Federal Reserve Bank of Chicago

Call for Papers

2000 Conference on Bank Structure and Competition

Fourth Quarter 1999

Economic perspectives

Economic perspectives

Economic Perspectives is published by the Research Department of the Federal Reserve Bank of Chicago. The views expressed are the authors' and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

Single-copy subscriptions are available free of charge. Please send requests for single- and multiple-copy subscriptions, back issues, and address changes to the Public Information Center, Federal Reserve Bank of Chicago, P.O. Box 834, Chicago, Illinois 60690-0834, telephone 312-322-5111 or fax 312-322-5515.

Economic Perspectives and other Bank publications are available on the World Wide Web at http://www.frbchi.org.

WWW.FRBCHI.ORG

ISSN 0164-0682

Contents

Fourth Quarter 1999, Volume XXIII, Issue 4

Is the EMU a viable common currency area? A VAR analysis of regional business cycles

Michael A. Kouparitsas

Introduction and summary

In January 1999, 11 European countries bravely launched into a common currency area known as the European Monetary Union (EMU). By joining the common currency area, member countries have agreed to keep the value of their national currency fixed in terms of the currencies of the other EMU countries for an indefinite period. Consumers and businesses in these countries will, however, find that very little has changed. The most noticeable change will not occur until 2002 when national currencies are replaced by a common currency known as the *euro*. In the intervening period, prices will be denominated in terms of existing national currencies and euros. Consumers using cash will pay the national currency price, while consumers using credit cards (including U.S. visitors to the euro zone) will notice that their transactions are carried out in euros.

Although they might disagree about the exact size of the gains, most economists would agree that the EMU will yield significant microeconomic benefits through lower transactions and hedging costs. According to the European Commission, the gains from carrying out transactions in a single currency could be as high as 0.5 percent of European Union gross domestic product (GDP) per year. However, many economists are skeptical about the long-run viability of the EMU. Euro-zone members have given up the right to set their own interest rates and the option of moving their exchange rates against each other. The widespread view is that this loss of flexibility may involve significant costs (in the form of persistent high unemployment and low output growth) if their economies do not behave as one or cannot easily adjust in other ways. The ultimate concern is that for some countries, these macroeconomic costs will eventually outweigh the microeconomic benefits and lead them to abandon the EMU.

How well the EMU performs along the macro dimension will depend on how closely it fits the notion of an "optimal currency area" (OCA). Beginning with Mundell (1961), economists have long agreed that the following four criteria must be met for a region to be an optimal currency area: 1) countries should be exposed to similar sources of disturbances (common shocks); 2) the relative importance of these common shocks should be similar (symmetric shocks); 3) countries should have similar responses to common shocks (symmetric responses); and 4) if countries are affected by country-specific sources of disturbance (idiosyncratic shocks), they need to be able to adjust quickly. The basic idea is that countries satisfying these criteria would have similar business cycles, so a common monetary policy response would be optimal.

How far the euro zone is from an OCA is an open question for research, as is the more important question of whether the apparent deviation from an OCA is sufficient to question the long-run viability of the EMU. On the surface, the data seem to support the skeptics' view that the EMU is not an OCA. First, euro-zone countries have experienced frequent and often large idiosyncratic shocks over recent years. A well-known example is German reunification, which many argue led to the breakdown of the precursor to the EMU known as the European Monetary System (EMS) in 1992.[1] Second, persistently high unemployment rates throughout Europe suggest that EMU economies (especially their labor markets) are slow to adjust to all economic disturbances.

Michael A. Kouparitsas is an economist at the Federal Reserve Bank of Chicago. *The author would like to thank David Marshall and Jonas Fisher for useful comments on an earlier draft.*

The purpose of this article is to formally assess the long-run viability of the EMU. I do this by comparing the sources and responses to economic shocks to the EMU with those from a well-functioning currency union, the U.S. My working hypothesis is that if the EMU is as close to an OCA as the U.S. is, based on the criteria outlined above, it may well be a viable currency union in the long run. If, on the other hand, the EMU is less like an OCA than the U.S. is, one might question the long-run viability of this monetary union.

Despite all the effort that has gone into the EMU debate, there is little in the way of empirical research on the sources and responses to economic shocks to this region. I use a statistical technique known as a structural vector autoregression (VAR) to extract these components from the data. My analysis suggests that U.S. regions are highly symmetric. U.S. regions face common sources of disturbance, to which they respond in a similar way. In contrast, the EMU countries can be grouped into a symmetric *center* and a clearly asymmetric *periphery*. Center countries are Austria, Belgium-Luxembourg (treated as one country for data purposes), France, Germany, Italy, the Netherlands, Portugal, and Spain, while the periphery countries are Finland and Ireland. Center countries display many of the characteristics of U.S. regions when compared on OCA criteria. Periphery countries appear to have quite different sources of disturbance from the center. In addition, they seem to respond to common shocks in a different way from the center countries. I conclude on the basis of this statistical analysis that the EMU will be a viable currency union for the center countries, but question the viability of a union with countries in the periphery.

Previous empirical analysis of the EMU

The EMU has spawned a number of empirical papers aimed at understanding the nature of regional business cycles and the regional impact of fiscal and monetary polices. The approaches vary considerably. For example, Carlino and DeFina (1998b) examine the regional effects of monetary policy within the EMU. Their approach is indirect. In earlier work, Carlino and DeFina (1998a) estimated the effects of U.S. monetary policy on the 48 contiguous U.S. states (and eight Bureau of Economic Analysis [BEA] regions). They build on this analysis in the later paper by estimating the cross-sectional relationship between the long-run regional output response to monetary policy and industry structure. Their findings suggest that monetary policy has a larger impact on more industrial-oriented U.S. regions, such as the Great Lakes. They use these cross-sectional U.S. findings and the industry structure of EMU countries to speculate on the long-run regional impact of monetary policy within the EMU. Their results suggest that monetary policy will have a differential impact on EMU countries. This implies that the EMU is not an OCA, since it fails to meet the symmetric responses criterion. In a competing study, Dornbusch, Favero, and Giavazzi (1998) test this hypothesis directly using time-series methods and find that the effect of monetary policy is not statistically different across EMU countries. Their study suggests EMU countries have similar responses to monetary policy shocks, which is necessary for a region to be an OCA. An obvious limitation of this work is that it is silent on the incidence of other disturbances affecting the EMU countries and the broader question of whether the EMU will be viable in the long run.

Eichengreen has approached the question of whether the EMU is an OCA from a number of interesting directions. Eichengreen (1992) joins others in gauging the importance of country-specific shocks by computing the variability of bilateral EMU real exchange rates, for example, the real exchange rate between Germany and France. The basic idea is that these relative price fluctuations reflect shifts in demand and supply affecting one EMU country relative to another, so countries with more highly correlated disturbances will have less volatile bilateral real exchange rates. The typical approach of this type of study is to compare the volatility of bilateral EMU real exchange rates with the volatility of relative output prices of U.S. BEA regions. A common finding is that the bilateral real exchange rates of EMU countries are considerably more volatile than the relative output prices of U.S. regions. This suggests that the EMU is further than the U.S. is from being an OCA. An obvious weakness of this type of analysis is that it does not directly compare the EMU and the U.S. using the OCA criteria outlined earlier.

Observing this limitation, Eichengreen and Bayoumi (1993) approach the issue in a more direct way. They estimate individual models for U.S. BEA regions and EMU countries using a technique developed by Blanchard and Quah (1989), which allows them to extract unobserved components from the data that describe so-called demand and supply shocks. Demand and supply shocks are distinguished by the fact that demand shocks are assumed to have a temporary impact on the economy, while supply shocks are assumed to have a permanent effect on the economy. Eichengreen and Bayoumi (1993) then compare the correlation coefficients of German supply (and demand) shocks and those of other EMU countries with the correlation coefficients of U.S. Mideast supply (and demand) shocks and those of other U.S. regions. They show that U.S. regional supply (and demand) shocks tend to be more highly correlated than EMU regional supply (and demand) shocks. The final step of their analysis is to compare regional responses to demand and supply shocks. Their results suggest that the response functions of U.S. regions are more alike than those of EMU countries. On the basis of this analysis, they conclude that the EMU is further than the U.S. is from being an OCA, which leads them to argue that the EMU may find it more difficult than the U.S. to operate a monetary union.

My empirical analysis builds on Eichengreen and Bayoumi (1993) along two dimensions. First, I update their work by analyzing more recent data. Eichengreen and Bayoumi's data spanned the years from 1963 to 1986, while I consider data covering the years from 1969 to 1997. These data are likely to be more informative about the behavior of countries under the EMU, since they include a greater number of years over which the EMU countries were part of the forerunner to the EMU, the EMS. Second, I adopt a different way of decomposing the data that allows me to directly measure the extent to which regional business cycles are driven by common and country-specific shocks. My conclusions differ from Eichengreen and Bayoumi's. In contrast to their findings, I show that with the exception of two relatively small countries, Finland and Ireland, the euro zone shares many of the regional business cycle characteristics of the U.S. In other words, the EMU comes as close to being an OCA as the U.S. does. I argue on the basis of these results that the long-run viability of the EMU is similar to that of the U.S. monetary union.

A weakness of all the foregoing empirical research is that historical data may be an unreliable guide to the way euro-zone countries will behave under the EMU. This observation is a simple application of the *Lucas critique*. The basic idea is that historical data may be uninformative since the structure of euro-zone economies (and possibly the world economy) will likely undergo significant change after the EMU adopts a common currency. Frankel and Rose (1998) find empirical support for this proposition by showing that one form of structural change that may occur under the single currency, greater trade flows between countries, leads to more highly correlated business cycles. A consequence of their work for all EMU studies is that countries that may appear from historical data to be poor candidates for inclusion in the euro zone may indeed turn out to be suitable candidates after joining the union. This clearly has implications for earlier work that argued against the long-run viability of the EMU. I argue that the EMU will be viable in the long run, so Frankel and Rose's results merely reinforce my conclusions.

How similar are EMU country business cycles?

A simple and direct way of assessing the similarity of regional business cycles is to calculate the correlation between aggregate and individual region business cycles. High correlations are indicative of common sources and responses to disturbances. In figure 1, I plot cyclical movements in U.S. aggregate and regional real income.[2] The underlying data are BEA annual state personal income from 1969 to 1997. These data are deflated by the national consumer price index.[3] I use personal income rather than gross state product because the former span a longer period.[4] The eight BEA regions are the Great Lakes, Plains, New England, Mideast, Southeast, Southwest, Rocky Mountains, and Far West.[5] The lowest correlation between a region and the U.S. aggregate is 0.76 for the Southwest, with the highest at 0.98 for the Southeast and Great Lakes. This suggests that common shocks explain a large share of the variation in U.S. regional income.

I repeat this exercise for the EMU. Figure 2 plots the cyclical fluctuations of aggregate and regional EMU income. The underlying data are International Monetary Fund (IMF) estimates of real annual GDP from 1969 to 1997. The correlations between regional and aggregate activity can easily be divided into two groups. The first group—Austria, Belgium-Luxembourg, France, Germany, Italy, the Netherlands, Portugal, and Spain—resemble the U.S. regions, with correlations ranging from 0.72 (Spain) to 0.90 (Germany and Italy). The second group of Finland and Ireland, with correlations of 0.45 and 0.58, respectively, appear to have business cycles that are quite different from the rest of the euro zone.

With the exception of Finland and Ireland, the coherence between EMU regional business cycles appears to be as high as that of U.S. BEA regions. On the basis of these results, a subset of the EMU can not be ruled out as a viable currency union. An obvious weakness of this approach is that it does not allow for a comparison of the sources of disturbances or responses to disturbances across regions. Next, I describe a statistical technique that overcomes this limitation. Using these results, I can more closely gauge the extent to which the EMU and the U.S. meet the OCA criteria described earlier.

Are the sources of shocks and responses to them similar across EMU countries?

Methodology

My starting point for isolating the sources of regional shocks and responses to them is recent work analyzing the regional effects of U.S. monetary policy. The typical approach is to use a structural vector autoregression (VAR). A VAR is a statistical method that allows one to estimate how an unpredictable change (or disturbance) in one variable affects other variables in the economy. For example, one of the questions raised by theoretical research is whether a change in monetary policy has a stronger effect on regions that devote a larger share of activity to industrial production. A VAR allows one to estimate the way that an unpredicted change in monetary policy

FIGURE 1

Cyclical movements of U.S. regional output
(percent)

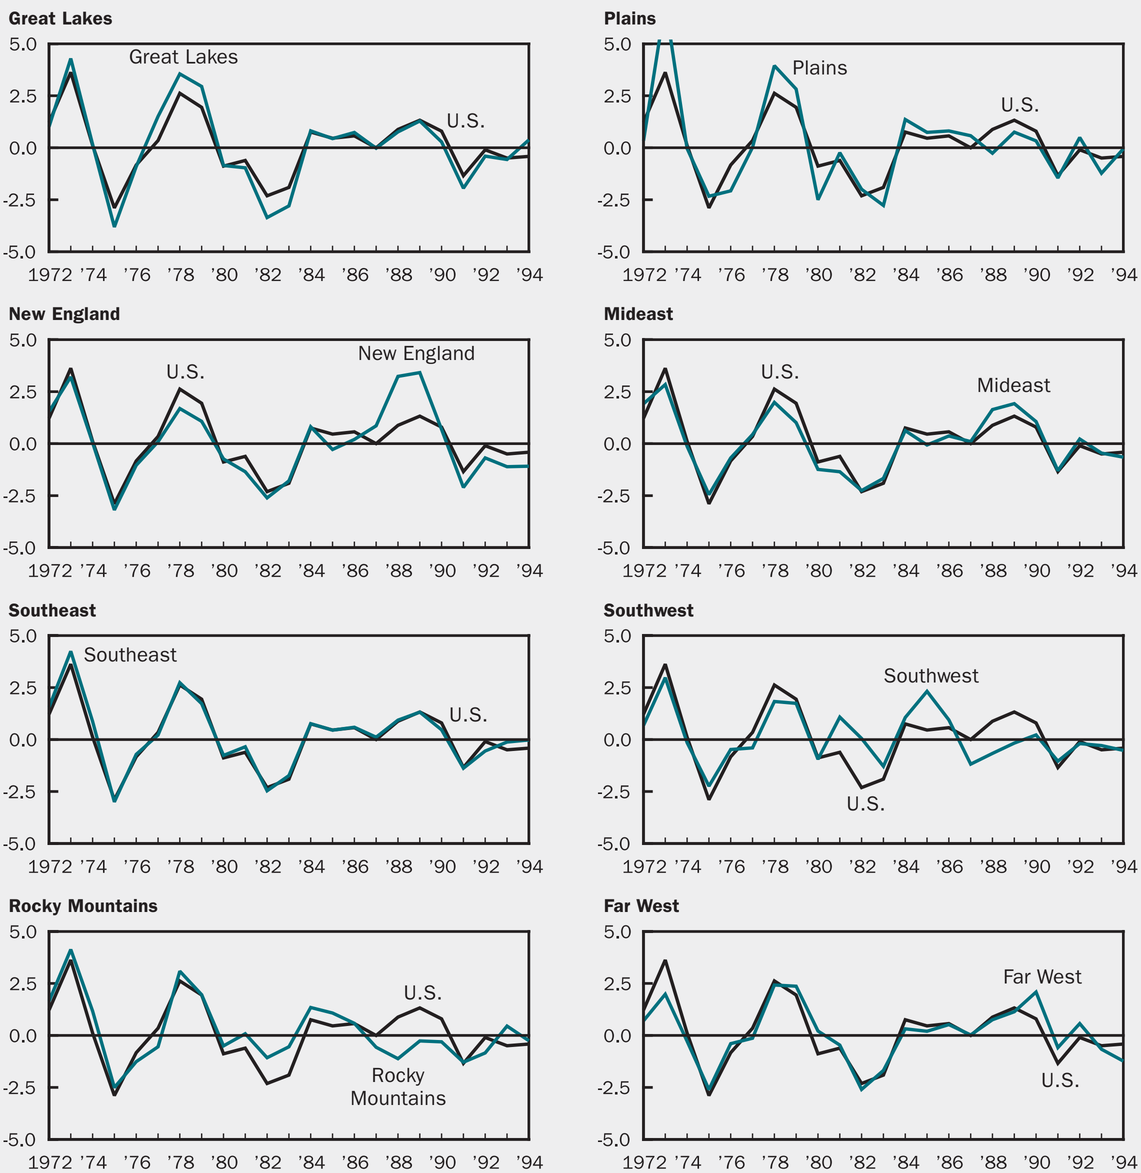

Note: Personal income data are filtered using the annual "business cycle" band-pass filter described in Baxter and King (1995).

Source: Author's calculations from U.S. Department of Commerce, Bureau of Economic Analysis, 1969–97, "State personal income," database.

FIGURE 2

Cyclical movements of EMU regional output
(percent)

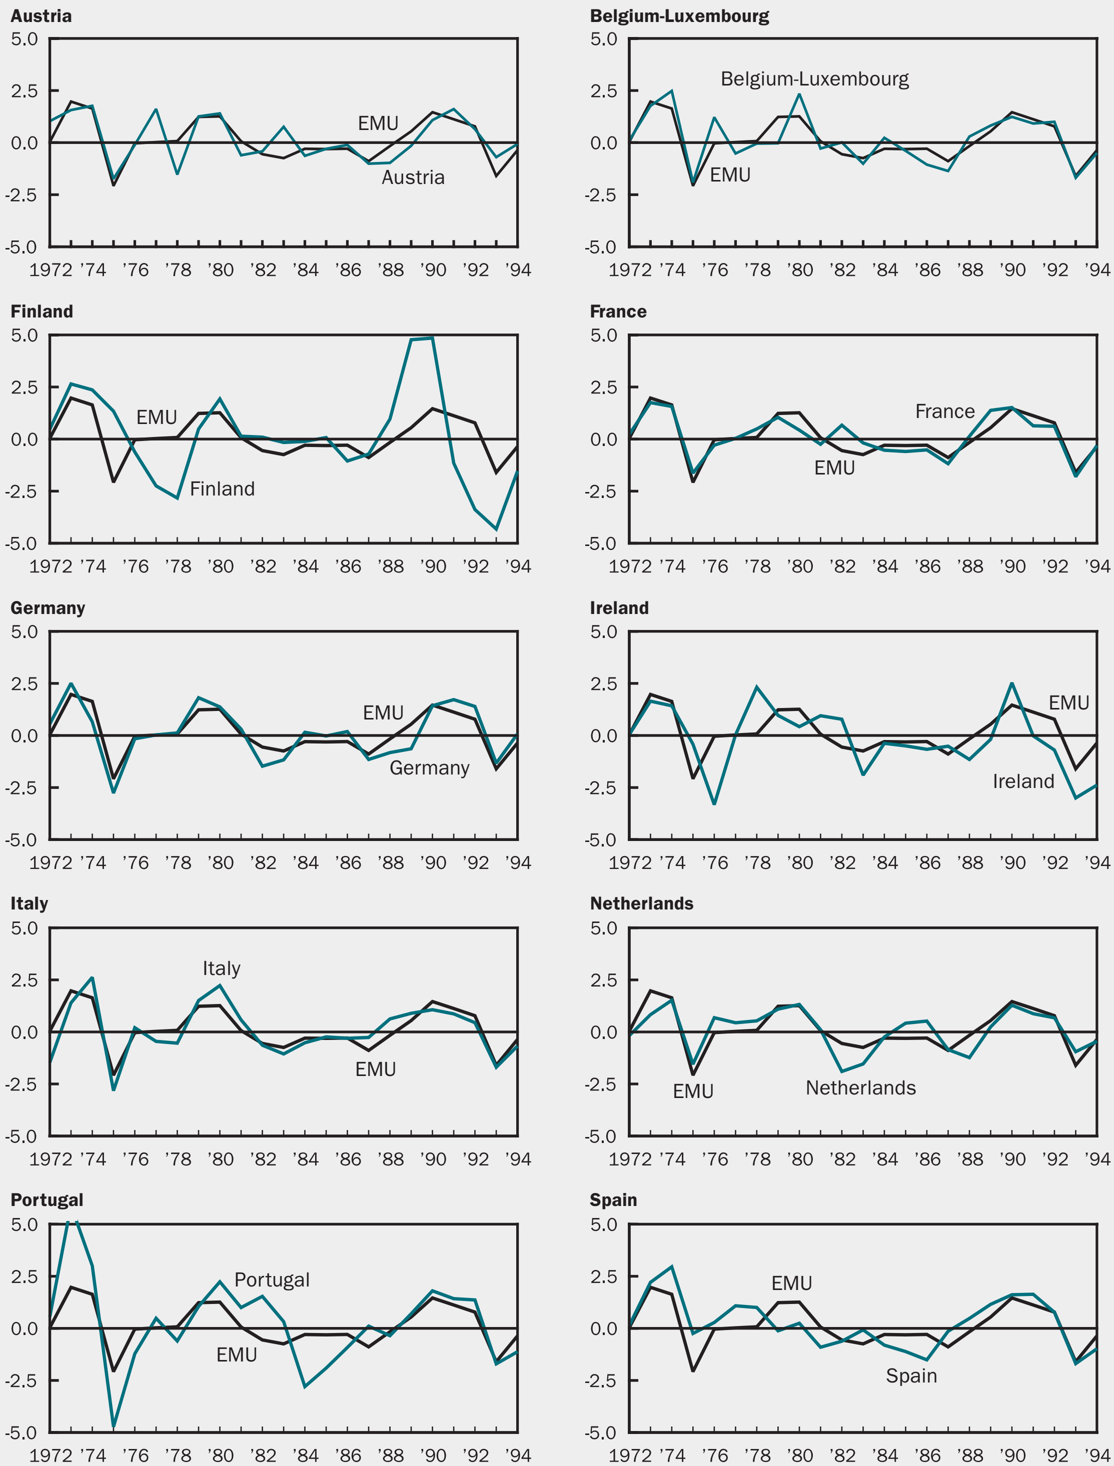

Note: Gross domestic product filtered as described in figure 1.
Source: Author's calculations from International Monetary Fund, gross domestic product data.

affects the output of regions with relatively large and small industrial sectors.

There is a wide range of variables one can use in analyzing regional business cycles. I follow the approach of Carlino and DeFina (1998a) by limiting the analysis of U.S. regional business cycles to eight VARs, which essentially study interaction between the U.S. and a given region, in this case the eight BEA regions. I adopt a slightly different structural model by drawing on the approach of Christiano, Eichenbaum, and Evans (1994) in their work on identifying and measuring the aggregate effects of U.S. monetary policy shocks. Each U.S. regional VAR is designed to study how unpredicted changes in world oil prices, aggregate U.S. and regional income, and U.S. monetary policy (U.S. federal funds rate) affect the region's income.

VAR studies of international business cycles take a somewhat similar approach to the U.S. regional business cycle literature. International research has focused almost exclusively on the relationship between U.S. and G-7 (Group of Seven) business cycles under different exchange rate regimes.[6] This type of analysis is generally restricted to bilateral VARs involving the U.S. and a G-7 country. I adapt this approach to the EMU. I employ 10 VARs. Just as in the U.S. regional case, each EMU VAR is designed to study how unpredicted changes in world oil prices, aggregate EMU and country of interest income, and EMU region monetary policy (German short-term interest rate) affect the EMU country's income.

I estimate the U.S. and EMU VARs using annual data over a common period spanning 1969 to 1997. I limit the U.S. and EMU VARs so that they estimate relationships between the four variables (world oil prices, aggregate income, regional income, and a regional short-term interest rate) with data from the last two years. In other words, I estimate the link between movements in aggregate and regional income that occurred within the last two years.

Before I can shed light on the nature of regional disturbances and responses to them, I need to impose some structure on the system of equations described by the VARs. There are numerous forms of identifying restrictions in the literature. In their work on the EMU, Eichengreen and Bayoumi (1993) impose long-run restrictions on the data motivated by a theoretical model. I use a recursive structure popularized by Sims (1972). This approach imposes restrictions on the covariance structure of the disturbances of the model. In particular, structural disturbances are identified by imposing a recursive information ordering. Throughout the analysis, I impose the following information ordering: world oil prices; aggregate regional income; indicator of regional monetary policy; and regional or country income. This approach assumes, as in Christiano, Eichenbaum, and Evans (1994) that the monetary authority chooses the value of the monetary policy instrument after observing contemporaneous movements in oil prices and aggregate output.[7] In this setting I can conveniently refer to the structural disturbances as an oil price or global shock, aggregate output shock, monetary policy shock, and region- or country-specific output shock.

With these models in hand, I am able to assess the similarity of EMU and U.S. regional business cycles along two dimensions. First, by studying the sources of regional economic disturbances in the U.S. and EMU, I can determine the extent to which fluctuations are caused by common and idiosyncratic shocks. In the U.S. case, common shocks include unpredicted changes to world oil prices, aggregate U.S. income, and U.S. monetary policy (U.S. federal funds rate). Similarly, in the case of the EMU, aggregate shocks include unpredicted changes to world oil prices, aggregate EMU income, and EMU monetary policy (German short-term interest rate). Idiosyncratic shocks are captured by U.S. region-specific and EMU country-specific output shocks. The relative importance of the various sources of disturbance will be revealed by the share of the one-step-ahead forecast error of U.S. region or EMU country income that is due to unpredicted changes in the disturbance. In a perfectly symmetric case, regions would have none of their forecast error explained by region-specific shocks and the same shares for the various common shocks.

Second, by studying the responses to economic disturbances, I can assess whether regions have similar responses to common shocks and determine the time it takes regions to respond to idiosyncratic shocks. The way that region and country income responds to various disturbances will be embodied in the estimated parameters of the VAR and revealed through the shape and size of the model's impulse response function. For a description of the methodology in greater detail, see the appendix.

Do U.S. regions have similar economic disturbances?

Tables 1 and 2 report decompositions of the forecast errors of income for U.S. regions and EMU countries, respectively. These decompositions indicate the share of the error attributable to a particular disturbance for a given forecast horizon. The one-step-ahead errors are informative about the similarity of disturbances across regions within a currency area, while step lengths of greater than one contain joint information about the similarity of disturbances and responses to disturbances.

Table 1 reveals that a large share of the disturbance to U.S. regions is due to common shocks (that is, unanticipated shocks to world oil prices, aggregate U.S. income, and U.S. monetary policy). For example, common disturbances explain a large share of the variation in the Southeast, Great Lakes, Mideast, and Far West's one-step-ahead forecast error (84 percent to 95 percent). The Rocky Mountains and Plains appear to have the largest region-specific influences, with 60 percent and 64 percent, respectively, of the variation in their one-step-ahead forecast errors explained by common disturbances. New England and the Southwest fall somewhere in between, with common disturbances explaining a little more than 70 percent of the variation in their one-step-ahead forecast errors. The relative importance of different common shocks is also similar across U.S. regions. Shocks to aggregate U.S. income are a more important source than shocks to world oil prices and

TABLE 1

Forecast error variance decompositions for real personal income of U.S. regions

Great Lakes

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	Great Lakes income
1	35	58	0	6
2	39	53	5	3
5	21	21	57	1
10	26	20	51	3

Plains

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	Plains income
1	16	47	0	36
2	25	54	2	18
5	18	33	33	15
10	23	29	29	19

New England

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	New England income
1	35	36	0	29
2	38	14	5	44
5	33	4	26	37
10	33	8	29	29

Mideast

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	Mideast income
1	12	74	1	14
2	16	42	11	31
5	24	15	33	27
10	26	17	31	25

Southeast

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	Southeast income
1	41	54	0	5
2	58	36	2	4
5	39	14	37	10
10	38	14	39	9

Southwest

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	Southwest income
1	2	72	0	26
2	1	68	2	30
5	3	50	16	31
10	2	48	26	24

Rocky Mountains

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	Rocky Mtns. income
1	20	40	0	40
2	24	30	2	44
5	10	17	32	40
10	9	19	46	26

Far West

Years ahead	Percentage of forecast error due to: Oil prices	U.S. income	Fed funds rate	Far West income
1	26	57	1	16
2	40	42	0	18
5	42	32	7	18
10	43	31	13	13

Notes: Each panel describes the decomposition of the forecast error for the region of interest's income. The first column in each block refers to the number of years (s = 1, 2, ..., 10) ahead for the forecast. Columns indicate the percentage of the s-step-ahead forecast error arising from a particular structural disturbance.

Source: Calculations from author's statistical model, using the following annual data series: IMF—world crude oil prices; BEA—personal income by state; and Federal Reserve Board of Governors—federal funds rate.

U.S. monetary policy. Overall, these results suggest that U.S. regions have similar sources of economic disturbances.

Table 1 also provides some indication of the similarity of responses to disturbances. Looking at horizons of greater than one year, the relative importance of

TABLE 2

Forecast error variance decompositions for real gross domestic product of EMU countries

Austria

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Austrian GDP
1	17	43	1	39
2	12	50	15	24
5	13	23	50	13
10	22	13	56	9

Belgium-Luxembourg

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Bel-Lux GDP
1	24	56	0	20
2	19	52	13	16
5	17	21	55	7
10	26	18	49	7

Finland

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Finnish GDP
1	3	0	4	93
2	1	4	2	94
5	17	12	5	66
10	19	14	7	60

France

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	French GDP
1	1	80	0	20
2	15	57	5	23
5	20	21	48	11
10	16	23	47	13

Germany

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	German GDP
1	1	77	0	22
2	1	59	15	25
5	6	34	43	17
10	10	35	42	14

Ireland

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Irish GDP
1	0	3	2	95
2	0	7	9	85
5	2	2	16	80
10	1	2	7	91

Italy

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Italian GDP
1	15	33	13	39
2	14	32	15	39
5	17	10	49	25
10	19	8	44	28

Netherlands

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Dutch GDP
1	12	49	6	33
2	6	41	20	33
5	3	18	41	37
10	2	20	42	35

Portugal

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Portuguese GDP
1	1	47	14	38
2	1	44	9	46
5	5	18	46	32
10	7	21	41	30

Spain

	Percentage of forecast error due to			
Years ahead	Oil prices	EMU GDP	EMU interest rate	Spanish GDP
1	2	45	15	38
2	1	36	27	36
5	6	22	46	25
10	11	21	45	24

Notes: Each panel describes the decomposition of the forecast error for the country of interest's GDP. The first column in each block refers to the number of years (s = 1, 2, ..., 10) ahead for the forecast. Columns indicate the percentage of the s-step-ahead forecast error arising from a particular structural disturbance.

Source: Calculations from author's statistical model, using the following annual data series: IMF—world crude oil prices, interest rates, and gross domestic product.

common and idiosyncratic disturbances is largely unchanged. This suggests that responses are fairly similar. A common finding is that unanticipated shocks to aggregate U.S. income are less important at longer horizons.

Are EMU country economic disturbances more alike than those of U.S. regions?

Table 2 reports forecast error decompositions for the income of EMU countries. Concentrating on the one-step-ahead forecast error, countries fall into three groups. Common shocks explain about 80 percent of the one-step-ahead forecast errors of income in Belgium-Luxembourg, France, and Germany. This share is a little above 60 percent for Austria, Italy, the Netherlands, Portugal, and Spain. The outliers are Finland and Ireland, where this share falls below 10 percent.

The decompositions of the first EMU group are similar to the U.S. group comprising the Great Lakes, Southeast, Mideast, and Far West. The second EMU group has forecast error decompositions that are close to those of the U.S. Rocky Mountain and Plains regions. In both cases, oil price shocks are relatively less important than in their U.S. counterpart, while interest rate shocks are relatively more important than in the U.S. regions. Just as in the U.S., aggregate income shocks are the most important economic disturbance to EMU country income. The findings suggest that, with the exception of Finland and Ireland, EMU country economic disturbances are as alike as those of U.S. regions.

Again, ignoring Finland and Ireland, the long-horizon picture of EMU disturbances is also similar to the U.S. This suggests that EMU responses to disturbances may well be as alike as U.S. responses.

Do U.S. regions have similar responses to economic disturbances?

Figures 3–6 describe in detail the responses of the eight BEA regions to common and idiosyncratic shocks. The black lines trace the impulse response functions of regional income: the way regional income responds over time to a one standard deviation shock to world oil prices, aggregate output, U.S. monetary policy, and regional income, respectively. (The colored lines are the 95 percent confidence bands of these impulse response functions.) These figures show that U.S. regions have similar responses to common disturbances (unanticipated shocks to world oil prices, aggregate U.S. output, and U.S. monetary policy) and that they adjust to idiosyncratic shocks over a period of about two years.

Figure 3 shows that an unanticipated increase in the growth rate of world oil prices has a significant negative impact on the income of seven of the eight U.S. regions, which persists for about one year. The exception is the Southwest, which is the largest oil producing region of the U.S. Although the result is not statistically significant, an increase in the growth rate of world oil prices raises Southwest real income.

In contrast, figure 4 reveals that an unexpected positive shock to aggregate U.S. income has an immediate positive impact on the income of all U.S. regions. The effect of this shock on regional income is generally not statistically significant beyond two years. The only exception is the Southwest, where the aggregate income shock has a statistically significant effect six years after the shock.

Figure 5 shows that an unexpected tightening of U.S. monetary policy (an unexpected rise in the U.S. federal funds rate) tends to have a statistically significant effect on U.S. regional income two years after the shock. The exceptions are the Southwest and Far West. In both cases, the impulse response function is virtually identical to those of other U.S. regions, but not statistically different from the zero line.

Turning to idiosyncratic shocks, figure 6 reveals that U.S. regions adjust quickly to region-specific disturbances. The regions can be divided into two groups. The first group, consisting of the Great lakes, Plains, Southeast, and Far West, have responses that are not statistically significant beyond the year in which the shock occurs. The second group, comprising New England, Mideast, Southwest, and Rocky Mountains, have responses that are statistically significant for no more than three years after the shock.

Do EMU countries have responses that are more alike than those of U.S. regions?

Figures 7–10 (pages 15–18) describe in detail the response functions of the EMU countries to common and idiosyncratic disturbances. These figures suggest that, with the clear exceptions of Finland and Ireland, the response functions of EMU countries are at least as alike as those of U.S. regions. In addition, the response functions imply that contrary to the general view, EMU countries adjust to idiosyncratic shocks at the same speed or faster than U.S. regions.

In contrast to the U.S. result, figure 7 shows that an unexpected positive shock to the change in world oil prices does not have a statistically significant effect on the income of all EMU countries.

However, figure 8 shows that an unanticipated positive shock to aggregate EMU output has a statistically significant positive effect on the output of most EMU countries that dies out one year after the shock. Again, the exceptions are Finland and Ireland, where the effects of the aggregate output shock are not statistically significant.

Turning to the regional monetary shock, we see in figure 9 that EMU responses are not only similar across countries, but also quite similar to the U.S. response functions. As in the U.S., an unanticipated tightening in regional monetary policy (an unanticipated increase in the German overnight money market rate) leads to a contraction in regional income two

FIGURE 3

U.S. output response: Shock to world oil prices
(percent)

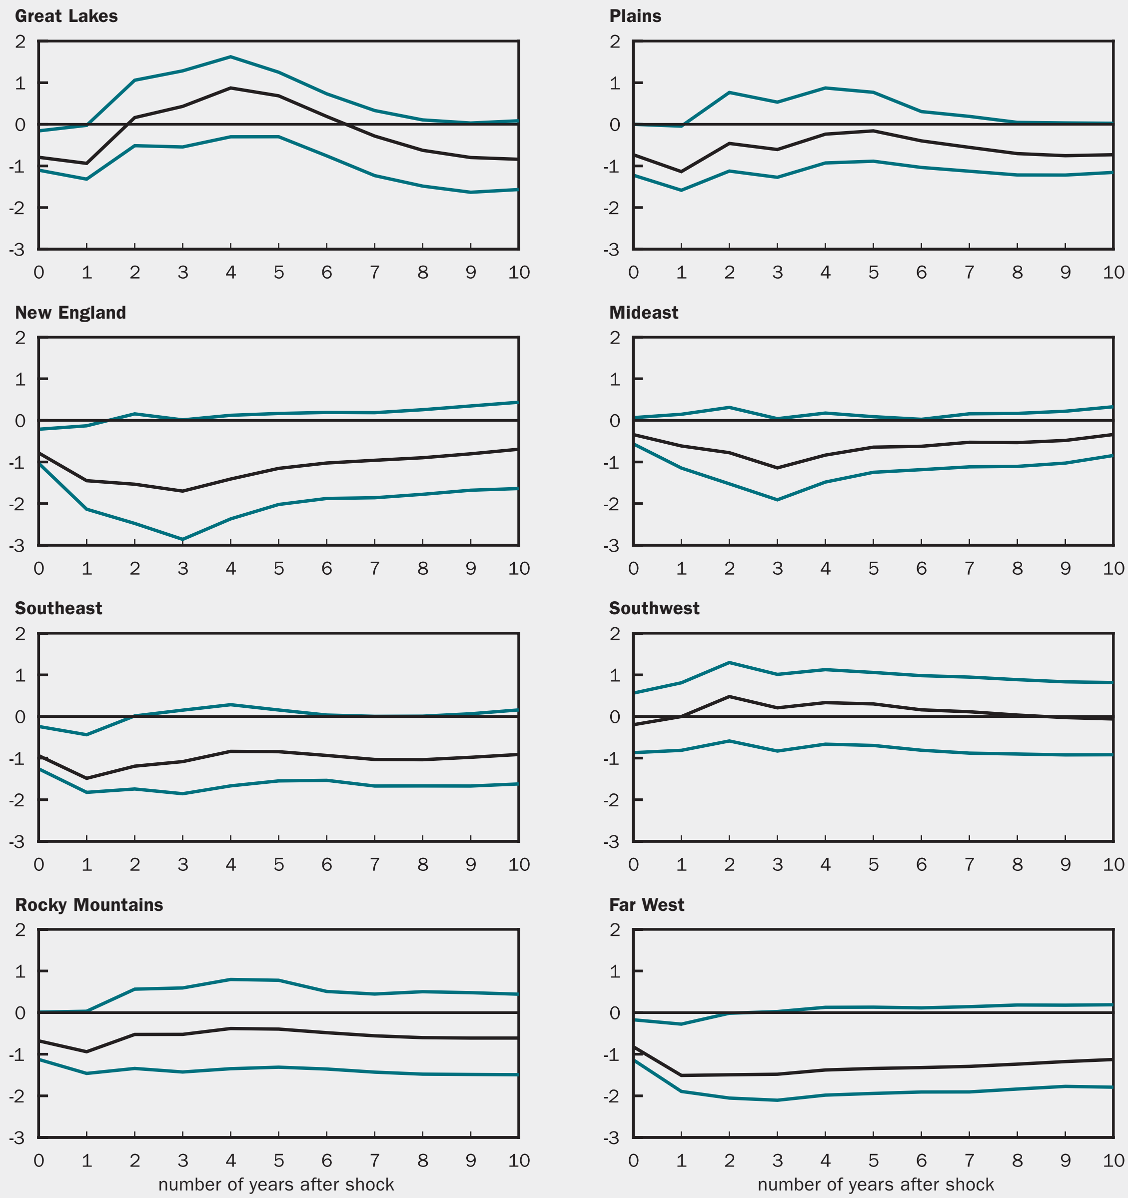

Notes: All figures report changes in the region of interest's real personal income following a one standard deviation shock to the given variable. The black line represents the point estimates of the impulse response function for the region of interest. The colored lines are the 95 percent confidence bands, computed by Monte Carlo simulation using 1,000 independent draws.

Source: Calculations from author's statistical model, using the following annual data series: International Monetary Fund—world crude oil prices; Bureau of Economic Analysis—personal income by state; and Federal Reserve Board of Governors—federal funds rate.

years after the shock. It is important to note that Finland and Ireland have similar responses to the rest of the EMU, but their responses are not statistically different from zero.

Finally, figure 10 describes the rate at which EMU countries adjust to country-specific shocks. Ignoring Finland and Ireland, there are essentially two groups, just as there are in the U.S. case. The first group, consisting of Austria, Belgium-Luxembourg, France, Germany, and Italy, have response functions that are not statistically different from zero a year after the shock. The second group, the Netherlands, Portugal, and Spain, adjust in under three years. The response functions of Finland and Ireland display considerably longer adjustment periods. In the case of Ireland, idiosyncratic shocks appear to be highly persistent.

FIGURE 4

U.S. output response: Shock to aggregate U.S. income
(percent)

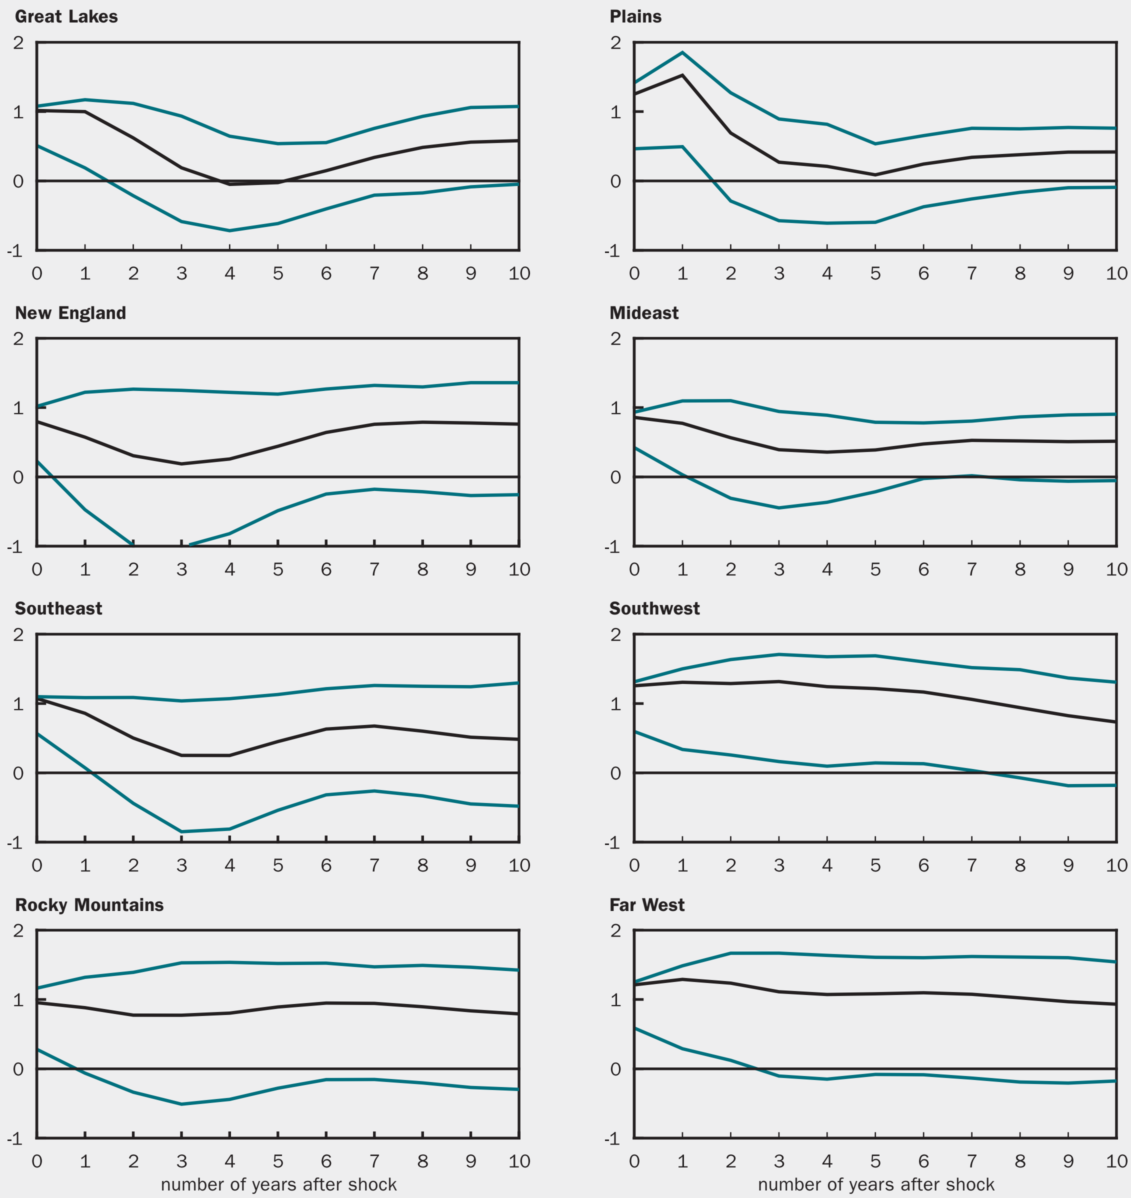

Notes and source: See figure 3.

The lessons learned from the simple business cycle analysis of the previous section carry over to the VAR analysis. The EMU is characterized by a highly symmetric center—Austria, Belgium-Luxembourg, France, Germany, Italy, the Netherlands, Portugal, and Spain—and an asymmetric periphery—Finland and Ireland. As noted earlier, the center countries have highly correlated business cycle fluctuations. The VAR analysis shows that these correlations are supported by common sources of disturbance and similar responses to these shocks. The VAR analysis also reveals that EMU countries and U.S. regions behave similarly along both these dimensions. Finally, in contrast to anecdotal evidence, the VAR analysis suggests that EMU countries adjust to idiosyncratic shocks at roughly the same speed as U.S. regions.

FIGURE 5

U.S. output response: Shock to U.S. federal funds rate
(percent)

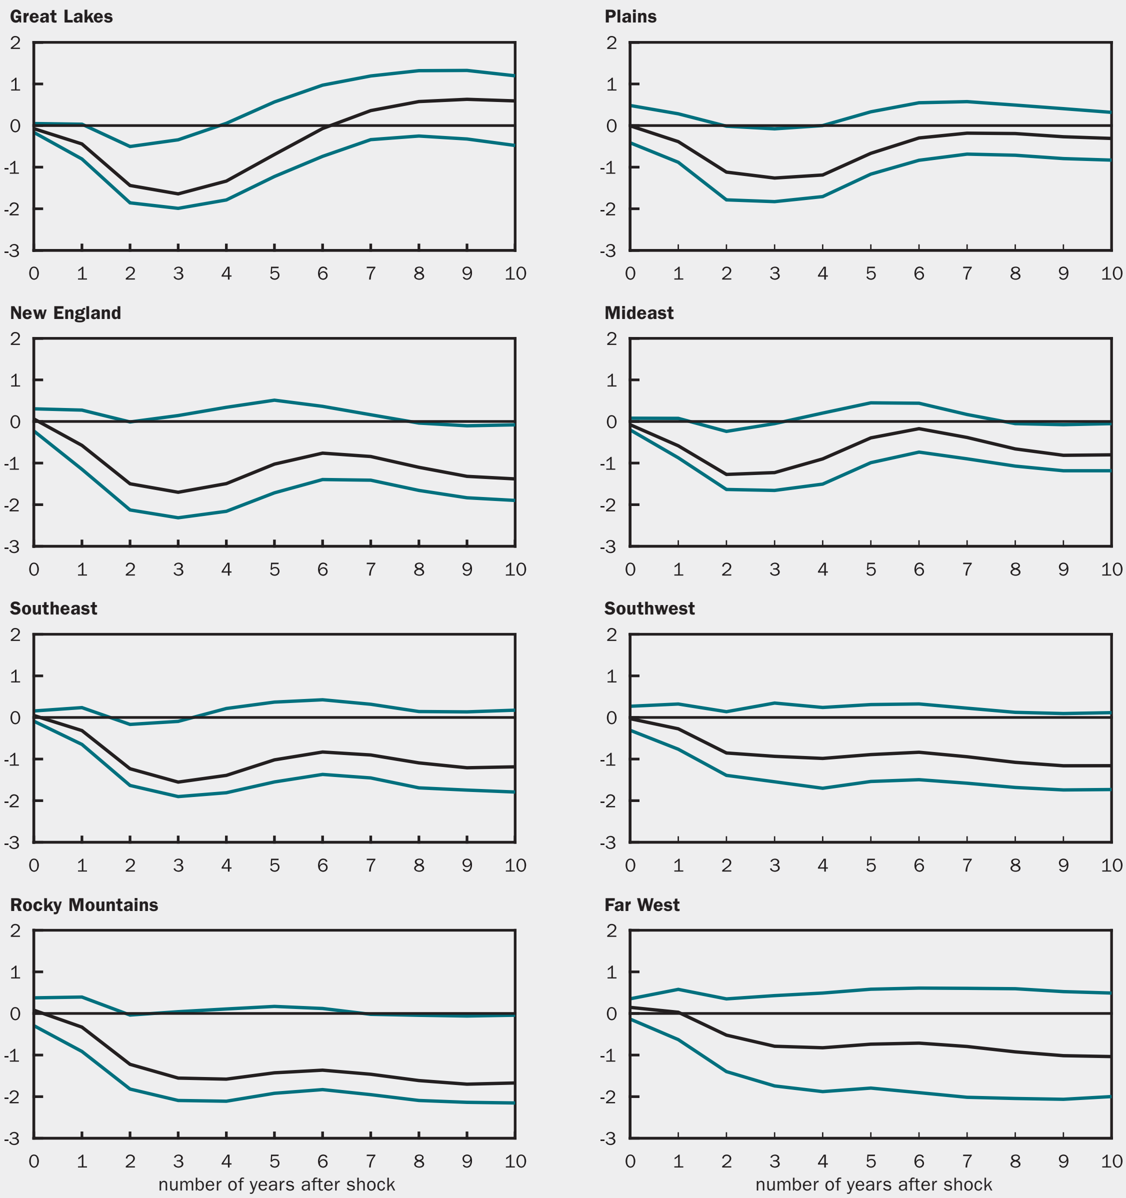

Notes and source: See figure 3.

Conclusion

The answer to the question of whether a currency union will be viable in the long run depends to a large extent on how far the union is from being an OCA. With this in mind, I assess the long-run viability of the EMU by comparing the EMU with a viable currency union (the U.S.) based on critical OCA criteria. My working hypothesis is that if the EMU is as close as the U.S. is to being an OCA, then there could be no presumption that the EMU would not be viable in the long run. Alternatively, if the EMU is much further from being an OCA than the U.S. is, then the adoption of a single currency could be problematic for some EMU countries and would call into question the viability of this monetary union. My analysis suggests that the behavior of countries at the center of the EMU is very similar to that of U.S. regions for all OCA criteria. In contrast, I find that

FIGURE 6

U.S. output response: Shock to U.S. regional income
(percent)

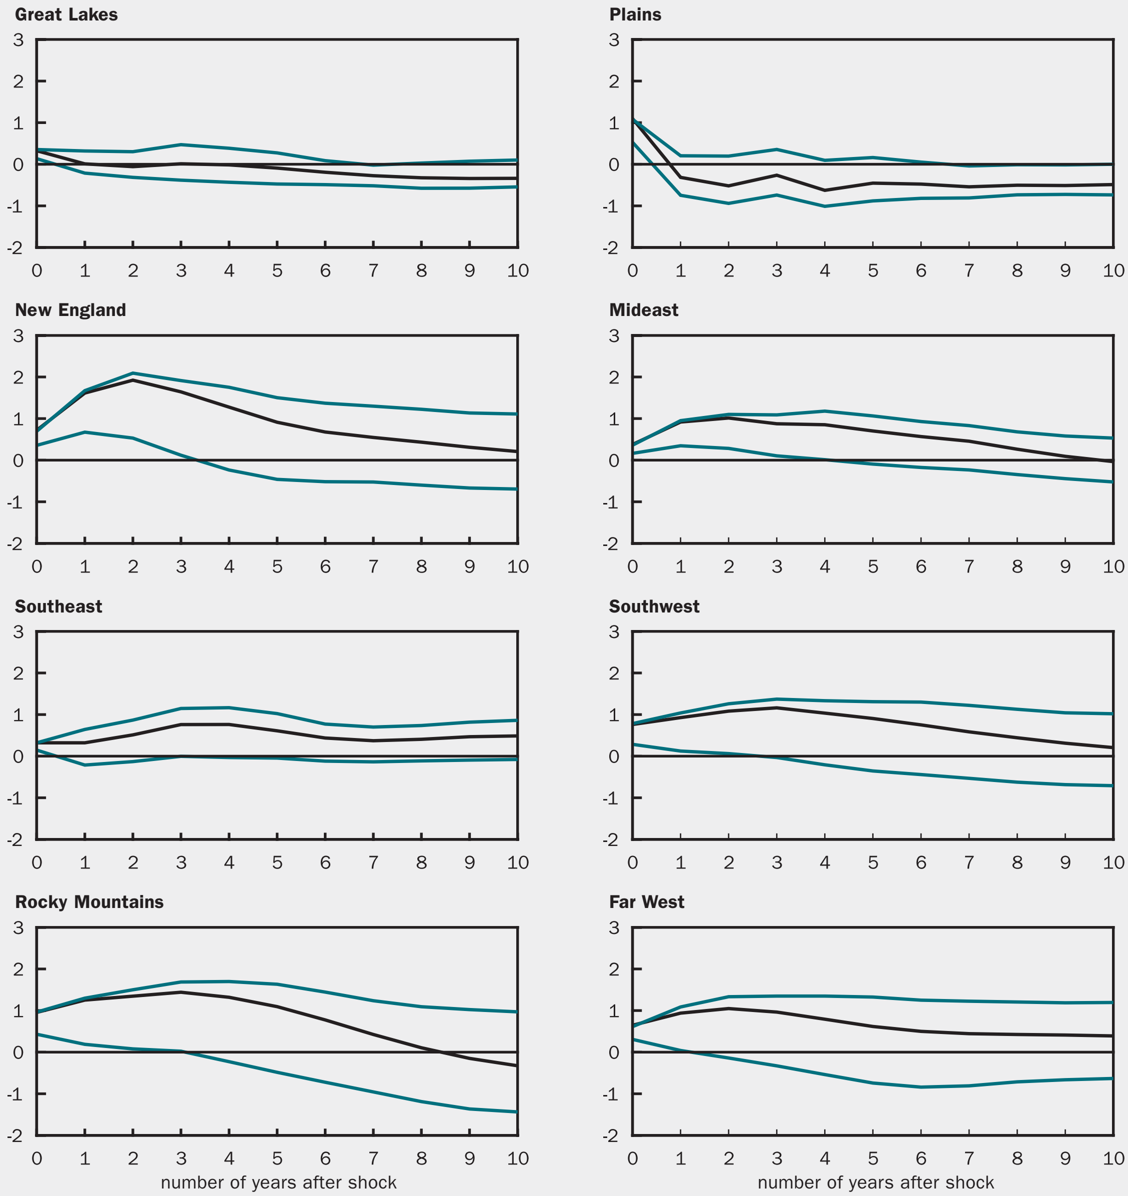

Notes and source: See figure 3.

FIGURE 7

EMU output response: Shock to world oil prices
(percent)

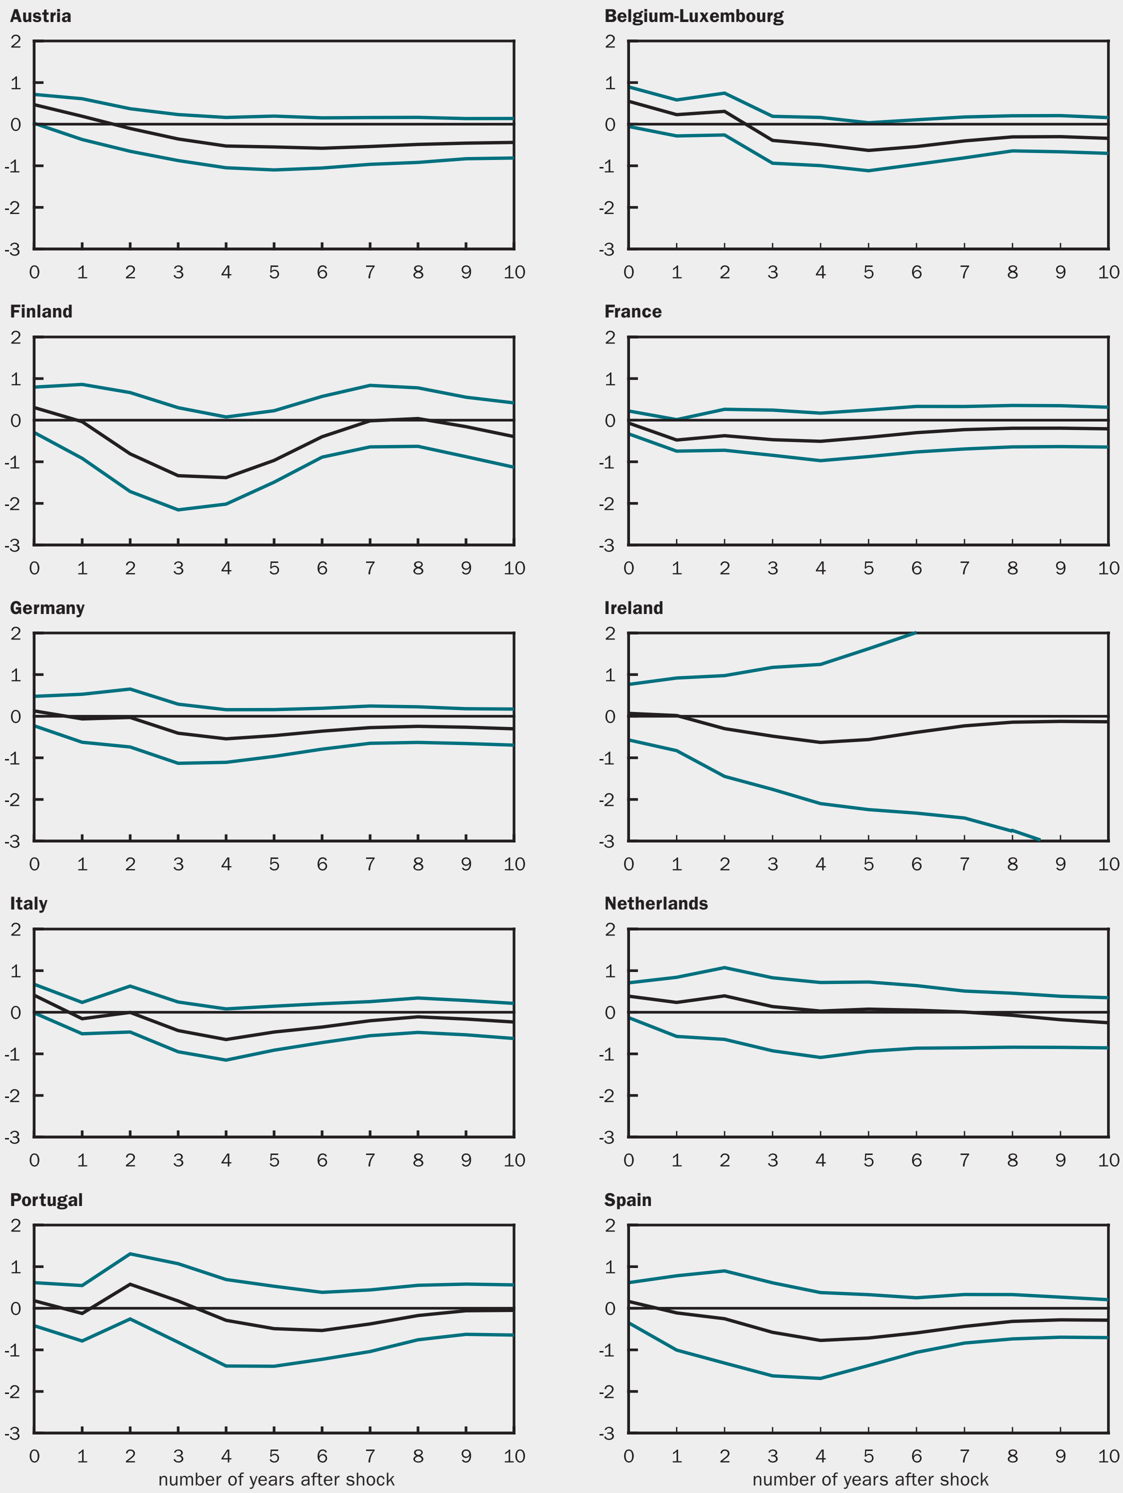

Notes: See figure 3.
Source: Calculations from author's statistical model, using the following annual data series:
International Monetary Fund—world crude oil prices, short-term interest rates, and gross domestic product.

FIGURE 8

EMU output response: Shock to aggregate EMU output
(percent)

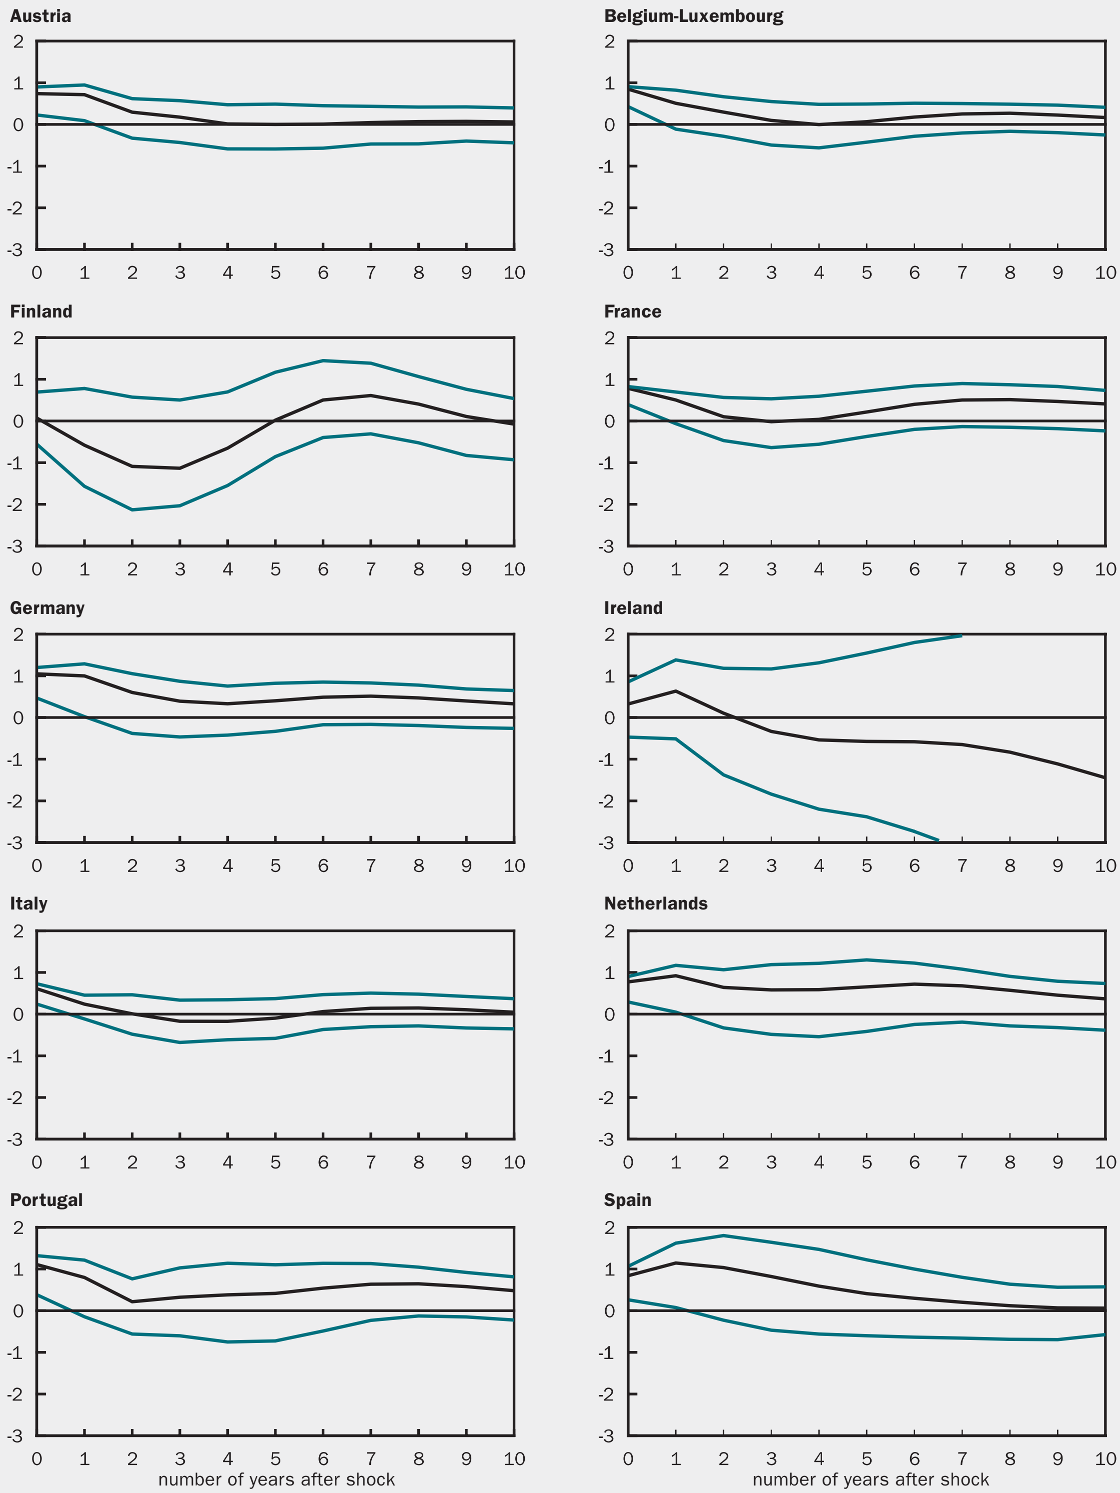

Notes: See figure 3.
Source: See figure 7.

FIGURE 9

EMU output response: Shock to EMU interest rates
(percent)

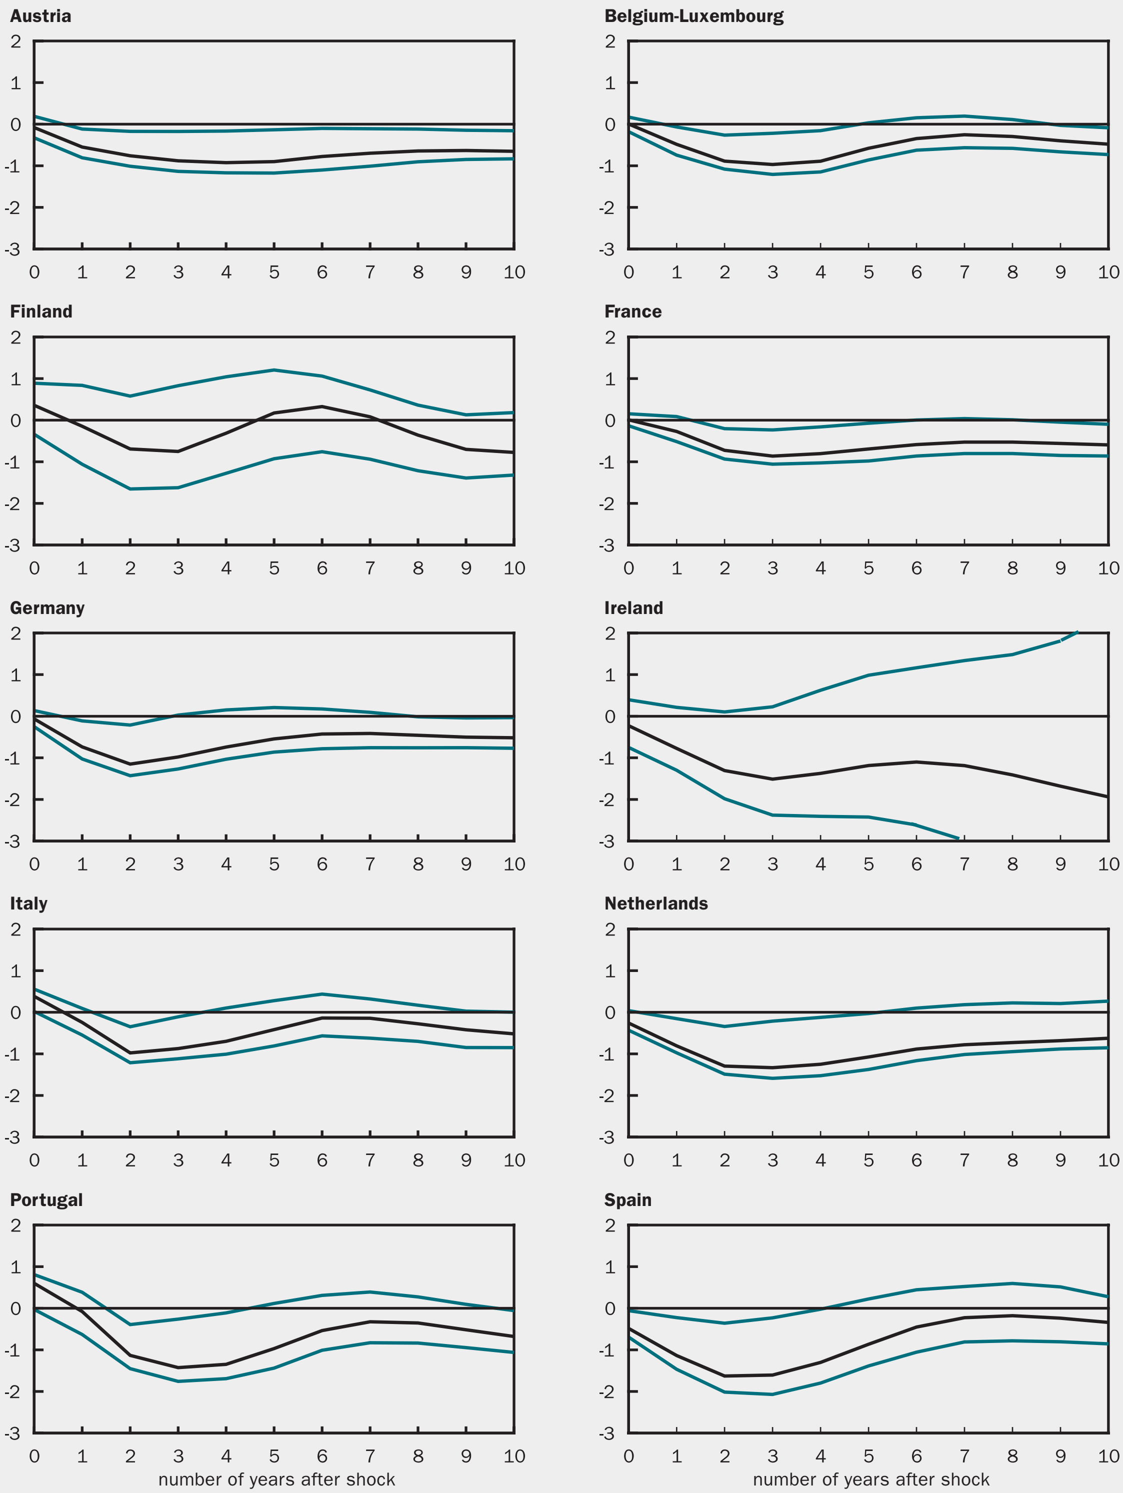

Notes: See figure 3.
Source: See figure 7.

FIGURE 10

EMU output response: Shock to EMU country income
(percent)

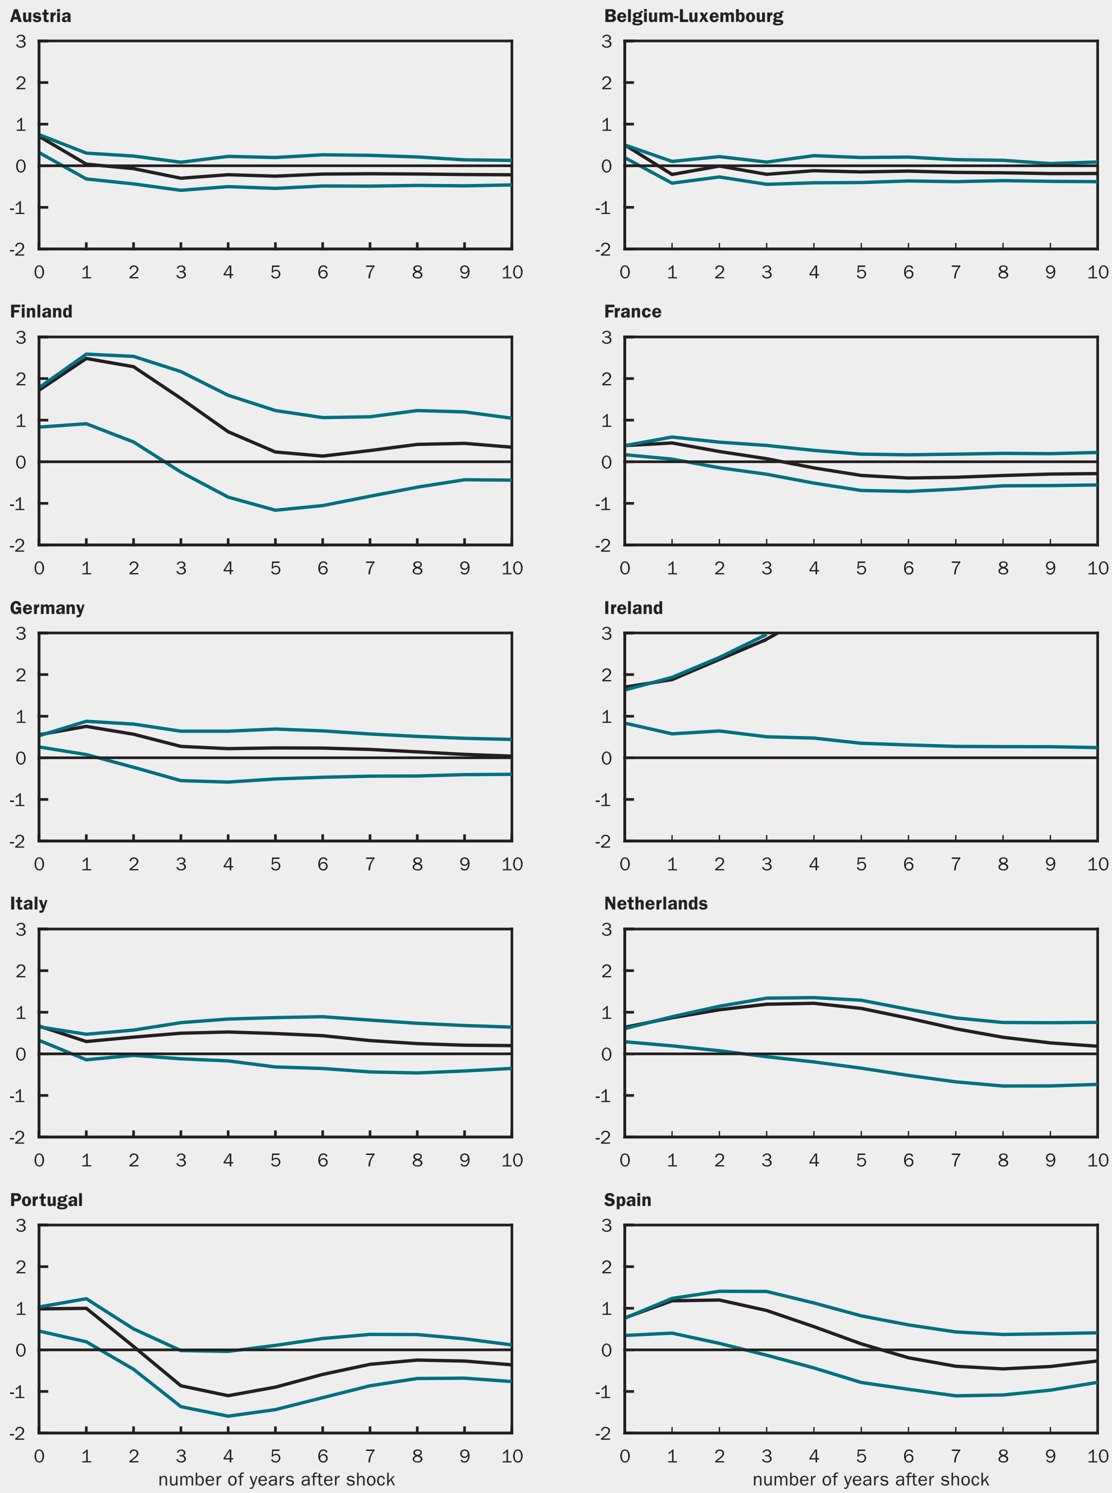

Notes: See figure 3.
Source: See figure 7.

countries in the periphery of the EMU, Finland and Ireland, are quite different from their EMU partners with regard to the OCA criteria. On the basis of this statistical analysis, I conclude that the EMU will likely be a viable currency union for the center countries, but question the viability of a union with countries in the periphery.

APPENDIX

A VAR analysis of regional business cycles

This appendix describes my methodology in greater technical detail. To isolate the various exogenous shocks, including monetary policy shocks, I use the vector autoregression (VAR) procedure developed by Christiano, Eichenbaum, and Evans (1994). Let Z_t denote the 4 × 1 vector of all variables in the model at date t. This vector includes changes in the log of world oil prices (*POIL*), log levels of aggregate U.S. (or euro-zone) income (*YA*), log levels of one of the eight U.S. regions (or 10 euro-zone countries) income (*YR*), and the level of the U.S. federal funds (or German overnight money market) rate (*R*), which I assume is the U.S. (or euro-zone) monetary policy indicator. The order of the variables is:

1) $Z_t = (POIL_t, YA_t, R_t, YR_t).$

I assume that Z_t follows a second-order VAR:

2) $Z_t = A_0 + A_1 Z_{t-1} + A_2 Z_{t-2} + u_t,$

where A_0, A_1, and A_2 are 4 × 4 coefficient matrices, and the 4 × 1 disturbance vector u_t is serially uncorrelated. I assume that the fundamental exogenous process that drives the economy is a 4 × 1 vector process $\{\varepsilon_t\}$ of serially uncorrelated shocks, with a covariance matrix equal to the identity matrix. The VAR disturbance vector u_t is a linear function of a vector ε_t of underlying economic shocks, as follows:

$$u_t = C\,\varepsilon_t,$$

where the 4 × 4 matrix C is the unique lower-triangular decomposition of the covariance matrix of u_t:

$$CC' = E[\,u_t u_t'\,].$$

This structure implies that the jth element of u_t is correlated with the first j elements of ε_t, but is orthogonal to the remaining elements of ε_t.

In setting policy, the U.S. Federal Reserve (or the euro-zone member central banks) both reacts to and affects the economy; I use the VAR structure to capture these cross-directional relationships. I assume that the feedback rule can be written as a linear function, Ψ, defined over a vector, Ω_t, of variables observed at or before date t. That is, if I let R_t denote the U.S. federal funds rate (or German overnight money market rate), then U.S. (or euro-zone) monetary policy is completely described by:

3) $R_t = \Psi(\Omega_t) + c_{3,3}\varepsilon_{3t},$

where ε_{3t} is the third element of the fundamental shock vector ε_t, and $c_{3,3}$ is the (3, 3) element of the matrix C. (Recall that R_t is the third element of Z_t.) In equation 3, $\Psi(\Omega_t)$ is the feedback-rule component of U.S. (or euro-zone) monetary policy, and $c_{3,3}\varepsilon_{3t}$ is the exogenous U.S. (or euro-zone) monetary policy shock. Since ε_{3t} has unit variance, $c_{3,3}$ is the standard deviation of this policy shock. Following Christiano, Eichenbaum, and Evans (1994), I model Ω_t as containing lagged values (dated $t-1$ and earlier) of *all* variables in the model, as well as time t values of those variables the monetary authority looks at contemporaneously in setting policy. In accordance with the assumptions of the feedback rule, an exogenous shock ε_{3t} to monetary policy cannot contemporaneously affect time t values of the elements of Ω_t. However, lagged values of ε_{3t} can affect the variables in Ω_t.

I incorporate equation 3 into the VAR structure described by equations 1 and 2. Variables *POIL* and *YA* are the contemporaneous inputs to the monetary feedback rule. These are the only components of Ω_t that are not determined prior to date t. With this structure, I can identify the right-hand side of equation 3 with the third equation in VAR equation 2: $\Psi(\Omega_t)$ equals the third row of $A_0 + A_1 Z_{t-1} + A_2 Z_{t-2}$, plus $\Sigma_{i=1}^{2} c_{3,i}\,\varepsilon_{it}$ (where $c_{3,i}$ denotes the (3, i) element of matrix C, and ε_{it} denotes the ith element of ε_t). Note that R_t is correlated with the first three elements of ε_t. By construction the shock $c_{3,3}\varepsilon_{3t}$ to U.S. (or euro-zone) monetary policy is uncorrelated with the monetary policy feedback rule Ω_t.

I estimate matrices A_0, A_1, A_2 and C by ordinary least squares. The response of any variable in Z_t to an impulse in any element of the fundamental shock vector ε_t can then be computed by using equations 1 and 2.

The standard error bounds in figures 3 through 10 are computed using the following bootstrap Monte

Carlo procedure. First, I construct 1,000 time series of the vector Z_t each of length T, where T denotes the number of observations in my data sample. Let $\{\xi_t\}_{t=1}^{T}$ denote the vector of residuals from the estimated VAR. I construct 1,000 sets of new time series of residuals, $\{\xi_t(j)\}_{t=1}^{T}$, $j = 1, ..., 1{,}000$. The tth element of $\{\xi_t(j)\}_{t=1}^{T}$ is selected by drawing randomly, with replacement, from the set of estimated residuals vectors $\{\xi_t\}_{t=1}^{T}$. For each $\{\xi_t(j)\}_{t=1}^{T}$, I construct a synthetic time series Z_t, denoted $\{Z_t(j)\}_{t=1}^{T}$, using the estimated VAR and the historical initial conditions on Z_t. Next, I reestimate the VAR using $\{Z_t(j)\}_{t=1}^{T}$ and the historical initial conditions and calculate the implied impulse response functions for $j = 1, ..., 1{,}000$. For each lag, I calculate the 25th lowest and 975th highest value of the corresponding impulse response coefficient across all 1,000 synthetic impulse response functions. The boundaries of the confidence intervals in the figures correspond to a plot of these coefficients.

NOTES

[1]See Corden (1993), chapters 7–9, for an extended discussion of the EMS and events surrounding the 1992 breakdown of the system.

[2]In general, time-series data are nonstationary. Nonstationary data do not have well-defined standard deviations or correlations. One way of overcoming this problem is to filter the data using a filter that removes the nonstationary components and renders the data stationary. There is a range of filtering techniques available, including linear time trends and first differencing. Baxter and King (1995) have designed a filter that isolates components of the data that policy analysts are interested in, the so-called business cycle frequencies of one and a half to eight years. I use a Baxter–King filter to isolate cyclical movements in U.S. and EMU time series.

[3]Consumer price indexes do exist for metropolitan areas in the various BEA regions. However, there is a very high degree of correlation in consumer price fluctuations across these metropolitan areas. In addition, using region-specific price series would impose a further limit on the analysis since many metropolitan indexes are not available after 1986.

[4]The gross product by state is available from 1977 to 1997.

[5]See Carlino and DeFina (1998a), appendix A, for a listing of states by BEA region.

[6]For examples, see references in Kouparitsas (1998).

[7]Carlino and DeFina (1998a) assume a similar recursive information ordering in their analysis of the regional impact of U.S. monetary policy.

REFERENCES

Baxter, M., and R. G. King, 1995, "Measuring business cycles: Approximate band-pass filters for economic time series," National Bureau of Economic Research, working paper, No. 5022.

Blanchard, O., and D. Quah, 1989, "The dynamic effects of aggregate demand and supply disturbances," *American Economic Review*, Vol. 79, No. 4, pp. 655–673.

Carlino, G. A., and R. DeFina, 1998a, "The differential regional effects of monetary policy," *The Review of Economics and Statistics*, Vol. 80, No. 4, pp. 572–587.

__________, 1998b, "Monetary policy and the U.S. states and regions: Some implications for European monetary union," Federal Reserve Bank of Philadelphia, working paper, No. 98-17.

Christiano, L. J., M. Eichenbaum, and C. L. Evans, 1994, "Identification and the effects of monetary policy shocks," Federal Reserve Bank of Chicago, working paper, No. 94-7.

Corden, W. M., 1993, *Economic Policy, Exchange Rates, and the International System*, Chicago: University of Chicago Press.

Dornbusch, R., C. A. Favero, and F. Giavazzi, 1998, "The immediate challenges for the European Central Bank," National Bureau of Economic Research, working paper, No. 6369.

Eichengreen, B. J., 1992, "Is Europe an optimum currency area?," reprinted in *European Monetary Unification: Theory, Practice, and Analysis*, B. J. Eichengreen (ed.), Cambridge: MIT Press, pp. 51–71.

Eichengreen, B. J., and T. Bayoumi, 1993, "Shocking aspects of European monetary unification," reprinted in *European Monetary Unification: Theory, Practice, and Analysis*, B. J. Eichengreen (ed.), Cambridge: MIT Press, pp. 73–109.

Frankel, J. A., and A. K. Rose, 1998, "The endogeneity of the optimum currency area criteria," *Economic Journal*, Vol. 108, July, pp. 1009–1025.

Kouparitsas, M. A., 1998, "Are international business cycles different under fixed and flexible exchange rate regimes?" *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 22, No. 1, pp. 46–64.

Mundell, R. A., 1961, "A theory of optimum currency areas," *American Economic Review*, Vol. 51, No. 4, pp. 657–665.

Sims, C., 1972, "Money, income, and causality," *American Economic Review*, Vol. 62, No. 4, pp. 540–552.

Regional employment growth and the business cycle

Ellen R. Rissman

Introduction and summary

The purpose of this article is to study the sources of regional employment fluctuations in the U.S. and to shed light on the interactions of these regional fluctuations with the aggregate economy. Many studies of regional employment growth have analyzed the effect of regional differences in a number of underlying factors, such as local government expenditures and tax policy, while controlling for aggregate economic activity. My analysis focuses alternatively on the role of regional fluctuations in determining aggregate economic activity.

Macroeconomists have tended to concentrate on the impact of changes in aggregate factors in determining the business cycle.[1] Such aggregate factors have included, for example, fiscal and monetary policy, the role of consumer confidence, aggregate supply and demand, and productivity. Yet there is a growing literature that suggests that aggregate disturbances are the result of a variety of influences.[2] In the work introduced here, I explicitly consider the role of regional employment fluctuations in determining the business cycle. I do not specifically identify the sources of such regional shocks. They could be the result of changing federal governmental policies, for example, immigration or defense spending, that impinge upon certain areas of the country more than others. They could also reflect changes in local welfare programs or shifts in local fiscal and tax policy.

The analysis is complicated by the fact that while regional fluctuations may have aggregate repercussions, aggregate factors influence regional growth as well. For example, general productivity shocks are likely to have broad consequences across a variety of industries and geographical areas that are reflected in regional employment growth. Ascertaining what movements in employment growth are common across regions and what are region-specific would be helpful for policymakers. If, for example, regional employment growth is largely unrelated to employment growth in other regions, a more regional policy focus might be appropriate. Examples of more localized policy would include differential taxation and spending programs that are coordinated within a region or a more geographically targeted approach to federal government spending. If, however, most regional employment growth is common across regions, a more centralized policy process is warranted.

The business cycle has been conceptualized as "expansions occurring at about the same time in many economic activities, followed by similarly general recessions, contractions, and revivals which merge into the expansion phase of the next cycle."[3] Thus, the business cycle is characterized by comovements among a variety of economic variables and is observable only indirectly. Only by monitoring the behavior of many economic variables simultaneously can one quantify the business cycle. For example, recessions are typically associated with declining output and employment across broadly defined industries. It is this notion of comovement that has supplied the foundation for measuring cyclical activity. This is the practice behind the widely publicized National Bureau of Economic Research's (NBER) dating of business cycles and Stock and Watson's (1988) index of coincident economic indicators.

Ellen R. Rissman is an economist in the Economic Research Department of the Federal Reserve Bank of Chicago. *The author would like to thank Ken Housinger for his research assistance. She is particularly indebted to Ken Kuttner for his insight and for providing the basic statistical programs. Dan Sullivan, David Marshall, Joe Altonji, and Bill Testa provided many thoughtful comments. The author would also like to thank the seminar participants at the Federal Reserve Bank of Chicago for their patience and suggestions.*

While most analyses of the business cycle focus on the notion of comovement in employment or output across industries, a great deal of comovement exists across geographical regions as well. Yet, until recently this regional cyclicality has gone largely unexplored, with a few notable exceptions such as Altonji and Ham (1990), Blanchard and Katz (1992), Clark (1998), and Clark and Shin (1999). The reason for the lack of interest in the regional cycle has largely been the belief that whatever cyclicality a geographical region experiences is due in large part to its industrial mix and to common aggregate shocks. In fact, regional shocks are typically not considered in assessing the business cycle.

Altonji and Ham (1990) investigate the effect of U.S., Canadian national, and sectoral shocks on Canadian employment fluctuations at the national, industrial, and provincial level. They find that sectoral shocks account for only one-tenth of aggregate variation, with two-thirds of the variation attributable to U.S. disturbances and one-quarter to Canadian shocks. The relatively small importance of sectoral fluctuations in describing aggregate variation in Canadian data suggests that regional shocks have little effect on the business cycle. The conclusion holds true for Canada but the study does not necessarily apply to the U.S. economy, in which external shocks presumably play less of a role.

In a model similar to Altonji and Ham (1990), Clark (1998) attempts to quantify the roles of national, regional, and industry-specific shocks on regional employment growth for U.S. data. Contrary to the traditional view that regional fluctuations are unimportant in determining the aggregate and the results of Altonji and Ham (1990) for Canada, Clark finds that "roughly 40 percent of the variance of the cyclical innovation in any region's employment growth rate is particular to that region."[4] He goes on to show that these regional shocks tend to propagate across regions. Clark's conclusion is that heterogeneous regional fluctuations have possibly important implications for business cycle study. Although valuable, the methodology he employs does not permit the construction of actual estimates of regional disturbances, which hampers his ability to clarify the underlying causes of the regional shocks.

In this article, I develop and estimate a model of regional employment growth aimed at understanding the role of the aggregate economy. Each region's employment growth is assumed to depend upon a common factor, thought of here as the business cycle.[5] This common factor is not directly observable, but is inferred through the comovements of employment growth across a number of regions simultaneously. This does not mean that each region responds in the same manner to cyclical fluctuations. Some regions will be more cyclically sensitive while others are less. Accordingly, the methodology permits the cycle to have a differential impact on regional employment growth.

The methodology I employ is similar to that in Rissman (1997) and utilizes a statistical technique known as the *Kalman filter*. The research here is akin to Clark's in that it is an attempt to isolate the effects of the business cycle and regional disturbances on regional employment growth. However, I expressly model the business cycle as a common factor affecting all regions and some more than others. A measure of the business cycle develops naturally from the estimation of the model and is based solely upon the comovements in employment growth across census regions. In addition, I estimate regional employment shocks, which are useful for elucidating the reasons behind regional differences in economic growth.

In summary, while aggregate fluctuations are an important force behind regional employment growth, local disturbances contribute significantly as well. The role of such local shocks is not uniform across regions. My estimates indicate that almost 60 percent of the steady state variance in employment growth in the West South Central region is attributable to local fluctuations. This compares with only about 10 percent in East South Central, where aggregate conditions are the driving force.

My results suggest that regional employment growth can be described remarkably well by a simple model in which a common business cycle has a differential impact upon the various regions. Measures of the business cycle from this approach are quite consistent across models and agree quite well with more typical measures of the business cycle. The main difference between this measure and other such measures is that this one relies upon regional employment data alone, while other measures may take into consideration a wide variety of other factors, such as productivity.

Interestingly, errors made in forecasting employment growth in the West South Central region appear to have some predictive content for forecasting employment growth in most other regions. This suggests that there is something unique about this region's economy that is not currently captured by the model but that does have aggregate repercussions. This might be due to the region's reliance on the oil industry. My analysis implies that regional policies may be an important tool in managing the economy. However, more research on the nature of the spillovers across regions would be required to support economic policy targeting specific regions.

Data

In formulating a model of regional employment growth, a necessary first step is to observe the patterns in the data. The Bureau of Labor Statistics (BLS) collects regional employment statistics from its *Employment Survey* for the following nine census regions: New England, Mid-Atlantic, East North Central, West North Central, South Atlantic, East South Central, West South Central, Mountain, and Pacific.[6] Figure 1 shows annualized quarterly employment growth for each of the nine census regions from 1961:Q1 to 1998:Q2. (The construction is explained in box 1.) It is clear from the figure that some regions consistently exhibit high employment growth (for example, South Atlantic, East South Central, West South Central, Mountain, and Pacific), while other regions consistently exhibit below-average employment growth (New England, Mid-Atlantic, East North Central, and West North Central).[7]

In addition to differences in mean employment growth, regional employment growth exhibits an apparent cyclical pattern. Typically, employment growth declines during a recession (shaded areas in figure 1) and increases in an expansion.[8] This cyclical pattern shows up quite clearly in all regions but is less pronounced in some. Specifically, the Pacific and Mountain states appear to be less affected by the business cycle than a more typical Rust Belt region such as East North Central. This is not to say that employment growth does not decline here as well, but in these regions contractions are associated with smaller declines.

Closer inspection of figure 1 shows that regional employment growth appears to have a random component in addition to a cyclical one. For example, the West South Central region experienced a marked decline in employment growth in the mid-1980s. This decline was echoed in a few other regions, but was nowhere as pronounced as in West South Central. In fact, regions such as the Mid-Atlantic, East North Central, South Atlantic, and Pacific experienced relatively little negative impact at that time.

In modeling the effect of the business cycle on regional employment growth it is useful to know how the business cycle affects the regional economy through other less-direct avenues. For example, the cycle may affect the distribution of employment across regions. Figure 2 exhibits regional employment growth net of aggregate employment growth. A negative number for a region indicates that that region's employment share of the aggregate is shrinking. Conversely, a positive number shows that that region's employment is growing relative to the aggregate. The figure shows that trends in employment growth seem to persist for long periods. For example, the Rust Belt New England region experienced below national average employment growth for most of the earlier part of the data period. This decline was temporarily reversed in the 1980s—the much-vaunted "Massachusetts miracle." However, the New England recovery was short-lived, as shown by the subsequent pronounced decline in New England's employment share. The Mid-Atlantic states lost ground as well over most of the period. In contrast, employment growth in the Mountain states was above the national average, with the exception of a brief period in the mid-1960s and again in the mid-1980s.

The employment shares in figure 2 do not appear, at least by casual observation, to behave cyclically. It is *not* the case that a given region's relative importance in the composition of aggregate employment is affected systematically by the business cycle. This is in direct contrast to the evidence on industries, where the composition of total employment shifts away from goods-producing and toward service-producing industries during contractions. Although regions show periods of expansion and contraction, at first blush the timing of these "regional cycles" is unlike the timing of the familiar business cycle. If a business cycle is described by comovements in a number of series, it is difficult to describe what these comovements might be from looking at net regional employment growth alone.

At times, statistical relationships can be difficult to ascertain by casual observation of the data at hand. To investigate a more complex model of the cyclicality of net regional employment growth, I perform a

BOX 1

Annual employment growth and net annual employment growth

Employment growth in region i at time t, y_{it}, is calculated as:

$$y_{it} \equiv \log(e_{it} / e_{it-4}) \times 100,$$

where e_{it} is employment in region i at time t. Define net employment growth n_{it} as the difference between regional employment growth and aggregate employment growth. Specifically,

$$n_{it} \equiv y_{it} - y_t$$
$$n_{it} = [\log(e_{it} / e_{it-4}) - \log(e_t / e_{t-4})],$$

where e_t is defined as aggregate employment at time t and y_t is aggregate employment growth.

FIGURE 1

Regional employment growth, 1961:Q1–98:Q2
(percent)

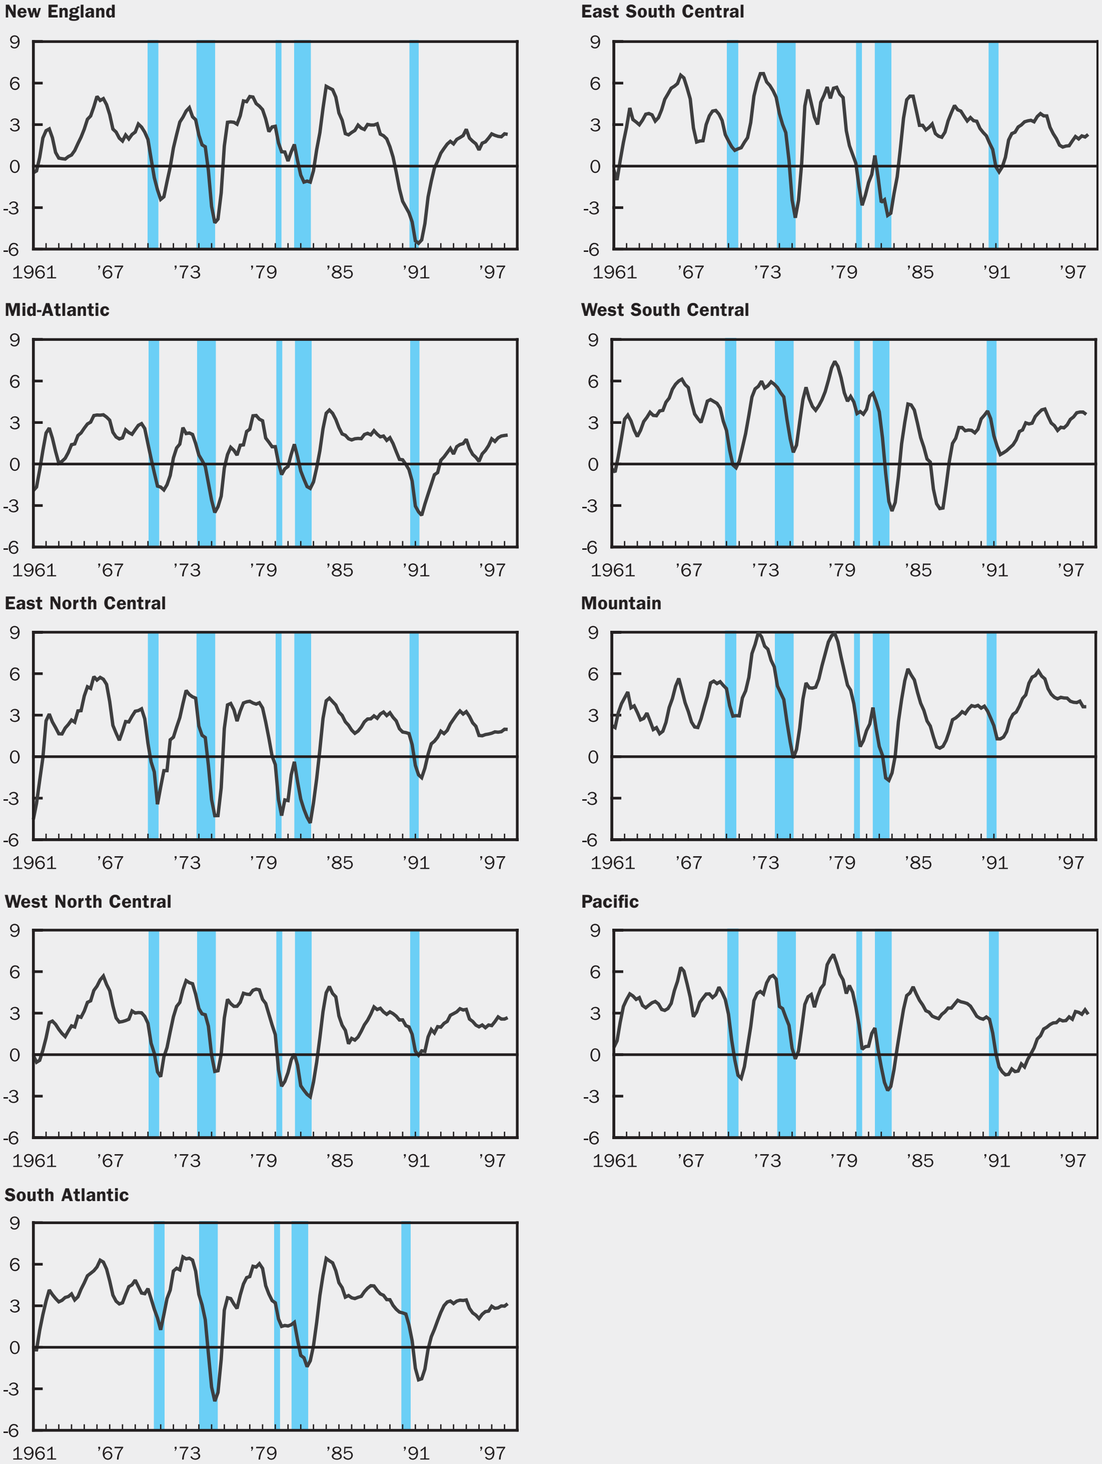

Notes: See box 1 for details of calculations. Shaded areas indicate recessions, as defined by the National Bureau of Economic Research.
Source: U.S. Department of Labor, Bureau of Labor Statistics, 1960–98, employment database available at ftp://ftp.bls.gov/pub/time.series and author's calculations.

FIGURE 2

Employment growth, regional less aggregate, 1961:Q1–98:Q2

(percent)

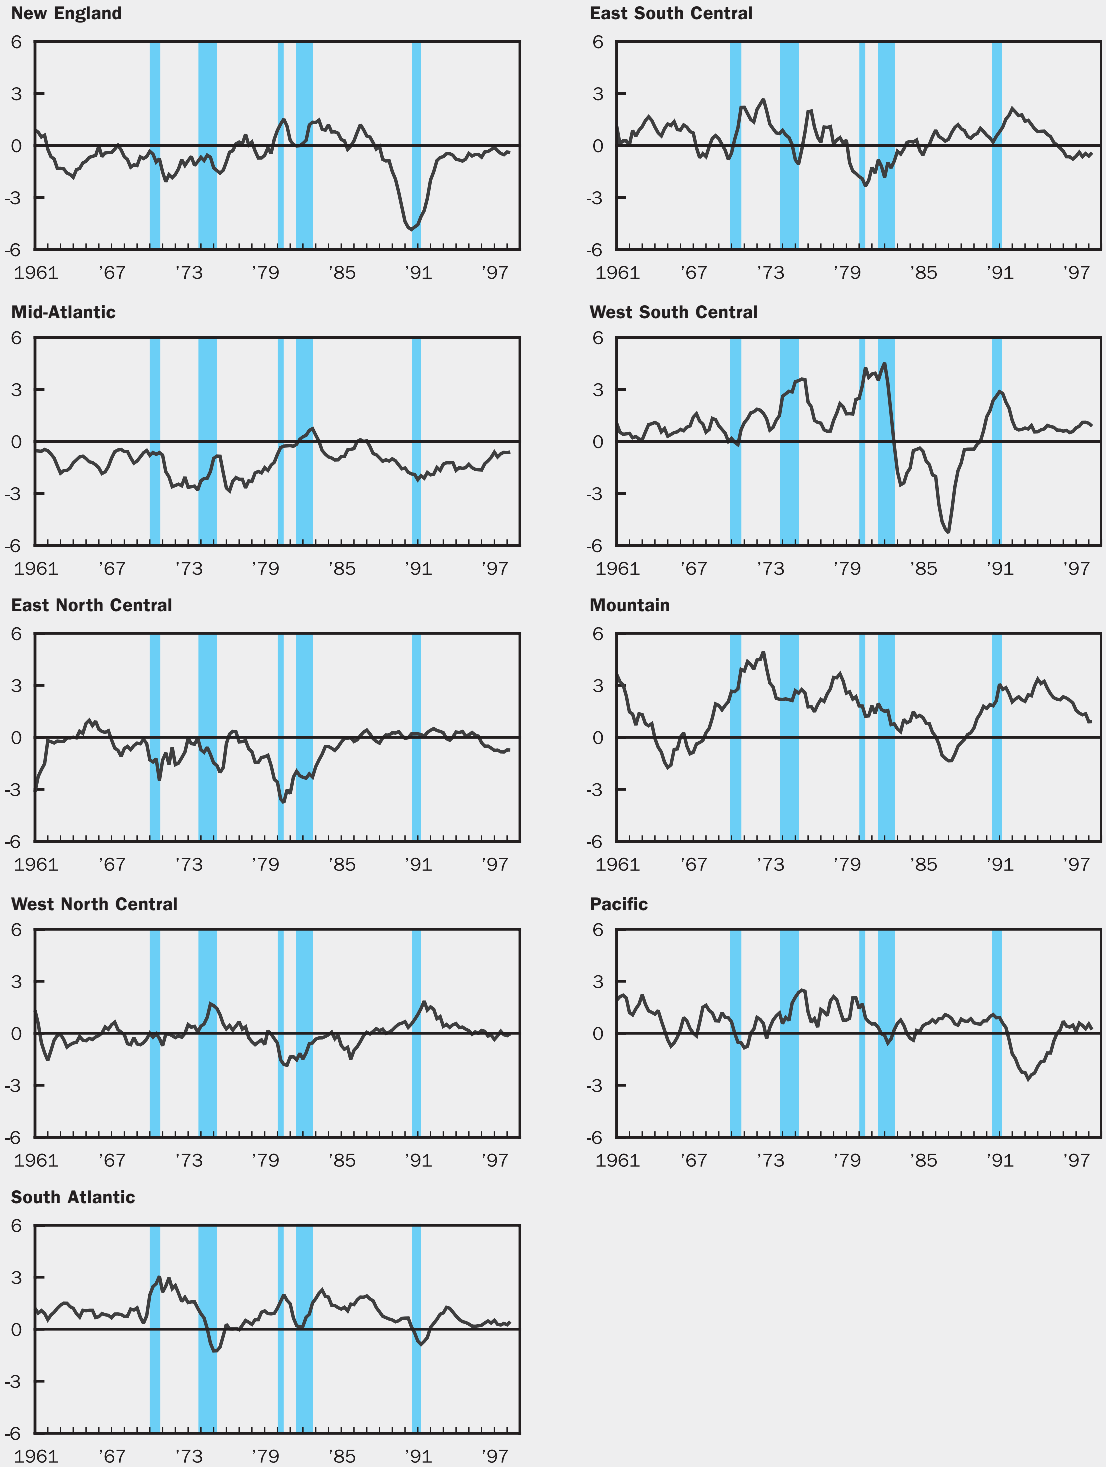

Notes: Regional employment growth less total employment growth, quarterly from previous year. Growth rates measured as four-quarter log differences. See box 1 for details. Shaded areas indicate recessions, as defined by the National Bureau of Economic Research.

Source: U.S. Department of Labor, Bureau of Labor Statistics, 1960–98, employment database available at ftp://ftp.bls.gov/pub/time.series and author's calculations.

regression exercise in which net regional employment growth is assumed to depend upon lags of net regional employment growth and whether the economy is in a contraction as defined by the NBER. (The form of the regression is shown in box 2.) Table 1 shows the results of these simple ordinary least squares (OLS) regressions. A significant negative or positive number in the *CONTRACT* column indicates that, even after accounting for dynamics through lags of own-region net employment growth, the state of the aggregate economy has an additional impact upon net employment growth. In the case of a negative number, the region's employment share shrinks during a contraction. Conversely, a positive number suggests that the region's employment share expands during a contraction.

From table 1, clearly business cycle contractions as defined by the NBER are not particularly good at explaining regional net employment growth after accounting for serial correlation in the dependent variable. Most of the estimates are not significantly different from zero. The exceptions are Mid-Atlantic, East North Central, and Mountain. In the East North Central region, comprising Ohio, Indiana, Michigan, Wisconsin, and Illinois, employment shares typically decline in a recession. Furthermore, the estimated effect for East North Central is quite large compared with the other regions. In the Mid-Atlantic and Mountain regions, employment shares tend to rise during a contraction. The $\bar{R}^2$ statistic is a measure of the *fit* of the regression. The closer this number is to unity, the better the data fit the estimated equation. The high values of $\bar{R}^2$ suggest that most of the variation in net regional employment growth is accounted for by lags in the dependent variable.

Industry effects

To summarize, the data on regional employment growth suggest that the business cycle affects regional employment growth directly and to a far lesser extent through its effect on the distribution of employment across regions. It has long been observed that the business cycle systematically affects the distribution of employment across industries.[9] One possible explanation for the cyclicality of regional employment growth is that certain regions are dominated by specific industries. To the extent that this is true, then the regional cycles found in employment growth merely mirror the effects of the business cycle on the regional industry mix and, thus, there is relatively little role for regional fluctuations or shocks to explain the patterns in the data. Box 3 shows how state industry employment data can be used to evaluate this issue.

Changes in state employment are dominated by two effects. First, there is the effect of shifting industry employment on employment within the state, holding the contribution of the state in employment within the industry constant. The second effect measures the importance of shifting the state's contribution to each industry, holding aggregate industry employment

BOX 2

OLS regression testing effect of contractions on net employment growth

Let *CONTRACT* be a dummy variable taking on the value 1 during an NBER contraction and 0 elsewhere. The OLS regression equation is of the form:

$$n_{it} = c + a(L)n_{it-1} + b * CONTRACT + \varepsilon_{it}.$$

Four lags of the dependent variable have been included and are generally enough to ensure that the error term is serially uncorrelated.

TABLE 1

Effect of timing of NBER contractions on regional employment growth less aggregate employment growth, OLS

Region	*CONTRACT*	$\bar{R}^2$
New England	0.1170	0.9281
Mid-Atlantic	0.1345**	0.8851
East North Central	–0.3581***	0.8541
West North Central	0.0272	0.8207
South Atlantic	0.0540	0.8686
East South Central	–0.0458	0.8393
West South Central	0.1444	0.9394
Mountain	0.1561*	0.9267
Pacific	–0.0308	0.8436

Notes: The regression equation estimated by OLS is:

$$n_{it} = c + a(L)n_{it-1} + b * CONTRACT + \varepsilon_{it},$$

where *CONTRACT* takes on the value of 1 during an NBER contraction and is 0 otherwise; *a*(*L*) is a polynomial in the lag operator with a maximum lag length of four. ***Indicates significance at the 1 percent level; **indicates significance at the 5 percent level; and *indicates significance at the 10 percent level.

Source: Author's calculations based on data from the U.S. Department of Labor, Bureau of Labor Statistics, database at ftp://ftp.bls.gov/pub/time.series and the National Bureau of Economic Research database available on the Internet at www.nber.org.

constant. The first effect can be thought of as an industry effect while the second can be thought of as a state effect. If state effects are not important, then an analysis of employment growth by geographical region is unlikely to yield any insight into business cycles. If, however, a significant portion of the change in employment within a state is state-specific, a regional analysis is likely to provide further information.

Table 2 shows the relative importance of each of these two factors for all states except Hawaii. Specifically, the table shows the portion of the normalized change between 1985:Q1 and 1998:Q2 in employment in state s attributable to changing industry employment and changing employment shares, respectively.[10] The industry categories are mining, construction, manufacturing, trade, services, transportation and public utilities, government, and finance, insurance, and real estate. The goal is to analyze how important state and industry effects are in explaining state employment changes. A full set of data on all states with the exception of Hawaii is available from 1982:Q1 forward. To avoid evaluating employment over two

BOX 3

Effect of industry composition on state employment

Define $e_i^s(t)$ as employment in industry i in state s at time t. Define

$$k_i^s(t) \equiv \frac{e_i^s(t)}{e_i(t)}$$

as the share of industry i's employment in state s. These numbers sum to unity over all states. The larger the share in a given state, the more important that state is in the employment of that particular industry. Employment in state s at time t, $e^s(t)$, can be calculated as:

$$e^s(t) = \sum_i k_i^s(t) e_i(t),$$

which says that total state employment is the sum of employment in each industry within that state.

Now define the difference operator Δ^τ as:

$$\Delta^\tau x(t) \equiv x(t) - x(t-\tau).$$

Applying the difference operator to the expression for state employment yields:

$$\Delta^\tau e^s(t) = \sum_i \Delta^\tau e_i(t) k_i^s(t) + \sum_i \Delta^\tau k_i^s(t) e_i(t) - \sum_i \Delta^\tau k_i^s(t) \Delta^\tau e_i(t).$$

From this expression, the change in state employment between periods t–τ and t can be separated into three different effects. The first term to the right of the equal sign reflects the effect of changing industry employment while keeping the share of industry i's employment in state s constant. An example will help clarify this construct. Suppose aggregate manufacturing employment declines, this effect calculates the effect of declining aggregate manufacturing employment on employment within a given state, holding the share of that state's contribution to total manufacturing employment constant. No secondary effects are permitted whereby the distribution of manufacturing across states has been altered.

The second term captures the effect of changing employment shares in industry i in state s while keeping total industry employment constant. Suppose that employment remains constant over time but that the importance of a given state in its contribution to the total changes. This second term calculates the effect of this shift on employment within that state. Finally, the third term is an interaction term that permits both state industry employment shares and industry employment to vary together. Because it is calculated by multiplying together two changes, it is smaller in magnitude than the first two effects and will be dominated by the first two terms in the expression.

Rearranging terms,

$$\Delta^\tau e^s(t) + \sum_i \Delta^\tau k_i^s(t) \Delta^\tau e_i(t) = \sum_i \Delta^\tau e_i(t) k_i^s(t) + \sum_i \Delta^\tau k_i^s(t) e_i(t)$$

or

$$1 = \frac{\sum_i \Delta^\tau e_i(t) k_i^s(t) + \sum_i \Delta^\tau k_i^s(t) e_i(t)}{\Delta^\tau e^s(t) + \sum_i \Delta^\tau k_i^s(t) \Delta^\tau e_i(t)}.$$

This expression says that the normalized sum of the two effects should be unity.

TABLE 2

Changes in employment in state *s*, 1985:Q1–98:Q2

	Industry effect	State effect
Alabama	0.80	0.20
Alaska	1.20	–0.20
Arizona	0.58	0.42
Arkansas	0.63	0.37
California	1.20	–0.20
Colorado	0.76	0.24
Connecticut	7.74	–6.74
Delaware	0.74	0.26
Florida	0.73	0.27
Georgia	0.65	0.35
Hawaii	n.a.	n.a.
Idaho	0.57	0.43
Illinois	1.31	–0.31
Indiana	0.76	0.24
Iowa	0.84	0.16
Kansas	0.82	0.18
Kentucky	0.67	0.33
Louisiana	1.78	–0.78
Maine	1.07	–0.07
Maryland	1.48	–0.48
Massachusetts	3.85	–2.85
Michigan	0.91	0.09
Minnesota	0.80	0.20
Mississippi	0.73	0.27
Missouri	0.99	0.01
Montana	0.91	0.09
Nebraska	0.89	0.11
Nevada	0.48	0.52
New Hampshire	1.06	–0.06
New Jersey	2.54	–1.54
New Mexico	0.81	0.19
New York	36.32	–35.33
North Carolina	0.63	0.36
North Dakota	1.15	–0.15
Ohio	1.03	–0.03
Oklahoma	1.34	–0.34
Oregon	0.60	0.40
Pennsylvania	1.87	–0.87
Rhode Island	4.89	–3.89
South Carolina	0.67	0.33
South Dakota	0.70	0.30
Tennessee	0.67	0.33
Texas	0.89	0.11
Utah	0.53	0.47
Vermont	1.15	–0.15
Virginia	0.82	0.18
Washington	0.61	0.39
West Virginia	1.11	–0.11
Wisconsin	0.73	0.27
Wyoming	1.76	–0.75

Notes: See box 3 for the exact calculations. n.a. indicates not available.
Source: Author's calculations based on data from the U.S. Department of Labor, Bureau of Labor Statistics, database at ftp://ftp.bls.gov/pub/time.series.

different phases of the business cycle, I analyze changes in state employment between 1985:Q1 and 1998:Q2.

The evidence provided in table 2 supports Clark's (1998) contention that location-specific shocks are important. For example, about 58 percent of the increase in employment in Arizona is attributable to within-industry employment growth. However, the remaining 42 percent of the increase is the result of a shifting industrial mix within the state. Although the effect of changing aggregate industrial employment dominates, the importance of the changing industrial composition within the state is not insignificant in most instances, most often leading to increases in state employment.

Some states, notably Alaska, California, Connecticut, Illinois, Louisiana, Maryland, Massachusetts, New Jersey, New York, North Dakota, Oklahoma, Pennsylvania, Rhode Island, Vermont, West Virginia, and Wyoming, would have experienced an even larger increase in employment between 1985:Q1 and 1998:Q2 except that employment shares shifted adversely. New York appears to be somewhat of an outlier with employment gains being offset to a large extent by shifts in employment shares: Manufacturing employment as a share of total state employment fell precipitously, while employment in finance, insurance, and real estate grew quickly.

The state industry employment data suggest that employment growth is only partly explained by industry effects and that a good portion of state employment changes results from location-specific factors. It follows that changes in local employment do not simply reflect the local industrial mix, but also have a significant location-specific component. This adds another dimension to our understanding of regional employment growth.

The model

The evidence above indicates that regional employment growth is driven in large part by a common business cycle. Furthermore, regional shocks are important even after accounting for changing aggregate industrial composition. Let annual employment growth in region *i*, y_{it}, have the following specification:

$$y_{it} = \alpha_i + \beta_0{}^i C_t + \beta_1{}^i C_{t-1} + \beta_2{}^i C_{t-2} + \gamma_i y_{it-1} + \varepsilon_{it}\,,$$

where α_i is a constant, C_t is a variable meant to capture the business cycle, β_0^i, β_1^i, β_2^i are coefficients measuring the effect on y_{it} of current and lagged values of the business cycle (that is, $\beta^i(L)C_t = \beta_0^i C_t + \beta_1^i C_{t-1} + \ldots + \beta_p^i C_{t-p}$, where $p = 2$), γ_i is a coefficient on lagged own-region employment growth, and ε_{it} is an independent and identically distributed random variable with mean 0 and variance σ_i^2, $i = 1, ..., I$.

The business cycle is assumed to affect each region differently in terms of both timing and magnitude. This differing effect is captured parsimoniously by the coefficients β_0^i, β_1^i, β_2^i. Those regions that are less cyclical have values of the β_j^i parameters that are closer to 0. Those regions that lag the cycle have estimates of β_j^i that are insignificantly different from 0 for small j.

Finally, I assume that one cannot observe the business cycle directly, but instead must infer it through its effects on regional employment growth across all regions simultaneously.[11] I assume that the cycle follows an AR(2) specification so that:

$$C_t = \phi_1 C_{t-1} + \phi_2 C_{t-2} + u_t.$$

The error term u_t is assumed to be serially independent and identically distributed with mean 0 and variance of σ_u^2. The imposition of an AR(2) process for the business cycle provides a succinct way of allowing for a business cycle that is characterized by recessions followed by expansions.

To completely specify the model, it is necessary to assume something about the two types of shocks, u_t and ε_{it}, where u_t can be thought of as a business cycle shock and ε_{it} is a regional disturbance. Specifically, I assume that the cyclical shock and the regional disturbances are mean 0, serially uncorrelated, and uncorrelated with each other. Box 4 provides a detailed discussion of the estimation.

Results

As currently specified, the model is not identified without additional restrictions.[12] Neither the scale nor sign of the business cycle is defined. To see this, suppose that the common cycle C_t is rescaled by multiplying it by some constant b, and define $C_t^* = bC_t$. Then $C_t^* = \phi_1 C_t^* + \phi_2 C_t^* + u_t^*$, where $u_t^* = bu_t$ and $\mathrm{var}(u_t^*) = b^2\sigma_u^2$. I fix the scale by setting σ_u^2 to 1 and choose the sign so that β_0 is positive in the East North Central region. In fact, the parameter β_0 turns out to be positive in all regions. This is the natural normalization because we define a boom to be a state when economic activity is high.

Additional assumptions are required to pin down the timing of the cycle. Following Stock and Watson (1989), I normalize by restricting the business cycle to enter only contemporaneously in at least one region j, that is, $\beta^j_1 = \beta^j_2 = 0$. This region has been set arbitrarily as East North Central.[13]

The results reported in table 3 are for the model described above, in which two lags of C_t are included (that is, $\beta^i(L)$ is second order). The estimation uses quarterly data from 1961:Q2 to 1998:Q3 for the nine census regions.[14]

According to the model, movements in the regional employment growth rate reflect macroeconomic conditions, local dynamics, and idiosyncratic fluctuations that are specific to the region. What kind of growth rates should the regions experience over the long term in the absence of cyclical fluctuations and regional shocks? The expected long-term regional growth rate depends upon both the constant α_i and the coefficient on the lagged dependent variable γ_i. Specifically,

$$E(y_i) = \frac{\alpha_i}{(1-\gamma_i)}.$$

From this computation, the West South Central, South Atlantic, and Mountain regions have had the highest growth rates on average, with mean growth over this period of 3.05 percent, 3.09 percent, and 3.74 percent, respectively. The Rust Belt regions of New England, Mid-Atlantic, and East North Central have had the lowest employment growth, recording annual percentage increases of 1.51 percent, 1.01 percent, and 1.98 percent, respectively.

The parameter β_0^i reflects the contemporaneous effect of the business cycle on region i's employment growth. These estimated coefficients (reported in column 2 of table 3) are positive and significant for all regions. The East North Central and East South Central regions are the most cyclically sensitive, exhibiting the largest estimated values for β_0. The West South Central region is by far the least cyclically sensitive contemporaneously with an estimated β_0 of only 0.8406, so that an increase in C_t of one unit is associated with a less than 1 percent increase in regional employment growth contemporaneously.

Technically, the Kalman filter and maximum likelihood estimation provide a way to obtain estimates of the business cycle, C_t, conditional on information prior to time t. I apply a Kalman smoothing technique that uses all available information through the end of the sample period to generate smoothed estimates of C_t. These estimates of the cycle are also referred to as two-sided estimates since they reflect both past and future data.[15]

The process generating the business cycle is estimated as

$$C_t = \underset{(0.1029)}{0.6036} C_{t-1} + \underset{(0.0822)}{0.0123} C_{t-2} + u_t$$

BOX 4

Estimation details

The Kalman filter is a statistical technique that is useful in estimating the parameters of the model specified above. These parameters include α_i, β_k^i, γ_i, ϕ_1, ϕ_2, σ_u^2, σ_i^2 for $i = 1, ..., I$ and for $k = 1, ..., p$. In addition, the Kalman filter enables the estimation of the processes u_t and ε_{it} and the construction of the unobserved cyclical variable C_t. The Kalman filter requires a state equation and a measurement equation. The state equation describes the evolution of the possibly unobserved variable(s) of interest, z_t, while the measurement equation relates observables y_t to the state.

The vector y_t is related to an $m \times 1$ state vector, z_t, via the measurement equation:

$$y_t = Cz_t + D\varepsilon_t + Hw_t,$$

where $t = 1, ..., T$; C is an $N \times m$ matrix; ε_t is an $N \times 1$ vector of serially uncorrelated disturbances with mean zero and covariance matrix I_N; and w_t is a vector of exogenous, possibly predetermined variables with H and D being conformable matrices.

In general, the elements of z_t are not observable. In fact, it is this very attribute that makes the Kalman filter so useful to economists. Although the z_t elements are unknown, they are assumed to be generated by a first-order Markov process as follows:

$$z_t = Az_{t-1} + Bu_t + Gw_t$$

for $t = 1, ..., T$, where A is an $m \times m$ matrix, B is an $m \times g$ matrix, and u_t is a $g \times 1$ vector of serially uncorrelated disturbances with mean zero and covariance matrix I_g. This equation is referred to as the transition equation.

The definition of the state vector z_t for any particular model is determined by construction. In fact, the same model can have more than one state space representation. The elements of the state vector may or may not have a substantive interpretation. Technically, the aim of the state space formulation is to set up a vector z_t in such a way that it contains all the relevant information about the system at time t and that it does do by having as small a number of elements as possible. Furthermore, the state vector should be defined so as to have zero correlation between the disturbances of the measurement and transition equations, u_t and ε_t.

The Kalman filter refers to a two-step recursive algorithm for optimally forecasting the state vector z_t given information available through time t–1, conditional on known matrices $A, B, C, D, G,$ and H. The first step is the prediction step and involves forecasting z_t on the basis of z_{t-1}. The second step is the updating step and involves updating the estimate of the unobserved state vector z_t on the basis of new information that becomes available in period t. The results from the Kalman filtering algorithm can then be used to obtain estimates of the parameters and the state vector z_t employing traditional maximum likelihood techniques.[1]

The model of regional employment growth proposed above can be put into state space form defining the state vector $z_t = (C_t, C_{t-1}, C_{t-2})'$; $y_t = (y_{1t}, ..., y_{It})'$. The system matrices are given below:

$$A = \begin{bmatrix} \phi_1 & \phi_2 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \qquad C = \begin{bmatrix} \beta_0^1 & \beta_1^1 & \beta_2^1 \\ \beta_0^2 & \beta_1^2 & \beta_2^2 \\ \vdots & & \\ \beta_0^9 & \beta_1^9 & \beta_2^9 \end{bmatrix}$$

$$D = \begin{bmatrix} \sigma_1 & 0 & \cdots & 0 \\ 0 & \sigma_2 & \cdots & 0 \\ & & \ddots & \\ 0 & \cdots & & \sigma_9 \end{bmatrix} \qquad B = \begin{bmatrix} \sigma_u \\ 0 \\ 0 \end{bmatrix}$$

$$H = \begin{bmatrix} \alpha_1 & \gamma_1 & 0 & \cdots & 0 \\ \alpha_2 & 0 & \gamma_2 & \cdots & 0 \\ \vdots & & & \ddots & \\ \alpha_9 & 0 & 0 & \cdots & \gamma_9 \end{bmatrix} \qquad G = 0$$

$$\varepsilon_t = (\varepsilon_{1t}\ \varepsilon_{2t} \cdots \varepsilon_{9t})' \qquad w_t = (1\ y_{1t-1}\ y_{2t-1} \cdots y_{9t-1})'.$$

The Kalman filter technique is a way to optimally infer information about the parameters of interest and, in particular, the state vector z_t, which in this case is simply the unobserved cycle, C_t, and its two lags. The cycle as constructed here represents that portion of regional employment growth that is common across the various regions, while allowing the cycle to differ in its impact on industry employment growth in terms of timing and magnitude through the parameters of $\beta^i(L)$. The model is very much in the spirit of Burns and Mitchell's (1946) idea of comovement but the estimation technique permits the data to determine which movements are common and which are idiosyncratic.[2]

[1]The interested reader may obtain further details in Harvey (1989) and Hamilton (1994).

[2]Stock and Watson (1989) is a recent illustration of the Kalman filtering technique for constructing the business cycle.

TABLE 3

Regional employment growth model with lagged dependent variable

Region	Constant	Current cycle	Cycle 1 quarter ago	Cycle 2 quarters ago	Lagged regional employment growth	Standard deviation of regional shock
New England	0.3711**	1.1428***	–0.1332	–0.5495***	0.7535***	1.1183***
	(0.1605)	(0.1207)	(0.1650)	(0.1218)	(0.0559)	(0.0710)
Mid-Atlantic	0.3520**	1.1286***	–0.4275***	–0.0985	0.6529***	0.9122***
	(0.1668)	(0.1092)	(0.1528)	(0.0980)	(0.0719)	(0.0633)
East North Central	1.2952***	1.8330***	0.0000	0.0000	0.3457***	1.1718***
	(0.4102)	(0.1450)	—	—	(0.0822)	(0.0869)
West North Central	1.8999***	1.0579***	0.5853***	0.0050	0.1164	0.8563***
	(0.4076)	(0.1025)	(0.1570)	(0.0632)	(0.0948)	(0.0614)
South Atlantic	1.8717***	1.2708***	0.0000	0.0000	0.3939***	0.9549***
	(0.3251)	(0.1157)	—	—	(0.0514)	(0.0696)
East South Central	2.2168***	1.7102***	0.4925**	–0.2914**	0.1411	0.9026***
	(0.5117)	(0.1359)	(0.2760)	(0.1401)	(0.1198)	(0.0757)
West South Central	0.7077***	0.8406***	–0.2395*	–0.1880*	0.7683***	1.2456***
	(0.2087)	(0.1207)	(0.1544)	(0.1204)	(0.0526)	(0.0750)
Mountain	1.3379***	1.0218***	–0.1161	–0.3058***	0.6418***	1.2519***
	(0.2923)	(0.1271)	(0.1715)	(0.1246)	(0.0642)	(0.0766)
Pacific	1.1861***	1.0751***	–0.2514*	–0.0295	0.5606***	1.3016***
	(0.2921)	(0.1327)	(0.1720)	(0.1366)	(0.0721)	(0.0809)

Notes: The dependent variable is measured as annualized quarterly regional employment growth rates. Regional employment growth is assumed to depend upon a constant, the current and two lags of the state of the economy, and a single lag of own-region employment growth. Maximum likelihood estimates are reported. Standard errors are in parentheses. ***Indicates marginal significance below 1 percent; **indicates marginal significance below 5 percent; and *indicates marginal significance below 10 percent. The mean log-likelihood is 6.48760 at the maximum.

Source: See table 2.

and is shown in figure 3 for the smoothed estimates. The estimated employment cycle roughly corresponds to the timing of the NBER business cycle in the sense that contractions occur at approximately the same time as the NBER recessions. Interestingly, business cycle peaks as measured here typically precede the NBER-dated peaks and recoveries tend to precede the NBER-dated recoveries. This is particularly notable in light of the fact that the measure of cyclical activity constructed here is based upon employment data alone. It is a well-known empirical regularity that employment lags the business cycle. This can be seen from carefully comparing real gross domestic product (GDP) growth and aggregate employment growth in figure 4. So cyclical measures constructed from employment data alone might be reasonably expected to lag as well. As figure 3 shows, however, this hypothesis is not supported by the data.

Given the high real GDP growth rates of recent quarters, as shown in figure 4, we might expect the business cycle to be abnormally high over this period.

FIGURE 3

Smoothed estimates of business cycle, 1961:Q4–98:Q2

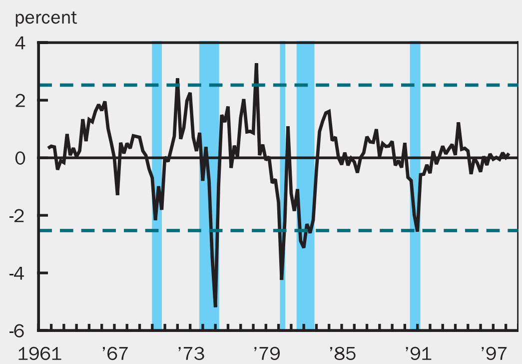

Notes: See box 2 for details of calculations. The horizontal lines represent a band of plus or minus two standard deviations. Shaded areas indicate recessions, as defined by the National Bureau of Economic Research.

Sources: U.S. Department of Labor, Bureau of Labor Statistics, 1960–98, employment database available at ftp://ftp.bls.gov/pub/time.series and author's calculations.

FIGURE 4

Alternative measures of the business cycle, 1961:Q4–98:Q2

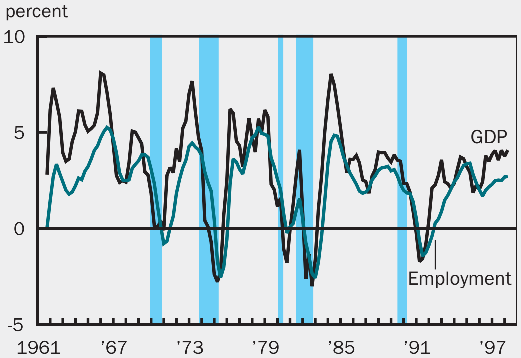

Notes: Growth rates were calculated as four quarter log differences in the respective variable. Gross domestic product (GDP) data in 1992 chain-weighted dollars; employment data for total civilian nonfarm payroll employment. All data are seasonally adjusted. Shaded areas indicate recessions, as defined by the National Bureau of Economic Research.

Sources: U.S. Department of Commerce, Bureau of Economic Analysis, 1961–98, *National Income and Product Accounts*, and U.S. Department of Labor, Bureau of Labor Statistics, 1961–98, employment database available at ftp://ftp.bls.gov/pub/time.series.

Instead, the estimated cycle suggests business conditions are currently hovering around neutral. The reason for the apparent disparity is quite simple. The business cycle as constructed here depends solely upon comovements in regional employment growth. However, employment growth has recently been close to its long-term average, as is also apparent in figure 4. The employment-based measure of the business cycle constructed here reflects this trend employment growth as implying neutral economic conditions.

GDP has exhibited such strong growth in recent quarters because of the increase in productivity of the economy and not because of any substantive increase in employment growth. High productivity growth will tend to increase output without a concomittant rise in employment. This is what appears to have happened in the latter part of the sample. Conversely, when productivity growth is low and employment growth remains stable, output-based measures of the cycle are likely to show deeper recessions than employment-based measures.

What happens to regional employment growth when the economy experiences an aggregate one-time shock, that is, a change in the common shock u_t? A positive cyclical shock of one standard deviation in magnitude increases the cycle by a unit of 1 at the time it occurs. This, in turn, affects regional employment growth contemporaneously. The following quarter the shock disappears but its effects linger and are felt in two ways. First, the shock has an evolving effect on the business cycle through its autoregressive structure.[16] This effect translates into movements in regional employment growth that also evolve over time. Second, the shock affects regional employment growth through the lag of regional employment growth (feedback).

Figure 5 traces the effect of a one standard deviation one-time aggregate business cycle shock on the cycle and also on regional employment growth. The effect of the aggregate disturbance on the business cycle itself dissipates smoothly over time. The regions' responses show more complicated dynamics, with the largest impact being felt at the same time the disturbance occurs and one quarter thereafter. The effect then fades over time. (In the West South Central region, the shock's initial effect is smaller but the effect lingers slightly longer than in other regions.)

In East North Central, for example, the cyclical shock contemporaneously increases employment growth by 1.75 percent per annum relative to its long-term average. The following quarter as these other feedbacks influence regional employment growth, the effect remains about the same at 1.71 percent, despite the value of the shock returning to 0. However, as time progresses, the cyclical shock's effect fades so by the seventh quarter following the shock, employment growth in the East North Central region is only 0.14 percent higher per annum than it would have been in the absence of the disturbance.

Recall that the variance of the cyclical shock has been scaled to equal unity. Because the current state of the economy depends upon past realizations of the business cycle as well as the aggregate shock, its variance will reflect these dynamics. The variance of C_t is computed as

$$\mathrm{var}(C_t) = \frac{(1-\phi_2)}{(1+\phi_2)\left[(1-\phi_2)^2 - \phi_1^2\right]} = 1.596.$$

Consequently, a one unit increase in u corresponds approximately to a one standard deviation shift in the cycle of $(1.596)^{1/2} = 1.263$.

Table 4 illuminates the relative importance of the business cycle and the regional idiosyncratic shocks in explaining the variance of each region's employment growth. (The calculations are shown in box 5.) Clearly, regional shocks are more important in some regions than in others. In West South Central, for example, the regional shock accounts for almost 60 percent of the variance in the region's employment growth rate. Regional idiosyncratic shocks account for a somewhat smaller but still sizable proportion of

FIGURE 5

Effect of one standard deviation cyclical shock
(percent)

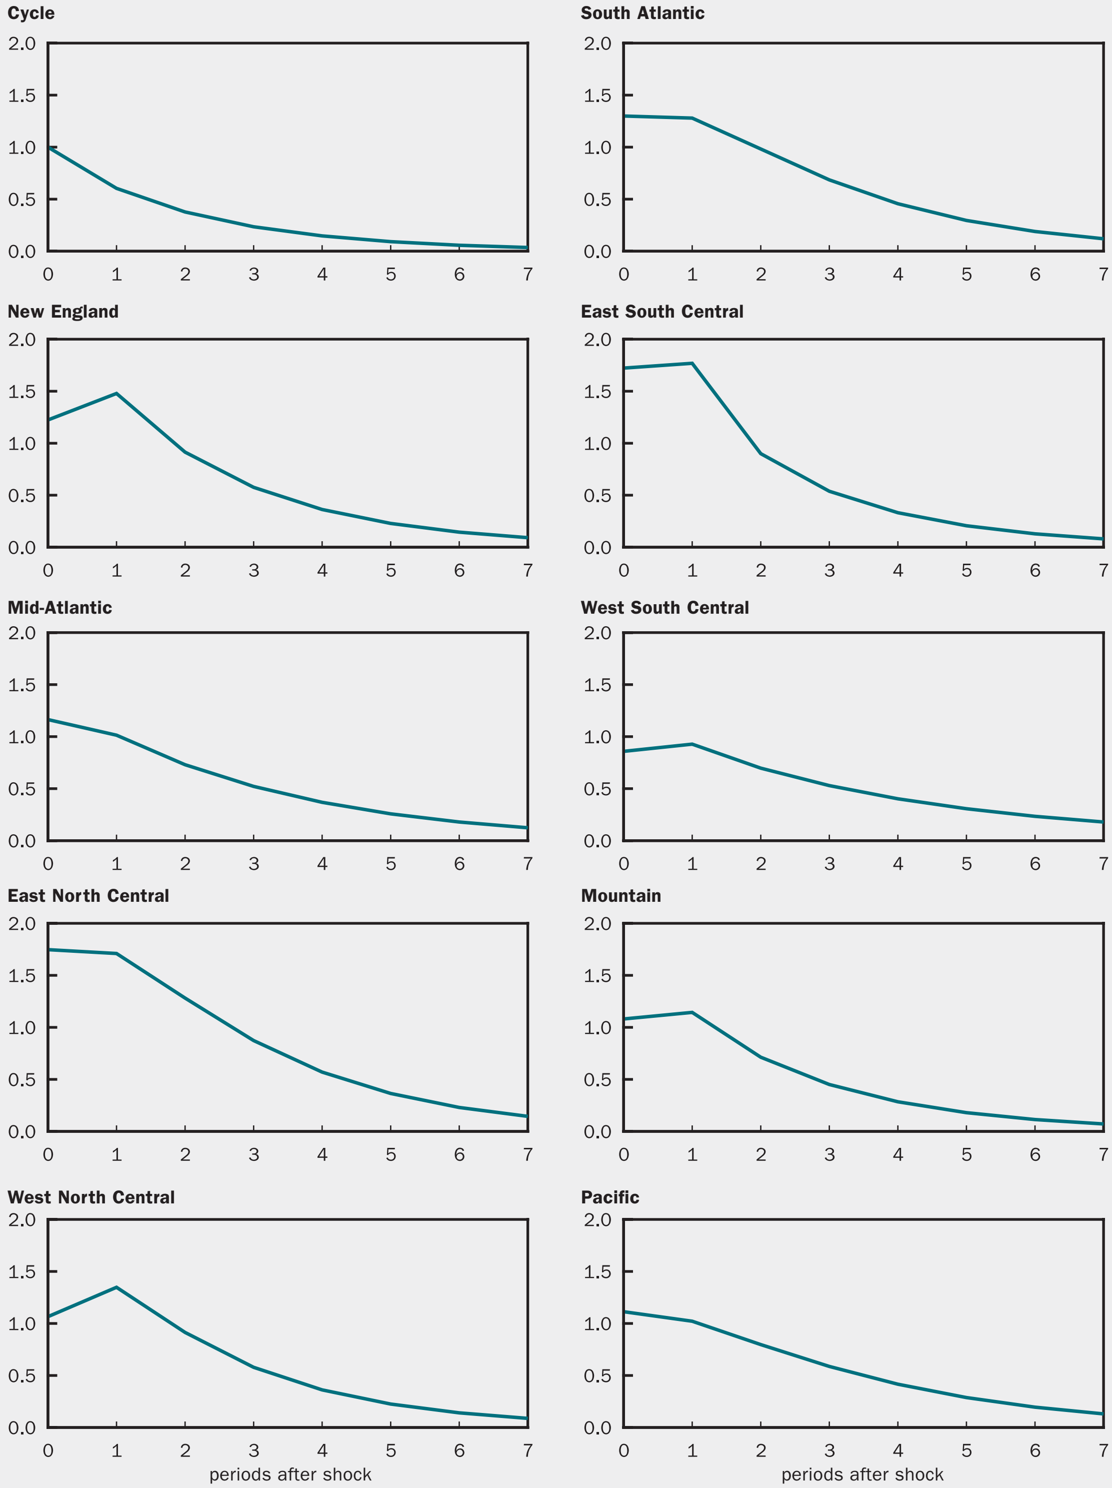

Note: The panels trace the effect of a one standard deviation shock of +1 in the cyclical disturbance on the cycle and each region separately, taking into account the dynamics of the cycle and the dependence of current regional employment growth on lagged regional employment growth.

Source: Author's calculations.

the total variance in New England, Mid-Atlantic, Mountain, and Pacific. This compares with East South Central, where almost 90 percent of the region's total variance is attributable to variance in the aggregate shock. The East North Central, West North Central, and South Atlantic regions appear to be influenced in large part by the aggregate shock.

The model has been estimated under the assumption that the regional disturbances are uncorrelated with each other for all leads and lags and are serially uncorrelated. This is a strong assumption and a test is useful to assess the validity of the estimated model. According to the model estimated above, all comovement is ascribed to the common cyclical shock. If the model is true, then errors made in forecasting regional employment growth in one region should not be useful for predicting regional employment growth in another region. One can construct a simple diagnostic test in which the estimated one-step-ahead forecast errors in a region's employment growth are regressed against lags of the one-step-ahead forecast errors in other regions.[17] If the model describes the data well, lags of another region's forecast errors should not be significantly different from 0 in these regressions. In other words, errors made in forecasting another region's employment growth should not significantly aid in the prediction of a given region's employment growth.

In table 5, *p*-values are reported for the regressions described above, testing for the significance of forecast error lags. If the model fits the data well, the *p*-values should be large. Small *p*-values indicate that the independent variable has some predictive content for the dependent variable. Because of natural variation, we would expect about 10 percent of the regressions (that is, eight or nine) to have *p*-values of less than 0.100 even if the hypothesis was true. Table 5 shows that, in fact, ten of the regressions show significantly low *p*-values. More significantly, most of these low *p*-values are in regressions involving the predictive content of forecast errors in the West South Central region.

One obvious reason why the West South Central region may wield such influence in regional employment growth stems from the industrial composition of the area. The West South Central states are heavily dependent on oil and gas production. Disturbances to these industries, in turn, have repercussions for other industries and regions of the country. My results imply that, in addition to the common cyclical factor affecting all regions, there might be another factor involved in explaining regional employment growth patterns. This factor is likely related to oil price shocks. Further research is necessary to test this hypothesis.

The main advantage of estimating a Kalman filter model of the sort presented here is its ability to obtain estimates of the underlying cyclical and regional disturbances, as shown in figure 6. The analysis suggests that New England experienced some positive shocks in the late 1970s and early 1980s, coinciding roughly with well-documented growth in technology and business services at that time. However, some time in the late 1980s, the region experienced a series of large negative shocks. These shocks correspond to the timing of the S&L crisis and the credit crunch. At about this time, computers were making the transition from mainframe to desktop and some larger New England employers were cutting back their labor force in large numbers. Employment growth in New England has recovered to some extent and is approximately in line with what is predicted by the model.[18]

The Mid-Atlantic region is heavily influenced by New York. Regional employment growth has held fairly steady, with the stock bust of 1987 causing lower employment growth. The East North Central region experienced a large negative disturbance during the period surrounding the first oil price shock and smaller negative ones in 1978 and in 1980. For much of the 1980s through mid-1990s, employment growth shocks in this area were small and tended to be positive. This likely reflects the bottoming out of the farm crisis in 1986 and strong export growth. The farm crisis also appears to have had an effect on employment

TABLE 4

Steady state regional employment growth variance due to cycle and shock, 1961:Q2–98:Q3

Region	Steady state employment growth variance	Percent of variance from cyclical shock	Percent of variance from regional shock
New England	7.3284	60.5	39.5
Mid-Atlantic	4.6298	64.8	31.3
East North Central	10.7283	85.5	14.5
West North Central	4.9939	85.1	14.9
South Atlantic	5.9623	81.9	18.1
East South Central	8.0531	89.7	10.3
West South Central	6.3430	40.3	59.7
Mountain	5.9528	55.2	44.8
Pacific	5.7963	57.4	42.6

Note: See box 5 for a discussion of the calculations.
Source: See table 2.

BOX 5

How important are regional shocks?

The steady state variance of regional employment growth reported in table 4 is, in general, a complicated function depending upon the variance of the idiosyncratic shock, the variance of the cyclical disturbance, the cross-correlation structure between regions, and the dynamics of the model. To construct a measure of the steady state variance of regional employment growth, first rewrite the model in terms of a vector AR(1) process. Specifically, let $\underline{z}_t = (y_{1t}, y_{2t}, \ldots, y_{9t}, C_{t+1}, C_t, C_{t-1})'$ and rewrite the model as:

$$\underline{z}_t = \Pi \underline{z}_{t-1} + \underline{v}_t,$$

where $\underline{v}_t = (\varepsilon_{1t}, \varepsilon_{2t}, \ldots, \varepsilon_{9t}, u_t, 0, 0)'$ and the matrix Π is formed as follows:

$$\Pi = \left[\begin{array}{c|c} \Gamma_{9x9} & C_{9x3} \\ \hline 0_{3x9} & A_{3x3} \end{array}\right],$$

where the matrix Γ has $\gamma_1, \ldots, \gamma_9$ along the diagonal and 0 elsewhere, and A is defined in box 4. Let the variance–covariance matrix of $\underline{v}_t$ and $\underline{z}_t$ equal Σ and Ω, respectively. Then

$$\Omega = \Pi\Omega\Pi' + \Sigma,$$

which has the following solution:

$$\text{vec}(\Omega) = [\text{I} - (\Pi \otimes \Pi)]^{-1}\text{vec}(\Sigma).$$

In this case the total steady state variance of a region's employment growth is the sum of two terms, one reflecting the variance of the idiosyncratic regional shock, and the other reflecting the variance of the cyclical disturbance. Calculating the percentage attributable to each of the two shocks follows easily.

growth in the West North Central region. The West South Central region appears to have more volatility, and experienced a large negative disturbance in the mid-1980s. This shock is most likely the result of the oil price bust, followed by a recovery in the industry. Finally, the Pacific region was hit by a series of negative shocks in the early 1990s due to cutbacks in defense spending.[19] The Pacific region seems to have recovered to a large extent.

Conclusion

The business cycle is not observable directly. Instead, it must be inferred from observing many data series simultaneously. Casual observation suggests

TABLE 5

Significance of lagged regional employment growth forecast errors

$j \downarrow$ $i \rightarrow$	New England	Mid-Atlantic	East North Central	West North Central	South Atlantic	East South Central	West South Central	Mountain	Pacific
New England	**0.083**	0.288	0.652	0.650	0.381	0.639	0.189	0.762	0.336
Mid-Atlantic	0.423	**0.063**	0.699	0.450	0.551	0.385	0.639	0.500	0.786
East North Central	0.294	0.863	0.161	0.304	**0.074**	0.438	0.316	0.200	0.678
West North Central	0.997	0.973	0.769	0.834	0.273	0.885	0.250	0.878	0.839
South Atlantic	0.693	0.766	0.735	0.767	0.698	0.987	0.860	0.854	0.330
East South Central	0.934	0.612	0.410	0.209	0.931	0.721	0.693	0.970	0.651
West South Central	0.219	0.214	**0.007**	**0.070**	**0.003**	**0.008**	0.749	**0.031**	**0.048**
Mountain	0.706	0.380	0.538	0.942	0.713	0.885	0.403	0.599	0.345
Pacific	0.501	**0.026**	0.271	0.339	0.744	0.692	0.190	0.480	0.382

Notes: The table reports p-values for OLS regressions of the form:

$$e_{it} = c + \beta_1 e_{jt-1} + \beta_2 e_{jt-2} + \ldots + \beta_6 e_{jt-6} + \upsilon_t,$$

where e_{it} and e_{jt} are the estimated one-step-ahead forecast errors at time t for regional employment growth and $i, j = 1, \ldots, 9$. The p-values reported in the table are the significance levels for the test of the null hypothesis that the β coefficients are 0. Low p-values indicate that the hypothesis is not consistent with the data. Numbers in bold indicate a p-value less than 0.100.

Source: See table 2.

FIGURE 6

Estimates of cyclical and regional shocks, 1961:Q4–98:Q2
(percent)

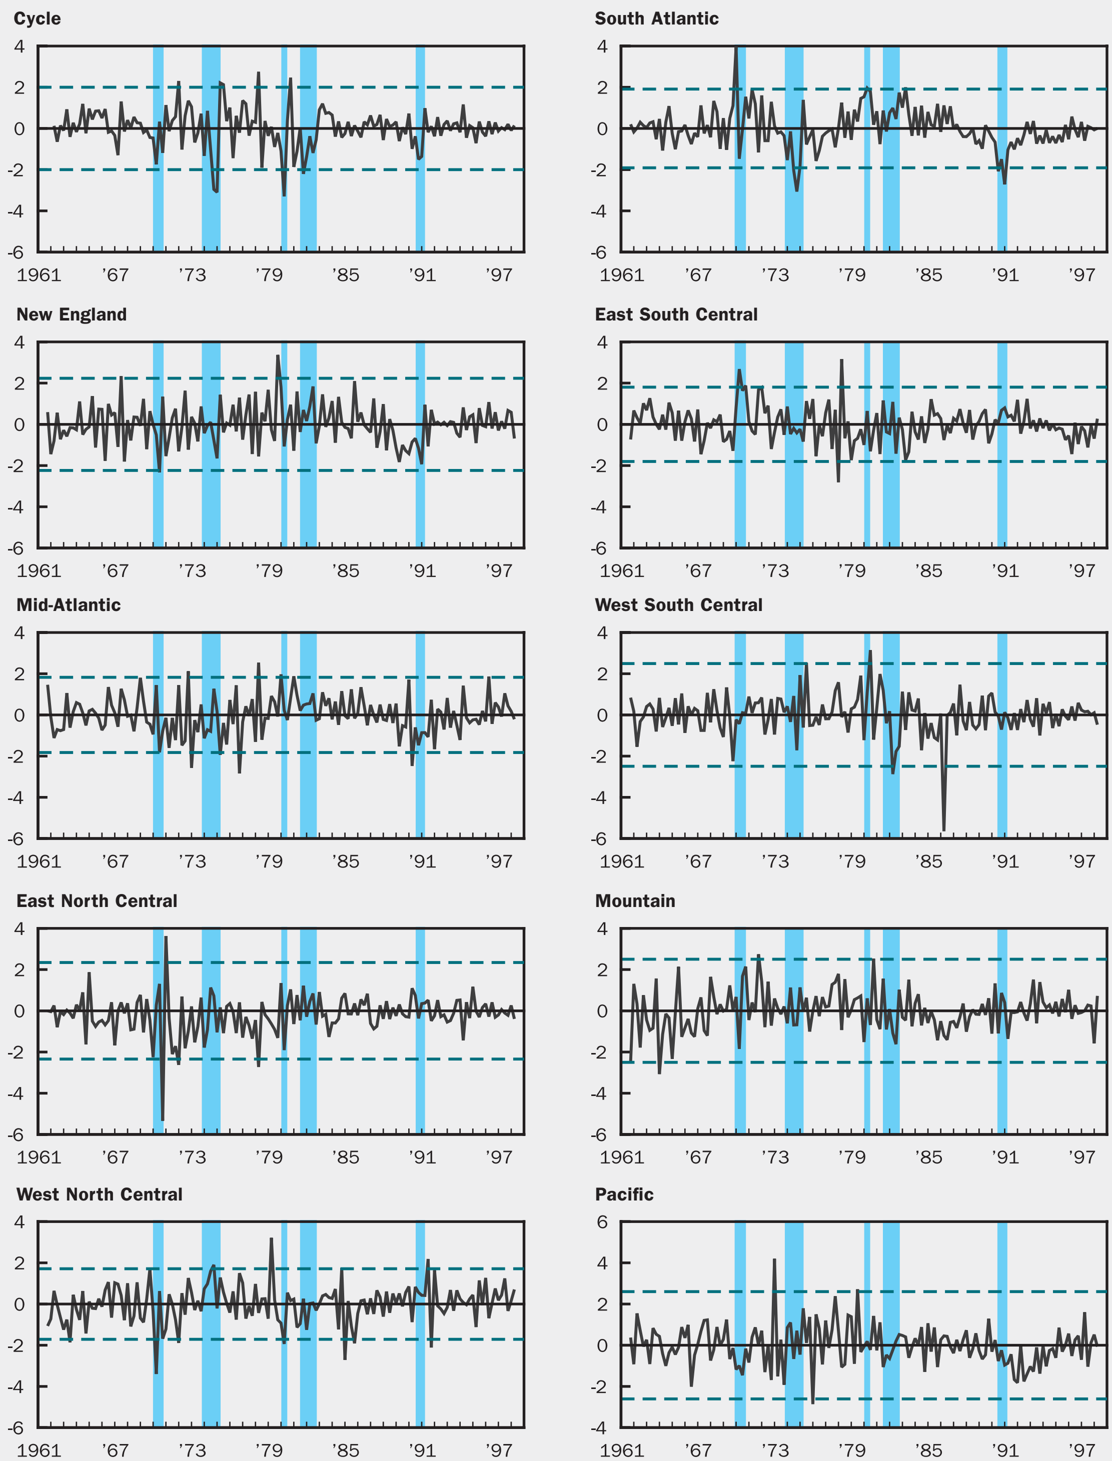

Notes: The horizontal lines represent two standard error bands. Shaded areas indicate recessions, as defined by the National Bureau of Economic Research.
Source: U.S. Department of Labor, Bureau of Labor Statistics, 1961–98, employment database available at ftp://ftp.bls.gov/pub/time.series.

that all regions experience some cyclicality in employment growth, despite the fact that some regions show above-average employment growth over long periods and other regions consistently report below-average employment growth. The fact that these regions move more or less in tandem over time provides a way to construct a measure of the business cycle.

In this article, I define the business cycle as comovements in regional employment growth. I estimate the cycle using the Kalman filter and maximum likelihood techniques. The estimates of the cycle obtained from the model are quite consistent and conform with more traditional measures of the business cycle, for example, GDP growth or the unemployment rate.

Because employment growth is distinct from productivity growth, the estimates of the cycle do not exhibit the large expansion in the most recent period that output-based measures do. In fact, current estimates of the business cycle show that the economy is well balanced, in the sense that there are no cyclical shocks that seem to be expanding or contracting regional employment growth above or below long-term averages. If employment growth contributes to inflation, this balance in the economy seems to imply that, despite high output growth, inflation is under control.

Sectoral disturbances appear to be an important determinant of regional employment growth—at least in some regions. This is particularly true for the West South Central, Mountain, Pacific, New England, and Mid-Atlantic states. Regional shocks play a far less important role in explaining regional employment growth in the East North Central, West North Central, South Atlantic, and East South Central regions, where most of the movements are related to aggregate fluctuations.

There are obviously many ways one could define the business cycle. The tack taken here is to define it relative to regional employment growth patterns. This is not to say that all other information should be excluded from the analysis. However, the focus on an employment-based measure helps shed light on regional issues. Furthermore, a comparison of an employment-based cyclical measure versus an output-based measure may aid in our understanding of productivity.

Finally, the methodology employed permits the recovery of a series of regional employment shocks. The timing of such disturbances may be helpful for assessing what factors may explain regional declines or expansions that are not anticipated by long-term patterns or cyclical influences. Although speculative, it appears that oil shocks and defense contracts might help explain the origin of regional shocks. The model estimated here is somewhat simplistic, in that it does not allow for regional spillovers that are not accounted for by the aggregate shock. By examining the regional disturbances that the model estimates and formulating a better notion of the underlying economics behind these regional shocks, one could develop a richer understanding of regional dynamics.

NOTES

[1]A comprehensive list is outside the scope of this article. A few references include Barro (1977, 1978), Mishkin (1983), Gordon and Veitch (1986), and Litterman and Weiss (1985).

[2]Blanchard and Watson (1986).

[3]Mitchell (1927).

[4]Clark (1998), p. 202.

[5]A more appropriate nomenclature might be the "employment cycle" since it is constructed by filtering out the common movements in employment across regions. In contrast, the "business" cycle is typically modeled as comovements in less narrowly focused series. For example, Stock and Watson (1989) construct their Coincident Economic Index with reference to industrial production, total personal income less transfer payments in 1982 dollars, total manufacturing and trade sales in 1982 dollars, and employees on nonagricultural payrolls.

[6]The New England states are Maine, New Hampshire, Vermont, Massachusetts, Connecticut, and Rhode Island. Mid-Atlantic contains New York, Pennsylvania, and New Jersey. East North Central comprises Wisconsin, Michigan, Illinois, Indiana, and Ohio. South Atlantic contains Maryland, Delaware, Virginia, West Virginia, North Carolina, South Carolina, Georgia, and Florida. East South Central states are Kentucky, Tennessee, Alabama, and Mississippi. West South Central contains Oklahoma, Arkansas, Louisiana, and Texas. The East North Central states are Minnesota, Iowa, Nebraska, Kansas, North Dakota, South Dakota, and Missouri. The Mountain states are Montana, Idaho, Wyoming, Nevada, Utah, Colorado, Arizona, and New Mexico. Pacific contains Alaska, Hawaii, Washington, Oregon, and California.

[7]These trends have been noted by previous researchers, including Blanchard and Katz (1992).

[8]The timing of the cyclical upturns and downturns in regional employment growth is somewhat different from that proposed by the NBER dating. It is well known that employment reacts with a small lag to cyclical events so, for example, the trough of the recessions is typically a short time after the NBER dating of the trough.

[9]This observation was made by Mitchell (1927).

[10]Seasonally unadjusted data are reported monthly by the BLS and are available on the BLS Labstat website. Calculations were carried out using quarterly data that have been seasonally adjusted using the PROC X11 procedure. Hawaii has been omitted from the calculations due to a lack of data for mining.

[11]A richer model might incorporate other cyclical series as well, such as gross domestic product (GDP) or industry employment. However, because the objective is to describe regional employment patterns, the business cycle is constructed by looking at comovements in regional employment patterns alone.

[12]The discussion here follows Harvey's (1989) analysis of common trends.

[13]A more subtle point is raised in Stock and Watson (1989). Given three data series that are serially uncorrelated but are correlated with each other, it is always possible to restructure the model with a single index. This common factor captures the covariance of the three series. Over-identification occurs when there are more than three observable variables (there are nine here) or when the variables are serially correlated.

[14]The BFGS algorithm was used in maximizing the likelihood function. In practice, numerical difficulties arose in which the Hessian matrix failed to invert when the model was estimated with the sole restriction that lags of the cycle do not enter into the East North Central Region. The problem was resolved by restricting the South Atlantic region to depend solely upon the contemporaneous cycle as well.

[15]Details of this procedure can be found in chapter 4 of Harvey (1989).

[16]The evolution of the business cycle following a temporary one standard deviation shock is found in the first panel of figure 5.

[17]The one-step-ahead forecast error is simply defined as:

$$\hat{e}_{it} \equiv y_{it} - \hat{y}_{it|t-1},$$

where the forecast error $\hat{e}_{it}$ is calculated as the difference between the actual regional employment growth rate at time t and the model's prediction of regional employment growth based upon information up to time $t - 1$.

[18]Bradbury (1993) examines employment over the 1990–91 recession and the recovery in New England.

[19]See Gabriel et al. (1995) for a discussion of migration trends in California.

REFERENCES

Altonji, Joseph G., and John C. Ham, 1990, "Variation in employment growth in Canada: The role of external, national, regional, and industrial factors," *Journal of Labor Economics*, Vol. 8, No. 1, Part 2, January, pp. S198–S236.

Barro, Robert, 1978, "Unanticipated money, output, and the price level in the United States," *Journal of Political Economy*, Vol. 86, August, pp. 549–580.

__________, 1977, "Unanticipated money growth and unemployment in the United States," *American Economic Review*, Vol. 67, March, pp. 101–115.

Blanchard, Olivier Jean, and Lawrence F. Katz, 1992, "Regional evolutions," *Brookings Papers on Economic Activity*, Vol. 1, Washington: Brookings Institution, pp. 1–61.

Blanchard, Olivier, and Mark Watson, 1986, "Are business cycles all alike?," in *The American Business Cycle*, Robert Gordon (ed.), Chicago: University of Chicago Press, pp. 123–156.

Bradbury, Katharine L., 1993, "Shifting patterns of regional employment and unemployment: A note," *New England Economic Review*, Federal Reserve Bank of Boston, September/October, pp. 3–12.

Burns, Arthur F., and Wesley C. Mitchell, 1946, *Measuring Business Cycles*, New York: National Bureau of Economic Research, Inc.

Clark, Todd E., 1998, "Employment fluctuations in U.S. regions and industries: The roles of national, region-specific, and industry-specific shocks," *Journal of Labor Economics*, January, pp. 202–229.

Clark, Todd E., and Kwanho Shin, 1999, "The sources of fluctuations within and across countries," in *Intranational Macroeconomics*, Gregory D. Hess and Eric van Wincoop (eds.), Cambridge, New York, and Melbourne: Cambridge University Press, chapter 9.

Gabriel, Stuart A., Joe P. Mattey, and William L. Wascher, 1995, "The demise of California reconsidered: Interstate migration over the economic cycle," *Economic Review*, Federal Reserve Bank of San Francisco, No. 2, pp. 30–45.

Gordon, Robert J., and John Veitch, 1986, "Fixed investment in the American business cycle, 1919–83," in *The American Business Cycle*, Robert J. Gordon (ed.), Chicago: University of Chicago Press, pp. 267–335.

Hamilton, James D., 1994, *Time Series Analysis*, Princeton, NJ: Princeton University Press.

Harvey, Andrew C., 1989, *Forecasting Structural Time Series Models and the Kalman Filter*, Cambridge, New York, and Melbourne: Cambridge University Press.

Litterman, Robert, and Lawrence Weiss, 1985, "Money, real interest rates and output: A reinterpretation of postwar U.S. data," *Econometrica*, Vol. 53, January, pp. 129–156.

Mishkin, Fredrick, 1983, *A Rational Expectation Approach to Macroeconomics: Testing Policy Ineffectiveness and Efficient Market Models*, Chicago: University of Chicago Press.

Mitchell, Wesley, C., 1927, *Business Cycles: The Problem and Its Setting*, New York: National Bureau of Economic Research, Inc.

Rissman, Ellen R., 1997, "Measuring labor market turbulence," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 21, No. 3, May/June, pp. 2–14.

Stock, James H., and Mark W. Watson, 1989, "New indexes of coincident and leading economic indicators," in *NBER Macroeconomics Annual 1989*, Cambridge, MA: Massachusetts Institute of Technology Press, pp. 351–394.

__________, 1988, "A probability model of the coincident economic indicators," National Bureau of Economic Research, working paper, No. 2772.

The 36th Annual Conference on Bank Structure and Competition

Call For Papers

The Federal Reserve Bank of Chicago invites the submission of research and policy-oriented papers for the 36th annual Conference on Bank Structure and Competition, to be held May 3–5, 2000, in Chicago. Since its inception, the purpose of the conference has been to encourage an ongoing dialogue and debate on current public policy issues affecting the financial services sector.

The Changing Financial Industry Structure And Regulation:

The theme for 2000 will return us to the original roots of the conference: financial industry structure and competition. With the recent elimination of geographic barriers to bank expansion, there has been a significant increase in the number and size of bank mergers in the U.S. Similar activity has been occurring in other countries. We have also seen an increase in cross-border mergers. Why? Is the consolidation occurring for the "right" reasons? In the right way? Why do banking consultants and academics frequently disagree on the potential benefits and costs associated with mergers? What determines a successful bank merger?

What are the implications of mergers for industry competition? For current antitrust methodologies? For regulation? Do the answers differ for "megamergers"? What impact might the development of Internet banking have on industry structure? On market definitions and related regulatory issues, such as the CRA? With all these changes, has the role of smaller banks been enhanced or diminished?

Perhaps even more importantly, in addition to bank mergers we are also seeing cross-industry mergers and affiliations—a convergence of commercial banks, investment banks, and insurance firms into modern financial service providers. While still "separate" in name, in reality, the boundaries between these firms have been significantly decreased or eliminated. What's driving this trend? Are there significant efficiency gains from universal banking or one-stop shopping? Should consideration also be given to allowing commercial banks to take equity positions in nonfinancial firms?

With respect to both merger trends, is regulation in the financial services industry keeping pace? Have the regulatory goals and structure adjusted appropriately? Most would argue that the SEC's mission and method of operation concerning investment bank regulation remains very different from that of commercial bank regulators. Furthermore, the insurance industry is regulated on a state-by-state basis. Are these differences desirable for the future? What should the goals be for financial regulation in the new environment? Is it possible for investment and commercial bank regulators to have conflicting mandates? Which, if either, takes precedent? Is there a need for a single financial regulator for all financial products or should there be several competing regulators each covering a broad range of financial services? Should regulators be organized around function or different customer classes, such as wholesale versus retail? Should they be organized around different public policy goals, such as consumer protection versus safety and soundness? Is there a need for an "umbrella" regulator? Given the accelerating rate of globalization in financial services, is there a need for more international harmonization of regulations? Similarly, should there be more international coordination among regulatory agencies?

The 2000 conference will focus on these and related questions. Although much of the program will be devoted to the primary conference theme, there will be a number of additional sessions on industry structure and regulation such as:

- Basel capital standards;
- measuring and managing risk, particularly for transnational/global financial services companies;
- alternative approaches to dealing with financial crises;
- the role and structure of the lender of last resort;
- financial modernization in the U.S.;
- safety net reform;
- reforming the international financial institutions (i.e., the IMF, World Bank, BIS, and regional development banks);
- the implications of Internet banking;
- technology and payment innovations;
- risk-based pricing of loans (particularly mortgage lending); and
- fair lending issues.

Continuing the format of recent years, the final session of the conference will feature a panel of industry experts who will discuss the purpose, structure, problems, and proposed changes associated with an important banking regulation. A record of the panel discussion will be included in the Proceedings of the Conference. Past topics discussed at this session include bank antitrust analysis, bank capital regulation, optimal regulatory structures, and the appropriate role of the lender-of-last-resort. Proposals for this session are also welcome.

Bridging States, Countries, And Industries In The New Millennium

If you would like to present a paper at the conference, please submit four copies of a completed paper or a detailed abstract (the more complete the paper the better) with your name, address, affiliation, telephone number, and e-mail address, and those of any co-authors, by December 17, 1999.
Correspondence should be addressed to:

Conference on Bank Structure and Competition
Research Department
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, Illinois 60604-1413

For additional information contact:
Douglas Evanoff at 312-322-5814 (devanoff@frbchi.org) or
Portia Jackson at 312-322-5775 (portia.jackson@chi.frb.org).

May 3–5, 2000

Child care costs and the return-to-work decisions of new mothers

Lisa Barrow

Introduction and summary

Women's labor force participation has nearly doubled over the last 50 years, from 31.0 percent in January 1948 to 60.6 percent by March 1999 (based on monthly data from the *Current Population Survey*). For women with young children, the increases have been even more dramatic. From 1947 to 1996, the labor force participation rate of women with preschool-aged children increased by more than a factor of five, rising from 12.0 percent to 62.3 percent (U.S. House of Representatives, 1998). The rapid increase in participation of women with young children indicates that women are spending less time out of the labor force for child bearing and rearing. Indeed, looking at new mothers in the *National Longitudinal Survey of Youth* (NLSY), of those who were working prior to the birth of their first child, three-quarters were back at work within a year of the birth.

An important consequence of the trend toward more rapid reemployment of new mothers is that recent generations of women will have more actual labor market experience (at each age) than their predecessors.[1] In labor economics, a standard analysis of the relationship between wages and education and age (reflecting potential experience) shows that wages increase with years of potential experience. For women, potential experience is likely to exceed actual experience by more than for men. Thus, the increase in women's actual work experience should be reflected in a narrowing of the gender earnings gap. In fact, despite the growing wage inequality of the 1980s, the male–female earnings gap has been closing steadily since the late 1970s. From 1978 to 1990, the ratio of female to male earnings rose from 0.73 to 0.85 for whites and from 0.60 to 0.70 for African-Americans.[2] According to O'Neill and Polachek (1993), about one-quarter of the closing of the male–female wage gap over the 1976–87 period can be attributed to changes in the actual labor force experience of women and an additional 50 percent can be accounted for by changes in returns to experience for women relative to men. Realistically, working women who choose to have children will have to take some time off of work either by taking family, sick, or vacation leave or by exiting the labor market entirely. However, given the importance of experience in determining wages, the faster women return to work following childbirth, the closer their actual experience will be to their potential experience and the smaller the average earnings penalty for women who have children.

In this article, I examine the economic determinants of a woman's decision to return to work quickly following childbirth. I consider three key factors in this decision: the opportunity cost of taking time out of the labor force (that is, the potential wage rate available to a woman), the wealth effect of other family income, and most particularly, the opportunity cost of working outside the home in terms of child care costs.

I first describe a simple theoretical model of a new mother's return-to-work decision. The model predicts that the decision to return to work will depend on a woman's wage net of hourly child care costs and other family income (including spouse or partner income). I then test the theoretical model as closely as possible. In order to get a measure of child care costs faced by women as they decide whether to return to work, I calculate average child care worker wages across states and over time to proxy for variation in child care cost across states and over time. I find that women with higher wages are significantly more

Lisa Barrow is an economist in the Economic Research Department of the Federal Reserve Bank of Chicago. *The author would like to thank Paul Huck, David Marshall, Dan Sullivan, and seminar participants at the Federal Reserve Bank of Chicago for helpful comments and suggestions.*

likely to return to work, and women facing higher child care costs or having greater other family income are significantly less likely to return to work after first birth. I also find that older women, women with more education, and women whose adult female role model was working when they were teenagers are more likely to return to work.

Additional interest in women's labor force participation has been generated by the reforms to welfare programs that have a primary goal of getting recipients off of welfare and into the work force. Because the majority of welfare recipients are women with children, child care costs may have important effects on getting these women into the labor force. Therefore, I look for greater sensitivity to child care costs among women with less than a high school education who are not married or do not have a spouse present. I find no evidence that these women's labor force participation decisions are more sensitive to child care costs. Additionally, I find that for these women the decision to return to work is also no more sensitive to the unemployment rate of their home county than for other women.

While this study was not designed to test alternative policies, several inferences may be drawn. First, the results suggest that delayed child bearing may have a greater impact on increasing labor force participation of women with young children than increases in wages or decreases in child care costs. Second, while access to reliable child care is likely to be a necessity for successfully moving mothers from welfare to the labor force, this research shows no evidence that welfare recipients will be more responsive to changes in child care costs than other women. Finally, the increased probability of a woman working after childbirth associated with her female role model having worked suggests that we should expect to see continuing increases in the labor force participation rate of women, thus increasing the size of the labor force.

Previous research

Much of the previous literature on the labor supply behavior of women with young children has focused on the effect of child care costs.[3] Looking at Census Bureau estimates from the Survey of Income and Program Participation (SIPP) 1988 Panel, employed mothers spend an average of $73.30 per week on child care, while employed women with at least one child under age one spend an average of $88.60 per week.[4] Since, on average, these women work about 36 hours per week, child care costs represent a $2.00 to $2.50 per hour "tax" on the work effort of mothers with young children.

Anderson and Levine (1998) provide a good overview of much of the empirical literature examining the relationship between child care and mothers' employment decisions. They note that while many studies find the expected negative relationship between child care costs and women's labor force participation decisions, there is much variability among the estimates in how responsive women are to changes in child care costs.

My approach builds on several of the earlier studies using the relatively detailed information available in the NLSY. Although some of the earlier studies—Blau and Robins (1991), Leibowitz, Klerman, and Waite (1992), and Klerman and Leibowitz (1990)—use NLSY data as well, their data are less current and hence less representative of women at first birth. In addition, I use the subset of new mothers who were working in the period before their first birth in order to focus specifically on the return-to-work decision.[5]

Women's labor force participation

As mentioned in the introduction, women's labor force participation rate has increased dramatically in the last 50 years. Labor force participation rates for women, men, and subgroups of women with children

FIGURE 1

Labor force participation rates

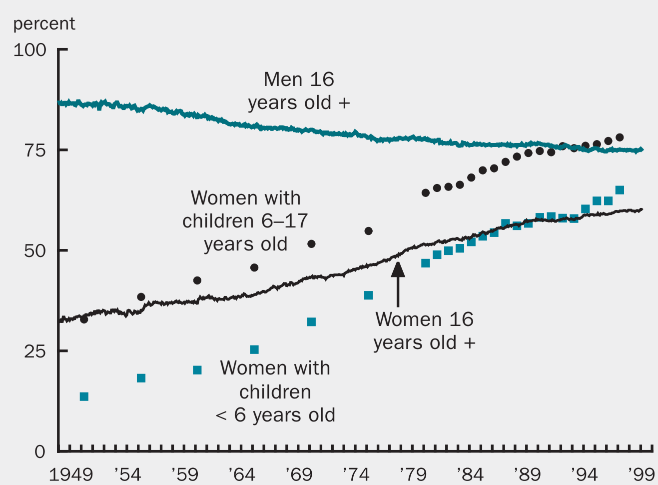

Source: U.S. Department of Labor, Bureau of Labor Statistics, 1948–99, *Current Population Survey*, Washington, DC.

are displayed in figure 1. The labor force participation rate for all women ages 16 and over has nearly doubled from 32 percent in 1948 to 60 percent in 1999. In comparison the labor force participation rate for men ages 16 and over decreased from 87 percent in 1948 to 75 percent in 1999. Over the same period, participation rates for women with preschool-aged children and women with school-aged children have increased even more dramatically. For women with children under six years old, labor force participation increased from 14 percent in 1950 to 65 percent in 1997. Similarly, women with children ages six to 17 increased labor force participation from 33 percent in 1950 to 78 percent in 1997.[6]

Child care worker wages as a measure of child care costs

Because I cannot observe the actual price of child care faced by the women in my sample, I use average child care worker wages across states and over time as a proxy for child care costs. Child care worker wages are likely to be a major portion of the cost of providing child care. One would expect to see differences in the cost of child care across states due to differences in minimum wage levels and in the supply of low-wage labor, among other factors. Because these differences may change over time, I calculate measures of child care costs by state and year. Differences in child care costs across states could also arise because of differences in demand for child care. However, if states in which more women work have higher child care costs because there is more demand for child care, this will bias the estimates against finding the expected negative effect of child care cost on the probability a woman returns to work after first birth.

I calculate average hourly wages for child care workers by state and year for 1979 to 1993 from the National Bureau of Economic Research's *Current Population Survey* (CPS), Labor Extract, Annual Earnings File Extracts (National Bureau of Economic Research, 1979–93). The average is the weighted average of hourly earnings of all surveyed workers who report a three-digit occupation code for child care workers, private households, or for child care workers, except private households.[7] Hourly earnings are calculated as edited hourly earnings when paid hourly and edited or computed usual weekly earnings divided by edited usual weekly hours otherwise. Hourly earnings less than $0.50 and above the 99th percentile in each year are dropped.[8]

Nationally, real average child care worker wages increased over the period 1979–93. Average child care worker wages and average wages for all women are shown in figure 2. Wages for child care workers and average wages for all women both increased in real terms from 1979 to 1993. From figure 2, one can see that average child care worker wages were increasing faster than average wages for women, particularly over 1984–91. From 1979 to 1993, average women's wages increased by 9 percent, adjusted for inflation, while average child care worker wages rose by 22 percent.[9]

FIGURE 2

Average hourly wages

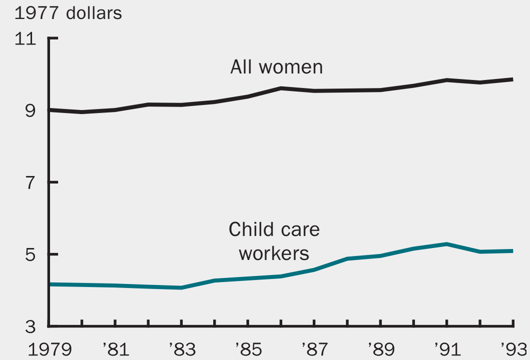

Source: U.S. Department of Labor, Bureau of Labor Statistics, 1979–94, *Current Population Survey*, Washington, DC.

Table 1 lists average child care worker wages by state for 1979–93. As one might expect, states or districts that had state minimum wages above the federal minimum wage throughout the 1980s such as the District of Columbia, Alaska, and Connecticut have higher than average child care worker wages over the period. Hawaii, Massachusetts, Rhode Island, and California did not raise their state minimum wages above the federal minimum wage until 1988, but they, too, have above-average child care worker wages over the period. Likewise, it is not surprising to find that West Virginia, Indiana, Idaho, and North Dakota, where wages are relatively low, have below-average child care worker wages.

Model description

To model women's return-to-work decisions, I assume that each woman has a reservation wage, that is, a "threshold" wage at which she would be willing to go back to work.[10] The probability that a woman returns to work is the probability that her wage offer net of child care costs exceeds her reservation wage. Thus, higher child care costs and lower wage offers will decrease the probability that a woman will go back to work. In addition, assuming that increases in

TABLE 1

Average child care worker wages by state, 1979–93

State	Average wage	State	Average wage
District of Columbia	6.45	Mississippi	4.35
Alaska	6.19	Vermont	4.34
Hawaii	5.82	Kentucky	4.33
New Jersey	5.79	Minnesota	4.32
Massachusetts	5.44	Arizona	4.28
Rhode Island	5.42	Tennessee	4.25
Connecticut	5.38	Alabama	4.22
New York	5.37	Missouri	4.20
California	5.31	Ohio	4.19
Nevada	5.23	Utah	4.13
New Hampshire	4.98	Arkansas	4.08
Maryland	4.98	Virginia	4.07
Georgia	4.94	Kansas	4.04
Florida	4.94	Michigan	3.99
Texas	4.80	Oregon	3.91
Oklahoma	4.78	South Dakota	3.91
Illinois	4.64	Maine	3.78
Delaware	4.60	Montana	3.72
New Mexico	4.56	Wisconsin	3.62
Pennsylvania	4.54	Nebraska	3.61
Louisiana	4.53	Iowa	3.50
Washington	4.46	West Virginia	3.48
Wyoming	4.43	Indiana	3.44
Colorado	4.42	Idaho	3.40
South Carolina	4.39	North Dakota	3.38
North Carolina	4.35	All states	4.58

Notes: Averages are reported in real 1997 dollars. Averages are the weighted average by state (or over all states) of hourly earnings of all surveyed workers in the 1979–93 NBER *CPS Annual Earnings File Extracts* who report a three-digit occupation code for child care workers, private household or for child care workers, except private households. Hourly earnings less than $0.50 and above the 99th percentile for each year are excluded.
Source: Author's estimates from National Bureau of Economic Research, 1979–93, *CPS Annual Earnings File Extracts*.

income increase the number of hours of leisure a person wants to consume, higher other family income will also decrease the probability of returning to work.

My empirical strategy is to study the determinants of the return-to-work decision for new mothers who were working prior to the birth of their first child. I limit the sample to women giving birth to their first child for simplification of the return-to-work decision. This group is more uniform in the sense that all mothers face a first birth but not all will face a subsequent birth. Additionally, these women are all facing the decision to return to work with the need to hire child care for a child under age one only, not for multiple children at various ages. Limiting the sample to women who worked in the year before birth defines a more homogenous group of women, since they all exhibit at least some attachment to the labor force prior to their first birth. This also allows me to use pre-birth wage information as a proxy for post-birth offered wages.

Data and estimation

NLSY data

The original NLSY sample contains 5,842 women, excluding the military sample that was dropped in 1985.[11] In this study, I primarily use the 1994 NLSY child file, which provides detailed information on the children of the original NLSY sample women, including some relevant information on their mothers. In addition, I use the 1993 NLSY youth file to get geographic and family income information for the mothers. According to the 1994 child file, there are 3,468 women whose first child was born between 1979 and 1994 and resided in the mother's household the first year of birth.[12] Characteristics of these women are reported in the first column of table 2.[13]

The NLSY reports the number of weeks before and after birth that a woman left and resumed employment. The women of the NLSY have high employment rates before giving birth; 76 percent of all mothers were working within 51 weeks prior to their first child's birth. Although the participation rate is high relative to the overall participation rate for women, this reflects in part the relatively young age of the NLSY women and, more generally, the age of women at the time of their first birth. The national rates are calculated for women ages 16 years and over, while the average age at first birth for NLSY women is 23 years. Nationally, the labor force participation rate for women in their early twenties is around 73 percent.[14]

Means and standard deviations for characteristics of the regression sample are presented in column 2 of table 2. The sample is limited to women who were working before the birth of their first child and women with complete data on variables used in the regression analysis. The women who were working prior to giving birth tend to have higher other family income and are older (24 versus 21 years old) and better educated (12.9 versus 11.2 years of education).

As shown by the variable in row 2 of table 2, 76 percent of the mothers who were working returned to work within 51 weeks following their child's birth. A more detailed picture of the process is provided in figure 3, which shows the fraction of the sample from column 2 of table 2 who were working in each week before and after childbirth. Expectant mothers gradually withdraw from employment in the months before their delivery and then gradually return.[15] The pattern for the full sample of NLSY women in column 1 of table 2 is very similar to that of the regression sample.

In addition to the standard variables included in a labor force participation equation—wages, unemployment rates, age, education, and race—I include an indicator for the mother having had a working female role model when she was 14 and one for the presence of a woman's parent, step-parent, or grandparent in the household around the birth year. The role model variable is intended to help capture a woman's attitude about being a working mother. Although a woman may have different feelings about

TABLE 2

Mean characteristics for returners and non-returners

Description	Full sample	Regression sample	Return: Yes	Return: No	t-value
Worked within 51 weeks before first birth	0.765 [0.424]	1 [0]	1 [0]	1 [0]	—
Working within 51 weeks after first birth	0.616 [0.486]	0.762 [0.426]	1 [0]	0 [0]	—
State average wage for child care workers	4.506 [0.858] N = 3,302	4.559 [0.890]	4.560 [0.897]	4.555 [0.869]	0.1
Hourly wage fourth quarter before birth	9.274 [5.040] N = 2,237	9.221 [4.954]	9.666 [5.177]	7.797 [3.836]	8.4***
Spouse or partner present	0.781 [0.414]	0.820 [0.384]	0.839 [0.368]	0.760 [0.428]	3.6***
Spouse or partner income	19,430 [32,625] N = 3,207	23,840 [35,558]	23,970 [32,623]	23,422 [43,677]	0.3
Mother's age in years at child's birth	23.234 [4.201]	24.190 [3.990]	24.512 [3.918]	23.159 [4.047]	6.4***
Mother's education in years by birth year	12.416 [2.293] N = 3,466	12.942 [2.129]	13.143 [2.118]	12.298 [2.038]	7.7***
Adult female role model worked when mother was 14	0.524 [0.499]	0.537 [0.499]	0.552 [0.497]	0.491 [0.500]	2.3**
Parent, step-parent, or grandparent of mother resides in household in birth year	0.300 [0.458] N = 3,395	0.245 [0.430]	0.223 [0.417]	0.313 [0.464]	3.7***
African-American	0.228 [0.420]	0.194 [0.395]	0.195 [0.397]	0.189 [0.392]	0.3
County unemployment rate in year following birth	8.066 [3.327] N = 3,159	7.793 [3.288]	7.620 [3.166]	8.343 [3.598]	3.9***
Observations	3,468	1,956	1,490	466	

Notes: All means are unweighted. The number of observations, N, is noted where different from the base sample size. Wages and income are in real 1997 dollars. Standard deviations are in brackets. ***Indicates statistically different from 0 at the 1 percent significance level; and ** indicates statistically different from 0 at the 5 percent significance level.

Source: Author's calculations using data from the Center for Human Resource Research, 1993 and 1994, *National Longitudinal Survey of Youth*, Columbus, OH working when she has young children versus when her children are teenagers, this is the only information available on whether a woman lived in a household with a working female role model. The "grandparent" indicator is included to reflect a woman having greater access to low-cost child care. As shown in rows 9 and 10 of table 2, 52 percent of the NLSY women's role models worked when they were 14, and 30 percent of the overall sample of new mothers lived with their own parent, step-parent, or grandparent.

Columns 3 and 4 of table 2 show the characteristics of women in the regression sample who were and were not back at work within a year of childbirth. A simple comparison across the columns suggests that women with higher wages, those with a spouse or partner, older women, those with more education, and those whose mother worked are more likely to return to work quickly. Column 5 presents absolute t-values for the hypothesis that the means in column 3 equal the means in column 4. As predicted by the model, women who return to work have higher wages on average; however, differences in average child care costs and in average other family income for returners and non-returners are not statistically significant. The differences in age, education, working female role model, and unemployment rates are statistically significant. Women who return to work are older, more educated, more likely to have had a working role model, less likely to live with a parent or grandparent, and are living in counties with lower average unemployment rates.

The employment pattern illustrated by figure 3 suggests estimating a more "dynamic" model of weeks to return to work following birth such as a *tobit* or *hazard* model. The results from estimating a tobit model of weeks to return to work censored at 52 weeks, although not reported in this article, are consistent with the *probit* estimates discussed below. Women with higher wages and more education return to work more quickly following birth, and women facing higher child care costs and having higher other family income delay their return to work longer after birth. This should not be surprising, however, since none of the variables vary over the weeks following birth.

FIGURE 3

Labor force participation, NLSY women

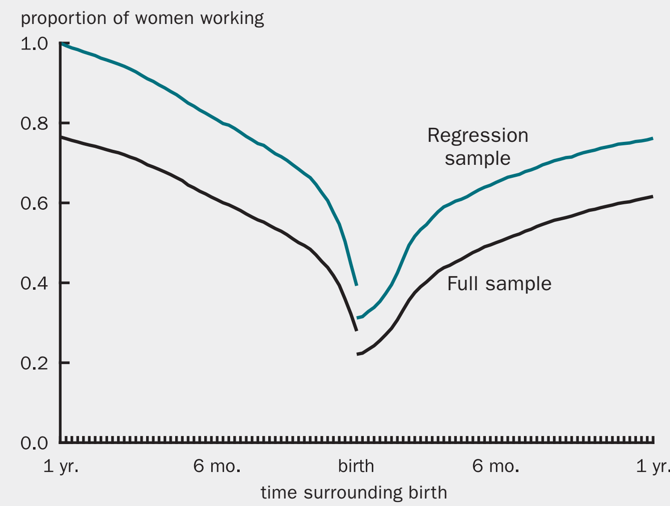

Source: Author's calculations using data from the Center for Human Resource Research, 1993 and 1994, *National Longitudinal Survey of Youth*, Columbus, OH.

Probit estimation of the probability a woman returns to work following first birth

As discussed above, I assume each woman has a reservation wage at which she is willing to go back to work. As modeled, the offered wage and child care costs affect the net wage and thus the probability that the net offered wage exceeds the reservation wage, while some of the other characteristics are expected to affect a woman's reservation wage. The *probit* model estimates the probability of returning to work as a function of offered wage, child care costs, other family income, and demographic and labor market characteristics. The estimation equation is as follows:

$$1)\quad Pr[\textit{working 1 year after birth}] = \beta_0 + \beta_1 \textit{wage} + \beta_2 C + Z\beta_3 + \beta_4 UR - \varepsilon,$$

where *wage* is the wage in the fourth quarter before birth,[16] C is the child care cost variable, Z is a matrix including age, education, other family income, and indicator variables for having a spouse or partner, having a working female role model, being African-American, and having one of the child's grandparents in the household, and UR is the county unemployment rate in the year following the birth year.

First, I estimate the model specified in equation 1. These results are presented in table 3. I report the change in probability of returning to work within one year of birth associated with a change in each independent variable.[17] For example, increasing the average child care worker wage by $1 decreases the probability that the average woman will return to work within one year of her child's birth by 0.038, from 0.778 to 0.740.[18] Thus, as predicted by the simple utility

TABLE 3

Probit estimates of labor force participation model

Independent variable	Associated change in probability of returning to work within 1 year
Child care worker wage	–0.038***
	(0.012)
Pre-birth wage	0.017***
	(0.003)
Spouse or partner income divided by 10,000	–0.012***
	(0.003)
Indicator for spouse or partner	0.089***
	(0.035)
Mother's age in birth year	0.007**
	(0.003)
Mother's education at birth year	0.017***
	(0.006)
Role model work	0.041**
	(0.019)
Grandparent	0.001
	(0.027)
African-American	0.049*
	(0.026)
Unemployment rate in year following birth	–0.007**
	(0.003)

Note: The dependent variable is an indicator for returning to work within one year of giving birth to the first child. The probability of returning to work predicted at the mean characteristics of the women in the sample is 0.778. The reported estimate is the change in probability of returning to work associated with a one unit change in a given variable, evaluated at the mean of the characteristics. For example, a $1 increase in the average child care worker wage is associated with a 0.038 decrease in the probability a woman returns to work, a decrease from 0.778 to 0.740. There are 1,956 observations. Standard errors are in parentheses. ***Indicates statistically different from 0 at the 1 percent significance level; **statistically different from 0 at the 5 percent significance level; and *statistically different from 0 at the 10 percent significance level.

maximizing model described above, women who live in states with higher child care costs, proxied by child care worker wages, are significantly less likely to return to work within one year of giving birth to their first child. In addition, lower wage women are less likely to return to work within one year of giving birth, as are women with higher partner or spouse income, controlling for the presence of a spouse or partner.[19] Older women, women with more education, and those who had a working female role model are all more likely to return to work after giving birth.

The theoretical model predicts that offered wage and hourly child care price should have coefficients equal in magnitude and opposite in sign. In comparing the wage and cost coefficients, the wage is measured in pretax dollars while child care expenditures are in after-tax dollars. In addition, the child care cost measure is the hourly child care worker wage rather than the hourly price. Given the Census Bureau estimates from the SIPP cited above, one would expect hourly child care costs to be at most 54 percent of average child care worker wages.[20] Assuming that the hourly cost of child care equals 54 percent of average child care worker wages, the tax rate would have to be in excess of 75 percent to generate the observed change in probability associated with a $1 change in the offered wage. This result can be partially reconciled if other costs of working are correlated with child care costs. If other costs of working are positively correlated with child care costs, then the effect of child care costs on the probability of returning to work is overstated.

Spouse/partner income affects women's probability of returning to work as predicted by the model: The higher a woman's spouse/partner income, the less likely she is to return to work. If other income is allowed to enter separately for women with spouses and women with partners, the decrease in probability associated with a $10,000 increase in spouse income is 0.011 with a standard error of 0.003; that is, the probability a woman will return to work falls from 0.778 to 0.767. Similarly, a $10,000 increase in partner income is associated with a decrease in the probability of returning to work from 0.778 to 0.752. Finally, 66 women with spouses or partners have other income calculated to be $0. When these observations are excluded, average child care worker wages becomes slightly more important. The change in probability associated with a $1 change in child care worker wages falls to –0.040 with a standard error equal to 0.012; that is, a decrease in probability from 0.778 to 0.738 is associated with a $1 increase in the average child care worker wage. The changes associated with other income, the spouse/partner indicator, age, female role model, the grandparent indicator, and African-American increase in magnitude, and the education coefficient decreases slightly.

The results presented in table 4 explore the possibility that women who are most like welfare recipients may differ from other women in their sensitivity to child care costs as well as to other economic variables, in particular, the unemployment rate. I try two measures for similarity to welfare recipients: education less than 12 years at child's birth and the combination of both being unmarried and having fewer than 12 years of education at child's birth. Columns 1 and 2 of table 4 list probit estimates using the education indicator only, while columns 3 and 4 use the joint indicator of education and marital status. Columns 1 and 3 present the results allowing for differing sensitivity to child care costs. In both specifications there is little evidence that either less educated women or less educated women without a spouse present are any more sensitive to child care costs than all women in the sample. While the estimated change in probability of returning to work associated with a $1 change in hourly child care costs for the women most like welfare recipients is smaller than for all other women, the difference is not statistically significant at conventional levels. Similarly, their probability of returning to work is not significantly more responsive to higher unemployment rates as shown in columns 2 and 4.

The calculated child care cost, wage, and family income elasticities of employment provide one way to compare the results of this study to others.[21] The elasticity is the percent change in probability associated with a 1 percent change in a given variable. The specification of table 3 implies a child care cost elasticity of –0.23. In other words, a 1 percent increase in child care cost is associated with a 0.23 percent decrease in the probability of returning to work.[22] This estimate is similar to the average price elasticity of employment of –0.20 estimated by Connelly (1992a), but somewhat smaller than estimates from many other studies. Blau and Robins (1988) calculate a price elasticity of employment of –0.38 over a range of child care costs, Kimmel (1993) calculates an elasticity of –0.31 for married women using her preferred child care cost measure, and Powell (1997) calculates an elasticity of –0.38 for married women using predicted cost of child care. The elasticities calculated by Anderson and Levine (1998) for women with children under six years are also much larger, between –0.46 and –0.59. Ribar (1995) calculates a much smaller elasticity of –0.09, while that of Ribar (1992) is much higher at –0.74. The wage elasticity of labor force participation is much smaller, at 0.21,

TABLE 4

Probit estimates of labor force participation model, by education and marital status

	No high school diploma		No spouse and no high school diploma	
Indicator	–0.194 (0.167)	–0.073 (0.078)	–0.203 (0.219)	–0.060 (0.112)
Child care worker wage	–0.040*** (0.013)	–0.036*** (0.012)	–0.036*** (0.013)	–0.036*** (0.012)
Child care worker wage interacted with Indicator	0.021 (0.030)	—	0.010 (0.036)	—
Unemployment rate in year following birth	–0.007** (0.003)	–0.007** (0.003)	–0.007** (0.003)	–0.006** (0.003)
Unemployment rate interacted with Indicator	—	–0.001 (0.007)	—	–0.008 (0.011)
Pre-birth wage	0.017*** (0.003)	0.017*** (0.003)	0.017*** (0.003)	0.017*** (0.003)
Spouse or partner income divided by 10,000	–0.011*** (0.003)	–0.011*** (0.003)	–0.011*** (0.003)	–0.011*** (0.003)
Indicator for spouse or partner	0.086*** (0.035)	0.085*** (0.035)	0.055 (0.036)	0.054 (0.036)

Note: The dependent variable is an indicator for returning to work within one year of giving birth to the first child. Each equation also includes the additional covariates listed in table 3. The reported estimate is the change in probability of returning to work associated with a one unit change in a given variable, evaluated at the mean of the characteristics. Columns 1 and 3 present the results allowing for differing sensitivity to child care costs. There are 1,956 observations. Standard errors are in parentheses. ***Indicates statistically different from 0 at the 1 percent significance level; and ** indicates statistically different from 0 at the 5 percent significance level.

than those estimated by Ribar (1992 and 1995) of 0.68 and 0.53, Kimmel (1993) of 0.58, Powell (1997) of 0.85, and Anderson and Levine (1998) of 0.58, but larger than the 0.04 calculated by Michalopoulos, Robins, and Garfinkel (1992).[23] Finally, the other income elasticity of –0.04 is very similar to the estimates of Michalopoulos, Robins, and Garfinkel (1992) and Ribar (1995), of –0.01 and –0.05, respectively.

Although more education seems to increase the probability that a woman will return to work after first birth, this result has several possible interpretations. It may be that women who get more education do so because they are more committed to the labor force and thus are more likely to go back to work. Alternatively, it may be that women with more education are more likely to hold jobs from which they can take leave as opposed to having to quit and, hence, they face lower costs of returning to work after birth.[24] Finally, this may be reflecting part of the wage effect due to the high correlation of education with wages and possible measurement error in the wage variable.

I include the working female role model variable to capture the idea that women may have different views about the appropriateness of working when they have children. Although a woman may view working when she has a young child differently than when she has a child aged 14, this is the only role model information available. Across all estimated equations, this variable has a consistent positive and significant coefficient. One might be concerned that this variable is reflecting an inter-generational correlation in income status rather than a role model effect per se. For example, poor women may be more likely to work, and their children may be more likely to be poor and, hence, also more likely to work. However, including other family income should help control for wealth, and the role model coefficient remains virtually unchanged when unearned income is excluded.

As for other variables in the model, older women are more likely to return to work after birth, although again this may partially be picking up the wage effect. Contrary to expectations, having a parent or grandparent in the household does not seem to affect the reemployment rate, suggesting that parents and grandparents may not serve as a major source of child care. While having a parent or grandparent in the home and the decision to return to work may be simultaneously determined, omitting the grandparent indicator does not change the coefficient estimates significantly. A better indicator of access to lower cost child care would be a measure of having relatives in close proximity, but this information is only available for one year of the NLSY. Finally, at the 10 percent level of significance, African-American women in this sample are more likely than other women to go back to work, and higher county unemployment rates reduce the probability that a woman returns to work after first birth.

Implications of the estimates

Using the table 3 results to explore some of the implications of the estimates, I simulate the effects of various factors on the probability of returning to work. Based on SIPP data, weekly expenditures on child care for families with a preschool-aged child increased 23 percent from 1986 to 1993. Considering a potential increase in child care subsidization that would reduce hourly costs by 20 percent, the probability of returning to work increases by 3 percentage points. If I assume these results hold for all women of childbearing age, this would lead to an expected increase in the labor force of 1.8 million workers.[25]

Next, as women delay child bearing they are more likely to return to work quickly, holding wages constant. Since wages generally increase over those years of delayed child bearing, older mothers will have an additional tendency to return to work quickly due to the higher opportunity cost of not working. On average the probability of returning to work is 0.78. The probability of returning for a 24-year-old (the median age at first birth) earning the average wage of 24-year-old mothers in this sample is 0.77. For a 27-year-old mother (the seventy-fifth percentile age at first birth in the sample) earning average wages for a 27-year-old in this sample, the probability increases to 0.83.[26]

From 1988 to 1991 the proportion of preschoolers being cared for by their fathers rose from 15 percent to 20 percent.[27] This number fell back to 16 percent in 1993, according to the most recent census report.[28] As suggested by the Census Bureau, this temporary rise in the percentage of children being cared for by their fathers in 1991 may be attributed to higher unemployment and underemployment of fathers. This is consistent with the possibility that worsening employment opportunities for women's spouses and partners during part of the sample period encouraged more women to return to work sooner after childbirth. For a high-wage woman (wage at the seventy-fifth percentile) with a high level of other family income (at the seventy-fifth percentile), the probability of returning to work in the first year is 0.80. If instead she faces low other family income (in the twenty-fifth percentile), the probability she returns within a year rises to 0.83.

Finally, from January 1992 to January 1999, t he unemployment rate in the U.S. dropped from 7.3 percent to 4.3 percent. The probability the average woman returns to work when the unemployment rate is 7.3 percent is 0.78. When the unemployment rate drops to 4.3 percent, the probability of returning to work rises to 0.80.

These estimates suggest that delayed child bearing will play a much more important role in increasing women's labor force participation shortly after childbirth, and, hence, their overall actual work experience accumulation, than small increases in child care cost subsidization, the effects of changing employment opportunities for their spouses and partners, or decreases in the overall unemployment rate. Another interesting long-term implication of the increased labor force participation of mothers today is that their daughters may also be more likely to participate in the labor force. Thus, we should expect to see continued participation rate increases with new generations of women entering the labor force.

Conclusion

This article examines the effects of child care costs, potential wages, and other family income on a woman's decision to return to work shortly following the birth of her first child. Utility maximization predicts that child care costs and other family income will have a negative effect on the probability of returning to work, while potential wages will have a positive effect. A simple comparison of means of cost, wages, and other income for returners and non-returners shows differences as predicted by the model that are significant for the wage measure. Further multivariate analysis confirms these results for wages and indicates that child care costs and other family income also have statistically significant effects on the probability of returning to work. The estimates suggest that the elasticity of the reemployment rate for new mothers with respect to child care costs is about –0.23, while the elasticity with respect to other family income is about –0.04. Finally, the elasticity with respect to the mother's wage is about 0.21. Additionally, age and education, having a spouse or partner, having had a working female role model, and living in areas with lower unemployment rates have statistically significant, positive effects on the probability that a woman returns to work.

As mentioned in the introduction, the results of this study have implications for evaluating policy. The results suggest that delayed child bearing may have a greater impact on increasing labor force participation of women with young children than increases in wages or decreases in child care costs. Additionally, while access to reliable child care is likely to be a necessity for successfully moving mothers from welfare to the labor force, this research shows no evidence that welfare recipients will be more responsive to changes in child care costs than other women. Moreover, the overall estimate of responsiveness to changes in child care costs does not indicate that such changes will lead to large changes in labor force participation. Thus, increasing subsidization of child care without additional programs and incentives is not likely to have large effects on labor force participation among the welfare population. Finally, the increased probability of a woman working after childbirth associated with her female role model having worked suggests that we should expect to see continuing increases in the labor force participation rate of women, thus increasing the size of the labor force.

APPENDIX

Theoretical model

I model a woman's return-to-work decision as a utility maximization problem with child care expenditures entering the budget constraint and, hence, affecting the employment decision. First, I assume a woman makes her labor force participation decision by maximizing her utility, taking her husband's labor force participation and income as given.[1] Her problem is to maximize:

$$U(X, D, L) \text{ s.t. } \begin{array}{l} (a)\, p_x X + p_d D \le wH + Y \\ (b)\, H + L = T \\ (c)\, 0 \le H \le T,\ 0 \le L \le T, \end{array}$$

where X is a composite good excluding day care and leisure, p_x is the price of X, D is the hours of day care demanded, p_d is the hourly price of day care, H is the number of hours the woman works, w is the wage rate, Y is her husband's income plus other unearned income, T is the total time constraint, and L is the number of leisure hours.[2] In modeling the decision this way, I am implicitly assuming that maternal and market child care are good substitutes.

Assuming additionally that $H < T$ and $D=H$, the optimization problem can be written,[3]

$$2)\quad \mathscr{L} = U(X,\ L) - \lambda[p_x X + (w - p_d)L - ((w - p_d)T + Y)] + \delta(T - L),$$

with the associated conditions:

$$(a)\quad U_1 - \lambda p_x = 0,$$
$$(b)\quad U_2 - \lambda(w - p_d) - \delta = 0,$$
$$(c)\quad \lambda[p_x X + (w - p_d)L - ((w - p_d)T + Y)] = 0, \text{ and}$$
$$(d)\quad \delta(T - L) = 0,$$

where $\lambda > 0$ is the marginal utility of wealth and δ is a non-negative slack variable associated with the woman's hours of work decision. From condition (b), $w - p_d = U_2/\lambda - \delta/\lambda$. Calling U_2/λ the reservation wage, $w^*(H)$, the first-order condition can be rewritten as $w - p_d = w^*(H) - \delta/\lambda$. If the woman works, $\delta = 0$, the net wage exceeds the reservation wage evaluated at $H = 0$, and hours of work are chosen such that $w - p_d = w^*(H)$ when $H > 0$.

For simplicity, I assume a utility function consistent with linear labor supply,

$$3)\quad H_i = \beta_1(w_i - p_{di}) + \beta_2 Y_i + Z_i\beta_3 + \gamma_i,$$

for individual i, where Z is a vector of demographic characteristics and γ is an error term. The linear labor supply function restricts the coefficient on the wage net of child care costs to be the same regardless of the level of the wage. This is the easiest form to model empirically; however, given that my measure of cost is an index of the true cost of child care, I do not impose the additional restriction during estimation that the coefficients on wages and costs are equal. Substituting the budget constraint into equation 3 and solving for the reservation wage,

$$w_i^*(0) = \alpha_1 Y_i + Z_i\alpha_2 + \mu_i,$$

where $\alpha_1 = -\beta_2/\beta_1$, $\alpha_2 = -\beta_3/\beta_1$, and $\mu_i = -\gamma_i/\beta_1$. The probability that a woman works can be represented by

$$\Pr(H_i > 0) = \Pr\left[(w_i - p_{di}) > w_i^*(0)\right] =$$
$$\Pr\left[\mu_i < (w_i - p_{di}) - \alpha_1 Y_i - Z_i\alpha_2\right].$$

Thus, higher child care costs and lower wages decrease the probability that a woman will go back to work. Assuming that leisure is a normal good, higher other family income also decreases the probability of returning to work.

An important consideration is that there may be unobserved taste shifters that have not been specified in the model. For example, let τ reflect taste for work and enter the model by affecting the marginal rate of substitution between leisure and money, that is, let $U = U(X, \tau^{-1}L)$. Condition (b) of equation 2 then becomes $w - p_d = (\tau^{-1})U_2/\lambda - \delta/\lambda$, where $\delta = 0$ if a woman works. The greater the taste for work (the greater τ), the lower the net wage needed to exceed $(\tau^{-1})U_2/\lambda$. Thus, correlations between τ and wages or child care costs can lead to biased estimates of their effects on the probability of returning to work.

Data

Child care cost measure

The state average child care worker wage is the weighted average by state and year of hourly earnings of all surveyed workers in the 1979–93 NBER CPS Annual Earnings File Extracts who report a three-digit occupation code for child care workers, private households, or for child care workers, except private households. Hourly earnings are calculated as hourly earnings where reported and as edited usual weekly earnings divided by edited usual weekly hours, otherwise. Hourly earnings less than \$0.50 and above the 99th percentile for the year are dropped. Weights used are the earnings weights provided in the CPS data.

NLSY data

The wage and employment data before and after birth and mother's age at birth come from the NLSY 1994 child file and were constructed or measured in relation to the birth of the child. The pre-birth wage is the wage recorded for the fourth quarter before birth, and the post-birth wage is the wage recorded for the fourth quarter after birth. All wages are in real 1997 dollars. Wages less than \$1 and greater than \$160 are recoded to missing. Other variables are from the youth file and relate to the survey year which may or may not match up well with the birth year, depending on the month of birth. For determining the usual residence of the child, I count the child as living with the mother if his or her usual residence is coded as in the mother's household either in the survey year of the birth year or in the survey year after the birth year. Similarly, a spouse or partner or mother's mother, grandmother, stepmother, father, grandfather, or stepfather is present if the mother reports so either in the birth year or in the survey year following the birth year. Mother's education is the highest grade completed in the survey year of the birth year or the most recent available record from previous years, since the variable is missing unless the status has changed from the previous year. If highest grade completed is ungraded, it is considered missing.

The unemployment rate data in the youth geographic data are county unemployment data from the *County and City Data Book*. The unemployment rate at birth is measured as the unemployment rate in the birth year, and the unemployment rate after birth is measured as the unemployment rate in the survey year after the birth year. The state of residence is the residence reported in the survey year of the birth year unless the code is missing, in which case it is the state reported in the survey year following the birth year. The child care cost variable is then matched by these state codes.

From 1979 to 1989, respondents were asked for total income for their partner in the previous year. After 1989 respondents were asked for partner income broken down into several categories. Spouse income for all years is reported broken down into several categories. Other income for women with partners from 1979 to 1989 is partner income as reported in the following survey year. Other income for women with spouses for all years is calculated as annual spouse income from wages and salary, plus any farm or own business income, plus spouse unemployment compensation, plus respondent or spouse income from food stamps and other sources. Other income for women with partners from 1990 to 1993 is calculated as total partner income from wages and salary, plus any farm or own business income, plus partner's total welfare income. To minimize the loss of observations from missing information, other income is used as calculated for the year of the birth or the year after birth. All income is top-coded at $75,001 for 1979–84 and at $100,001 for 1985–93. Income is in real 1997 dollars.

---

[1]The validity of this assumption is certainly debatable, and future analysis could model the labor supply decisions of a woman and her spouse/partner as a joint decision.

[2]Below, I assume a linear labor supply function. See Stern (1986) for a discussion of the form of the utility function and the implications of the assumption.

[3]I assume that day care is specifically purchased to cover hours worked and that a woman's leisure time includes time she spends caring for her children. Certainly, women may hire child care during their leisure hours, but I consider these nonwork child care hours to be a separate good included in the composite good.

NOTES

[1]Shapiro and Mott (1994) provide some evidence that labor force participation surrounding first birth is an important predictor of a woman's later labor force participation behavior, and hence greater actual work experience at all points in life.

[2]Blau and Kahn (1992).

[3]See Nakamura and Nakamura (1992) for a review of some of the literature analyzing the effect of children on female labor supply more generally. See Leibowitz and Klerman (1995) for a more recent paper looking at the effects of children on married mothers' labor supply over time.

[4]U.S. Department of Commerce, Bureau of the Census (1992). Mean expenditures are conditional on making positive child care payments and have been converted to real 1997 dollars.

[5]Much of this article is based on Barrow (1999).

[6]While it appears that women with school-aged children have higher labor force participation rates than men, this is a function of the difference in the age distribution of all men versus women with school-aged children. The participation rate for men with school-aged children is 93 percent (U.S. Department of Labor, Bureau of Labor Statistics, unpublished data).

[7]Weights used are the earnings weights provided in the CPS data.

[8]205 observations were dropped, leaving 20,080 wage observations for child care workers in 50 states and one district over 15 years.

[9]Approximately 95 percent of child care workers in the CPS data are women.

[10]See the appendix for a more formal description of the model.

[11]The NLSY is a nationally representative sample of 12,686 men and women who were between the ages of 14 and 21 in 1979, including a military sample and an oversample of African-Americans, Hispanics, and poor non-African-Americans and non-Hispanics. See Center for Human Resource Research (1989 and 1993) for more information on the survey.

[12]For the 918 women with first births before 1979, there are no birth year data available.

[13]The appendix contains more details of how the dataset is constructed.

[14]U.S. Department of Commerce, Bureau of the Census (1998), table No. 645. In 1997 the participation rate for women ages 16 to 19 was 51.0 percent, the rate for women 20 to 24 was 72.7 percent, and the rate for women 25 to 34 was 76.0 percent.

[15]Although a larger percentage of NLSY women return to work after first birth, the employment patterns are very similar to those of the *National Longitudinal Survey of Young Women* presented in McLaughlin (1982)

[16]Pre-birth wage is the best approximation I have of the wage women actually face when making their return-to-work decision. Because I am looking at these women over such a short time frame, I assume that there is minimal wage erosion.

[17]For continuous variables, this is the change in probability associated with an infinitesimal change in the independent variable, while for discrete variables it is the change associated with a one unit change in the independent variable.

[18]0.778 is the predicted probability of returning to work for a woman with the characteristics of the average woman in the sample. The predicted change in probability is calculated at this mean.

[19]Very few observations are affected by the income top-coding, and including an indicator for the presence of a top-coded income measure has no important effects on the results; however, women whose spouse or partner income is top-coded are significantly less likely to return to work.

[20]As noted above, U.S. Department of Commerce, Bureau of the Census (1992) estimates women with at least one child under age one spend an average of $88.60 on child care per week and work an average of 36 hours per week. This $2.46 per hour cost in 1997 dollars is 54 percent of the mean state average child care worker wage of $4.58 per hour.

[21]Elasticities are only available from a subset of the studies for a subset of the elasticities of interest. In the text I cite all studies for which an elasticity calculation is available.

[22]Elasticities are calculated at the mean employment rate and the mean average child care worker wage across observations.

[23]Even if mother's age and education at child's birth are omitted from the estimation, the wage coefficient is never large enough to generate an elasticity as large as the cited studies.

[24]The Family and Medical Leave Act of 1993 became effective after most of the women in the NLSY sample gave birth to their first child. This act allows workers at companies with more than 50 employees to take up to 12 weeks of "job-protected" leave to care for a child or other immediate family member, lowering the cost for many women of returning to the labor force after childbirth.

[25]I use census population estimates of approximately 60.1 million women aged 15 to 44 as of April 1, 1999.

[26]The probabilities are evaluated at the mean values for all covariates other than the ones being changed for the simulations.

[27]U.S. Department of Commerce, Bureau of the Census (1994).

[28]U.S. Department of Commerce, Bureau of the Census (1996).

REFERENCES

Anderson, Patricia M., and Phillip B. Levine, 1998, "Child care and mothers' employment decisions," paper presented at the "Labor Markets and Less-Skilled Workers Conference," sponsored by the Joint Center for Poverty Research, Washington, DC, November 5–6.

Barrow, Lisa, 1999, "An analysis of women's return-to-work decisions following first birth," *Economic Inquiry*, Vol. 37, No. 3, July, pp. 432–451.

Blau, David M., and Philip K. Robins, 1991, "Child care demand and labor supply of young mothers over time," *Demography*, Vol. 28, No. 3, August, pp. 333–351.

__________, 1988, "Child care costs and family labor supply," *Review of Economics and Statistics*, Vol. 70, No. 3, August, pp. 374–381.

Blau, Francine D., and Lawrence M. Kahn, 1992, "Race and gender pay differentials," in *Research Frontiers in Industrial Relations and Human Resources*, David Lewin, Olivia S. Mitchell, and Peter D. Sherer (eds.), Madison, WI: Industrial Relations Research Association, pp. 381–416.

Center for Human Resource Research, 1993, *NLS Users' Guide 1993*, Columbus, OH: Ohio State University.

__________, 1989, *NLSY Child Handbook 1989*, Columbus, OH: Ohio State University.

Connelly, Rachel, 1992, "The effect of child care costs on married women's labor force participation," *The Review of Economics and Statistics*, Vol. 74, No. 1, February, pp. 83–90.

Kimmel, Jean, 1993, "Child care and the employment behavior of single and married mothers," W. E. Upjohn Institute for Employment Research, Kalamazoo, MI, working paper, No. 93-14.

Klerman, Jacob Alex, and Arleen Leibowitz, 1990, "Child care and women's return to work after childbirth," *American Economics Association Papers and Proceedings*, May, pp. 284–288.

Leibowitz, Arleen, and Jacob Alex Klerman, 1995, "Explaining changes in married mothers' employment over time," *Demography*, Vol. 32, No. 3, August, pp. 365–378.

Leibowitz, Arleen, Jacob Alex Klerman, and Linda J. Waite, 1992, "Employment of new mothers and child care choice," *Journal of Human Resources*, Vol. 27, No. 1, Winter, pp. 112–133.

Michalopoulos, Charles, Philip K. Robins, and Irwin Garfinkel, 1992, "A structural model of labor supply and child care demand," *Journal of Human Resources*, Vol. 27, No. 1, Winter, pp. 166–203.

McLaughlin, Steven D., 1982, "Differential patterns of female labor-force participation surrounding the first birth," *Journal of Marriage and the Family*, pp. 407–420.

Nakamura, Alice, and Masao Nakamura, 1992, "The econometrics of female labor supply and children," *Econometric Reviews*, Vol. 11, No. 1, January, pp. 1–71.

O'Neill, June, and Solomon Polachek, 1993, "Why the gender gap in wages narrowed in the 1980s," *Journal of Labor Economics*, Part I, January, pp. 205–228.

Powell, Lisa M., 1997, "The impact of child care costs on the labour supply of married mothers: Evidence from Canada," *Canadian Journal of Economics*, Vol. 30, No. 3, August, pp. 577–594.

Ribar, David C., 1995, "A structural model of child care and the labor supply of married women," *Journal of Labor Economics*, Vol. 13, No. 3, July, pp. 558–597.

__________, 1992, "Child care and the labor supply of married women: Reduced form evidence," *Journal of Human Resources*, Vol. 27, No. 1, Winter, pp. 134–165.

Shapiro, David, and Frank L. Mott, 1994, "Long-term employment and earnings of women in relation to employment behavior surrounding the first birth," *Journal of Human Resources*, Vol. 29, No. 2, Spring, pp. 248–275.

Stern, Nicholas, 1986, "On the specification of labour supply functions," in *Unemployment, Search, and Labour Supply*, Richard Blundell and Ian Walder (eds.), New York: Cambridge University Press, pp. 143–189.

U.S. Department of Commerce, Bureau of the Census, 1998, *Statistical Abstract of the United States: 1998*, available on the Internet at www.census.gov/prod/3/98pubs/98statab/cc98stab.htm.

__________, 1997, *Statistical Abstract of the United States: 1995*, 115th edition, Washington, DC: U.S. Government Printing Office.

__________, 1996, "Who's minding our preschoolers? (data from the Survey of Income and Program Participation), *Current Population Reports*, Washington, DC: U.S. Government Printing Office, No. P70-53.

__________, 1994, "Who's minding the kids? Child care arrangements: Fall 1991 (data from the Survey of Income and Program Participation)," *Current Population Reports*, Washington, DC: U.S. Government Printing Office, No. P70-36.

__________, 1992, "Who's minding the kids? Child care arrangements: Fall 1988 (data from the Survey of Income and Program Participation)," *Current Population Reports*, Washington, DC: U.S. Government Printing Office, No. P70-30.

U.S. House of Representatives, 1998, *The Green Book*, Washington, DC: U.S. Government Printing Office.

Will a common European monetary policy have asymmetric effects?

Luigi Guiso, Anil K Kashyap, Fabio Panetta, and Daniele Terlizzese

Introduction and summary

The launch of the euro has been accompanied by a vigorous debate. On the one hand, supporters of a common monetary policy (for example, Lamfalussy, 1997) have argued that the move to a single currency is necessary to fully exploit the obvious advantages of a single market. On the other hand, skeptics have argued that European Union (EU) economies are too dissimilar to be subjected to a common monetary policy. Feldstein (1997) goes so far as to predict that the political tensions created by the common monetary policy could lead to another European war.

The debate boils down to a disagreement over how hard it will be to effectively run a common monetary policy. There are at least three conditions that must be met for a common policy to succeed without causing frictions among the members of the coalition. First, members must agree on the ultimate goals to be achieved through monetary policy. This issue was formally settled through the 1992 Maastricht Treaty and the ensuing ratification process by national parliaments, leading to the adoption of a goal of price stability as the primary objective for the European Central Bank (ECB).

Second, the common policy will be easier to implement if the member countries' business cycles are aligned. Monetary policy instruments are macroeconomic variables that work across the board and, therefore, cannot simultaneously be tailored to diverging cyclical conditions in the area of their jurisdiction. However, if different countries or sizable regions are at different points in the inflation cycle, then assessing the appropriate monetary policy stance becomes a much more difficult task. Large countries such as the U.S. constantly confront this problem, but the degree of economic integration and the availability of alternative policy instruments to redistribute the burden of the adjustment are likely to be poorer in the euro area than in the U.S.[1]

A third and perhaps more subtle issue is whether the transmission mechanism operates in a similar fashion across all the countries in the union. In particular, even if shocks hit all countries equally, their business cycles are aligned, and there is no disagreement over whether a response is needed, differences in the transmission mechanism could mean that the appropriate size and timing of the response are difficult to assess. Moreover, if the burden of adjustment is not equally shared across countries, sizable distribution differences are likely to create political tension.

The issue of how much the transmission mechanism differs across the member states of the monetary union is just beginning to attract interest. One obvious difficulty with addressing the question is the possibility of a regime switch that could have occurred with the creation of the ECB. It is possible that all past evidence on the transmission mechanism is no longer relevant because beliefs about policy will now differ.

While we concede that this is possible, we doubt that this institutional change has brought about behavioral changes in a sharp, discontinuous fashion.

Luigi Guiso is a professor of economics at the Università di Sassari, Italy, and scientific coordinator at Ente Luigi Einaudi for monetary, banking, and financial studies. He is also a research fellow at the Centre for Economic Policy Research, London. Anil K Kashyap is a professor of economics at the University of Chicago, Graduate School of Business, consultant to the Federal Reserve Bank of Chicago, and research associate at the National Bureau of Economic Research (NBER). Fabio Panetta and Daniele Terlizzese are deputy directors of the Research Department at the Bank of Italy. *Mr. Kashyap thanks the National Science Foundation for a research grant, administered through the NBER, and the University of Chicago Graduate School of Business for research support. The views in this article do not necessarily reflect those of the Bank of Italy.*

There is abundant evidence that people adjust their behavior gradually. In this case, collecting evidence on how agents operated in the past regime should provide some information on how they will behave in the present one.

Even in the absence of structural breaks, however, trying to conduct the relevant cross-country aggregate comparisons in the transmission mechanism is fraught with difficulties. Research on how to identify the response of a single economy to monetary disturbances in a convincing and robust fashion is just becoming available for some countries. There has been very little work on doing this for multiple countries using a common framework. In particular, to study the effects of how a common monetary policy might matter, one needs to impose a uniform monetary policy reaction function across countries and to constrain exchange rate movements.

Our reading of the existing literature is that this type of study has yet to be done. As a result, we are left with a set of only partially comparable findings, which prevents us from drawing any strong conclusions about the similarities of the transmission mechanisms across European countries. A full investigation of this type would be quite valuable but is beyond the scope of this article.

We believe, however, that the evidence from studies conducted at the aggregate level should be supplemented by systematic comparisons at the micro level. The richness of the information available at the micro level should allow us to identify differences in behavior among different groups of agents in the same country and similar groups of agents in different countries. This is important because aggregate differences could arise for a variety of reasons. One possibility is that similar firms and individuals in different countries could behave differently. In this case, one might believe that as institutional arrangements converge, and the single market is fully realized, the differences could abate. Alternatively, similar firms and individuals might act similarly, but the mix of these agents across countries might differ.

Disentangling true behavioral differences from differences that are the result of compositional effects is important for several reasons: first, because doing so is likely to enhance our understanding of the causes of the differences; second, because this should lead to a better assessment of the likely persistence of any differences; and finally, because this might help identify policy actions that could be used to partially alleviate the differences. Of course, a full investigation of these issues will require several detailed studies. Here, we take a first step and present a sort of "feasibility analysis," aimed at assessing whether what appear to be large structural differences in the economic and financial structures of the various countries in the euro area can be expected to lead to differences in the transmission mechanism.

Our analysis follows three logical steps. First, we try to identify the types of microeconomic heterogeneity that different theories of monetary transmission suggest could be important. The goal here is not to compile any evidence on which of these theories is most important, but rather to use the union of the theories to guide our selection of which cross-country data we ought to compare.

Next, we collect a number of indicators available for multiple countries to demonstrate that, along the dimensions identified in the previous step, there are sharp cross-country differences in the underlying microeconomic landscape of the different EU countries. Theoretically, these firm-level and institutional differences could alter the aggregate impact of monetary policy.

Finally, having identified many potentially important factors suggestive of differences in the transmission mechanism, we turn to data on one specific country, Italy, to see what these factors say about business cycle dynamics and the response to monetary policy shocks in that country. If they were to possess explanatory power in one country, we would read this result as corroborating the basic idea that the structural characteristics of the various economies are relevant factors in explaining cross-country differences and similarities of the transmission mechanism. While the analysis is still preliminary and does not go much beyond a descriptive level, our findings suggest that microeconomic characteristics of Italian firms do seem to have considerable predictive power regarding cyclical fluctuations.

Summing up, we draw three main conclusions from our analysis. First, there are several good reasons why previous attempts to uncover the likely effects of the shift to the common monetary policy have been inconclusive. Second, looking at the micro data from different countries can help resolve some of the questions left unanswered by the studies that have focused on aggregate data. Finally, in the Italian recession that followed Italy's exit from the Exchange Rate Mechanism (ERM) in 1992, a number of suggestive differences in investment rates and profitability of different sets of firms emerge, in line with existing theories. The next step in our research will be to study these differences further by refining our indicators, controlling for the correlation among them, and dealing better with endogeneity problems.

Prior studies comparing monetary transmission mechanisms in Europe

A number of recent papers have attempted to gauge the differences and similarities among the monetary transmission channels of the EU countries. Almost all these studies rely on aggregate data and analyze the response to a monetary policy shock displayed by macroeconometric models of the different economies.[2]

An obvious, preliminary issue is whether anything at all can be learned from the past research. Indeed, it is certainly possible that the final move to European Economic and Monetary Union (EMU) is such a big regime shift that past experience is no longer a reliable guide.[3] However, there is no clear evidence—nor is it likely on *a priori* grounds—that the regime shift will lead to a sharp discontinuous break in relations in the economy. As behavior tends to adjust gradually, past relationships are likely to retain some of their predictive value for the near term.

Even in the absence of structural breaks, however, considerable care is required to translate the knowledge of the (past) differences and similarities among monetary policy transmission mechanisms in the EU countries into an assessment of the (future) transmission mechanisms of the single monetary policy in the different countries. The move to a single currency changes significantly the conditions under which monetary policy operates, making it difficult to interpret most of the empirical evidence on the past transmission mechanisms.

The "ideal" study, based on past experience, which would be informative about differences across countries in the transmission mechanism of a single monetary policy, would consider the response of the various EU economies to the *same temporal sequence of monetary policy shocks*, *holding fixed the exchange rate among them*. In addition, as stressed by Dornbusch, Favero, and Giavazzi (1998), in the ideal circumstances it should also be possible to test the statistical significance of any difference found in the transmission mechanism. With this benchmark in mind, we can survey the existing empirical literature on the European monetary transmission channel.

Studies based on large-scale macroeconometric models

The existing literature can roughly be classified into two main groups, depending on whether the evidence is obtained from models of the various economies that do or do not have a common structure. The primary findings involving models that do not necessarily have the same structure come from the comprehensive Bank for International Settlements (BIS) project on the transmission mechanisms in the principal industrialized countries (Bank for International Settlements, 1995). The project simulated the response of the central banks' macroeconometric models to a common, standardized monetary policy shock (an increase of the policy rate by 1 percentage point for two years, with the rate returning to the baseline path immediately afterwards).

Importantly, the BIS research protocol envisaged the simulations to be conducted both under unchanged exchange rates and allowing the exchange rates to react to the move in the interest rate. In the latter case, two variants were agreed upon: one allowing for an independent response of each currency and a second involving a coordinated response of the ERM currencies, with a common pattern of the exchange rate vis-à-vis the rest of the world. Thus, in principle, the evidence produced within the BIS study complies with the two main requirements of the "ideal" experiment.

Unfortunately, however, not all countries in the study implemented the protocol in its strict form. Specifically, the variant corresponding to a coordinated response of the ERM countries—precisely the exercise that would have been necessary to comply with the "ideal" defined above—is missing for Germany, Spain, and the UK.[4] In addition to this limitation, since the BIS study makes use of "traditional" large-scale macroeconometric models, it is subject to the standard criticisms of those models.

In particular, the sheer size of the models and the lack of fully articulated and consistent foundations in optimizing behavior can lead to simulation results that are difficult to interpret. Moreover, one can argue that many of the equations in these models would fail statistical tests aimed at assessing their specification. Similarly, the modeling of the instruments of monetary policy is often done in an ad hoc way. Collectively, these problems could distort the picture of how monetary policy operates. Finally, the BIS study does not allow formal statistical testing of the differences found, since the models are estimated independently.

Bearing these caveats in mind, the evidence from the BIS study—summarized for the main euro area countries in table 1[5]—points to some differences, particularly among large and small countries. In particular, the gross domestic product (GDP) response is considerably more pronounced in the larger countries. Among them, Italy exhibits a slightly larger and definitely longer lasting response. A second relevant difference concerns the price response, which is initially non-negligible only in Italy, the Netherlands, and Belgium; in Germany, it becomes sizable only after the first two years, and keeps increasing over the period; in Austria, the price response is basically nil.

TABLE 1

Compilation of simulation data from BIS study

		First year	Second year	Peak effect	Last year (7th)
Italy	GDP	–0.18	–0.44	–0.44	–0.12
	PGDP	–0.13	–0.38	–0.51	0.07
France	GDP	–0.18	–0.36	–0.36	0.05
	PGDP	–0.04	–0.19	–0.31	–0.21
Germany[a]	GDP	–0.15	–0.37	–0.37	0.11
	PGDP	0.03	–0.02	–0.53	–0.53
Netherls.	GDP	–0.10	–0.18	–0.18	0.02
	PGDP	–0.08	–0.36	–0.47	–0.16
Belgium	GDP	–0.03	–0.12	–0.23	0.02[b]
	PGDP	–0.13	–0.51	–0.84	–0.55[b]
Austria	GDP	–0.08	–0.14	–0.14	0.01
	PGDP	0.02	–0.01	–0.05	0.00

[a]German data are not strictly comparable because the exchange rate was not handled in exactly the same way as for the other countries.
[b]Fifth year after the shock.
Note: Responses of real GDP and the GDP deflator (PGDP) to a 100 basis point increase in the policy rate for two years, followed by return of the rate to the normal level (fixed exchange rate vis-à-vis ERM countries; deviations from baseline in percentage points).
Source: Bank for International Settlements (1995).

Overall, given that the BIS study comes somewhat close to satisfying two of the three conditions characterizing the "ideal" empirical study, the differences identified in this study should be taken seriously. Moreover, the model used in the BIS study (central banks' models) represents the "insider wisdom" of the monetary policy authorities, which is interesting in itself. However, the lack of a common structure in the models raises the question of whether any differences one observes are simply an artifact of different and arbitrary modeling choices.

Studies imposing a common structure on the models for different countries

The second group of papers studying the transmission channels in Europe is more heterogeneous. These studies include evidence from structural vector autoregressions (Gerlach and Smets, 1995; Barran, Coudert, and Mojon, 1996; Ehrmann, 1998; Kieler and Saarenheimo, 1998; Ramaswamy and Sloek, 1998; and Dedola and Lippi, 1999); from small structural models with a common structure (Britton and Whitley, 1997); from relatively large multicountry models (the U.S. Federal Reserve multicountry model in the BIS study; the models in International Monetary Fund [IMF], 1996, and the European Commission; and Roeger and In't Veld, 1997); and from prediction equations for output, estimated for different countries (Dornbusch, Favero, and Giavazzi, 1998; and Peersman and Smets, 1998).

The papers using structural vector autoregressions (SVAR) all try to determine how a change to one of the variables being analyzed influences the other variables under consideration (for instance, how interest rates might influence investment).[6] These papers run into two problems in this context. First, the shocks to the models typically differ across countries, both in terms of size and time path. These differences make it impossible to make legitimate comparisons among the responses. This problem is exacerbated because most models embody different assumptions about the way in which the monetary authority responds to new developments (that is, the endogenous component of monetary policy). Thus, even on the off-chance that the same initial disturbance is analyzed, the monetary policy responses would not be harmonized so that a symmetric response across countries would not be expected. Instead, the differences in the assumed monetary reactions would generate different economic responses, even if the underlying structure of the economies were similar.[7]

The second problem in the SVAR literature is a failure to properly account for the lock-in of the parities among the currencies in the euro area, which implies a common response of the exchange rate. Indeed, the SVARs often do not include the exchange rate; when they do, the shocks are often inferred in dubious manner. For instance, the studies we have seen always assume that a disturbance to interest rates does not simultaneously influence exchange rates (or vice versa). Such shocks are hard to imagine since they imply a "free lunch," whereby investors could move money towards high-interest countries without expecting to see some of the interest rate gains eroded by changes in exchanges rates. With the shocks having been identified in this fashion, it is very likely that the so-called "monetary policy shock" is in fact a combination of shocks, including the endogenous response to movements of the exchange rate.

As a result of these two problems, much of the evidence produced by the SVAR literature is of only limited relevance for the issue at hand, as it does not appropriately represent the situation that is likely to prevail in the monetary union. A vivid example of the difficulties in interpreting the SVAR results is the Gerlach and Smets (1995) study, in which the responses to both a one standard deviation, one-period shock (reported in table 2 as variant 1), and a 100 basis point, two-year sustained increase of the interest rate (variant 2) are presented. In the first case the response of GDP looks similar across Germany, France, and Italy, while in the second case, German output moves by almost twice as much as that of the other two major countries of the euro area; in the latter case the German result is also much more persistent (although this is masked in the table).

Even taking the SVAR evidence at face value, the results are often ambiguous.[8] While many of the studies tend to conclude that the differences in the transmission mechanism are not large, the differences they identify do not seem to be particularly robust: As summarized in table 2, different studies present somewhat different rankings of the potency of monetary policy.

The main regularities that do seem to emerge are that Germany is almost always the country in which monetary policy is most powerful, often followed by France, and that monetary policy is always seen as being more potent in Germany than in Italy, where monetary policy appears to have the mildest effect on output. These conclusions are almost the opposite of the findings from the aforementioned BIS project. One potential reconciliation is offered by Kieler and Saarenheimo (1998), who show the extreme indeterminacy of the SVAR results: A very large set of widely different responses of output to monetary policy, each equally supported by the available data, can be produced by varying the assumptions used to identify shocks. Restricting the identifying assumptions to those that yield impulse responses bounded within a sort of "window of plausibility" (for example, the initial output and price response to a contractionary shock should not be too positive) still leaves open a very wide range of possibilities.

Looking at small structural models and multicountry models, both with essentially the same structure across countries, none of the studies quite comply with the requirements set out above. In particular, the common response of the exchange rate has not been implemented. The evidence extracted from simulations of these models points to relatively small differences in the transmission channels across countries. Aside from the U.S. Federal Reserve multicountry model (which generates a much stronger initial response for Germany and France than for Italy), the other models show little or no difference in the impact on GDP. Of course, the identifying assumptions that underlie these models are subject to the same criticisms leveled at the national macroeconometric models.

Finally, the studies based on "prediction equations for output" have the advantage of having been devised precisely to provide the sort of ideal evidence described above. The estimated equations allow the path of both the monetary policy shock and the exchange rate to be common across countries, and the estimation is done jointly so that formal statistical testing is possible. On the other hand, the ad hoc nature of these equations limits one's ability to interpret the results, and doubts can be raised about the identification of the monetary policy shock. Dornbusch, Favero, and Giavazzi (1998) jointly estimate an equation for output growth in each country. The specifications predict output growth in each country as a function of its past own values and of past and present values of growth in the other countries, expected and unexpected components of interest rates,[9] and the bilateral exchange rates with the dollar and the deutschemark (DM). The specification of the

TABLE 2

Effect of monetary policy on output, using SVARs

Study	Effect on GDP one year after shock: Germany	France	Italy	UK	Sweden	Netherlands	Strength of responses[a]
Ramaswamy and Sloek (1998)	–0.6	–0.4	–0.5	–0.5	–0.3	–0.6	S<F<I=UK<G=NL
Barron, Coudert and Mojon (1996)	–0.6	–0.4	–0.3	–0.4	–0.4	–0.3	I=NL<F=UK=S<G
Gerlach and Smets (1995), variant 1	–0.3	–0.2	–0.2	–0.6	n.a.	n.a.	I=F<G<UK
Gerlach and Smets (1995), variant 2	–1.0	–0.5	–0.5	–0.7	n.a.	n.a.	I=F<UK<G
Ehrmann (1998)	–0.9	–0.5	–0.1	0.2[b]	–0.1	0.0	NL<I=S<F<G
Dedola and Lippi (1999)[c]	–2.2	–1.4	–1.1	–1.4	n.a.	n.a.	I<UK=F<G

[a]These orderings rank the responses according to their magnitude in each study.
[b]Data are not comparable.
[c]Figures refer to the maximum elasticity to the shock of industrial production.
n.a. indicates data not available.

output equations in Peersman and Smets (1998) is similar, but they include the German real interest rate and the differential with the German real rate instead of the expected and unexpected components of interest rates, and they replace the bilateral exchange rate against the dollar with the bilateral exchange rate between Germany and the U.S.; in addition, they allow no contemporaneous relationships. While the quantitative results differ in the two papers, they both point to significant differences in the output responses of Italy, on one side, and Germany and France, on the other. In particular, the Italian response is stronger, a result that is similar to that in the BIS study but sharply in contrast with the SVAR evidence.[10]

Summary

The main lesson we draw is that the evidence so far available is not quite appropriate to assess whether the single monetary policy will have a differential impact on the euro area countries. Moreover, the results are not robust: Methodological differences (such as which variables are included in the models and how shocks are identified) change the conclusions quite substantially. With the relevant exception of the "output equations," one regularity is that models with a similar structure tend to yield small differences in the transmission mechanisms, whereas models with a more idiosyncratic structure tend to show larger differences. However, it is unclear whether, on the one hand, the similarities in the former case are forced by the choice to impose the same structure on (truly) different economies or whether, on the other hand, the differences in the latter case result from the choice of modeling as different economies that are (truly) similar. It should nonetheless be acknowledged that, though far from being conclusive, the two pieces of evidence that most closely comply with the requisites for the "ideal" experiment—namely the BIS study and the output equations—provide roughly consistent results and point to noticeable differences in the transmission mechanisms.

Microeconomic evidence on the structure of European economies

The ambiguity of the macroeconometric findings on differences in the transmission mechanism undoubtedly stems, at least in part, from the poor design of the existing studies. Further work to remedy these problems should help to substantially clarify matters. We believe, however, that one additional reason for the inconclusive findings of these studies is their reliance on aggregate data. Relevant differences in the response to a monetary shock might be observed among different groups of agents in the same country, similar groups of agents in different countries, or both. However, the relative weights of these groups could differ across countries, in which case aggregation problems will confound attempts to make sense of the evidence.

Therefore, we propose to supplement the macro-level analysis with an exploration conducted at the micro level. Focusing on micro data has two further advantages. First, by identifying the behavioral responses of sets of agents that have been grouped according to different structural characteristics, this approach provides the information needed to uncover the causes of whatever differences might be present at the macro level. Second, it might help identify possible policy interventions or natural mutations which, by altering the "microeconomic landscape" in the relevant ways, could lead to more uniform effects of the common monetary policy.

We consider four different theories of how monetary policy can affect the economy. These theories identify the characteristics of the various economies that should determine the potency of monetary policy. While we recognize that these theories of monetary transmission share some common features—for instance, most require that prices do not instantly adjust to changes in monetary conditions—we consider it useful to highlight the differences among the theories rather than the similarities. Once we have identified the salient characteristics, we can see whether the member countries of the monetary union differ along these dimensions.

Theories of monetary policy transmission

The textbook model of monetary transmission supposes that open market operations matter because, in the presence of temporarily fixed prices, altering the mix of money and bonds changes the real value of the money supply. This leads to a shift in interest rates to clear the money market and, subsequently, to changes in spending on interest sensitive items. Since this mechanism operates in a host of models ranging from the IS/LM to cash-in-advance or limited-participation models, we refer to it as the *conventional* mechanism. We take its central prediction to be that the potency of monetary policy across countries will depend on the cross-country variation in the interest sensitivity of spending (see Kakes, 1999, for further discussion).

A second theory of monetary transmission builds on the interest rate mechanism by assuming that financing difficulties can amplify the impact of the initial change in interest rates. Capital market distortions induce lenders to require collateral before they will make funds available. Because any interest rate increase lowers the value of future cash flows, collateral is influenced by open market operations, and this is assumed to alter the availability of funds and ultimately spending. We call this the *borrower-net-worth* mechanism (see Bernanke and Gertler, 1995). We take the central prediction of this theory to be that debt capacity will depend on borrowers' net worth and this will drive spending.[11]

A third, and closely related, theory emphasizes the role of banks. This theory posits that an open market sale matters because it removes reserves from the banking system; this in turn impairs banks' ability to make loans. For some customers a cut in bank lending is assumed to translate into reduced spending. Thus, the theory requires that both banks and bank customers have financing problems that are exacerbated when a monetary tightening is undertaken—see Stein (1998) for a formal model and Kashyap and Stein (1997) for a discussion in the context of the EMU. This channel is really a special case of the borrower-net-worth channel since it focuses on the importance of the availability of funds from banks; to highlight this we call it the *bank-lending* channel. We take its central prediction to be that the potency of monetary policy will depend on the degree to which banks are able to raise alternative funds to offset reserve fluctuations and the extent to which consumers and firms must rely on banks for their financing.

A final mechanism, which has a long history in discussions of monetary policy transmission, focuses on the non-price methods of allocating credit. For instance, Roosa (1951) argued that monetary policy could be quite potent without moving interest rates by influencing the availability of credit. The net-worth and bank-lending mechanisms described above are special cases of this theory, in that they assume that contracting difficulties influence credit allocations in a particular way. Alternative versions of the *credit-rationing* hypothesis would permit factors beyond net worth and collateral to influence credit availability.

For example, in the seminal Stiglitz and Weiss paper (1981), equilibria in which credit is rationed are possible because of asymmetric information between borrowers and lenders that leads to problems of moral hazard and adverse selection. Williamson (1987) studies the implications for lending of an imperfect ability to monitor borrowers. He shows that a rationing equilibrium may exist in which interest rates are no longer allocative; instead lenders adjust to shocks by changing the amount of credit they extend.

Working out the precise implications for monetary policy transmission is difficult because the credit allocations can differ depending on the modeling assumptions. However, one robust prediction from these models is that credit rationing becomes increasingly likely and widespread in economies with less efficient legal systems, more "opaque" borrowers' activities, and weak enforcement of contracts.[12] Thus, we also report data comparing the economies along these dimensions.

Microeconomic data describing different economies in Europe

Collectively, these theoretical considerations suggest a number of structural features that would be useful to compare across the European economies that are operating with a common monetary policy (or, in the case of the UK, are considering joining the union). Finding comparable data on the relevant indicators for all 11 countries that adopted the euro is quite difficult, so our preliminary exploration focuses on seven countries with readily available data.[13] The proxies shown in table 3 are intended to provide some evidence on the differences in interest sensitivity, collateral positions, importance and availability of bank loans, and the costs of contract enforcement. First, we review the findings for the different indicators. Then we draw some tentative conclusions about individual countries.

One factor that is common to all the theories is some form of imperfect price adjustment. If prices adjust more quickly to monetary impulses in some countries rather than others then this would lead to different patterns of output adjustment. Thus, an obvious starting point for comparisons would be the degree of price rigidity across countries.

A major problem with this tack is the uncertainty over how pricing practices may change once prices in the euro area are quoted in the same units. One of the benefits often cited by the advocates of the single currency is that it will increase competitiveness of product markets, which will tend to equalize prices and price-setting practices across countries. To the extent this is true it raises questions about how much faith to put in past evidence on pricing policies—this is one case where a sharp change in behavior seems possible.

Nevertheless, we can probably gain some insight into the price rigidity issue by looking at labor market frictions. Labor costs account for a major portion of total costs and it is generally agreed that legislation governing the hiring and firing of workers in Europe makes wages relatively rigid. Moreover, the move to a single currency will not directly (or immediately) change the contractual framework governing the labor market. Thus, we report data on labor markets as a first measure of structural differences.

The first row in table 3 shows summary information on employment protection legislation in different countries. Taken from the June 1999 Organization for Economic Cooperation and Development (OECD) *Employment Outlook*, the data represent a weighted average of indicators pertaining to

TABLE 3

Selected characteristics for European countries

	Country						
Variable	**UK**	**Germany**	**Italy**	**France**	**Spain**	**Netherlands**	**Belgium**
Employment protection[a]	0.9	2.6	3.4	2.8	3.1	2.2	2.5
(rank, 26 OECD countries)	(2)	(20)	(23)	(21)	(22)	(13)	(16)
Capital output ratio[b]	1.99	4.0	3.2	3.0	n.a.	n.a.	3.0
(Investment/GDP)	(0.154)	(0.223)	(0.180)	(0.191)	(0.212)	(0.197)	(0.181)
Fraction of financing that is short term[c]	0.960	0.593	0.838	0.893	0.925	0.620	0.882
Exports outside EU-15 relative to GDP[d]	0.47	0.44	0.45	0.38	0.29	0.25	0.29
Firms' leverage (median) %[e]	63.1 (60.5)	52.0 (61.0)	52.3 (62.5)	46.3 (49.1)	53.5 (56.4)	43.9 (63.7)	51.4 (58.4)
Median number of employees per firm[f]	1,128	406	251	357	267	205	363
Household indebtness[g]	1.020	0.779	0.314	0.510	0.580	0.649	0.415
Months to repossess[h]	12	15	48	11	36	2.5	24
Repossession cost as % of house value[i]	4.75	6	19	15	10	11	19.5
% of firms with single bank[j]	22.5	14.5	2.9	4	1.5	14.3	0
Market capitalization relative to GDP[k]	1.65	0.48	0.46	0.65	0.69	1.53	0.94
Average bank size, billions of dollars[l]	24.9	12.8	12.3	20.1	10.2	32.1	22.3
% of total deposits in 5 largest banks[m]	27.0	14.0	40.4	68.8	39.8	81.3	61.0

[a]OECD (1999b), summary indicators of strictness of employment protection, table 2.5.

[b]Stock of capital at current prices divided by value added at current prices in 1996. The stock of capital is computed by the perpetual inventory method from OECD, 1999; the investment to GDP ratio is calculated from the IMF's International Financial Statistics, using the reported data on gross investment and GDP, in current dollars, averaged from 1992 to 1996.

[c]Ratio of current liabilities to total liabilities minus equity in 1996 from Enria (1999).

[d]Openness of EMU members from Favero and Giavazzi (1999)

[e]Firms' leverage is total debt divided by total debt plus net capital in 1996 using the sample of firms from Amadeus from Enria (1999).

[f]Median of the mean of industry-level employment built by Kumar, Rajan, and Zingales (1999) using raw data from Eurostat.

[g]1994 total household liabilities as a fraction of disposable income from BIS (1995).

[h]Number of months (as of 1990) necessary to repossess collateral in case of default on a mortgage from European Mortgage Federation.

[i]Legal costs to repossess collateral in case of default on a mortgage as a percentage of the value of the house in 1990 from European Mortgage Federation.

[j]Share of firms entertaining only one bank relation from Ongena and Smith (2000).

[k]Market value of firms listed on major exchanges as of year-end 1998 divided by GDP from Federation of European Stock Exchanges Annual Report, with GDP data from the OECD.

[l]IBCA *Bankscope* database for European banks; figures pertain to total assets as of 1997 year-end.

[m]Share of deposits of five biggest credit institutions in 1996 from European Central Bank (1999).

n.a. indicates data not available.

regular labor contracts, temporary contracts, and collective dismissals. The levels of these averages therefore have no direct economic interpretation, but the rankings for the main 26 OECD members are informative.

The data confirm the well-known finding that labor market institutions in the UK are much more flexible than in the rest of Europe. The amount of employment protection in the other countries (except possibly in the Netherlands) is fairly similar. If one believes that labor market frictions are going to be a key determinant of future cross-country differences in wage and price flexibility, it would appear that the differences among the continental economies will not be too large.[14]

Turning to the specific theories, trying to find evidence on interest sensitivity of spending one runs into many of the same econometric difficulties discussed in the last section. In particular, determining whether results are driven by ad hoc specification choices or true behavioral differences is not easy. Therefore, the evidence we provide should only be considered a first pass at the issue. We try, however, to assess the robustness of any inferences that we might draw by providing several indicators that should be closely related to interest sensitivity.

One measure we consider is the ratio of fixed capital to output. Countries with high levels of capital to output will (assuming they are close to a long-run desired level) have higher investment requirements. We expect that interest rate changes should matter more in high-investment countries. Looking at the data in the table we find three groups of countries: Germany, which has a very high level of capital; the UK, which has a relatively low level; and the remaining countries, which lie in between (although they are closer to Germany than to the UK). The numbers in parentheses below the capital-to-output ratio are average levels of investment to GDP between 1992 and 1996 from national income account data. These numbers essentially confirm that the British and German differences are not due to the vagaries involved in estimating the stock of capital. By this metric, monetary policy should have strong output effects in Germany, while it should have much more modest effects in the UK. The other countries, except possibly the Netherlands, should be in between.

As a second indicator, we look at data on the maturity structure of debt. Countries with mostly short-term debt can expect changes in interest rates to affect borrowing costs more rapidly than countries with mostly long-term debt. The data again show that Germany and the UK are the two polar cases, although the ranking of monetary policy potency is reversed, with German firms having much more long-term debt than British firms.[15] Aside from the Netherlands, which also has a relatively low fraction of short-term debt, most of the other countries' debt-maturity structures are closer to the UK than to Germany.

The negative correlation between the debt maturity and the capital-to-output ratio is not too surprising. If there are any frictions in borrowing and lending, then it may be desirable to match the maturity of any debt to the life of the asset. Therefore, it makes sense that in Germany, with its higher level of fixed (long-term) assets, the fraction of long-term debt is also higher.

A slight extension of the conventional model would allow interest rates to be important because of their impact on exchange rates. With a single monetary policy this channel no longer directly matters for trade within the euro area. However, it will retain its relevance if there are differences in trading patterns with countries outside the euro area. Data on the ratio of exports to GDP outside of the 15 countries in the EU are reported in table 3. It appears that the four large countries are much more likely to trade outside of the EU than the smaller countries. This pattern is probably going to persist and should mean that, all else equal, monetary policy should have more potency in the larger countries than in the smaller countries.[16]

The net-worth channel suggests that we look for differences in collateral levels. We consider three proxies. One measure is the leverage of firms—in particular, the ratio of debt to debt plus equity. The data in the table show that there is relatively little variation across countries in this dimension. Except for France, the median firm has a leverage ratio of between 0.56 and 0.64. The French firms have less debt, and one possible interpretation of this observation is that they have more borrowing capacity. Alternatively, the lack of debt may reflect problems with contract enforcement; we discuss this interpretation below.

The data on leverage are for a sample of large firms, including those listed on public stock markets. It is quite plausible that borrowing frictions are more important for smaller, non-publicly traded companies. Therefore, we also report data from Kumar, Rajan, and Zingales (1999) on firm size (in which firms are weighted according to the total employment in enterprises of a given size).[17] In terms of the size of the median firm, there are three groups of countries. The typical UK enterprise is much larger than those found on the Continent. The Italian, Dutch, and Spanish firms are relatively small, while the remaining countries have middle-sized firms. These figures suggest that collateral considerations should be strong in Italy, the Netherlands, and Spain and much weaker in the UK.

The last of the proxies we consider is household debt levels, more specifically the ratio of household liabilities to disposable income. Once again, the UK stands out, with borrowing levels far exceeding those found elsewhere. Italy stands out as the country with the lowest household borrowing, although Belgium also shows quite low levels.

One possible interpretation of these data is that Italian and Belgian households should at least be able to borrow to make up any income shortfalls. But the alternative interpretation is that households in these countries are less willing to borrow. Past research analyzing cross-country savings patterns, however, favors the former interpretation (Jappelli and Pagano, 1989, and Guiso, Jappelli, and Terlizzese, 1994).

Furthermore, two proxies related to contract enforcement suggest these patterns reflect differences in the efficiency of credit markets, rather than differences in households' willingness to borrow. One of these indicators is the number of months needed to repossess collateral in the event of a default. The second measure is the estimated legal costs of repossessing a house in the event of a mortgage default (expressed as a percentage of the value of the house). Both variables suggest that enforcement costs are high in Italy and low in the UK.

Thus, one would expect much less mortgage debt in Italy than in the UK and, hence, much lower overall borrowing. These considerations lead us to interpret the debt data as a measure of the depth of local capital markets. On the one hand, the Italians are less able than the British to smooth out shocks to consumption, since their capital markets are not as well developed and will not be able to rely as much on borrowing. On the other hand, being less leveraged than the British, the Italians are less vulnerable to shocks to interest rates.

Belgium and, to some extent, Spain also appear to be countries where contract enforcement is relatively costly. Interestingly, the Belgian, Italian, and Spanish legal systems are all derived from the French legal system. As La Porta, Lopez-de-Silanes, Shleifer, and Vishny (1997) note, creditors' rights to reorganize or liquidate firms are relatively weak in the French system. In contrast, Germany appears to be relatively efficient by these measures—which also accords with La Porta et al.'s findings. This suggests that credit rationing is more likely to occur in Belgium, Italy, and Spain than in Germany or the UK. However, as mentioned earlier, this could strengthen or weaken the impact of monetary policy.[18]

Finally, as proxies for the bank lending channel we report several measures of bank loan demand and loan supply (see Cecchetti, 1999, for further data). Our data show that in all the countries, it is typical for large firms to have several banks. This should help insulate them from a credit crunch that might result if an individual bank gets into trouble. Smaller firms appear to be more likely to rely on a single bank, although, to the best of our knowledge, it is not possible to get comparable data for small firms. Therefore, the previously described data on the variation in average firm size will be relevant for the lending channel too. From the lending channel perspective, this suggests that the reliance on bank funding is likely to be highest in Belgium, Italy, and Spain and lowest in the UK.

A second indicator of the importance of banks for the funding of businesses is the size of the capital market. Judging by the ratio of the value of shares traded on the major public stock exchanges to GDP, there is striking variation in the depth of capital markets across countries. Particularly in the UK, but also in the Netherlands, there are many huge publicly traded companies. These companies almost always have access to some types of nonbank finance. In contrast, in Germany and Italy the stock market capitalization is relatively low, a feature supporting the commonly held view that the banks dominate the financial system in these countries.

In terms of bank loan supply, Kashyap and Stein (1999) find that in the U.S. smaller banks' lending is more closely tied to monetary policy than that of large banks. This suggests that shifts in bank loan supply are more likely if a country's banking system consists mostly of small rather than large banks. One way to make this comparison is to look at differences in the absolute size of banks in the different countries. Table 3 shows the average size of the banks in the IBCA *Bankscope* database for European banks in each country in 1997. This database provides information on the largest banks in each country, covering institutions that grant between 80 percent and 90 percent of domestic credit. By this yardstick the Belgian, British, Dutch, and French banks are best positioned to insulate borrowers from changes in credit availability; the German, Italian, and Spanish banks are relatively small and therefore may not be so well able to guarantee funding for their clients.

The data in table 3 also show the share of total banking deposits in the top five banks. Focusing on concentration may be appropriate if one believes that the lack of integration of the banking markets is likely to persist, and if the largest banks in each country are expected to be able to attract funds during a credit squeeze, even if some of the banks may not be large in an absolute sense. Interestingly, except for the UK, this size measure suggests the same classification of countries as implied by the absolute measure of size; in the UK the many nonbanking financing options and the large absolute size of the leading British banks lead us to suspect that shifts in bank loan supply would be relatively less important.

Summary

Obviously, the data in table 3 are open to multiple interpretations, and the connections between some of our proxies and the ideal variables suggested by theory are sometimes loose, but we feel that several general conclusions are warranted. First, there do seem to be fairly strong differences across the countries in several respects. Moreover, the indicators do not seem likely to change quickly. Therefore, if these features do matter for monetary transmission, it seems likely that the differences will be in place for several years.

The Italian economy appears to be one in which several of the theories would predict a strong effect of monetary policy on the economy. In relative terms, Italy has a fairly high fixed-capital stock, poor contractual enforcement, lots of small firms, rigid labor markets, and many small banks operating within a financial system that has been bank-dominated. All of these factors suggest comparatively strong effects of monetary policy.

The UK looks to be almost the opposite of the Italian case. There is relatively little fixed capital, good contract enforcement, very flexible labor markets, and many large firms with genuine alternatives to nonbank financing. The only common feature between the two countries is that they both do a significant amount of trading with non-European countries.

Most of the other countries sit in the middle, with characteristics that, according to which theory of monetary transmission one considers, indicate stronger or weaker effects of monetary policy. For instance, in Germany firms are relatively large and contract enforcement is pretty good, which should help to insulate firms from monetary policy. However, Germany also has a high level of investment, fairly rigid labor markets, and exports a significant amount of goods to countries outside of Europe. France has more large banks and a more developed stock market than Germany, but corporate leverage and household borrowing in France are much lower, and it is fairly costly to repossess collateral.

Cross-firm differences in cyclical performance in Italy

Ultimately, it will take a number of studies and a considerable amount of work to determine which of the factors identified above are most important for the transmission of monetary policy. As a first step, with the intent of providing a sort of benchmark and, at the same time, assessing whether the characteristics highlighted above do indeed matter, we explore how firms that differ along those dimensions have fared in the wake of a monetary tightening. We focus on the one country, Italy, in which *a priori* we are most likely to observe strong effects of monetary policy. We believe that subsequent work can try to narrow the alternatives and, more importantly, pinpoint whether the factors that may have been significant in Italy are also relevant in other countries.

Macroeconomic conditions in Italy in the 1990s

Before we investigate the microeconomic evidence in Italy it is necessary to describe the macroeconomic environment. Table 4 shows a set of macroeconomic indicators for 1989–97, the period for which we have good firm-level data. The period is marked by considerable volatility, much of which is attributable to the developments leading up to the adoption of a common monetary policy. The year 1992 was a watershed year. Growth in the three preceding years had been relatively rapid, although the economy was gradually slowing down. While the primary deficit had improved, the overall deficit was still around 10 percent of GDP. In 1991 the total deficit deteriorated slightly and reached 10.8 percent in 1992. This situation put downward pressure on the exchange rate (which was fixed as part of the ERM).

Over the next year a number of policy changes aimed to help ease the pressure on the lira. In July the government adopted a 30,000 billion lire (about 2 percent of 1992 GDP) fiscal tightening, which ultimately proved to be insufficient to ease pressure on the exchange rate. In September, the government decided to abandon the attempt to maintain parity with the DM and the exchange rate started floating freely: It jumped from 756 lire to the DM in August to 806 lire in September and 882 in October, a devaluation of 15 percent from the previous central parity. From then on the exchange rate continued to fall, though the devaluation had, overall, relatively small effects on the price level.

To stabilize the exchange rate, interest rates were sharply increased and (perhaps more importantly) a second, remarkably large set of fiscal measures were announced at the end of 1992. Collectively, these changes reduced spending by approximately 92,000 billion lire (6 percent of GDP). The fiscal adjustment marked a clear break: In 1993 the primary deficit climbed to 2.6 percent of GDP. This was also a year of deep recession, with industrial production falling

TABLE 4

Macroeconomic conditions in Italy, 1989 to 1997

Variable	1989	1990	1991	Full year 1992	October 1992	1993	1994	1995	1996	1997
Lira/DM exchange rate (% depreciation)	729.7 (–1.54)	741.6 (1.63)	747.7 (0.82)	790.0 (5.67)	881.92 (16.1)	950.7 (20.33)	994.7 (4.63)	1,138.0 (14.41)	1,026.3 (–9.82)	982.2 (–4.29)
Real GDP growth, %	2.9	2.2	1.1	0.6	n.a.	–1.2	2.2	2.9	0.7	1.5
3-month Treasury bill rate, %	12.65	12.28	12.66	14.48	15.51	10.47	8.84	10.73	8.61	6.40
Domestic credit growth, %	14.85	13.14	12.67	11.71	11.75	7.60	6.22	5.10	4.68	4.21
Government primary deficit/GDP, %	1.1	1.7	–0.1	–1.9	n.a.	–2.6	–1.8	–3.9	–4.5	–6.7
Total government deficit/GDP, %	9.8	11.1	10.1	9.6	n.a.	9.5	9.2	7.7	6.6	2.7

Notes: The exchange rate devaluation in October 1992 is with respect to the exchange rate in August 1992.
Credit growth for October 1992 is relative to October 1991. n.a. indicates not applicable.
Sources: Bank of Italy, 1997 and 1992, *Annual Report*.

by 2.4 percent and GDP down 1.2 percent. However, recovery began quickly; in 1994 industrial production increased by 5.2 percent and GDP by 2.2 percent.

Due to the combination of the sharp devaluation (which greatly benefited export-oriented firms) and the tight fiscal policy (which heavily affected firms with a domestic market), the recession and the subsequent recovery were unevenly distributed. This is relevant in interpreting some of the latter results. As table 4 makes clear, 1993 also saw a marked slowdown in credit availability. Total credit to the economy grew by 7.6 percent, almost two-thirds its growth rate in the previous year. Though this slowdown can partly be explained by a reduction in demand, it is likely that access to credit became more difficult.[19] The recovery continued in 1995, while at the same time the exchange rate depreciated sharply. As the dollar tumbled in the wake of the Mexican crisis, and concerns arose over the domestic political situation, the lira depreciated sharply in February and March. Interest rates were then increased temporarily. The two subsequent years saw a marked slowdown followed by a mild recovery. At the same time, under pressure to fulfill the Maastricht criteria for admission to the monetary union, the government tried to speed up Italy's fiscal adjustment and, in 1997, the primary surplus reached 6.7 percent of GDP, allowing a total deficit of 2.7 percent.

Firm-level comparisons over the last ten years in Italy

To further examine the potential importance of microeconomic heterogeneity in the monetary transmission mechanism we report some simple diagnostics about investment and profitability for different sets of Italian firms. On the one hand, this task is complicated by the odd mixture of shocks, described above, that have hit the Italian economy since 1992. On the other hand, the shocks were very large and, therefore, have the potential to yield some clearly visible results. Ultimately, much more work will be needed to carefully identify and quantify these disturbances and to keep track of their impact on firms' performances. In the meantime, we hope that these exploratory tabulations may provide some guidance about which contrasts deserve further investigation.

The data that we analyze are drawn from the Italian Company Accounts Database, a large dataset collecting balance sheet information and other items on a sample of over 30,000 Italian firms. The data, available since 1982, are collected by Centrale dei Bilanci, an organization established in the early 1980s jointly by the Bank of Italy, the Association of Italian Banks, and a pool of leading banks to gather and share information on borrowers. Besides reporting balance-sheet items, the database contains detailed information on firm demographics (including year of foundation, location, type of organization, ownership status, structure of control, and group membership), employment, and flow of funds. It also reports a firm's credit score, computed directly at the Centrale dei Bilanci to help banks in screening borrowers. Balance sheets for the banks' major clients (defined according to the level of their borrowing) are collected by the banks.

The focus on the level of borrowing skews the sample toward larger firms (which also means that trade and service sector firms are underrepresented, while manufacturing firms are overrepresented). Furthermore, because most of the leading banks are in the northern part of the country, the sample has more firms headquartered in the North than in the South. Finally, since banks are most interested in firms that are creditworthy, firms in default are not in the dataset, so the sample is also tilted towards higher than average quality borrowers. Despite these biases, the sample still has much broader coverage than most datasets analyzed by economists since it includes thousands of unlisted companies and many very small firms—for example, the median firm in the sample in the early 1990s had only 26 employees.

The first panel in table 5 shows the evolution of investment and return on assets (ROA) for the median firm in the full sample. The major macroeconomic developments described in the last section are clearly reflected in this Company Accounts Database. In particular, profitability and investment were highest in the late 1980s and early 1990s. The 1993 recession also is easy to spot, as investment plunged and profitability sagged. By the end of the period investment had recovered, although profitability remained depressed.

However, the data for the median firm mask some stark differences across segments of the economy. The "size" panel in the table contrasts small firms (defined as having fewer than 50 employees) and large firms (more than 500 employees). Small firms generally have higher profit rates, as measured by return on assets (ROA), than large firms—this is not surprising given the larger failure rates of such firms. The smaller firms also have a lower investment rate, partly because these firms are less likely to be in capital-intensive industries.

For our purposes, however, the differences around the 1993 recession are most relevant. For large firms the recession was rather mild; the investment rate fell by about 20 percent and profitability dipped slightly. For small firms the declines were much steeper: Investment dropped by more than 40 percent and ROA also declined by more than 1 percentage point. As late as 1996, small firms' ROA had not returned to the 1992 level, whereas large firms' profitability had recovered by 1995. Thus, it appears that smaller firms fared worse than larger firms in this episode.

The "export propensity" panel of the table compares firms based on their exports as a fraction of their sales. Interestingly, prior to 1992 there was virtually no difference in profitability (ROA) between the high export sensitivity firms (whose exports account for more than 30 percent of sales) and the low export sensitivity firms—although the investment rates were higher for the high-export firms. The two groups, however, fared quite differently during the recession. For the typical low-export firm, investment virtually ceased in 1993 and was down nearly 25 percent in 1994; profits also dropped sharply. For the 10 percent of firms that were heavily export-oriented, profits were unchanged and investment dropped a bit but had fully recovered by 1994.

Given the large devaluation it is not too surprising that the exporters outperformed the domestic sellers, but we find the magnitude of these differences surprising. We explore these differences further below. Note that the strong exchange rate effects reinforce the concerns raised earlier about the importance of properly accounting for the impact of the single currency on the exchange rate when studying the transmission mechanism.

Another obvious contrast to consider is the degree to which firms are dependent on banks for their funding. The interest rate spike in the fall of 1992 and the subsequent recession severely affected the strength of Italian banks' balance sheets. For instance, the percentage of nonperforming loans rose from about 14.6 percent in 1992 to 22.5 percent in 1993 and then peaked at 31.1 percent in 1994, before dropping back to pre-crisis levels. Given the degree of the banking problems and the usual lending channel considerations, studying borrowers' bank dependence seems particularly appropriate.

Unfortunately, the institutional arrangements in Italy make developing a measure of bank dependence difficult. The standard approach in most studies is to compare firms that have access to public capital markets (for example, firms that are listed on a stock exchange or have publicly traded bonds) with firms that have little or no access. However, the underdevelopment of Italian capital markets means that essentially all firms have been bank dependent (for example, less than 0.5 percent of the firms in the sample are listed and these firms account for less than 8 percent of total sales in the sample). Thus, any measure of the amount of bank borrowing scaled by firm size tends to uncover relatively profitable and creditworthy firms rather than high-risk firms that are extremely reliant on banks. One challenge for further work on monetary transmission in Italy and other countries with underdeveloped capital markets will be to find better proxies to study bank dependence.[20]

The bank-dependence indicator we use in this study is whether a firm belongs to a corporate group. These alliances are quite important in Italy. Our working definition of a group member is whether the firm reports that it is controlled by a holding company. The holding companies for these groups typically have access to reliable funding through large banks and the capital markets, and operate an internal capital market for their group members. For instance, Bianco et al. (1999) find that member firms' investment is less sensitive to cash flow than that of nonmember

TABLE 5

Investment and profitability for different sets of Italian firms

(data for median firm)

Category		1988	1989	1990	1991	1992	1993	1994	1995	1996
All firms	NF	34,379	36,009	37,436	37,326	36,883	39,280	42,814	34,772	32,114
	I/A	2.29	2.02	2.05	2.07	1.77	1.12	1.43	2.00	2.30
	ROA	8.52	8.52	8.14	7.35	7.39	6.44	6.22	7.18	6.53
Size										
Small	NF	23,933	25,330	26,312	26,023	25,618	27,178	29,828	21,421	18,849
	I/A	1.62	1.37	1.55	1.59	1.35	0.77	1.06	1.41	1.66
	ROA	8.62	8.72	8.38	7.60	7.67	6.60	6.33	7.27	6.63
Large	NF	780	804	797	842	793	782	780	798	847
	I/A	4.83	4.74	4.18	3.85	3.43	2.79	2.84	3.74	3.72
	ROA	7.88	7.36	6.53	6.16	6.00	5.76	5.58	6.43	6.17
Export propensity										
High	NF	3,656	4,153	4,363	4,243	3,608	3,438	4,259	4,380	4,744
	I/A	3.98	2.82	2.63	2.73	2.40	2.33	2.75	3.83	3.41
	ROA	8.49	8.49	8.03	7.37	7.59	7.63	7.50	8.95	7.32
Low	NF	30,723	31,856	33,073	33,083	33,275	35,842	38,555	30,392	27,370
	I/A	2.20	1.91	1.98	1.99	1.70	0.10	1.30	1.75	2.10
	ROA	8.53	8.52	8.15	7.35	7.37	6.31	6.05	6.92	6.39
Group membership										
Nonmember	NF	6,764	7,753	8,633	9,091	9,762	10,616	11,131	8,930	8,415
	I/A	3.14	2.62	2.58	2.53	2.09	1.62	2.06	2.84	2.79
	ROA	9.39	9.16	8.78	7.86	7.88	6.91	6.69	7.90	6.95
Member	NF	5,184	5,683	6,344	6,732	7,134	7,906	8,383	8,003	7,385
	I/A	3.15	3.03	2.75	2.53	2.17	1.53	1.76	2.20	2.36
	ROA	8.36	8.21	7.48	6.69	6.54	5.74	5.78	6.76	6.25
Interest coverage										
High	NF	28,701	29,585	29,641	28,451	27,032	29,156	34,141	27,910	25,765
	I/A	2.59	2.31	2.36	2.40	2.12	1.37	1.72	2.43	2.67
	ROA	9.41	9.41	9.08	8.41	8.60	7.53	7.08	8.18	7.43
Low	NF	5,678	6,424	7,795	8,875	9,851	10,124	8,673	6,862	6,349
	I/A	1.02	0.99	1.04	1.19	0.97	0.57	0.57	0.71	1.08
	ROA	2.04	2.63	2.53	1.85	2.17	1.42	1.01	1.70	1.79
Interest sensitivity										
High	NF	10,189	10,653	11,118	11,092	11,046	11,140	11,643	8,826	7,831
	I/A	2.51	2.32	2.28	2.23	1.76	1.30	1.52	2.31	2.58
	ROA	8.54	8.72	8.38	7.61	7.48	6.50	6.10	7.11	6.54
Low	NF	11,070	11,459	11,894	11,823	11,598	12,170	13,440	11,168	10,358
	I/A	3.11	2.74	2.72	2.69	2.35	1.51	1.90	2.37	2.71
	ROA	8.19	8.05	7.70	7.07	7.14	6.41	6.15	6.84	6.22
Location										
North	NF	23,247	24,279	24,931	24,828	24,801	26,465	29,254	24,486	22,988
	I/A	2.57	2.32	2.36	2.29	1.94	1.27	1.58	2.27	2.50
	ROA	8.85	8.72	8.23	7.34	7.34	6.54	6.38	7.54	6.70
South	NF	4,590	4,958	5,193	5,120	4,943	5,212	5,567	4,203	3,648
	I/A	1.57	1.24	1.42	1.57	1.19	0.63	0.94	1.21	1.70
	ROA	7.47	7.66	7.25	6.72	6.78	5.36	5.03	5.28	5.53

Notes: I/A is investment in fixed capital during the year divided by year-end assets; ROA is return on assets; and NF is the number of firms. Sample splits are defined in the text.

Source: Authors' calculations based on data from the Italian Company Accounts Database.

firms. Thus, group membership may be an indirect proxy for firms that are not susceptible to a bank credit crunch. Conversely, the firms that classify themselves as independent are likely to be highly reliant on bank financing.

The "group membership" panel in table 5 compares member firms with nonmember firms.[21] In terms of investment, the typical member and nonmember firms are almost identical until 1993; only in the last three years of the sample do any differences appear and in these years the member firms invest less. The member firms also show consistently lower ROA than the nonmember firms. However, it does not appear that the member/nonmember distinction explains very much of the movement in ROA around the 1993 recession. For both sets of firms, ROA drops (by fairly similar percentages) and recovers by 1995. Overall, it does not appear that splitting the sample based on group membership is very informative.

One reading of the borrower-net-worth theory is that balance-sheet conditions should determine the cyclical sensitivity of different firms. We separated the firms whose required interest payments exceed their operating income (and operating income is positive)—the most extreme evidence of an impaired financial condition.[22] When we compare them with the remaining firms, the distressed firms show low levels of investment and ROA—undoubtedly these firms have some real problems with operating efficiency beyond their financial troubles. The recession was particularly harsh for the firms that had interest coverage problems. Investment dropped by more than 40 percent, while profitability was down by more than one-third. Certainly, this is consistent with the predictions of the net-worth models, but these firms having been hit by real shocks (perhaps the same ones driving the business cycle) might also be a plausible explanation.

According to the traditional theory of monetary transmission, interest sensitivity is the key indicator of which firms will adjust the most during a monetary tightening. As a crude proxy for interest sensitivity, we sort firms according to their industry. We classify firms in the construction sector or that produce capital goods, durable goods, and intermediate goods used in the production of investment goods as highly interest sensitive. The low interest sensitivity firms produce nondurable consumption goods or intermediate goods needed for nondurable consumption goods. We exclude agricultural firms, service sector firms, utilities, and other firms for which we could not make a clear classification based on their industrial code.

The "interest sensitivity" panel in table 5 shows investment and profitability for these firms. There do not appear to be noticeable differences for these two sets of firms around the recession. For both types of producers, investment and ROA drop noticeably in 1993. In percentage terms, the drop in investment is larger for the low-sensitivity firms, but the opposite is true for ROA. Furthermore, in the next year investment recovers more for the nondurables producers, while the ROA drop is again bigger for the durable good producers. By 1996, investment for both sets of firms had moved back to early 1990s levels. Overall, we see no clear pattern to the changes for these firms.

The "location" panel of table 5 compares firms based on whether their headquarters are in the northern or southern part of the country.[23] The southern firms are generally considered to operate in an environment that is less conducive to efficiency, are more generally dependent on government subsidies, and are typically less export-oriented. We would expect the combination of the fiscal contraction and high interest rates during the recession to have a more potent effect in the South than the North. The data confirm our conjectures. The southern firms begin with lower ROA and a lower investment rate, and show extreme drops in investment and profitability in 1993. The ROA for the southern firms remains low through 1996.

While these simple comparisons can be misleading, we believe we can safely draw several overall conclusions from table 5. First, information on export sensitivity seems essential to understand the 1993 Italian recession. More than any other factor, export sensitivity appears to isolate the firms that suffered the most. In addition, firm size appears to be important. In line with many theories, small firms had a more difficult time managing the recession. Similarly, firm location seems to matter. For the other factors, we consider the results rather mixed.

The obvious next step is to jointly control for the various features that we have identified. A full-blown regression analysis will eventually be needed; at this point, we prefer to keep the analysis simpler and shorter. As a robustness check and first step towards simultaneously allowing for alternative factors, we report several four-way sample splits. We first control for export propensity and then separate the firms along other dimensions. These tabulations allow us to see the extent to which all the table 5 results may be driven by export patterns.

The results in table 6 confirm that while exports are indeed important, they do not seem to be the whole story—to save space the table only shows the four years around the recession. In particular, we draw five conclusions from this table. First, in all but one case (discussed further below) the high-export firms do noticeably better than comparable low-export firms. Second, among the low exporters, small firms fare worse than large firms. Hence, size is not simply standing in for exporting tendencies. Third, the previous ambiguous results involving the comparisons of durable goods and nondurables producers do not become any clearer after controlling for exports. Among the domestically focused firms, both the interest-sensitive and interest-insensitive firms experience comparable declines in investment and ROA. Fourth, the group membership results remain mixed. Perhaps one can conclude that the low export group member firms did slightly worse than comparable nonmember firms; however, these differences are not very pronounced.

Finally, table 6 indicates that the results for interest coverage appear to involve more interesting interactions with exporting patterns than the other comparisons. The high-export firms with coverage problems actually underperform the non-exporters in terms of ROA, though their investment is less affected by the recession. Also, the drop in investment among non-exporting firms is not too different in percentage terms between the high- and low-coverage firms. Further study of this interaction is needed.

Conclusion

Our three main findings are as follows. First, the existing attempts to assess the likely effects of the shift to a common monetary policy are not very informative. The main problem is that no one has conducted a careful examination of what would happen if the euro system countries were subjected to the same temporal sequence of monetary policy shocks, holding fixed the exchange rate among them. This is the key constraint that will be imposed by the common monetary policy, and we simply do not know how different the responses would be across countries. Some work to fill this gap in the literature would be quite valuable.

TABLE 6

Investment and profitability, controlling for export propensities
(data for median firm)

Category		1991	1992	1993	1994
Size					
Small firms					
High export	NF	2,191	1,932	1,612	1,861
	I/A	1.80	1.68	1.52	1.88
	ROA	7.69	7.98	8.23	7.90
Low export	NF	23,832	23,686	25,566	27,967
	I/A	1.56	1.32	0.72	1.00
	ROA	7.59	7.64	6.49	6.22
Large firms					
High export	NF	178	111	161	202
	I/A	4.91	3.96	3.36	3.63
	ROA	4.20	4.46	5.83	6.92
Low export	NF	664	682	621	578
	I/A	3.62	3.42	2.64	2.51
	ROA	6.39	6.29	5.73	5.22
Interest sensitivity					
High					
High export	NF	1,786	1,519	1,441	1,768
	I/A	2.84	2.36	2.27	2.52
	ROA	7.60	7.66	7.55	7.40
Low export	NF	9,306	9,527	9,699	9,875
	I/A	2.10	1.67	1.14	1.35
	ROA	7.62	7.45	6.32	5.86
Low					
High export	NF	1,502	1,529	1,148	1,472
	I/A	2.57	2.47	2.26	2.75
	ROA	7.19	7.57	7.71	7.72
High export	NF	10,321	10,339	11,022	11,968
	I/A	2.71	2.33	1.42	1.78
	ROA	7.03	7.05	6.22	5.95
Group membership					
Nonmember					
High export	NF	1,278	1,175	1,216	1,478
	I/A	3.13	2.66	2.56	2.92
	ROA	7.77	8.07	7.94	7.73
Low export	NF	7,813	8,587	9,400	9,653
	I/A	2.40	2.01	1.51	1.95
	ROA	7.88	7.85	6.78	6.53
Member					
High export	NF	976	835	970	1,321
	I/A	3.26	2.81	2.72	3.07
	ROA	6.90	7.13	7.10	2.38
Low export	NF	5,756	6,299	6,936	7,062
	I/A	2.41	2.05	1.37	1.50
	ROA	6.76	6.47	5.53	5.52
Interest coverage					
High					
High export	NF	3,655	3,164	3,156	3,983
	I/A	2.96	2.54	2.44	2.86
	ROA	8.04	8.17	8.17	7.87
Low export	NF	26,261	27,201	29,349	32,318
	I/A	2.27	1.9	1.16	1.49
	ROA	8.32	8.32	7.26	6.84
Low					
High export	NF	588	444	282	276
	I/A	1.43	1.29	1.21	1.42
	ROA	–0.99	–1.23	–2.84	–2.52
Low export	NF	6,822	6,074	6,493	6,237
	I/A	1.11	0.90	0.46	0.46
	ROA	1.34	–0.03	–1.05	–0.30

Notes: I/A is investment in fixed capital during the year divided by year-end assets; ROA is return on assets; and NF is the number of firms. Sample splits are defined in the text.

Source: Authors' calculations based on data from the Italian Company Accounts Database.

Second, there are good reasons to believe that looking carefully at microeconomic data across countries might provide some insights about the transmission mechanism. Looking at some of the microeconomic structural differences among several European countries, these countries appear to differ significantly along many dimensions that are potentially relevant for the transmission of monetary policy. For instance, conditions in Italy and the UK look to be very different.

Finally, drawing on micro data for a specific country during a particular episode, we find that differences among firms that are related to the observed differences across countries do matter for the cyclical pattern and the response to shocks, including monetary shocks. Our analysis is mainly descriptive. Further work needs to be done to improve the methodology and obtain better measures of a number of relevant firm characteristics. However, our exploratory findings suggest that similar exercises using micro data—possibly extended to households—from other countries could be quite valuable in helping us to understand the nuances of the monetary transmission mechanism.

NOTES

[1]See Kouparitsas (1999) and Carlino and DeFina (1998) for some statistical evidence on this point. Supporters of the monetary union argue that the launch of the euro will result in an increase in the degree of synchronization of the business cycles of the member countries. However, there are theoretical arguments suggesting that synchronization could increase or decrease. For example, Krugman (1991) shows how synchronization can depend on productive specialization. If the monetary union makes it easier for countries to specialize in production for certain sectors then countries may become less harmonized. Alternatively, if intra-industry trade increases this can lead to greater synchronization.

[2]Surveys of the literature can also be found in Kieler and Saarenheimo (1998), Dornbush, Favero, and Giavazzi (1998), Gambacorta (1999), and Kouparitsas (1999).

[3]For an interesting version of this argument, see Frankel and Rose (1998), who discuss how the changing trade linkages that might follow a shift to a single currency could alter the output correlations across countries.

[4]We include the UK in the analysis since it may join the union at a later date. The lack of comparable data forced us to drop Greece from the analysis.

[5]Data for Germany are not strictly comparable, as they refer to an experiment in which the exchange rate moves vis-à-vis all countries. However, owing to the specific pattern for the exchange rate assumed in the "ERM-coordinated" experiment, the changes in the effective exchange rate are roughly the same as in the other countries (stronger in the last years of the experiment). Spain is not included in table 1 as the changes in the effective exchange rate in the experiments performed are not comparable with those of the other countries.

[6]The SVAR relates a set of variables to lags of the variables. For instance, investment and interest rates could be assumed to be determined by past values of investment and interest rates. See Kouparitsas (1999) for a further discussion of how the inference is conducted.

[7]The article by Gerlach and Smets (1995), among the first on the subject, explicitly recognizes this point and complements the standard impulse responses with responses to a prespecified path for the interest rate (this is equivalent to hitting the model with a sequence of shocks appropriately chosen). However, aside from Kieler and Saarenheimo (1998), subsequent papers have ignored the issue. As we argue below, this can be quite important.

[8]We focus here on the output comparisons mainly for convenience; the price responses are often not reported. We would not, however, expect them to be any more uniform than the patterns for GDP.

[9]In the preferred equation, only the expected part of interest rates is retained. The expected rate is constructed to be near a target level which is a function of exchange rate, GDP, and inflation deviations from "target levels" that vary across countries.

[10]Peersman and Smets find the response in Belgium is also stronger than in other countries, contradicting the BIS study.

[11]This theory is sometimes called the credit channel (or the broad credit channel) of monetary transmission.

[12]It is possible that a monetary policy contraction will be more potent in countries with poor legal enforcement. For instance, in the Williamson (1987) setup, low monitoring costs increase the possibility that the equilibrium involves no rationing, and in these equilibria interest rates on loans change but quantities do not respond to monetary policy. In rationing equilibria, which are more likely with high monitoring costs, a tightening will affect loan quantities but not prices.

[13]See Cecchetti (1999) for a similar exercise that focuses more on financial and legal differences.

[14]There is considerable pressure and a countervailing strong amount of resistance to reforming labor market institutions in most European countries, including Spain, Italy, France, and Germany. Reform is moving slowly so that in relative terms the European labor markets are still fairly rigid. One factor for the slow adjustment is the tendency to temporarily suspend a general practice in a particular set of circumstances rather than completely rolling back the general practice.

[15]Rajan and Zingales (1995) show that the German treatment of pension obligations can inflate the liabilities figures for German firms. We do not believe that this effect is very important for this sample.

[16]For all the countries, the fact that some primary commodities are priced in dollars could mean that a change the euro/dollar exchange rate could cause fluctuations in input prices—of course, this has been true historically as well.

[17]These data are the medians across industries in each country. The industry average levels of employment are calculated by weighting firm size by the fraction of industry employment in each firm.

[18] Cecchetti (1999) conducts an intriguing exercise in which he relates the La Porta et al. measures of shareholder rights, creditor rights, and the ability to enforce contracts on measures of the impact of interest rates on output and inflation. He finds that variation in the legal code does seem to partially explain why the potency of monetary policy varies. One difficulty for our purposes is that the interest rate sensitivities he uses come from models that do not account for the exchange rate restrictions discussed in the last section. These correlations also involve non-European countries. Nevertheless, the findings suggest that enforcement costs and legal structure do matter for monetary transmission.

[19] An annual Bank of Italy survey on a sample of manufacturing firms collects information on the access to bank credit. Specific questions are asked as whether firms applied for loans and were rejected by the bank(s), even if they were willing to pay the market rate and possibly even accept an increase in the cost of credit. Guiso (1998) shows that the share of firms that were turned down at the end of 1992 and 1993 were 9 percent and 12.8 percent, respectively, compared with an average of about 3 percent in the previous years.

[20] One proxy that we experimented with is the number of banks with which a borrower has contact. In Italy it appears that firms with a single bank do exhibit the characteristics that one might expect for bank dependent borrowers. However, the propensity to use multiple banks is very high, so it is possible that this screen may not generalize to other countries. Within Italy using this variable is also complicated by the need to merge the company accounts data with another data source, which means many firms end up being dropped from the analysis.

[21] Unfortunately, many firms do not classify themselves as either belonging to a group or as being independent, so we exclude these firms from the comparison.

[22] The exact classification is that low-coverage firms have a positive level of gross operating margin, but a ratio of gross operating margin which is less than the interest payments on their outstanding debt.

[23] Northern firms are located in one of the following regions: Valle d'Aosta, Piemonte, Liguria, Lombardia, Veneto, Trentino Alto Adige, Friuli Venezia Giulia, and Emilia Romagna. Southern firms are from the following regions: Abbruzzo, Molise, Campania, Puglia, Basilicata, Calabria, Sicilia, and Sardegna. The remaining firms are in the central region and are excluded from this comparison.

REFERENCES

Bank for International Settlements, 1995, *Financial Structure and the Monetary Policy Transmission Mechanism,* Basel, Switzerland, March, No. C.B. 394.

Bank of Italy, 1998, *Annual Report*, Rome.

__________, 1997, *Annual Report*, Rome.

__________, 1996, *Annual Report*, Rome.

__________, 1995, *Annual Report*, Rome.

__________, 1994, *Annual Report*, Rome.

__________, 1993, *Annual Report*, Rome.

__________, 1992, *Annual Report*, Rome.

Barran, F., V. Coudert, and B. Mojon, 1996, "The transmission of monetary policy in the European countries," London School of Economics, Financial Markets Group, special paper, No. 86.

Bernanke, B., and M. Gertler, 1995, "Inside the black box: The credit channel of monetary policy transmission," *Journal of Economic Perspectives*, Vol. 9, pp. 27–48.

Bernanke, B., S. Gilchrist, and M. Gertler, 1999, "The financial accelerator in a quantitative business cycle framework," *Handbook of Macroeconomics*, New York: Elsevier Science Ltd., Vol. 1c, chapter 21.

Bianco, M., G. Ferri, and P. Finaldi Russo, 1999, "Condizioni di accesso e nuove esigenze per finanziamento dell piccole e medie imprese," in *Nuovi orizzonti per il sistema bancario Italiano*, I. Angeloni (ed.), Bologna, Italy: Il Mulino.

Britton, E., and J. Whitley, 1997, "Comparing the monetary transmission mechanism in France, Germany and the United Kingdom: Some issues and results," Bank of England, *Quarterly Bulletin*, May, pp. 152–162.

Carlino, G. A., and R. DeFina, 1998, "Monetary policy and the U.S. states and regions: Some implications for European Monetary Union," Federal Reserve Bank of Philadelphia, working paper, No. 98-17.

Cecchetti, S. G., 1999, "Legal structure, financial structure, and the monetary policy transmission mechanism," National Bureau of Economic Research, working paper, No. 7151.

Dedola, L., and F. Lippi, 1999, "Effetti settoriali della politica monetaria in cinque paesi OCSE," Bank of Italy, mimeo.

Dornbusch, R., C. Favero, and F. Giavazzi, 1998, "A red letter day?," Center for Economic Policy Research, London, discussion paper, No. 1804.

Ehrmann, M., 1998, “Will EMU generate asymmetry? Comparing monetary policy transmission across European countries,” paper presented at the conference “Monetary Policy of the ESCB: Strategic and Implementation Issues,” Università Bocconi, July.

Enria, A., 1999, “Leverage e composizione per scadenza del debito delle impresse nei paesi europi: Una prima analisi su dati di bilancio,” Bank of Italy, mimeo.

European Central Bank, 1999, *Possible Effects of EMU on the EU Banking System in the Medium to Long-Term*, Frankfurt, Germany: European Central Bank.

European Mortgage Federation, 1992, *Statistical Bulletin*, Brussels, Belgium.

Favero, C., and F. Giavazzi, 1999, “An evaluation of monetary policy transmission in the context of the European Central Bank,” report to the European Parliament, Strasbourg, France.

Federation of European Stock Exchanges, 1998, *Annual Report*, Brussels, Belgium.

Feldstein, M., 1997, “EMU and international conflict,” *Foreign Affairs*, Vol. 76, No. 6, November/December, pp. 60–73.

Frankel, J. A., and A. K. Rose, 1998, “The endogeneity of the optimum currency area criteria,” *The Economic Journal*, Vol. 108, July, pp. 1009–1025.

Gambacorta, L., 1999, “Unione monetaria e politicia monetaria ottima: Canali di trasmissione eterogenei e patto di stabilità,” University of Pavia, doctoral dissertation.

Gerlach, S., and F. Smets, 1995, “The monetary transmission mechanism: Evidence from the G-7 countries,” in *Financial Structure and the Monetary Policy Transmission Mechanism*, Basel, Switzerland: Bank for International Settlements, March, No. C.B. 394.

Guiso, L., 1998, “High-tech firms and credit rationing,” *Journal of Economic Behavior and Organization*, Vol. 35. No. 1, pp. 39–59.

Guiso, L., T. Jappelli, and D. Terlizzese, 1994, “Why is Italy’s saving rate so high?,” in *Saving and the Accumulation of Wealth*, A. Ando, L. Guiso, and I. Visco (eds.), Cambridge: Cambridge University Press.

IBCA, Group, 1997, Bankscope database, Bureau van Dijk.

International Monetary Fund, 1996, *World Economic Outlook*, Washington DC, October.

Jappelli, T., and M. Pagano, 1989, “Aggregate consumption and capital market imperfections: An international comparison,” *American Economic Review,* Vol. 79, No. 2, pp. 1088–1105.

Kakes, J. I., 1999, *Monetary Transmission: The Credit View*, Capelle aan den Ijssel, Netherlands: Labyrint Publications.

Kashyap, Anil K, and Jeremy C. Stein, 1999, “What do a million observations on banks have to say about the transmission of monetary policy?,” *American Economic Review*, forthcoming.

__________, 1997, “The role of banks in monetary policy: A survey with implications for the European Monetary Union,” *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 21, No. 5, September/October, pp. 2–18.

Kieler M., and T. Saarenheimo, 1998, “Differences in monetary policy transmission? A case not closed,” European Commission, economic paper, No. 132.

Kouparitsas, Michael A., 1999, “Is the EMU a viable common currency area? A VAR analysis of regional business cycles,” *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 23, Fourth Quarter, pp. 2–21.

Krugman, P., 1991, *Geography and Trade*, Leuven, Belgium: Leuven University Press.

Kumar, K., R. Rajan, and L. Zingales, 1999, “What determines firm size?,” University of Chicago, mimeo.

Lamfalussy, A., 1997, “Securing the benefits of EMU,” remarks before the forum on “The Future of European Monetary Integration” at the London School of Economics, March 10, available on the Internet at www.ecb.int/emi/key/key08.htm.

La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert W. Vishny, 1997, “The legal determinants of external finance,” *Journal of Finance*, Vol. 52, No. 3, July, pp. 1131–1150.

Organization for Economic Cooperation and Development, 1999a, *International Sectoral Database*, diskette, No. ISDB 98.

__________, 1999b, *Employment Outlook*, Paris: OECD, June, chapter 2.

Ongena, S., and D. C. Smith, 2000, "What determines the number of bank relationships? Cross country evidence," *Journal of Financial Intermediation*, forthcoming.

Peersman, G., and F. Smets, 1998, "The Taylor Rule: A useful monetary policy guide for the ECB?," paper presented at the conference "Monetary Policy of the ESCB: Strategic and Implementation Issues," Università Bocconi, July.

Rajan, R., and L. Zingales, 1995, "What do we know about capital structure? Some evidence from international data," *Journal of Finance*, Vol. 50, No. 5, pp. 1421–1460

Ramaswamy, R., and T. Sloek, 1998, "The real effects of monetary policy in the European Union: What are the differences?," *IMF Staff Papers*, Vol. 45, No. 2, pp. 374–396.

Roeger, W. E., and J. In't Veld, 1997, "A multi-country business cycle and growth model," European Commission, economic papers, No. 123.

Roosa, R., 1951, "Interest rates and the central bank," in *Money, Trade and Economic Growth*, New York: The Macmillan Company.

Stein, J. C., 1998, "An adverse-selection model of bank asset and liability management with implications for the transmission of monetary policy," *Rand Journal of Economics*, Vol. 29, Autumn, pp. 466–486.

Stiglitz, J. E., and A. Weiss, 1981, "Credit rationing in markets with imperfect information," *American Economic Review*, Vol. 71, No. 3, pp. 393–410.

Williamson, S. D., 1987, "Costly monitoring loan contracts and equilibrium credit rationing," *Quarterly Journal of Economics*, Vol. 103, No. 1, pp. 135–146.

Index for 1999

To order copies of any of these issues, or to receive a list of other publications, please telephone (312)322-5111 or write to:

Federal Reserve Bank of Chicago
Public Information Center, P.O. Box 834
Chicago, IL 60690-0834

Economic Perspectives and other Bank publications are available on the World Wide Web at www.frbchi.org.