Chapter Title: A Dynamic Index Model for Large Cross Sections

Chapter Author: Danny Quah, Thomas J. Sargent

Chapter URL: http://www.nber.org/chapters/c7195

Chapter pages in book: (p. 285 - 310)

7   A Dynamic Index Model for Large Cross Sections

Danny Quah and Thomas J. Sargent

In studying macroeconomic business cycles, it is convenient to have some low-dimensional characterization of economy-wide fluctuations. If an aggregate measure of economic activity—such as GNP or total industrial production—could be widely agreed on as accurately assessing the state of the economy, then all attention could focus on that measure alone. Econometric analysis and forecasting would then be, in principle, straightforward.

But any one given aggregate measure is likely affected by many different disturbances: understanding economic fluctuations then involves disentangling the dynamic effects of those different disturbances. For instance, much has been written on decomposing aggregate measures of output into different interpretable components. Some of this work uses multivariate time-series methods to analyze such decompositions. Little empirical work has been done, however, proceeding in the opposite direction, that is, using information from a wide range of cross-sectional data to shed light on aggregate fluctuations. In other words, in the business-cycle analysis triad of *depth, duration,* and *dispersion,* the third has typically been ignored.

We know that professional forecasters look at a variety of different indicators to predict aggregate activity. Interest rates, measures of consumer sentiment, the money supply, and the spectrum of asset prices are all candidates for helping predict economy-wide movements. There again, to understand

Danny Quah is reader in economics at the London School of Economics. Thomas J. Sargent is a senior fellow at the Hoover Institution, Stanford University, and professor of economics at the University of Chicago.

aggregate fluctuations, one needs to go beyond examining a single time series in isolation.

This paper takes the natural next step in these observations: it provides a framework for analyzing commonalities—aggregate comovements—in dynamic models and data structures where the cross-sectional dimension is potentially as large, and thus potentially as informative, as the time-series dimension. Why might this be useful for analyzing macroeconomic business cycles?

When the aggregate state of the economy affects and is in turn affected by many different sectors, then it stands to reason that *all* those sectors should contain useful information for estimating and forecasting that aggregate state. It should then simply be good econometric practice to exploit such cross-sectional information jointly with the dynamic behavior to uncover the aggregate state of the economy.

Standard time-series analysis, however, is not well suited to analyzing large-scale cross sections; rather, it is geared toward the study of fixed-dimensional vectors evolving over time. Cross-sectional and panel data analysis, on the other hand, are typically not concerned with the dynamic forecasting exercises that are of interest in macroeconomics. It is this gap in macroeconometric practice between cross-sectional and time-series analyses that this paper addresses. While there are many interesting issues to be investigated for such data structures, we focus here on formulating and estimating dynamic index models for them.

The plan for the remainder of this paper is as follows. Section 7.1 relates our analysis to previous work on index structures in time series. Section 7.2 shows how (slight extensions of) standard econometric techniques for index models can be adapted to handle random fields. Section 7.3 then applies these ideas to study sectoral employment fluctuations in the U.S. economy. Appendices contain more careful data description as well as certain technical information on our application of the dynamic index model to random fields.

7.1   Dynamic Index Models for Time Series

This paper has technical antecedents in two distinct literatures. The first includes the dynamic index models of Geweke (1977), Sargent and Sims (1977), Geweke and Singleton (1981), and Watson and Kraft (1984). The second literature concerns primarily the analysis of common trends in cointegrated models, as in, for example, Engle and Granger (1987), King et al. (1991), and Stock and Watson (1988, 1990). Because of the formal similarity of these two models—the dynamic index model and the common trends model—it is sometimes thought that one model contains the other. This turns out to be false in significant ways, and it is important to emphasize that these models represent different views on disturbances to the economy. We clarify

this in the current section in order to motivate the index structure that we will end up using in our own study.

Notice first that neither of these standard models is well suited to analyzing structures having a large number of individual time series. When the number of individual series is large relative to the number of time-series observations, the matrix covariogram cannot be precisely estimated, much less reliably factored—as would be required for standard dynamic index analysis.

Such a situation calls for a different kind of probability structure from that used in standard time-series analysis. In particular, the natural probability model is no longer that of vector-valued singly indexed time series but instead that of multiply indexed stochastic processes, that is, *random fields*. In developing an index model for such processes, this paper provides an explicit parameterization of dynamic and cross-sectional dependence in random fields. The analysis here thus gives an alternative to that research that attempts robust econometric analysis without such explicit modeling.

Begin by recalling the dynamic index model of Geweke (1977) and Sargent and Sims (1977). An $N \times 1$ vector time series X is said to be generated by a K-index model ($K < N$) if there exists a triple (U, Y, a) with U $K \times 1$ and Y $N \times 1$ vectors of stochastic processes having all entries pairwise uncorrelated and with a an $N \times K$ matrix of lag distributions, such that

$$(1) \qquad\qquad X(t) = a * U(t) + Y(t),$$

with $*$ denoting convolution. Although jointly orthogonal, U and Y are allowed to be serially correlated; thus (1) restricts X only to the extent that K is (much) smaller than N. Such a model captures the notion that individual elements of X correlate with each other only through the K-vector U. In fact, this pattern of low-dimensional cross-sectional dependence *is* the defining feature of the dynamic index model.

Such a pattern of dependence is conspicuously absent from the common trends model for cointegrated time series. To see this, notice that, when X is individually integrated but jointly cointegrated with cointegrating rank $N - K$, then its common trends representation is

$$(2) \qquad\qquad X(t) = aU(t) + Y(t),$$

where U is a $K \times 1$ vector of pairwise orthogonal random walks, Y is covariance stationary, and a is a matrix of numbers (Engle and Granger 1987; or Stock and Watson 1988, 1990). As calculated in the proof of the existence of such a representation, U has increments that are perfectly correlated with (some linear combination of) Y. It is not hard to see that there need to be no transformation of (2) that leaves the stationary residuals Y pairwise uncorrelated and uncorrelated with the increments in U. Thus, unlike the original index model, representation (2) is ill suited for analyzing cross-sectional dependence in general. In our view, it is not useful to call Y idiosyncratic or

sector specific when Y turns out to be perfectly correlated with the common components U.

7.2   Dynamic Index Structures for Random Fields

While model (1) as described above makes no assumptions about the stationarity of X, its implementation in, for example, Sargent and Sims (1977) relied on being able to estimate a spectral density matrix for X. In trending data, with permanent components that are potentially stochastic and with different elements that are potentially cointegrated, such an estimation procedure is unavailable.

Alternatively, Stock and Watson (1990) have used first-differenced data in applying the index model to their study of business-cycle coincident indicators—but, in doing so, they restrict the levels of their series to be eventually arbitrarily far from each series's respective index component. Stock and Watson chose such a model after pretesting the data for cointegration and rejecting that characterization. Thus, their non-cointegration restriction might well be right to impose, but, again, the entire analysis is thereafter conditioned on the variables of interest having a particular cointegration structure.

When the cross-sectional dimension is potentially large, standard cointegration tests cannot be used, and thus the researcher cannot condition on a particular assumed pattern of stochastic trends to analyze dynamic index structure. (For instance, if N exceeds T, the number of time-series observations, no [sample] cointegrating regression can be computed.) Our approach explicitly refrains from specifying in advance any particular pattern of permanent and common components: instead, we use only the orthogonality properties of the different components to characterize and estimate an index model for the data. It turns out that, by combining ingredients of (1) and (2), one obtains a tractable model having three desirable features: (1) the cross-sectional dependence is described by some (fixed) low-dimensional parameterization; (2) the data can have differing, unknown orders of integration and cointegration; the model structure should not depend on knowing those patterns; and (3), finally, the key parameters of the model are consistently estimable even when N and T have comparable magnitudes.

Let $\{X_j(t), j = 1, 2, \ldots, N, t = 1, 2, \ldots, T\}$ be an observed segment of a random field. We are concerned with data samples where the cross-sectional dimension N and the time-series dimension T take on the same order of magnitude. We hypothesize that the dependence in X across j and t can be represented as

$$(3) \qquad\qquad X_j(t) = a_j * U(t) + Y_j(t),$$

where U is a $K \times 1$ vector of orthogonal random walks; Y_j is zero mean, stationary, and has its entries uncorrelated across all j as well as with the increments in U at all leads and lags; and, finally, a_j is a $1 \times K$ vector of lag

distributions. Model (3) is identical to (1) except that (3) specifies U to be orthogonal random walks. A moment's thought, however, shows that the random walk restriction is without loss of generality: an entry in a_j can certainly sum to zero and thereby allow the corresponding element in U to have only transitory effects on X_j.

As is well known, model (3) restricts the correlation properties of the data: in particular, the number of parameters in X's second moments grows only linearly with N, rather than quadratically, as would be usual.[1] Further, it is useful to repeat that (3) differs from the common trends representation (2) in having the j-specific disturbances uncorrelated both with each other and with U; neither of these orthogonality properties holds for (2). Thus, (3) shares with the original index model (1) the property that all cross-correlation is mediated only through the lower-dimensional U—independent of N and T. Such a property fails for the common trends representation.

Since Y_j is orthogonal to U, (3) is a projection representation, and thus (relevant linear combinations of) the coefficients a_j can be consistently estimated by a least squares calculation: this holds regardless of whether X_j is stationary, regardless of the serial correlation properties of Y_j, and regardless of the distributions of the different variables. Notice further that, for this estimation, one does not require observations on U, but only consistent estimates of the conditional expectations of certain cross-products of U and X_j. Finally, since Y_j is uncorrelated across j's and the regressors U are common, equation-by-equation estimation is equivalent to joint estimation of the entire system. Many of these properties were already pointed out by Watson and Engle (1983); here we add that, if the parameters of (3) are consistently estimated as $T \to \infty$ independent of the value of N, then they remain consistently estimated even when N has the same order of magnitude as T.

The estimation problem is made subtle, however, by the U's not being observable. Nevertheless, using the ideas underlying the EM algorithm (see, e.g., Watson and Engle 1983; Wu 1983; or Ruud 1991), model (3) can easily be estimated. The key insight has already been mentioned above, and that is that, to calculate the least squares projection, one needs to have available only sample moments, not the actual values of U. When, in turn, these sample moments cannot be directly calculated, one is sometimes justified in substituting in their place the conditional expectation of these sample moments, conditional on the observed data X. Note, however, that these conditional expectations necessarily depend on all the unknown parameters. Thus, using them in place of the true sample moments, and then attempting to solve the first-order conditions characterizing the projection—even equation by equation— could again be intractable, for large N.

Instead, treat the projections in (3) as solving a quasi-maximum likelihood

---

1. If necessary, such statements concerning moments can be read as *conditional on initial values*. For brevity, however, such qualification is understood and omitted in the text.

problem, where we use the normal distribution as the quasi likelihood. Then the reasoning underlying the EM algorithm implies that, under regularity conditions, the algorithm correctly calculates the projection in (3).

This procedure has two features worthy of note: first, estimation immediately involves only least squares calculations, equation by equation, and thus can certainly be performed with no increased complexity for large N. To appreciate the other characteristic, one should next ask, Given this first property, how has the cross-sectional information been useful? The answer lies of course in the second feature of this procedure—the cross-sectional information is used in calculating the conditional expectation of the relevant sample moments.

To make this discussion concrete, we now briefly sketch the steps in the estimation. These calculations are already available in the literature; our contribution here has been just to provide a projection interpretation for the EM algorithm for estimating models with unobservables—thus, to distance the EM algorithm from any tightly specified distributional framework—and to note its applicability to the case when N is large.

Write out (3) explicitly as

$$(4) \qquad X_j(t) = \sum_{k=1}^{K} \alpha_{jk}(L)U_k(t) + Y_j(t);$$

in words, each X_j is affected by K common factors $U = (U_1, U_2, \ldots, U_K)'$ and an idiosyncratic disturbance Y_j. This gives a strong-form decomposition for each observed X_j into unobserved components: the first, in U, common across all j and the second Y_j specific to each j and orthogonal to all other X's. It will be convenient sometimes to add a term $d_j W(t)$ to (4), where W comprises purely exogenous factors, such as a constant or time trend.

The common factors in (4) have dynamic effects on X_j given by

$$\alpha_{jk}(L) = \sum_{m=0}^{M_a} a_{jk}(m)L^m, \quad M_a \text{ finite.}$$

Take U_k's to be integrated processes, with their first-differences pairwise orthogonal and having finite autoregressive representation

$$(5) \qquad \Gamma(L)\Delta U(t) = \eta_U(t),$$

where $\Gamma(L)$ is diagonal, with the kth entry given by

$$1 - g_k(1)L - g_k(2)L^2 - \ldots - g_k(M_g)L^{M_g} \quad M_g \text{ finite,}$$

and η_U is K-dimensional white noise having mean zero and the identity covariance matrix. As suggested above, the normalizations following (5) are without loss of generality because U affects observed data only after convolution with a. If X_j were stationary, the sequences a_j—provided that $M_a \geq 1$—

can contain a first-differencing operation to remove the integration built into (5).

Recall that the idiosyncratic or sector-specific disturbances Y_j are taken to be pairwise orthogonal as well as orthogonal to all ΔU. We now assume further that each Y_j has the finite autoregressive representation

$$(6) \qquad \beta_j(L)Y_j(t) = \varepsilon_j(t), \quad \text{for all } j,$$

with

$$\beta_j(L) = 1 - b_j(1)L - b_j(2)L^2 - \ldots - b_j(M_b)L^{M_b}, \quad M_b \text{ finite.}$$

Combining (4) and (5), and defining $\phi_{jk}(L) = \beta_j(L)\alpha_{jk}(L)$, we get

$$\beta_j(L)X_j(t) = \sum_{k=1}^{K} \phi_{jk}(L)U_k(t) + \varepsilon_j(t).$$

Notice that, because α and β are individually unrestricted, we can consider the model to be parameterized in β and ϕ, being careful to take the ratio ϕ/β whenever we seek the dynamic effects of U on X. When W—the set of exogenous variables—is lag invariant, as would be true for a time trend and constant, then $\beta_j(L)[d_jW(t)]$ spans the same space as just $W(t)$. Thus, without loss of generality, we can write the model as simply

$$(7) \qquad \beta_j(L)X_j(t) = \sum_{k=1}^{K} \phi_{jk}(L)U_k(t) + \varepsilon_j(t) + d_jW(t)$$

(possibly redefining d_j). Subsequent manipulations will exploit this representation (7) of the original strong-form decomposition of X given in (4).

Directly translating the model above into state space form—we do this in Appendix A below—gives

$$(8) \qquad X(t) = AZ(t) + dW(t) + Y(t),$$
$$Z(t + 1) = cZ(t) + \eta(t + 1).$$

The state vector Z, unfortunately, has dimension $O(N)$: this is computationally intractable for data structures with large N and T. Conceptually more important, however, such a state space representation (8) simply describes N time series in terms of an $O(N)$-dimensional vector process—certainly neither useful nor insightful. The reason for this is clear: from (7), lagged X_j's are part of the state description for each X_j owing to the β_j's being nontrivial, that is, owing to the serial correlation in each idiosyncratic disturbance.[2]

The solution to this is to notice that a large part of the state vector happens

---

2. This serial correlation pervades our model in particular and business-cycle data in general. It is this feature that distinguishes our work from many EM algorithm/unobservable common factor applications in finance, e.g., Lehmann and Modest (1988), where the data are either only a large cross section or, alternatively, taken to be serially independent.

to be directly observable; the unobservable part turns out to have dimension only $O(K)$, independent of N. Exploiting this structure gives Kalman smoother projections and moment matrix calculations that are, effectively, independent of N. (The assertions contained in this and the previous paragraph are detailed in Appendix A.)

The intuition then is that increasing the cross-sectional dimension N can only help estimate more precisely (the conditional expectations of) the (unobserved part of the) state and its cross-moments. This must imply the same property for estimates of the parameters of interest. Notice further that deleting entries of X leaves invariant the orthogonality properties on an appropriate reduced version of (8). Thus, if the model is correctly specified, estimators that exploit the orthogonality conditions in (8) remain consistent independent of N; at the same time, smaller N-systems must imply less efficient estimators. This observation is useful for building up estimates for a large N-system from (easier to estimate) smaller systems; additionally, it suggests a Wu-Hausman-type specification test for these index structures.[3]

Next, when the unknown distribution of (X, U, Y, W) generates a conditional expectation $E(U|X, W)$ that equals the linear projection of U on X and W, then standard Kalman smoother calculations yield the conditional expectations $E[Z(t)Z(t)'|X, W]$ and $E[Z(t)|X, W]$, taking as given a particular setting for the unknown parameters. We will assume that the underlying unknown distribution does in fact fall within such a class.[4]

Iteration on this scheme is of course simply the EM algorithm and, under weak regularity conditions, is guaranteed to converge to a point that solves the first-order (projection) conditions. Estimation of the dynamic index model for random fields with large N and T is thus seen to be feasible.

7.3   An Application: Sectoral Employment

This section gives a sectoral, employment-disaggregated description of U.S. economic fluctuations as interpreted by our index-model structure. We consider the behavior of annual full-time equivalent (FTE) employment across sixty industries.

The national income and product accounts (NIPA) report the number of full-time equivalent workers in different industries on an annual basis (NIPA, table 6.7b). Excluding the government sector and U.S. workers employed outside this country, there are sixty industrial sectors for which these data are

---

3. We do not follow up this testing suggestion below, leaving it instead for future work.

4. Elliptically symmetrical distributions would be one such class. Some researchers, such as Hamilton (1989), prefer to work with state vectors that have a discrete distribution—this of course need not invalidate the hypothesized equality of conditional expectations and linear projections, although it does make that coincidence less likely. In principle, one can directly calculate those conditional expectations using only laws of conditional probability. Thus, any simplification—whether by using linear projections or by discrete distributions—is not conceptual but only computational.
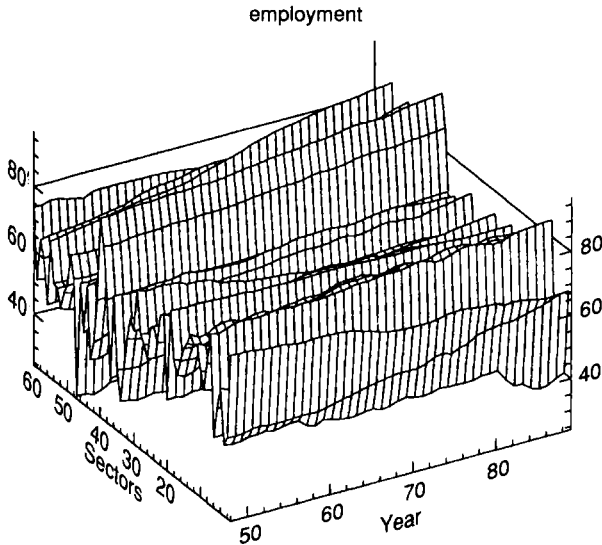
Fig. 7.1   Log (times ten) of employment, across sector and time

available from 1948 through 1989. The industries range from farming, metal mining, and coal mining through motion pictures, legal services, and private household services. (The entire set of sixty industries—complete with Citibase codes—is reported in Appendix B.) Figure 7.1 gives a three-dimensional depiction of these FTE employment data; the "sectors" axis in this graph arrays the different industries in the order given in NIPA table 6.7b and Appendix B. The vertical axis is the log (times ten) of employment in each sector. It is possible to give a more traditional time-series plot of these sixty series. Doing so, however, reveals very little information: the different time-series lines quickly merge, and the page is simply awash in black.

Since these employment figures are available only annually from 1948 through 1989, our data are larger in the cross-sectional than in the time-series dimension ($60 = N > T = 42$). No full rank (60×60) covariance matrix estimator is available; consequently, no full-rank spectral density estimator can be formed. From figure 7.1, it is also clear that the data are trending, with potentially differing orders of stochastic and deterministic permanent components. Again by $N > T$, no cointegrating regression could be calculated; no cointegration tests could be performed.

Figures 7.2–7.4 explore the extent to which the cross-correlation across sectors can be captured by two observable measures typically used by empirical researchers: first, (the log of annual) real GNP and, second, (the log times ten of) total—equivalently, average—employment across the sixty sectors. When we refer to total employment subsequently, we mean this second series, not total U.S. employment. These figures plot residual sample standard devia-

Standard·Deviation

Fig. 7.2   **Residual standard deviations projecting on own two lags plus GNP,
against those projecting on just own two lags**

tions from second-order autoregressions, fitted independently across individual sectors, over 1951–87, and always including a constant and time trend. (In a study with as high dimensionality as this one, presenting alternative specifications—varying lag lengths, e.g.—quickly becomes awkward; the signal-noise ratio in presentation falls rapidly and dramatically. Unless stated otherwise, the main conclusions hereafter should be taken as robust across small lag length increases.)

Figure 7.2 graphs the residual sample standard deviation when values for GNP at lag -2 through lag 2 are included as additional regressors (on its vertical axis) and when they are not (on the horizontal axis).[5] By least squares algebra, no point in figure 7.2 can lie above the forty-five-degree line. The further, however, that points fall *below* the forty-five-degree line, the more successfully does GNP—common to all sectors—explain employment in

---

5. Notice that these regressions include past, present, and future values of GNP. Below, we shall compare the residual variances from these regressions with residual variances from regressions on estimates of our indexes. Those estimates are projections of our indexes on past, present, and future values of all the employment series.
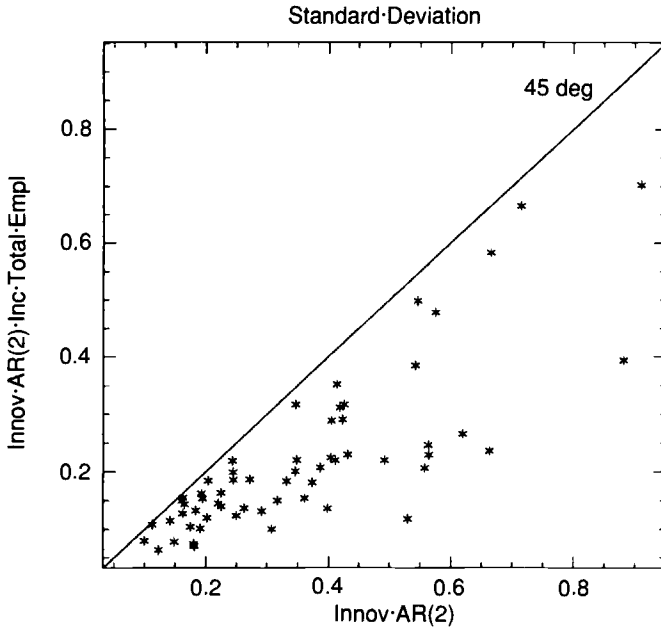
Fig. 7.3   Residual standard deviations projecting on own two lags plus total employment, against those projecting on just own two lags

each sector. From this graph, we conclude that aggregate GNP does appear to be an important common component in sectoral employment fluctuations.

Figure 7.3 is the same as 7.2, except GNP is replaced by total employment across the sixty sectors. The message remains much the same. There appear to be common comovements in sectoral employment, and those comovements are related to aggregate GNP and total employment movements. We emphasize that, in the regressions above, both lagged and lead aggregate measures enter as right-hand-side variables. The sample standard deviations increase significantly when lead measures are excluded.

Figure 7.4 compares these two measures of common comovements by plotting against each other the sample standard deviations from the vertical axes of figures 7.2 and 7.3. We conclude from the graph here that both GNP and total employment give similar descriptions of the underlying comovements in sectoral employment.

For the period 1948–89, we have estimated one- and two-index representations for sectoral employment.[6] (In the notation of the previous section, we take $M_a = 1$, $M_g = 1$, and $M_b = 2$; again, small increases in lag lengths do

---

6. All index-model and regression calculations and all graphs in this paper were executed using the authors' time-series, random-fields econometrics shell tsrf.

Fig. 7.4   Residual standard deviations projecting on own two lags plus GNP, against those projecting on own two lags plus total employment
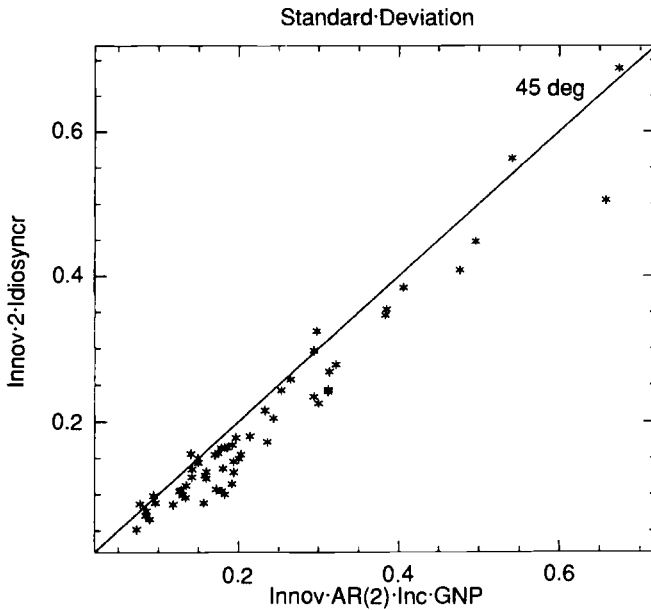

Fig. 7.5   Standard deviations of idiosyncratic disturbances in two-index model, against those of the residuals in the projection on own two lags plus GNP

not affect our conclusions materially.) Figure 7.5 plots standard deviations of the innovations in the idiosyncratic disturbances $\varepsilon_j(t)$, under the two-index representation, against the residuals in sector-by-sector projections including GNP (i.e., the vertical axis of fig. 7.2). In other words, the vertical axis describes the innovations on removing two common unobservable indexes and imposing extensive orthogonality conditions; the horizontal axis describes the innovations on removing the single index that is GNP and without requiring the resulting innovations to be orthogonal across sectors.[7] Since the models are not nested, there is no necessity for the points to lie in any particular region relative to the forty-five-degree line. Again, however, to the extent that these points fall below that line, we can conclude that the comovements are better described by the model represented on the vertical axis than that on the horizontal.

In this figure, twelve sectors lie marginally above the forty-five-degree line, six marginally below, and the remainder quite a bit below. Overall, we conclude that the two unobservable indexes provide a better description of underlying commonalities in sectoral employment than does aggregate GNP.

Figure 7.6 is the same as figure 7.5 except that GNP is replaced by total employment. We draw much the same conclusions from this as the previous graph.

Figure 7.7 replaces the horizontal axes of figures 7.5 and 7.6 with the standard deviation of idiosyncratic innovations from a single-index representation. Notice that the improvement in fit of the additional index is about the same order of magnitude as that of the two-index representation over either of the single observable aggregates.

We do not present here the calculations that we have performed comparing the single unobservable index with the two observable aggregates. The calculations show much what one would expect from the analysis thus far. Aggregate GNP and average employment are about as good descriptions of sectoral comovements as is the single unobservable index model.[8]

So what have we learned up until now? If one were concerned only about goodness of fit in describing the commonalities in sectoral employment fluctuations, then one might well simply use just total GNP or average employment.[9] In our view, index models should be motivated by a Burns-Mitchell kind of "pre-GNP accounting," dimensionality-restricted modeling. The idea

7. Thus, the vertical dimension of fig. 7.5 contains enough information to compute a normal quasi-likelihood value for the unobservable index model; for the horizontal dimension, however, the cross-correlations are nonzero and cannot, as a whole, be consistently estimated.

8. This finding is related to patterns detected in previous applications of unobservable index models to aggregate U.S. time series. For example, Sargent and Sims (1977) encountered so-called Heywood solutions at low frequencies: in those solutions, the coherence of GNP with one of the indexes is unity.

9. We should emphasize again that this is true only when one proxies for those common comovements using both leads and lags of these indicators. As already stated above, the projection residuals get considerably larger when future aggregates are excluded.
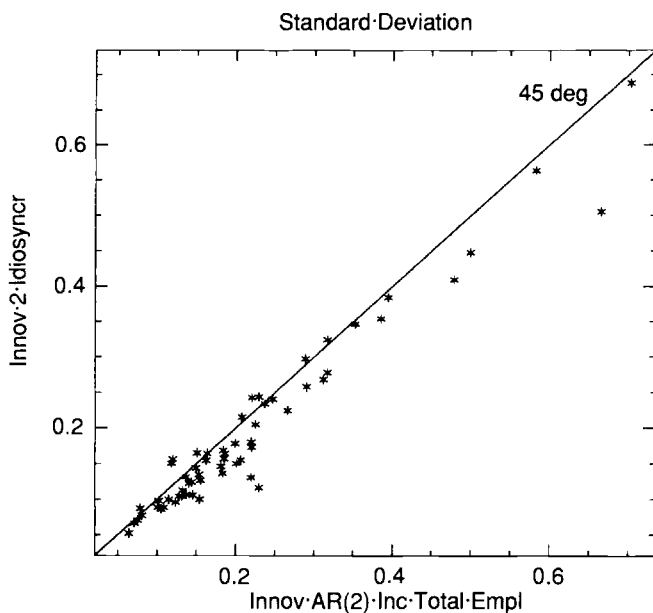
Fig. 7.6   Standard deviations of idiosyncratic disturbances in two-index model, against those of the residuals in the projection on own two lags plus total employment

is to allow a one-dimensional measure of "business activity" to emerge from analyzing long lists of different price and quantity time series. We have implemented one version of this vision and arrived at projection estimators of one- and two-index models. In computing these estimated indexes, we have used only employment series. Now we wish to see how, in particular dimensions, our estimated indexes compare with GNP, the most popular single "index" of real macroeconomic activity. GNP is, of course, constructed using an accounting method very different in spirit from that used by us and other workers in the Burns-Mitchell tradition.

Thus, we turn to some projections designed to explore this difference. A regression (over 1952–89) of GNP *growth rates* on a constant and first-differences of our fitted two indexes, lags 0–3, gives an R^2 of 83 percent. A similar regression using just the single index (estimated from the one-index model), again lags 0–3, gives an R^2 of 72 percent. GNP growth rates are thus highly correlated with the estimated index process. This correlation captures the sense in which a purely mechanical low-dimensional index construction yields an aggregate that closely tracks GNP growth.[10]

10. Note that sample versions of these indexes are estimated by a Kalman smoothing procedure and therefore use observations on sectoral employment over the entire sample, past and future. This is also why, in figs. 7.2 on, we always used future and past observable aggregates to make the comparison fairer.
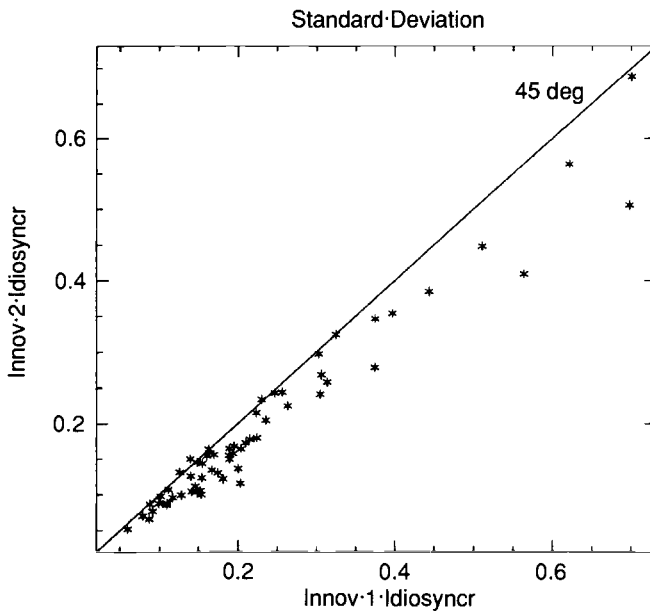
Standard·Deviation

Fig. 7.7   Standard deviations of idiosyncratic disturbances in two-index model, against those in one-index model

We have also experimented with vector autoregressions in GNP growth, average employment growth, and first-differences of the indexes. Exclusion tests of lag blocks of the indexes here, however, *cannot* be viewed as Granger-causality tests because the estimated indexes are themselves two-sided distributed lags of employment. It would be possible to use the estimated parameters of our index representations to construct one-sided-on-the-past index projections. These one-sided projections could then be used to conduct Granger-causality tests. We have not done this here, but we think that it would be a useful exercise.

In concluding this empirical section, we judge that the unobservable index application to the large cross section here has been, in the main, successful. The empirical results here encourage us to be optimistic about continued use of large cross sections for dynamic analysis. There are, however, dimensions along which relative failure might be argued. Most notably, the refinement in the description of commonalities given by the unobservable index model was not spectacular relative to that given by observable measures, such as aggregate GNP or total employment. While the two-index representation is a better description of sectoral employment fluctuations, relative to a single-index representation, not both indexes turn out to be equally important for predicting GNP. For this exercise, the tighter parameterization implicit in a single index appears to dominate the marginal increase in information from a two-index representation. These failures should be contrasted with our two princi-

pal successes: (i) the tractability of our extension of standard index model analysis (simultaneously to encompass differing nonstationarities, large cross-sectional dimensions, and extensive orthogonality restrictions) and (ii) the strong informational, predictive content for GNP of our estimated common indexes.

Unlike in research on interpreting unobservable disturbances—such as in VAR (vector autoregressive) studies—we do not attempt here to name the two unobservable factors. Rather, the goal of the present paper has been to carry out the methodological extensions in point i above and to examine the forecasting properties of the resulting common factors in point ii. Future work could, in principle, attempt the same exercise for the unobservable index models here as has been performed for VAR representations elsewhere.

7.4   Conclusion

We have provided a framework for analyzing comovements—aggregate dynamics—in random fields, that is, data where the number of cross-sectional time series is comparable in magnitude to the time length. We have shown how, on reinterpretation, standard techniques can be used to estimate such index models.

In applying the model to estimate aggregate dynamics in employment across different sectors, we discovered that a model with two common factors turns out to fit those data surprisingly well. Put differently, much of the observed fluctuation in employment in those many diverse industries is well explained by disturbances that are perfectly correlated across all the sectors.

The econometric structure that we have used here seems to us quite rich and potentially capable of dealing with many different interesting questions. Among others, this includes issues of the relative importance of different kinds of disturbances (e.g., Long and Plosser 1983; and Prescott 1986), convergence across economic regions (e.g., Barro and Sala-i-Martin 1992; and Blanchard and Katz 1992), aggregate and sectoral comovements (e.g., Abraham and Katz 1986; Lilien 1982; and Rogerson 1987), and location dynamics, conditioning on exogenous policy variables (e.g., Papke 1989). Previous work, however, has used measurement and econometric techniques that differ substantially from that which we propose here; clearly, we think our procedure is closer in spirit to the relevant economic ideas. Questions about the appropriate definition and measurement of inflation, comovement in consumption across economic regions, and the joint behavior of asset prices observed for many assets and over long periods of time all can be coherently dealt with in our framework. In future research, we intend to apply our techniques to these and similar, related issues.

Appendix A
Technical Appendix

This appendix constructs a state space representation for the model of section 7.2; this is needed to compute the Kalman-smoothed projections that are in turn used in applying the EM algorithm to our estimation problem.

Recall that $\phi_{jk} = \beta_j \, \alpha_{jk}$ and that the maximum lags on β_j and α_{jk} are M_b and M_a, respectively. Define $M_f = M_b + M_a$; this is the maximum lag on ϕ_{jk}. From M_g, the maximum lag on Γ in (5), define $M_h = \max(M_f, M_g)$, and let

$$\bar{U}(t) = [U(t)', U(t-1)', \dots, U(t-M_h)']'$$

and conformably

$$\bar{\eta}_U(t) = [\eta_U(t)', 0', \dots, 0']'.$$

Then write (5) in first-order form as

(A1) $$\bar{U}(t) = C \, \bar{U}(t-1) + \bar{\eta}_U(t),$$

where C is $(1 + M_h) \times K$ square, with the following structure: Call a collection of K successive rows (or columns) a K-row (or -column) block. The last M_h K-row blocks comprise simply zeroes and ones, in the usual way, forming identities. Consider then the kth row ($k = 1, 2, \dots, K$) in the first K-row block. The first K-column block of this row vanishes everywhere except in the kth position, where it contains $g_k(1) + 1$; the $(1 + M_g)$-th K-column block of this row vanishes everywhere except in the kth position, where it contains $-g_k(M_g)$. For K-column block m ($m = 2, 3, \dots, M_g$), the entries again vanish everywhere except in the kth position, where they equal $g_k(m) - g_k(m-1)$. This pattern of coefficients comes from (5) being in first-differences whereas (A1) is in levels.

Turning to the observables, write

$$X(t) = [X_1(t), X_2(t), \dots, X_N(t)]';$$

since this is observed at time t, we denote

$$\text{info}(t) = \{X(t), X(t-1), X(t-2), \dots\}.$$

Let

$$\tilde{X}(t) = [X(t)', X(t-1)', \dots, X(t-M_b+1)']'.$$

Also, write out $\phi_{jk}(L)$ explicitly as

$$\phi_{jk}(L) = \sum_{m=0}^{M_f} f_{jk}(m)L^m.$$

We can then rewrite equation (7) as

(A2) $X(t) = A \begin{bmatrix} U(t) \\ X(t - 1) \end{bmatrix} + dW(t) + \varepsilon(t) = aZ(t) + dW(t) + \varepsilon(t),$

where d has rows formed from d_j, and A has the following structure. Its jth row has the first $(1 + M_f)K$ entries given by

$$[f_{j1}(0), \ldots, f_{jK}(0), f_{j1}(1), \ldots, f_{jK}(M_f)];$$

the remaining entries in this row differ from zero only in M_b places. After the first $(1 + M_b)K$ entries (not the first $[1 + M_f]K$), there are M_b N-row blocks. Each such block vanishes except in the jth entry, which equals $b_j(m)$, for $m = 1, 2, \ldots, M_b$.

Now augment the transition equation (A1) with $\tilde{X}$ and W; that is, write

(A3) $Z(t + 1) = cZ(t) + \delta W(t) + \eta(t + 1),$

with

$$c = \begin{pmatrix} C & 0 \\ c_{21} & c_{22} \end{pmatrix},$$

where the first N rows of $(c_{21} \; c_{22})$ contain a, and the remaining rows comprise just zeros and ones; the matrix δ vanishes everywhere except in the N rows after the first $(1 + M_h)K$ (in those N rows it is composed of d); and, finally,

$$\eta(t + 1) = \begin{bmatrix} \bar{\eta}_U(t + 1) \\ \varepsilon(t) \\ 0 \end{bmatrix}.$$

Notice that (A3) contains within it (A2).

Since all but the first $(1 + M_h)K$ entries of $Z(t + 1)$ are observed at time t, it will be natural to let $P[Z(t + 1) \mid info(t)]$ equal $\tilde{X}(t)$ except in its first $(1 + M_h)K$ entries. Similarly, we will choose the conditional mean square $\text{Var}[Z(t + 1) \mid info(t)]$ to be zero everywhere except in the leading $(1 + M_h)K$ square diagonal block.

In summary, our state space representation is

$$Z(t + 1) = cZ(t) + \delta W(t) + \eta(t + 1),$$
$$X(t) = AZ(t) + dW(t) + \varepsilon(t),$$

with $\text{Var}[Z(1)|info(0)]$ initialized to vanish everywhere except in the leading $(1 + M_h)K$ diagonal block. The disturbance vector $[\eta(t + 1)', \varepsilon(t)']'$ is serially uncorrelated and has covariance matrix

$$\Omega = \begin{pmatrix} \Omega_\eta & \Omega_{\eta\varepsilon} \\ \cdot & \Omega_\varepsilon \end{pmatrix},$$

where Ω_η is singular, $\Omega_{\eta\varepsilon}$ contains Ω_ε, and Ω_ε is $N \times N$ diagonal. We write $(\Omega_{\eta\varepsilon})_1$ to denote the first $(1 + M_h)K$ rows of $\Omega_{\eta\varepsilon}$ and Ω_{η_1} to denote the leading $(1 + M_h)K$ diagonal block of Ω_η.

Partition $Z(t)$ as $[Z_1(t)', Z_2(t)']'$, where

$$Z_1(t) = \tilde{U}(t), \quad (1 + M_h)K \times 1,$$

and

$$c = \begin{pmatrix} c_{11} & c_{12} \\ c_{21} & c_{22} \end{pmatrix}, \quad \text{with } c_{11} = C(1 + M_h)K \text{ square}, c_{12} = 0.$$

Thus, $\eta(t)$ is also partitioned into $[\eta_1(t)', \eta_2(t)']'$, with $\eta_1 = \tilde{\eta}_U[(1 + M_h)K \times 1]$. We can now write the measurement equation

$$X(t) = (a_1 \ a_2) \begin{bmatrix} Z_1(t) \\ Z_2(t) \end{bmatrix} + \varepsilon(t),$$

where a_1 comprises the first $(1 + M_h)$ K columns of a. Note that $\text{Var}[Z_1(t)|\text{info}(s)]$ is thus only $(1 + M_h)K$ square and has dimensions that are invariant to N.

Appendix B
Data Appendix

Real GNP in the text refers to the annual version of GNP82 obtained from Citibase July 1991. This series is normalized to constant 1982 dollars and is no longer available—the new constant-dollar GNP series in Citibase is normalized to constant 1987 dollars and is available only from 1959. Since the employment data that we use go back through 1948, we decided to use the older GNP82.

Next, the sectoral description we study is from the national income and product accounts: employment (as listed below).

FTE Employees by Industry (annual)

Private Industries
  1. Farms (GAFAF1)
  2. Agriculture, forestry, and fisheries (GAFAF7)
  3. Metal mining (GAFM10)
  4. Coal mining (GAFM12)
  5. Oil and gas extraction (GAFM13)
  6. Nonmetallic minerals, except fuels (GAFM14)
  7. Construction (GAFCC)

Manufacturing: Durable Goods
  8. Lumber and wood products (GAFD24)
  9. Furniture and fixtures (GAFD25)

10. Stone, clay, and glass products (GAFD32)
11. Primary metal industries (GAFD33)
12. Fabricated metal products (GAFD34)
13. Machinery, excluding electric (GAFD35)
14. Electric and electronic equipment (GAFD36)
15. Motor vehicles and equipment (GAF371)
16. Other transportation (GAFD37)
17. Instruments and related products (GAFD38)
18. Miscellaneous manufacturing industries (GAFM39)

Manufacturing: Nondurable Goods

19. Food and kindred products (GAFN20)
20. Tobacco manufacturing (GAFN21)
21. Textile mill products (GAFN22)
22. Apparel and other textile products (GAFN23)
23. Paper and allied products (GAFN26)
24. Printing and publishing (GAFN27)
25. Chemicals and allied products (GAFN28)
26. Petroleum and coal products (GAFN29)
27. Rubber and miscellaneous plastic products (GAFN30)
28. Leather and leather products (GAFN31)

Transportation and Public Utilities

29. Railroad transportation (GAFT40)
30. Local and interurban passenger transit (GAFT41)
31. Trucking and warehousing (GAFT42)
32. Water transportation (GAFT44)
33. Transportation by air (GAFT45)
34. Pipelines, except gas (GAFT46)
35. Transportation services (GAFT47)

Communication

36. Telephone and telegraph (GAF481)
37. Radio and television broadcastings (GAF483)
38. Electric, gas, and sanitary services (GAFUT)

39. Wholesale trade (GAFW)
40. Retail trade (GAFR)

Finance, Insurance, and Real Estate

41. Banking (GAFF60)
42. Credit agencies other than banks (GAFF61)
43. Security and commodity brokers (GAFF62)
44. Insurance carriers (GAFF63)
45. Insurance agents and brokers (GAFF64)

46. Real estate (GAFF65)
47. Holding and other investment companies (GAF67F)

Services

48. Hotels and other lodging places (GAFS70)
49. Personal services (GAFS72)
50. Business services (GAFS73)
51. Auto repair services and garages (GAFS75)
52. Miscellaneous repair services (GAFS76)
53. Motion pictures (GAFS78)
54. Amusement and recreational services (GAFS79)
55. Health services (GAFS80)
56. Legal services (GAFS81)
57. Educational services (GAFS82)
58. Social services and membership organizations (GAFS86)
59. Miscellaneous professional services (GAFS89)
60. Private households (GAFS88)

References

Abraham, Katherine G., and Lawrence F. Katz. 1986. Cyclical unemployment: Sectoral shifts or aggregate disturbances. *Journal of Political Economy* 94 (3): 507–22.

Barro, Robert J., and Xavier Sala-i-Martin. 1992. Convergence. *Journal of Political Economy* 100 (2): 223–51.

Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. Regional evolutions. *Brookings Papers on Economic Activity* (1), 1–75.

Engle, Robert F., and Clive W. J. Granger. 1987. Cointegration and error correction: Representation, estimation, and testing. *Econometrica* 55 (2): 251–76.

Geweke, John. 1977. Labor turnover and employment dynamics in US manufacturing industry. In *New methods in business cycle research,* ed. Christopher A. Sims. Minneapolis: Federal Reserve Bank of Minneapolis.

Geweke, John, and Kenneth J. Singleton. 1981. Maximum likelihood "confirmatory" factor analysis of economic time series. *International Economic Review* 22 (1): 37–54.

Hamilton, James D. 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57 (2): 357–84.

King, Robert G., Charles I. Plosser, James H. Stock, and Mark W. Watson. 1991. Stochastic trends and economic fluctuations. *American Economic Review* 81 (4): 819–40.

Lehmann, Bruce N., and David M. Modest. 1988. The empirical foundations of the arbitrage asset pricing theory. *Journal of Financial Economics* 21 (2): 213–54.

Lilien, David M. 1982. Sectoral shifts and cyclical unemployment. *Journal of Political Economy* 90 (4): 777–93.

Long, John B., and Charles I. Plosser. 1983. Real business cycles. *Journal of Political Economy* 91 (1): 39–69.

Papke, Leslie E. 1989. Interstate business tax differentials and new firm location: Evidence from panel data. Working Paper no. 3184. Cambridge, Mass.: NBER, November.

Prescott, Edward C. 1986. Theory ahead of business cycle measurement. *Federal Reserve Bank of Minneapolis Quarterly Review* 10 (4): 9–22.

Rogerson, Richard. 1987. An equilibrium model of sectoral reallocation. *Journal of Political Economy* 95 (4): 824–34.

Ruud, Paul A. 1991. Extensions of estimation methods using the EM algorithm. *Journal of Econometrics* 49:305–41.

Sargent, Thomas J., and Christopher A. Sims. 1977. Business cycle modelling without pretending to have too much *a priori* economic theory. In *New methods in business cycle research,* ed. Christopher A. Sims. Minneapolis: Federal Reserve Bank of Minneapolis.

Stock, James H., and Mark W. Watson. 1988. Testing for common trends. *Journal of the American Statistical Association* 83 (404): 1097–1107.

————. 1990. Business cycle properties of selected U.S. economic time series. Working Paper no. 3376. Cambridge, Mass.: NBER, June.

Watson, Mark W., and Robert F. Engle. 1983. Alternative algorithms for the estimation of dynamic factor, MIMIC, and varying coefficient regression models. *Journal of Econometrics* 23:385–400.

Watson, Mark W., and Dennis F. Kraft. 1984. Testing the interpretation of indices in a macroeconomic index model. *Journal of Monetary Economics* 13 (2): 165–81.

Wu, C. F. Jeff. 1983. On the convergence of properties of the EM algorithm. *Annals of Statistics* 11 (1): 95–103.

Comment   John Geweke

The paper by Danny Quah and Thomas J. Sargent examines some possibilities for the application of index models of multiple time series in situations in which the number of cross sections (N) is large relative to the number of time periods (T). Since there are, in fact, many more of the former than there are of the latter, this is an interesting and worthwhile project. The paper argues that, when $N >> T$, one cannot estimate covariograms or autoregressions because the number of parameters grows quadratically with the number of observations but that one can estimate index models because the number of parameters grows linearly. In fact, stronger cases can be made for both approaches. In autoregressions, it is natural to regard sectors of the economy as exchangeable in a prior distribution, and either the Minnesota prior (Litterman 1986) or a hierarchical prior could be employed to yield a manageable inference problem. Bayesian inference for covariograms and spectra would follow at once. With regard to index models, it is clear that, given conventional normalizations and identifying restrictions, one can consistently extract the unob-

served indexes, as $N \rightarrow \infty$ with fixed T. Just how N must be related to T to do so as $(N, T) \rightarrow \infty$ would be a significant research question.

It would be useful to know how the authors resolved important identification questions. A necessary step in identifying the time series $U(t)$ is removal of the possibilities for rotating these latent factors through premultiplication by an orthonormal matrix. This problem arises in traditional factor analysis (Lawley and Maxwell 1971) and can be resolved in a variety of ways. If left unresolved, it must arise as a mechanical problem in application of the EM algorithm, but the technical appendix does not indicate what normalizations were used. Additional identification problems are introduced by the fact that $\Delta U(t)$ is a stationary time series. In particular, even after removal of conventional rotation possibilities, there are several ways to represent $\Delta U(t)$ as a linear combination of serially uncorrelated processes (Whittle 1983, chap. 2). Finally, dynamic-factor models that are "overfit" in the number of factors become underidentified in fundamental ways since a sector-specific disturbance can masquerade as a commonality. This has subtle but damaging effects on test statistics in the factor-analysis model (Geweke and Singleton 1980), and it is unlikely that these effects would be mitigated in a dynamic-factor model. It is difficult to see how a careful demonstration of the convergence of the EM algorithm to maximum likelihood estimates could proceed, without first dealing with these identification issues.

The development of an asymptotic distribution theory also depends on resolving identification questions: hence (presumably) the absence of standard errors, test statistics, or any other sort of formal statistical analysis in the paper. The development of an appropriate asymptotic foundation for inference is likely to be more difficult than resolution of identification questions. Asymptotics in T would require managing the unconventional unit root theory in the context of a nontrivial latent-variable model. Given the authors' motivation for examining these models, asymptotics in N is presumably of more compelling interest. Here the ratio of parameters to observations converges to a positive value, but the ratio of the number of latent variables to observations drops to zero: hence the opportunity to extract signals without error asymptotically. But, if T is fixed, and perhaps even if only $N/T \rightarrow \infty$, it may be sensible to abandon most of the time-series structure on the commonalities altogether, in the same way that it is not necessary to assume a time-series structure for disturbances in panel data when N is very large relative to T. There is a rich supply of nontrivial problems for the frequentist econometrician here.

For the authors' application, with $N = 60$ and $T = 37$, the presence or absence of such theory could turn out to be moot since no asymptotic theory might be applicable. It appears that the EM algorithm is used to estimate 540 parameters (nine for each of sixty sectors in the dynamic-factor model) and to extract forty latent variables, based on 2,400 initial degrees of freedom. For the two-factor model, the numbers rise to 780 parameters and eighty latent

variables. Without any distribution theory, it would be difficult to evaluate the claim in the concluding section that "much of the observed fluctuation in employment in those many diverse industries is well explained by disturbances that are perfectly correlated across all the sectors." Summary statistics on proportion of variance accounted for by commonalities would help, but there would still be no basis for discriminating between statistical artifact and properties of the population. Nevertheless, two aspects of the reported empirical results cast doubt on this claim.

First, from the information given in section 7.3, a conventional test of the hypothesis that real GNP does not affect sectoral employment may be applied using an $F(5, 28)$ distribution. The hypothesis fails to be rejected at the 5 percent level for all those industries that lie above a ray from the origin with a slope of .828 and at the 1 percent level with a slope of .684. The hypothesis is not rejected for about one-quarter of the industries at the 5 percent level and for about half the industries at the 1 percent level. More informative would be confidence intervals for the fraction of variation accounted for by real GNP, but this cannot be reconstructed on the basis of the information provided. The outcome is similar when real GNP is replaced by the multisector employment measure (fig. 7.3).

Second, the presentation in figures 7.5 and 7.6 does little to support the case for the dynamic-index model. As the authors point out, neither of the models compared is nested in the other. However, the time-series structure of the commonality in the index model (i.e., the autocorrelation function of $\phi_{jk}[L]U[t]$) is very close to that of real GNP given that annual log real GNP is a random walk to a rough order of approximation and that coefficients can unwind modest departures from a random walk. Beyond this, in the index model, one is free to choose the time-series realizations of the commonality. From a strictly heuristic point of view—all one has here, given that no theory has been worked out—this appears to give an edge to the index model that goes beyond the usual counting rules of thumb for parameters. Further, notice that figures 7.5 and 7.6 report results for two-index models, which have, by my count, thirteen parameters per sector, whereas the real GNP regressions have only ten. (At this point, the two-index model still gets an additional eighty realizations of the commonality unavailable to the regressions!) The issue is somewhat clouded by the fact that sectoral employment in the regressions apparently extended from 1951 through 1987, whereas in the index models it ranged from 1948 through 1989. In this situation, figures 7.5 and 7.6 do nothing to dislodge the null hypothesis that the estimated index models reflect nothing beyond the regressions on real GNP.

Index models for large numbers of cross sections may well become an important tool in the decomposition of aggregate output. If they are to become a reliable and discriminating tool, however, substantial further development of the theory of inference for such models is needed.

References

Geweke, J. F., and K. J. Singleton. 1980. Interpreting the likelihood ratio statistic in factor models when sample size is small. *Journal of the American Statistical Association* 75:133–37.

Lawley, D. N., and A. E. Maxwell. 1971. *Factor analysis as a statistical method.* New York: American Elsevier.

Litterman, R. B. 1986. Forecasting with Bayesian vector autoregressions—five years of experience. *Journal of Business and Economic Statistics* 4:25–38.

Whittle, P. 1983. *Prediction and regulation by linear least-square methods.* Minneapolis: University of Minnesota Press.

This Page Intentionally Left Blank