Why has the U.S. economy become less correlated with the rest of the world?

by JONATHAN HEATHCOTE AND FABRIZIO PERRI[1]

In this paper we do two things. First we document that over the last 40 years the United States business cycle has become less synchronized with the cycle in the rest of the world. Second we try to explain why this has happened. We use a general equilibrium model as a tool to discriminate between two alternative explanations: (i) a change in the nature of real shocks, and (ii) an increase in U.S. financial integration with the rest of the world. Our results indicate that financial integration has played the major role in producing the observed changes in international co-movement.

1. Measuring changes in business cycle synchronization

We are interested in measuring the change in business cycle correlation between two regions: the U.S. and the rest of the industrialized world. We identify the rest of the world as an aggregate of the 15 countries of the European Union plus Japan. Our period of interest is from the first quarter of 1960 to the second quarter of 2002; this is the longest period for which we could find comparable data series for both areas. For both regions we use data on GDP, gross fixed capital formation, consumption and civilian employment. Due to the presence of wars in our sample we subtract government consumption from consumption and from GDP.[2]

When measuring synchronization between two macroeconomic series it is important to define the range of cycles in which we are interested. This is usually done by filtering the series to remove cycles at frequencies that are not of interest. Two popular filters are the Hodrick and Prescott (HP) filter that approximately removes cycles of length longer than 8 years, and the first difference filter that emphasizes shorter cycles. Before we focus on a particular filter we document the change in cross-country correlation across our sample period for different ranges of cycles. To do so we use a high pass filter (see Marianne Baxter and Robert G. King 1999, and Lawrence J. Christiano and Terry J. Fitzgerald 1999) that can be used extract cycles of frequency shorter that an arbitrary cutoff. In figure 1 we show the inter-regional correlation of four high-pass-filtered macroeconomic series with cutoffs ranging from 6 to 12 years. The solid line reports the various correlations in the first half of our sample (1960.1-1981.1) while the crossed line reports the correlations in the second half of the sample (1981.2 2002.2).The fact that the crossed line lies always below the solid line reveals a general decline in business cycle synchronization from the first to the second sub-sample. The decline is quite large and significant for medium length cycles (with a cut-off between 8 and 10 years), especially for investment, consumption and employment.

FIGURE 1 APPROXIMATELY HERE

In table 1 we report the change in correlation in the two sub-samples obtained using the HP filter. These results confirm the findings of the high pass filter. The cross-country correlations for all variables have declined, with the largest decline seen for investment. In the last column of the table we report the correlation of productivity, which we take as a measure of the correlation of real shocks.[3] Interestingly productivity does not display any decline in correlation. This suggests that a change in the shock process did not play a major role in driving the decline in business cycle synchronization. In the theoretical section of the paper we explore this hypothesis in more detail.

TABLE 1 APPROXIMATELY HERE

Finally in figure 2 we provide some evidence that the decline in business cycle synchronization does not depend on how we split our sample. The picture reports the cross-country correlation of the four macro variables of interest (HP filtered) in successive 20 year windows, starting with 1960.1-1979.4 and ending with 1982.3-2002.2. The picture shows a pretty uniform decline in correlation starting in the early 1990s. In the next section we use a simple general equilibrium model to shed more light on the causes of this decline.

FIGURE 2 APPROXIMATELY HERE

2. Model

The modeling framework is the one developed by David K. Backus, Patrick J. Kehoe and Finn E. Kydland (1994) with the asset market structure introduced in Heathcote and Perri (2002b). There are two countries, each of which is populated by a continuum of identical, infinitely-lived households. The two countries are perfectly symmetric. Thus in what follows we focus primarily on the domestic economy; where necessary we use star superscripts to denote foreign variables.

Domestic households maximize expected discounted utility over consumption and leisure:

$$(1) \qquad E\sum_{t=0}^{\infty}\beta^t\left(\mu\ln(c_t)+(1-\mu)\ln(1-n_t)\right).$$

Households supply labor to perfectly-competitive intermediate-goods-producing firms. These firms use country-specific capital and labor to produce internationally-traded intermediate goods. The intermediate good produced in the domestic country is labeled a, while the good produced in the foreign country is labeled b.The production technology is given by

$$(2) \qquad F(z_t,k_t,n_t)=e^{z_t}k_t^{\theta}n_t^{1-\theta}$$

where z_t is an exogenous country-specific technology shock. The law of motion for the vector of shocks $\widehat{z}_t=[z_t,z_t^*]$ is given by

$$(3) \qquad \widehat{z}_{t+1}=A\widehat{z}_t+\widehat{\varepsilon}_{t+1}$$

where A is a 2×2 matrix, and $\widehat{\varepsilon}_{t+1}$ is a 2×1 vector of independently normally distributed random variables with variance-covariance matrix Σ.

Within each country Cobb-Douglas technologies aggregate goods a and b to produce a country-specific consumption / investment good. Thus

$$c_t+x_t=a_t^{\omega}b_t^{1-\omega}$$

where $\omega>0.5$ determines the size of the local input bias in the composition of domestically produced final goods.

We assume that intermediate-goods-producing firms hold capital and make investment decisions. They maximize expected discounted profits, given by

$$
\begin{aligned}
& E\sum_{t=0}^{\infty} Q_t d_t, \\
d_t &= q_{at}\left[F\left(z_t, k_t, n_t\right) - w_t n_t\right] - x_t
\end{aligned}
$$

where Q_t (discussed below) is the price the firm uses to discount dividends at t relative to consumption at date 0, and dividends in a given period are equal to output minus payments to labor minus investment (denoted x_t). Here q_{at} denotes the price of good a (the output of the domestic firm) relative to consumption, which in equilibrium is equal to the marginal product of good a in the production of the final consumption good.

The only assets traded internationally are assumed to be shares in the domestic and foreign firms. Dividend income from abroad is potentially taxed at rate τ. We further assume that all asset trade occurs in the first period. In Heathcote and Perri (2002b) we show that for different values for the tax on foreign dividends, this model nests the two extreme possibilities for international risk sharing. When the tax rate is zero, allocations are equivalent to those arising under a regime of unrestricted stock trade and are also equivalent to those arising when international financial markets are complete.[4] When the tax rate is high enough instead, allocations are equivalent to those under financial autarky (no international asset trade). For intermediate values for τ, partial risk-sharing is achieved.

In every period except the first, the household simply consumes the sum of labor income and any dividend income from its shareholdings. Thus for $t \geq 1$ the state by state budget constraint is given by

$$
c_t \leq q_{at} w_t n_t + \lambda d_t + \lambda^f r x_t (1-\tau) d_t^* + \psi_t. \tag{4}
$$

Here λ (λ^f) denotes the fraction of the domestic (foreign) firm held by the domestic household. Foreign dividend income is taxed locally and revenue ψ_t is redistributed to domestic households in a lump-sum fashion. Note that foreign dividends are in units of the foreign consumption good, so they are multiplied by the real exchange rate rx_t in the budget constraint.

At the start of period 0, the domestic household owns the entire domestic firm. In this period alone the household chooses purchases of domestic and foreign stocks subject to the budget constraint

$$(5) \qquad c_0 + P_0\lambda + rx_0 P_0^* \lambda^f \leq q_{a,0} w_0 n_0 + P_0 + d_0.$$

In this equation P_0 denotes the (ex-dividend) price of the domestic firm in units of period 0 domestic consumption, and P_0^* denotes the price of the foreign firm in units of period 0 foreign consumption.

At date 0, domestic households choose portfolio weights λ and λ^f and consumption and labor supply for all dates and states to maximize expected discounted utility (eq. 1) subject to eqs. 4 and 5. Note that since the domestic and foreign economies are assumed perfectly symmetric at date 0, in equilibrium $P_0 = P_0^*$, $rx_0 = 1$, and $\lambda^f = 1 - \lambda = \lambda^{h*} = 1 - \lambda^*$ where λ^* (λ^{h*}) denotes the fraction of the foreign (domestic) firm held by the foreign household.

We assume that firms price state-contingent dividends using a weighted sum of the domestic and foreign shareholders' stochastic discount factors. We set these weights equal to the fractions of the firm owned by the domestic and foreign shareholders (see Heathcote and Perri 2002b for more discussion of this issue).

3. Experiment

The goal of the paper is to try to understand what might have led to the changes in international business-cycle co-movement documented in the data section. The two prime candidates are (i) a change in the productivity shock process, and (ii) an increase in portfolio diversification.

Although, as reported in table 1, the correlation of productivity did not significantly change our measure of productivity might not be too accurate because we have only crude measures of inputs to production. We do not have quarterly capital stock data, and available series for hours do not cover the entire sample period for every country, so we have to rely on employment indices as a measure of the labor input. In addition, while international portfolio diversification has undoubtedly risen, it is hard to exactly quantify the size of the increase. For example, the value of foreign equity holdings in the U.S. increased dramatically in the 1990s, but in part this reflected surging world equity markets rather than increased diversification.[5]

In light of these measurement issues, we let the model discriminate between the two hypotheses. The moments we are primarily interested in explaining are the cross-country correlations of output, consumption and investment. We therefore set the values of parameters that are not particularly crucial for cross-country correlations to standard values and then perform a simple method of moments estimation exercise to let the model tell us what values for λ and for the key parameters in the productivity process generate the best fit with the empirical cross-country correlations. We do this exercise twice, once for each sample period.

We set the household's discount factor β to 0.99, capital's share in production α to 0.36 and the depreciation rate δ to 0.025. Consumption's share in utility μ is set so that households so that on average the representative household works around 30 percent of the time (the implied μ is 0.34) and the share of domestic intermediate inputs in production of the consumption / investment good ω is set to 0.85, so that on average imports are 15 percent of GDP.

The remaining parameter values are the tax rate τ and the parameters defining the productivity process. Since tax revenues are rebated lump-sum, the only place the tax rate appears in the system of equations characterizing equilibrium is in the first order condition defining the optimal choice for λ at date 0. Thus we can essentially treat λ as a free parameter, and characterize equilibrium allocations conditional on a choice for λ with no reference to τ. Given a value for λ in the interval $[0, 1]$, figure 8 in Heathcote and Perri (2002b) illustrates how this level of diversification can be supported in equilibrium by the appropriate choice for the tax rate τ. Proposition 1 in that paper shows that if $\tau = 0$ the equilibrium value for λ is given by

$$\lambda^* = 1 - \frac{1-\omega}{1+\theta-2\omega\theta}$$

which equals 0.7995 given our choices for ω and θ.

The joint auto-regressive process for productivity is described by eq. 3. In order to reduce the number of free parameters somewhat, we assume that the process is perfectly symmetric (so that $A_{11} = A_{22}$, $A_{12} = A_{21}$, $\Sigma_{11} = \Sigma_{22}$ and $\Sigma_{12} = \Sigma_{21}$). We also assume, consistently with a large amount of empirical evidence, that the productivity process is non-stationary (see, for example, Baxter and Mario J. Crucini 1995). In particular, given a value for A_{12}, the off-diagonal element of the A matrix, we assume that $A_{11} = 1 - |A_{12}|$.

At this point we are left with four independent parameter choices: λ, A_{12}, σ (where $\sigma = \sqrt{\Sigma_{11}}$, the standard deviation of the innovations) and ρ (where $\rho = \Sigma_{12}/\Sigma_{11}$, the correlation of the innovations). The moments target for each sub-sample of the data are the cross-country correlations of output, consumption and investment, and the percentage standard deviation of output. We first linearize the system of equations characterizing equilibrium around the non-stochastic steady state with diversification level λ. We then simulate the model 100 times, each time using a simulation length of 80 periods (the same length as our data sample).

4. Results

In table 2 we report the estimated values of the four parameters discussed above. The model tells us that what is required to fit the facts is a substantial increase in international portfolio diversification. The estimated value for λ declines across the two sample periods from 1.005 (essentially perfect home bias) in period 1 to 0.76 (close to perfect diversification). By contrast the estimated values for the parameters defining the shock process change very little; the main differences across the two sample periods is that correlation of innovations to productivity declines from 0.096 to 0.026. We take these results as evidence that international diversification is a key factor in explaining the changes in business cycle synchronization.

In table 3 we report business cycle statistics in each of the two subperiods for the data and for the model. For the second subperiod we report results when we estimate both the productivity parameters and the diversification parameter (see the row labeled 'model 1') and for an additional experiment in which we reduce λ to 0.76 but leave the productivity parameters equal to their estimated values for the first period (see the row labeled 'model 2').

First, note that model 1 reproduces exactly the targeted cross-country correlations in each sub-period. This is an interesting finding in its own light, since the original complete markets model of Backus, Kehoe and Kydland (1994) failed to deliver strong positive international co-movement and failed to deliver a cross-country correlation in output exceeding the correlation in consumption. Note also that model 2 is able to reproduce the correlations in the second period quite closely, confirming that increasing financial integration is indeed the key factor in explaining the observed correlation changes.

TABLES 2 and 3 APPROXIMATELY HERE

How well does the model do in terms of replicating empirical moments that were not directly targeted? The model does a reasonable job in both sub-samples in terms of replicating the volatilities of consumption and investment. The model is also consistent with the observed increase in the standard deviation of net exports, although it underpredicts the volatility of net exports in the first period. The model does not replicate the observed decline in the cross-country correlation of employment. If we explicitly target this moment rather than the cross-country consumption correlation, the estimated increase in diversification is reduced but remains large, the estimated reduction in the correlation of the shocks remains small, but now the predicted decline in the consumption correlation becomes very small.

This suggests that an additional source of shocks may be required to account for the entire set of changes in cross-country correlations. A natural candidate would be monetary shocks, since these were presumably much more strongly correlated in our first sub-sample, which includes a period of fixed exchange rates, than in the second when exchange rates were floating.

Which parameter values are important for particular moments? In order to account for the fact that the observed decline in the investment correlation was much larger than the decline in the output correlation, a decrease in λ is required. The reason is that by permitting inter-temporal specialization in production, increased diversification tends to reduce international co-movement in investment but has a relatively small effect on the output correlation, since the output correlation more closely follows the correlation of the shocks.

The gap between the correlation of output and the correlation of consumption is partly driven by A_{12}, the spill-over term in the productivity process. Small negative spill-overs reduce the consumption correlation relative to the output correlation since a good shock abroad leads agents to expect a negative future impact at home. Financial integration accounts for the fact that the decline in the cross-country correlation of consumption has been larger than the decline in the correlation of output. The surprising result that financial integration reduces the relative correlation of consumption is independent of the changes in the shock process (it holds even when we change only financial integration, as in model 2) and it is related to results discussed by Linda Tesar (1993), Athanasios V. Arvanitis and Anne Mikkola (1996) and Michael R. Pakko (1997). All these authors argue that agents use international financial markets for two functions: (i) to reduce fluctuations in the path for the total consumption bundle through time, and (ii) to reduce deviations from the optimal mix between home and foreign goods in this bundle. The first function implies that more integrated financial markets should lead to more strongly correlated consumption, but the second can imply that more integration leads to less synchronized consumption.[6]

References

[1] Arvanitis, Athanasios V. and Mikkola Anne "Asset market structure and international trade dynamics.", American Economic Review, May 1996 (Papers and Proceedings), 86(2), pp. 67-70.

[2] Backus, David K., Kehoe Patrick J. and Kydland Finn E. "Dynamics of the trade balance and the terms of trade: the J-curve?" American Economic Review, March 1994, 84(1), pp. 84-103.

[3] Baxter, Marianne and Crucini, Mario J. "Business cycles and the asset structure of foreign trade." International Economic Review, November 1995, 36(4), pp. 821-854.

[4] Baxter, Marianne and King Robert G. "Measuring business cycles: approximate band-pass filters for economic time series." Review of Economics and Statistics, November 1999, 81(4), pp. 575-593.

[5] Christiano Lawrence J. and Fitzgerald Terry J. "The band pass filter." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 7257, July 1999.

[6] Cole, Harold L. and Obstfeld, Maurice 1991, Commodity trade and international risk sharing, Journal of Monetary Economics, August 1991, 28(1), 3-24.

[7] Heathcote, Jonathan. and Perri, Fabrizio, "Financial autarky and international business cycles", Journal of Monetary Economics, April 2002a, 49 (3), 601-627.

[8] Heathcote, Jonathan. and Perri, Fabrizio, "Financial globalization and real regionalization." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 9292, October 2002b.

[9] Pakko, Michael R. "International risk sharing and low cross-country consumption correlations: are they really inconsistent?", Review of International Economics, August 1997, 5(3), pp. 386-400.

[10] Tesar Linda, "International risk sharing and non traded goods." Journal of International Economics, August 1993, 35(1), pp. 69-89.

Notes

[1]Heathcote: Department of Economics, Georgetown University, Washington, DC 20057. Perri: Department of Economics, Stern School of Business, New York University, New York, NY 10012, NBER and CEPR. We thank Hal Cole and Linda Tesar for their comments. Heathcote thanks the Economics Program of the National Science Foundation for financial support.

[2]The source of the data is OECD Quarterly National Accounts. More details on the data construction and the actual data used are available at http://pages.stern.nyu.edu/~fperri/research_data.htm

[3]We compute productivity as $\log(GDP) - \alpha \log(Employment)$ where α is set to 0.64.

[4]This result extends the finding of Harold L. Cole and Maurice Obstfeld, 1991 to an environment with a richer production structure.

[5]See Heathcote and Perri 2002b for a discussion of alternative measures of international diversification.

[6]Lower substitutability between home and foreign goods, a stronger bias in preferences towards home-produced goods, and higher willingness to substitute consumption inter-temporally all work to lower consumption co-movement when asset market integration increases.

Table 1. Cross-country correlations (HP filtered series)

	GDP	$Inv.^{(**)}$	$Cons.^{(*)}$	$Emp.^{(**)}$	$Prod.$
Period 1: 60.1-81.1	0.516	0.558	0.448	0.532	0.335
Period 2: 81.2-02.2	0.314	0.135	0.131	0.138	0.337

Note: (*) and (**) indicates that the hypothesis of equal correlation in the 2 subsamples is rejected at 5% and 1% respectively.

Table 2. Estimation results

	Parameter				
	A_{11}	A_{12}	σ	ρ	λ
Period 1	0.996	-0.004	0.016	0.096	1.005
Period 2	0.998	-0.002	0.012	0.026	0.760

Table 3. Business Cycles Statistics

		International correlations					
		GDP	*Inv.*	*Cons.*	*Emp.*	*Prod.*	
Period 1	Data	0.52	0.56	0.45	0.53	0.34	
60.1-81.1	Model 1	0.52	0.56	0.45	0.58	0.11	
Period 2	Data	0.31	0.13	0.13	0.14	0.33	
81.2-02.2	Model 1	0.31	0.13	0.13	0.52	0.04	
	Model 2	0.39	0.25	0.15	0.62	0.11	
		Percentage Standard Deviations					
		Absolute	Relative to GDP				
		GDP	*Inv.*	*Cons.*	*Emp.*	*Prod.*	*NX*
Period 1	Data	2.25	1.98	0.65	0.46	0.57	0.17
60.1-81.1	Model 1	2.25	2.58	0.49	0.36	0.88	0.01
Period 2	Data	1.71	2.25	0.61	0.51	0.52	0.24
81.2-02.2	Model 1	1.71	2.67	0.55	0.32	0.87	0.34
	Model 2	2.33	2.76	0.53	0.34	0.85	0.32

Figure 1. Correlation of high pass filtered series (US v/s Rest of the World) in two subsamples

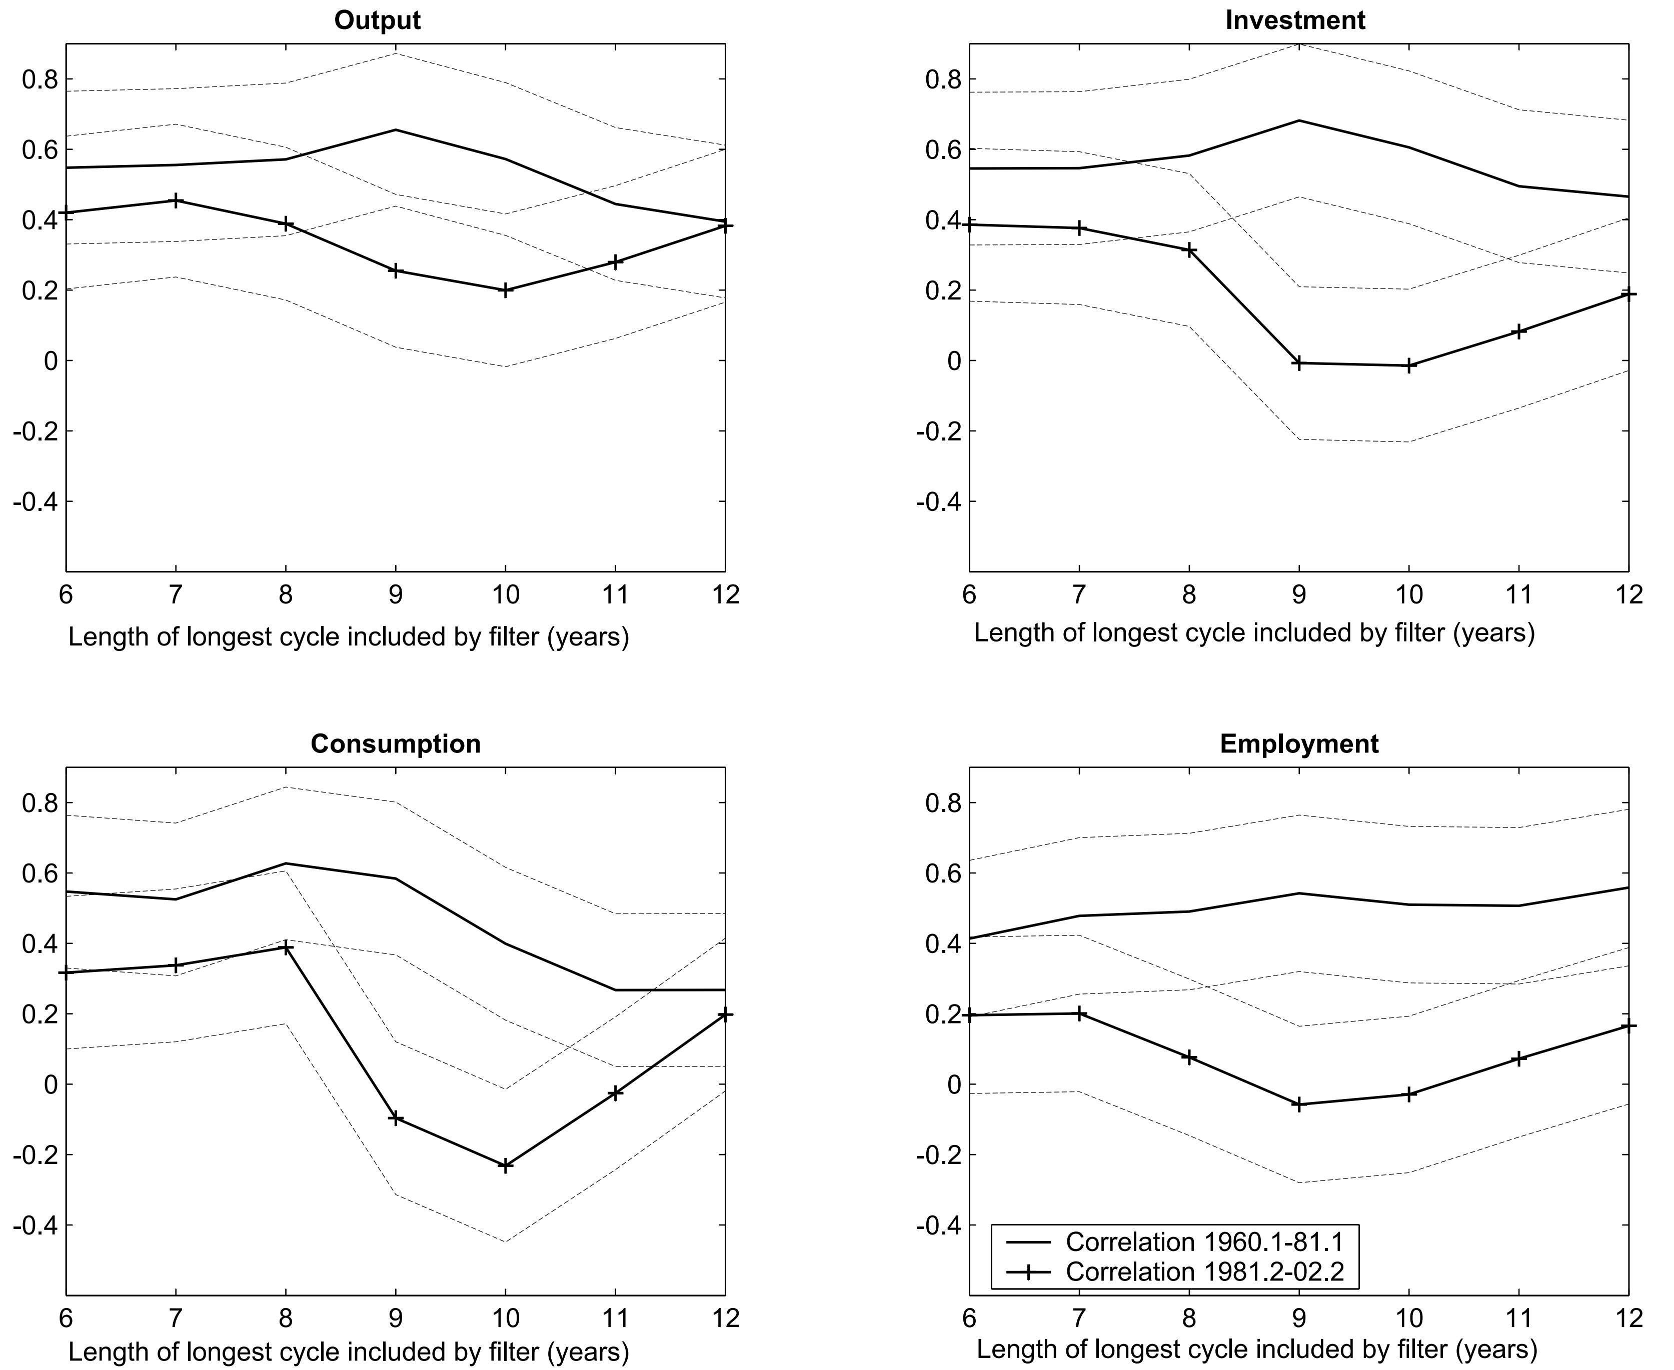

Note: The dashed lines are approximate two standard errors bands

Figure 2. Correlations (US v/s Rest of the world) of HP filtered series over 80 quarters rolling windows