

5 Measuring the dynamics of global business cycle connectedness

Francis X. Diebold and Kamil Yilmaz

5.1 Introduction

The empirical literature on business cycle measurement has been particularly vibrant since the early 1980s, with increased emphasis on international aspects in the 1990s, producing, for example, evidence of commonality in the macroeconomic fluctuations of industrial and developing countries. Indeed that finding is widespread, ranging from classical analyses of pairwise gross domestic product (GDP) correlations (e.g., Backus et al. (1995), Baxter (1995), *inter alia*) to more recent analyses of dynamic latent factor models (e.g., Kose et al. (2003), Canova et al. (2007), Kose et al. (2008), *inter alia*).

As the evidence of commonality in cross-country business cycle fluctuations accumulated, the literature started to focus on the effect of globalization on international business cycles. Kose et al. (2003), for example, find that the impact of the world factor on the cross-country correlation of macroeconomic aggregates (output, consumptions, and investment) increased from the 1990s onward.

Of course the evidence is not uniformly one-sided. Doyle and Faust (2005) find no evidence of increased output growth-rate correlations for the G-7 countries. Stock and Watson (2005) argue that that the comovement of macroeconomic aggregates actually *declined* during the globalization era of 1984–2002, but rather than linking their results directly to the globalization process, they conclude that their results are likely due to the diminished importance of common shocks among the G-7 countries. Eickmeier (2007) emphasizes that the impact of globalization on international propagation of macroeconomic shocks is unclear and needs to be studied further.

Here, we contribute by using a new connectedness measurement technology that is fundamentally grounded in modern network theory (Diebold and Yilmaz (2014)) to measure real output connectedness for a set of six developed countries, 1962–2011. We show that global connectedness is sizable and time-varying over the business cycle, and we study the nature of the time

variation vis-à-vis the changing nature of the global business cycle. We also show that connectedness corresponding to transmissions to others from the United States and Japan is disproportionately important.

In particular, we define and track connectedness of monthly industrial production for the G-6 countries (G-7 less Canada). Our approach is different from earlier studies of international business cycles in that, rather than finding a common world factor or indicator that measures international business cycles, our approach helps one identify how output shocks in one country affect output in other countries, with a lag. Thus we can study directional connectedness, both pairwise and in the aggregate. Time variation in connectedness is potentially of great interest as the intensity of business cycle connectedness is likely to vary over time. We show that business cycle connectedness among G-6 countries is important, that connectedness is indeed time-varying, and that the United States and Japan are the major transmitters of business cycle shocks to other countries.

In addition, our analysis differs from most earlier work in terms of the data used. We use industrial production indexes at monthly frequency rather than the quarterly data from the national income accounts. There are two reasons for this choice. First, the use of monthly data allows us to capture the connectedness of business cycle shocks much faster, as seen in the latest economic crisis. Second, the use of monthly data allows us to have more observations in calculating the connectedness index for each rolling sample window.

The structure of this paper is very simple. We introduce our framework for connectedness measurement in Section 5.2, discussing both population and sample (estimation) issues. Then we apply it to global business cycles in Section 5.3, emphasizing dynamic aspects. We conclude in Section 5.4.

5.2 Measuring connectedness

Here we introduce basic aspects of connectedness measurement as relevant for our subsequent multi-country business-cycle analysis. The connectedness measurement framework was originally developed in Diebold and Yilmaz (2009, 2012, 2014).

5.2.1 THE CONNECTEDNESS TABLE

The variance decomposition indicates the amount of information each variable contributes to the other variables in the autoregression. It determines how much of the forecast error variance of each of the variables can be explained by exogenous shocks to the other variables. Our approach to connectedness is

Table 5.1 Connectedness table schematic

	x_1	x_2	...	x_N	From others
x_1	d_{11}	d_{12}	...	d_{1N}	$\sum_{j=1}^N d_{1j}, j \neq 1$
x_2	d_{21}	d_{22}	...	d_{2N}	$\sum_{j=1}^N d_{2j}, j \neq 2$
\vdots	\vdots	\vdots	\ddots	\vdots	\vdots
x_N	d_{N1}	d_{N2}	...	d_{NN}	$\sum_{j=1}^N d_{Nj}, j \neq N$
To others	$\sum_{i=1}^N d_{i1}$ $i \neq 1$	$\sum_{i=1}^N d_{i2}$ $i \neq 2$...	$\sum_{i=1}^N d_{iN}$ $i \neq N$	$\frac{1}{N} \sum_{i,j=1}^N d_{ij}$ $i \neq j$

based on assessing shares of forecast error variation in various locations due to shocks arising *elsewhere*. This is intimately related to the familiar econometric notion of a variance decomposition: the H -step forecast error variance share d_{ij} is just the fraction of i 's H -step forecast error variance due to shocks in variable j . The full set of variance decompositions produces the core of what we call the “connectedness table.” All of our connectedness measures—from simple pairwise to system-wide—flow from the connectedness table.

The simple Table 5.1, the connectedness table, proves central to understanding the various connectedness measures and their relationships. Its main upper-left $N \times N$ block contains the variance decompositions. For future reference we call that upper-left block a “variance decomposition matrix,” and we denote it by $D = [d_{ij}]$. The connectedness table simply augments D with a rightmost column containing row sums, a bottom row containing column sums, and a bottom-right element containing the grand average, in all cases for $i \neq j$.

To understand and interpret the information conveyed by the connectedness table, it is helpful to cut through the notational clutter via a simple example, as in the example connectedness table, Table 5.2 with $N = 4$. The twelve off-diagonal entries in the upper-left 4×4 D matrix are the twelve pieces of the four forecast-error variance decompositions, d_{ij} . From a connectedness perspective, they measure “pairwise directional connectedness.” The 3,2 entry of 14, for example, means that shocks to x_2 are responsible for 14% of the H -step-ahead forecast error variance in x_3 . We write $C_{3 \leftarrow 2} = 14$. In general the pairwise directional connectedness from j to i is $C_{i \leftarrow j} = d_{ij}$. Note that in general $C_{i \leftarrow j} \neq C_{j \leftarrow i}$. Hence there are $N^2 - N$ separate pairwise directional connectedness measures. They are analogous to bilateral imports and exports for each of a set of N countries. Sometimes we are interested in *net* pairwise directional connectedness, in a fashion analogous to a bilateral trade balance. For example, for x_2 and x_3 we have $C_{23} = C_{3 \leftarrow 2} - C_{2 \leftarrow 3} = 14 - 1 = 13$. In general we have $C_{ij} = C_{j \leftarrow i} - C_{i \leftarrow j}$. There are $\frac{N^2 - N}{2}$ net pairwise directional connectedness measures.

Table 5.2 Connectedness table example

	x_1	x_2	x_3	x_4	From others
x_1	96	1	2	1	4
x_2	28	67	1	3	32
x_3	14	14	70	1	29
x_4	18	11	5	65	34
To others	60	26	8	5	25

In our subsequent analysis, however, we will be interested in more aggregative connectedness measures. The eight off-diagonal row and column sums, labeled “from” and “to” in the connectedness table, are the eight *total directional connectedness* (“directional connectedness” for short) measures. The value of 32 in the second entry of the rightmost column, for example, means that x_2 receives 32% of its variation from others (x_1 , x_3 , and x_4). We write $C_{2\leftarrow\bullet} = 28 + 1 + 3 = 32$. In general, the total directional connectedness from others to i is

$$C_{i\leftarrow\bullet} = \sum_{\substack{j=1 \\ j \neq i}}^N d_{ij},$$

and the total directional connectedness from j to others is

$$C_{\bullet\leftarrow j} = \sum_{\substack{i=1 \\ i \neq j}}^N d_{ij}.$$

Hence there are $2N$ total directional connectedness measures, N “to others” and N “from others,” analogous to “total exports” and “total imports” for each of a set of N countries. Just as with pairwise directional connectedness, sometimes we are interested in net total directional effects. For x_2 , for example, we have $C_2 = C_{\bullet\leftarrow 2} - C_{2\leftarrow\bullet} = 26 - 32 = -6$. In general, net total directional connectedness is $C_i = C_{\bullet\leftarrow i} - C_{i\leftarrow\bullet}$. There are N net total directional connectedness measures, analogous to the total trade balances of each of a set of N countries.

Finally, we consider the most aggregative connectedness measure, the grand total of the off-diagonal entries in D (equivalently, the sum of the “from” column or “to” row), which measures “total (system-wide) connectedness”. We typically express this total cross-variable variance contribution, given in the lower right cell of the connectedness table, as a percentage of total variation. Hence total connectedness in our example is $C = \frac{99}{400} \times 100 = \frac{99}{4} = 24.8$. Note that total variation is just one hundred times N , the number of variables in the system, because each of the N rows sums to one hundred. Conversion to percent eliminates the one hundred so that ultimately total connectedness

is simply average total directional connectedness (whether “from” or “to”). In general we have

$$C = \frac{1}{N} \sum_{\substack{i,j=1 \\ i \neq j}}^N d_{ij}.$$

There is just a single total connectedness measure, as total connectedness distills a system into a single number analogous to total world exports or total world imports (the two are of course identical).

The connectedness table makes clear how one can begin with the most disaggregated (e.g. microeconomic, firm-level pairwise directional) connectedness measures and aggregate them in various ways to obtain macroeconomic economy-wide total directional and total connectedness. Different agents may be disproportionately interested in one or another of the measures. For example, firm i may be maximally interested in total directional connectedness from others to i , $C_{i \leftarrow \bullet}$, whereas regulators might be more concerned with total directional connectedness from i to others, $C_{\bullet \leftarrow i}$, or in total connectedness C .

5.2.2 IDENTIFYING SHOCKS

As already emphasized, our approach is based on variance decompositions. An H -step forecast error variance decomposition, d_{ij} , answers an interesting and important question: what fraction of H -step forecast error variance of variable i is due to shocks in *another* variable j ? This must, of course, be done within the context of a model. If shocks in reduced-form vector autoregressions were orthogonal, variance decomposition calculations would be trivial. That is, variance decompositions are easily calculated in orthogonal VARs, because orthogonality ensures that the variance of a weighted sum is simply an appropriately-weighted sum of variances.

The much more realistic case involves correlated VAR shocks. Consider therefore a data-generating process with correlated shocks, with moving-average representation $x_t = \Theta(L)\varepsilon_t$, $E(\varepsilon_t\varepsilon_t') = \Sigma$. The variance decomposition calculations are more involved, because we first need to isolate the independent shocks that underlie the observed system. One way or another, we must transform the shocks to orthogonality to calculate variance decompositions. The orthogonalization can be handled in several ways, and we now sketch two.

Cholesky factorized variance decompositions

This time-honored orthogonalization method traces at least to Sims (1980). The correlated-shocks model above is mathematically identical to

the orthogonal-shocks model, $x_t = A(L)u_t$, $E(u_t u_t') = I$, $A(L) = \Theta(L)Q$, $u_t = Q^{-1}\varepsilon_t$, where the lower triangular matrix Q is the Cholesky factor of Σ ; that is, $QQ' = \Sigma$. Hence a simple Cholesky factor transformation orthogonalizes the system.

Variance decompositions based on Cholesky factorization may be sensitive to ordering, as has been well-appreciated at least since Sims (1980). Interestingly, we often find that total connectedness is robust to ordering (that is, the range of total connectedness estimates across orderings is often quite small), whereas directional connectedness often appears more sensitive to ordering. Hence so-called “generalized variance decompositions,” which we now introduce, may be more useful when studying directional connectedness.

Generalized variance decompositions

The generalized variance decomposition (GVD) framework of Koop et al. (1996) and Pesaran and Shin (1998) produces variance decompositions invariant to ordering. The GVD approach does not require orthogonalized shocks; rather, it allows and accounts for correlated shocks using the historically observed error distribution, under a normality assumption.

The GVD matrix $\Delta = [\delta_{ij}]$ has entries

$$\delta_{ij} = \frac{\sigma_{jj}^{-1} \sum_{h=0}^{H-1} (e_i' A_h \Sigma e_j)^2}{\sum_{h=0}^{H-1} (e_i' A_h \Sigma A_h' e_i)},$$

where σ_{jj} is the standard deviation of ε_j , e_j is a selection vector with j^{th} element unity and zeros elsewhere, A_h is the coefficient matrix multiplying the h -lagged shock vector in the infinite moving-average representation of the non-orthogonalized VAR, and Σ is the covariance matrix of the shock vector in the non-orthogonalized VAR.

Because shocks are not necessarily orthogonal in the GVD environment, sums of forecast error variance contributions (that is, row sums in GVD matrices) are not necessarily unity ($\sum_{j=1}^N \delta_{ij} \neq 1$). Hence we base our generalized connectedness indexes not on Δ , but rather on $\tilde{\Delta} = [\tilde{\delta}_{ij}]$, where $\tilde{\delta}_{ij} = \frac{\delta_{ij}}{\sum_{j=1}^N \delta_{ij}}$.

Note that, by construction, $\sum_{j=1}^N \tilde{\delta}_{ij} = 1$ and $\sum_{i,j=1}^N \tilde{\delta}_{ij} = N$. Armed with $\tilde{\Delta}$, we can immediately calculate generalized connectedness measures \tilde{C} , $\tilde{C}_{\bullet \leftarrow j}$, $\tilde{C}_{i \leftarrow \bullet}$, \tilde{C}_i , $\tilde{C}_{i \leftarrow j}$, $\tilde{C}_{j \leftarrow i}$, and \tilde{C}_{ij} .

In essence the GVD does not impose orthogonality of shocks. Therefore, in the GVD framework all variables in a system are subject to shocks simultaneously. This in return amounts to obtaining impulse responses and variance decompositions for each variable treating each variable as the leading variable in the VAR. GVDs moreover require normality, and hence may be more useful

for assessing connectedness of log volatilities, which are well-approximated as Gaussian, than for returns, which are not.

5.2.3 SAMPLE CONNECTEDNESS

Clearly C depends on x , H , and $A(L)$, so we have written $C(x, H, A(L))$. In reality A is unknown and must be approximated (e.g., using a finite-order vector autoregression). Recognizing the centrality of the approximating model adopted, we write $C(x, H, A(L), M(\theta))$, where θ is a finite-dimensional parameter.

However, everything we have written thus far is in *population*, whereas in reality we must use an approximating model estimated using data $1 : T$, so we write $\widehat{C}_t(x, H, A_t(L), \widehat{M}_{1:T}(\hat{\theta}_t))$. To economize on notation we henceforth drop $A(L)$, because it is determined by nature rather than a choice made by the econometrician, relying on the reader to remember its relevance and writing $\widehat{C}_t(x, H, \widehat{M}_{1:T}(\hat{\theta}_t))$.

Hence we now consider estimation of connectedness. In addition, we want to allow for time-variation of connectedness, effectively allowing the connectedness table and all of its elements to vary over time, so we write $C_t(x, H, A_t(L), M(\theta_t))$.

Estimation

Choice of x —the object of interest to be studied—has important implications for the appropriate approximating model; for example, x may (or may not) be strongly serially correlated, conditionally heteroskedastic, or highly disaggregated. Here we study real output growth across countries, for which serial correlation, in particular, is surely important.

Connectedness measurements are defined only with respect to a reference universe (specific set of x 's). In general they will not—and should not—be robust to choice of reference universe. Hence, given a decision as to the type of x to be examined, a second important issue is precisely which (and hence how many) x 's to use. For example, in cross-country analyses we may want to use sufficiently many countries to ensure that we have good global coverage. Whether this requires a small or large number of x 's depends on the distribution of activity across countries.

The predictive horizon, H , is also important, and one must take a stand on it as well. The horizon is important particularly because it is related to issues of dynamic connectedness (in the fashion of contagion) as opposed to purely contemporaneous connectedness. To take a simple pairwise example, shocks to j may impact the forecast error variance of i only with a lag, so that $C_{i \leftarrow j}$ may be tiny for small H but nevertheless large for larger H . Intuitively, as the

horizon lengthens there may be more chance for things to become connected. Such dynamic phenomena, and the rich patterns that are possible, are closely related to aspects of multi-step Granger causality, as treated, for example, in Dufour and Renault (1998), Dufour and Taamouti (2010), and the references therein.

In a sense, varying H lets us break down connectedness into “long-run,” “short-run,” etc. More precisely, as H lengthens we obtain a corresponding sequence of conditional prediction error variance decompositions for which the conditioning information is becoming progressively less valuable. In the limit as $H \rightarrow \infty$, we obtain an *unconditional* variance decomposition.

The remaining issue—choice of empirical approximating model $\widehat{M}_{1:T}(\hat{\theta}_t)$ —is the obvious elephant in the room. For real activity, conditional mean dynamics in $M(\theta)$ will surely be important and will need to be modeled. We will use the obvious workhorse model, a vector autoregression, $VAR(p)$. We emphasize, however, that almost any model could be used, including much more deeply structural models, so long as underlying orthogonal shocks can be identified.

Time-variation

Connectedness is just a transformation of system coefficients. Hence if the coefficients are time-varying, so too will be connectedness. Tracking (“now-casting”) real-time connectedness movement is of central interest. One can capture parameter variation by using a rolling estimation window; we write $\hat{C}(x, H, \widehat{M}(\theta; w))$, where w denotes window width. We will take this approach, estimating the model repeatedly, at each time using only the most recent w observations.

There are of course both benefits and costs of a rolling-window approach. Benefits include both tremendous simplicity and coherence with a wide variety of possible data-generating processes (DGPs) involving time-varying parameters. Costs include that, given a particular DGP, rolling is generally sub-optimal, and it requires selection of w . Too large a w produces “oversmoothing,” and too small a w produces “undersmoothing,” in a manner precisely analogous to bandwidth choice in density estimation.

5.3 Global business cycle connectedness

To study business cycle connectedness we use monthly seasonally adjusted industrial production (IP) for the G-7 countries less Canada (G-6), from January 1958 to December 2011, as shown in Figure 5.1. We exclude Canada because its IP growth is very highly correlated with that of the United States.

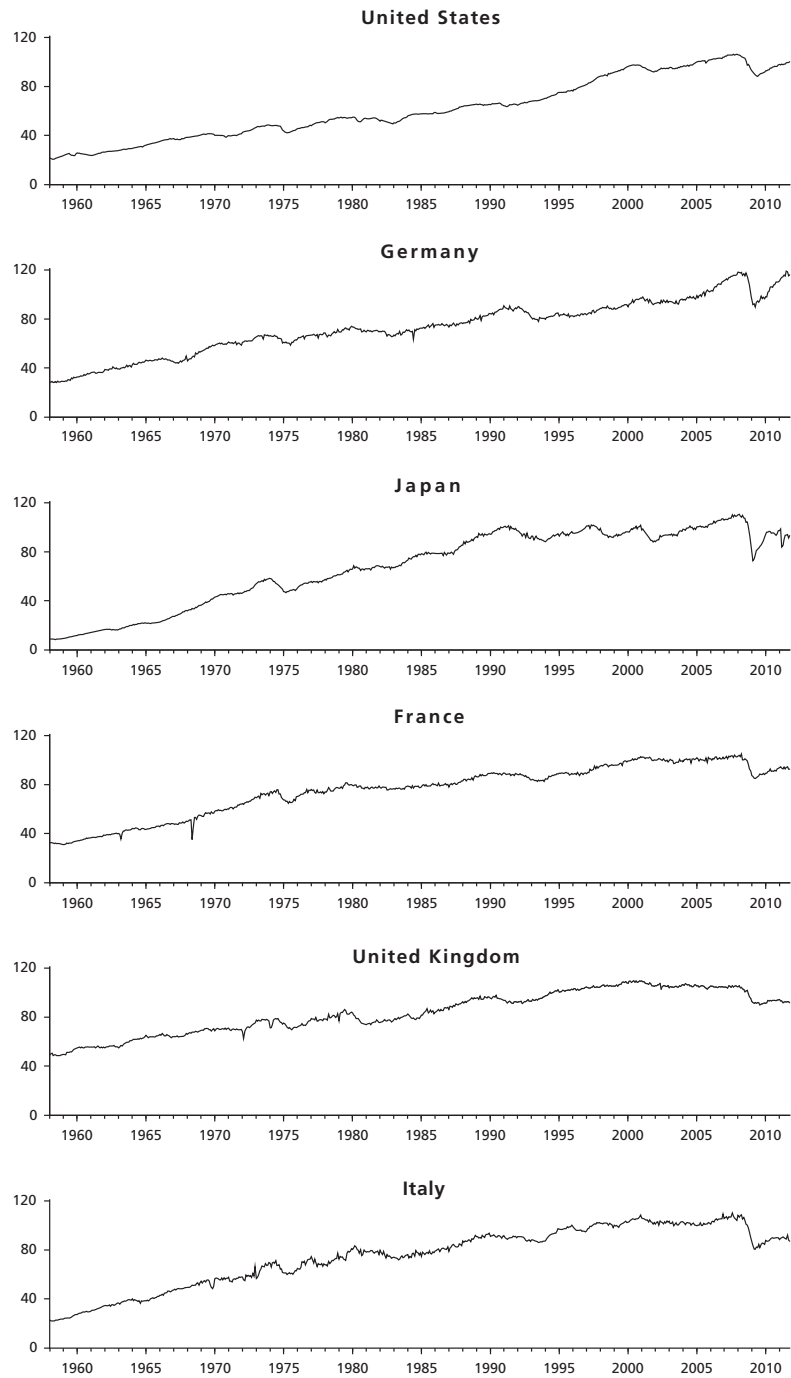


Figure 5.1. G-6 industrial production, January 1958 to December 2011

Indeed year-on-year industrial production growth rates for the U.S. and Canada have a correlation of 0.87, which is much higher than the correlations for other G-6 country pairs.

We begin in Section 5.3.1 by assessing the cointegration status of the IP data, which has implications for the structure of our subsequent analysis in Sections 5.3.2, 5.3.3, and 5.3.4. In Section 5.3.2 we examine static connectedness, in Section 5.3.3 we examine dynamic total connectedness, and in Section 5.3.4 we examine dynamic directional connectedness.

5.3.1 COINTEGRATING RELATIONSHIPS

Before proceeding to examine business cycle connectedness, we must ascertain the likely integration and cointegration status of our G-6 (log) IP data. We first tested for unit roots in each log IP series, using augmented Dickey-Fuller (ADF) tests with augmentation lag length selected using Akaike information criterion (AIC), and allowing for linear trend under the alternative. To conserve space we do not show the unit root test results. There was no evidence against the unit root in any log IP series, and overwhelming evidence against the unit root in every differenced log IP series (allowing for nonzero mean under the alternative).

We now assess cointegration status. We show the results of Johansen's maximum eigenvalue and trace tests in Table 5.3. The hypothesis of 0 cointegrating relationships is clearly rejected, the hypothesis of at most 1 cointegrating relationship is not rejected at the 5% level (but it is rejected by the trace test at the 10% level), and there is no evidence against the hypothesis of at most R cointegrating relationships for $R > 1$. All told, then, it appears that there is only one cointegrating relationship among the G-6 IP series. Hence we proceed to adopt a vector error-correction specification for our approximating model, although we note that it is not far from a simple VAR in first differences (i.e., 0 cointegrating relationships).

Table 5.3 Tests of the number of cointegrating relationships

CI rank	Max eval	(P-value)	Trace	(P-value)
0	48.9	(0.004)	115.6	(0.0011)
At most 1	29.1	(0.168)	66.7	(0.0865)
At most 2	17.4	(0.550)	37.6	(0.3191)
At most 3	14.7	(0.310)	20.3	(0.4054)
At most 4	5.4	(0.687)	5.6	(0.7473)
At most 5	0.12	(0.723)	0.12	(0.7230)

Note: We assume linear trends in industrial production series and intercepts in any cointegrating relationships.

5.3.2 STATIC CONNECTEDNESS

In the empirical analysis of business cycle connectedness we first estimate the vector error correction (VEC) model for the full sample and report the connectedness index and the directional connectedness in Table 5.4 along with the underlying generalized variance decomposition. The connectedness index for the full sample period is 29.1%, indicating that less than one-third of the total variance of the forecast errors for the G-6 countries is explained by the connectedness of shocks across countries, whereas the remaining 70.9% is explained by idiosyncratic shocks.

It is important at this stage to note that the connectedness index for the whole sample is sensitive to the inclusion of new observations in the sample. The connectedness index for the period from January 1958 to December 2008 is only 27%. When the sample is extended to May 2009, the connectedness index for the full sample jumps to 69%. Finally, the inclusion of observations from June 2009 to December 2011 lowers the index to 29.1%.

In terms of the directional connectedness to others (measured by $\tilde{C}_{\bullet \leftarrow i}^H$) throughout the full sample, Japan is the country that contributed the most to other countries' forecast error variance (58.1 points, which is close to 10% of the total forecast error variance to be explained), followed by France (35.4 points). According to the full sample directional connectedness measures, the U.S., Germany, and the U.K. contributed at similar rates (29.2, 23.9, and 17.1 points, respectively), followed by Italy (10.7 points).

Table 5.4 Static connectedness, G-6 industrial production, January 1958 to December 2008

	U.S.	Germany	Japan	France	U.K.	Italy	From others
U.S.	89.2	0.8	3.5	3.3	2.2	1.1	10.8
Germany	5.6	55.7	23.9	8.4	6.0	0.5	44.3
Japan	8.1	5.6	77.4	6.0	1.9	1.0	22.6
France	3.5	10.1	11.0	64.2	4.2	7.0	35.8
U.K.	7.7	3.3	4.2	3.8	80.0	1.0	20.0
Italy	5.9	2.0	14.3	13.2	4.0	60.6	39.4
To others	30.9	21.7	56.8	34.7	18.3	10.5	
Net =							Index =
To – from	20.0	-22.5	34.1	-1.1	-1.7	-28.8	28.8%

Note: Each cell in the upper-left 6×6 matrix reports the relative (percent) contribution of the "column" country to the variance of the forecast error for the "row" country. The "directional from others" column reports the total forecast error variance shares of the row countries attributable to shocks from other countries. The "directional to others" row reports the total contributions of each column country to all other countries' forecast-error variance. Each cell in the "net directional connectedness" (to – from) row reports the difference between the corresponding cells in the "directional to others" row and the ones in the "directional from others" column. The total connectedness index in the lower right is the average of the elements of the "directional from others" column (or equivalently, the "directional to others" row), multiplied by 100%.

In terms of the directional connectedness received from others (measured by $\tilde{C}_{i\leftarrow\bullet}^H$), the U.S. appears to be the country that received the lowest percentage of shocks from other countries (10.7 points, equivalent to just 1.8% of the total forecast error variance to be explained), followed by the U.K. (22.4 points) and Japan (22.6 points). Germany received the highest percentage (43.5 points) of shocks from other countries, followed by Italy (39.4 points) and France (35.7 points).

Finally, we calculate the difference between the column-wise sum (the “contribution from others”) and the row-wise sum (“contribution to others”) to obtain the “net directional connectedness” given by \tilde{C}_i^H . Japan (35.4 points) and the U.S. (18.4 points) are net transmitters of industrial production shocks to other countries, while the U.K. (−5.3 points) and France (−0.3 points) received very low percentage of business cycle shocks in net terms. Italy (−28.7 points) and Germany (−19.6 points), on the other hand, are definitely the leading net recipients of business cycle shocks over the full sample.

5.3.3 DYNAMIC TOTAL CONNECTEDNESS

The connectedness table for the full sample provides important clues as to how the connectedness index is calculated and interpreted. However, we want to focus more on the dynamics of business cycle connectedness over time. The fact that the inclusion of new observations in the sample leads to significant jumps in the connectedness index definitely highlights the need to study the dynamics of connectedness over time.

As the VEC is the correct model for the full sample, the dynamic analysis of connectedness also relies on the variance decomposition from the VEC model estimated over rolling five-year windows. Here is how the connectedness plot is obtained: we estimate the VEC model for the first five-year sub-sample window (April 1958–April 1963) and obtain the value for the generalized variance decomposition-based connectedness index (from now on, the connectedness index). Moving the sub-sample window one month ahead, we estimate the VEC model again and calculate the connectedness index for the new sub-sample, and so on. Graphing the connectedness index values for all sub-sample windows leads to the connectedness plot.

So far we have discussed the connectedness plot based on the underlying VEC(1) model, estimated over five-year rolling windows with a twelve-month forecast horizon. Next we want to discuss the appropriateness of our assumptions and the robustness of our results to these assumptions. Let us start with the underlying VEC(1) model. Earlier we reported unit root and cointegration test results for the full sample from January 1958 to December 2011. While the test results indicated that the correct underlying model is the VEC(1) for the full sample, this does not necessarily imply that VEC(1) is the

correct model for each five-year rolling window. For that reason, we repeated the unit root and cointegration tests for all five-year rolling windows considered. The ADF test fails to reject the presence of unit roots in the log IP series for an overwhelming majority of the windows considered for all countries. In the case of first-differenced log IP series, the ADF test rejects the presence of unit roots for almost all countries. The only exception is Japan. In the case of Japan, for a non-negligible number of rolling windows the ADF test fails to reject the presence of unit roots in the first-differenced log IP series. While this is a cause for concern, we do not test for unit roots in further differenced IP series. Instead, we proceed with the Johansen cointegration test over the rolling windows. In the case of the U.K., the ADF test rejects the unit root in the first-differenced log IP series so strongly for all windows that the p-value is very close to zero.

Similarly, both the trace and maximum test statistics reject the null of no cointegration relationship among the six log IP series at the 5% level for an overwhelming majority of rolling windows considered. This means that the test prefers a VEC model to a VAR in first differences. In contrast, the trace statistic, in particular, fails to reject the null of at most one cointegration equation linking all six log IP series. Therefore, the Johansen cointegration tests indicate that there is either one or two cointegration equations among the six IP series. Based on these results, we expand the connectedness index analysis to rolling windows based on the VEC(1) model. Later on, we will show the differences in the connectedness indexes for different models.

The dynamic connectedness index based on the VEC(1) model is plotted in Figure 5.2. We also calculate an alternative connectedness index based on the Cholesky variance decomposition. Even though we do not plot it here, we can report that the two indexes move in tandem, with the difference between the two indexes seldom exceeding ten percentage points. Therefore, it would be sufficient to focus on the generalized VD-based connectedness index for the rest of the paper.

Turning to Figure 5.2, the first thing one observes about the connectedness plot is the absence of a long-run trend. The connectedness plot clearly shows that while there are periods during which shocks to industrial production are substantially transmitted to others, there are yet other periods during which the connectedness of business cycle shocks was much less important. Actually, during or after all U.S. recessions (indicated by the shaded bars in Figure 5.2), the connectedness index recorded significant upward movements. The only exception is the 1969–70 recession, during which the index moved down. In addition, the index went up in late 1993, and after a brief correction in late 1994, it went up again in 1995. While there was no U.S. recession during this period, France, Germany, Italy, and Japan experienced recessions ending in late 1993 or early 1994 (see the Economic Cycles Research Institute's website <<http://www.businesscycle.com/resources/cycles/>>). As a result, the upward




Figure 5.2. Dynamic total connectedness, G-6 industrial production, January 1958 to December 2011

Note: We use a 5-year rolling window. The shaded areas indicate U.S. recessions. See Section 5.3.3 for details.

movement in the connectedness index is most likely due to the connectedness originating from these countries.

Second, while the connectedness index fluctuates over time, it is possible to differentiate between several trends. First, during the 1973–75 recession the connectedness index increases by almost 20–25 percentage points and fluctuates around 50% after the 1981–82 recession. Starting in 1984, the connectedness index declines all the way to 33%. This result is consistent with the findings of McConnell and Perez-Quiros (2000) and Blanchard and Simon (2001) that the volatility of the U.S. GDP declined after 1984 (the great moderation). As the volatility of GDP declines, the connectedness index declines to pre-1973 levels.

Third, after the great moderation of the late 1980s, the behavior of the connectedness index reflects the influence of globalization. From 1989 onward, the band within which the connectedness index fluctuates starts to move upward with the current wave of globalization that started in earnest in the early 1990s. As the sample windows are rolled to include 1996, the index reaches 60%, but declines to 40% as the data for the late 1990s and 2000 are included. The index starts to increase again toward the end of the mild recession of 2000–01, reaching 60% by the end of 2002. However, as the other G-6 countries followed the quickly recovering U.S. economy to a major expansion, the connectedness index reached 65% in 2004Q2. The index then declines to 60% again as the window is rolled to include 2004H2, and then gradually moves down, reaching its bottom around 40% from 2006Q4 until 2008Q1.

During the era of globalization, from the late 1980s to 2007, the connectedness index followed three distinct cycles. Each cycle lasted longer and had a

larger bandwidth than the previous one. During the first cycle, which lasted from 1989 to the end of 1992, the index fluctuated between 33% and 53%, while in the second cycle, which lasted from 1993 to 1999, the index fluctuated between 37% and 60%. Finally, during the third cycle from 2001 to 2007, the index fluctuated between 44% and 65%.

This result is consistent with Kose et al.'s (2003) finding that with the globalization process, business cycles have become more synchronized. It basically indicates that the comovement of industrial production fluctuations has tended to be more significant since the late 1980s. In other words, when there is a shock to industrial production in one or more countries in the G-6 group, its tendency to be transmitted across other countries increases as one move from 1989 toward 2007. This result can also be interpreted as consistent with Doyle and Faust's (2005) conclusion that the correlation coefficients among the industrial production series have not increased much since the late 1980s. The output fluctuations tend to move together during periods of high connectedness indexes, compared to the periods with low connectedness indexes. When one analyzes the period since the late 1980s as a whole, one may not obtain high correlation coefficients. Actually, for the period from 1989 to 2007 the connectedness index is only 36%.

Next, we focus on the behavior of the connectedness index since June 2008 (see Figure 5.3). We want to focus on its most recent behavior, not only because it provides us with more clues about the business cycle connectedness since the beginning of the sub-prime crisis in the U.S., but also because, in 2008 and 2009, the index recorded the biggest jump in its history. The index increased sharply from 41% in May 2008 to 53% in July, to 68% in September, and then to 80% in December. With the inclusion of January 2009 through October 2009 in




Figure 5.3. Dynamic total connectedness, G-6 industrial production, January 2000 to December 2011

Note: We use a 5-year rolling window. The shaded areas indicate U.S. recessions. See Section 5.3.3 for details.


the analysis, the index declined slightly to 71%. As the economic recovery went underway in the G-6 countries in 2009H2, the index moved slightly upward again, reaching 75% by December 2009.

The behavior of the index during the Great Recession of 2007–09 is in stark contrast to its behavior in previous recessions. It increased thirty seven points from April to December 2008. The jump in the index during the Great Recession is an indication of how the G-6 countries were pulling each other down. To give an example, during the recession following the first oil price hikes, in a matter of three and a half years from 1972 to 1976, the connectedness index recorded a relatively smaller increase, from a low of thirty two in August 1972 to a high of sixty four in April 1974.

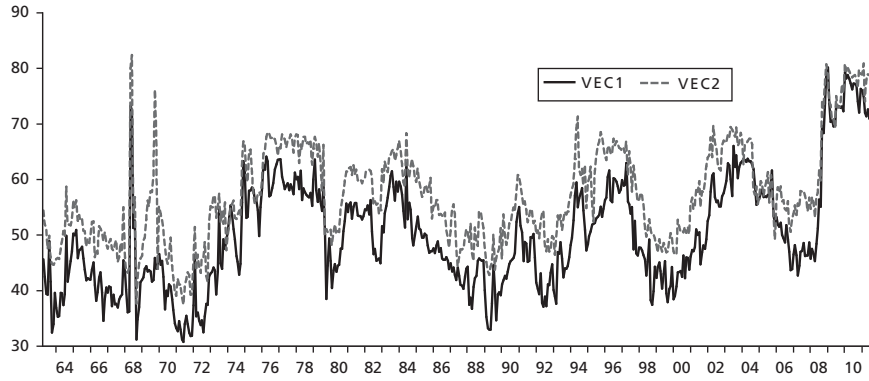
So far we have only discussed the total connectedness index. However, as we argued in the introduction, the analysis of directional connectedness provides us with quite interesting results to discuss in some detail. However, before going ahead with the analysis of directional connectedness plots, we want to make sure that the results we obtained from the total connectedness analysis are not due to some special characteristics of the VAR framework we use. For that reason, we now report the robustness of the total connectedness index with respect to the model choice, the window width, the forecasting horizon, and the ordering of variables.

Let us start the robustness analysis with the choice of the underlying model. Along with the single cointegration equation (VEC1) model, we calculated the total connectedness index under a VAR in first differences (DVAR), a VEC2, and a VEC5 model. We present the dynamic connectedness plots in Figure 5.4. The dynamic connectedness plot obtained from the underlying DVAR model differs substantially from the one obtained from the VEC1 model. Given the fact that the null hypothesis of no cointegration equation was very strongly rejected by Johansen's cointegration test using both trace and maximum eigenvalue statistics, it is not very surprising that the dynamic connectedness index obtained from the DVAR model is quite different from the one obtained with the VEC1 model.


Cointegration test results, in general, preferred the VEC1 model to the VEC2 and VEC5 models, but there were many instances, especially in the case of maximum eigenvalue statistics, where the null of at most one cointegration equation was rejected in favor of two or more cointegration equations. Given that the Johansen test results did not indicate an overwhelming preference in favor of the VEC1 model, we decided to compare the behavior of the total connectedness index from the underlying VEC1 model with the ones obtained from the VEC2 and VEC5 models. As can be seen, there is a level difference between the connectedness index obtained from the VEC1 model and the ones obtained from the VEC2 and VEC5 models. As the level difference is not changing substantially over time, there is little difference between the time



(a) One cointegrating vector vs. VAR in growth rates



(b) One vs. two cointegrating vectors



(c) One cointegrating vector vs. five cointegrating vectors

Figure 5.4. Robustness of estimated dynamic total connectedness to VECM specification

Note: In each sub-panel, a solid line denotes our VECM with one cointegrating vector, and a dashed line denotes an alternative model.

behavior of the VEC1-based connectedness index and the ones obtained from the VEC2 and VEC5 models. Based on these results, we decided to use the VEC1 as the main underlying model.

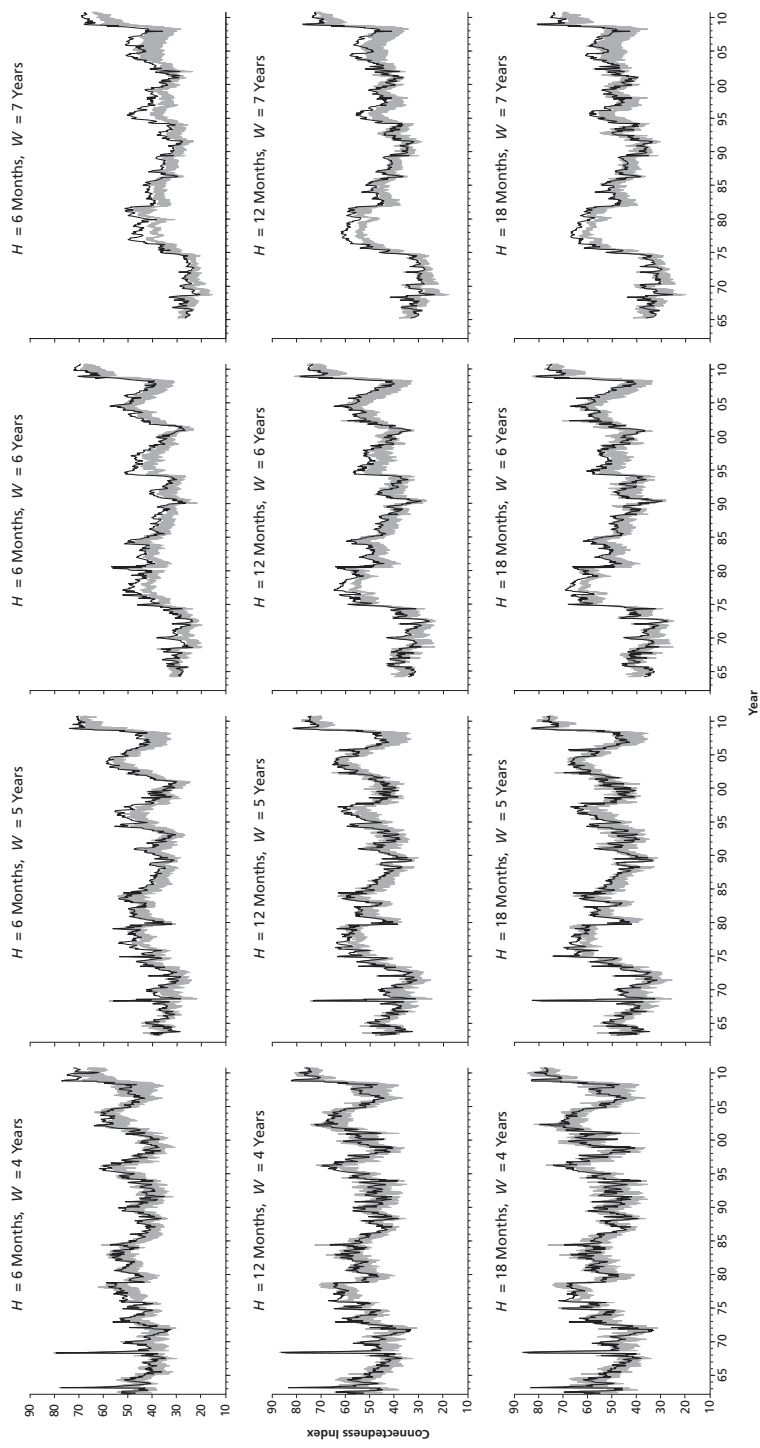
Next, we present the robustness checks with respect to the window width and the forecast horizon in Figure 5.5. In this robustness exercise we consider four, six, and seven years as the alternatives to our benchmark window width of five years. In the case of the forecast horizon, we consider six, and eighteen months in addition to our benchmark forecast horizon of twelve months. In total we plot the dynamic connectedness index in twelve sub-graphs. In addition to the connectedness index, which is based on generalized variance decomposition, we plot the tenth and ninetieth percentile values of the Cholesky-based connectedness index out of one hundred random orderings.

Irrespective of the forecast horizon and window width considered, the generalized and Cholesky variance decomposition-based connectedness indexes follow very similar patterns. This comparison assures us that the use of the generalized variance decomposition-based connectedness index leads to quite sensible results. Figure 5.5 also assures us that the result we obtained for the benchmark values of the window width and the forecast horizon carries over when we use other values for these two important parameters of our connectedness index methodology.

5.3.4 DYNAMIC DIRECTIONAL CONNECTEDNESS

Following the analysis of the total connectedness index, we can now focus on the directional connectedness of business cycles across countries. Directional connectedness indexes are critical in understanding the respective roles of each of the G-6 countries in spreading shocks to their local industrial outputs to other countries. In Figure 5.6 we present all three indexes of connectedness: “connectedness to others,” “connectedness from others,” and “net connectedness to others,” which we will discuss in some detail.

Throughout the 1970s, Japan was the most important source country of net connectedness (Figure 5.6), followed by France and Germany. During the second half of the 1970s, the gross connectedness from Japan to others reached as high as 180%, whereas the connectedness received by Japan from others was only around 40–50%, leading the net connectedness from Japan to reach as high as 150%. While Germany had high directional connectedness to others in the late 1960s, early 1970s, and late 1970s, its net connectedness was negative in the mid-1970s, immediately after the first oil price shock. France, on the other hand, had significant directional connectedness to others after the first oil price hikes in 1973–74 as well as the second half of the 1970s. The U.S., on the other hand, had negative directional connectedness to others, and, therefore, was a net recipient of business cycle shocks over most of the 1970s.



H is forecast horizon. *W* is estimation window width.
 [line] Koop-Pesaran-Potter-Shin generalized identification; (grey box) Cholesky factor identification, (10%, 90%) band based on one hundred randomly selected orderings.

Figure 5.5. Robustness of estimated dynamic total connectedness to window width, forecast horizon, and VECM identification

Note: We explore window widths of four, five, six, and seven years (across columns), and horizons of six, twelve, and eighteen months (across rows). Each sub-figure contains a solid line based on generalized identification, and a 90% region based on one hundred randomly selected Cholesky factor identifications. See Section 5.3.3 for details.

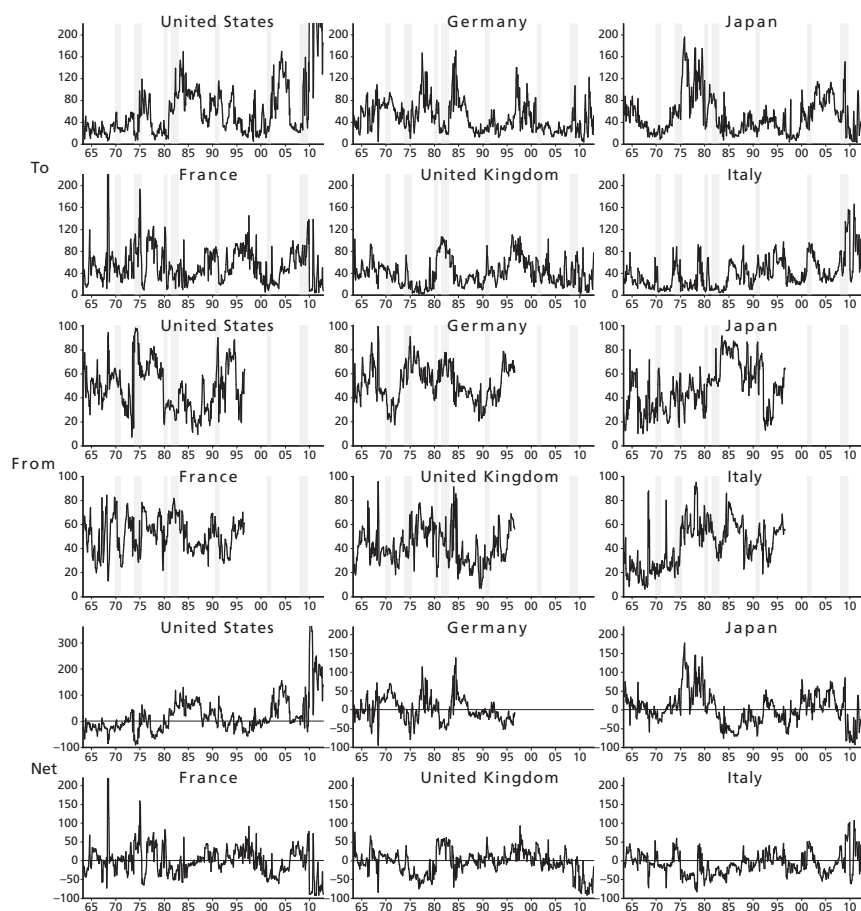


Figure 5.6. Dynamic directional connectedness, “to,” “from,” and “net”

Note: We show transmissions to others in rows one and two, receipts from others in rows three and four, and net flows (to–from) in rows five and six. The shaded areas indicate U.S. recessions.

The roles were reversed in the 1980s: the U.S. became the major net transmitter of shocks, while Japan became the net recipient of the business cycle shocks. Following the 1981–82 U.S. recession, gross connectedness transmitted by the U.S. to others jumped above 160%, and net connectedness from the U.S. fluctuated between 50 and 120%. Japan’s net connectedness, on the other hand, declined to as low as –80% after the 1982 recession and stayed at low levels until the end of 1987. Germany and the U.K. also had positive net directional connectedness to others after the 1981–82 recession, but their roles were rather secondary compared to that of the U.S.

Throughout the 1990s, Japan’s net directional connectedness was positive, but it was rather low. This fact is consistent with the decade-long recession

Japan suffered while the other G-6 countries continued to attain higher growth rates. Neither the U.S. nor Germany was one of the countries that had net directional connectedness in the 1990s. Rather, France and the U.K. had sizable net directional connectedness in the 1990s, even though their net connectedness was not as significant and as persistent as the ones the U.S., Japan, and Germany attained in the 1970s and 1980s. The role these countries played during the 1990s is closely related to the aftermath of the ERM crisis of 1992 and the ensuing slowdown in these economies.

Moving forward in time, the U.S. and Japan returned to their locomotive roles in the 2000s. In particular, with a high net connectedness to others, the U.S. was a net transmitter of business cycle shocks after the 2001 recession. In response to the slowdown in the economy in early 2001, the Federal Reserve lowered the fed funds rate from 5.5% in January 2001 all the way down to 2% in November 2001. This aggressive policy stance was effective in stimulating domestic demand. As a result, industrial production grew at a monthly rate of between 0.5 and 1.0% in 2002H1. The 2001 recession lasted for only eight months, from March to November 2001. With this rapid turnaround, the U.S. started to generate substantial connectedness in 2002H1, with the net connectedness from the U.S. reaching 100%. After a brief lull in 2003, the net connectedness from the U.S. increased again to surpass 130%. Being the driver of worldwide demand, the U.S. had an impact on other countries until the end of 2006, as the net connectedness of the U.S. declined to almost zero. Japan also had a positive net connectedness in the first seven years of the 2000s; its net connectedness fluctuated around 50%. Germany's net directional connectedness was negative throughout the 2000s and during the global recession of 2008–09. France, Italy, and the U.K. were also net recipients of business cycle shocks before the global recession.

Lately, with a net connectedness measure lower than –50% since 2007, Japan has become a net recipient rather than a net transmitter of business cycle connectedness. In the meantime, the net connectedness from the U.S. has gradually increased with the intensification of the sub-prime crisis since mid-2007. As emphasized above, from April to December 2008, the total connectedness index jumped up substantially to reach 80%. The U.S. was the most important contributor to the increase in business cycle connectedness, with a net connectedness contribution of more than 150%. The gross directional connectedness from the U.S. jumped close to 180% following the collapse of Lehman Brothers in September 2008.

While the U.S. was the major net transmitter of shocks to others, France and Italy had also become important net transmitters during the global recession of 2008–09. Actually, the net connectedness of both countries stayed high after late 2008, fluctuating between 50 and 100%. According to Figure 5.6, since the global financial crisis, the net connectedness of Germany, the U.K., and Japan declined rapidly, dropping all the way to –100% in early 2010.

5.3.5 USING COUNTRY REAL-ACTIVITY FACTORS

In order to check the robustness of our results obtained from the industrial production indexes, in this short section we use a set of alternative measures of the behavior of each of the G-7 economies over the business cycle. Recently, using data on major macroeconomic variables at monthly and quarterly frequency Aruoba et al. (2011) estimated dynamic factor models for the G-7 countries and derived country factors for each of the countries using thirty-seven monthly indicators. Aruoba et al. (2011) showed that the country factors captured the main macroeconomic developments over a period of forty years and their behavior over time were fairly consistent with the business cycle narrative for each of the countries.


Applying the connectedness methodology to monthly country factors obtained by Aruoba et al. (2011), we calculated the total connectedness index for different window lengths. In Figure 5.7 we present the total connectedness index obtained from a VAR of the country factors over a sample window of five years through seven years, along with total connectedness indexes obtained from the VEC1 model of monthly industrial production indexes.

When we use five-year and six-year rolling windows to calculate the connectedness index, the behavior of the country-factor and industrial-production-based connectedness indexes is quite similar. In the case of seven-year window there is a level difference between the two indexes. The industrial-production-based index tends to be higher than the country-factor-based index. Despite that level difference, however, the two indexes behave quite similarly over time. Based on Figure 5.7, we can conclude that the connectedness of business cycles across the G-6 countries are well captured by the use of industrial production data.


5.3.6 INTERNATIONAL TRADE AND DIRECTIONAL CONNECTEDNESS

Germany has been the biggest economy and the manufacturing powerhouse of Europe. It is therefore not easy to reconcile some of the above dynamic directional connectedness results with the image of Germany as the engine of growth in the E.U. Now, let us discuss the logic behind this result in some detail. Trade flows play a key role in the transmission of shocks across countries. When there is a shock to domestic demand in country i , this shock is transmitted to other countries through the trade channel. As the aggregate demand in country i takes the hit, the demand for imports is affected as well. As a result, the domestic shock is likely to be transmitted to other countries that are major exporters to that country.


As can be seen in Table 5.5, from 1999 to 2008, Germany's average trade surplus in manufacturing vis-à-vis the other five countries was equivalent to



(a) Five-year window



(b) Six-year window



(c) Seven-year window

Figure 5.7. Total connectedness as assessed using country factors vs. industrial production

Note: We show the country-factor version as a dashed line and the industrial production version as a solid line. See Section 5.3.5 for details.

Table 5.5 Bilateral manufacturing trade balance relative to local manufacturing production, 1999–2008 average

	U.S.	Germany	Japan	France	U.K.	Italy
U.S.	–	2.4	2.7	1.4	1.6	0.6
Germany	–0.6	–	0.1	–2.4	–5.0	–1.7
Japan	–1.2	–0.2	–	–0.1	–1.6	–0.1
France	–0.1	1.6	0.02	–	0.1	0.8
U.K.	–0.1	1.5	0.3	0.7	–	0.9
Italy	–0.2	1.1	0.02	–0.2	–1.2	–
Total	–2.0	6.5	3.1	–0.5	–6.1	0.5

Note: Each cell shows the manufacturing trade balance of the column country with the row country, divided by the industrial production of the column country. For example, Germany's manufacturing trade surplus with the U.S. is 2.4% of German industrial production, the corresponding U.S. manufacturing trade deficit with Germany is 0.6% of U.S. industrial production.

Source: Authors' calculations based on OECD data.

6.5% of its industrial output. Over the same period the U. K., the U. S., and France ran manufacturing trade deficits, while Japan and Italy ran manufacturing trade surpluses vis-à-vis the other G-6 countries. Germany happens to be the most important exporter of manufacturing goods to France, the U.K., and Italy, and it ranks only second or third among exporters to the U.S. and Japan. As a result, when there is a negative shock to industrial production in one or more of the G-6 countries, this shock is likely to be transmitted, first and foremost, to Germany and then to the other countries. From this perspective, it is logical for Germany to have a higher connectedness from others compared to its connectedness to others.

In order to better understand the possible link between the trade balance and the business cycle connectedness, we undertake a linear regression analysis. In the regressions, the dependent variable is logarithmic net directional connectedness ($\text{Log}(C_{\bullet \leftarrow i} / C_{i \leftarrow \bullet})$). On the right-hand side we include only the logarithmic trade balance ($\text{Log}(Exports/Imports)$), with a lag of twelve months. Regression results for the full sample (April 1963 to February 2010) as well as for the sub-sample that covers the period from January 1990 to the end of the sample are presented in Table 5.6.

The results for the full sample (April 1963 to February 2010) are not encouraging. Only the U.S. has a statistically and economically meaningful coefficient estimate with the expected negative sign. The estimated coefficient implies that when the ratio of one-year lagged U.S. exports to U.S. imports increases by 1%, the ratio of the connectedness from the U.S. to the connectedness received by the U.S. will decline by 1.76%. The estimated coefficient for the U.K. (6.43) is also statistically significant, and its positive sign implies that a decrease in the export–import ratio will lead to a decrease in the connectedness from the U.K. relative to the connectedness received by the U.K.

With the globalization process underway, trade flows became more and more important in the 1990s. Therefore, it makes sense to focus on the recent

Table 5.6 Directional connectedness and the trade balance

	April 1963–February 2010		January 1990–February 2010	
	Coefficient	R^2	Coefficient	R^2
France	2.45 (2.91)	0.70	-4.95 (4.36)	0.80
Germany	-3.85 (2.63)	0.70	-6.31** (2.4)	0.61
Italy	0.59 (1.42)	0.68	-4.10** (1.17)	0.56
Japan	-1.16 (1.34)	0.75	-6.98** (2.46)	0.74
U.K.	6.43* (2.61)	0.74	6.46+ (3.45)	0.71
U.S.	-1.76** (0.49)	0.76	-2.59** (0.91)	0.63

Note: We regress logarithmic net directional connectedness ($\text{Log}(C_{i \rightarrow j} / C_{j \rightarrow i})$) on the logarithmic trade balance ($\text{Log}(\text{Exports} / \text{Imports})$), lagged twelve months. Standard errors are presented in parentheses. **, * and + indicate statistical significance at one, five, and ten percent levels, respectively.

decades. For that reason, we restrict the sample to the post-1990 period. Five out of six estimated coefficients for the post-1990 period had negative signs, as expected. For four of these countries (France is the exception), the negative coefficient estimates are statistically different from zero at the 1% significance level. The estimated elasticity for the U.K. is still positive, but it is different from zero at the 10% level of significance only. The U.K. runs chronic deficits in merchandise trade, which is financed by chronic trade surpluses in services. Being a deficit country in merchandise trade, the U.K. is likely to be a net transmitter of shocks to other countries. However, given the large size of its trade surplus in services (service exports revenue reached \$249 billion in 2009 compared to \$117 billion in merchandise exports), the U.K. is likely to be a net receiver of shocks in services sectors from other countries. As a result, when we regress the log of the ratio of connectedness transmitted and received by the U.K. on the log of the export–import ratio for goods, the coefficient turns out to be positive but statistically insignificant.

5.4 Concluding remarks

We have applied the connectedness methodology to the study of business cycles in several leading economies. In doing so, we have made several contributions to the literature on international business cycles.

First, the connectedness methodology is different from, and complements, the empirical approaches typically used. While the factor model approach aims to obtain a world business cycle measure, the connectedness framework

distinguishes between idiosyncratic shocks to industrial production and the connectedness of industrial production shocks across countries. Furthermore, the connectedness index is based on a multivariate VECM, which is better able to capture the increased comovement of business fluctuations in more than two countries compared to analyses based on pairwise correlations.

Second, the analysis sheds new light on the nature of business cycles, clearly showing that the cross-country comovement of business fluctuations is not constant over time, and that it does follow an upward trend. Rather, business cycle connectedness fluctuates substantially over time. However, the band within which the connectedness index fluctuates has increased since 1984. This result is consistent with the findings of both Kose et al. (2003) and Doyle and Faust (2005): when shocks in individual countries are small, they cannot be expected to be transmitted to other countries regardless of the integration among countries, but when the shocks are large enough they are transmitted, and the cross-country correlation of macroeconomic aggregates increases.

Third, we use directional connectedness measures to identify each country as a gross and/or net transmitter of business cycle shocks to other countries as well as a gross/net recipient of business cycle shocks from other countries over different time periods. The directional connectedness measures show that the U.S. (1980s and 2000s) and Japan (1970s and 2000s) are the major net transmitters of shocks to other countries, while Germany is the major net receiver of shocks in the 2000s.

Fourth, with an unprecedented jump between May and December 2008, the business cycle connectedness index captures well the global nature of the current recession. The quick and large increase in cross-country connectedness suggests that recovery from “the great global recession” may require coordinated policy actions among the major industrial and emerging market economies.

Fifth, we showed that the sign of a country’s net business cycle connectedness is closely related to the country’s trade balance. Those countries that run trade surpluses tend to be net recipients of shocks, whereas countries that run trade deficits are likely to be transmitters of shocks.

■ ACKNOWLEDGMENTS

For helpful guidance we thank the editors, Siem Jan Koopman and Neil Shephard, and two anonymous referees. For support we thank the Real Time Data Research Center at the Federal Reserve Bank of Philadelphia. Yilmaz thanks the Turkish Scientific and Technological Research Council (TUBITAK) for Grant No. 111K500. The usual disclaimer applies.