigonometric seasonal model and choose parameters based on spectral considerations.

Local AR Trend: Inclusion of a local AR component may help in a number of ways. By sopping up local variability, it may cause both the trend and the seasonal term to be smoother. Note that inclusion of an AR term will involve more difficult optimizations.

Outlier Models: Only a couple of possible outlier models were used in this study. It is worthwhile investigating whether more complex models offer any significant improvement. In particular, experimentation needs to be done in terms of modeling seasonal breaks. Several of the series (e.g., BGMRRI) seemed to exhibit a seasonal break, and this appeared to cause problems in fitting the models.

The MING procedure had difficulty in detecting one or two outliers when multiple level shifts were present in a series (see section 8.2). This could be solved by constraining the prior probabilities to be greater than a certain value.

General Non-Gaussian Models: We have focused on some very specific non-Gaussian disturbances which are commonly observed in practice. Our core model, however, is always assumed to be Gaussian. An fruitful line of inquiry would be the explore non-Gaussian core models.

The MING program can handle one class of such models: a Gaussian mixture distribution can be used for the core process in addition to or in place of the outlier process.

Robust Initial Parameters: The initial parameters for the optimization of MING were obtained by first running X-12-ARIMA and then running REGCMPNT. For MING to be useful as a stand-alone routine, a good robust but fast method needs to be developed to estimate initial parameters for the optimizer.

Better Parameterizations: Fitting time series structural models by maximum likelihood is often a difficult task due to the flatness of the likelihood. It may be possible to "tune" the optimizer to obtain successful convergence. A more fundamental solution to this problems lies in finding alternative transformations of the parameters for easier optimization.

A Fitting Details for MING

The likelihoods are maximized using the quasi-Newton nonlinear optimizer of Gay (1979) (see also Dennis et al. (1981)). The optimizer uses a trust region approach with a double dogleg step. Finite difference gradients are used with the BFGS secant update to the hessian.

The initial values to the optimizer for the BSM are obtained by fitting the BSM Gaussian model. To ensure robust initial estimates, the outlier identification scheme of X-12-ARIMA are used to first identify AO's and LS's. These are included in the model as fixed regression effects. The fits are done using the program REGCMPNT (Monsell and Otto (1991)), which is more efficient than MING for purely Gaussian models. The TRIG-1 model was fit with the initial values derived from the maximum likelihood estimates for the BSM. The TRIG-6 model was fit with the initial values derived from the maximum likelihood estimates for TRIG-1.

Prior adjustment is done for trading days based on the REGCMPNT procedure. While MING accommodates fitting trading day variables, this involves nonlinear optimization over six parameters, greatly increasing the computations. Some examples indicate that further refinement of the estimates for trading day effect is not important (see section 9.3).

Convergence Criteria

The optimizer is considered to have converged successfully if little improvement has been achieved in the objective function from the previous iteration. This is known as "relative function convergence", and is satisfied if

$$\frac{|\log \mathrm{L}(j) - \log \mathrm{L}(j-1)|}{|\log \mathrm{L}(j)| + |\log \mathrm{L}(j-1)|} \leq 0.00005 \tag{15}$$

where $\log \mathrm{L}(j)$ is the log-likelihood on the j-th iteration. Alternatively, the optimizer converges if the change in the estimated parameters is small. This is known as "relative X-convergence" and is satisfied if

$$max_{i=1,\dots,p} \left\{ \frac{|\hat{\alpha}_i(j) - \hat{\alpha}_i(j-1)|}{|\hat{\alpha}_i(j)| + |\hat{\alpha}_i(j-1)|} \right\} \leq 0.005. \tag{16}$$

where $\hat{\alpha}$ is a vector of the scaled parameters.

B Description of detrending procedures

Detrending Procedure for the SSI Plot

Let T_t, O_t, and L_t, be the trend, outlier, and level shift components of the X-12-ARIMA decomposition. To obtain the points in the SSI-plots, we detrend the data as follows:

1. Subtract the level shifts and outliers from the trend: $T_t^0 = T_t - O_t - L_t$.

2. Smooth T_t^0 to obtain a smooth trend:

$$T_t^1 = \sum_{j=-11}^{11} \left(\frac{12 - |j|}{144} \right) T_t^0$$

3. Add back in the level shifts to the smoothed trend to obtain our detrending sequence $\tilde{T}_t = T_t^1 + L_t$.

4. Subtract $\tilde{T}_t$ from the data to obtain the observations plotted: $\tilde{Y}_t = Y_t - \tilde{T}_t$.

This above method was adopted since we are comparing several decompositions with possibly sharp breaks in the trend.

Detrending Procedure for the Periodogram Plot

The periodogram of the data is based on the $\tilde{Y}_t$ as computed above for the SSI Plot.

The periodogram for X-12-ARIMA is based on the detrended seasonally adjusted data $\tilde{A}_t$. This is obtained as above except we also subtract the seasonal (and calendar component if present):

$$\tilde{A}_t = Y_t - S_t - \tilde{T}_t$$

The periodograms for the structural models BSM, TRIG-1, and TRIG-6 are similarly obtained, except that the appropriate trend is substituted for the X-12-ARIMA trend in step 1) above.

References

Alspach, D. L. and Sorenson, H. W. (1972). Nonlinear Bayesian estimation using Gaussian sum approximations. *IEEE Transactions on Automatic Control*, AC-17:439–448.

Ansley, C. F. and Kohn, R. (1985). Estimation, filtering, and smoothing in state space models with incompletely specified initial conditions. *The Annals of Statistics*, 13(4):1286–1316.

Bell, W. (1986). A computer program for detecting outliers in time series. In *American Statistical Association Annual Conference.*

Bell, W. and Hillmer, S. (1983). Modeling time series with calendar variations. *Journal of the American Statistical Association*, 78:526–534.

Bell, W. and Hillmer, S. (1987). Initializing the Kalman filter in the nonstationary case. Research Report CENSUS/SRD/RR-87/33, Statistical Research Division, Bureau of the Census, Washington, D.C. 20233.

Bell, W. R. and Hillmer, S. C. (1984). Issues involved with the seaosnal adjustment of economic time series. *Journal of Business and Economic Statistics*, 2.

Box, G. E. P., Hillmer, S. C., and Tiao, G. C. (1978). Analysis and Modeling of Seasonal Time Series. In Zellner, A., editor, *Seasonal Analysis of Economic Time Series*, pages 309–334. U. S. Bureau of the Census, Washington, D. C.

Bruce, A. G. and Cordera, M. (1992). Notes on Two Filter Smoother Formulae. Technical report, Statistical Sciences, Inc., 1700 Westlake Ave N., Seattle, WA 98109. In preperation.

Bruce, A. G. and Jurke, S. (1992a). Empirical comparison of two methods for non-Gaussian seasonal adjustment. Research Report CENSUS/SRD/RR-92/11, Statistical Research Division, Bureau of the Census, Washington, D.C. 20233.

Bruce, A. G. and Jurke, S. (1992b). Empirical comparison of two methods for non-Gaussian seasonal adjustment: a book of plots. Technical report, Victoria University, Wellington, New Zealand.

Bruce, A. G. and Martin, R. D. (1992). Tree based robust Bayesian estimation of time series structural models. Technical Report 107, Deptartment of Statistics, University of Washington, Seattle, WA, 98195.

Chang, I. and Tiao, G. C. (1983). Estimation of time series parameters in the presence of outliers. Technical Report 8, University of Chicago, Graduate School of Business.

Chang, I., Tiao, G. C., and Chen, C. (1988). Estimation of time series parameters in the presence of outliers. *Technometrics*, 30:193–204.

Cleveland, W. S. and Devlin, S. J. (1980). Calendar effects in monthly time series: detection by spectrum analysis and graphical methods. *Journal of the American Statistical Association*, 75:487–496.

Cleveland, W. S. and Terpenning, I. J. (1982). Graphical methods for seasonal adjustment. *Journal of the American Statistical Association*, 77(377):52–62.

Dagum, E. B., , Huot, G., and Morry, M. (1988). Seasonal adjustment in the eighties: some problems and solutions. *Canadian Journal of Statistics*, 16:109–126.

Dagum, E. B. (1978). Modelling, forecasting and seasonally adjusting economic time series with the X-11 ARIMA metohd. *The Statistician*, 27(3,4):203–216.

Dagum, E. B. (1980). The X-11-ARIMA seasonal adjustment method. Research paper, Seasonal Adjustment and Time Series Staff, Statistics Canada, 25-A, R.H. Coats Bldg., Ottawa K1A 0T6.

Dagum, E. B. and Morry, M. (1984). Basic issues on the seasonal adjustment of the Canadian consumer price index. *Journal of Business and Economic Statistics*, 2(3):250–259.

De Jong, P. (1988). The likelihood for a state space model. *Biometrika*, 75(1):165–169.

De Jong, P. (1991). The diffuse Kalman filter. *The Annals of Statistics*, 19(2):1073–1083.

den Butter, F. A. G., , Coenen, R. L., and van de Gevel, F. J. J. S. (1985). The use of ARIMA models in seasonal adjustment. *Emprical Economics*, 10:209–230.

den Butter, F. A. G. and Mourik, T. J. (1990). Seasonal adjustment using structural time series models: an application and a comparison with the Census X-11 method. *Journal of Business and Economic Statistics*, 8(4):385–394.

Dennis, J. E., Gay, D. M., and Welsch, R. E. (1981). An adaptive nonlinear least-squares algorithm. *ACM Transactions on Mathematical Software*, 7(3):348–383.

Findley, D. F., Monsell, B. C., Otto, M. C., Bell, W. R., and Pugh, M. G. (1988). Toward X-12-ARIMA. In *Proceedings of the Fourth Annual Research Conference, U.S. Bureau of the Census*, pages 591–622.

Findley, D. F., Monsell, B. C., Shulman, H. B., and Pugh, M. G. (1990). Sliding-spans diagnostics for seasonal and related adjustments. *Journal of the American Statistical Association*, 85(410):345–355.

Fox, A. J. (1972). Outliers in time series. *Journal of the Royal Statistical Society, Series B*, 34:350–363.

Gay, D. M. (1979). SUMSOL: A general quasi-Newton nonlinear function optimizer.

Gersch, W. and Kitagawa, G. (1983). The prediction of time series with trends and seasonalities. *Journal of Business and Economic Statistics*, 1:253–264.

Gordon, K. and Smith, A. F. M. (1990). Modeling and monitoring biomedical time series. *Journal of the American Statistical Association*, 85(410):328–337.

Hannan, E. J. (1964). The estimation of a changing seasonal pattern. *Journal of the American Statistical Association*, 59:1063–1077.

Hannon, E. J., Terrell, R. D., and Tuckwell, N. E. (1970). The seasonal adjusmtnet of economic time series. *International Economic Review*, 11(1):24–52.

Harrison, P. J. and Stevens, C. F. (1971). A Bayesian approach to short-term forecasting. *Operational Research Quarterly*, 22:341–362.

Harrison, P. J. and Stevens, C. F. (1976). Bayesian Forecasting. *Journal of the Royal Statistical Society, Series B*, 38:205–247.

Harvey, A. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press, Cambridge.

Harvey, A. C. (1984). A unified view of statistical forecasting procedures. *J. Forecast.*, 3:245–275.

Harvey, A. C. and Todd, P. J. J. (1983). Forecasting economic time series with structural and Box-Jenkins models: a case study. *Journal of Business and Economic Statistics*, 1(4):299–315.

Harvey, A. C. and Valls Pereira, P. L. (1989). Trend, seasonality and seasonal adjustment. In *Proceedings of the Seminar held in Mar del Plata*, pages 163–199, Argentina.

Hillmer, S. C., Bell, W. R., and Tiao, G. C. (1983). Modeling considerations in the seasonal adjustment of economic time series. In Zellner, A., editor, *Applied time series analysis of economic data*, pages 74–100. U. S. Bureau of the Census, Washington, D. C.

Jain, R. K. (1989). The seasonal adjustment procedures for the consumer price indexes: some empirical results. *Journal of Business and Economic Statistics*, 7(4):461–469.

Jain, R. K. (1992). Structural Model-Based Seasonal Adjustments of the Bureau of Labor Statistics Series and Their Benefits. Technical report, Bureau of Labor Statistics, 600 E St., NW, Room 4013, Washington, DC 20212. Draft.

Kitagawa, G. (1987). Non-Gaussian state space modeling of nonstationary time series. *Journal of the American Statistical Association*, 82(400):1032–1063.

Kitagawa, G. (1988). Numerical Approach to Non-Gaussian smoothing and its applications. In *20th Symposium on the Interface: Computing Science and Statistics*, pages 82–91.

Kitagawa, G. (1990). The two-filter formula for smoothing and an implementation of the Gaussian-sum smoother. Technical report, The Institute of Statistical Mathematics, 4-6-7 Minami-Azabu, Minato-ku Tokyo, Japan 106.

Kitagawa, G. (1991). Unpublished FORTRAN implementation of the Gaussian sum two filter smoother.

Kitagawa, G. and Gersch, W. (1984). A smoothness priors-state space approach modeling of time series with trend and seasonality. *Journal of the American Statistical Association*, 79:378–389.

Maravall, A. (1985). On structural time series models and the characterization of components. *Journal of Business and Economic Statistics*, 3:350–355.

Monsell, B. C. (1983). Using the Kalman smoother to adjust for moving trading day. Research Report SEC/RR-83/04, Statistical Research Division, U.S. Bureau of the Census, Room 3000-4, Washington, D.C. 20233.

Monsell, B. C. (1990). Documentation for the current prototype of the Census X-12-ARIMA seasonal adjustment program: Version 1.0. Technical report, Statistical Research Division, U.S. Bureau of the Census, Room 3000-4, Washington, D.C. 20233.

Monsell, B. C. and Otto, M. C. (1991). Documentation for the REGARIMA program for IBM compatible 386 PC's. Technical report, Statistical Research Division, U.S. Bureau of the Census, Room 3000-4, Washington, D.C. 20233.

Ozaki, T. and Thomson, P. J. (1992). Transformation and seasonal adjustment. (In preparation.) Institute of Statistics and Operations Researc h, Victoria University of Wellington.

S-PLUS (1991). *S-PLUS User's Manual.* Statistical Sciences, Inc., 1700 Westlake Ave., Suite 500, Seattle, WA 98109.

Shiskin, J., Young, A. H., and Musgrave, J. C. (1967). The X-11 variant of the Census method II seasonal adjustment program. Technical Paper 15, U.S. Bureau of the Census, U.S. Government Printing Office, Washington, D.C.

Shulman, H. B. and McKenzie, S. K. (1984). A Study of Pre-Adjustment Transformation. Statistical Research Division Report Series CENSUS/SRD/RR-84/34, U.S. Bureau of the Census, Room 3000, FOB #4, Washington, D.C. 20233.

Smith, A. and West, M. (1983). Monitoring renal transplants: an application of the Multiprocess Kalman filter. *Biometrics*, 39:867–878.

Tsay, R. S. (1988). Outliers, level shifts, and variance changes in time series. *J. Forecast.*, 7:1–20.

Young, P. C. (1988). Recursive Extrapolation, Interpolation and Smoothing of Nonstationary Time-Series. In *Proc. IFAC Symposium o Identification and System Parameter Estimation*, pages 33–44, Beijing, China.

Young, P. C. and Ng, C. N. (1990). Recursive Estimation and Forecasting of Nonstationary Time Series. *J. Forecast.*, 9:173–204.