

Camacho, Maximo; Leiva-Leon, Danilo

Working Paper

The Propagation of Industrial Business Cycles

Bank of Canada Working Paper, No. 2014-48

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Camacho, Maximo; Leiva-Leon, Danilo (2014) : The Propagation of Industrial Business Cycles, Bank of Canada Working Paper, No. 2014-48, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2014-48>

This Version is available at:

<https://hdl.handle.net/10419/123732>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2014-48

The Propagation of Industrial Business Cycles

by Maximo Camacho and Danilo Leiva-Leon

Bank of Canada Working Paper 2014-48

October 2014

The Propagation of Industrial Business Cycles

by

Maximo Camacho¹ and Danilo Leiva-Leon²

¹Universidad de Murcia
mcamacho@um.es

²International Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
Corresponding author: leiv@bankofcanada.ca

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

Acknowledgements

We are thankful to Gabriel Perez Quiros for his comments. M. Camacho acknowledges support from project ECO2010-19830. All remaining errors are our responsibility.

Abstract

This paper examines the business cycle linkages that propagate industry-specific business cycle shocks throughout the economy in a way that (sometimes) generates aggregated cycles. The transmission of sectoral business cycles is modelled through a multivariate Markov-switching model, which is estimated by Gibbs sampling. Using nonparametric density estimation approaches, we find that the number and location of modes in the distribution of industrial dissimilarities change over the business cycle. There is a relatively stable trimodal pattern during expansionary and recessionary phases characterized by highly, moderately and lowly synchronized industries. However, during phase changes, the density mass spreads from moderately synchronized industries to lowly synchronized industries. This agrees with a sequential transmission of the industrial business cycle dynamics.

JEL classification: C22, E27, E32

Bank classification: Business fluctuations and cycles; Domestic demand and components; Econometric and statistical methods

Résumé

Les auteurs analysent les liens entre les cycles d'activité qui permettent aux chocs survenus au sein d'un cycle sectoriel d'atteindre d'autres secteurs et d'enclencher (parfois) des cycles pour l'ensemble de l'économie. La transmission des cycles sectoriels est exprimée dans un modèle de Markov multivarié avec changement de régime, estimé au moyen de l'échantillonnage de Gibbs. Le nombre et l'emplacement des modes dans la distribution des différences sectorielles changent au cours du cycle d'activité, comme le révèle l'emploi de méthodes non paramétriques d'estimation de la densité. Durant les phases d'expansion et de récession se dégage une structure trimodale relativement stable caractérisée par la coexistence de secteurs hautement, modérément et peu synchronisés. Toutefois, pendant les transitions, le centre de la distribution des densités se déplace, passant des secteurs modérément synchronisés aux secteurs peu synchronisés. Cette observation corrobore la transmission par séquence de la dynamique des cycles d'activité sectoriels.

Classification JEL : C22, E27, E32

Classification de la Banque : Fluctuations économiques et cycles; Demande intérieure et composantes; Méthodes économétriques et statistiques

1 Introduction

In practice, people do not know the state of the business cycle, which is especially uncertain around turning points. This could be because “the state of the economy” depends on the behaviour of many interdependent industries that do not necessarily all “boom” when the aggregate economy is prosperous or “bust” when the economy is in recession. Accordingly, although the aggregate business cycle could be described at a macro level as a series of distinct recession and expansion phases, it could never be understood at that level. Although recessions are not all alike, some lessons can be learned regarding their propagation when tracking the micro foundations of cycles through a variety of interconnected market dynamics at industry levels.

The main purpose of this paper is to understand (1) when microeconomic idiosyncratic business cycle shocks lead to aggregate fluctuations, (2) which industries manifest the first signs of the phase changes, and (3) how the interconnections across industries lead to cascade effects that propagate the idiosyncratic shocks to the rest of the economy. Our analysis contributes to several strands in the existing literature. First, Gabaix (2011) and Acemoglu et al. (2012) postulate that when there are significant asymmetries in the roles that industries play as suppliers to others, idiosyncratic firm-level shocks can explain an important part of aggregate movements. Foerster, Sarte and Watson (2011) show that the role of sectoral shocks increased considerably after the Great Moderation. We complement these analyses by focusing on the transmission of shocks over distinct business cycle phases.

Second, the sharp downturn in the economy experienced in 2008 and the subsequent jobless recovery increased concerns for security for asset holders (Malmendier and Nagel, 2011) and for people seeking work (Urquhart, 1981). Our analysis provides assessments of the industries that are both less and more sensitive to the aggregate cycle in bad times, which may represent useful information for investors and workers. For this purpose, we characterize salient features of the propagation of industrial business cycles.

Third, although the business cycle analysis has largely focused on national-level phases, there is a growing interest in state-level data (Owyang, Piger and Wall, 2005, and Hamilton and Owyang, 2012) and in city-level data (Owyang et al. 2008). The analyses of state-level data find that the disparities in regional business cycles can partially be attributed to differences in the industrial composition of the regions and the analyses of city-level data find that the low-phase growth is related only to the relative importance of manufacturing. Undoubtedly, these analyses could be complemented by examining the business cycle at the industry level.

Despite this interest in the literature, little attention has been paid to analyzing business cycle dynamics at the industry level. Without being exhaustive, some examples include Christiano and Fitzgerald (1998), who gauge the extent of co-movement across a range of disaggregated sector categories and the total by computing the square of the correlation between their business cycle components in hours worked, which are the outcome of band-pass filters. Carlino and DeFina (2004) focus on the analysis of growth cycles by examining the pairwise correlation between the sectorial cycles in band-pass filtered employment. Recently, Goodman and Mance (2011) analyze the per cent change in industry employment data during recessions, as determined by the National Bureau of Economic Research (NBER).

One important drawback of these analyses is that they rely on static measures, such that changes over time can only be captured by splitting the samples into sub-periods. The problem with this approach is that it provides pictures of the cycle linkages that rely on specific date breaks, the location of which is sometimes controversial. To our knowledge, no attention has been paid to examining the time-varying dynamics of these interactions, which is crucial to understand the transmission of business cycle shocks across industries. To fill this gap, this paper aims at establishing a rigorous procedure to characterize the propagation of business cycle shocks at the industry level.

This paper uses an extension of the Markov-switching model of Hamilton (1989) that computes inferences of the industry cycles and characterizes the evolution of the interactions across their cycles by combining the framework developed by Bengoechea, Camacho and Perez Quiros (2006) and Leiva-Leon (2014). In particular, each of the industrial cycles is viewed as a continuous-valued Markov-switching variable whose transition between two distinct phases defines the states of its business cycles. The synchronization of two industries in a bivariate specification is viewed as a time-varying combination between two extreme cases: (i) two independent Markov processes, which indicate completely independent industries; and (ii) a unique Markov process shared by both industries, which indicates perfect synchronization. The shifts between these extreme regimes is governed by the outcome of an unobserved Markov chain.

By means of purely statistical techniques, such as nonparametric density estimation and bootstrap multimodality tests, the number of modes in the time-varying distributions of pairwise business cycle dissimilarities is tested. This is useful to uncover distinct business cycle dynamics for different population subgroups of industries and to assess how these subgroups evolve across the distinct phases of the business cycle. Also, multi-dimensional scaling techniques are used to

understand the formation of these subgroups and their intra-distribution transitional dynamics.

We report two major findings. First, we find that, at a micro level, the U.S. business cycle is a more elusive phenomenon than we would have expected at the macro level. While some industries seem to “stick together” and show business cycle experiences that are similar to those of the nation, there are many that do not. Goods-producing industries, complementary businesses, and wholesale and retail industries are among the first to fall at the onset of recessions. However, durable goods industries, professional and technical services, and industries providing transportation and warehousing and storage for goods do not experience job cuts until some time after the beginning of recessions. For their part, businesses engaged in providing education and training, health care and social assistance and industries providing utility services and public goods are less sensitive to national recessions, particularly the 2001 recession. These results agree with Peterson and Strongin (1996), who examine the cyclicalities of industries, finding that durable goods industries are three times more cyclical than non-durable goods industries.

The distribution of the ergodic business cycle dissimilarities has a big hump at the high end, representing industries that are less sensitive to the national business cycle. However, it also shows a prominent middle mass and a smaller low end hump, representing those industries typically considered to be procyclical.

Second, we detect a thought-provoking recurrent business cycle pattern. Over the past three decades, a salient characteristic of the U.S. cycle dynamics is that the distribution of business cycle dissimilarities across industries shifts over time. During expansions and recessions, the distribution is characterized by three clusters of highly, moderately and lowly synchronized industries, yielding a trimodal distribution. However, during periods of shifts in business cycle phases or turning points, the moderately synchronized industries enjoy downward synchronization mobility and shift over to the lowly synchronized cluster, yielding a bimodal distribution. Once the transitions end, the trend is reversed and the economy stabilizes in the new business cycle phase.

The structure of this paper is organized as follows. Section 2 describes the framework used to compute inferences on the industrial business cycle dynamics. Section 3 presents the empirical results. Section 4 concludes and proposes some lines for future research.

2 Assessing industrial business cycles

Two things are required to study co-movement across industries over the business cycle: a measure of economic activity in the industries and a precise definition of their business cycles. In this paper, the economic activity at date t in a given sector is measured by the annual growth rates of total employees in the sector. The definition of business cycles relies on the recognized empirical fact that although series on employment present trends, they are not monotonic curves, but rather exhibit sequences of upturns and downturns. During periods known as recessions, the value of employment growth rates are usually lower (and sometimes negative) than during periods of expansion. A natural approach to model this particular non-linear dynamic behaviour is the regime-switching model proposed by Hamilton (1989).

2.1 Univariate model

Following Hamilton (1989), we assume that the switching mechanism of the k th sector's employment growth at time t , $y_{k,t}$, is controlled by an unobservable state variable, $s_{k,t}$. Owyang, Piger, and Wall (2005) specify a simple switching model that captures this non-linear dynamic:

$$y_{k,t} = \mu_{s_{k,t}} + \varepsilon_{k,t}, \quad (1)$$

where the growth rate of employment in sector k at date t has the mean $\mu_{s_{k,t}} = \mu_{k0} + \mu_{k1}s_{k,t}$, which is allowed to switch between two distinct regimes. At time t , one can label $s_{k,t} = 0$ as expansions and $s_{k,t} = 1$ as recessions. Deviations from this mean growth rate are created by $\varepsilon_{k,t}$, which is an *i.i.d.* Gaussian stochastic disturbance with a mean of zero and variance σ_k^2 . Therefore, employment is expected to exhibit high (usually positive) growth rates in expansions and low (usually negative) growth rates in recessions. The variable $s_{k,t}$ is assumed to evolve according to a first-order Markov chain, whose transition probabilities are defined by

$$p(s_{k,t} = j_k | s_{k,t-1} = i_k, s_{k,t-2} = r_k, \dots, \tilde{y}_{k,t-1}) = p(s_{k,t} = j_k | s_{k,t-1} = i_k) = p_{ij}^k, \quad (2)$$

where $i, j = 1, 2$, and $\tilde{y}_{k,t} = (y_{k,t}, y_{k,t-1}, \dots)'$.

2.2 Bivariate model

Although univariate Markov-switching models provide information about the timing of regime changes, they are inappropriate for drawing inferences about the synchronization of business cycles. Camacho and Perez Quiros (2006) show that the business cycle analyses based on

univariate processes are biased to show relatively low values of business cycle synchronization. Therefore, using univariate models would be particularly inappropriate for industries that exhibit highly synchronized cycles.

Phillips (1991) shows that the univariate model can be slightly modified to examine the business cycle transmission in a two-sector setup. Let us assume that we are interested in measuring the degree of business cycle synchronization between two industries, a and b . In this case, their employment growths are driven by two (possibly dependent) Markov-switching processes, $s_{a,t}$ and $s_{b,t}$, which share the statistical properties of the previous latent variable $s_{k,t}$. The bivariate state-dependent model is given by

$$\begin{aligned} y_{a,t} &= \mu_{s_{a,t}} + \varepsilon_{a,t}, \\ y_{b,t} &= \mu_{s_{b,t}} + \varepsilon_{b,t}, \end{aligned} \tag{3}$$

where $(\varepsilon_{a,t}, \varepsilon_{b,t})'$ is an identically and independently distributed bivariate Gaussian process with zero mean and covariance matrix Ω_{ab} . To complete the dynamic specification of the process, one can define a new state variable $s_{ab,t}$ that characterizes the regime for date t in a way that is consistent with the previous univariate specification. The basic states of $s_{ab,t}$ are

$$s_{ab,t} = \begin{cases} 1 & \text{if } s_{a,t} = 0 \text{ and } s_{b,t} = 0 \\ 2 & \text{if } s_{a,t} = 1 \text{ and } s_{b,t} = 0 \\ 3 & \text{if } s_{a,t} = 0 \text{ and } s_{b,t} = 1 \\ 4 & \text{if } s_{a,t} = 1 \text{ and } s_{b,t} = 1 \end{cases}, \tag{4}$$

which encompass all the possible combinations.

Bengoechea, Camacho and Perez Quiros (2006) postulate that this specification allows for two extreme kinds of interdependence between the business cycles of two industries. The first case characterizes industries for which their individual business cycle fluctuations are completely independent. In the opposite case of perfect synchronization (or dependence), both industries share the state of the business cycle. In this case, their business cycles are generated by the same state variable, so $s_{a,t} = s_{b,t}$. In empirical applications, the business cycles of two industries usually exhibit an intermediate degree of synchronization that is located between these two extreme possibilities in the sense of a weighted average. The authors consider the actual business cycle synchronization to be δ_{ab} times the case of independence and $(1 - \delta_{ab})$ times the case of perfect dependence, where $0 \leq \delta_{ab} \leq 1$. The weights δ_{ab} may be interpreted as measures of business cycle desynchronization since they evaluate the proximity of their business cycles to

the case of complete independence. This suggests that an intuitive measure of business cycle co-movement is $1 - \delta_{ab}$. Using pairwise comparisons, the collection of $1 - \delta_{ab}$ for each sector a and b catches a glimpse of the business cycle synchronization across the collection of industries. However, one important limitation of this approach is that the propagation of business cycle shocks throughout the interconnected industries could be examined only by splitting the sample.

To overcome this drawback, Leiva-Leon (2014) suggests that independence and perfect dependence constitute two distinct business cycle situations, with the shifts between these extreme regimes governed by the outcome of an unobserved first-order Markov chain, $v_{ab,t}$, whose transition probabilities are given by

$$p(v_{ab,t} = j | v_{ab,t-1} = i, s_{t-2} = r, \dots, \tilde{y}_{ab,t-1}) = p(v_{ab,t} = j | v_{ab,t-1} = i) = p_{ij}^{ab}, \quad (5)$$

where $i, j = 1, 2$, and $\tilde{y}_{ab,t} = (y_{a,t}, y_{a,t-1}, \dots, y_{b,t}, y_{b,t-1}, \dots)'$. Here, the state variable $v_{ab,t}$ reflects the actual state of the business cycle synchronization between industries a and b at time t . In what follows, $v_{ab,t} = 0$ indicates that industries a and b have independent cycles, while $v_{ab,t} = 1$ indicates that their cycles are fully synchronized (or perfectly dependent). Accordingly, $p(v_{ab,t} = 0)$ measures the probability of independent cycles. Within this framework, one can easily examine the evolution of the intersectoral business cycle linkages by collecting $p(v_{ab,t} = 0)$ for all a, b and t .

2.3 Inferences

We collect the parameters that fully characterize the model in a vector θ . It is convenient to define a new state variable that governs the individual business cycles and their degree of synchronization,

$$s_{ab,t}^* = \begin{cases} 1 & \text{if } s_{a,t} = 0, s_{b,t} = 0, \text{ and } v_{ab,t} = 0 \\ 2 & \text{if } s_{a,t} = 1, s_{b,t} = 0, \text{ and } v_{ab,t} = 0 \\ 3 & \text{if } s_{a,t} = 0, s_{b,t} = 1, \text{ and } v_{ab,t} = 0 \\ 4 & \text{if } s_{a,t} = 1, s_{b,t} = 1, \text{ and } v_{ab,t} = 0 \\ 5 & \text{if } s_{a,t} = 0, s_{b,t} = 0, \text{ and } v_{ab,t} = 1 \\ 6 & \text{if } s_{a,t} = 1, s_{b,t} = 0, \text{ and } v_{ab,t} = 1 \\ 7 & \text{if } s_{a,t} = 0, s_{b,t} = 1, \text{ and } v_{ab,t} = 1 \\ 8 & \text{if } s_{a,t} = 1, s_{b,t} = 1, \text{ and } v_{ab,t} = 1 \end{cases}, \quad (6)$$

which also follows a first-order Markov chain.¹ Using an extended version of the procedure described in Hamilton (1989), inferences on the business cycle states are calculated as a by-product of an algorithm, which is similar in spirit to a Kalman filter. Briefly, the procedure is based on the iterative application of the following two steps.²

STEP 1: *Computing the likelihoods.* At time t , the method adds the observation $y_{ab,t} = (y_{a,t}, y_{b,t})'$ to $\tilde{y}_{ab,t-1}$ and accepts as input the forecasting probabilities

$$p(s_{ab,t}^* = i | \tilde{y}_{ab,t-1}, \theta) \quad (7)$$

for $i = 1, 2, \dots, 8$. In this case, the likelihood of $y_{ab,t}$ is

$$f_{ab}(y_{ab,t} | \tilde{y}_{ab,t-1}, \theta) = \sum_{i=1}^8 f_{ab}(y_t | s_{ab,t}^* = i, \tilde{y}_{ab,t-1}, \theta) p(s_{ab,t}^* = i | \tilde{y}_{ab,t-1}, \theta), \quad (8)$$

where $f_{ab}(\bullet)$ is the conditional Gaussian bivariate density function.

To perform inference, the joint probabilities can be obtained from the marginal probabilities as

$$p(s_{ab,t}^* = i | \tilde{y}_{ab,t-1}, \theta) = p(s_{ab,t} = j | v_{ab,t} = h, \tilde{y}_{ab,t-1}, \theta) p(v_{ab,t} = h | \tilde{y}_{ab,t-1}, \theta), \quad (9)$$

with $i = 1, \dots, 8$, $j = 1, \dots, 4$ and $h = 0, 1$. The way in which the model computes inferences on the four-state unobservable variable $s_{ab,t}$ depends on the business cycle synchronization between industries a and b . Suppose that each of these industries follows independent regime-shifting processes, i.e., $v_{ab,t} = 0$. Then, the four-state probability term of $s_{ab,t}$ is

$$p(s_{ab,t} = j | v_{ab,t} = 0, \tilde{y}_{ab,t-1}, \theta) = p(s_{a,t} = j_a | \tilde{y}_{ab,t-1}, \theta) p(s_{b,t} = j_b | \tilde{y}_{ab,t-1}, \theta), \quad (10)$$

with $j = 1, \dots, 4$.

In contrast, if the two industries exhibit perfectly correlated business cycles, which occurs when $v_{ab,t} = 1$, they can be represented by the same state variable since $s_{a,t} = s_{b,t}$. Therefore, one can define a new four-state variable $\varsigma_{ab,t}$ as in (4), where states 2 and 3 never occur and the two industries share the cycle in states 1 and 4. In this case, the probability term is

$$p(s_{ab,t} = j | v_{ab,t} = 1, \tilde{y}_{ab,t-1}, \theta) = p(\varsigma_{ab,t} = j | \tilde{y}_{ab,t-1}, \theta), \quad (11)$$

¹The probabilities of occurrence of states 6 and 7 are zero by definition.

²We focus on the case of industries switching between two regimes. In principle, the analysis can be extended, allowing for three regimes to explore a scenario under extreme situations. However, in such a case, the analysis of synchronization becomes cumbersome and the output would be less easy to interpret.

with $j = 1, \dots, 4$ and $p(\varsigma_{ab,t} = 2|\tilde{y}_{ab,t-1}, \theta) = p(\varsigma_{ab,t} = 3|\tilde{y}_{ab,t-1}, \theta) = 0$. The transition probabilities of $\varsigma_{ab,t}$ are

$$p(\varsigma_{ab,t} = j|\varsigma_{ab,t-1} = i, \varsigma_{ab,t-2} = h, \dots, \tilde{y}_{ab,t-1}) = p(\varsigma_{ab,t} = j|\varsigma_{ab,t-1} = i) = q_{ij}^{ab}. \quad (12)$$

STEP 2: *Updating the forecasting probabilities.* Using the data up to time t , the optimal inference on the state variables can be obtained in the following way:

$$\begin{aligned} p(s_{k,t} = j_k|\tilde{y}_{ab,t}, \theta) &= f_k(y_{k,t}|s_{k,t} = j_k, \tilde{y}_{ab,t-1}, \theta) p(s_{k,t} = j_k|\tilde{y}_{ab,t-1}, \theta) / f_k(y_{k,t}|\tilde{y}_{ab,t-1}, \theta), \\ p(v_{ab,t} = h|\tilde{y}_{ab,t}, \theta) &= f_{ab}(y_{ab,t}|v_{ab,t} = h, \tilde{y}_{ab,t-1}, \theta) p(v_{ab,t} = h|\tilde{y}_{ab,t-1}, \theta) / f_{ab}(y_{ab,t}|\tilde{y}_{ab,t-1}, \theta), \\ p(\varsigma_{ab,t} = j|\tilde{y}_{ab,t}, \theta) &= f_{ab}(y_{ab,t}|\varsigma_{ab,t} = j, \tilde{y}_{ab,t-1}, \theta) p(\varsigma_{ab,t} = j|\tilde{y}_{ab,t-1}, \theta) / f_{ab}(y_{ab,t}|\tilde{y}_{ab,t-1}, \theta), \\ p(s_{ab,t} = j|\tilde{y}_{ab,t}, \theta) &= f_{ab}(y_t|s_{ab,t} = i, \tilde{y}_{ab,t-1}, \theta) p(s_{ab,t} = j|\tilde{y}_{ab,t-1}, \theta) / f_{ab}(y_t|\tilde{y}_{ab,t-1}, \theta), \end{aligned}$$

where $f_k(\bullet)$ is the conditional Gaussian univariate density function of industry j_k , $h = 1, 2$, $j = 1, \dots, 4$, and $k = a, b$.

Finally, one can form the forecasts of how likely the processes are in period $t + 1$, using the observations up to date t . These forecasts can be computed by using the following expressions:

$$p(s_{k,t+1} = j_k|\tilde{y}_{ab,t}, \theta) = \sum_{i_k=0}^1 p(s_{k,t} = i_k|\tilde{y}_{ab,t}, \theta) p_{i_k j_k}, \quad (13)$$

$$p(v_{ab,t+1} = h|\tilde{y}_{ab,t}, \theta) = \sum_{i=0}^1 p(v_{ab,t} = i|\tilde{y}_{ab,t}, \theta) p_{ab,ij}, \quad (14)$$

$$p(\varsigma_{ab,t+1} = j|\tilde{y}_{ab,t}, \theta) = \sum_{i=1}^4 p(\varsigma_{ab,t} = j|\tilde{y}_{ab,t}, \theta) p_{ij}, \quad (15)$$

$$p(s_{ab,t+1} = j|\tilde{y}_{ab,t}, \theta) = \sum_{i=1}^4 p(s_{ab,t} = i|\tilde{y}_{ab,t}, \theta) p_{ab,ij}. \quad (16)$$

The joint probabilities $p(s_{ab,t+1}^* = i|\tilde{y}_{ab,t}, \theta)$ can then be updated by using (9) and be used to compute the likelihood for the next period, as described in the first step.

In the classical approach, the estimate of the parameters of the model is obtained by maximizing the likelihood function (8) by numerical optimization. In this context, performing inferences based on maximum likelihood becomes computationally cumbersome because of the complicated nature of the joint likelihood function. In Appendix A, we describe a multi-move Gibbs-sampling method that makes the Bayesian analysis approach easy to implement. In short, both the parameters of the model, θ , and the Markov-switching variable, $\tilde{s}_{ab,T}^* = s_{ab,1}^*, \dots, s_{ab,T}^*$, are treated

as random variables. This Monte Carlo Markov chain method is achieved by sequentially generating a realization of θ^j from the distribution of $\theta|\tilde{y}_{ab,T}, \tilde{s}_{ab,T}^{*j-1}$, followed by a realization of $\tilde{s}_{ab,T}^{*j}$ from the distribution of $\tilde{s}_{ab,T}^*|\tilde{y}_{ab,T}, \theta^j$. Thus, the marginal distributions of the state variables and the parameters of the model can be approximated by the empirical distribution of the simulated values. The descriptive statistics regarding the sample posterior distributions are then based on 12,000 draws, where the first 2,000 draws are discarded to mitigate the effect of initial conditions.

3 Empirical analysis

3.1 Data description

The data used to measure the industry-level business cycles are the seasonally adjusted, year-over-year growth rates of employment from the U.S. Bureau of Labor Statistics Current Employment Statistics survey. To classify the data, we follow the three-digit industry format of the North American Industry Classification System (NAICS). The effective sample period is January 1991 to May 2013 and the list of 86 industries included in the analysis appears in Table 1.³

Figure 1 plots the annual growth rates of U.S. unemployment since January 1991, where the dates of economic recessions, as determined by the NBER, are indicated with shaded regions. Although employment grows 0.94% during this sample period, the average growth rate in recessions is -1.22% , rising to 1.20% in expansions, which agrees with the well-known procyclical behaviour of employment at a macro level. In addition, this figure shows that changes in employment usually occur a few periods after changes in the economy as a whole.

However, the picture of employment showing micro data is much more complicated. According to the within-phases averages shown in Table 1, employment growth varied a great deal across industries over the sample period. Undoubtedly, not all of the industries boom when the aggregate economy is prosperous and bust when the economy is in recession in a synchronous manner. For example, 31 out of the 86 three-digit industries, mostly related to agriculture and manufacturing, exhibit negative average growth rates over the sample. Filardo (1997) postulates that this might be related to the increasing shift from goods production to services. In addition, agriculture, manufacturing and construction are all among those industries that contract the

³Some industries were not included in the analysis owing to data-availability issues.

fastest during recessions. The cross-industry differences in growth rates during expansions are also large, with agriculture, mining and manufacturing industries even losing employees (Barker, 2011).

Figure 2 plots the nonparametric Gaussian kernel estimates of the densities of industrial employment growth in the NBER recessions and in the NBER expansions. As in the aggregate, the mean of the recession distribution is negative while the mean of the expansion distribution is positive. However, there is a large region of considerable overlapping between these two distributions. This indicates that there are many industries for which employment is falling rather than rising during a national expansion and rising rather than falling during a national recession. Therefore, understanding the business cycle seems to be more complicated than simply analyzing at the aggregate level.

3.2 Univariate analysis

We first conduct an analysis of each industry individually to examine the periods of advance or delay with which these business cycle co-movements might appear. Accordingly, we fit a univariate model like (1) for each of the identified industries and compute the corresponding filtered probabilities of low-mean states, which appear in the choropleth maps displayed in Figures 3 and 4.⁴ The charts are divided into rectangles that show the relative size of employment in each of the 86 industries in the total.⁵ Light colours indicate low evidence of recession, while the darker the shade, the stronger the statistical confidence that the indicated industry was in recession at that time. As one moves to the right, the charts show how the business inferences vary in each quarter, from pre-recession to recession and to the first stages of recovery, as dated by the NBER.

Several conclusions emerge from the analysis of these choropleth charts. First, recessions are marked by widespread contractions in many sectors of the economy. Second, goods-producing industries, complementary businesses, and wholesale and retail industries are among the first to fall at the onset of recessions. However, durable goods industries, professional and technical services, businesses that operate facilities or that provide services to meet varied cultural, entertainment and recreational interests, and industries providing transportation and warehousing and storage for goods do not experience job cuts until some time after the beginning of re-

⁴To facilitate the exposition, the monthly figures have been converted to quarterly by averaging over the respective quarter. The monthly analysis, available upon request, reveals qualitatively similar results.

⁵To simplify the charts, we use the average relative size of industries over the sample period.

cessions. Third, businesses engaged in providing education and training, health care and social assistance and industries providing utility services and public goods are less sensitive to national recessions, especially to the 2001 recession.

3.3 Ergodic linkages

A glimpse of the business cycle linkages across the industries over the sample can be obtained by collecting the pairwise ergodic probabilities of the Markov chain that governs the strength of business cycle synchronization, $v_{ab,t}$, for all industries a, b . Therefore, we begin by calculating the ergodic probabilities,

$$\pi_i^{ab} = \left(1 - p_{jj}^{ab}\right) / \left(2 - p_{00}^{ab} - p_{11}^{ab}\right), \quad (17)$$

that each industry pair exhibits independent ($i = 0, j = 1$) or perfectly synchronized ($i = 1, j = 0$) cycles. Since the ergodic probabilities can be viewed as the unconditional probability of each of the different states, the matrix of ergodic synchronizations provides insights on the unconditional business cycle linkages across the industries. In this paper, we focus on the analysis of business cycle distances, π_0^{ab} .

Although this approach is appealing, a difficulty with it is that there are many such measures. With a set of N industries, there are $\eta = N(N - 1)/2$ different possible business cycle distances. It is therefore a challenge to organize and present the results in a coherent way. To overcome this drawback, we take nonparametric density estimation approaches to examine the distribution of the business cycle distances. These techniques allow us to provide complete information on the entire distribution and have the advantage of letting the data speak for themselves. In this framework, for a given bandwidth h , the kernel distribution of business cycle distances estimated from the ergodic dissimilarities π_0^{ab} is

$$f_h(\pi_0) = \frac{1}{\eta h} \sum_{a=1}^N \sum_{b>a}^N K\left(\frac{\pi_0 - \pi_0^{ab}}{h}\right), \quad (18)$$

where $K(\bullet)$ is the Gaussian kernel.

The density estimate of the cross-industry distributions of pairwise business cycle distances is plotted in Figure 5. A preliminary inspection of the estimated density reveals that this is a multimodal distribution, which shows at least two clear distinct local maxima. The left tail is made up of industries that exhibit large degrees of business cycle synchronization (small distances), whereas the right tail is pretty much exclusively made up of industries with idiosyncratic cycles (big distances). Although most of the distribution's mass is located in the right tail, the

industries experiencing such idiosyncratic cycles tend to be smaller, in terms of total share of U.S. employment, than the industries associated with the left tail, which experience high synchronization. Between these two modes, one might detect a third peak, at around $\pi_0^{ab} = 0.5$. The interpretation of this multimodality is that there is a mixed distribution containing two or three subpopulations of industries with different degrees of business cycle synchronization.

The nonparametric density estimation approach enables us to explicitly test for the number of modes of the underlying distribution. If confirmed, multimodality would point to population heterogeneity, implying the existence of separate population groups. To test for multimodality, we follow the lines suggested by Silverman (1981), who proposed a simple way to assess the p -value that a density is at most k -modal against the alternative that it has more than m modes. Since the number of modes in a normal kernel density estimate does not increase as h increases, let h_m be the minimum bandwidth for which the kernel density estimate is at most m -modal. Let τ_0^{ab} be a resample drawn from the estimated business cycle proximities

$$\tau_0^{ab} = (1 + h_m^2/s^2)^{-0.5} (\pi_0^{ab} + h_m u^{ab}), \quad (19)$$

where s^2 is the sample variance of the data, and u^{ab} is an independent sequence of standard normal random variables. Let h_m^* be the smallest possible h producing at most m modes in the bootstrap density estimate

$$f_h^*(\tau_0) = \frac{1}{\eta h} \sum_{a=1}^N \sum_{b>i}^N K \left(\frac{\pi_0 - \tau_0^{ab}}{h} \right). \quad (20)$$

Repeated many times, the probability that the resulting critical bandwidths h_m^* are larger than h_m , which is equivalent to the proportion of occurrences in which $f_{h_m}^*(\tau_0)$ has more than m modes, can be used as the p -value of the test.

Computed from 1,000 replications, Table 2 displays the critical window widths and the p -values of the null hypothesis that the underlying density has at most m modes against the alternative that it has more than m modes, with $m = 1, 2, 3, 4$. The tests should be applied successively for an increasing number of modes until, for a certain a number, the null is accepted. Clearly, unimodality is rejected for all significance levels (p -value of 0.00), which suggests distinct business cycle distribution dynamics for different population subgroups of industries. In addition, the p -value corresponding to the null of bimodality versus trimodality is 0.27, which indicates that the global distribution of ergodic business cycle distances is bimodal. It exhibits one hump in the very low end representing the industries with a high level of business cycle synchronization and then a larger hump representing those with idiosyncratic cycles.

Notably, the distribution shows a sizable concentration of mass in the middle range. This could explain why the p -value for the test of three modes versus more than three modes falls to 0.12, which is less conclusive. Using a significance level of 0.05, which is the most common cut-off for p -values, the distribution is bimodal. However, using more conservative significance levels, such as 0.15, the distribution would be trimodal (the p -value of four modes rises to 0.26). Therefore, there are signals that the modes located at the tails of the distribution are not well separated and the test does not exclude the possibility that the high-end range of the distribution could be split into two subgroups.

Although useful, the kernel density estimation approach does not allow us to understand the business cycle affiliations detected across the set of industries. To address this deficiency, we employ clustering techniques and classical multi-dimensional scaling (see Timm, 2002, among others) to the pairwise business cycle distances. Collecting the distances, π_0^{ab} , in the symmetric matrix D , the goal of cluster analysis is to develop a classification scheme of our set of industries in several distinct groups, since they present homogeneous business cycles. For this purpose, we make use of dendograms, which are tree-structured graphs used to visualize the result of a hierarchical clustering calculation. The end-points of the dendrogram depicted in Figure 6, whose numbers appear in Table 1, represent the original industries. Clusters are successively combined, forming the tree's branches until the top of the graph. Although it is not easy to interpret, the height of the tree represents the level of dissimilarity at which observations or clusters are merged. Big jumps to join groups occur when there are high intergroup dissimilarities. Therefore, a reasonable number of final groups is often obtained by cutting the dendrogram at those junctures. In line with the results obtained with the kernel approach, the dendrogram shows that two (cutting at around 2) or three (cutting at around 1.5) clusters could be enough to explain the groups that form among the industries.

The multi-dimensional scaling map (see Appendix B) of business cycle similarities is reported in Figure 7, whose plotted numbers refer to the industries listed in Table 1.⁶ Notably, the industries grouped in the two/three different clusters of the dendrogram belong to two/three concentric circles, whose radius lengths reflect the business cycle dissimilarities from the centre to the periphery. The U.S. economy appears in the centre or mass of the distribution of cyclical similarities, reflecting that some industries seem to stick together in the map and show business cycle experiences that are similar to those of the nation. Although they are a relatively small

⁶In these maps, the axes are meaningless and the orientation of the picture is arbitrary.

number of industries, they represent 46.5% of total U.S. employment. However, the map also shows an intermediary zone and a peripheral zone, indicating that other industries appear away from the attractor and do not seem to be as closely related to the nation in terms of business cycle as the industries at the core. The intermediary and peripheral zones are composed of industries representing 20.6% and 32.9% of total employment, respectively.

Let us have a deeper look at these business cycle affiliations. The core, in which the total U.S. employment is also included to facilitate comparison, is plotted in the centre of the map and includes goods-producing industries, which typically experience the largest declines in unemployment during recessions (Bureau of Labor Statistics, 2012), such as construction and textiles, wood, furniture, and electronic products manufacturers. According to Goodman and Mance (2011), complementary businesses that may suffer from ripple effects, such as furniture and food stores, accommodations, appraisal services, motor vehicles, parts manufacturing, and rental and leasing services, are also included. Finally, this core is also formed by other procyclical industries (Bureau of Labor Statistics, 2012), such as wholesale and retail trade and personal services, support activities and business services, especially administrative and waste services.⁷

The contrast between the national business cycle attractor and those industries plotted in the intermediary zone of the perceptual map is a telling indication of their lower (albeit some) business cycle concordance. In this middle circle, we observe some manufacturing industries that may be subject to labour hoarding. Although they depend on the national business cycle, their synchronization could be diminished. According to Clark (1973), examples are durable goods industries, such as chemical, rubber, plastic, primary metal and machinery manufacturing, electrical equipment and building materials. In addition, Rotemberg and Summers (1990) find that industries with a large ratio of nonproduction workers to employment also tend to hoard labour.⁸ Examples are those businesses engaged in providing services in producing and distributing information and cultural products and leisure activities. In addition, this cluster is also formed by most of the industries providing transportation and related facilities, and warehousing and storage for goods. Interestingly, Christiano and Fitzgerald (1998) find that most of the industries belonging to this cluster exhibit strong channels for intermediate goods.

The last cycle cluster is formed by some peripheral industries, which are less closely associated to the U.S. cycle. These industries are plotted further away from the business cycle centre,

⁷Conlon (2011) documented that the payrolls of administrative and waste services shrank by more than 1 million positions during the Great Recession.

⁸Parsons (1986) also documented a stronger tendency to hoard nonproduction labour.

which reflects their low sensitivity to the national cycle. In addition, they appear separate from each other, which indicates that their business cycle shocks are idiosyncratic. This cluster is mainly formed by those industries classified by Berman and Pflieger (1997) as “not coincidentally cyclical” industries. For some of these industries, the consequences of a negative demand shock are relatively reduced since their product cannot normally be postponed. Examples are businesses engaged in providing education and training, health care and social assistance, and industries providing utility services, such as electric power, natural gas, steam supply, water supply and sewage removal.⁹

In addition, this cluster includes sectors that depend highly on international shocks, such as mineral extraction and its related supporting activities, as well as gasoline stations, or on international competition, such as wholesale electronic markets.¹⁰ Also belonging to this cluster are industries providing financial services, not because they are not cyclical, but because they typically lead the national cycle.¹¹ Finally, we find in this cluster monetary authorities, and federal, state and local government services. These industries provide necessities or public goods and demand for these goods remains relatively strong throughout lows in the economy.

3.4 Evolution of linkages

How have industrial business cycle linkages evolved over time? Traditionally, the literature would address this question by dividing a full sample into several sub-periods. The problem with this approach is that it would provide pictures of the cycle linkages that rely on specific date breaks, the location of which is sometimes controversial. To overcome this drawback, the pairwise probabilities of cycle dependence $p(v_{ab,t} = i)$ for all industries a, b , are collected for all periods t . Kernel density estimates, multimodality tests and multi-dimensional scaling maps can then be calculated for each month of the sample. Accordingly, the pictures of Kernel density and multi-dimensional scaling become animated videos that allow one to easily identify which industries manifested the first signs of phase changes and to examine how the interconnections

⁹Goodman (2001) finds that private education and health care services are countercyclical. In fact, employment in these industries has decreased in only 1 of the 12 NBER recessions that have occurred since 1945 (Bureau of Labor Statistics, 2012).

¹⁰Groshen and Potter (2003) find that oil and gas extraction firms are countercyclical.

¹¹Christiano and Fitzgerald (1998) find that the business cycle components of the finance, insurance and real estate industries exhibit low *contemporaneous* co-movement with aggregate employment. Goodman and Mance (2011) show that employment in financial activities peaked one year before the official start of the Great Recession.

among industries propagate the business cycle shocks across industries.¹²

To examine the intra-distribution evolution of business cycle proximities, Figure 8 and Figure 9 plot the overlaid families of empirical kernel distributions across several months around the 2001 and 2008 recessions, respectively. When interpreting these charts, it is worth emphasizing that the horizontal axis measures pairwise business cycle dissimilarities. Therefore, each right-hand horizontal movement represents absolute losses in pairwise synchronization among industries.

Figure 8 shows that in some months before the 2001 recession, the densities exhibited a trimodal distribution of business cycle distances, agreeing with the pattern obtained in the multi-dimensional scaling analysis. That is, we find a core composed of industries highly synchronized with each other (left-hand mode), a group of moderately synchronized industries (middle mode) and some idiosyncratic industries that follow independent cyclical patterns (right-hand mode). However, as the economy approached the recession, the distribution tended to reshape to bimodal. This occurred because industries did not fall into the recession simultaneously but sequentially, as indicated in Figure 3. The first industries to fall into the recession were those in the core, while the industries that belonged to the middle mass remained in expansion. This reduced the synchronization and pushed the middle mass of the distribution to the right-hand side.

During the recession, more industries changed the phase cycle, which implied that a third mode appeared again in the middle of the distribution. However, when the trough took place, the cyclical position was reversed back. The core changed the phase cycle first, accentuating the business cycle discrepancies with the middle mass, which shifted again to the right. Therefore, the distribution was reshaped to bimodal as it did in the peak. Finally, the middle mode appeared again when more industries initiated the recovery phase after the core.

A similar but more accentuated pattern occurred during the 2008 recession. Figure 9 shows that before the recession, the distribution seemed to be characterized by three modes. When the peak occurred, the distribution became bimodal, since industries fell into recession sequentially, providing evidence of a cascade effect. Once the economy was in recession, the trimodal pattern in the distribution was recovered until the peak, when only the industries in the core initiated the recovery and the distribution became bimodal again. The cycle ended when the economy returned to the stable expansionary phase, and the distribution presented three modes, which

¹² The full animated graphs for this paper can be found at <http://www.um.es/econometria/Maximo>.

remained until the next turning point.

In sum, we find that the propagation of micro-level shocks to national shocks is enhanced when the mass shifts that characterize the turning points occur. A formal statistical test of this pattern is provided by applying the modality tests in the density distributions of business cycle similarities from January 1991 to May 2013. According to the plots of the kernel densities, the nulls of unimodality (not shown here to save space) were clearly rejected for all months, since the p -values were always quite close to zero. Figure 10 plots the p -values of the null of two versus more than two modes. To facilitate the analysis, the figure includes shaded areas that refer to the NBER-referenced recessions and a dashed line that refers to the 0.05 significance level. The figure shows that bimodality is rejected during national expansions and recessions while it cannot be rejected at any reasonable level of significance at the turning points.¹³ Notably, the lagged business cycle behaviour that characterizes employment implies that bimodality appears with some lags with respect to the NBER turning points. Therefore, the Silverman test confirms the mass shift in turning points documented above. In other words, the time-varying p -values reaching the 0.05 threshold could be interpreted as indicators of national turning points in (employment) cycles.

Which industries are involved in these large changes in the distribution? To address this question, Figure 11 captures the polarization tendencies around turning points documented in the density estimate analysis. For this purpose, the figure shows three representative months out of the 269 multi-dimensional scaling maps computed in our sample period. Following the NBER classification, the maps capture the across-industry business cycle distances in an expansion, June 2000, and in a recession, March 2009. In both cases, the maps refer to months for which the modality test detected trimodal distributions. The figure also shows the map dated September 2002, which is a month in which the modality test detected only two modes.

The within-expansion and within-recession maps look similar to the map computed from the ergodic probabilities. According to their corresponding trimodal distributions of business cycle distances, they show concentric circles of industries that exhibit highly, moderately and lowly synchronized cycles with each other and with the national business cycle. However, the map that refers to the period of transition from an expansion to a recession reflects a higher dispersion across industries, which agrees with a bimodal distribution of business cycle distances. According to this figure, it appears that the polarization could be owing to the fact that some

¹³Although it is not shown to save space, when bimodality is rejected, trimodality could not be rejected.

industries in the core were engaged in an expansionary phase, while others in the core and the middle circle continued in a contractionary phase.

Let us focus on peaks since troughs behave symmetrically. The movements that occurred during the transitions are initiated by some industries that are in the core and many industries that are in the middle mass, which did not experience the severe job losses that the core did in the initial stages of the recessions. Over the course of the national downturns, these industries start to experience the deepening of the economic downturns and the pace of declines accelerated as employment losses spread to them. Among these industries, we find durable goods industries, private services industries, professional and technical services, recreational services, and industries providing transportation, warehousing and storage for goods.

These industries fall when there is no way back in a national recession, clearly indicating a phase shift. We postulate that this salient characteristic of peaks could be caused by at least one of the following determinants. First, since investments in durable goods can usually be postponed, buyers can benefit from waiting to receive more information about how the economy develops. Therefore, the contagion appears only when the recession becomes more evident. Second, although labour hoarding guarantees that employee talent will be available, the practice implies high risks, since it reduces companies' profitability during difficult times. When the bad times are prolonged, their negative consequences become unavoidable. Third, some of the output of the consumption goods sector, which is in the core, is also used as intermediate goods in the production of durable goods, such as machinery and equipment, which are in the middle mass. During the recession, the durable goods industries are seriously exposed to ripple effects, which come from the core with some lags.

4 Concluding remarks

This paper is part of a growing empirical literature that analyzes the sources of interindustry co-movements. It differs from this literature in several respects. First, the approach is in the mould of "measurement without theory". Using employment data, we ask whether industry-level business cycles are coherent not only with the national business cycle, but also with each other. Second, the filter used to compute the business cycle inferences is an extension of the Markov-switching filter that allows for time-varying business cycle interdependence. Third, nonparametric density estimation techniques are applied to assess the degree of population het-

erogeneity and to examine the changes in the business cycle distribution. Finally, heuristic techniques of classical multi-dimensional scaling and clustering are used to understand the industry movements going in and out of recessions, which helps us to identify changes in cyclical affiliations.

Our main results are the following. First, there is a large heterogeneity in the distribution of business cycle similarities, implying the existence of population groups that follow distinct distributional dynamics. Second, there is not a monotone movement toward the emergence of an increasingly cohesive national business cycle core. The positions of the lower mode, which comprises extremely synchronized industries, and the cluster at the high end of the distribution, which represents industries with idiosyncratic cycles, are relatively stable over time. However, the position of a third, middle mode when the economy is in expansionary or recessionary phases jumps up substantially during the period of transition from one phase to another, switching from pairwise business cycle distances of just over 0.5 to almost one. Therefore, the proposed framework is able to provide assessments of when a national turning point takes place and how the business cycle shocks propagate across industries.

The model used in this paper provides a solid foundation for starting a line of research that seeks to explain the determinants of the business cycle affiliations across industries. Various factors have been put forward in the literature that may affect business cycle synchronization, ranging from the proportion of fixed and variable costs, industry concentration, product differentiation and dependence on external finance. However, the modifications of the model used in the paper to capture the changes in affiliations would be substantial; therefore, this task is left to future research.

Appendix A

This appendix describes the estimation of the parameters in vector θ and the inference on $\tilde{s}_{ab,T}^*$, which is performed through a multi-move Gibbs-sampling procedure. The distribution of the parameters can be approximated by the empirical distributions of simulated values, by iterating the following steps.

STEP 1. The Gibbs sampler is started with arbitrary starting values for the parameters of the model, θ^0 , which are used to generate $\tilde{s}_{ab,T}^{*1}|\tilde{y}_{ab,T}, \theta^0$. For this purpose, we run the Markov-switching filter described in Section 2 and obtain the filtered probabilities $p(s_{ab,t}^*|\tilde{y}_{ab,t}, \theta^0)$. To draw the state variables, we employ the following result:

$$p(\tilde{s}_{ab,T}^*|\tilde{y}_{ab,T}, \theta^0) = p(s_{ab,T}^*|\tilde{y}_{ab,T}, \theta^0) \prod_{t=1}^{T-1} p(s_{ab,t}^*|s_{ab,t+1}^*, \tilde{y}_{ab,t}, \theta^0). \quad (\text{A1})$$

The last iteration of the Markov-switching filter provides us with $p(s_{ab,T}^*|\tilde{y}_{ab,T}, \theta^0)$, from which $s_{ab,T}^*$ is generated. To generate $s_{ab,t}^*$, with $t = 1, \dots, T-1$, we use

$$p(s_{ab,t}^*|s_{ab,t+1}^*, \tilde{y}_{ab,t}, \theta^0) = p(s_{ab,t+1}^*|s_{ab,t}^*) \propto p(s_{ab,t}^*|\tilde{y}_{ab,t}, \theta^0), \quad (\text{A2})$$

where $p(s_{ab,t}^*|s_{ab,t+1}^*)$ refers to the transition probabilities, which are included in θ^0 . Using this expression, it is straightforward to generate $s_{ab,t}^*$ by computing the probability of state i from

$$p(s_{ab,t}^* = i|s_{ab,t+1}^*, \tilde{y}_{ab,t}, \theta^0) = \frac{p(s_{ab,t+1}^*|s_{ab,t}^* = i)p(s_{ab,t}^* = i|\tilde{y}_{ab,t}, \theta^0)}{\sum_{j \neq i} p(s_{ab,t+1}^*|s_{ab,t}^* = j)p(s_{ab,t}^* = j|\tilde{y}_{ab,t}, \theta^0)}. \quad (\text{A3})$$

Using random numbers from a uniform distribution between 0 and 1, $s_{ab,t}^{*1}$ is set to a particular state i by comparing the probability of this state with the random numbers. Following a similar reasoning, one can also generate $\tilde{s}_{k,T}^1 = s_{k,1}^1, \dots, s_{k,T}^1$, $\tilde{v}_{ab,T}^1 = v_{ab,1}^1, \dots, v_{ab,T}^1$ and $\tilde{\varsigma}_{ab,T}^1 = \varsigma_{ab,1}^1, \dots, \varsigma_{ab,T}^1$, for any industry k and any pair a and b at any time $t = 1, \dots, T$.

STEP 2. The generated state variables are used to draw the transition probabilities p_{ij}^k , p_{ij}^{ab} and q_{ij}^{ab} . Since these parameters are drawn in a similar way, we focus only on p_{ij}^k to save space. Conditional on $\tilde{s}_{k,T}$, the transition probabilities are independent of the data set $\tilde{y}_{ab,T}$ and the model's other parameters. Given $\tilde{s}_{k,T}^1$, let n_{ij}^k , $i, j = 0, 1$ be the total number of transitions from state i to state j in industry k . By taking the beta family of distributions as conjugate priors,

$$p_{ii}^k \sim \text{beta}(u_{ii}^k, u_{ij}^k), \quad (\text{A4})$$

where u_{ii}^k and u_{ij}^k are known parameters of the priors, it can be shown that the posterior distributions of p_{ii}^k are given by

$$p_{ii}^k | \tilde{s}_{k,T}, \tilde{y}_{ab,T} \sim \text{beta}(u_{ii}^k + n_{ii}^k, u_{ij}^k + n_{ij}^k), \quad (\text{A5})$$

from which p_{ii}^{k1} is drawn. In particular, we set $u_{00}^k = 8$, $u_{01}^k = 2$, $u_{11}^k = 0$ and $u_{10}^k = 1$ for all k .

STEP 3. Conditional on the covariance matrix Ω_{ab} , the generated state variables and transition probabilities are used to draw the means. Let $\mu_{ab} = (\mu_{a0}, \mu_{a1}, \mu_{b0}, \mu_{b1})'$ be the vector of means for which we assume a normal prior,

$$\mu_{ab} \sim N(\mu_{ab}^*, V_{ab}^*), \quad (\text{A6})$$

where the expected values μ_{ab}^* and the covariance matrix V_{ab}^* are known. The model can now be expressed as

$$\begin{pmatrix} y_{a,t} \\ y_{b,t} \end{pmatrix} = \begin{pmatrix} 1 & s_{a,t} & 0 & 0 \\ 0 & 0 & 1 & s_{b,t} \end{pmatrix} \begin{pmatrix} \mu_{a0} \\ \mu_{a1} \\ \mu_{b0} \\ \mu_{b1} \end{pmatrix} + \begin{pmatrix} \varepsilon_{a,t} \\ \varepsilon_{b,t} \end{pmatrix}, \quad (\text{A7})$$

or

$$y_{ab,t} = D_{ab,t} \mu_{ab} + \varepsilon_{ab,t}, \quad (\text{A8})$$

with $\varepsilon_{ab,t} \sim N(0, \Omega_{ab})$. According to the large business cycle heterogeneity across industries documented in the empirical analysis, we estimate the univariate models by maximum likelihood and use the estimated state-dependent means to specify the parameters μ_{ab}^* of the priors. To check for robustness, we also tried with $\mu_{i0} = y_i^{\min}$ and $\mu_{i1} = y_i^{\max} - y_i^{\min}$, where y_i^{\min} and y_i^{\max} are the minimum and maximum values of employment growth in the i th industry, with $i = a, b$, but the results were unchanged. For the covariance matrices, we set $V_{ab}^* = I$ for all a, b .

The posterior distribution of μ_{ab} is given by

$$\mu_{ab} | \tilde{s}_{ab,T}^{*1}, \tilde{y}_{ab,T}, \Omega_{ab} \sim N(\mu_{ab}^+, V_{ab}^+), \quad (\text{A10})$$

where

$$V_{ab}^+ = \left(V_{ab}^{*-1} + \sum_{t=1}^T D'_{ab,t} \Omega_{ab}^{-1} D_{ab,t} \right)^{-1}, \quad (\text{A11})$$

$$\mu_{ab}^+ = V_{ab}^+ \left(V_{ab}^{*-1} \mu_{ab}^* + \sum_{t=1}^T D'_{ab,t} \Omega_{ab}^{-1} D_{ab,t} y_{ab,t} \right). \quad (\text{A12})$$

STEP 4. Conditional on the generated state variables, transition probabilities and state-dependent means, the parameters of the covariance matrix are drawn. For this purpose, we use the Wishart distribution as the conjugate prior of the inverse covariance matrix,

$$\Omega_{ab}^{-1} \sim W(\Sigma_{ab}^{*-1}, r_{ab}^*), \quad (\text{A13})$$

where Σ_{ab}^* and r_{ab}^* are known. In particular, we set $\Sigma_{ab}^* = I$ and $r_{ab}^* = 0$. Then, the posterior distribution is

$$\Omega_{ab}^{-1} | \tilde{s}_{ab,T}^{*1}, \tilde{y}_{ab,T}, \mu_{ab}^1 \sim W(\Sigma_{ab}^{+-1}, r_{ab}^+), \quad (\text{A14})$$

where

$$r_{ab}^+ = T + r_{ab}^*, \quad (\text{A15})$$

$$\Sigma_{ab}^{+-1} = \Sigma_{ab}^{*-1} + \sum_{t=1}^T (y_{ab,t} - D_{ab,t}^1 \mu_{ab}^1) (y_{ab,t} - D_{ab,t}^1 \mu_{ab}^1)'. \quad (\text{A16})$$

Steps 1 through 4 can be iterated $L + M$ times, where L is large enough to ensure that the Gibbs sampler has converged. Thus, the marginal distributions of the state variables and the parameters of the model can be approximated by the empirical distribution of the M simulated values.

Appendix B

This appendix describes the main steps followed to compute dendograms and multi-dimensional scaling maps. Detailed descriptions of these methods can be found in Timm (2002).

To compute the dendograms, we begin the analysis with $N(N - 1)/2$ clusters, each containing only one industry. Using the matrix of business cycle distances, $D = [d_{ij}]$, the algorithm searches for the “most similar” pairs of industries, so that industry a and b are selected. In this respect, we follow the most similar criterion that is based on the minimum increase in the within-group variance of distances. Industries a and b are now combined into a new cluster, called p , which reduces the total number of clusters by one. Then, dissimilarities between the new cluster and the remaining clusters are computed again following the most similar criterion. For instance, the distance from the new cluster p to, say, industry q , is computed according to

$$d_{p,q} = \frac{n_a + n_q}{n_p + n_q} d_{a,q} + \frac{n_b + n_q}{n_p + n_q} d_{b,q} - \frac{n_q}{n_p + n_q} d_{a,b}, \quad (\text{B1})$$

where n_a , n_b , n_p and n_q are the number of industries included in the respective clusters, and $d_{a,b}$, $d_{a,q}$, and $d_{b,q}$ are the business cycle distances. Finally, these steps are repeated until all industries form a single cluster.

The second classification technique used in this paper is multi-dimensional scaling. To compute these maps, we project the business cycle distances among the N industries in a map in such a way that the Euclidean distances among the industries plotted in the plane approximate the business cycle dissimilarities. In the resulting map, industries that exhibit large business cycle dissimilarities have representations in the plane that are far away from each other. Given the matrix of business cycle distances, D , the technique searches for the so-called $(N \times 2)$ configuration matrix X that contains the position in two orthogonal axes to which each industry is placed in the map. Following Timm (2002), we define

$$B = \frac{1}{2} (I - N^{-1}O) D^2 (I - N^{-1}O), \quad (\text{B2})$$

where I is the identity matrix and O is a $(N \times N)$ of ones. We then compute the (2×2) diagonal matrix Λ with the two largest eigenvalues of B on the main diagonal, and P , the $(N \times 2)$ matrix of its corresponding eigenvectors. The classical metric scaling coordinates correspond to

$$X = P\Lambda^{1/2}. \quad (\text{B3})$$

References

- [1] Acemoglu, D., V. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012), The network origins of aggregate fluctuations. *Econometrica* 80: 1977-2016.
- [2] Barker, M. (2011), Manufacturing employment hard hit during the 2007–09 recession. *Monthly Labor Review*, April: 28-33.
- [3] Bengoechea, P., M. Camacho, and G. Perez Quiros (2006), A useful tool for forecasting the Euro-area business cycle phases. *International Journal of Forecasting* 22: 735-749.
- [4] Berman, J. and J. Pflieger (1997), Which industries are sensitive to business cycles? *Monthly Labor Review*, February: 19-25.
- [5] Bureau of Labor Statistics (2012), The recession of 2007-2009. *BLS Spotlight on Statistics*, February.
- [6] Camacho, M., and G. Perez Quiros (2006), A new framework to analyze business cycle synchronization. In: Milas, C., Rothman, P., and van Dijk, D. *Nonlinear Time Series Analysis of Business Cycles*. Elsevier’s Contributions to Economic Analysis series. Chapter 5 (pp. 133-149). Elsevier, Amsterdam.
- [7] Carlino, G., R. H. DeFina (2004), How strong is co-movement in employment over the business cycle? Evidence from state/industry data. *Journal of Urban Economics* 55: 298-315.
- [8] Christiano, L., and T. Fitzgerald (1998), The business cycle: it’s still a puzzle. *Economic Perspectives*, Fourth Quarter: 56-83.
- [9] Clark, S. (1973), Labor hoarding in durable goods industries. *American Economic Review* 63: 811-824.
- [10] Conlon, F. (2011), Professional and business services: employment trends in the 2007-09 recession. *Monthly Labor Review* April: 34-39.
- [11] Filardo, A. (1997), Cyclical implications of the declining manufacturing employment share. Federal Reserve Bank of Kansas City Economic Review II: 63-87.

- [12] Foerster, A., P. Sarte, and M. Watson (2011), Sectorial versus aggregate shocks: a structural factor analysis of industrial production. *Journal of Political Economy* 119: 1-38.
- [13] Gabaix, X. (2011), The granular origins of aggregate fluctuations. *Econometrica* 79: 733-772.
- [14] Goodman, W. (2001), Employment in services industries affected by recessions and expansions. *Monthly Labor Review* October: 3-11.
- [15] Goodman, W., and S. Mance (2011), Employment loss and the 2007-09 recession: an overview. *Monthly Labor Review* April: 3-12.
- [16] Groshen, E., and S. Potter (2003), Has structural change contributed to a jobless recovery? *Current Issues in Economics and Finance* 9 (August).
- [17] Hamilton, J. (1989), A new approach to the economic analysis of nonstationary time series and the business cycles. *Econometrica* 57: 357-384.
- [18] Hamilton, J., and M. Owyang (2012), The Propagation of Regional Recessions. *The Review of Economics and Statistics*, 94(4): 935-947.
- [19] Leiva-Leon, D. (2014), A New Approach to Infer Changes in the Synchronization of Business Cycle Phases. Working Paper, Bank of Canada, 2014-38.
- [20] Malmendier, U., and S. Nagel (2011), Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics* 126: 373-416.
- [21] Owyang, M., J. Piger, and H. Wall (2005), Business cycle phases in the U.S. states. *The Review of Economics and Statistics* 87: 604-616.
- [22] Owyang, M., J. Piger, H. Wall, and C. Wheeler (2008), The economic performance of cities: A Markov-switching approach. *Journal of Urban Economics* 64: 538-550.
- [23] Parsons O. (1986), The employment relationship: Job attachment, work effort and the nature of contracts. *Handbook of Labor Economics*. Ashenfelter O, Layard R (eds.), Elsevier: New York.
- [24] Peterson, B., and S. Strongin (1996), Why are some industries more cyclical than others? *Journal of Business and Economic Statistics* 14(2) : 189-198.

- [25] Phillips, K. (1991), A two-country model of stochastic output with changes in regime. *Journal of International Economics*, 31 (1-2) 121-142.
- [26] Rotemberg J, and L. Summers (1990), Inflexible Prices and Procyclical productivity. *Quarterly Journal of Economics* 105: 851-874.
- [27] Silverman, B. (1981), Using kernel density estimates to investigate multimodality. *Journal of the Royal Statistical Society, Series B* 43: 97-99.
- [28] Timm, N. (2002), Applied multivariate analysis. Springer texts in Statistics.
- [29] Urquhart, M. (1981), The services industry: Is it recession-proof? *Monthly Labor Review* October: 12-18.

Table 1. Properties of the sectorial business cycles

Industry, 3 digits	Averaged growth rates		
	Total	Expansion	Recession
Forestry (1)	-2.17	-1.52	-7.60
Oil and Gas Extraction (2)	0.15	-0.38	4.58
Mining, except Oil and Gas (3)	-1.26	-1.23	-1.53
Support Activities for Mining (4)	4.03	3.74	6.39
Construction of Buildings (5)	-0.35	0.62	-8.38
Heavy and Civil Engineering Construction (6)	0.47	1.02	-4.10
Specialty Trade Contractors (7)	0.93	1.79	-6.18
Food Manufacturing (8)	-0.10	-0.05	-0.44
Textile Mills (9)	-6.08	-5.30	-12.51
Textile Product Mills (10)	-3.01	-2.40	-8.02
Apparel Manufacturing (11)	-7.66	-7.26	-10.96
Wood Product Manufacturing (12)	-1.83	-0.58	-12.20
Paper Manufacturing (13)	-2.35	-2.10	-4.37
Printing and Related Support Activities (14)	-2.45	-2.04	-5.79
Petroleum and Coal Products Manufacturing (15)	-1.27	-1.49	0.51
Chemical Manufacturing (16)	-1.20	-1.05	-2.48
Plastics and Rubber Products Manufacturing (17)	-0.96	-0.29	-6.48
Nonmetallic Mineral Product Manufacturing (18)	-1.55	-0.84	-7.42
Primary Metal Manufacturing (19)	-2.23	-1.59	-7.56
Fabricated Metal Product Manufacturing (20)	-0.43	0.17	-5.35
Machinery Manufacturing (21)	-0.96	-0.54	-4.42
Computer and Electronic Product Manufacturing (22)	-2.35	-2.11	-4.34
Electrical Equipment, Appliance, and Component Manufacturing (23)	-2.30	-1.96	-5.06
Transportation Equipment Manufacturing (24)	-1.52	-0.71	-8.20
Furniture and Related Product Manufacturing (25)	-2.21	-1.18	-10.71
Miscellaneous Manufacturing (26)	-1.98	-1.90	-2.64
Merchant Wholesalers, Durable Goods (27)	0.07	0.54	-3.78
Merchant Wholesalers, Nondurable Goods (28)	0.21	0.43	-1.61
Wholesale Electronic Markets and Agents and Brokers (29)	2.28	2.63	-0.55
Motor Vehicle and Parts Dealers (30)	0.76	1.41	-4.58
Furniture and Home Furnishings Stores (31)	0.27	1.19	-7.39
Electronics and Appliance Stores (32)	0.63	1.18	-3.94
Building Material and Garden Equipment and Supplies Dealers (33)	1.29	1.88	-3.61
Food and Beverage Stores (34)	0.17	0.23	-0.36
Health and Personal Care Stores (35)	1.13	1.20	0.56
Gasoline Stations (36)	-0.31	-0.12	-1.95
Clothing and Clothing Accessories Stores (37)	0.42	0.81	-2.80
Sporting Goods, Hobby, Book, and Music Stores (38)	1.05	1.29	-0.91
General Merchandise Stores (39)	0.97	1.14	-0.40
Miscellaneous Store Retailers (40)	0.46	0.90	-3.16
Nonstore Retailers (41)	0.31	0.61	-2.14
Air Transportation (42)	-0.62	-0.55	-1.16
Rail Transportation (43)	-0.68	-0.47	-2.43
Water Transportation(44)	0.55	0.73	-0.89
Truck Transportation (45)	0.94	1.52	-3.82
Transit and Ground Passenger Transportation (46)	2.34	2.42	1.68
Pipeline Transportation (47)	-1.27	-1.70	2.31
Scenic and Sightseeing Transportation (48)	2.59	2.79	1.00
Support Activities for Transportation (49)	2.18	2.46	-0.12
Couriers and Messengers (50)	1.69	2.19	-2.44

Warehousing and Storage (51)	2.41	2.80	-0.75
Utilities (52)	-1.26	-1.48	0.55
Publishing Industries, except Internet (53)	-0.71	-0.40	-3.23
Motion Picture and Sound Recording Industries (54)	1.93	2.47	-2.50
Broadcasting, except Internet (55)	0.05	0.27	-1.73
Telecommunications (56)	-0.59	-0.53	-1.07
Data Processing, Hosting, and Related Services (57)	0.85	1.24	-2.37
Other Information Services (58)	5.54	6.11	0.84
Monetary Authorities - Central Bank (59)	-1.47	-1.73	0.76
Credit Intermediation and Related Activities (60)	0.34	0.69	-2.56
Securities, Commodity Contracts, and Other Financial Investments and Related Activities (61)	2.73	3.04	0.22
Insurance Carriers and Related Activities (62)	0.69	0.71	0.47
Funds, Trusts, and Other Financial Vehicles (63)	2.24	2.22	2.33
Real Estate (64)	1.14	1.35	-0.59
Rental and Leasing Services (65)	0.10	0.59	-3.95
Lessors of Nonfinancial Intangible Assets, except Copyrighted Works (66)	2.58	2.94	-0.43
Professional, Scientific, and Technical Services (67)	2.60	2.83	0.66
Management of Companies and Enterprises (68)	0.89	0.98	0.11
Administrative and Support Services (69)	2.69	3.76	-6.16
Waste Management and Remediation Services (70)	2.24	2.47	0.32
Educational Services (71)	3.14	3.06	3.77
Ambulatory Health Care Services (72)	3.70	3.74	3.42
Hospitals (73)	1.42	1.30	2.40
Nursing and Residential Care Facilities (74)	2.47	2.41	2.97
Social Assistance (75)	4.16	4.22	3.67
Performing Arts, Spectator Sports, and Related Industries (76)	1.94	2.20	-0.25
Museums, Historical Sites, and Similar Institutions (77)	3.20	3.39	1.62
Amusement, Gambling, and Recreation Industries (78)	2.73	3.03	0.20
Accommodation (79)	0.57	0.89	-2.07
Food Services and Drinking Places (80)	1.97	2.19	0.14
Repair and Maintenance (81)	0.80	1.20	-2.45
Personal and Laundry Services (82)	0.76	0.85	0.03
Religious, Grantmaking, Civic, Professional, and Similar Organizations (83)	1.46	1.45	1.53
Federal Government (84)	-0.55	-0.58	-0.30
State Government (85)	0.72	0.62	1.55
Local Government (86)	1.14	1.09	1.54
United States	0.94	1.20	-1.23


Notes. Industries by NAICS Code. Numbers with which they appear in multi-dimensional scaling maps are in brackets. The last three columns refer to the average growth rate of employment (total, within the NBER expansions and within the NBER recessions). The shading pattern refers to the classification of industries at two digits.

Table 2. Bootstrap multimodality tests

N. modes under null	Critical bandwidth	<i>p</i> -value
1	0.141	0.00
2	0.031	0.27
3	0.028	0.12
4	0.021	0.26


Notes. Each row shows the number of modes under the null, the critical bandwidth and the corresponding *p*-value of the Silverman (1981) test.

Figure 1. U.S. employment




Notes. The figure plots annual growth rates. Shaded areas indicate NBER recessions.

Figure 2. Kernel distribution over the cycle



Notes. The figure plots the kernel distribution of year-over-year employment growth across industries. Expansions and recessions are obtained from the NBER business cycle dates.

Figure 3. Univariate Markov-switching analysis




Notes. The figure plots the choropleth map of industries for different quarters surrounding the 2001 recession (bold red letters). Each chart corresponds to one quarter; the partitions of each chart represent U.S. industries; and the size of the partitions represents the industries' share of national employment. The darker an industry's partition, the higher the probability of recession for that industry during that quarter.


Figure 4. Univariate Markov-switching analysis

2008 recession


2007Q2




2007Q3




2007Q4




2008Q1




2008Q2




2008Q3




2008Q4




2009Q1




2009Q2




2009Q3




2009Q4




2010Q1





2010Q2



2010Q3




2010Q4



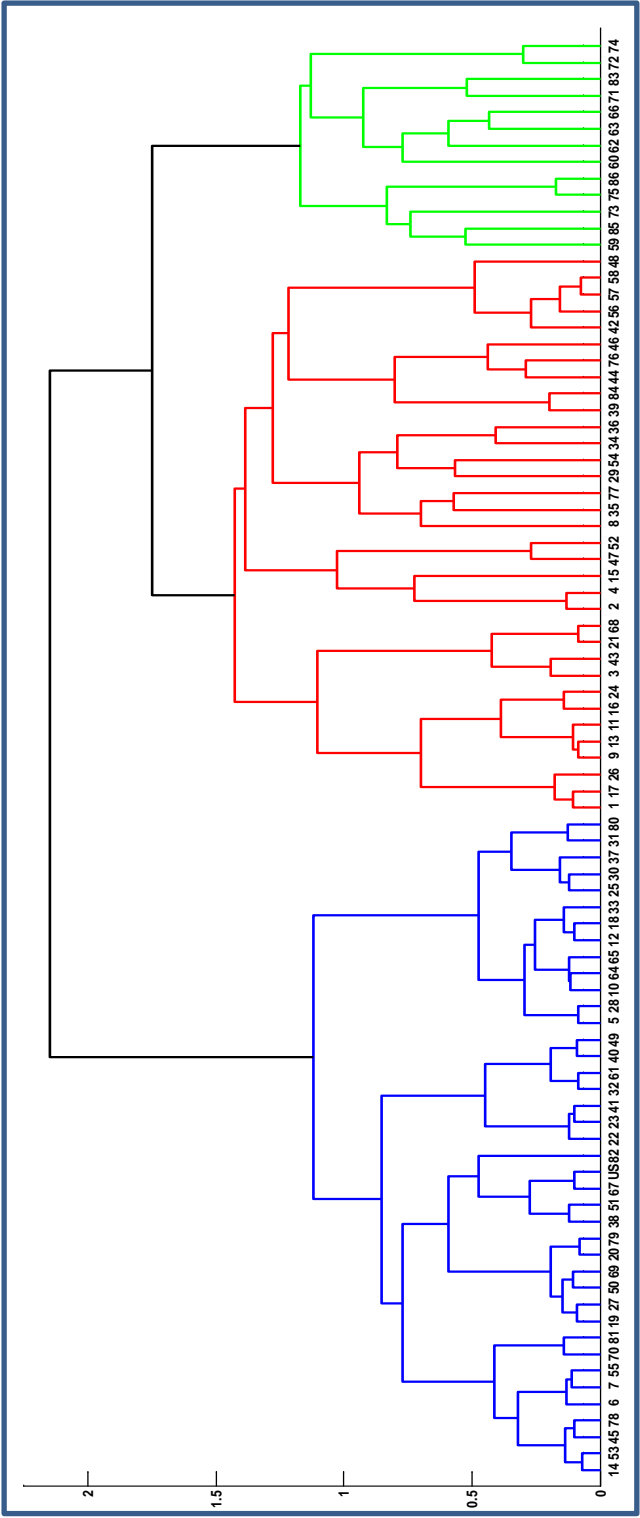
Notes. The figure plots the choropleth map of industries for different quarters surrounding the 2008 recession (bold red letters). Each chart corresponds to one quarter; the partitions of each chart represent U.S. industries; and the size of the partitions represents the industries' share of national employment. The darker an industry's partition, the higher the probability of recession for that industry during that quarter.

Figure 5. Kernel density estimation from ergodic probabilities




Notes. Density estimates of the cross-industry distributions of pairwise business cycle synchronization using ergodic probabilities.

Figure 6. Cluster analysis with ergodic probabilities



Notes. The dendrogram's heights represent the level of dissimilarity at which observations or clusters are merged. Numbers used to represent the industries are provided in Table 1.


Figure 7. Multi-dimensional scaling map of ergodic probabilities



Notes. The map plots in a two-dimensional scale the business cycle distances across the industries. Numbers used to represent the industries are provided in Table 1.

Figure 8. Kernel density estimation with time-varying probabilities

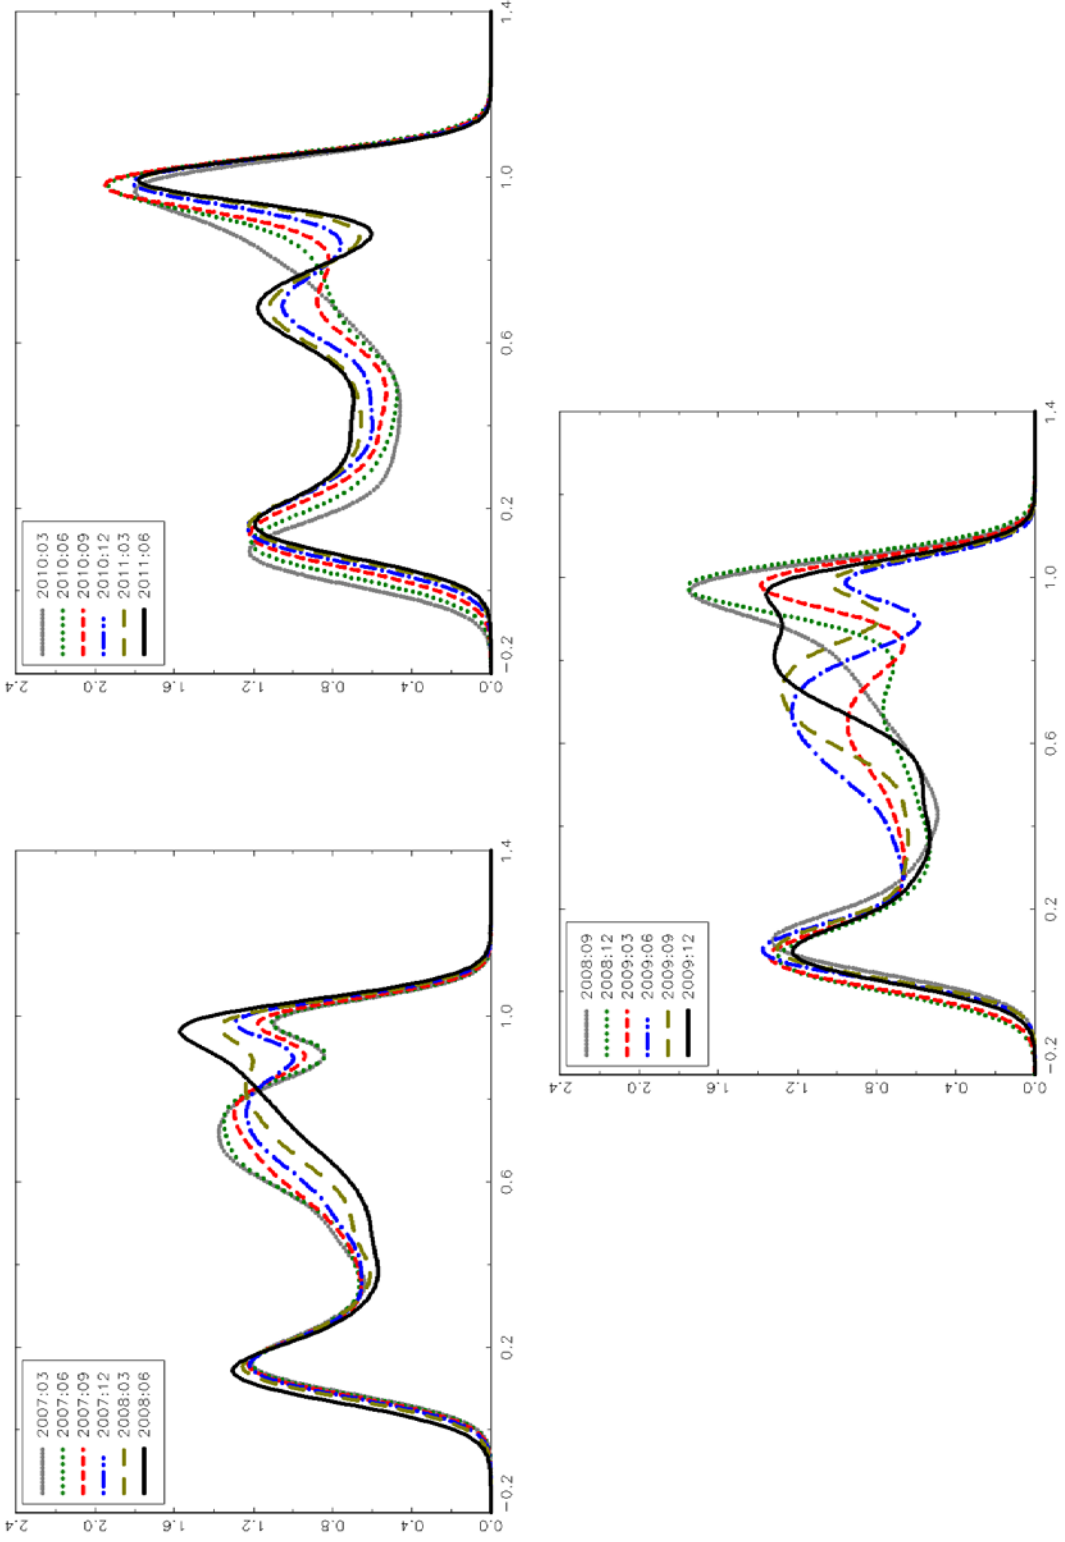
2001 recession



Notes. Density estimates of the cross-industry distributions of pairwise business cycle synchronization for different months surrounding the 2001 recession.


Figure 9. Kernel density estimation with time-varying probabilities

2008 recession




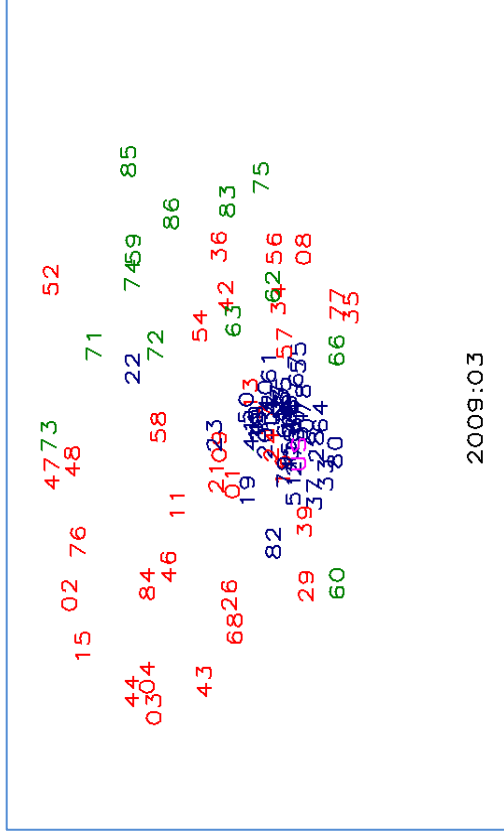
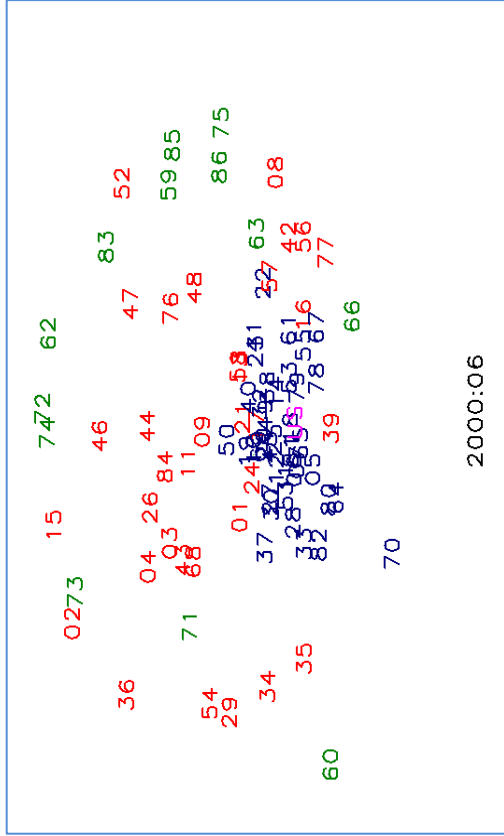
Notes. Density estimates of the cross-industry distributions of pairwise business cycle synchronization for different months surrounding the 2008 recession.

Figure 10. Multimodality test



Notes. The figure plots the p -value of the null of at most two modes against the alternative of more than two modes.

Figure 11. Kernel density estimation with time-varying probabilities



Notes. See notes to Figure 7.