

IDENTIFYING THE COMMON COMPONENT OF INTERNATIONAL ECONOMIC FLUCTUATIONS: A NEW APPROACH*

Robin L. Lumsdaine and Eswar S. Prasad

In this paper, we develop an aggregation procedure using time-varying weights for constructing the common component of international economic fluctuations. The methodology for deriving time-varying weights is based on some stylised features of the data documented in the paper. The model allows for a unified treatment of cyclical and seasonal fluctuations and also accommodates the dynamic propagation of shocks across countries. We find evidence for a ‘world business cycle’ as well as evidence for a distinct European common component. We also find some evidence that macroeconomic fluctuations have become more closely linked across industrial economies in the period after 1973.

The world economy has become more closely integrated in recent years due to increasing trade and financial flows across countries. This has spurred interest in the question of how the ongoing phenomenon of ‘globalisation’ has affected the transmission and propagation of business cycle fluctuations across national borders. An important question in this context is whether a substantial fraction of economic fluctuations are country-specific or if there exists a ‘world business cycle’ which might be defined as fluctuations that are common across all countries. More generally, the comovement of macroeconomic aggregates across different countries has become a topic of increasing interest in both academic and policy circles.

These issues have implications in a number of dimensions. An empirical analysis of the relative importance of country-specific versus common international fluctuations is of relevance for determining the structure of certain classes of international business cycle models. From a policy perspective, strong comovement of output across countries at business cycle frequencies implies that external trade would be unlikely to play a significant role in dampening fluctuations. Domestic policies aimed at affecting the real exchange rate and thereby attempting to boost net exports in the short run would then tend to have limited impact. The degree of synchronisation of business cycles across countries also has important implications for short-run international policy coordination and for assessing the feasibility of monetary unions. Therefore, identifying and analysing the common component of international economic fluctuations is relevant from a number of different perspectives.

* Part of this research was conducted while Lumsdaine was a Visiting Scholar, IMF and a National Fellow, Hoover Institution. We are grateful to Joe Beaulieu, Peter Christofferson, Frank Diebold, Rob Engle, Sharon Kozicki, Adrian Pagan, Lucrezia Reichlin, Julio Rotemberg, Mark Watson and numerous other colleagues for discussions and helpful suggestions. We also received useful comments from seminar participants at UCLA, the University of British Columbia, UC Davis, UC San Diego, the IMF, the Federal Reserve Bank of San Francisco and meetings of the Econometric Society and the European Economic Association. This paper has benefited greatly from the comments of the editor and three anonymous referees. The opinions expressed in this paper are those of the authors and do not represent those of the IMF or Deutsche Bank and its affiliates. This article has a technical appendix that is only available on the Royal Economic Society website: www.res.org.uk

The objective of this paper is to estimate the common component in international economic fluctuations and to examine its properties. One strand of related literature has attempted to shed light on common fluctuations by looking at bivariate correlations of business cycle indicators and examining changes in these correlations over different time periods; see, e.g., Baxter and Stockman (1989) and Backus and Kehoe (1992). Another strand of literature has focused on using time series models to analyse the sources of economic fluctuations. Previous literature in this latter area has focused on trying to identify aggregate, country-specific and industry-specific shocks separately. For instance, Stockman (1988) and Bayoumi and Prasad (1997) use an error components methodology while Altonji and Ham (1990), Stock and Watson (1989, 1993), Forni and Reichlin (1996), Norrbin and Schlagenhauf (1996) and Gregory *et al.* (1997) use dynamic factor models. A key issue in this literature is the propagation mechanism that allows for lagged feedback effects across various shocks. Although dynamic factor models are able to allow for such feedback effects, this comes at the cost of having to estimate a large number of parameters and restricting the covariance properties of these shocks. In addition, the procedure followed in most of the literature implicitly weights all units of the disaggregated data equally in all periods.

One method for relaxing the equal-weights assumption is to weight by some measure of each country's relative size in total world output. Following this approach, we first examine a measure of the common component of international fluctuations obtained by using a fixed PPP-adjusted weight to aggregate seasonally adjusted industrial production growth rates. The correlations between industrial production growth in each country and this common component are strongly positive for most countries, supporting the notion of a 'world business cycle'. The fixed-weight measure of the common component is, however, inadequate in many respects. One is that relative economic sizes of countries change over time and the weights should reflect this dynamic nature. Another is that countries experience idiosyncratic shocks; these shocks, by definition, should not affect the common component (even though purely country-specific fluctuations, especially for large countries, may affect total world output in an accounting sense). Fixed weights do not allow for different types of shocks in different periods; all shocks are presumed to have the same influence.

To address these limitations, in this paper we propose a new methodology that incorporates a time-varying weighting scheme for constructing the common component. The modelling strategy that we employ involves estimating univariate models of time-varying conditional variances for output growth fluctuations in each country. The time-varying weights for each country are then derived as a function of the estimated conditional variances.¹

¹ The notion of aggregating using time-varying weights has been used in models of combining forecasts; for example, Deutsch *et al.* (1994) use rolling regressions to estimate time-varying weights. Also see Diebold and Pauly (1990). Christofferson and Giorgianni (2000) develop a methodology for extracting time-varying weights used in constructing basket currencies (where the weights are not publicly known) and show how to account for the risk that time-varying weights add to the standard basket-hedge position.

The weighting scheme is motivated by two empirical regularities that are documented in this paper. The first is the negative relationship between country size and the average volatility of industrial production growth rates. The second feature is the presence of conditional heteroscedasticity in monthly industrial production growth rates for all countries in the sample. We use these two features to determine time-varying weights by noting that each country's volatility relative to that of other countries provides a measure of the degree of idiosyncrasy in the observed shocks. The weighting scheme developed in this paper implicitly assigns a lower weight to a country when it is subject to a large country-specific shock but leaves the weights unchanged if a common shock occurs. The extent to which the methodology downweights outliers provides a way of endogenously distinguishing between idiosyncratic and common fluctuations in the construction of the common component.

The objective of the methodology developed in this paper is to identify the common component rather than to distinguish among different sources of shocks (global or country-specific). Hence, the methodology is designed to implicitly accommodate the effects of the dynamic propagation of shocks across countries but without placing restrictions on the propagation of shocks across countries, unlike in the case of dynamic factor models that require restrictions on the feedback effects among different shocks. In addition, since we are interested in identifying the common component of fluctuations across countries, idiosyncratic fluctuations are assigned a lower weight, even if they occur in large countries that may account for a relatively large share of total world output. As an empirical matter, however, shocks to large countries do tend to get propagated across countries and these effects are then reflected in the common component.

Another aspect of economic fluctuations that has gained prominence recently is the importance of seasonal fluctuations and the relationship between seasonal and business cycle fluctuations.² The methodology developed in this paper can, in principle, eliminate the effects of idiosyncratic seasonal fluctuations on the common component. On the other hand, common seasonal fluctuations and the part of seasonal variation correlated with the business cycle do enter into the construction of the common component. Thus, the aggregation procedure allows for a unified treatment of seasonal and business cycle fluctuations.

However, unlike in the case of techniques that involve orthogonal decompositions, the approach in this paper does not permit us to identify the shocks themselves. In particular, we can not tell apart global shocks from country-specific shocks or shocks that might initially appear to be country-specific but that eventually propagate to other countries through trade or other links. Since we are interested in identifying the common component of international fluctuations, irrespective of the sources of shocks to this common component, this distinction is not important for our purposes. But it does limit the relevance of the technique in this paper for understanding the relative importance of different sources of shocks. In addition, since we do not put much structure on cross-country propagation

² See, e.g., Beaulieu and Miron (1992), Beaulieu *et al.* (1992), Canova and Ghysels (1994), Cecchetti and Kashyap (1996), and Cecchetti *et al.* (1997).

mechanisms, we can not shed light on the channels through which shocks are propagated across countries.

An alternative approach, which could more directly capture propagation dynamics and address some of the other issues described above, would be to estimate a vector autoregression with multivariate or univariate GARCH. However, identifying such a model would require other strong assumptions about the propagation dynamics as well as about the ordering of countries in the VAR. Furthermore, such approaches rapidly become unwieldy, and the required identifying restrictions become more onerous, as the number of countries included in the estimation increases. The approach developed in this paper has the virtue of being able to easily handle aggregation across a large number of series.

The paper proceeds as follows. Section 1 motivates the use of time-varying weights in constructing the common component and describes the econometric procedure for estimation of these weights. Section 2 examines the properties of the estimated time-varying weights and compares the properties of the common component constructed using these weights to that of a benchmark fixed-weight common component. Section 3 extends the results in two ways: (a) by investigating potential structural change in our specification after 1973 and (b) by estimating a European common component. The sensitivity of the aggregation procedure to the treatment of deterministic seasonal effects is also examined. Section 4 concludes.

1. Aggregation Using Time-Varying Weights

This Section first sketches a time series model that clarifies the identification issues involved in measuring the common component of fluctuations across countries. The methodology introduced here is more broadly applicable to situations where construction of an aggregate from a collection of individual time series is desirable, but we will focus on the details in the context of fluctuations in output growth. Evidence is then presented on some empirical regularities that could be exploited to devise a procedure for constructing time-varying weights. The econometric procedure used to derive these weights and construct the resulting common component is then described.

Our choice of aggregation methodology is motivated by three considerations. First, an ideal weighting scheme should be capable of distinguishing between country-specific and common fluctuations. In principle, the weights chosen for constructing the aggregate measure should reflect fluctuations only in the common components in each series. The relative weight of a particular country should decrease when that country experiences a largely idiosyncratic shock. If, on the other hand, a country's shock is of the common component type, its relative weight should remain unchanged. If it were possible to identify the two types of shocks for each country separately, we could compute time-varying weights which took into account both the relative across-country weight and the relative within-country weight (between common and idiosyncratic shocks). Because these are not observable, however, it is necessary to determine a mechanism for distinguishing between these two effects without having to impose unwieldy restrictions.

Second, another important consideration in estimating the common component of international fluctuations is to allow for the propagation of shocks across countries. Error component models typically ignore this issue while dynamic factor models attempt to capture this phenomenon by allowing for feedback effects across country-specific and aggregate fluctuations. This comes at the cost, however, of having to estimate a large number of parameters and having to impose stringent restrictions on the covariance properties of the shocks. In addition, the structure of the transmission mechanism for these shocks is generally assumed to remain unchanged over time. An alternative approach is the common trends and common cycles method developed by Engle and Kozicki (1993), although this methodology requires restrictions on the factor loadings of the common cycles in order to allow for additional idiosyncratic behaviour.³

Third, monthly industrial production data typically display a high degree of seasonality, an aspect that could potentially complicate econometric work. We prefer to remain agnostic on the appropriate characterisation of seasonal variation in the data. We recognise that patterns of seasonal variation could change over time. In addition, as noted by Beaulieu *et al.* (1992), seasonal cycles may be correlated with business cycles. In a similar vein, Cecchetti and Kashyap (1996) and Cecchetti *et al.* (1997) have documented that, in the OECD economies, patterns of seasonal fluctuations in industrial production vary with the state of the business cycle. Furthermore, care must be taken not to remove a potential common seasonal component; Engle and Hylleberg (1996), for instance, find evidence of common seasonal patterns in unemployment among some OECD countries. For these reasons, rather than attempting to remove the entire seasonal component, we are interested in eliminating seasonality only to the extent that it interferes with our ability to measure the common component of fluctuations.

The above discussion suggests a role for time-varying weights in the construction of a common component. In what follows, we propose a methodology for constructing these time-varying weights.

1.1. A Basic Model

Consider the following time-series representation for output growth:

$$y_{it} = \mu_i + \beta(L)e_{it} + \sum_{j \neq i} \gamma_j(L)e_{jt} + \eta(L)e_t \quad (1)$$

where y_{it} indicates the growth rate of industrial production in country i at time t , μ_i is a country specific mean, e_{it} represents a country-specific shock to country i , e_{jt} represents a country-specific shock to country j , and e_t is a global shock. The lag polynomials β , γ_j , and η capture the effects of propagation of the different types of shocks. Distinguishing between these shocks in a reduced-form framework is clearly a difficult task, especially if β , γ_j , and η differ across countries.

³ Lippi and Reichlin (1994) provide a useful discussion of alternative concepts of co-movements of variables in the short run and the long run when different trend-cycle decompositions are considered.

Error component models typically impose an assumption of orthogonality between country-specific shocks and global shocks while dynamic factor models make a similar assumption but also allow for dynamic effects of these shocks. In either case, e_{jt} and e_t cannot be identified separately except under very restrictive assumptions about the propagation structure. Further, there is no reason to believe that the lag polynomials governing the propagation of shocks are the same across all countries or that these are constant over time. Also note that the formulation in (1) implies that feedback from lagged values of e_{jt} could complicate the identification of common shocks.

We take a different approach since our aim is not to identify the shocks themselves but to construct a measure of the common component that could include global shocks as well as country-specific shocks that eventually propagate to other countries. The approach, described in greater detail in Subsection 1.3 below, involves the construction of a weighted average measure of output fluctuations in each country, where the weights are allowed to vary over time. To do this, we first measure the conditional volatility of output growth for each country. Second, we interpret a single country's specific increase in conditional volatility as arising from a country-specific shock and, consequently, reduce the weight attributed to such a fluctuation when computing the common component. Thus, the methodology in this paper does not require strong assumptions about the correlation structure across different types of shocks or about the propagation mechanisms for different shocks.

One issue that arises here is whether taking an average is in fact an appropriate approach for constructing the common component. We discuss the intuition here and provide a more formal illustrative example in Appendix A on the website. Consider the case where both e_{it} and e_{jt} have zero mean and are drawn from distributions with similar second moments but are serially and mutually uncorrelated for all i, j . Also assume, for the moment, that the lag polynomials in (1) are all equal to unity. In this case, it is fairly easy to see that a simple average would in fact yield the common component e_b if the sample contained a sufficiently large number of countries, so that the sum of e_{it} and e_{jt} over all these countries was equal to their respective unconditional means of zero.

If each country's shocks had different variances or were correlated, a simple average weighting scheme would not necessarily be optimal. The optimal weights would then be inversely related to the unconditional volatility of these shocks. The intuition behind this is similar to why, in a regression, GLS is efficient relative to OLS when the errors are heteroscedastic. Similarly, if the coefficients in the lag polynomial η were not equal to unity then a simple average of the y_{it} would not yield e_b , while a weighted average of y_{it} , where the weights were inversely related to these coefficients, would.

In economic terms, it is plausible that country-specific shocks hitting larger countries are more likely to be eventually propagated to smaller countries than vice versa. Similarly, global shocks would tend to have a larger impact on smaller economies, especially since smaller industrial economies are generally more open to international trade than larger industrial economies. Therefore, we would expect the coefficients β and γ_j to be inversely related to some

measure of country size. In fact, as discussed below, we find a strong negative relationship between country size and the standard deviations of industrial production growth and exploit this empirical relationship in developing our methodology.

1.2. *Some Stylised Facts*

We begin by documenting the relationship between the fixed OECD weights (W_i), which are interpretable as a measure of relative country size, and the standard deviations of the individual industrial production growth rates (std_i).⁴ This relationship is summarised in the following regression (standard errors are in parentheses):

$$std_i = \frac{0.059}{(0.0064)} - \frac{0.0014}{(0.0006)} W_i \quad \text{and} \quad R^2 = 0.25. \quad (2)$$

There is clearly a strong negative relationship between country size and volatility in industrial production growth rates.⁵ This result is consistent with the view that larger economies tend to be more diversified, thereby tending to have lower aggregate volatility, and are also less affected by external shocks emanating from other economies; see Gerlach (1988). The methodology developed in this paper is motivated by this negative cross-sectional relationship between country size and business cycle volatility. The above observation also suggests, however, that if volatility in individual industrial production growth rates were constant over time, the use of fixed weights (that are related to country size, such as the OECD weights) might be justified.

We therefore investigate whether the individual industrial production growth series display evidence of time-varying volatility, in particular, conditional heteroscedasticity; such evidence would motivate the need for time-varying weights. One way to test for this is to use the Box-Pierce Q -statistic to test for autocorrelation in the *squared* residuals from a regression of industrial production growth rates on a constant and twelve lags. Results from the computation of this statistic are given in the last column of Appendix Table B1 on the website; for all countries, we reject the null hypothesis of no autocorrelation (conditional homoscedasticity of the squared residuals) in favour of the alternative. In all cases, autocorrelations up to order 12 were used for the computation of the statistic; under the null hypothesis, this is distributed as a $\chi^2(12)$ random variable. The corresponding 1% critical value is 26.2.

⁴ The OECD weights are derived from gross domestic product originating in the industrial sector and the GDP purchasing power parity for 1990.

⁵ Although the explanatory power of this regression is not large, it is in fact rather striking since the estimated relationship is an unconditional one that does not control for any other exogenous factors. The results were similar when we excluded the US and/or other outliers such as Luxembourg. We obtained virtually identical results using 1985 OECD weights (earlier weights were not available). In related work, we have also examined this relationship for US states using annual real gross state product over the period 1977–92. We find a similar, although less strong, negative relationship between the standard deviation of annual gross state product and relative state size. Head (1995) documents a similar negative relationship between country size and the variance of real GDP.

1.3. *The Methodology for Constructing Time-Varying Weights*

Since all series display evidence of conditional heteroscedasticity, we estimate univariate GARCH(1,1) models for each series and use the predicted values of the conditional variance to construct time-varying weights for the aggregate series. The GARCH model developed by Bollerslev (1986) is a variant of the autoregressive conditional heteroscedasticity (ARCH) model introduced in Engle (1982). The GARCH(1,1) model expresses the conditional variance as a function of lagged squared residuals and past conditional variance. We select this model because it has been shown empirically to capture the volatility dynamics in a wide variety of data and because quasi-maximum likelihood estimators of this model are consistent and asymptotically normal (Lumsdaine, 1996). The precise specification, for each country i , is as follows:

$$y_{it} = c_i + \varepsilon_{it}, \quad \varepsilon_{it}|I_{t-1} \sim N(0, h_{it}), \quad (3a)$$

$$h_{it} = w_i + \alpha_i \varepsilon_{it-1}^2 + \beta_i h_{it-1}, \quad (3b)$$

where y_{it} represents industrial production growth in country i at time t , c_i is a country-specific mean, and I_t denotes information available at time t . The parameters w_i , α_i , and β_i are constrained to be positive; the likelihood is also penalised to ensure that $\alpha_i + \beta_i \leq 1$, a constraint that never binds in the estimation. In addition, the unconditional mean and variance of country i 's industrial production growth rate are chosen as starting values for c_i and w_i , respectively, and the initial value of the conditional variance, h_{i0} , is 1.⁶

We estimate model (3) and compute \hat{h}_{it} for each series, $i = 1, \dots, 17$. Based on the stylised fact summarised in (2), $\hat{h}_{it}^{-1/2}$ can then be interpreted as a time-varying measure of the contribution of the fluctuations in a particular country to fluctuations in the international common component. Alternatively, we could use factor analysis to decompose the conditional variance into the sum of a common component conditional variance and an idiosyncratic component. This approach also requires restrictive orthogonality assumptions. Instead, we use the conditional variance for a given country relative to the average across countries as a measure of the idiosyncratic variance. Based on the empirical motivation given earlier for our weighting scheme, the time-varying weights W_{it} are then related to the inverse of the estimated conditional standard deviations and are expressed as a fraction of the total weight, so that

$$W_{it} = \frac{1}{\sqrt{\hat{h}_{it+1}}} \bigg/ \sum_{i=1}^{17} \frac{1}{\sqrt{\hat{h}_{it+1}}}. \quad (4)$$

Note that h_{it+1} is in the information set I_t . The aggregate series representing the common component of international fluctuations is then constructed as

⁶ The above parameter restrictions are standard in the estimation of GARCH models. Given these restrictions, as long as the initial value of h_{it} is assumed to be drawn from the stationary distribution, dependence on this initial value diminishes exponentially. None of the results reported below were sensitive to the choice of starting values. All of the estimation in this paper was done in RATS. The Berndt, Hall, Hall and Hausman optimisation algorithm was used in the GARCH estimation.

$$Z_t^G = \sum_{i=1}^{17} W_{it} y_{it}. \quad (5)$$

The key assumption underlying our methodology is that the relative conditional standard deviation is a measure of the degree of commonality among fluctuations shared across countries. This differs from the assumptions underlying factor models and error components models (which assume orthogonality between the common and idiosyncratic components). In this context, it is worth re-emphasising that our objective is to estimate the common component in fluctuations rather than to identify a global ‘shock’ that is orthogonal to all country-specific shocks. Appendix A presents a simple and highly stylised illustrative example that provides further motivation for the methodology used here to construct time-varying weights.⁷

In constructing the common component using time-varying weights, we have not specified the transmission mechanism between fluctuations in the aggregate series and in individual countries. We interpret country-specific increases in conditional volatility as reflecting country-specific fluctuations. Thus, holding other shocks constant, a shock that hits only one country would increase that country’s conditional volatility alone. This would result in a decline in the weight attributed to that country in constructing the common component for that period. If the shock propagated to other countries over time, however, the conditional volatility of fluctuations in other countries would increase, and the weights would then depend on how widely and over what time horizon the shock was propagated across countries. Thus, the methodology implicitly accommodates the propagation of shocks across countries without imposing much structure on the dynamics of this propagation.⁸

In the working paper version of this paper, we provide illustrative numerical examples of how the weights adjust to allow for the propagation of shocks. It is also important to note that a more restrictive time series model such as an ARCH(1) specification could capture contemporaneous transmission but would not allow for the dynamic cross-country propagation of shocks. In contrast, the GARCH model provides a flexible functional form that is capable of reflecting propagation dynamics and that allows for persistence in the weights via the coefficient β in (3).

Further, since we use conditional volatilities in constructing these weights, positive and negative shocks that are specific to a particular country are treated symmetrically since both these shocks would increase country-specific conditional

⁷ Forni and Reichlin (1996) use a dynamic factor approach and show that the optimal weights in such a framework are the eigenvalues corresponding to the maximum eigenvector. This fixed-weight approach implicitly assumes that the variance of the idiosyncratic component is a constant proportion of the variance of the total. Even with ‘optimal weights,’ however, their approach does not allow these relationships to change over time.

⁸ An alternative approach would be to estimate a multivariate GARCH model. To make such a model more tractable would, however, require additional assumptions on the conditional correlations of the shocks; see, e.g., Diebold and Nerlove (1989) and Bollerslev (1990).

volatility, thereby resulting in a lower weight for that country in the construction of the common component.⁹

The endogeneity between the aggregate series and the individual countries is captured in the conditioning information of the GARCH model; in particular, since $h_{t+1} \in I_t$, the time-varying weights are in the conditioning information set and can thus be thought of as known at time t . Therefore, the GARCH model also provides a mechanism for forecasting future relative fluctuations.

To summarise, our time-varying weighting scheme has the following characteristics: (i) the weights vary over time in a manner that minimises the impact of country-specific fluctuations on the common component; (ii) the weights, on average, are related to relative country size; (iii) the methodology allows for a unified treatment of seasonal and business cycle fluctuations; and also (iv) implicitly allows for the propagation of shocks across countries without placing restrictions on the transmission mechanism for the shocks.

A possible concern is that, if a particular country were to be relatively more sensitive to global shocks, our methodology would result in that country being assigned a lower weight on account of its relatively high observed volatility. This despite the fact that fluctuations in such a country could be more informative about the common component. We acknowledge this conceptual point but would note that it is unlikely to be empirically important for this particular application of our methodology for constructing time-varying weights. In Appendix B on the website, we present some suggestive numerical evidence of how important a problem this could be. The degree of potential distortion is related to the extent of country output fluctuations that are attributable to the common component and also to the relative volatility of the common component – our conclusion is that, based on the results reported in this and related papers, the distortion is unlikely to be significant for this application.

2. Results

The dataset used in this paper contains seasonally unadjusted monthly indices of industrial production for 17 OECD economies over the period 1963–94. On average, industrial production accounts for only about one-third of total output in these economies. However, this index tends to be highly correlated with the aggregate domestic business cycle and, since it represents output in the traded goods sector, is more relevant for examining the transmission and propagation of business cycles across countries. In addition, real GDP is available only at a quarterly frequency, which is inadequate for the implementation of our empirical methodology given the available span of the data. The data are transformed into logarithms, first differenced to achieve stationarity, and regressed on 12 monthly

⁹ We note that there is a literature which has explored asymmetries in business cycle variation, e.g., Hamilton (1989) and Pagan (1997). Our methodology could, in principle, be extended to allow for asymmetric effects of positive and negative shocks. However, for the purposes of identifying the common component, the interpretation of such asymmetries is much less straightforward. There is little evidence that the propagation of positive and negative shocks across countries is different or that positive and negative common shocks have different effects. Nevertheless, this is an interesting topic that we leave for future research.

dummy variables. Reasons for this choice of transformation, along with descriptive statistics and a discussion of seasonality and other issues related to the data, are given in Appendix C on the website.

We first examine the correlations of fluctuations in individual country industrial production growth rates with a benchmark fixed-weight common component. Some properties of the time-varying weights estimated using the univariate GARCH estimates are then discussed, followed by a more detailed analysis of the common component constructed using these weights.

2.1. Fixed-Weight Common Component

To construct a benchmark common component, we use the 1990 OECD weights as given in the first column of Table 1 to aggregate the data into a single series. The second column of Table 1 summarises the correlations of this fixed-weights benchmark common component with industrial production growth rates of the individual countries. Not surprisingly, many of the countries with large weights are also highly correlated with the aggregate series, but there is also substantial correlation with countries that have low weights but high levels of variability. For instance, Luxembourg has a correlation of around 0.5, higher than the correlation for the US.

In addition, the correlation between the benchmark and the individual countries does not appear to be constant; for example, industrial production growth in Finland and France is negatively correlated with the benchmark before 1973

Table 1
Correlations with the Common Component of International IP Growth Fluctuations

	OECD Weights	Fixed-Weight Measure of the Common Component			Time-Varying Weights			Time-Varying Measure of the Common Component		
		Full Sample	1963–73	1974–94	Mean	Min.	Max.	Full Sample	1963–73	1974–94
Austria	1.1	0.25	0.29	0.22	3.89	0.98	6.27	0.24	0.29	0.21
Belgium	1.1	0.08	0.07	0.08	2.25	1.00	3.32	0.18	0.31	0.12
Canada	3.1	0.28	0.25	0.29	14.71	9.04	21.77	0.62	0.65	0.61
Finland	0.5	-0.07	-0.38	0.09	3.95	0.35	12.00	0.14	0.01	0.20
France	6.6	0.00	-0.15	0.10	2.82	0.43	5.14	0.18	-0.04	0.33
Greece	0.3	0.23	0.20	0.24	3.31	0.90	5.31	0.24	0.21	0.24
Germany	10.9	0.45	0.43	0.46	3.13	0.80	4.93	0.36	0.43	0.33
Italy	7.1	0.67	0.77	0.60	0.29	0.14	0.45	0.20	0.26	0.17
Japan	19.5	0.63	0.63	0.61	14.62	2.63	27.00	0.41	0.35	0.42
Luxembourg	0.1	0.53	0.67	0.46	1.27	0.58	1.96	0.28	0.33	0.26
Norway	0.5	0.45	0.36	0.50	0.63	0.14	0.97	0.35	0.21	0.42
Netherlands	1.6	0.40	0.33	0.43	3.83	1.85	6.41	0.45	0.50	0.43
Portugal	0.7	0.57	0.61	0.54	0.39	0.09	0.60	0.22	0.29	0.18
Spain	3.3	0.69	0.77	0.64	0.31	0.19	0.47	0.26	0.14	0.27
Sweden	1.0	0.06	-0.09	0.10	0.78	0.09	1.25	0.13	0.37	0.14
UK	6.3	0.31	0.23	0.35	4.91	1.23	7.67	0.20	0.37	0.45
US	36.2	0.35	0.19	0.43	38.92	14.84	51.99	0.63	0.45	0.71

Notes: The OECD weights reported in the first column are 1990 relative industrial production weights constructed using purchasing power parity exchange rates. The fixed weights are normalised to sum to 100, as are the time-varying weights. The construction of the aggregate components using these weights is described in the text.

(column 3) and is positively correlated after 1973 (column 4). While some European countries witnessed a post-1973 decline in correlation with this fixed-weight benchmark, many countries in fact experienced an increase. These results differ from those of Baxter and Stockman (1989), who conclude that cross-country correlations of industrial production growth rates have declined markedly in the 'post-Bretton Woods' period. However, they base their conclusions on bilateral correlations with US industrial production growth rates, while the benchmark measure used here is more comprehensive.

One problem with the fixed-weights measure of the common component, as noted earlier, is that it might in fact partly reflect country-specific fluctuations. In particular, large idiosyncratic fluctuations experienced by countries even with relatively small weights would tend to unduly influence the fixed-weight common component. Hence, we now turn to an examination of the time-varying weights.

2.2. *Time-Varying Weights*

Table 1 (centre panel) presents summary statistics for the estimated time-varying weights for each country. The weights are volatile and generally quite skewed. Nevertheless, the means and the ranges of the weights are of some interest.

In comparing the averages (over time) of the time-varying weights to the fixed OECD weights used in the benchmark model, the time-varying weights attribute much less importance to smaller, more highly volatile countries such as France and Spain, and relatively more importance to the US and Canada. In a few cases, the time-varying weights may at first glance be surprising. In particular, Italy has the smallest weight in the aggregate series; this is due to large seasonal fluctuations (in higher moments) associated with the vacation structure in Italy.¹⁰ Because of this, Italy's fluctuations are inherently more idiosyncratic. The time-varying weights model implicitly accounts for the importance of idiosyncratic shocks relative to common shocks when determining the weights, something the benchmark model cannot do (unless the share of idiosyncratic to total shocks remains constant over time). The other surprising case is that of Germany, which has a small weight relative to its fixed OECD weight. Note that Austria and Belgium have larger average time-varying weights than their OECD fixed weights, suggesting that these countries may pick up part of the 'German business cycle' since these economies are closely related to that of Germany and face similar shocks.¹¹ The average weights are somewhat misleading as the weights tend to be very volatile. For instance, in the case of the US, the weights attain a minimum as low as 14.8 and a maximum of 52.0% of the total. The weights for other countries also exhibit a wide range of variation.

¹⁰ Note that the seasonal adjustment procedure used in this paper eliminates seasonal fluctuations only in the conditional mean of each series. Idiosyncratic seasonal fluctuations in the variance, as in the case of Italy, are important for the identification of our time-varying weights. Seasonal fluctuations that are common to all countries will have no effect on the weights with this structure. Consequently, common seasonal fluctuations, if any, would be reflected in the time-varying aggregate.

¹¹ Both Belgium and Austria have relatively strong positive correlations with Germany, suggesting the presence of a common cycle in these countries. Pairwise correlations among all countries are given in Appendix Table A2 on the website.

The time-varying weights in each time period are principally determined by the *relative* fluctuations in industrial production growth across countries. A common seasonal fluctuation will have little effect on the relative weights in a given time period, whereas an idiosyncratic seasonal component (as in the case of Italy) will receive a smaller weight and will, therefore, have a smaller influence on the fluctuations of the overall aggregate. This is apparent in Figure 1, which plots the deseasonalised log differences of monthly industrial production and the estimated time-varying weights for Italy. The deseasonalising procedure leaves a significant amount of residual higher moment seasonality, which leads to downward spikes in the time-varying weights. Figure 2, which shows the deseasonalised log differences of industrial production and the time-varying weights for the US illustrates that such seasonal effects are absent in this case.

Both Figures demonstrate that the time-varying weights are volatile. In mid-1974, the US weight has a downward spike, apparently reflecting the sharp effect of the oil price shock on the US economy. The mirror image of this, of course, is an increase in the relative weights of most other countries, including Italy, in this period. Note, however, that the US weight rises quickly thereafter, revealing the propagation of this shock to other countries.

This episode reveals an important feature of the weights. Although the US accounts for a large fraction of world output, the effect of the oil shock on the common component is estimated to be smaller in the initial period since the US is affected before other countries. As the shock propagates to other countries, however, the weights of all countries return to their 'normal' levels and the full effects of the oil shock are then apparent on the common component. Note that, in our methodology, a country-specific shock (in this case, to US output) that propagates to other countries is observationally equivalent to a common shock that influences output growth in all countries but with different lags. Thus, as discussed earlier, our methodology can not disentangle the different sources of shocks.

Figure 3 shows a plot of the estimated common component.¹² The top panel of this Figure shows the common component constructed as described in (5), while the lower panel shows a cumulated measure of this component. The lower panel provides a clear indication of how the common component reflects, for instance, the global recession in 1974-5, around the time of the first OPEC oil shock, and the recession in the early 1980s. The effects of the post-1973 productivity slowdown are reflected in the slower trend increase in the cumulated common component after 1973. Furthermore, there are no seasonal patterns evident in this common component, indicating the absence of common seasonal patterns in fluctuations in industrial output among the OECD economies.¹³

¹² Since it dampens the effects of idiosyncratic shocks, the common component constructed using time-varying weights has an average volatility, as measured by the standard deviation, that is about 40% lower than the average volatility of the fixed-weight aggregate.

¹³ Regressions of the common component on seasonal dummies confirmed this visual observation. We also found no evidence of residual ARCH in the estimated common component. The Box-Pierce Q-statistic, computed using twelve autocorrelations, was 16.03, well below even the 10% critical value for rejecting the null hypothesis of conditional homoscedasticity of the residuals.




Fig. 1. *Italy*

Notes: The log differences of monthly industrial production shown above are residuals from a regression on a constant and eleven seasonal dummies.

Correlations between the time-varying weighted aggregate series and the individual countries' industrial production growth rates are reported in the last panel of Table 1. There are a few countries for which the correlations are different when




Fig. 2. *United States*

Notes: The log differences of monthly industrial production shown above are residuals from a regression on a constant and eleven seasonal dummies.

compared to the correlations with the fixed-weight aggregate. For instance, the correlation of US fluctuations with the time-varying common component is much higher than its correlation with the fixed-weight common component. On the

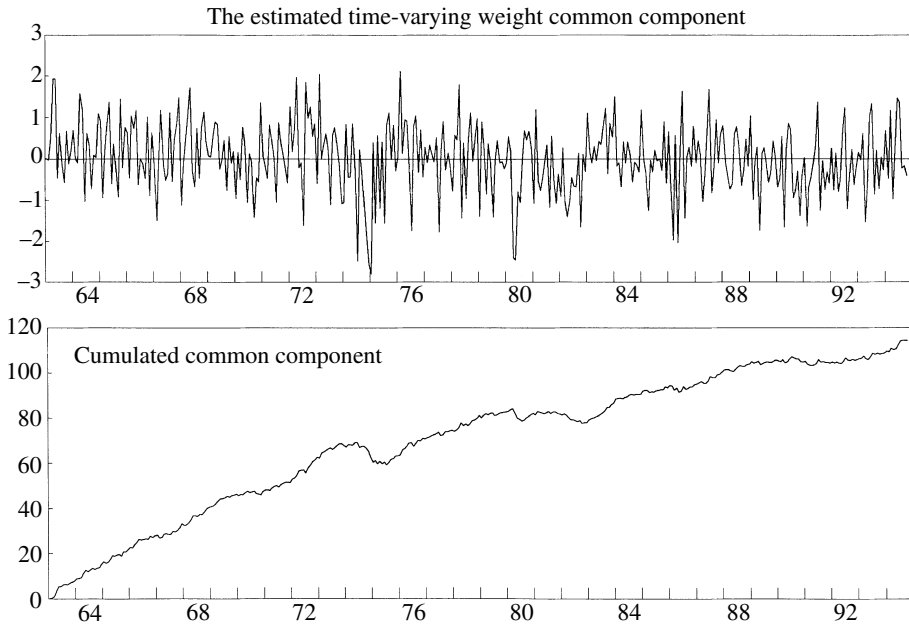


Fig. 3. *The Estimated Time-varying Weight Common Component*

Notes: We have cumulated the common component in the lower panel by adding a fixed weight average of the unconditional mean growth rates for the countries in the sample.

other hand, the correlation for Italy drops sharply when using the time-varying rather than the fixed-weight common component. This reflects the (substantially) lower average weight of Italy in constructing the time-varying common component, which reduces the effect of its idiosyncratic seasonal fluctuations on the common component. In the case of France, however, the full sample correlation with the time-varying component is much higher than with the fixed-weight common component, even though the average time-varying weight for France is much lower than its fixed OECD weight.

A question that might arise at this juncture is the relative importance of global versus country-specific shocks for macroeconomic fluctuations. Unlike in frameworks (such as error components models) that impose the assumption of orthogonality between global and country-specific shocks, however, we cannot directly answer this question in our framework. In particular, we are interested in estimating the component that is common to *all* countries. Thus, there could still be significant correlations between subsets of countries. Most previous literature has focused on identifying common ‘shocks’ and the relative importance of different shocks. We cannot separately identify the nature of individual countries’ shocks but instead attempt to identify the extent to which shocks of any type – seasonal, business cycle, etc. – are common across countries. Nevertheless, the strong positive correlations between individual country industrial production growth fluctuations and the common component

Table A2
Cross-Country Correlations of Industrial Production Growth Rates: Full Sample

	AT	BE	CA	FI	FR	GR	GE	IT	JA	LX	NO	NT	PO	SP	SW	UK	US
Austria	1.00																
Belgium	0.14	1.00															
Canada	-0.03	0.25	1.00														
Finland	0.03	0.27	0.30	1.00													
France	-0.28	0.10	0.24	0.17	1.00												
Greece	0.12	0.14	0.23	0.15	-0.13	1.00											
Germany	0.28	0.27	0.23	0.07	0.08	0.03	1.00										
Italy	0.31	0.04	0.02	-0.16	-0.32	0.33	0.12	1.00									
Japan	0.14	-0.10	0.16	0.01	0.05	0.24	0.09	0.37	1.00								
Luxembourg	0.21	0.05	0.12	-0.12	-0.20	0.23	0.14	0.68	0.28	1.00							
Norway	0.13	-0.17	0.03	-0.04	0.03	0.09	0.07	0.27	0.19	0.20	1.00						
Netherlands	0.32	0.23	0.26	0.15	-0.06	0.26	0.23	0.29	0.15	0.23	0.29	1.00					
Portugal	0.25	0.04	0.08	-0.13	-0.34	0.31	0.08	0.80	0.34	0.65	0.27	0.37	1.00				
Spain	0.17	-0.05	0.07	-0.16	-0.25	0.38	0.13	0.84	0.43	0.63	0.39	0.29	0.77	1.00			
Sweden	0.19	0.12	0.05	0.42	-0.04	0.10	-0.04	-0.03	0.15	-0.12	-0.01	0.16	-0.06	-0.09	1.00		
UK	-0.08	-0.02	0.24	-0.05	0.15	0.00	-0.02	-0.14	0.14	0.02	0.23	0.05	-0.09	-0.02	-0.01	1.00	
US	0.16	-0.06	0.33	-0.04	0.28	-0.07	0.16	-0.11	0.27	-0.04	0.33	0.11	-0.09	-0.05	-0.01	0.32	1.00

are consistent with other evidence that global shocks are quantitatively quite important.¹⁴

2.3. *Evaluating the Time-Varying Weight Common Component*

This sub-section uses two approaches to further evaluate the features of the time-varying weight common component. First, to better understand the comovement between individual country fluctuations and the common component, we regress each country's IP growth rate on a constant and a measure of the common component. The slope coefficient from this regression can be interpreted as a country's 'beta' (analogous to this concept in the finance literature) in that it measures the sensitivity of a country's IP growth to movements in the common component. The first two columns of Table 2 report the full sample results based on regressions with the fixed-weight and time-varying weight common components, respectively. Not surprisingly, in column 2, most countries (13 of 17) have betas that exceed one, confirming that the time-varying aggregate is less volatile than IP growth in individual countries. Further, countries that experience large seasonal fluctuations – including Italy, Norway and Spain – have correspondingly high betas. It also appears that more of the estimated betas in column 2 are closer to unity compared to those in column 1. Standard F-tests confirmed that the null hypothesis of a slope coefficient equal to unity could not be rejected for only 4 countries when the fixed-weight common component is used, compared to 8 when the time-varying weight common component is used. Thus, the greater degree of comovement obtained using the time-varying weights suggests that these weights provide a better measure of the common component.

Next, to characterise the dynamic relationship between the time-varying weight common component (CC) and fluctuations in individual countries, we estimated a set of simple bivariate VARs, using a standard Cholesky decomposition to orthogonalise the shocks. In other words, the structural assumption underlying the VARs is that a shock to country i does not have a contemporaneous effect on the CC, but a shock to the CC can have a contemporaneous effect on the country. Rather than reporting a plethora of results, we only summarise the main features of the results here.¹⁵

We cumulated the impulse responses to measure the level responses of the CC and the individual country IP index to shocks. Interestingly, for all countries other than the US, the effects of shocks to individual country IP growth on the CC were relatively small and transitory (and, in general, not statistically significant). This was true even for large countries such as Germany and Japan. On the other hand, the CC has large and persistent effects on the levels of the IP indices for all countries,

¹⁴ A principal components analysis of our dataset indicated that the first common component obtained using this technique had an R^2 contribution of about 0.25.

¹⁵ Bivariate VARs were run separately for each country using a constant and 12 lags each of the respective country's IP growth rate and the time-varying weight common component. Standard errors for the impulse responses and variance decompositions were computed using Monte Carlo techniques with 1,000 replications.

Table 2
Regressions of Each Country's IP Growth on Common Components

	Fixed Weight Common Component	Time-Varying Weight Common Component		
	Full Sample	Full Sample	1963–73	1974–94
Austria	0.64 (0.12)	1.02 (0.21)	1.52 (0.43)	0.80 (0.24)
Belgium	0.23 (0.14)	0.90 (0.24)	1.68 (0.44)	0.54 (0.29)
Canada	0.40 (0.07)	1.54 (0.10)	1.99 (0.20)	1.33 (0.11)
Finland	-0.23 (0.17)	0.81 (0.28)	0.06 (0.50)	1.15 (0.35)
France	-0.22 (0.15)	0.86 (0.25)	-0.26 (0.56)	1.36 (0.24)
Greece	0.72 (0.15)	1.25 (0.26)	1.22 (0.48)	1.22 (0.31)
Germany	1.22 (0.12)	1.65 (0.22)	2.10 (0.38)	1.49 (0.27)
Italy	4.26 (0.24)	2.23 (0.54)	3.70 (1.17)	1.61 (0.58)
Japan	0.98 (0.06)	1.10 (0.12)	0.95 (0.22)	1.12 (0.51)
Luxembourg	2.41 (0.20)	2.17 (0.38)	2.63 (0.65)	2.06 (0.48)
Norway	2.60 (0.26)	3.43 (0.46)	1.92 (0.78)	4.30 (0.58)
Netherlands	1.14 (0.13)	2.19 (0.22)	2.27 (0.34)	2.18 (0.29)
Portugal	3.15 (0.23)	2.11 (0.47)	3.69 (1.03)	1.45 (0.49)
Spain	4.09 (0.22)	2.65 (0.49)	3.30 (1.06)	2.39 (0.54)
Sweden	0.29 (0.26)	1.18 (0.44)	0.75 (0.45)	1.38 (0.63)
UK	0.75 (0.12)	1.73 (0.19)	1.43 (0.31)	1.93 (0.24)
US	0.30 (0.04)	0.90 (0.06)	0.66 (0.11)	1.01 (0.06)

Notes: The coefficients reported above are from regressions of each country's IP growth rate on the respective common component and a constant. Standard errors are reported in parentheses.

including Japan and the US, indicating the importance of the CC for domestic fluctuations in all of the industrial economies. These impulse responses were nearly all statistically significant (i.e., the two standard error bands did not include zero) at horizons of up to 12 quarters and, for most countries, at longer horizons as well.

We also examined the forecast error variance decompositions from the VARs. Over horizons of 12 to 24 months, the contribution of individual country IP growth rates to the forecast error variance of the CC was quite small and generally less than 10%, even for relatively large countries such as Germany and Japan. The maximum contribution was about 15% in the case of the US. The results were

quite similar at longer forecast horizons. On the other hand, the relative importance of the CC for the forecast error variance of IP growth rates over a 1 to 2 year horizon was much larger and in the range of 15% for large countries such as Germany and Japan as well as many of the smaller countries (these results were statistically significant at the 5% level). The maximum was for the US, at about 50% (also statistically significant). Interestingly, for countries such as Italy and Spain that have large idiosyncratic seasonal fluctuations, the relative importance of domestic fluctuations for the forecast error variance of the CC was barely 5%, similar to that of far smaller countries. Further, the CC explained a relatively small fraction of the forecast error variance of IP growth in these economies.

As another check on our results, we also compared the CC with those estimated by other authors. For instance, after some appropriate transformations to enable such a comparison, we determined that the correlation between the CC and the corresponding common factor for the G-7 estimated by Gregory *et al.* (1997), who use entirely different techniques, was strongly positive (about 0.6). Restricting our sample to the G-7 increased this correlation significantly (about 0.7).¹⁶

Overall, the time-varying weight common component that we have constructed appears to have reasonable properties.

3. Extensions

This Section extends and explores the sensitivity of the results discussed in the previous Section. First, we separately examine the properties of the time-varying weights common component over different sub-periods. This enables us to address the question of whether the correlation of business cycles across countries has changed significantly in the post-1973 period. However, with the reduced-form approach adopted here, we can document these stylised facts but cannot directly attribute changes in the patterns of these correlations to changes in exchange rate regimes or other factors.

Second, we construct a measure of the European common component and examine its properties. There has been growing interest in the relative importance of common economic fluctuations, particularly in the context of European Economic and Monetary Union (EMU). The exchange rate plays a potentially useful role as an adjustment mechanism in response to country-specific shocks. Hence, the relationship between country-specific and common fluctuations could have important implications for the success of a currency union. Finally, we examine the sensitivity of the results to our choice of deseasonalising procedure. In particular, the time-varying weights methodology implicitly accounts for common seasonal fluctuations. Thus, the effects of deseasonalising should be less important with our time-varying aggregate than with the benchmark aggregate. In addition, residual seasonality should also be lower.

¹⁶ To enable these comparisons, we took quarterly changes of our (cumulated) common component and restricted the sample to 1970–93. We are grateful to Allen Head for providing the Gregory *et al.* world common factor.

3.1. *Sub-Sample Analysis*

Around 1973, the Bretton Woods system of fixed exchange rates met its demise and a major oil price shock occurred. This period also heralded the beginning of a distinct slowdown in productivity growth among industrial countries. Although we have little to say directly about these shocks and how they influence the common component, it is nevertheless of interest to see if our results differ across these two periods (which are often referred to as the Bretton Woods and post-Bretton Woods periods). Table A1 on the website documents that industrial production growth has slowed in all countries during the post-1973 period. Based on standard deviations of the data, however, there does not seem to be a systematic commensurate change in volatility. We investigate this more thoroughly in this Section. Failure to control for the mean change could result in misleading inference about the conditional variance (Lumsdaine and Ng, 1999) which, in turn, could affect the accuracy of the time-varying weights. To investigate this possibility, we estimate a modified version of (3):

$$y_{it} = c_i + c_{1i}\mathbf{1}(t > 1973 : 12) + \varepsilon_{it}, \quad \varepsilon_{it}|I_{t-1} \sim N(0, h_{it}) \quad (6a)$$

$$h_{it} = w_i + \alpha_i \varepsilon_{it-1}^2 + \beta_i h_{it-1}, \quad (6b)$$

where $\mathbf{1}(A)$ is an indicator variable equal to 1 if event A is true and 0 otherwise. That is, in the deseasonalised data, we allow for a change in mean after 1973.¹⁷

The means of the associated time-varying weights estimated using this specification were similar across the pre- and post-1973 periods. The correlations of individual country industrial production growth fluctuations with the time-varying weight common component for these two periods are reported in the last two columns of Table 1. For most countries, the correlations are similar across the two subperiods. The US and certain European countries including Finland, France, Norway, and Spain have more strongly positive correlations with the common component in the second period. On the other hand, the correlations with the common component decline in this period for some countries such as Belgium, Germany, Portugal, and Sweden.

Of particular interest is the comparison of the betas (as before, these are the coefficients from regressions of individual country IP growth on the common component and a constant) between the two periods, as shown in the last two columns of Table 2. In the post-1973 period, 13 of the 17 countries have betas that are closer in 12 cases or equally close in 1 case to unity than the betas in the earlier period. This suggests that, in this period, fluctuations in IP growth in industrial countries have become more synchronised with the common component,

¹⁷ Alternatively, we could estimate separate GARCH(1,1) models for the two subperiods; such a procedure is problematic due to the diminished number of observations. Accurate estimation of the GARCH(1,1) model typically requires a large number of observations; see, e.g., Hong (1987) and Lumsdaine (1995).

supporting the notion that macroeconomic fluctuations in industrial countries have become more closely linked since 1973; see, e.g., Gerlach (1988). We do not find evidence to support the notion that economic fluctuations have become substantially more country-specific after 1973; see, e.g., Baxter and Stockman (1989). In our view, the main conclusion to be drawn from these results is that virtually all countries have a strong positive correlation with the common component in international fluctuations, particularly after 1973, confirming the existence of a 'world business cycle'.

3.2. *European Common Component*

The methodology developed in this paper can easily be applied to examine common cycles among different groups of countries. Once the basic model is estimated, all that is required is a renormalisation of the weights; individual country fluctuations for the relevant group can then be aggregated to compute the group-specific common component.

This section examines measures of the common component in European economic fluctuations, constructed using all countries in the sample except Canada, Japan and the US. The fixed weight component uses the same OECD 1990 weights discussed earlier while the time-varying component is constructed using (4) and (5); both sets of weights are normalised to sum to 100 for the European countries in each time period.

Table 3 reports summary statistics for the time-varying weights and the correlations of each country's industrial production growth rate with both the fixed and variable weight measures, for the full sample and also for the pre- and post-1973 subsamples. As in the case of the world common component, Italy and Spain apparently experience many idiosyncratic shocks and thus receive substantially less weight using our time-varying method than in the fixed-weight aggregate.

The correlations of individual country fluctuations with the European common component are strongly positive for virtually all of the European countries. The last column of Table 3 indicates that this result is more evident in the post-1973 period and confirms the existence of a 'European business cycle'.¹⁸ For most European countries, the full sample correlation with the European common component is significantly stronger than the correlation with the world common component. An interesting finding is that, despite their relatively large weights in the construction of the European common component, both France and the UK have higher correlations with the world common component than with the European common component. Fluctuations in the US were negatively correlated with the European common component before 1973 but are positively correlated after 1973. Fluctuations in Japan and Canada are positively correlated with the European component in both periods. Also, perhaps not surprisingly, the aggregate constructed with time-varying weights is more highly correlated with the time-varying world common component than with the fixed-weight counterpart.

¹⁸ Artis and Zhang (1999) reach a similar conclusion after examining bivariate cross-country correlations of output fluctuations using a number of different detrending techniques.

Table 3
Correlations with the Common Component of European IP Growth Fluctuations

	Fixed-Weight Measure of the Common Component				Time-Varying Weights			Time-Varying Measure of the Common Component		
	OECD weights	Full Sample	1963–73	1974–94	Mean	Min.	Max.	Full Sample	1963–73	1974–94
Austria	2.60	0.29	0.28	0.29	12.29	3.67	17.13	0.46	0.41	0.49
Belgium	2.60	0.20	0.20	0.20	7.13	3.51	10.92	0.45	0.54	0.39
Finland	1.30	-0.03	-0.30	0.13	12.07	1.19	22.37	0.29	0.14	0.36
France	16.10	-0.09	-0.22	0.02	8.89	1.34	13.19	-0.04	-0.25	0.11
Greece	0.80	0.30	0.32	0.29	10.51	3.03	15.51	0.47	0.53	0.42
Germany	26.40	0.51	0.48	0.54	9.92	2.68	14.31	0.44	0.54	0.38
Italy	17.40	0.81	0.86	0.76	0.91	0.44	1.23	0.45	0.48	0.43
Luxembourg	0.30	0.61	0.77	0.52	4.03	1.78	5.17	0.45	0.56	0.39
Norway	1.30	0.40	0.27	0.47	2.01	0.51	3.00	0.30	0.07	0.41
Netherlands	3.90	0.45	0.38	0.50	12.06	6.55	17.05	0.67	0.70	0.65
Portugal	1.80	0.68	0.68	0.68	1.22	0.34	1.61	0.44	0.47	0.43
Spain	8.00	0.79	0.83	0.75	0.98	0.57	1.35	0.43	0.42	0.45
Sweden	2.30	0.02	-0.15	0.08	2.46	0.29	3.57	0.24	0.08	0.29
UK	15.30	0.19	0.16	0.21	15.52	4.52	22.58	0.25	0.21	0.27
Canada		0.21	0.26	0.16				0.32	0.48	0.20
Japan		0.38	0.43	0.35				0.21	0.14	0.22
US		0.02	-0.11	0.09				0.06	-0.12	0.14
World common component		0.54	0.52	0.56				0.69	0.71	0.68

Notes: The European common component was constructed using all the countries in the sample excluding Canada, Japan, and the US. The fixed OECD weights and the time-varying weights were normalised to sum to 100 for the European countries.

We also examined various other regional common components and note a couple of the main results here. We computed common components for EMU (excluding Finland, Norway, Sweden and the UK from the 14 European countries in the sample) and core EMU countries (Austria, Belgium, Germany, France, Italy, Luxembourg and the Netherlands). These two common components were strongly positively correlated mutually (0.92) and with the broader European common component (0.78 and 0.81, respectively). However, the correlations of fluctuations in the UK with these common components were significantly lower than with the overall European common component.

3.3. Seasonal Adjustment

As discussed in Section 1, the procedure for deseasonalising unadjusted data could potentially have a large impact on the empirical results. The time-varying weights methodology developed in this paper should, in principle, discriminate between country-specific and common seasonal fluctuations and adjust each country's weights accordingly. But, as noted earlier, we removed seasonal means from each country's data by regressing on a set of monthly dummies in order to avoid the problems that could result from the misspecification of conditional means. To

examine the sensitivity of the results to this procedure, we recomputed the time-varying weights and the international common component using unadjusted data. The use of unadjusted data may be viewed as allowing for common deterministic seasonal fluctuations to be reflected in the common component.

To conserve space, we summarise only the main results here. The relative ranking in terms of average weights was roughly similar to that in Table 1 although there were some differences. The mean weight for the US was higher at 53.1% while the weights for Canada and Japan were smaller, suggesting that the deterministic components of seasonal fluctuations in the latter two countries are idiosyncratic. The correlations between individual country fluctuations and the common component were generally higher than those reported in Table 1, indicating that part of the fluctuations that are captured by deterministic seasonal dummies is similar across countries. We are reluctant to make too much of these results because of the possible misspecification problems that could arise from the use of unadjusted data. Nevertheless, the principal result about the existence of a substantial common component in international fluctuations is confirmed by these correlations.

4. Concluding Remarks

This paper has proposed a new methodology for estimating the common component of international economic fluctuations. The methodology accounts for relative country size and also implicitly accommodates the effects of the cross-country propagation of shocks, without imposing a formal structure on the dynamic propagation of these shocks across countries. In addition, it provides a unified treatment of seasonal and business cycle fluctuations, allowing for correlations between these fluctuations while eliminating the impact of idiosyncratic seasonal variation on the common component.

The methodology is based on two properties of fluctuations in industrial production growth rates that were documented in this paper. The first is the negative relationship between country size and the volatility of industrial production growth rates among OECD industrial countries. The second property is that industrial production growth rates exhibit evidence of conditional heteroscedasticity. Combining these two features suggests a time-varying weighting scheme for measuring the common international component where the time-varying weights are inversely proportional to the relative conditional variance of industrial production growth rates for each country.

The methodology has potential applications for aggregation in a wide variety of other contexts where conditional volatility provides a natural stochastic specification with which to form time-varying weights. Possible further applications include the construction of stock market indices and aggregate price indices. An important virtue of the methodology, which is relevant in this context, is that it easily handles aggregation across a large number of series. Other aggregation approaches, including dynamic factor models, rapidly become unwieldy (or require the use of strong ordering assumptions) as the number of individual units in the analysis increases.

Another interesting extension of this approach would be to examine if there is systematic variation in the conditional volatility of output growth over different phases of the business cycle; see the discussion in Diebold and Rudebusch (1996). This could have implications for business cycle modelling as well as forecasting. The model developed in this paper could also be extended to test for business cycle asymmetries, although, as noted earlier, we do not see compelling reasons why the weighting scheme itself should treat positive and negative shocks asymmetrically.

In the empirical example considered here, we found that industrial production growth fluctuations in virtually all countries in the sample have strong, positive correlations with the common component of international fluctuations constructed using time-varying weights; these correlations are generally stronger after 1973. Similar results were obtained when we constructed a time-varying measure of the common component in European economic fluctuations. Virtually all European countries in the sample had strong, positive correlations with this common component, which was distinct from the world common component. These results are consistent with other research showing the importance of common international influences in driving business cycle fluctuations in industrial economies.

The methodology can also easily be extended to examine common components of fluctuations for different groups of countries. This can be used to show, for instance, how synchronous the fluctuations are across a group of countries and, therefore, how well-suited such a group might be (in terms of patterns in historical data) for a currency union, which implies the loss of the exchange rate as an adjustment mechanism. We find, for instance, strong positive correlations of fluctuations in individual EMU countries with the common component for the EMU region. Similarly, correlations with the common component could help assess a given country's suitability for entering a currency union in terms of a key criterion—the commonality of fluctuations.

Deutsche Bank and NBER Research Department, IMF

Date of receipt of first submission: June 2000

Date of receipt of final transcript: February 2002

Technical Appendix is available for this paper: www.res.org.uk

Dataset is available for this paper: www.res.org.uk

References

- Altonji, J. and Ham, J. C. (1990). 'Variation in employment growth in Canada: the role of external, national, regional and industrial factors', *Journal of Labor Economics*, vol. 8, pp. S198–236.
- Artis, M. J. and Zhang, W. (1999). 'Further evidence on the international business cycle and the ERM: is there a European business cycle?', *Oxford Economic Papers*, vol. 51, pp. 120–32.
- Backus, D. K. and Kehoe, P. J. (1992). 'International evidence on the historical properties of business cycles', *American Economic Review*, vol. 82, pp. 864–88.
- Baxter, M. and Stockman, A. C. (1989). 'Business cycles and the exchange rate regime: some international evidence', *Journal of Monetary Economics*, vol. 23, pp. 377–400.
- Bayoumi, T. and Prasad, E. S. (1997). 'Currency unions, economic fluctuations, and adjustment: some new empirical evidence', *IMF Staff Papers*, vol. 44, pp. 36–58.

- Beaulieu, J. J. and Miron, J. A. (1992). 'A cross-country comparison of seasonal cycles and business cycles', *ECONOMIC JOURNAL*, vol. 102, pp. 772–88.
- Beaulieu, J. J., MacKie-Mason, J. K. and Miron, J. A. (1992). 'Why do countries and industries with large seasonal cycles also have large business cycles?', *Quarterly Journal of Economics*, vol. 107, pp. 621–56.
- Bollerslev, T. (1986). 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics*, vol. 31, pp. 307–27.
- Bollerslev, T. (1990). 'Modeling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model', *Review of Economics and Statistics*, vol. 72, pp. 498–505.
- Canova, F. and Ghysels, E. (1994). 'Changes in seasonal patterns: are they cyclical?', *Journal of Economic Dynamics and Control*, vol. 18, pp. 1143–71.
- Cecchetti, S. G. and Kashyap, A. K. (1996). 'International cycles', *European Economic Review*, vol. 40, pp. 331–60.
- Cecchetti, S. G., Kashyap, A. K. and Wilcox, D. W. (1997). 'Interactions between the seasonal and business cycles in production and inventories', *American Economic Review*, vol. 87, pp. 884–92.
- Christofferson, P. F. and Giorgianni, L. (2000). 'Interest-rate arbitrage in currency baskets: forecasting weights and measurement risk', *Journal of Business and Economic Statistics*, vol. 18, pp. 242–53.
- Deutsch, M., Granger, C. W. J. and Teräsvirta, T. (1994). 'The combination of forecasts using changing weights', *International Journal of Forecasting*, vol. 10, pp. 47–57.
- Diebold, F. X. and Nerlove, M. (1989). 'The dynamics of exchange rate volatility: a multivariate latent factor ARCH model', *Journal of Applied Econometrics*, vol. 4, pp. 1–21.
- Diebold, F. X. and Pauly, P. (1990). 'Use of prior information in forecast combination', *International Journal of Forecasting*, vol. 6, pp. 503–8.
- Diebold, F. X. and Rudebusch, G. D. (1996). 'Measuring business cycles: a modern perspective', *Review of Economics and Statistics*, vol. 78, pp. 67–77.
- Engle, R. F. (1982). 'Autoregressive conditional heteroskedasticity', *Econometrica*, vol. 50, pp. 987–1008.
- Engle, R. F. and Hylleberg, S. (1996). 'Common seasonal features: global unemployment', *Oxford Bulletin of Economics and Statistics*, vol. 58, pp. 615–30.
- Engle, R. F. and Kozicki, S. (1993). 'Testing for common features', *Journal of Business and Economic Statistics*, vol. 11, pp. 369–95.
- Forni, M. and Reichlin, L. (1996). 'Dynamic common factors in large cross-sections', *Empirical Economics*, vol. 21, pp. 27–42.
- Franses, P. H., Hylleberg, S. and Lee, H. S. (1995). 'Spurious deterministic seasonality', *Economics Letters*, vol. 48, pp. 249–56.
- Gerlach, S. (1988). 'World business cycles under fixed and flexible exchange rates', *Journal of Money, Credit, and Banking*, vol. 20, pp. 621–32.
- Ghysels, E. (1993). 'Seasonal adjustment and other data transformations', Université de Montréal Working Paper.
- Gregory, A. W., Head, A. C. and Raynauld, J. (1997). 'Measuring world business cycles', *International Economic Review*, vol. 38, pp. 677–702.
- Hamilton, J. D. (1989). 'A new approach to the economic analysis of nonstationary time series and the business cycle', *Econometrica*, vol. 57, pp. 357–84.
- Head, A. C. (1995). 'Country size, aggregate fluctuations and international risk sharing', *Canadian Journal of Economics*, vol. 28, pp. 1096–119.
- Hong, C. (1987). 'The integrated generalized autoregressive conditional heteroskedasticity model: the process, estimation and some Monte Carlo experiments', unpublished manuscript, University of California, San Diego.
- Lippi, M. and Reichlin, L. (1994). 'Common and uncommon trends and cycles', *European Economic Review*, vol. 38, pp. 624–35.
- Lumsdaine, R. L. (1995). 'Finite sample properties of the maximum likelihood estimator in GARCH(1,1) and IGARCH(1,1) models: a Monte Carlo investigation', *Journal of Business and Economic Statistics*, vol. 13, pp. 1–10.
- Lumsdaine, R. L. (1996). 'Consistency and asymptotic normality of the IGARCH(1,1) and covariance stationary GARCH(1,1) models', *Econometrica*, vol. 64, pp. 575–96.
- Lumsdaine, R. L. and Ng, S. (1999). 'Testing for ARCH in the presence of a possibly misspecified conditional mean', *Journal of Econometrics*, vol. 83, pp. 257–59.
- Nelson, C. R. and Plosser, C. I. (1982). 'Trends and random walks in macroeconomic time series: some evidence and implications', *Journal of Monetary Economics*, vol. 10, pp. 139–62.
- Norrbin, S. and Schlagenhauf, D. (1996). 'The role of international factors in the business cycle: a multicountry study', *Journal of International Economics*, vol. 40, pp. 85–104.
- Pagan, A. (1997). 'Towards an understanding of some business cycle characteristics', *The Australian Economic Review*, vol. 30, pp. 1–15.

- Stock, J. H. and Watson, M. W. (1989). 'New indexes of coincident and leading economic indicators', in (O. Blanchard and S. Fischer, eds), *NBER Macroeconomics Annual*, Cambridge, MA: MIT Press, pp. 351–94.
- Stock, J. H. and Watson, M. W. (1993). 'A procedure for predicting recessions with leading indicators: econometric issues and recent experience', in (J.H. Stock and M.W. Watson, eds), *Business Cycles, Indicators and Forecasting*, Chicago: University of Chicago Press, pp. 255–84.
- Stockman, A. (1988). 'Sectoral and national aggregate disturbances to industrial output in seven European countries', *Journal of Monetary Economics*, vol. 21, pp. 387–409.