Journal of Econometrics 55 (1993) 129–133. North-Holland

Discussion*

The importance of seasonality in inventory models

Jean-Marie Dufour
C.R.D.E., Université de Montréal, Montréal, Québec, Canada H3C 3J7

The article by Nerlove, Ross, and Willson in this issue pulls together three topics on which Nerlove and his collaborators have already done important work in the past: the statistical analysis of qualitative survey data, inventories, and seasonal fluctuations. As we can expect, it is of high quality and the authors should be congratulated for this interesting piece.

In this comment, I will first summarize the main findings of the paper (section 1) and then discuss three issues that relate to the econometric methodology employed and to possible extensions: namely, the measurement of the seasonal effects (section 2), the two-stage estimation procedure (section 3), and the problem of serial dependence (section 4). Section 5 contains a few additional remarks.

1. Findings

The Swiss Business Survey is a rich microlevel database covering 14 years of a monthly survey which includes a large number of questions. The authors are certainly among the first ones to study these data carefully with econometric methods. Of course, they can only exploit a small portion of the information in the survey. In this paper, they study the behavior of four inventory variables in relation with demand and production variables. The data are analyzed in two different ways: first, by the estimation of (pairwise) polychoric correlation coefficients, through a two-stage procedure; second, by

*This work was supported by the Centre de Recherche et Développement en Économique (C.R.D.E., Université de Montréal), the Natural Sciences and Engineering Research Council of Canada, the Social Sciences and Humanities Research Council of Canada, and the Government of Québec (Fonds FCAR).

0304-4076/93/$05.00

the estimation of an equation derived from a simple dynamic net inventory investment model.

The analysis of the polychoric correlation coefficients is strongly suggestive and yields several interesting findings:

(1) Changes in unfilled orders are strongly correlated with demand (new orders) and production changes, so that they appear to play a central role as a primary buffer stock.
(2) Similarly, but to a smaller extent, work-in-process and inventory changes are strongly correlated with production and demand changes, suggesting that they also play an important role in production and inventory dynamics.
(3) Distinguishing between expected and unexpected changes in demand and production appears to be important, as these two types of variables often have quite different correlations with the inventory variables.
(4) There is a strong positive correlation between changes in expected demand and changes in expected production.
(5) Controlling for seasonal effects appears to make little difference to the results, an observation from which the authors conclude that firms do not use inventories to smooth seasonal movements in demand.

The results from the net investment model confirm the importance of expected demand changes in explaining expected production changes. Contrary to an assumption often made in inventory studies, they show also that unfilled orders and finished goods inventories do not have symmetrically inverse effects on expected production changes. Further, seasonal dummies are highly significant.

2. Measurement of seasonal effects

An important objective of the paper is to study whether inventory behavior indicates that firms smooth their production in response to seasonal demand fluctuations. The presence of seasonal smoothing is however appreciated in a very indirect way: the authors estimate the polychoric coefficients while allowing the thresholds (associated with the three observed 'levels' of variable changes) to vary across seasons (e.g., months), and they compare the correlations obtained in this way with those obtained when the thresholds are restricted to be the same for all seasons. Changing thresholds between seasons, however, may well capture seasonal movements in the variables studied, such as inventories, and provide information on seasonal smoothing. This suggests comparing thresholds obtained for different seasons for each variable and between variables, as done in fact by one of the authors in earlier work [Nerlove (1988)]. For example, the seasonal thresholds for finished-goods inventories may be compared with those for new orders. Even though the authors mention (in the latest version of their paper) that they calculated the thresholds for ΔU (changes in unfilled orders) and ΔD (changes in new orders) and found 'little relationship in the seasonal movements of the thresholds', it is clear that this type of evidence should be better documented (with appropriate statistical tests) and that a larger set of variable pairs should be considered.

It is also of interest to note that the thresholds are monotonic transformations of the marginal probabilities associated with the different responses:

$$a_{11} = \Phi_1^{-1}[\mathrm{P}(z_1 = 1)], \qquad a_{21} = \Phi_1^{-1}[1 - \mathrm{P}(z_1 = 3)],$$
$$a_{12} = \Phi_1^{-1}[\mathrm{P}(z_2 = 1)], \qquad a_{22} = \Phi_1^{-1}[1 - \mathrm{P}(z_2 = 3)],$$

where Φ_1^{-1} is the inverse of the cumulative N(0, 1) distribution function. In the two-stage procedure used by the authors, the marginal probabilities $\mathrm{P}(z_1 = j)$ and $\mathrm{P}(z_2 = j)$ are first estimated by the marginal frequencies $n_{j\cdot}/n$ and $n_{\cdot j}/n$ and then transformed with Φ_1^{-1} to get estimates of the thresholds. Clearly seasonal movements in the thresholds could as well be analyzed by looking directly at these marginal frequencies (for different seasons). All this analysis could be based on binomial distributions and contingency table methods.

Along similar lines, one may also allow the correlation coefficients to vary across seasons, and then estimate the latter from the different subsamples corresponding to the different seasons. Again seasonal effects can be analyzed by comparing the coefficients across seasons.

3. Two-stage estimation procedure

In order to estimate the polychoric coefficients, the authors use a two-stage procedure: first, the thresholds are estimated from the marginal frequencies; second, for each pair of variables, the polychoric correlation is estimated by maximum likelihood (ML) taking the estimated thresholds as the true ones. This raises two difficulties. First, the separate estimation of the different correlations does not guarantee that the correlation matrix will be positive semidefinite. Second, since the uncertainty in the estimation of the thresholds is not taken into account, standard errors tend to be underestimated, so that t-tests are too liberal even asymptotically [see Durbin (1970), Gong and Samaniego (1981), and Pagan (1984, 1986)].

An obvious way of dealing with the latter problem consists in estimating jointly the thresholds and the correlations. Since each correlation is obtained by considering only two variables at a time (instead of the whole vector), computing the likelihood function requires evaluating bivariate normal distribution functions. Is it the case that this is infeasible here?

If indeed full ML estimation is too costly, a possible solution would be to replace the Wald-type tests (asymptotic t-tests) by Neyman's (1959) $C(\alpha)$ tests. The $C(\alpha)$ test can be viewed as a generalization of Rao's score test (or the Lagrange multiplier test) where the restricted estimator of model parameters need only be root-n-consistent (under the null hypothesis). So one does not need to find the restricted ML estimator of the parameters. For example, let $\theta = (\rho_{12}, a')'$ be the parameter vector involved in the comparison of the latent variables z_1^* and z_2^*, where $a = (a_{11}, a_{21}, a_{12}, a_{22})'$ is the vector of thresholds parameters and ρ_{12} is the correlation of interest. Then a simple restricted root-n-consistent estimator is given by $\tilde{\theta} = (0, \hat{a}')'$, where $\hat{a}$ is the vector of estimated thresholds obtained by the authors in their first stage. The $C(\alpha)$ statistic can be written

$$C(\tilde{\theta}) = \frac{1}{n}\left(D'\tilde{I}^{-1}\tilde{D} - \tilde{D}_2'\tilde{I}_{22}^{-1}\tilde{D}_2\right),$$

where $\tilde{D} = D(\tilde{\theta})$, $D(\theta) = \partial(\ln L)/\partial\theta$, $\tilde{I} = \hat{I}(\tilde{\theta})$, $\hat{I}(\theta)$ is a consistent estimator of the information matrix associated with the log-likelihood function $\log L$, and $\tilde{D}_2$ and $\tilde{I}_{22}$ are the submatrices of $\tilde{D}$ and $\tilde{I}$ corresponding to the nuisance parameters a. Under standard regularity conditions, the asymptotic distribution of $C(\tilde{\theta})$ is $\chi^2(1)$. For further discussion and references, see Dagenais and Dufour (1991).

4. Serial dependence

As pointed out by the authors (section 4), the observations may not be independent because in the pooled sample 'some (but not all) firms respond in adjacent months'. In other words, specific factors (or shocks) affecting individual firms may be serially correlated. Note also that correlations between the observations can also originate from common factors that fluctuate across seasons (e.g., business cycle fluctuations).

To deal with autocorrelation, the authors suggest that the 'total pooled sample size' (of 25,000 observations) could be replaced by an effective sample size corresponding to the 'average monthly sample size' (approximately 500 observations), in order to compute standard errors: this leads to dividing the t-statistics by approximately 7. (Note that this correction is not applied in the reported results.) I must say that I fail to see why this is an appropriate correction: no statistical justification is presented.

A possibly more rigorous way of checking the robustness of the results to serial dependence may be to split the sample according to seasons and check whether the results remain the same. This may also yield interesting information on possible structural changes. Of course, one can also try to model the serial dependence: autocorrelations between observations (e.g., taken from identical firms at different points in time) may be estimated in a way similar to the polychoric correlation coefficients. But obviously this is a more complex operation, both conceptually and computationally.

5. Other remarks

Correlations between observations may also become industry-specific factors (shocks common to firms of the same industry). Again this may affect the validity of standard errors computed from the pooled sample. Further seasonal behavior and correlations can differ across industries. This suggests studying subsamples corresponding to different industries. In the same vein, industry-specific dummies may be a useful addition to the equation derived from the simple model of net inventory investment behavior.

Finally, in section 5, the authors point out that the Swiss survey contains information on cost shocks, such as 'actual and expected changes in raw materials prices', which however they 'found unrelated to changes in production plans, or movements in the various inventory stocks'. Since this type of observation may be relevant to the evaluation of business cycle theories (e.g., theories that emphasize real factors), it would certainly be useful to document this type of observation more fully.

References

Dagenais, M. and J.-M. Dufour, 1991, Invariance, nonlinear models and asymptotic tests, Econometrica 59, 1601–1615.

Durbin, J., 1970, Testing for serial correlation in least squares regression when some of the regressors are lagged dependent variables, Econometrica 38, 410–421.

Gong, G. and F.J. Samaniego, 1981, Pseudo maximum likelihood estimation: Theory and applications, Annals of Statistics 9, 861–869.

Nerlove, M., 1988, Analysis of business-test survey data by means of latent variable models, in: W. Franz, W. Gaab, and J. Wolters, eds., Theoretische und angewandte Wirtschaftsforschung: Heinz Koenig zum 60. Geburstag (Springer-Verlag, Berlin) 241–259.

Nerlove, M., D. Ross, and D. Willson, 1993, The importance of seasonality in inventory models: Evidence from business survey data, Journal of Econometrics, this issue.

Neyman, J., 1959, Optimal asymptotic tests of composite statistical hypotheses, in: U. Grenander, ed., Probability and statistics: The Harald Cramer volume (Almqvist and Wiksell, Uppsala) 213–234.

Pagan, A.R., 1984, Econometric issues in the analysis of regressions with generated regressors, International Economic Review 25, 221–247.

Pagan, A.R., 1986, Two stage and related estimators and their applications, Review of Economic Studies 53, 517–538.