Discussion Paper 35

Institute for Empirical Macroeconomics
Federal Reserve Bank of Minneapolis
250 Marquette Avenue
Minneapolis, Minnesota 55480

December 1990

RATIONAL EXPECTATIONS MODELING WITH SEASONALLY ADJUSTED DATA

Christopher A. Sims
Economics Department
Yale University

ABSTRACT

In a world where time series show clear seasonal fluctuations, rational agents will take account of those fluctuations in planning their own behavior. Using seasonally adjusted data to model behavior of such agents throws away information and introduces possibly severe bias. Nonetheless it may be true fairly often that rational expectations modeling with seasonally adjusted data, treating the adjusted data as if it were actual data, gives approximately correct results; and naive extensions of standard modeling techniques to seasonally unadjusted data may give worse results than naive use of adjusted data. This paper justifies these claims with examples and detailed arguments.

JEL Numbers: 023, 132, 311

This material is based on work supported by the National Science Foundation under Grant No. SES-8722451. The Government has certain rights to this material.

Any opinions, findings, conclusions, or recommendations expressed herein are those of the author(s) and not necessarily those of the National Science Foundation, the University of Minnesota, the Federal Reserve Bank of Minneapolis, or the Federal Reserve System.

Rational Expectations Modeling with Seasonally Adjusted Data

by Christopher A. Sims

Revised 12/90

ABSTRACT

In a world where time series show clear seasonal fluctuations, rational agents will take account of those fluctuations in planning their own behavior. Using seasonally adjusted data to model behavior of such agents throws away information and introduces possibly severe bias. Nonetheless it may be true fairly often that rational expectations modeling with seasonally adjusted data, treating the adjusted data as if it were actual data, gives approximately correct results; and naive extensions of standard modeling techniques to seasonally unadjusted data may give worse results than naive use of adjusted data. This paper justifies these claims with examples and detailed arguments.

I. The Prima Facie Case Against Seasonally Adjusted Data

In order to conserve scarce analytical capacity, quantitative economists usually work on the assumption that the models they are estimating include, under some setting of their parameters, the truth. Under this assumption, it is clear that use of seasonally adjusted data in estimating practically any model is a mistake. It will often be true that observations on economic behavior related to seasonal frequencies could be informative about the unknown parameters we are trying to estimate, in which case using seasonally adjusted data at best amounts to throwing away observations and at worst could severely bias results. Even when seasonal variation is mostly noise, unrelated to the phenomena being studied, a properly specified model should be able to allow for this. Using seasonally adjusted data will then be at best an approximation to the correct treatment of seasonal noise.

Despite these self-evident truths, Wallis [1974] and I [1974] showed that in distributed lag regression models with exogenous regressors, proper use of seasonally adjusted data could produce estimates with small approximation bias under broad regularity conditions. Roughly speaking, one can use a parametric form for the lag distribution B that enforces smoothness of its Fourier transform $\tilde{B}$ across seasonal bands. With seasonally adjusted data, this will produce small bias outside seasonal bands so long as the parameterization is accurate outside seasonal bands. To obtain similarly accurate results in non-seasonal bands with seasonally unadjusted data would require that seasonal dynamics in the lag distribution be modeled as accurately as non-seasonal dynamics. This would ordinarily require a much more elaborate parameterization. If our interest focuses on non-seasonal (say business cycle) frequencies, there may be a tendency not to model the seasonal frequencies carefully enough to avoid strong bias in results.

It has been argued that, whatever the validity of the foregoing argument in the case of distributed lag regression with exogenous regressors, it no longer applies in rational expectations models. There is some basis for this argument -- only rather artificial rational expectations models fit into the exogenous regressor distributed lag regression framework. A typical rational expectations model imposes cross-equation restrictions, for example that the coefficients in a demand equation are related to coefficients in an equation generating forecasts of an exogenous variable. Instead of theory's making predictions about the lag distribution in a single regression equation, theory predicts relations among equations of a complete dynamic model for several variables.

In the regression framework, we deal with two variables (or vectors of variables) y and x. The relation between them, in the frequency domain, can be written as

$$\tilde{y}(\omega) = \tilde{b}(\omega)\tilde{x}(\omega) + \tilde{\upsilon}(w), \ \omega \text{ in } (0,2\pi) \qquad (1)$$

We assume that the disturbance process υ is uncorrelated with the x process at all leads and lags. If y and x are both filtered with the same filter c, so that Y=c*y and X=c*x (with "*" indicating convolution), then the filtered variables are related by

$$\tilde{Y}(\omega) = \tilde{b}(\omega)\tilde{X}(\omega) + \tilde{c}(\omega)\tilde{\upsilon}(w) . \tag{2}$$

The filtered error term remains orthogonal to X, and the lag distribution b is the same as that in (1). Thus filtering, whether for seasonal adjustment or other purposes, leaves the probability limit of estimates of b unaffected. The effects of filtering can be discussed entirely in terms of effects on sampling error and on approximation error.

When we are estimating a complete dynamic system, there is a single observed vector stochastic process y, with a moving average representation

$$y(t) = A*\varepsilon(t) \ . \tag{3}$$

Filtering y through c so that Y(t)=c*y(t) produces

$$Y(t) = c*A*\varepsilon(t) \ . \tag{4}$$

If $\tilde{c}$ has no zeros in the lower half of the complex plane (which is called the "minimum delay" property and is almost the same as requiring c to have a one-sided inverse under convolution) then ε in (4) is the innovation in Y, just as in (3) it was the innovation in y. However the moving average operator A has been replaced by c*A. More generally, c may fail to be minimum delay, in which case ε as well as the moving average operator will be different if we treat Y rather than y as our data series. Thus in contrast to the result in the regression case -- where b is unaffected by filtering -- for a complete dynamic system the basic estimated parameter A is altered by filtering.

II. Approximation Error for Dynamic Systems

Suppose we consider a class of autoregressive filters $B(\cdot;\alpha)$, n-vector-valued functions on the non-negative integers with k-dimensional parameter α. Our data are the $Y(t)$, defined in (4) based on the underlying true model (3). We assume $B(0;\alpha)=I$ for all α. It is natural to estimate α by fitting the equation

$$Y(t) = -B^{+}(\alpha)*Y(t) + \eta(t) , \tag{5}$$

minimizing with respect to α the sample average of $\eta(t)\Sigma^{-1}\eta(t)$, where $\Sigma=\mathrm{var}(\eta(t))$. ($B^{+}(t)$ is defined by $B^{+}(0)=0$, $B^{+}(t)=B(t)$, $t>0$.) We suppose that the true moving average representation for Y is

$$Y(t) = G*\xi(t) . \tag{6}$$

If c is minimum-delay, $\xi=\varepsilon$ and $G=c*A$. If the model (5) contains the truth, say when $\alpha=\alpha_0$, then $B(\alpha_0)^{-1}=G$ and $\eta(t)=\xi(t)$.

Assume for simplicity that G^{-1} exists and call it B_0. Also, to simplify notation, assume $\mathrm{var}(\xi(t))=I$. Then the prediction error from using $B^{+}(\alpha)$ in (5) to predict Y is

$$\eta(t;\alpha) = \left(B^{+}(\alpha)-B_0^{+}\right)*Y(t) + \xi(t) . \tag{7}$$

Since $\xi(t)$ is uncorrelated with $Y(s)$ for $s<t$,

$$E\left[\eta(t;\alpha)'\eta(t;\alpha)\right] = \mathrm{tr}\left(I+\left(B^{+}(\alpha)-B_0^{+}\right)*R_Y*\left(B^{+}(\alpha)-B_0^{+}\right)'(0)\right). \tag{8}$$

This means that the limiting large-sample value of α arrived at by the usual fitting criterion will be the one which minimizes (translating the last part of (8) into the frequency domain)

$$\mathrm{tr}\int_0^{2\pi}\left(\tilde{B}^{+}(\omega;\alpha)-\tilde{B}_0^{+}(\omega)\right)S_Y(\omega)\left(\tilde{B}^{+}(\omega;\alpha)-\tilde{B}_0^{+}(\omega)\right)'d\omega . \tag{9}$$

This approximation error metric is formally similar to that which arises in the distributed lag regression problem. It is less easily applied to studying the effects of filtering the data, however, because filtration here affects not only S_Y, but also B_0.

Observe that $S_Y=\tilde{B}_0^{-1}\tilde{B}_0^{-1\prime}=\tilde{c}S_Y\tilde{c}'$. If c is a seasonal adjustment filter that reduces power in seasonal bands, we can expect S_Y and hence $|\tilde{B}_0|$ to be small in the seasonal bands.[1] As $\|\tilde{c}(\omega)\|\to 0$ for ω in the seasonal bands, for any fixed $\tilde{B}(\alpha)$ the two outer terms in the integrand in (9) are dominated by $\tilde{B}_0^+$. Furthermore, $\tilde{B}_0^+(\omega)-\tilde{B}_0(\omega)\equiv I$, whereas $\tilde{B}_0$ itself goes to infinity as $\|\tilde{c}(\omega)\|\to 0$. Thus the integrand in (9) approaches the identity matrix as $\|\tilde{c}(\omega)\|\to 0$. Since the forecast error covariance matrix using B_0 itself is the identity by assumption, the contribution to overall prediction error of approximation error in the seasonal band due to fixing $\tilde{B}(\omega;\alpha)$ arbitrarily in that band approaches, as $\|\tilde{c}(\omega)\|\to 0$ in seasonal bands, proportionality to the aggregate length of the seasonal bands as a fraction of the full $(0,2\pi)$ interval.

We conclude, then, that we can get an accurate picture of the behavior of $\tilde{B}_0$ outside the seasonal bands by using a parameterization $B(\cdot;\alpha)$ that

----

[1]One might think that "good" seasonal adjustment would only eliminate peaks at seasonals in S_y, not create dips at seasonals in S_y. However analysis of seasonal adjustment as a signal extraction problem shows that optimal seasonal adjustment is likely to generate dips in log S_Y of the same size and shape as the peaks in log S_y. Furthermore, here as in the distributed lag model, it appears that the effects of seasonal adjustment are most easily understood and controlled when there is "overadjustment" -- reduction of power in seasonal bands to near zero.

i) is capable of matching $\tilde{B}_0$ closely outside seasonal bands and

ii) cannot vary sharply over intervals as short as the seasonal bands,

and by seasonally adjusting the data to make power close to zero in seasonal bands. Our argument relies on the assumption that the seasonal bands are narrow enough to constitute a small part of the full $(0,2\pi)$ interval. Under these conditions, the best approximating $\tilde{B}(\cdot;\alpha)$ will fit well at non-seasonal frequencies and interpolate across the seasonal bands.

III. Application to Seasonal Rational Expectations Models

How much comfort we should take from the results of the preceding section depends on the extent to which conclusions about parameters of interest can be based on fitting the dynamic system accurately at non-seasonal frequencies and interpolating across seasonal bands. We can distinguish several cases.

One case, where seasonally adjusted data will never be useful, arises when the behavioral mechanism being studied itself generates some of the seasonality in the data. For example, if we were estimating a model of the construction industry with the aim of estimating the magnitude of cost increases for winter construction, the parameters of interest are related primarily to the seasonal bands of power in the data. Seasonal adjustment removes the part of the data that is informative about the behavior we wish to study.

We might instead have a situation where the behavior being studied does not inherently generate seasonality, but in which all variables can emerge with seasonal variation if exogenous variables have seasonal variation. Examples would be standard permanent income models and standard investment models with costs of adjustment. If the exogenous endowment process in a permanent income model contains seasonal variation, then consumption may as well. If the exogenous demand process in an investment model contains seasonal variation, then investment will as well. But neither model will generate seasonal variation when the exogenous driving process contains none.

In cases like these, where the source of the seasonality in endogenous variables is behavioral response to exogenous seasonality, the effects of seasonality may be confined to the seasonal bands. Suppose the agent in the model can separately observe the seasonal and nonseasonal components of the data, even though the econometrician cannot. This is perhaps not unrealistic when seasonality is due to some aspect of the weather, for example. If the behavioral model is linear-quadratic, the responses of endogenous variables to exogenous variables will be linear distributed lags, and these linear responses can be calculated independently for the separately observable exogenous forcing variables. Since the seasonal component of exogenous variation has its power concentrated in seasonal bands, the econometrician can (to a good approximation) correctly recover the structure of variation in the nonseasonal component of the exogenous variables by seasonally adjusting the data and computing autocorrelations or spectral densities from the adjusted data. Seasonally adjusting the endogenous variables will (to a good approximation) leave, at nonseasonal frequencies, just the responses to nonseasonal variation in the exogenous variables. Use of adjusted data to determine structural parameters will therefore give fairly accurate results.

On the other hand, even if endogenous seasonality is entirely due to responses to exogenous seasonality,seasonal variation will affect the nature of nonseasonal variation if agents cannot distinguish seasonal and nonseasonal components. Their reaction to a nonseasonal disturbance in this case is likely to be affected by their uncertainty as to whether it represents movement in the seasonal or nonseasonal component of the data. For example, in a permanent income model, if the nonseasonal component of income is highly persistent, then consumption will react almost one-for-one to random disturbances in income. If there is also a random seasonal component, however, the reaction of consumption to an income disturbance will be damped; the consumer will not know whether the income disturbance is persistent, coming from the nonseasonal component, or persistent only at a seasonal periodicity.

Note that under some conditions agents need not be able to observe the underlying forces generating seasonality in order to be able to distinguish it from nonseasonal variation. When seasonality is purely deterministic, or even stochastic but very slowly changing, it may contribute so little to the forecast error in the series that agents can observe the disturbances in the nonseasonal component with high accuracy. For example when the seasonal is purely deterministic, if the model is linear-quadratic, the behavioral equations can be written in terms of deviations from the deterministic seasonal pattern, and the presence of deterministic seasonality has no effect on the solution. Clearly use of "seasonally adjusted" data, where the adjustment is done by regression on seasonal dummy variables, creates no bias here. This is not to say it is always advisable -- the model will generally imply relations among the seasonal patterns in different variables. Using these implications (when the model is correct) will generally sharpen parameter estimates and strengthen tests.

What is true for strictly deterministic seasonals will also be true for stochastic seasonals that change so little or so slowly that errors in projecting the seasonal pattern are unimportant in overall short-term forecast error -- so long as seasonal adjustment is done by a method that properly accounts for the highly predictable nature of the seasonal.

IV. Dangers of Modeling Seasonality Directly

Using seasonally adjusted data can only be justified if there is substantial danger of making mistakes in attempting to model the unadjusted data. There is such danger. Because a seasonal component of variation is by definition one in which the same pattern of variation within the year tends to persist from year to year, apparently "good" forecasts of a seasonal component are easily available simply by extrapolating seasonal patterns from the last year or two. On the other hand, there may be little statistical evidence to distinguish among forecasting formulas for seasonal components even when the coefficients in those formulas appear very different. In quarterly data, for example,

$$\hat{Y}(t)=.5\left(Y(t-4)+Y(t-8)\right) \tag{10}$$

and

$$\hat{Y}(t)=.05\sum_{s=1}^{\infty}.95^{4s}Y(t-4s) \tag{11}$$

may have similar forecasting performance in macroeconomic time series samples, as both use a sort of average of past seasonal patterns in forming forecasts.

This situation, which is much like that arising in modeling trend components of series, tends to lead to practical modeling procedures in which the true weak identification of the stochastic structure of the seasonal is masked by strong assumptions of convenience about the form of seasonal variation. These assumptions of convenience will not harm forecasting performance much if incorrect, indeed they may improve forecasting performance by reducing "overparameterization". But they will in general have effects on the estimated form of the model which can strongly distort model interpretations, particularly in rational expectations models.

V. Examples

We illustrate the argument of the paper with some examples. We consider a simple standard linear-quadratic permanent income model, and focus attention on the implied regression coefficient of endowment (or "labor income") innovations on consumption innovations. When this coefficient is large, it implies that consumption is smoother than income, reacting less to a given disturbance than does income.

Suppose that agents maximize

$$E\left[\sum_{t=1}^{\infty}\left(C_t-.5C_t^2\right)\beta^t\right] \tag{12}$$

subject to

$$C_t+W_t = \rho W_{t-1} + Y_t \quad . \tag{13}$$

We follow convention in assuming that solutions exploding at a rate faster than $\beta^{-t/2}$ are ruled out, and for simplicity we consider the case $\beta=\rho^{-1}$. As is well known, this leads to the first order condition

$$E_t\left[C_{t+1}\right] = C_t \quad . \tag{14}$$

Note that equation (14) holds for the seasonally unadjusted data, even if seasonality is present in Y_t. The theory can be tested using (14), therefore, without any requirement that seasonality in Y be modeled explicitly. This situation is generic. In models where the behavioral mechanism is not generating seasonality, there are usually Euler equations implied by the theory that can be tested using unadjusted data without any explicit modeling of seasonality. However, as has been recognized in the literature on the permanent income model, such tests are weak, as they do not use the all the theory's implications about the relation between variation in endogenous variables (here C) and exogenous variables (here Y).

Suppose there is deterministic seasonality in Y, in particular a quarterly pattern of the form $1/\sqrt{2},0,-1/\sqrt{2},0$. Suppose the nonseasonal part of Y satisfies

$$Y_t = 1.8Y_{t-1} - .81Y_{t-1} + \varepsilon_t \ , \tag{15}$$

where $E_t\varepsilon_{t+1}\equiv 0$ and we normalize the variance of this nonseasonal part of Y at 1. Equation (15) gives Y's characteristic equation two roots of .9 and makes Y highly persistent. Y is so persistent, in fact, that the innovation in C (which by (14) is just ΔC_t) moves almost exactly one-for-one with the innovation in Y_t (which by (15) is ε_t). In fact, the regression of ε_t on ΔC_t is .99.

What happens if we generate data from this model but fit the wrong model of Y? In particular, say we do not include seasonal dummies in our model for Y but instead use a model implying all seasonality is stochastic, with Y a 9th order autoregression. The resulting estimated autoregressive model for Y will, in large samples, tend to the form

$$\begin{aligned} Y_t = {} & 1.792Y_{t-1}-1.1156Y_{t-2}+.5775Y_{t-3}+.0464Y_{t-4}-.5775Y_{t-5} \\ & -.0464Y_{t-6}+.5775Y_{t-7}-.2084Y_{t-8}-.0590Y_{t-9}+\varepsilon_t . \end{aligned} \tag{16}$$

Clearly this equation differs from (15), and in a way that allows it to correctly forecast most of the determistic seasonal variation in Y despite the absence of seasonal dummy variables. This equation also implies that a non-trivial part of the innovation in Y is due to a stochastic seasonal component, resulting in a theoretical value for the regression of ε_t on ΔC_t of 1.26. That is, the misspecified model implies substantially less reaction of consumption to income than the correct model. The use of the incorrect model (16) will not bias our estimate of the regression of the innovation on Y on the innovation in C, however. The model's structure implies that ΔC is the innovation in C, regardless of the autoregressive equation for Y. Using (16) to predict Y will produce suboptimal forecasts, but the forecast errors from (16) will consist of the true innovation in Y plus an additional component orthogonal to the true innovation. According to the model, ΔC and the true innovation are exactly collinear, so a regression of the suboptimal forecast errors from (16) on ΔC will produce the same coefficient as a regression of the residuals from (15) on ΔC.

We might be led to conclude, then, that the actual regression coefficient relating ε to ΔC in the data is smaller than that implied by the theory, since the true coefficient of .99 will be recovered from the data even if we mismeasure ε by estimating it from the 9th order AR with no seasonal dummies.

Note that in this model we would get correct results also from seasonally adjusted data, so long as we did not allow our predictive model for Y to "unravel the seasonal adjustment". Estimating a second-order AR for Y on the seasonally adjusted data would give approximately correct results, for example, so long as the seasonal bands removed by adjustment were not too wide.

Thus we have verified that in a model with a highly predictable seasonal, use of unadjusted data together with a model that cannot match the highly predictable nature of the seasonal, results in substantial bias. Use of seasonally adjusted data in this same situation would result in little bias and would not require precise knowledge of the structure of seasonal variation.

Of course, we can construct a model with precisely the reverse lesson. In fact, the biased 9th-order AR that would asymptotically be recovered from the data in the previous example was taken as a true data-generating mechanism to generate a second example. Now the coefficient of 1.26 in the regression of ε_t on ΔC_t represents the truth. In an infinitely large sample, *none* of the stochastically varying seasonality implied by this model could be captured in seasonal dummy variables. In a sample of size 100, however, nearly all of it will be captured by seasonal dummies with high probability. The dummies can capture seasonality that changes little. Eventually, the model implies the seasonal will change. Within a sample of size 100, the model implies a low probability of much change in the seasonal.

A second-order autoregression with 4 seasonal dummies was used as a model of Y and fit to 100 random samples from the data generating mechanism described above, which actually has no deterministic seasonality. The median estimate of the theoretical regression coefficient for ε on ΔC was 1.64, with interquartile range (1.26,2.14). Thus the median bias for using a simple deterministic seasonal adjustment scheme on these data containing stochastic seasonality was biased upward by a factor of 1.30, about the same as the asymptotic bias factor (1.26) when a stochastic seasonality model is used with unadjusted data containing a deterministic seasonal.[2]

---

[2]While it is interesting to note that both biases are in the same

VI. Conclusion

Our aim in this paper has been to show that seasonality is important, and that there is no way to avoid confronting the true depths of our ignorance about it. It is true that use of seasonally adjusted data in estimating rational expectations models can produce large bias, but in some situations the bias may be small. It is also true that use of unadjusted data and a correctly specified model of seasonal variation is always the best option. But conventional approaches to modeling stochastic seasonals may often produce incorrectly specified models of seasonal variation that appear to perform well. In such cases the bias from use of unadjusted data can be large, even though use of adjusted data would produce small bias.

----

direction, it is not generically true or even commonly true that any use of the wrong model for Y implies a larger regression coefficient for the regression of ε on ΔC. In fact, it is easy to generate examples where false imposition of a unit root on the model for Y, despite its being easily accepted by conventional statistical tests, produces very strong downward bias in the theoretical implied coefficient for the regression of ε on ΔC.

REFERENCES

Sims, Christopher A. [1974] "Seasonality in Regression," *J. American Statistical Association, 69*, September, 618-626.

Wallis, Kenneth [1974]. "Seasonal Adjustment and Relations Between Variables", *Journal of the American Statistical Association 69*, March, 18-32.