A Multivariate Time Series Analysis of the Data Revision Process for Industrial Production and the Composite Leading Indicator *

Norman R. Swanson(+), Eric Ghysels (++), and Myles Callan (+)
(+) Pennsylvania State University
and
(++) Pennsylvania State University and Centre Interuniversitaire de Recherche et Analyse des Organisations (CIRANO), Montréal

October 1998

Abstract

In this paper we try to enhance our understanding of preliminary data releases of industrial production (IP) and the composite leading indicator (CLI) by investigating several time series properties of their data revision processes. In particular, we examine moments, autocorrelation functions, and integratedness properties of IP and CLI revision processes. We also construct and estimate univariate and multivariate regression models in order to assess the "efficiency" of IP and CLI revisions. Our findings based on regressions which include both IP and the CLI suggest that multivariate information "matters" in the revision process. For example, previously available IP revisions are useful for explaining CLI revisions, suggesting that releases of the CLI do not fully incorporate newly available IP data. In addition, IP revisions can be predicted from past CLI revisions, suggesting that a kind of causal feedback characterizes the revision processes for IP and the CLI. Finally, we conduct a series of real-time forecasting experiments in order to provide further evidence that there is useful univariate as well as multivariate information in the revision processes of these variables.

JEL classification: C82, C53, C22.
Keywords: data revision process, efficiency, predictability, real-time data.

* Corresponding Author: Norman R. Swanson, Department of Economics, Pennsylvania State University, University Park, PA 16827, nswanson@psu.edu, http://econ.la.psu.edu/˜swanson/index.htm.
The authors are grateful to two anonymous referees and the editors for useful comments, and to Greg Orlosky and Scott Briggs for excellent research assistance. Swanson thanks the National Science Foundation (Grant No. SBR-9730102) and the Research and Graduate Studies Office at Pennsylvania State University for research support.

1 Introduction

In this paper we introduce and examine a real time dataset consisting of all releases of seasonally adjusted and unadjusted U.S. industrial production, and the U.S. composite leading indicator. This is done by forming various forecasting models the revision processes of our variables. Although forecasting is clearly are area of research which Clive Granger holds dear, the topic of real-time data may not appear to be an appropriate topic with which to pay tribute to the work of Clive Granger. However, the subject is one to which Oscar Morgenstern - one of Clive's earliest collaborators - devoted an entire monograph. Further, Oskar Morgenstern's importance in Clive's early academic development, and Clive's great respect for him, should not go unnoticed. To illustrate this point, consider that it was Morgenstern who invited Clive, upon receiving the prestigious Harkness Scholarship of the Commonwealth Fund, to join a time series research project at Princeton University. While at Princeton, Clive began his work on spectral analysis, leading to one of his first monographs (see Granger and Hatanaka (1964)), and also resulting at least one of his many seminal contributions to the analysis of time series data (i.e. Granger (1969)).[1]

Today, Morgenstern would most likely write a somewhat different monograph. For example, he would likely examine preliminary data with many of the time series techniques and tools which Clive and many of his collaborators have developed over the last four decades. These might include careful examination of the integratedness properties of the revision process (see e.g. Granger and Newbold (1974), Granger and Ding (1996), Granger and Ermini (1993), Granger and Hallman (1991), and Granger and Swanson (1997)), cointegration analysis, (see e.g. Granger(1983, 1986), Granger and Engle (1987), Granger and Gonzalo (1995), Granger and Lee (1990), Granger and Swanson (1996), and Granger and Weiss (1983)), causal analysis (see e.g. Granger (1969, 1980, 1988), Granger and Lin (1995), and Swanson and Granger (1997)), and forecasting experimentation (see e.g. Granger (1992, 1996), Granger and Bates (1968), Granger and Nelson (1979)), for example. At this juncture, our goal is not to produce such a monograph. We offer only a modest attempt to reflect again on the issue of preliminary data releases which is so important in many regards, and which has perhaps yet not received as much attention as it deserves. Of course, this subject has not been completely ignored since the publication of Morgenstern (1963). A very limited number

[1]Further details of Clive's early research appear in the Econometric Theory interview of Clive Granger by Peter Phillips (1997).

of recent articles in this area (from which many other important references can be obtained) are: Boschen and Grossman (1982), Conrad and Corrado (1979), Ghysels (1982), Hamilton and Perez-Quiros (1996), Harvey et al. (1993), Kavajecz and Collins (1995), Keane and Runkle (1990), Koenig and Emery (1994), Mankiw and Shapiro (1986), Mankiw et al. (1984), Maravall and Pierce (1986), Mariano and Tanizaki (1995), Patterson (1995), Pierce (1981), Swanson (1996), and Swanson and White (1995, 1997a,b).

There are many examples in applied economics which illustrate the need to take a closer look at questions pertaining to the quality of preliminary data releases. For example, many econometric forecasting models are routinely constructed using "currently available" data. In the US., data are often downloaded from CITIBASE, and used without giving too much thought to the "timing" of the data. However, it is well known that many CITIBASE series are formed by combining various different "vintages" of economic data (e.g. preliminary data and data which have been revised a number of times).[2] Along the same lines, consider that it is most often preliminary data which are used by policymakers while, *post mortem*, their actions are scrutinized based on the use of revised data.[3] Further, forecasters typically use a mixture of revised and preliminary data in real time settings, and their predictions are initially appraised against preliminary releases. *Ex post* or *in sample* benchmarking of forecasting performance, however, is usually based on final figures. One natural question which arises in these types of scenarios is: "Which vintages of data are used and/or should be used by policy setters and forecasters, and are these data the same as those that are usually used in the construction of standard econometric models and forecasts¿'

In our analysis, we classify economic data into three categories:

(1) "*Preliminary*", "*First Reported*", or "*Unrevised*" Data. These types of data consist of the first reported datum for each variable at each point in time. Thus, a series of this type has had no revisions to any observations at any point in time. Swanson and White (1996, 1997) use unrevised data to construct real-time or ex-ante forecasts of a group of macroeconomic variables, and find that professional forecasts (which are necessarily real-time) are sometimes dominated by econometric models, based of a number of model selection criteria such as mean square forecast error and

[2]There are various exceptions to this practice. For example, Fair and Shiller (1990) acknowledge the issue of data revision quite explicitly, although they do not address it in their analysis.

[3]One of the few explicit treatments of data errors and its transmission into policy decisions is Maravall and Pierce (1986).

directional change forecasting ability.

(2) "*Partially Revised*" or "$Real - Time$" Data. These types of data are difficult to collect, as they are made up of a full vector of observations at each point in time for each variable. For example, if constructing a real-time data set for money, say M2, then for January 1990 a complete sequence of data, from say 1959, to January 1990 must be collected. Furthermore, the data must be collected as if one were in January 1990, so that no revisions of any kind made after January 1990 are incorporated into the dataset at time January 1990. Then, a whole new sequence of data from 1959 to February 1990 is collected, representing all of the information which was available in February of 1990, and thus including unrevised figures for February 1990, "once" revised figures for January 1990, and so on. This procedure is continued for each observational period in the sample. This real-time data collection strategy ensures that "future information" due to the use of information which is temporally antecedent to the date under consideration is not (accidentally) used in the construction of revised data. Also, this type of data avoids many of the types of problems associated with seasonal revisions, benchmark revisions, and definitional changes, for example, and can be thought of as truly real-time in the sense that it is the data set which is available to real-time forecasters and policy-setters at any given point in time. As an illustration of the potential for problems of this sort to arise, note that when dealing with seasonally adjusted data, it is has been observed by Pierce (1981) that revisions are mostly due to the adjustment process, as it involves two-sided filters. But these two-sided filters necessarily involve a mixing of data vintages (e.g. future data seep into the revised values of past data), and as such it is clearly difficult to ensure the future information does not enter into the revision process. For more on the features of the most commonly used X-11 seasonal adjustment procedure, see for instance Ghysels, Granger and Siklos (1996) among others. In this paper, we focus on real-time datasets, examining the revision process across more than three decades for our series.[4] One of the real-time datasets which we examine in this paper an updated version of the real-time CLI dataset used by Diebold and Rudebusch (1991) in their analysis of the forecasting ability of the CLI for fully revised industrial production.

(3) "*Fully Revised*" or "*Final*" Data. It is quite possible that true "final" data will never be available for many economic series. This is because benchmark and definitional changes are ongoing and may continue into the indefinite future, for instance. In practice, by final data we usually mean

[4] Sargent (1989) is one of the few examples where revision process errors (in his case, "final" data errors) is treated explicitly, examining how they affect the estimation of econometric models.

the revision of a data observation after which no more revisions will be made. This is the type of data that academics often have in mind when conducting economic time series studies, perhaps simply because it is data which is not subject to revision, and it is felt that if one could adequately forecast a "final revised" figure, then there is no need for further modeling. However, this data is clearly not easy to obtain, as data are generally subject to revision for indefinite lengths of time. Indeed, most datasets which are constructed by applied economists clearly consist of a mixture of preliminary data, partially revised data, and final revised data. In the following, we consider real-time data, where $X_{t+i}(t)$ refers to the $t+i^{th}$ release date of data pertaining to calendar date t, and X is the growth rate of the original series.

Our primary aim in this paper is to underscore this potential shortcoming of dataset construction by examining real time datasets, noting their time series characteristics, and examining univariate and multivariate regression models of the data revision process. This approach allows us to assess whether there are predictable patterns in the data revision process. We indeed find that this is the case. For example, our findings based regressions which include both IP and the CLI suggest that multivariate information "matters" in the revision process. In particular, previously available IP revisions are useful for explaining CPI revisions, suggesting that releases of CPI do not fully incorporate newly available IP data. In addition, real-time forecasting experiments suggest that there is useful univariate as well as multivariate information in the revision processes of IP and the CLI. Given these sorts of findings, we suggest that there is a need for careful examination of real-time data when asking a variety of standard questions, including: "Are two variables Granger causal for one another?", and "Is variable x useful for constructing forecasts of variable y, in a real-time forecasting scenario, say, such as that faced by policy makers in Washington as well as decision makers on Wall street?".

The rest of the paper is organized as follows. In Section 2, we examine basic statistical properties of our data, including integratedness, cointegratedness, and autoregressive behavior. In Section 3 we perform a series of regression and ex ante forecasting experiments in order to ascertain whether or not revisions to our industrial production series are useful for forecasting revisions to the composite leading indicator, and *vice-versa*, thereby examining the "efficiency" of data revision within the context of providing optimal forecasts of "fully revised" data. The fourth section concludes and makes recommendations for future research.

2 The Data

The three variables for which we collect real-time data are U.S. seasonally adjusted industrial production (1950:4 to 1996:2), unadjusted industrial production (1950:4 to 1996:2), and the composite leading indicator (1968:10 to 1996:2). A typical months' release of data for these variables comprises of a first, or preliminary, release for the previous month, and 4 to 6 months of revisions to data previously released. In addition, more comprehensive benchmark and baseyear revisions occur from time to time for each of the variables.

Turning first to our industrial production data, the following details are perhaps worth noting. The seasonally adjusted industrial production and unadjusted industrial production series are compiled by the Federal Reserve Board. The main source for seasonally adjusted industrial production data is the Federal Reserve Bulletin. For unadjusted industrial production data we use the Federal Reserve Bulletin for data up to 1959:12, and the Survey of Current Business for data after 1959:12. Additional data for each series are obtained from Federal Reserve monthly statistical releases.[5] Also, for three of the major (benchmark) revisions to the variables, the Federal Reserve Board released separate publication - (1) Industrial Production 1957-59 Base, (2) Industrial Production 1971, and (3) 1976 Revision. Since 1980, benchmark revisions to both industrial production series have occurred approximately every 2 years, while prior to 1980 such revisions were less frequent. Benchmark revisions to industrial production usually result in updates of around 2 years of previously available data, and occur 5 times during our sample period: December 1953, December 1959, July 1971, July 1985, and April 1990.[6] These 5 base year revisions incorporate re-weighting of, and changes to the components of the index. For both series there are 3 missing entries due to two major revisions, they are 1953:11,12 and 1985:03. We replaced each missing observation with the first available data for that period (which in each case is a second release).

Our other variable, the composite leading indicator was compiled by the Department of Commerce until 1994:12. It is currently released by The Conference Board. Our CLI dataset up until 1988:12 was made available to us by Glenn Rudebusch (see e.g. Diebold and Rudebusch (1991)).

[5]Federal Reserve releases for industrial production are referred to as G.12.3 for dates up to 1990:4, and G.17 thereafter. These releases are typically published on the 15th of each month and contain the preliminary data for the previous month, as well as revisions to data for earlier months.

[6]Details of the base year changes are available in the relevant Federal Reserve Bulletin, see e.g. pp. 1247-1279 - December 1953, pp. 1451-1466 and pp. 552-573 - July 1971, pp. 447-497 - July 1985, and pp. 187-204 - April 1990.

We augment this dataset by including data from the Business Consumers Digest, up to 1990:12, and the Survey of Current Business, from 1991:1 to 1994:12. Benchmark revisions to the CLI occur every 12-18 months, and are revisions to the whole series. These revisions incorporate changes in methodology for computing the index, updated statistical factors, and historical revisions in component data. Base year revisions to the CLI occur three times - August 1970, January 1989, and October 1993.[7] For our analysis, we examine revisions of monthly growth rates, and hence data not rebased. This approach allows us to avoid problems associated with level shifts in our CLI series. We address what might be called "variance shifts" by examining sub-samples of our datasets.

3 Basic Statistics and Data Analysis

In our examination of the real-time datasets discussed above, we examine revisions to the data process. In particular, we consider what we shall call fixed width revisions, which are defined as $X_{t+i}(t) - X_{t+i-1}(t)$, where X is the growth rate of the original levels data, formed by taking log differences. Here, the subscript refers to the release date of the data, and the bracketed index denotes the date to which the release pertains. Thus, by varying the index, i, while keeping t fixed, we examine various different "releases" of data for a given time period, for example. We also consider what we call increasing width revisions, which are defined as $X_{t+i}(t) - X_t(t)$.[8] The increasing width revisions represent the accumulated fixed width revisions, and for $i = 24$ we obtain the difference between what we will here call the "final" datum, and the first available datum. Our examination of this type of data allows us to assess whether or not there is significant systematic bias in accumulated revisions, when there is no significant systematic bias in individual revisions. One reason why such information is of interest is that we can then ascertain whether there is significant bias in the difference between first and final releases of our variables.

Tables 1a-2b contain summary statistics for fixed and increasing width revisions of unadjusted and seasonally adjusted industrial production and the CLI, for the entire datasets (Tables 1a-1b)

[7]See Business Conditions Digest, pp. iii - August 1970, and the Survey of Current Business, pp. 23-27 - January 1989, and pp. 44-52 - October 1993, for details of the revision.

[8]All time series formed in our analysis were tested for nonstationarity using augmented Dickey-Fuller tests with an intercept, and with the number of lagged dependent variables used in the regressions chosen by starting with 20 lags, and decreasing the number of lags used until the last included lag had a coefficient significantly different from zero at a 95% level of confidence. In all cases, the I(1) null hypothesis was rejected at a 5 percent significance level.

and for a smaller sample from 1975:10-1996:2 (Tables 2a-2b)[9]. We consider revisions over horizons $i = 1, \ldots, 12, 18,$ and 24. Each table contains three panels corresponding to our three variables. The first panel pertains to seasonally adjusted IP. We report the mean, variance, skewness and kurtosis for both the fixed and increasing width revisions. We observe that the mean of fixed length revisions is significantly different from zero at a 95% level of confidence, for $i = 1, 3,$ and 6. This suggests that there is systematic bias in revisions of adjusted industrial production, which could be used to increase the accuracy of preliminary releases. Interestingly, this feature also characterizes unadjusted IP and, to a lesser extent, the CLI. This result holds for our full sample periods (Table 1a-1b), as well as for the smaller sample period (Table 2a-2b), which we use in our "efficiency" tests described in the next section. Notice also that the results in Tables 1a and 2a show some clear departures from normality as both the skewness and kurtosis statistics usually differ from the values associated with the Gaussian distribution, but only for the fixed revision case.

Perhaps not surprisingly, nonzero mean revisions also arise for increasing width revisions. Indeed, we expect that fixed length revision mean bias associated with lower values of i leads to mean bias for subsequent values of i, when examining increasing widths revisions. Interestingly, the number of significant nonzero means is much greater for the columns associated with increasing width revisions than for those associated with fixed widths. For example, based on increasing revisions with $i = 24$, the difference between the "final" and initial releases display significant mean bias for a number of variables and sample periods (see Tables 1a- 2b). In these cases, a correction could be made to all releases of the variables, prior to their final release, which improves the accuracy of all of our preliminary data.

An important caveat to the above discussion is that the summary statistics for our data series reported in Tables 1a and 2a are reported for the entire data series. Hence, no effort has been made to assess the impact of benchmark revisions on the findings, for example. In an effort to address this issue, we also include two additional tables, (Tables 1b and 2b) which contain summary statistics based on datasets from which observations associated with benchmark revisions have been removed. Interestingly, the patterns of significant revision biases are essentially the same as those

[9]The smaller sample period was selected because in late 1975 there appears to be a "variance shift" in the CLI and in IP (see Figures 1-3). In addition, it is worth noting that there was a 25% level shift in the CLI in late 1975. However, this level shift does not affect our data, as we are looking at growth rate revisions

reported above, although bias estimates are generally lower, as might be expected given that we have removed observations which might be viewed as outliers. Another feature of our datasets which is worth mentioning is that major revision points are usually associated with base year and benchmark revisions. For IP, benchmark revisions occurred on 1953:11, 1959:12, 1967:10, 1971:8, 1976:7, 1985:6, 1990:3, 1994:1, 1997:1, and 1997:11, while baseyear revisions occurred on 1953:11, 1959:12, 1962:10, 1971:8, 1985:6, 1990:3, and 1997:1. For the CLI, benchmark revisions occurred on 1973:9, 1975:4, 1975:10, 1979:2, 1983:1, 1993:9, and 1996:10, while baseyear revisions occurred on 1968:12, 1970:7, 1976:10, 1988:12, and 1993:9. Casual examination of these dates, however, suggests that there are no obvious links between major revisions of IP and major revisions of CLI.

The statistics reported in Tables 1b and 2b are complemented by a set of figures displaying the data without benchmark revisions (see Figures 1 to 3). Each figure shows plots of the first, second, sixth, and twelfth fixed width revisions, as well as the first, second, eighteenth, and twenty-fourth increasing width revisions. We notice some common patterns in Figures 1 and 2 (Industrial Production). First, there appears to be much more variability in first revisions than in later revisions. This finding corresponds to the results presented in Tables 1a and 1b, where we primarily find significant mean bias for fixed length revisions constructed using small values of i. This finding does not impact on our results concerning increasing length revisions, though, as in these cases even small errors accumulate over time. Indeed, the increasing width revisions plotted in Figures 1 and 2 are not only highly variable but are also frequently nonzero, regardless of the value of i. Second, the revision patterns are somewhat similar, with most outliers occurring in the same time periods, regardless of whether the data have been seasonally adjusted or not. This might be viewed as somewhat surprising, given that seasonally adjusted data are contaminated by future information "leakage" caused by the use of two-sided moving average filters (see above discussion). On the other hand, the result makes sense given that benchmark revisions which tend to drive the outliers do not follow a seasonal pattern. Figure 3 contains plots CLI revisions. Overall the features of these data are similar to those of the IP series, and there is no obvious structural break, although there is an apparent decrease in fixed revision variability associated not only with 1st revisions versus 2nd revisions, say, but also with increasing calendar time for any given revision series. Next, we turn to an examination of the predictability of our variables using information in past revisions. Any evidence of this type of predictability which we find will be taken as evidence that the revision process is "inefficient". Thus our use of the word of efficiency is used only in the context of predictability.

4 Efficiency of Data Revisions

The main question we are concerned with in this section is whether there are predictable patterns in the revision process. We attempt to answer this question by constructing various different regressions and by examining the autocorrelation functions of the variable. Before turning to our regression results, consider the autocorrelation functions plotted in Figures 4 and 5. Each figure consists of four three-dimensional plots. On one axis we display the order of the autocorrelations, while the other axis corresponds to different values of i. Thus, each row of autocorrelation bars, corresponding to the revision index, is based on the examination of an individual revision time series, i.e. $\rho[(X_{t+i}(t) - X_{t+i-1}(t)), (X_{t-j+i}(t-j) - X_{t-j+i-1}(t-j))]$ for fixed width revisions, and $\rho[(X_{t+i}(t) - X_t(t)), (X_{t-j+i}(t-j) - X_{t-j}(t-j))]$ for increasing width revisions, where ρ denotes the autocorrelation, i is the revision index, and j is the lag order. We first discuss the results appearing in Figure 4, which are based on an examination of seasonally adjusted and unadjusted Industrial Production. The left two panels correspond to fixed width, while the right two panels are for increasing width revisions. We note two features of importance. First, the autocorrelations in the seasonally adjusted data are rather similar in magnitude to those associated with unadjusted data, although there are *more* fixed width autocorrelations which are significantly different from zero when adjusted data are examined. Second, there is much more persistence in the increasing width revisions. The second feature is consistent with the accumulation of errors in increasing width revisions, while the first feature provides weak evidence in support of a hypothesis of increasing autocorrelation associated with the use of two-sided seasonal adjustment filters (i.e. a hypothesis that there is information leakage when seasonally adjusted data are revised). In summary, though, the incidence of nonzero autocorrelations in this first casual examination of autocorrelation functions suggests that past revisions may be useful for predicting future revisions, and hence future releases of the actual data. Thus, the revision process may be inefficient. Even cursory examination of Figure 5 suggests that these findings also apply to the CLI revisions data. Moreover, it is interesting to note that the CLI autocorrelation functions are similar in the upper and lower panels, suggesting that the our shorter sample has prediction patterns similar to our longer sample.

In order to investigate the statistical significance of any predictable patterns in our data, we began by running a number of autoregressions. In particular, we ran autoregressions based on all of the data series plotted in Figures 4 and 5 (i.e. $i = 1, \ldots, 24$), and with lags selected based on the Schwarz Information Criterion. The regressions are formed by projecting j^{th} revisions of data pertaining to time t onto the j^{th} revisions pertaining to time $t-i$, $i = 1, \ldots, 24$. Tables 3a-3b contain our findings for the full and shorter sample periods, where results are only tabulated for regressions in which at least one explanatory variable is significantly different from zero at a 90% level of confidence.[10] Our overall conclusion based on this univariate analysis is clearly that the inefficiency in the revision process is primarily due to mean bias (note the number of significant constants in the reported regressions), with little information in autoregressive variables being relevant. This can be seen most clearly by examining the results contained in Table 3b. In addition, note that reported adjusted R^2 values are quite small, suggesting that there is little predictive content, at least when we attempt to explain a first revision today with a first revision yesterday, say, as is done in the regressions.

In order to further examine the issue of efficiency in a univariate context, we consider regressions involving different vintages of revisions of the same variable, for a fixed time period. Hence, instead of running autoregressions with time period $t-i$, $i = 1, \ldots, 24$, used to construct the explanatory variables, we fix $i = 0$, and the regressions are thus formed by projecting j^{th} revisions of data pertaining to time t onto $(j-i)^{th}$ revisions pertaining to time t, $i = 1, \ldots, 24$, where the number of regressors in each regression is constrained by the condition that $j > i$. This condition is needed in order to ensure that present and/or future information is not used in and of the explanatory variables. Tables 4a-4b contain our findings, where again only coefficients significantly different from zero at a 90% level of confidence are reported. We find that future revisions (e.g. $X_{t+j}(t) - X_{t+j-1}(t)$), can often be predicted from past revisions when j is small, regardless of which sample period is used. For instance, for the CLI (small sample, see Table 4b) various coefficients associated with regressions, where the dependent variable is $X_{t+j}(t) - X_{t+j-1}(t)$ for $j = 2, 4, 6, 8, 9, 10, 20, 21,$ and, 22, are significant. Such predictable patterns also occur, but to a lesser extent, for revisions in unadjusted and adjusted IP. Thus, while our autoregressions reported in Tables 3a-3b show that inefficiencies often appear through mean biases when examining individual

[10] The values in parentheses are the numerical values of the significant coefficient estimates. All estimates are based on OLS and are evaluated using HAC standard error estimates.

revision vintages (allowing t to vary), our results in Tables 4a-4b suggest that inefficiencies also occur *between* different vintages (for fixed t).

Given these findings, it should also be of interest to assess whether information in lags of one of our variables is useful for predicting revisions in other variables. We first address this question by running multivariate regressions of the type reported on in Table 5. In all, 16 regressions - with up to 23 exogenous variables in each - were run, and, as above, those with significant coefficients are reported on in the table. In all regressions, the dependent variable is the first revision of a given variable, say $Y_{t+1}(t) - Y_t(t)$. Panel A of the table contains regression results based on fixed width revisions, while Panel B contains results for increasing width revisions. In the fixed width regressions, the explanatory variables are first revisions available during previous time periods, $X_{t-i+1}(t-i) - X_{t-i}(t-i)$, for $i = 1, \ldots, 24$. In the increasing width regressions, the explanatory variables are accumulated revisions available at time t regarding previous periods' data (i.e. $X_t(t-i) - X_{t-i}(t-i)$, for $i = 1, \ldots, 24$). The first column in the table lists the dependent variable, while the second shows the independent variable. The third column lists the significant coefficients from each regression. The results are interesting for a number of reasons. First, there are a number of significant regressors in our regressions, suggesting that multivariate information may "matter" in the revision process, particularly in increasing width revision cases. In addition, note the relatively high $\bar{R}^2$ values for some of the regressions. Second, the usefulness of previously available IP revisions for explaining CLI revisions suggests that releases of the composite leading indicator do not fully incorporate newly available IP data. For example, we might expect a one period lag in the transfer (across government agencies) of updated IP information to the CLI. However, we find that information available prior to time $t-1$ is useful for predicting CLI revisions in 2 CLI regressions. This multivariate result is rather significant, as it casts some doubt on the accuracy of early CLI releases. In addition, note that this result is not affected by the use of our shorter sample period. Finally, it is worth pointing out that CLI revisions may be useful for predicting IP revisions, although the relationship appears weaker than that from IP to CLI revisions.

In order to shed additional light on the above findings, we also report on a different variety of multivariate regression (see Table 6). In these regressions, previous updates of the CLI are used to model current updates of IP, for example. The most surprising result from this table is that past revisions of a particular calendar date for CLI are useful for predicting the current IP revision for the same calendar date, in many case. In addition, there appears to be causality from IP to CLI, but the linkages are weaker than for the CLI to IP case. Thus, our earlier conclusions based on an examination of the results in Table 5 remain intact when a different variety of multivariate regression is run. In summary, our multivariate analysis appears to point to a number of potentially interesting types of inefficiencies inherent to the revision process which are not readily apparent when univariate data are examined.

In the spirit of much of Clive Granger's work on out-of-sample forecasting, we close our discussion of the data revision process with an examination of the relevance of our regression findings within the context of a real-time forecasting experiment. In particular, for a variety of the univariate and multivariate fixed width regressions reported on in Tables 3a-3b and 5 we construct a sequence of 1-step ahead forecasts based on regressions formulated using increasing samples of data with coefficients re-estimated at each point in time (model specifications were fixed beforehand using the SIC). A representative sample of our findings is reported in Figure 6. Of particular note is that the out-of-sample R^2 values (see e.g. Swanson and White (1995) for further explanation of out-of-sample R^2 values) based on all of the reported experiments (where the ex-ante period is 1992:1-1996:2) are greater than zero, with many values greater than 0.10. In addition, the highest R^2 values are associated with experiments based on univariate models.

5 Conclusions and Recommendations

In this paper we have undertaken a multivariate time series analysis of the data revision process for industrial production (IP) and the composite leading indicator (CLI). We offer the following conclusions. First, there appears to be mean bias in the revision process for both seasonally adjusted and unadjusted IP, and for the CLI. This type of inefficiency is interesting, because merely adding constants to preliminary releases can correct such systematic bias. Second, based on univariate autoregressive models we find some evidence that past revisions of our variables have predictive content for current and future revisions. This in turn suggests that past revisions can be used to improve upon our estimates (or preliminary releases) of final data. Third, we present findings based multivariate regressions which suggest that multivariate information "matters" in the revision process. For example, previously available IP revisions are useful for explaining CLI revisions, suggesting that releases of CLI do not fully incorporate newly available IP data. In addition, real-time forecasting experiments suggest that there is useful univariate as well as multivariate information in the revision processes of IP and the CLI.

This work represents a starting point, and many issues remain to be explored, both theoretical and empirical. From a theoretical perspective, for example, it is of interest to characterize the data revision process within the context of a macroeconomic model of policy decision-making. From an empirical perspective, it is of interest to assess whether knowledge of the revision process be used in a real-time context to improve economic forecasts of fully-revised data. Also, within a multivariate real-time data context, many questions concerning Granger causality and model specification remain unanswered. Finally, it remains to assess whether the types of inefficiencies which we find in this paper characterize real-time data in general, or are specific to our datasets.

6 Bibliography

Boschen, J.F. and Grossman, H.I. (1982), Tests of Equilibrium Macroeconomics Using Contemporaneous Monetary Data, *Journal of Monetary Economics,* 10, 309-333.

Conrad, W. and C. Corrado (1979), Application of the Kalman Filter to Revisions in Monthly Retail Sales Estimates, *Journal of Economic Dynamics and Control,* 1, 177-198.

Diebold, F.X. and Rudebusch, G.D. (1991), Forecasting Output With the Composite Leading Index: A Real Time Analysis, *Journal of the American Statistical Association,* 603-10.

Engle, R.F. and C.W.J. Granger, (1987), Co-integration and Error Correction: Representation, Estimation, and Testing, *Econometrica,* 55, 251-76.

Fair, R.C. and Shiller, R.J. (1990), Comparing Information in Forecasts from Econometric Models, *American Economic Review,* 80, 375-89.

Ghysels, E. (1982), Tijdreeksanalyse en Fouten in BNP-Cijfers - Een Theoretische Korrektie en Herinterpretatie, *Cahiers Economiques de Bruxelles,* 96, 489-495.

Ghysels, E., C.W.J. Granger and P. Siklos (1996), Is Seasonal Adjustment a Linear or Nonlinear Data-Filtering Transformation?, *Journal of Business and Economic Statistics,* 14, 374-386.

Granger, C.W.J. (1969), Investigating Causal Relations by Econometric Models and Cross Spectral Methods, *Econometrica,* 37, 428-438.

Granger, C.W.J. (1980), Testing For Causality, a Personal Viewpoint, *Journal of Economic Dynamics and Control,* 2, 329-352.

Granger, C.W.J. (1983), Co-integrated Variables and Error-correcting Models, Discussion Paper 83-13, University of California, San Diego.

Granger, C.W.J. (1986), Developments in the Study of Co-integrated Economic Variables, *Oxford Bulletin of Economics and Statistics,* 48, 213-228.

Granger, C.W.J. (1988), Some Recent Developments in a Concept of Causality, *Journal of Econometrics,* 39, 199-211.

Granger, C.W.J. (1992), Forecasting Stock Market Prices - Lessons for Forecasters, *International Journal of Forecasting,* 8, 3-13.

Granger, C.W.J. (1996), Can We Improve the Perceived Quality of Economic Forecasts?, *Journal of Applied Econometrics,* 11, 455-473.

Granger, C.W.J. and J. Bates (1969), The Combination of Forecasts, *Operational Research Quarterly,* 20, 451-468.

Granger, C.W.J. and Z. Ding (1996), Varieties of Long Memory Models, *Journal of Econometrics,* 73, 61-77.

Granger, C.W.J. and L. Ermini (1993), Some Generalizations of the Algebra of I(1) Processes, *Journal of Econometrics,* 58, 369-384.

Granger, C.W.J. and J. Gonzalo (1995), Estimation of Common Long Memory Components in Cointegrated Systems, *Journal of Business and Economic Statistics,* 13, 27-35.

Granger, C.W.J. and J. Hallman (1991), Nonlinear Transformations of Integrated Time Series, *Journal of Time Series Analysis,* 12, 207-224.

Granger, C.W.J. and M. Hatanaka (1964), *Spectral Analysis of Economic Time Series,* Princeton University Press.

Granger, C.W.J. and T.-H. Lee (1990), Multicointegration, in: *Advances in Econometrics,* ed. T. Fomby.

Granger, C.W.J. and J.-L. Lin (1995), Causality in the Long Run, *Econometric Theory,* 11, 530-536.

Granger, C.W.J. and H.L. Nelson (1979), Experience with Using the Box-Cox Transformation when Forecasting Economic Time Series, *Journal of Econometrics,* 9, 57-69.

Granger, C.W.J. and P. Newbold (1974), Spurious Regressions in Econometrics, *Journal of Econometrics,* 2, 111-120.

Granger, C.W.J. and N.R. Swanson (1996), Further Developments in the Study of Cointegrated Economic Variables, *Oxford Bulletin of Economics and Statistics,* 58, 537-553.

Granger, C.W.J. and N.R. Swanson (1997), An Introduction to Stochastic Unit Root Processes, *Journal of Econometrics,* 80, 35-62.

Granger, C.W.J. and A.A. Weiss (1983), Time Series Analysis of Error-correction Models, in: *Studies in Econometrics, Time Series and Multivariate Statistics,* in honor of T.W. Anderson, ed. S. Karlin, T. Amemiya, and L.A. Goodman, Academic Press.

Hamilton, J.D. and G. Perez-Quiros (1996), What Do The Leading Indicators Lead?, *Journal of Business,* 69, 27-49.

Harvey, A.C., C.R. McKenzie, D.P.C. Blake and M.J. Desai (1993), Irregular Data Revisions, in A. Zellner (Ed.), *Applied Time Series Analysis of Economic Data,* Bureau of the Census, 329-347.

Kavajecz, K.A. and S. Collins (1995), Rationality of Preliminary Money Estimates, *Review of Economic Studies,* 62, 32-41.

Keane, M.P. and Runkle, D.E. (1990), Testing the Rationality of Price Forecasts, *American Economic Review,* 80, 714-35.

Koenig, E.F. and K.M. Emery (1994), Why the Composite Index of Leading Indicators Does Not Lead, *Contemporary Economic Policy,* 12, 52-66.

Mankiw, N.G., D.E. Runkle, and M.D. Shapiro (1984), Are Preliminary Announcements of the Money Stock Rational Forecasts?, *Journal of Monetary Economics,* 14, 15-27.

Mankiw, N.G. and M.D. Shapiro (1986), News or Noise: An Analysis of GNP Revisions, *Survey of Current Business,* 66, 20-25.

Maravall, A. and D. Pierce (1986), The Transmission of Data Noise into Policy Noise in U.S. Monetary Control, *Econometrica,* 54, 961-979.

Mariano, R.S. and Tanizaki, H. (1995), Prediction of Final Data with Use of Preliminary and/or Revised Data, *Journal of Forecasting,* forthcoming.

Morgernstern, O. (1963), *On The Accuracy of Economic Observations,* 2nd Ed., Princeton University Press, Princeton.

Patterson, K.D. (1995), An Integrated Model of the Data Measurement and Data Generation Process with an Application to Consumers' Expenditure, *The Economic Journal,* 105, 54-76.

Pierce, D.A. (1981), Sources of Error in Economic Time Series, *Journal of Econometrics,* 17, 305-321.

Phillips, P.C.B. (1997), ET Interview: Clive Granger, *Econometric Theory,* 13, 253-305.

Sargent, T.J. (1989), Two Models of Measurements and the Investment Accelerator, *Journal of Political Economy,* 97, 251-287.

Swanson, N.R. (1996), Forecasting Using First-Available Versus Fully Revised Economic Time-Series Data, *Studies in Nonlinear Dynamics and Econometrics,* 1, 47-64.

Swanson, N.R. and C.W.J. Granger (1997), Impulse Response Functions Based on Causal Approach to Residual Orthogonalization in Vector Autoregressions, *Journal of the American Statistical Association,* 92, 357-367.

Swanson, N.R. and H. White (1995), A Model Selection Approach to Assessing the Information in the Term Structure Using Linear Models and Artificial Neural Networks, *Journal of Business and Economic Statistics,* 13, 265-275.

Swanson, N.R. and H. White (1997a), A Model Selection Approach to Real-Time Macroeconomic Forecasting Using Linear Models and Artificial Neural Networks, *Review of Economics and Statistics,* 79, 540-550.

Swanson, N.R. and H. White (1997b), Forecasting Economic Time Series Using Flexible Versus Fixed Specification and Linear Versus Nonlinear Econometric Models, *International Journal of Forecasting,* 13, 439-461.