

THE FEDERAL RESERVE BANK of KANSAS CITY
ECONOMIC RESEARCH DEPARTMENT

A Chronology of International Business Cycles Through Non- parametric Decoding

Hsieh Fushing, Shu-Chun Chen, Travis J. Berge, and Òscar Jordà

October 2010

RWP 11-13



RESEARCH WORKING PAPERS

A Chronology of International Business Cycles Through Non-parametric Decoding*

Abstract

This paper introduces a new empirical strategy for the characterization of business cycles. It combines non-parametric decoding methods that classify a series into expansions and recessions but does not require specification of the underlying stochastic process generating the data. It then uses network analysis to combine the signals obtained from different economic indicators to generate a unique chronology. These methods generate a record of peak and trough dates comparable, and in one sense superior, to the NBER's own chronology. The methods are then applied to 22 OECD countries to obtain a global business cycle chronology.

- *Keywords:* decoding, hierarchical factor segmentation, network analysis, business cycles.
- *JEL Codes:* C14, C54, E32, F42

Hsieh Fushing
Department of Statistics
University of California, Davis
e-mail: fushing@wald.ucdavis.edu

Shu-Chun Chen
Institute of Mathematics
Academia Sinica, Taiwan
e-mail: d92621201@ntu.edu.tw

Travis Berge
Department of Economics
University of California, Davis
e-mail: tjberge@ucdavis.edu

Òscar Jordà (*corresponding author*)
Department of Economics
University of California, Davis
e-mail: ojorda@ucdavis.edu

*Fushing is grateful for partial support by the NFS under Grant No: BCS 0527766, HSD 0826844 and DMS-1007219 (cofunded by the Cyber-enabled Discovery and Innovation –CDI– program), and by the NIH under Grant No: 1R01AG025218-01A2. Jordà acknowledges the hospitality of the Federal Reserve Bank of San Francisco during the preparation of this manuscript.

1 Introduction

Even before the onset of the recent financial crisis—the most severe global crisis since the Great Depression—an extensive collection of papers investigated the evolution of global economic fluctuations, contagion, convergence phenomena, and the synchronization of international economic activity (to cite a few Kose, Prasad and Terrones, 2003; Amber, Cardia and Zimmermann, 2004; Stock and Watson, 2005; Kose, Otrok and Whiteman, 2008; and Kose, Otrok and Prasad, 2008). This fascination with the manner in which economies are interlinked emanates from the remarkable growth of trade and capital flows experienced since the early 1980's (e.g. Kose, Prasad and Terrones, 2003). However, in virtually every paper the approach is the same: Are there common factors (in the statistical meaning of the word) that explain fluctuations of economic activity across several countries?

In this paper we shift the focus from economic fluctuations to business cycles. It is important to clarify the difference between the two:

Business Cycles are not merely fluctuations in aggregate economic activity. The critical feature that distinguishes them from commercial convulsions of earlier centuries or from seasonal or other short term variations of our own age is that the fluctuations are widely diffused over the economy—its industry, its commercial dealings, and its tangles of finance. The economy of the Western World is a system of closely interrelated parts.

Arthur F. Burns in the Introduction to Wesley C. Mitchell (1951) **What Happens During Business Cycles: A Progress Report**. National Bureau

of Economic Research: New York.

Building on the work by Burns and Mitchell (1946) in the U.S., the National Bureau of Economic Research (NBER) formally established a Business Cycle Dating Committee (BCDC) in 1978 (although the NBER, founded in 1920, had published its first business cycle dates in 1929 following the work of Mitchell, 1927). The BCDC produces a formal chronology of expansions and recessions that now dates as far back as the trough of 1854. In 2002 the Center for Economic Policy Research followed suit and founded the Euro Area Business Cycle Committee. Similar official chronologies of business cycle turning points are not easily found for many other countries.

It is instructive to quote here the definition of recession provided by the NBER:

A recession is a significant decline in economic activity spread across the economy, lasting more than a few months, normally visible in production, employment, real income, and other indicators.

–Determination of the December 2007 Peak in Economic Activity, December 11, 2008. Business Cycle Dating Committee of the National Bureau of Economic Research

This classification of economic activity into expansions and recessions recognizes that these two states are fundamentally different from each other. Declines in economic activity tend to be sudden but short-lived, whereas expansions tend to be more gradual and last for a longer period. For example, in the U.S., the average duration of a recession is 11 months, whereas expansions last on average four years (see Berge and Jordà, 2010).

Our goal in this paper is to first provide intuitive and robust methods to construct country-specific business cycle histories, and then to produce the foundation for regional or global business cycle chronologies. We think such an exercise is of considerable interest to policy-makers for several reasons. First, our focus is on the timing of turning points across countries, and not on the correlation of an economic indicator across countries and time. Focusing on the classification of economic activity itself is a fundamentally different problem than looking for associations in economic fluctuations –one that we feel is central to effective policy-making. Second, regional chronologies offer an alternative measure of contagion that may be more conducive to analyzing how international shocks propagate across economic regions. Third, and as the current crisis shows, international institutions such as the International Monetary Fund (IMF) and the World Bank are naturally concerned about the impact of simultaneous fiscal expansions in many countries to the extent that such packages can have implicit protectionist effects. For these reasons, information about the state of an economy is in many ways more constructive than examining the fluctuations of an economy per se.

Classification of economic activity is difficult since the underlying state of the economy is intrinsically unobservable. Moreover, the definitions used by dating committees are often not operational. It is not clear what is meant by “economic activity,” or how one should construct such a variable. Moreover, even if such a variable could be constructed, dating committees do not sufficiently clarify how one then decides on what is an expansion and what is a recession, and very often simple ‘rules-of-thumb,’ such as defining a recession as

two consecutive quarters of negative GDP growth, classify economic activity quite poorly (see, e.g. Berge and Jordà, 2010).

An important contribution of our paper is to introduce a novel statistical approach to the problem of classifying economic activity. The question of whether or not the economy is experiencing a recession is viewed as a computational learning or pattern recognition problem, more specifically, a decoding problem. Decoding is most often referred to in information theory as an algorithm for recovering a sequence of code words or messages from a given sequence of noisy output signals (Geman and Kochanek, 2001). Since our approach differs from traditional practices in economics, it is useful to provide some context here before we introduce the methods formally in the next section.

Perhaps the best-known statistical approach to uncovering the latent state of the economy is based on specifying the state-space with a hidden-Markov mixture model, of which Hamilton’s (1989) well-known regime-switching model is an example. These models specify regime-specific stochastic processes for a given variable (usually, but not exclusively, GDP), where the transition between regimes is regulated by a latent Markov process. Regime-switching models have the disadvantage of being parametrically intensive, and require that one impose assumptions regarding the evolution of economic activity that may or may not be sensible. Instead, the procedures that we propose are non-parametric, parsimonious and computationally simple. We want to avoid having to make structural assumptions about state-space dynamics. To do so, we utilize techniques that are generally referred to as “non-parametric decoding;” in particular, we use the Hierarchical Factor Segmentation (HFS)

algorithm introduced in Fushing, Hwang, Lee, Lang and Horng (2006).

HFS applies maximum entropy methods to derive the recurrence time distribution of certain pre-specified events observed in the data. That is, HFS algorithmically partitions the data into segments where the intensity of the recurrence distribution in adjacent segments differ significantly. Then, using model selection criteria, one can select the best partition of the data and use the likelihood ratio principle to compare this partition to the null that the recurrence distribution has a constant intensity. The idea of looking at recurrence times dates back at least to Poincaré (1890). In this way, the decoding problem is transformed into an event intensity change-point problem where there is no prior knowledge about the number of change-points.

Applying HFS to economic indicators produces patterns of event intensity (expansions and recessions). As an example, consider the pre-specified event of negative GDP growth. Observed U.S. GDP growth was negative in about 17 percent of the quarters between 1947:I and 2009:IV. Under the null that the recurrence distribution has a constant intensity, one would expect to observe a negative quarter of GDP growth every 6 or 7 quarters. In the data, however, we instead observe that negative GDP growth comes in clusters, each of which we usually identify as a recession. It is in this sense that the HFS algorithm searches for patterns—the algorithm is able to partition a sequence of data by exploring the recurrence time of a particular event.

Our search for recessions and expansions internationally is guided by the definition of recession given by the NBER’s dating committee. Since ‘economic activity’ cannot be de-

scribed by simply one observable (GDP), we consider several indicators for many different countries. The HFS algorithm is applied to each indicator separately. In order to combine information obtained in this manner from different indices, we then apply network analysis (e.g. Watts and Strogatz, 1998) to determine with more precision the onset of cyclical phases. Network analysis has the flavor of a majority voting rule, a commonly used tool in pattern recognition problems (see, e.g. Hastie, Tibshirani and Friedman, 2009).

We next describe in detail the techniques used to classify economic activity. In the interest of checking the robustness of our findings, we then compare the chronology that results from our procedure to the chronologies available for the U.S. by the NBER. This serves as a benchmark for readers to assess how our methods work in practice. We then apply HFS to a collection of 22 OECD countries. We aim to generate a chronology of business cycles for each individual country that can serve as a benchmark for those countries that do not have business cycle dating committees, as well as a single global business cycle chronology since these countries represented over 60% of world output in 2009.

2 Statistical Design

Dating business cycle turning points is a classic pattern-recognition problem, as the definition of what a recession is quoted in the previous section shows. However, as Berge and Jordà (2010) discuss, there is considerable tension within the NBER's BCDC when grappling with the concept of *economic activity*, in particular, how to combine signals from output and employment indicators. Evidence of this tension was apparent in a brief statement issued by

the seven member BCDC on April 12, 2010, in which the committee states that “it would be premature to set a date marking the end of the last economic contraction and the beginning of an expansion” (see, e.g., the related story in the Los Angeles Times, April 13, 2010)¹ . The statement was followed by a dissenting (and surprisingly candid) press release by one of its members, Robert J. Gordon.

The committee decided not to declare that the recession was over and thus did not devote serious attention to determining the trough month. However, the data are clear, once we focus on measures of production rather than measures of employment, which historically have lagged the recovery of production.

–Robert J. Gordon, quoted in the article “*Dissent on Recession’s End*” by Catherine Rampell in the New York Times, April 12, 2010.²

Since then, the end of the recession was announced to have been June 2009, but that pronouncement was made a few months after this quote, on September 20, 2010.

Even if there were to be agreement on what economic activity is, the underlying state of the economy (expansion/recession) is not directly observable and has to be inferred. In fact, Hamilton (2005) suggests that the inability to define a business cycle as a fundamental attribute of the data-generating process is inherent in any time series model. One reason is that the recurrence of cyclical phenomena is itself stochastic. Burns and Mitchell (1946) used graphs to summarize the heterogeneity of business cycle patterns. Subsequently, Stock

¹ <http://www.latimes.com/business/la-fi-recession13-2010apr13,0,3381428.story>

² <http://economix.blogs.nytimes.com/2010/04/12/dissent-on-recessions-end/>

(1987, 1988) argued that this way of thinking about the data necessarily implies a nonlinear data generating process. In order to reclaim some regularity to the recurrence of business cycles, he proposed the idea of time-deformation.

In this paper, we abandon the idea of fitting an economic indicator via statistical modeling in either the original or a re-scaled temporal axis. Instead we use a computational approach: a diagnostic testing mechanism combined with a network. This approach takes advantage of the synchronicity of patterns across economic indicators as suggested by the NBER's definition of recession, and then summarizes them via a network's connectivity properties. Because we will be looking at several indicators and several countries simultaneously, the high-dimensionality of the data makes it impractical to use structural hidden-Markov models such as Hamilton's (1989) regime-switching model.

2.1 The Null Distribution When the State-Space Trajectory is Constant

Let $\mathcal{X}_n = \{X_t\}_{t=1}^n$ denote the observed time series for an economic indicator (it will help the reader to think of X_t as the annual growth rate in the industrial production index, for example). Under the null hypothesis the trajectory of \mathcal{X}_n is determined by a unique state. Under the alternative hypothesis, two state-space trajectories for \mathcal{X}_n are determined depending on the state variable $\mathcal{S}_n = \{S_t\}_{t=1}^n$, $S_t \in \{0, 1\}$ where, say, $S_t = 0$ for expansion; 1 for recession. The task of recovering \mathcal{S}_n from \mathcal{X}_n is called decoding. When no assumptions are imposed on state-space dynamics, this task is called non-parametric decoding.

Under the view that the evolution of \mathcal{X}_n is driven by two state-space trajectories, consider

a separating event A that is more likely to be associated with one state than the other. For example, when Hsieh, Chen and Pollard (2009) decode DNA sequences, an event corresponds to the observation of a CG di-nucleotide, which marks genes often associated with certain diseases. In Hsieh (2010) the event is based on the volatility of the returns of a financial asset. Here we choose observations of \mathcal{X}_n such that $X_t < h$ where h is pre-specified in a manner that we discuss momentarily.

The choice of h is best understood with an example. Figure 1 shows kernel density estimates for the empirical distribution of the annual growth in the industrial production index (IP) for the U.S. (i.e. the log difference of a given quarter from that quarter in the previous year), using the NBER's chronology of recessions. The idea is that an observation of IP growth at the bottom of the 30th percentile (a precise cut-off is not necessary as we shall see) of the overall empirical distribution of IP is more likely to be associated with a recession than an expansion. Smaller values of h generate a lower rate of false positives (or (1 - *sensitivity*), i.e. $\Pr[X_t < h | S_t = 0]$), but also result in a lower true positive rate (or *sensitivity*, i.e. $\Pr[X_t < h | S_t = 1]$). Using the NBER recession dates, IP turns out to have an area under the ROC curve (AUC) of 0.88 (the ROC curve is a tool used in classification evaluation where a value of $AUC = 1$ implies perfect classification and a value of $AUC = 0.5$ no better than coin-toss classification ability, see Berge and Jordà, 2010). The AUC is a Mann-Whitney statistic that measures the distance between the recession and expansion distributions. For $h = 30^{th}$ percentile, the true positive rate is $TP(h) = 0.78$; and the false positive rate is $FP(h) = 0.20$, which strikes a nice balance between successfully detecting

recessions often, with a relatively low rate of false positives.

Although a precise choice of h may appear critical to obtain a proper partition of the sample into expansions and recessions, in practice this is not the case for reasons that will become clear momentarily. We have experimented with Monte Carlo simulations not reported here and found that values of h representative of the 30-40th percentile distribution of \mathcal{X}_n strike a good balance between specificity and sensitivity (or true positives and false positives). We use the 30th percentile in our applications.

The number of periods between consecutive separating events is called the recurrence time. That is, let $\mathcal{C}_n = \{C_t^{(0)}\}_{t=1}^n$ such that $C_t^{(0)} = 1$ if $X_t < h$, 0 otherwise with $\sum_{t=1}^n C_t^{(0)} = n(1)$, which counts the number of events for which $X_t < h$, and hence define $\tau_k^{(1)} = \inf\{t | \tau_{k-1}^{(1)} < t < n \text{ such that } C_t^{(0)} = 1\}$ where k denotes the k^{th} event. The use of the superscript (j) for $j = 0, 1$ will become clear momentarily but it essentially denotes the segmentation level. Then the recurrence times are $R_k^{(1)} = \tau_k^{(1)} - \tau_{k-1}^{(1)}$ so that $\mathcal{R}_{n(1)} = \{R_k^{(1)}\}_{k=1}^{n(1)}$. When \mathcal{X}_n is generated by a unique state-space trajectory, the recurrence time distribution for $\mathcal{R}_{n(1)}$ can be shown to be geometric $G(R_k; \lambda)$ with intensity parameter λ . This result is based on the maximum entropy principle of Jaynes (1957a, b), where one can also show that the intensity parameter is a function of the average recurrence time.

Specifically, let the probability distribution of $\mathcal{R}_{n(1)}$ be denoted by $p_i = \Pr[R_k^{(1)} = i]$, $i = 0, 1, \dots$ with entropy $H(\mathcal{R}_{n(1)}) = -\sum_{i=0}^{n(1)} p_i \log p_i$ such that $E[R_k^{(1)}] \equiv r = \sum_{i=0}^{\infty} i p_i$ and with the obvious constraint that $\sum_{i=0}^{\infty} p_i = 1$. Then, Fushing, Chen and Hwang (2010a, b)

solve for the maximum of

$$Q(p, \lambda) = \sum_{i=0}^{\infty} p_i \log p_i + \eta_1 \left(\sum_{i=1}^{\infty} i p_i - c \right) + \eta_2 \left(\sum_{i=1}^{\infty} p_i - 1 \right)$$

where η_1 and η_2 are Lagrange multipliers associated with the constraints, and show that

$$p_1 = 1/r;$$

which is easily estimated from

$$\hat{r} = \frac{\sum_{k=1}^{n(1)} R_k^{(1)}}{n(1)}.$$

In addition,

$$p_i = 1/r e^{\lambda(i-1)}$$

with intensity

$$\lambda = \log \left\{ \frac{r}{1-r} \right\},$$

Fushing, Chen and Hwang (2010a, b) denote this geometric distribution as $G(R_k; \lambda)$, where λ denotes its intensity.

In the next section we present the HFS algorithm, which partitions the data into segments with recurrence times generated by the optimal mixture of geometric distributions and which can be compared to the recurrence time distribution from a conventional mixture model using the likelihood ratio principle, as we will show.

2.2 The Hierarchical Factor Segmentation (HFS) Algorithm

We introduce some additional notation to provide intuition for the method we present in this section. Let \mathcal{P}_0 denote the distribution of X_t when $S_t = 0$ and \mathcal{P}_1 when $S_t = 1$ under

the view that \mathcal{X}_n is generated by a mixture process. A key feature of the HFS algorithm is that it does not construct likelihood functions for \mathcal{P}_0 and \mathcal{P}_1 or the transition probabilities, as a typical hidden-Markov model would do.

The HFS algorithm partitions the data \mathcal{X}_n according to h using two sequential segmentations of the resulting empirical recurrence times meant to non-parametrically seek the optimal mixture. The basic idea is to generate alternating sequences of high/low recurrence intensity and then choose that partition that minimizes a noise-to-signal criterion to be explained below.

The HFS algorithm can therefore be summarized as the following collection of steps, some of which have been discussed already:

HFS-1 Transform \mathcal{X}_n into a 0-1 digital string $\mathcal{C}_n = \{C_t^{(0)}\}_{t=1}^n$ such that $C_t^{(0)} = 1$ if $X_t < h$, 0 otherwise; with $\sum_{t=1}^n C_t^{(0)} = n(1)$, the number of identified events in which $X_t < h$.

HFS-2 Let

$$\tau_k^{(1)} = \inf\{t | \tau_{k-1}^{(1)} < t < n \quad \text{s.t.} \quad C_t^{(0)} = 1\}$$

where $\tau_k^{(1)}$ is event-time for the k^{th} occurrence and hence define the recurrence times

$$R_k^{(1)} = \tau_k^{(1)} - \tau_{k-1}^{(1)} \text{ so that } \mathcal{R}_{n(1)} = \{R_k^{(1)}\}_{k=1}^{n(1)}.$$

HFS-3 Let ρ_1 denote the upper ρ_1^{th} - percentile of the empirical distribution of $\mathcal{R}_{n(1)}$ and

hence let the 0-1 digital string $\mathcal{C}_{n(1)} = \{C_k^{(1)}\}_{k=1}^{n(1)}$ be defined such that $C_k^{(1)} = 1$ if

$R_k^{(1)} > \rho_1$ and note that $\sum_{k=1}^{n(1)} C_k^{(1)} = n(2)$, the number of events in which $R_k^{(1)} > \rho_1$.

HFS-4 Let

$$\tau_k^{(2)} = \inf\{\tau_k^{(1)} | \tau_{k-1}^{(2)} < \tau_k^{(1)} < n(1) \quad \text{s.t.} \quad C_k^{(1)} = 1\}$$

and hence define the recurrence times $R_k^{(2)} = \tau_k^{(2)} - \tau_{k-1}^{(2)}$ so that $\mathcal{R}_{n(2)} = \{R_k^{(2)}\}_{k=1}^{n(2)}$.

HFS-5 Let ρ_2 denote the upper ρ_2^{th} -percentile of the empirical distribution of $\mathcal{R}_{n(2)}$ and hence

let the 0-1 digital string $\mathcal{C}_{n(2)} = \{C_k^{(2)}\}_{k=1}^{n(2)}$ be defined such that $C_k^{(2)} = 1$ if $R_k^{(2)} > \rho_2$.

HFS-6 $\mathcal{C}_{n(2)}$ can then be mapped into \mathcal{C}_n (since $\tau_k^{(2)}$ maps into $\tau_k^{(1)}$ and $\tau_k^{(1)}$ into t) or \mathcal{X}_n as a

partition of $\sum_{k=1}^{n(2)} C_k^{(2)} = m$ segments on the time span $[1, n]$. Denote this partition as

$\mathcal{N}(\mathcal{X}_n) = \{[P_i, T_i]; [T_i, P_i]\}_{i=1}^m$ where P_i denotes ‘‘peak’’ and T_i denotes ‘‘trough’’ to use

the same nomenclature as the NBER’s dating so that the first set of segments refers

to recessions (peak to trough) and the second set of segments to expansions (trough

to peak).

Thus, the partition $\mathcal{N}(\mathcal{X}_n)$ separates high event-intensity from low-intensity segments.

Segments $[P_i, T_i)$ corresponding to recessions will have a high proportion of ones in the code-

sequence \mathcal{C}_n whereas segments $[T_i, P_i)$ corresponding to expansions will have a low proportion

of ones. The partition therefore gives rise to the recession indicator $\mathcal{S}_{n(1)} = \{S_k^{(1)}\}_{k=1}^{n(1)}$ with

$S_k^{(1)} \in \{0, 1\}$ to indicate whether $R_k^{(1)}$ is a recurrence time belonging to an expansion segment

($S_k^{(1)} = 0$) or a recession segment ($S_k^{(1)} = 1$).

Computationally speaking, the number of partitions $\mathcal{N}(\mathcal{X}_n)$ associated to all thresholds

(ρ_1, ρ_2) is much less than 2^n so that the set of candidate partitions is relatively small. This

feature of the HFS algorithm facilitates computation of all the relevant change-points.

2.3 Determining the Optimal Thresholds (ρ_1, ρ_2)

In Section 2.1 we showed that for a given event, the code-sequence $\mathcal{C}_n^{(0)}$ generates event recurrence times $\mathcal{R}_{n(1)}$ characterized by a geometric distribution $G(R_k; \lambda)$ under the null that the state-space trajectory is unique. The code sequences $\mathcal{C}_{n(j)}$ for $j = 1, 2$ are generated with the thresholds (ρ_1, ρ_2) and result in the partition $\mathcal{N}(\mathcal{X}_n)$ where the event recurrence time in the segments $[P_i, T_i)$ is possibly different than for the segments $[T_i, P_i)$ for $i = 1, \dots, m$ and we use $S_k^{(1)} \in \{0, 1\}$ to denote whether $R_k^{(1)}$ is a recurrence time observed in an expansion segment or a recession segment.

It is worth summarizing some of the features implied by our set-up:

A.1 [No cross-over]: Each $R_k^{(1)}$ comes from a single segment.

A.2 [Independence]: The recurrence time $R_k^{(1)}$ is stochastically independent of $R_h^{(1)}$ for any $k \neq h$ in which k and h belong to different segments.

A.3 [Geometric Distribution]: $\Pr[R_k^{(1)} | S_k^{(1)} = s] = G(R_k; \lambda_s); s = 0, 1$ where the intensity parameters $\lambda_s, s = 0, 1$ are unknown.

Features A.1 and A.2 are quite mild. Feature A.3 requires that each regime segment be stationary but with limited serial correlation. Specifically, if \mathcal{X}_n is an exchangeable process in the sense of Chang et al. (2010) then one can use the maximum entropy principle under the i.i.d. of $R_k^{(1)}$.

Under features A.1-A.3 and the maximum entropy principle, then the large-sample distribution of $\mathcal{R}_{n(1)}$ is characterized by a geometric mixture with intensity parameters given

by

$$\lambda_s^M = \log \left\{ \frac{r_s}{r_s - 1} \right\}; s = 0, 1$$

where

$$\hat{r}_s = \frac{\sum_{k=1}^{n(1)} (1-s) R_k^{(1)} (1 - S_k^{(1)}) + s R_k^{(1)} S_k^{(1)}}{s \sum_{k=1}^{n(1)} S_k^{(1)} + (1-s) (n(1) - \sum_{k=1}^{n(1)} S_k^{(1)})}; s = 0, 1.$$

Thus the likelihood for the mixture is calculated, delivering intensity parameter estimates $\hat{\lambda}_0^M$ and $\hat{\lambda}_1^M$. Then, using a grid-search over different values for $\theta = (\rho_1, \rho_2)$, we minimize the noise-to-signal ratio given by

$$\hat{\theta} = \arg \min_{\theta} \left[\frac{|\hat{\lambda}_0^M - \hat{\lambda}_0^{\theta}|}{\hat{\lambda}_0^M} + \frac{|\hat{\lambda}_1^M - \hat{\lambda}_1^{\theta}|}{\hat{\lambda}_1^M} \right]$$

which implicitly generates intensity estimates $\hat{\lambda}_0^{\theta}, \hat{\lambda}_1^{\theta}$. We have found that this procedure works very well in practice, and is superior to the use of information criteria, such as AIC or SIC (see Hsieh, 2010).

3 Evaluating the HFS Chronology against the NBER

The NBER compiles a chronology of business cycle turning points for the U.S. –a natural benchmark with which to evaluate the HFS algorithm. Several features complicate a direct comparison between methods, however. First, the NBER uses several economic indicators to generate a single series of recession dates. For this reason, we begin by discussing methods to properly combine information from several indicators to produce a single chronology based on the HFS algorithm. These methods are based on network analysis (see e.g. Watts and Strogatz, 1998). Second, the true state of the economy is unobservable and therefore

conventional methods cannot be used to evaluate the accuracy of classification from each approach. We discuss each of these issues in turn.

3.1 Statistical Learning from Multivariate Data Through Networks

The NBER’s definition of recession (see the quote in the introduction) focuses on fluctuations in “economic activity,” a loose term that captures some combination of output and employment variables that is not made specific. However, the NBER does provide the list and the transformations of the indicators used in the deliberations.³ The appendix contains the complete list along with the sources. They are: the industrial production index (IP); real personal income less transfers (PI); payroll employment (PE); household employment (HE); real manufacturing and trade sales (MTS); real gross domestic product (GDP); and real gross domestic income (GDI). All the indicators are available at a monthly frequency although the latter two are interpolated from quarterly data following the procedures described by the NBER (and described in the appendix in more detail).

For all of these series, we apply the HFS algorithm to 100 times the year-on-year log difference, which is approximately the annual growth rate in percentage terms. For a discussion of the effects of transformations other than the year-on-year transformation used here, see Berge and Jordà (2010). In experiments not reported here, we found that the year-on-year transformation provided the best balance between smoothing noise while highlighting the signal.

³ These are available at: <http://www.nber.org/cycles/dec2008.html>

Next, we apply the HFS algorithm to each individual series and therefore generate seven segmentations for the U.S. Figure 2 presents the recession dates for each series against the NBER recession dates in the U.S. Several results deserve comment. Most recessions are bracketed by the segmentation found for the collection of series although individually, it is easy to see that there may be slight differences in the peaks and troughs implied by the segmentation and those from the NBER. This is not surprising: it is well known that employment tends to lag output when coming out of recessions, for example. There appear to be some false positives –for some series HFS identifies some phantom recessions– as well as some false negatives –some recessions that are not detected in some of the series. However, recall that a recession is an event that spreads throughout the economy and is felt in production, employment and income, *simultaneously* – it is not surprising to find these discrepancies. For this reason, the natural next step is to describe how this information can be combined to come up with the dating of expansions and recessions. This we do with network analysis.

We construct a network for each period in the sample using seven nodes for the U.S., one for each economic indicator. Next, for each country we connect any two nodes whenever the HFS segmentation of each of the two candidate indices selects a particular period t as being a recession. Given this network, one can calculate the “wiring ratio,” which is the proportion of wired pairs among all possible pairs. That is, let q_t denote the number of series considered (the seven indicators in this application) and p_t the number of series out of the q_t that at time t are classified as being in the recession regime. Then the wiring ratio

is $w_t = p_t(p_t - 1)/(q_t(q_t - 1))$. The time profile of the wiring ratio is shown at the bottom of Figure 2. It is easy to see that high values of the wiring ratio match the NBER dates rather well. Notice that for the first couple of decades (in a sample that begins in 1920), only industrial production data are available and hence the wiring ratio is zero by construction.

The wiring ratio combines the information from the seven indicators, but in the end we would like to obtain a binary indicator of recession dates. In order to obtain an optimal threshold for the wiring ratio (that would best sort the data into expansions and recession) we choose the threshold that would maximize the net correct classification skill (see Baker and Kramer, 2007 and Berge and Jordà, 2010), that is, choose the threshold c that maximizes:

$$U(c) = 2\pi \left(TP(c) - \frac{1}{2} \right) - 2(1 - \pi) \left(FP(c) - \frac{1}{2} \right)$$

where $\pi = P(S = 1)$, $TP(c) = P(w > c | S = 1)$ and $FP(c) = P(w > c | S = 0)$ where recall that S denotes the binary indicator of the true underlying state with $S \in \{0, 1\}$ and 1 denotes recession, and w denotes the wiring ratio. Figure 3 displays the optimal threshold calculated using a combination of the *incidence rate* –the ratio of series indicating recession relative to total– when the number of indicators available is less than three, and the wiring ratio thereafter. The optimal threshold turns out to be 0.0741 suggesting that when 3 or more indicators signal recession, it should be concluded that the economy is in recession.

As a way to compare the chronology of recessions generated by the NBER and the HFS algorithm, we construct an aggregate activity index using the first principal component of the seven economic indicators described by the NBER and listed in the appendix. Therefore, the left panel of Figure 4 displays this index against the HFS recession dates, whereas the right

panel does the same but for the NBER dates. Further, the dates of peaks and troughs are listed at the end of Table 1. Notice that there is a strong agreement (the same recessions are detected after 1968, which is when data for all indicators is available to build a comprehensive aggregate activity index), although HFS tends to generate slightly longer lasting recessions.

As a final check, we examined the classification ability of HFS against the NBER using the receiver operating characteristic (ROC) curve, specifically, the area under this curve (AUC), which we described above. The area under the ROC curve (AUC) can be interpreted here as the probability that an observation from the aggregate activity index drawn from the recession distribution has a smaller value than an observation drawn from the expansion distribution – call this probability δ (see Green and Swets, 1966). Intuitively, we expect recessions to be associated with periods in which economic activity is dim and expansions with periods where it is bright. The AUC is a statistic that takes on the value of 0.5 when a given classification produces no better than a random chance δ , and takes the value of 1 when separation between the two distributions is perfect. This statistic has a Gaussian large sample distribution (see Hsieh and Turnbull, 1996). We calculated the AUC statistic for HFS and for the NBER using the aggregate activity index. The AUC for HFS is $AUC = 0.86$ whereas for the NBER it is $AUC = 0.91$. Both are statistically different from the null of random chance at better than 99% confidence level but one cannot reject the null that they are equivalent in the statistical sense at a 95% confidence level. Of course, this is not meant to suggest that therefore one should replace the BCDC with a computer, but it is meant to give assurance that the HFS algorithm has good statistical properties and that it produces

a classification of recessions that is sensible when compared with the NBER’s benchmark.

4 International Business Cycles

Results from the previous section suggest that HFS is a reasonable approach to constructing a chronology of business cycle turning points –the metric that we introduce here performs similarly for both procedures. Moreover, the application of network analysis to refine the construction of this chronology by combining the signals from different indicators meshes well with the philosophy spelled out by the NBER that recessions ought to be visible in several sectors of the economy *simultaneously*. In this section we present several results of interest. First, we construct a chronology of peaks and troughs for 22 OECD countries: Australia, Austria, Belgium, Canada, Switzerland, Chile, Denmark, Germany, Spain, Finland, France, Greece, Ireland, Italy, Japan, South Korea, Netherlands, Norway, New Zealand, Portugal, Sweden, and the U.K. We do this by generating HFS chronologies for three indicators: monthly industrial production (IP) indices and employment and, linearly interpolated quarterly real gross domestic product (GDP). The source for our data set is the IMF and the OECD, with transformations and samples available discussed in the appendix. Although each series is often available for samples of differing lengths, we use the following approach. For periods when only one series is available, we use the HFS chronology directly. For periods when two series are available, we use the rule that when one or two indicators signal recession, we call that period a recession. For periods when all three indicators are available, we use the rule that when two or more indicators signal recession, that period is classified as

a recession. The results of this exercise are reported in Table 1. For completeness, we have added the HFS-based chronology for the U.S., calculated in the previous section along with the NBER's official chronology.

The dates reported in Table 1 should, in and of themselves, prove of value to researchers as they provide a benchmark chronology for many countries for which there is no official dating of business cycle turning points. For this reason, we have chosen to report the dates obtained by our procedures without any further adjustment. We mention this because we have encountered a few instances in which two contiguous recessions are separated by only a few months and could therefore be conceivably considered as a single, longer running recession. For example, in the U.K. we date a recession that starts May 1973 and ends January 1974, with the next recession starting August 1974 and ending July 1975. Therefore, one could have merged these two recessions into one that started May 1973 and ended July 1975. To provide more clarity on this issue, we provide a plot of the yearly growth rate of GDP against the HFS chronologies in Figure 5 for the UK and in Figure 6 for the remaining OECD countries analyzed. Thus, Figure 5 shows the temporary jump in GDP between recessions. Another good example is the recession that began in February 1988 and ended in October 1989, which was then followed by the recession that began in May 1990 and ended April 1991. It is clear from the graph that the period between recessions represents a brief respite in the steep decline of GDP that began May 1990 and ended with the trough of 1991.

Figures 7, 8 and 9 display the HFS dates selected for each country per indicator: industrial production (IP); real gross domestic product (GDP); and employment. They are provided

for completeness and to show that for some indicators, the available sample is rather short. However, for many countries there exist relatively long samples for all three indicators. Thus, Figures 7-9 document the length of the samples available for each indicator and country HFS chronologies.

The last exercise in this section takes us to our goal of constructing a global business cycle indicator since our sample of countries accounts for about 60% of world output in 2009. We use the following approach to identify events that had international resonance, for which the IMF might have been particularly interested. We constructed the wiring ratio (defined above as consisting of counting the pairs of countries simultaneously experiencing a recession in a given period normalized by the total number of possible pairs) using the HFS dates of peaks and troughs presented in Table 1. We identify global recessions with the rule-of-thumb requiring nine or more countries to be involved. Nine countries corresponds to a wiring ratio of about 0.14, which is sufficient to focus on a few choice events. The value is a slightly higher value than we used for the U.S. (0.0741), but not by a large amount. This is, of course, arbitrary but it is meant to serve as an illustration. Also, we note that the wiring ratio is not weighted by the size of the economy, which is another modification that could be considered. The graph of the wiring ratio and the threshold 0.14 are presented in Figure 10 and result in 7 cyclical events tabulated in Table 2.

We comment on some of these global recession dates starting from the most recent. It is clear that the recent financial crisis had a widespread effect, engulfing the majority of countries in our sample. In the U.S. the beginning of the recession was dated December

2007 and the end June 2009. Our global recession indicator suggests that by February 2008, the recession that had started in the U.S. swept a large majority of developed economies in its wake. Some countries whose financial system withstood the near collapse of the American financial system after the fall of Lehmann Brothers began to recover so that by April 2009, only a few countries remained in recession (such as the U.S.). The next interesting event has to do with the recession that we date as starting in August 2000. This is interesting because it predates the onset of the recession in the U.S. (dated March 2001) by half a year, and ends in May of 2001 – half a year earlier than it did in the U.S. although one has to factor in the attacks of September 11 of that year in New York. The next event is short-lived –one month– and it is probably best classified as a non-recession but it is kept as an “honesty” check of the mechanical rule that we have self-imposed. It is clear that a slightly lower threshold would have produced a more regular looking event. After that we have another global event dated to have started November 1973 and ending in January 1975, which is easily identifiable with the first oil crisis that began in October 1973 with OPEC’s oil embargo in response to the U.S.’ decision to supply the Israeli military during the Yom Kippur (or Arab-Israeli) War. The May-August 1964 event pre-dates the American escalation in the Vietnam war after the August 2, 1964 events in the gulf of Tonkin, but involves a smaller number of countries. The May 1960 to July 1961 again seems to have been led by the U.S. where the recession started in April 1960 but ended in February 1961. We remark that the last global event detected to take place January to December 1957 is based on a smaller number of countries for which data was available and hence may not be as “global” an event as the preceding ones.

Of course, there are many interesting extensions that are possible using the tools that we have introduced in the paper. For example, one may be interested in regional recessions (say for countries in the European monetary union), or by grouping the U.S. and its major trading partners, just to name two that would be interesting. And of course, one may want to weigh the wiring ratio by the size of the economy (for the first example) or by the size of its exports to the U.S. (in the second example). There are extensions that are left for a different paper.

5 Conclusions

The traditional approach in modeling time series with switches in regimes consists in specifying the likelihood function of a model that allows for such dynamic transitions and that characterizes the dynamic behavior of the data within each regime. In other words, one is required to characterize all aspects of the dynamic evolution of the stochastic process under consideration. This paper uses an alternative philosophy to this problem: it takes a decidedly non-parametric approach so as to avoid the constraints that a fully specified model imposes on the researcher's ability to properly classify the data. Specifically, our approach simply looks at the relative incidence of tail events and then characterizes the duration distribution between such events. Under the null that observations are drawn from a unique data generating density, the entropy principle (Jaynes, 1957a, b) allows us to characterize the null duration distribution. But under the alternative, the duration between tail events under one regime will be considerably different than under the other, and such differences

can be formally tested. This principle uses basic ideas that have been around at least since Poincaré (1890) and is translated into practice with a simple algorithmic procedure: the hierarchical factor segmentation method of Fushing, Hwang, Lee, Lang and Horng (2006). The HFS algorithm can be thought of as straddling the statistics and the computer science literatures in that it can also be interpreted as a decoding algorithm. In fact, HFS was originally developed in biology but has since been applied to other disciplines.

HFS is one element of a broader strategy for finding and properly classifying cyclical phenomena in economic time series. Recessions are characterized by a slowdown in economic activity felt across different sectors that exhibit varying resilience to the economic downturn, and therefore a different chronology of peaks and troughs. Therefore, another important element of the strategy presented here consists in operationalizing the NBER's stated principle of combining information garnered from a variety of economic indicators. In following a more comprehensive approach to this problem than is conventional (say, with common dimension reduction techniques based on principal component analysis or with factor models), we consider methods borrowed from computer science and network analysis in particular. Such methods are particularly well suited for this problem as they focus on the temporal concurrence of recession signals rather than on their relative correlation patterns.

Our methods produce a business cycle chronology for the U.S. that rivals that produced by the NBER. Armed with this result, we embark in an analysis of a long list of OECD countries (22 plus the U.S.), something that would be far more arduous to do with existing methods. We hope that these 22 chronologies will at least serve as the basis for a bench-

mark registry of international business cycles. Undoubtedly, individual chronologies can be improved upon on a case-by-case basis by extending the menu of economic indicators considered for each country, and we hope that others will follow such an approach. In continuing work by us (Fushing, Chen, Berge and Jordà, 2010), we do precisely this for Taiwan where an official business cycle chronology is published by the Council of Economic Planning and Development (CEPD) and can be used as a basis for comparison.

A natural extension of our methods (based on the ability to generate business cycle chronologies for a long list of countries that represent about 60% of world output in 2009) is to document episodes of global downturns, where a considerable portion of countries experience a slowdown simultaneously. Such an analysis uses the same network techniques that we used for the U.S., but across countries rather than across indicators. Several extensions are natural to the analysis that we present here, such as weighing by the size of the economy, focusing on trading blocks, regional economic unions, and so on. This type of analysis has the advantage of extracting synchronicity in cyclical activity rather than looking at correlations across economic indicators that may not be of the same scale and hence difficult to associate (such as the growth rates of a developed and an emerging country). We recognize that several improvements to our analysis can be done and many of these improvements will be driven by the specific nature of the application considered, but we hope to have provided the fundamental tools for such an exploration.

6 Appendix: Data Sources

6.1 US Data

This is a summary of the economic indicators, transformations and data sources provided in the appendix of the December 11, 2008 press release of the Business Cycle Dating Committee of the National Bureau of Economic Analysis and available from their website (www.nber.org).

<i>Indicator</i>	<i>Sample Available</i>	<i>Source and Method</i>
Industrial Production	1919:1 - 2009:12	FRB index B50001
Real Personal Income less transfers	1959:1 - 2009:12	BEA Table 2.6, line 1 less line 14, both deflated by a monthly interpolation (see below) of BEA Table 1.1.9 line 1
Payroll Employment	1939:1 - 2009:12	BLS Series CES0000000001
Household Employment	1948:1 - 2009:12	BLS Series LNS12000000
Real Manufacturing and Trade Sales	1967:1 - 2009:12	BEA Table 2BU, line 1
Real Gross Domestic Product	1947:I - 2010:I	BEA Table 1.1.6, line 1
Real Gross Domestic Income	1947:I - 2010:I	BEA Table 1.10, line 1, divided by BEA Table 1.1.9, line 1

Websites:

- Federal Reserve Board industrial production index:

www.federalreserve.gov/releases/g17/iphist/iphist_sa.txt

- Bureau of Economic Analysis, U.S. Department of Commerce, all but sales:

www.bea.gov/national/nipaweb/SelectTable.asp?Selected=N

Sales: www.bea.gov/national/nipaweb/nipa_underlying/SelectTable.asp

- BLS payroll survey: <http://data.bls.gov/cgi-bin/surveymost?ce>

- BLS household survey: <http://data.bls.gov/cgi-bin/surveym>

6.2 International Data

Three economic indicators are used for each country: Real Gross Domestic Product, Industrial Production and Employment. The primary data source for all three series was the IMF's International Financial Statistics database. Employment data were supplemented with data from the OECD *OECD.Stat* database when data are missing or to provide a longer series. We use monthly indicators when available, then transform this indicator into a year-on-year growth rate, so that the indicator becomes $x_t = 1200 * \Delta_{12} \log X_t$. We interpolate quarterly observations into monthly observations using the same methodology that the Business Cycle Dating Committee of the NBER uses to interpolate Gross Domestic Product and Gross Domestic Income. In particular, the value of the index in the first month of the quarter is one third of the past quarter's value plus two-thirds of the current quarter's value. In the second month, it is the quarter's value. In the third month, it is two-thirds of the quarter's

value plus one third of the next quarter's value. The interpolated monthly observations are again transformed into growth rates.

<i>Country Isocode</i>	<i>GDP</i>	Industrial Production	<i>Employment</i>
AUS	1959:III - 2010:I	1957:III - 2010:I	1964:2 - 2009:12*
AUT	1964:I - 2010:I	1988:1 - 2009:12	1957:1 - 2009:12
BEL	1981:I - 2010:I	1957:1 - 2009:12	1999:I - 2009:IV
CAN	1957:I - 2010:I	1995:1 - 2009:12	1956:1 - 2009:12*
CHE	1966:I - 2009:IV	1960:I - 2009:IV	1976:I - 2010:I
CHL	1980:I - 2009:IV	1958:1 - 2009:12	1986:1 - 2009:12
DEN	1977:I - 2009:IV	1957:1 - 2009:12	1995:I - 2010:I
DEU	1960:I - 2009:IV	1958:1 - 2009:12	1962:I - 1991:I*; 1991:1 - 2009:12
ESP	1970:I - 2009:IV	1961:1 - 2009:12	1964:II - 2010:I
FIN	1970:I - 2009:IV	1957:1 - 2009:12	1964:1 - 2009:12*
FRA	1970:I - 2010:I	1957:1 - 2009:12	2003:I - 2008:IV

* Indicates data from the OECD database

<i>Country Isocode</i>	<i>GDP</i>	Industrial Production	<i>Employment</i>
GRC	1975:I - 1991:I; 2001:I - 2010:I	1995:1 - 2009:12	1998:I - 2009:IV
IRE	1997:I - 2009:IV	1975:7 - 2009:12	1998:I - 2010:I
ITA	1980:I - 2009:I	1957:1 - 2009:12	1959:I - 2009:IV
JPN	1957:I - 2010:I	1957:1 - 2009:12	1955:1 - 2009:12*
KOR	1960:I - 2010:I	1980:1 - 2009:12	1982:7 - 2009:12*
NLD	1977:I - 2009:IV	1957:1 - 2009:12	2001:1 - 2009:12*
NOR	1966:I - 2010:I	1957:1 - 2009:12	1972:I - 2010:IV
NZD	1982:II - 2009:IV	1987:2 - 2009:VI	1985:IV - 2010:1
PRT	1977:I - 2009:IV	1957:1 - 2009:12	1983:II - 2010:I
SWE	1986:I - 2010:I	1957:1 - 2009:12	1965:5 - 2009:12*
UK	1957:I - 2009:IV	1957:1 - 2009:12	1956:I - 1992:II; 1992:4 - 2009:12*

* Indicates data from the OECD database

References

- Ambler, Steve, Emanuela Cardia and Christian Zimmermann. 2004. International business cycles: What are the facts? *Journal of Monetary Economics*, 51(2): 257-276.
- Baker, Stuart G. and Barnett S. Cramer. 2007. Pierce, Youden, and Receiver Operating Characteristics Curves. *The American Statistician*, 61(4): 343-346.
- Berge, Travis J. and Òscar Jordà. 2010. Evaluating the Classification of Economic Activity into Expansions and Recessions. *American Economic Journal: Macroeconomics*, forthcoming.

- Burns, Arthur F. and Wesley C. Mitchell. 1946. *Measuring Business Cycles*. NBER Book Series Studies in Business Cycles N. 2. New York: NBER.
- Fushing Hsieh and Bruce W. Turnbull. 1996. Nonparametric and Semiparametric Estimation of the Receiver Operating Characteristics Curve. *Annals of Statistics*, 24: 25-40.
- Fushing Hsieh, Shu-Chun Chen and Katherine Pollard. 2009. A Nearly Exhaustive Search for CpG Islands on Whole Chromosomes. *International Journal of Biostatistics*, forthcoming.
- Fushing Hsieh. 2010. Discovering stock dynamics through multidimensional volatility-phases. U. C. Davis, mimeo.
- Fushing Hsieh, Chii-Ruey Hwang, T-K Lee, Yen-Chiu Lan and Shwu-Bin Horng. 2006. Exploring and reassembling patterns in female bean weevil's cognitive processing networks. *Journal of Theoretical Biology*, 238(4): 805-816.
- Fushing Hsieh, Shu-Chun Chen and Chii-Ruey Hwang. 2010a. Non-parametric Decoding on Discrete Time Series and Its Applications in Bioinformatics. *Statistics in Biosciences*, forthcoming.
- Fushing Hsieh, Shu-Chun Chen and Chii-Ruey Hwang. 2010b. Discovering stock dynamics through multidimensional volatility-phases. *Quantitative Finance*. doi:10.1080/14697681003743040.
- Geman, Stuart and Kevin Kochanek. 2001. Dynamic programming and the graphical representation of error-correcting codes. *IEEE Transactions on Information Theory* 47(2): 549-568.
- Green, David M. and John A. Swets. 1966. *Signal Detection Theory and Psychophysics*. Peninsula Publishing: Los Altos, California.
- Hamilton, James D., 1989. A New Approach to the Economic Analysis of Nonstationary Time Series Subject to Changes in Regime. *Econometrica*, 57(2): 357-384.
- Hamilton, James D., 2005. What's real about the business cycle?, Review, Federal Reserve Bank of St. Louis, issue July, pages 435-452.
- Hastie, Trevor, Robert Tibshirani and Jerome Friedman. 2009. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer: New York.
- Jaynes, Edwin T. 1957a. Information Theory and Statistical Mechanics I. *Physical Review*, 106: 620.
- Jaynes, Edwin T. 1957b. Information Theory and Statistical Mechanics II. *Physical Review*, 108: 171.

- Kose, M. Ayhan, Christopher Otrok and Eswar S. Prasad. 2008. Global Business Cycles: Convergence or Decoupling? NBER Working Papers 14292, National Bureau of Economic Research, Inc.
- Kose, M. Ayhan, Christopher Otrok and Charles H. Whiteman. 2008. Understanding the evolution of world business cycles. *Journal of International Economics*, 75(1): 110-130.
- Kose, M. Ayhan, Eswar S. Prasad and Marco E. Terrones. 2003. How Does Globalization Affect the Synchronization of Business Cycles? *American Economic Review*, 93(2): 57-62.
- Mitchell, Wesley C. 1927. *Business Cycles: The Problem and its Setting*. National Bureau of Economic Research: New York.
- Mitchell, Wesley C. 1951. *What Happens During Business Cycles: A Progress Report*. National Bureau of Economic Research: New York.
- Poincaré, J. Henri. 1890. Sur les Equations aux Dérivées Partielles de la Physique Mathématique. *American Journal of Mathematics*, 12(3): 211-294.
- Stock, James H. and Mark W. Watson. 2005. Understanding Changes In International Business Cycle Dynamics. *Journal of the European Economic Association*, 3(5): 968-1006.
- Stock, James H. 1987. Measuring Business Cycle Time. *Journal of Political Economy*, 95(6): 1240-61.
- Stock, James H. 1988. Estimating continuous-time processes subject to time-deformation: an application to postwar U.S. GNP. *Journal of the American Statistical Association*, 83(401): 77-85.
- Watts, Duncan J. and Stephen H. Strogatz. 1998. Collective dynamics of ‘small-world’ networks. *Nature*, 393(4): 440-442.

Table 1 – Country-specific dates of Peaks and Troughs detected with HFS

Australia		Austria		Belgium		Canada	
Aug-1958 to Dec-2009		Jan-1958 to Dec-2009		Jan-1955 to Dec-2009		Jan-1956 to Dec-2009	
<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>
Aug-1960	Jul-1961	Nov-1960	Oct-1961	Feb-1960	Sep-1961	Jun-1957	Jan-1958
Aug-1962	Apr-1963	Nov-1963	Nov-1965	Jun-1963	Jun-1965	Oct-1959	Jan-1961
Aug-1964	Jan-1965	Feb-1966	Jul-1966	May-1969	Jun-1971	Feb-1962	Jul-1963
Aug-1965	Jan-1966	Mar-1968	Aug-1969	May-1972	Nov-1973	Nov-1965	Jan-1967
Aug-1971	Jan-1972	Feb-1972	Jul-1972	Mar-1974	Jul-1975	Jul-1967	Jun-1968
Feb-1974	Jan-1975	Aug-1974	Apr-1975	Feb-1976	Aug-1977	Apr-1969	Mar-1970
Jun-1977	Oct-1977	Apr-1977	Apr-1978	Jul-1981	Nov-1982	Jun-1973	May-1974
Oct-1981	Oct-1982	Nov-1979	Jul-1981	Aug-1983	Jun-1985	Aug-1974	Mar-1975
Nov-1985	Jul-1986	Nov-1983	Nov-1985	Feb-1988	Jul-1988	May-1976	Apr-1977
Feb-1990	Jul-1990	May-1992	Jan-1993	Dec-1988	Jun-1990	May-1981	Jul-1982
Aug-1994	Jan-1995	Jan-1995	Jan-1996	Aug-1992	Jan-1993	Nov-1985	Oct-1986
May-1996	Oct-1996	Dec-1999	Apr-2000	Feb-1995	Jul-1995	Feb-1989	Feb-1991
Aug-2000	Jan-2001	Nov-2000	Feb-2001	Nov-1997	Apr-1998	Jan-1995	Jan-1996
Aug-2008	Jan-2009	May-2008	Jun-2009	Nov-1999	Mar-2001	Aug-2000	Jul-2001
				Aug-2008	Jan-2009	Nov-2007	Apr-2009
Switzerland		Chile		Denmark		Germany	
Jan-1955 to Dec-2009		Jan-1959 to Dec-2009		Feb-1958 to Dec-2009		Jan-1959 to Dec-2009	
<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>
May-1964	Apr-1965	Feb-1959	Sep-1960	Aug-1961	Jan-1962	Oct-1959	Oct-1961
Aug-1969	Apr-1971	Nov-1961	Jul-1964	Aug-1969	Jan-1970	Aug-1965	Sep-1965
Nov-1973	Apr-1975	Aug-1967	Jul-1969	Nov-1973	Jan-1975	Feb-1966	Apr-1967
May-1980	Jan-1981	Apr-1981	Sep-1982	Feb-1979	Jul-1979	Feb-1969	Jul-1969
Aug-1981	Jul-1982	May-1983	Jul-1987	Feb-1980	Jan-1981	Mar-1970	Jan-1972
Feb-1990	Apr-1990	Feb-1990	Jul-1990	May-1986	Jan-1987	Aug-1973	Jan-1975
Feb-1995	Oct-1995	Sep-1993	Oct-1993	Feb-1988	Jul-1988	May-1976	Apr-1977
Aug-2000	Jan-2002	Nov-1995	Mar-1996	Feb-1989	Jan-1990	Feb-1984	Jan-1985
Nov-2007	Apr-2009	Nov-1998	Apr-1999	Aug-1990	Jan-1991	Feb-1991	Jan-1992
		Oct-2008	Apr-2009	Nov-1992	Jul-1993	Apr-2008	Jan-2009
				May-1994	Jan-1996		
				Nov-1997	Apr-1998		
				Nov-2000	Apr-2001		
				Aug-2001	Oct-2001		
				Aug-2004	Jul-2005		
				Aug-2008	Apr-2009		

Table 1 – Country-specific dates of Peaks and Troughs detected with HFS (cont.)

Spain		Finland		France		Greece	
Jan-1962 to Dec-2009		Jan-1958 to Dec-2009		Jan-1958 to Dec-2009		Feb-1976 to Dec-2009	
<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>
Jul-1962	May-1963	Apr-1960	Nov-1960	Mar-1964	Jul-1964	May-1979	Oct-1979
Feb-1964	Jul-1965	Jan-1962	Jun-1962	May-1973	Jan-1975	Nov-1989	Oct-1990
Oct-1966	May-1968	Mar-1964	Nov-1964	Sep-1976	Oct-1977	Apr-1999	Oct-1999
Feb-1974	Jul-1975	Jan-1967	May-1967	May-1982	Apr-1983	Nov-2000	Nov-2001
Feb-1977	Oct-1977	Sep-1985	Apr-1986	Aug-1984	Jan-1985	Aug-2003	Jan-2004
Aug-1991	Oct-1991	Feb-1990	Mar-1990	Oct-1985	Dec-1986	Aug-2006	Jan-2007
Feb-1992	Jan-1993	Jul-1990	Sep-1991	Mar-1988	Sep-1988	Jul-2007	May-2009
Mar-2000	Nov-2000	Feb-2001	Jul-2001	Feb-1992	Jan-1993		
Aug-2007	Apr-2009	May-2008	Dec-2008	May-1995	Apr-1996		
				Jun-1998	Dec-1998		
				Nov-1999	Apr-2002		
				Nov-2007	Jan-2009		
Ireland		Italy		Japan		South Korea	
July-1976 to Dec-2009		Jan-1958 to Dec-2009		Jan-1956 to Dec-2009		Feb-1961 to Dec-2009	
<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>
May-1979	Oct-1979	May-1961	Apr-1962	Sep-1957	Jan-1958	Nov-1961	Apr-1962
Nov-1989	Oct-1990	Apr-1963	Jan-1965	Jan-1962	Jun-1962	May-1965	Oct-1965
Apr-1999	Oct-1999	Nov-1967	Oct-1968	Jun-1964	Apr-1965	May-1966	Oct-1966
Nov-2000	Nov-2001	Jan-1971	Apr-1972	Mar-1970	Apr-1971	Feb-1974	Jul-1974
Aug-2003	Jan-2004	Jul-1973	Apr-1975	Oct-1973	Jan-1974	May-1978	Oct-1978
Aug-2006	Jan-2007	Dec-1976	Nov-1977	Sep-1980	Oct-1980	Feb-1979	Apr-1980
Jul-2007	May-2009	Nov-1991	Jan-1993	May-1991	Oct-1992	Feb-1984	Feb-1984
		Apr-2008	Feb-2009	Feb-1997	Jul-1998	May-1987	Apr-1988
				Nov-2000	Oct-2001	Feb-1992	Jul-1992
				Feb-2008	Jan-2009	Feb-1994	Feb-1995
						Aug-1997	May-1998
						Aug-2000	Jan-2001
						Apr-2002	Jun-2003
						May-2004	Oct-2004
						Apr-2008	Dec-2008

Table 1 – Country-specific dates of Peaks and Troughs detected with HFS (cont.)

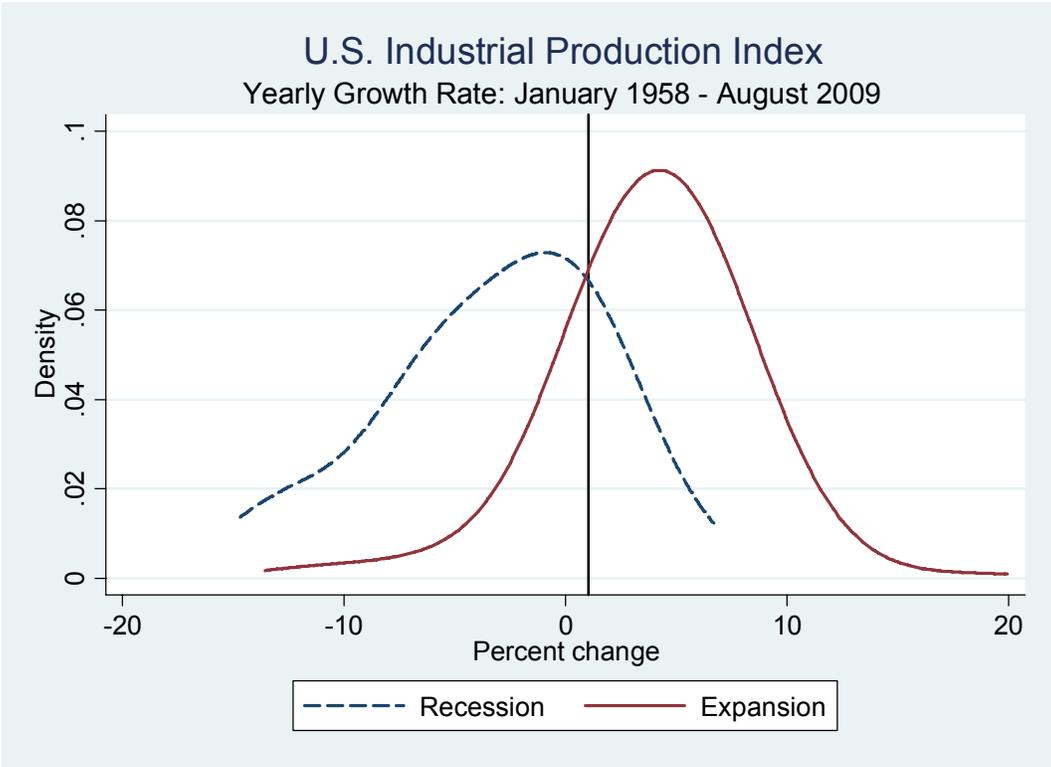
Netherlands		Norway		New Zealand		Portugal	
Jan-1958 to Dec-2009		Oct-1957-Dec-2009		May-1985 to Dec-2009		Jan-1958 to Dec-2009	
<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>
Feb-1978	Jan-1979	Nov-1967	Apr-1968	Feb-1984	Jul-1984	Jan-1959	Jul-1959
Feb-1980	Jan-1983	May-1970	Oct-1970	Nov-1987	Apr-1988	Jan-1961	Jul-1961
Jun-1984	Jul-1985	Nov-1972	Jan-1973	May-1995	Jan-1996	Oct-1963	Jul-1964
Aug-1986	Apr-1987	Feb-1980	Jul-1980	May-1997	Apr-1998	Dec-1965	Aug-1966
Feb-1991	Aug-1991	Aug-1987	Dec-1987	Nov-1999	Oct-2000	Dec-1968	Sep-1969
Jan-1997	Sep-1997	May-1991	Oct-1991	Nov-2004	Jan-2005	Dec-1973	Feb-1975
Nov-1999	Dec-2001	Aug-2002	Apr-2003	Nov-2007	Apr-2008	Feb-1978	Apr-1978
Jul-2008	Apr-2009	May-2008	Oct-2008			Dec-1978	Jun-1979
		Mar-2009	Jul-2009			May-1981	Apr-1982
						Nov-1982	Apr-1983
						Feb-1985	Jul-1985
						May-1988	Jul-1988
						Feb-1992	Apr-1992
						Aug-2008	Jan-2009
Sweden		UK		US (HFS)		US(NBER)	
Jan-1958 to Dec-2009		Feb-1957 to Dec-2009		Jan-1958 to Dec-2009		Jan-1958 to Dec-2009	
<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>	<i>Peak</i>	<i>Trough</i>
Dec-1959	Jul-1960	May-1957	Jan-1958	Mar-1962	Nov-1962	Apr-1960	Feb-1961
Sep-1965	Mar-1966	May-1960	Oct-1961	Feb-1966	Apr-1967		
Nov-1969	May-1970	Feb-1964	Jul-1965	Oct-1969	Feb-1970	Dec-1969	Nov-1970
Jun-1976	Jul-1978	May-1973	Jan-1974	Jul-1973	Jan-1975	Nov-1973	Mar-1975
Apr-1980	Dec-1981	Aug-1974	Jul-1975	Mar-1979	Oct-1979	Jan-1980	Jul-1980
Nov-1983	Jun-1984	Nov-1979	May-1980	Jul-1981	Sep-1982	Jul-1981	Nov-1982
Dec-1990	Dec-1991	May-1984	Jul-1984	May-1984	Apr-1985		
Aug-1994	Aug-1994	May-1985	Jan-1986	May-1990	Apr-1991	Jul-1990	Mar-1991
Feb-1996	Mar-1997	Feb-1988	Oct-1989	May-2000	May-2001	Mar-2001	Nov-2001
Feb-2001	Jan-2002	May-1990	Apr-1991	Sep-2007	May-2009	Dec-2007	Jun-2009
Aug-2006	Nov-2007	Mar-1998	Jul-1998				
Jul-2008	Dec-2009	May-2000	Jan-2001				
		Dec-2002	Dec-2002				
		Feb-2004	Jan-2005				
		Aug-2007	Oct-2007				
		Mar-2008	Apr-2009				

Table 2 – Dates of Peaks and Troughs for the Global Business Cycle

Peak	Trough
January 1957	December 1957
May 1960	July 1961
November 1973	January 1975
April 1980	April 1980
August 2000	May 2001
February 2008	April 2009

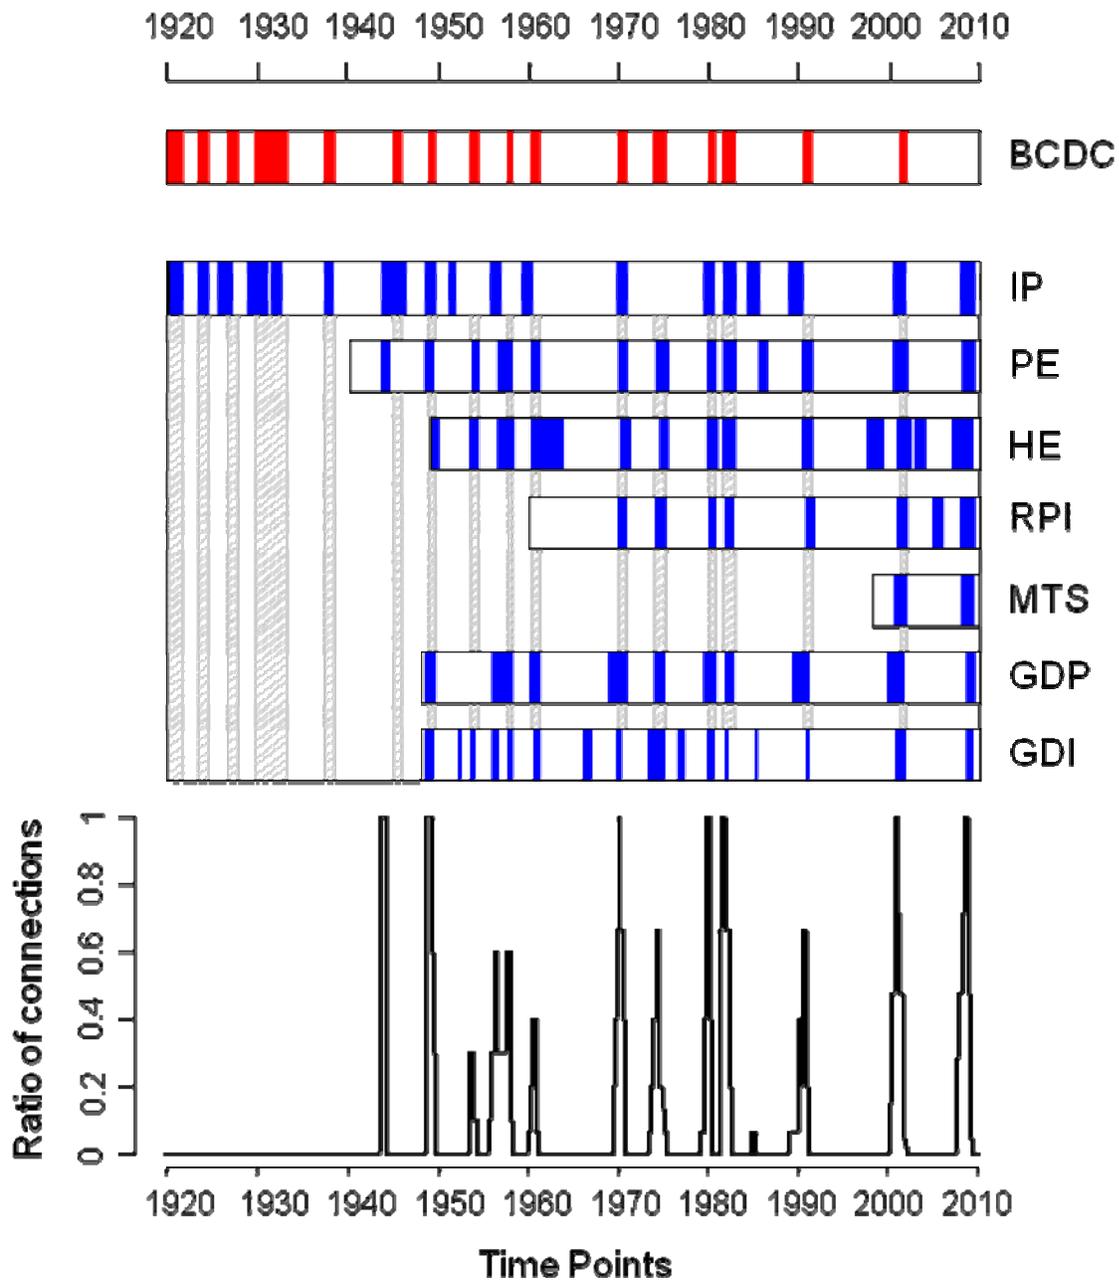
Notes: these dates correspond to episodes in which the wiring ratio exceeded 0.14.

Figure 1 – Choosing the event for the HFS Algorithm



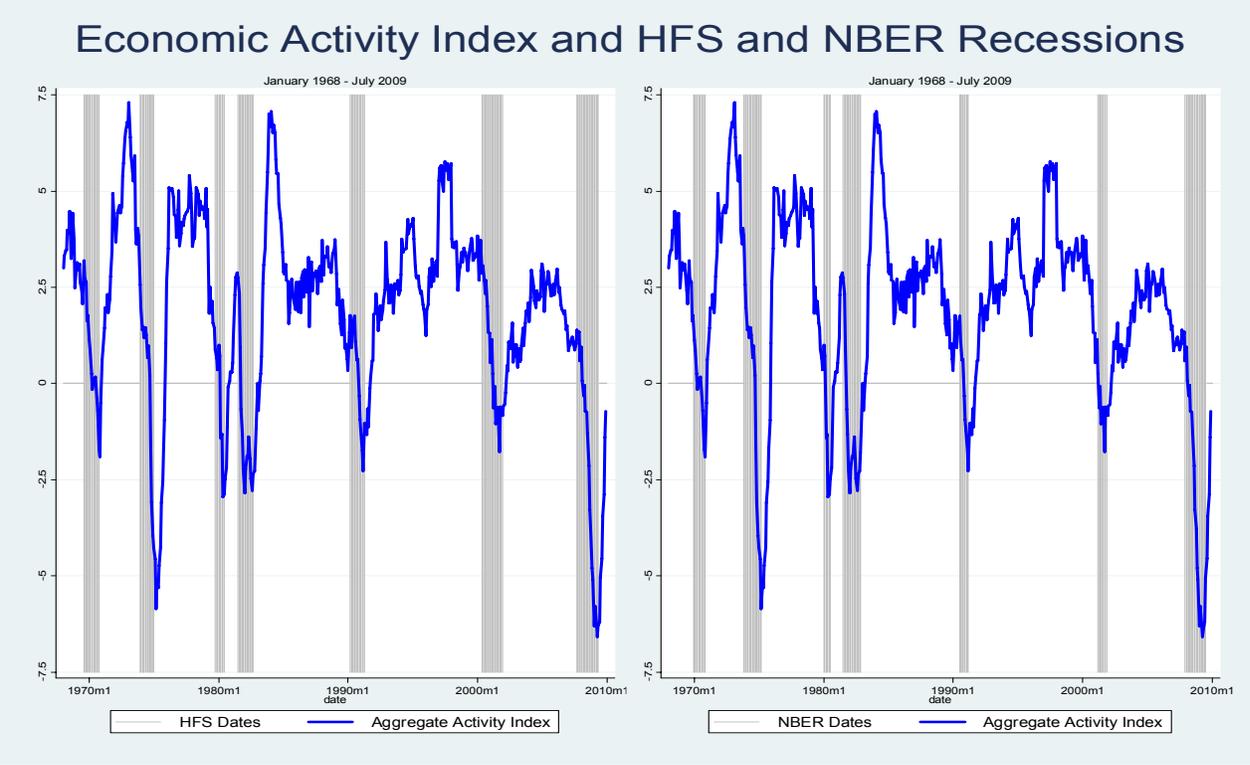
Notes: smooth kernel density plots of the empirical distribution of the U.S. industrial Production Index for expansion and recession periods as defined by the NBER. The vertical line indicates the 30th percentile of the distribution of the IP index. Using such a cut-off as a rule of thumb to determine when the economy is in expansion and when in recession generates a true positive rate of 78% and a false positive rate of 20%.

Figure 2 – Summary of individual recession dates calculated by HFS against the NBER’s BCDC.



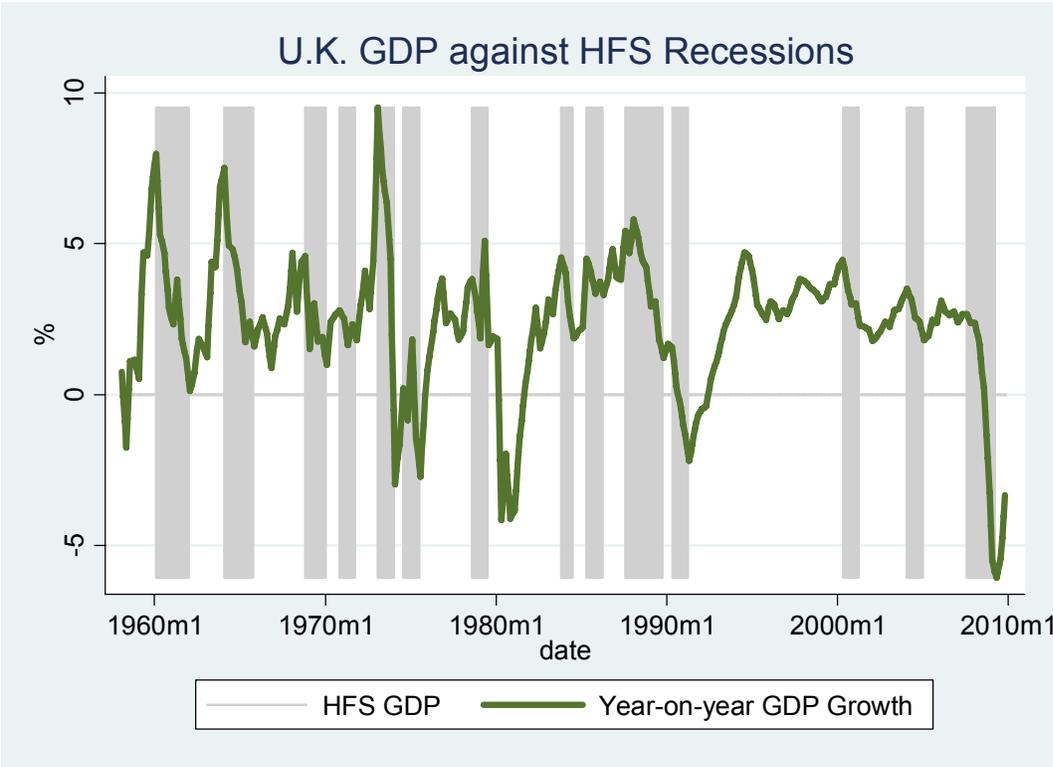
Notes: The top graph indicates for each series, periods where HFS classified the indicator as being in “recession.” The top row labeled “BCDC” is for reference and indicates the recession dates provided by the NBER. The bottom graph displays the wiring ratio associated with the ratio of pairs of indicators simultaneously classified as being in recession at a given point in time, relative to the total number of possible pairs.

Figure 4 – HFS versus NBER recessions against economic activity



Notes: The aggregate activity index is constructed as the first principal component of the seven economic indicators used by the NBER to determine peak and trough dates. The left panel displays the HFS recessions shaded in grey whereas the right panel displays NBER recessions.

Figure 5 – U.K. yearly growth rate of GDP against HFS recessions



Notes: yearly percentage growth in U.K. GDP. Shaded areas are recessions detected using the HFS algorithm.

Figure 6 – Yearly growth rates of GDP against HPS recessions, all remaining OECD countries

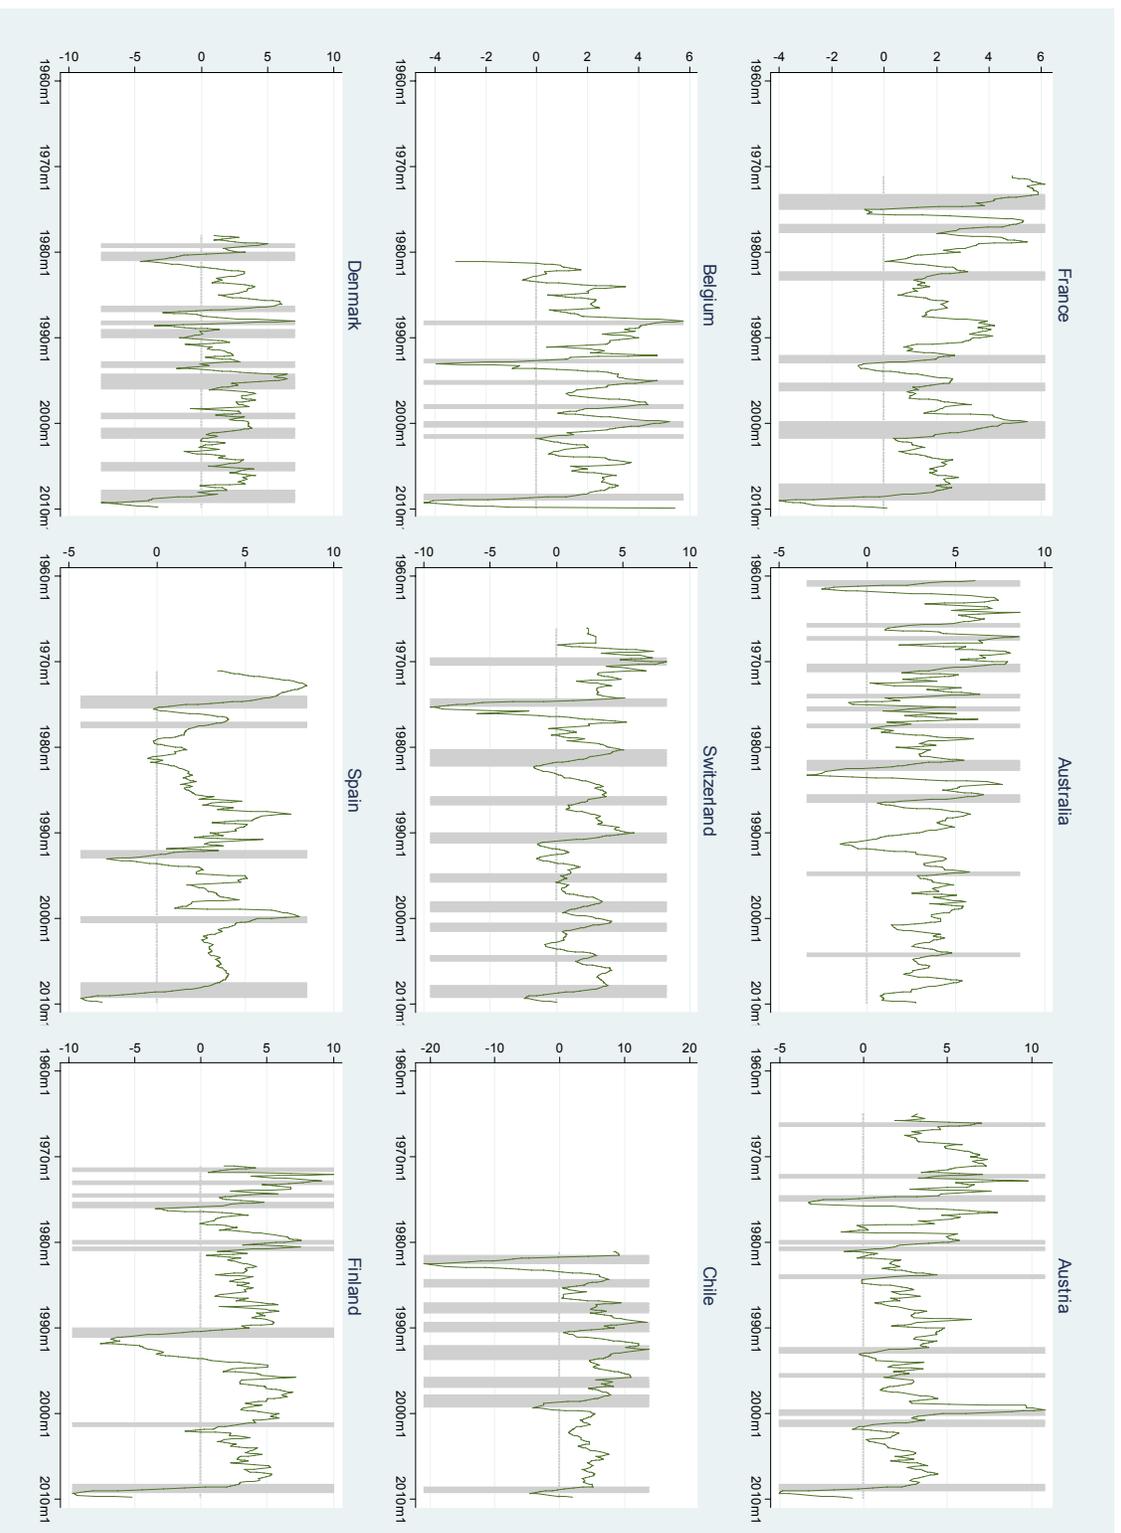


Figure 6 (cont.) – Yearly growth rates of GDP against HFS recessions, all remaining OECD countries

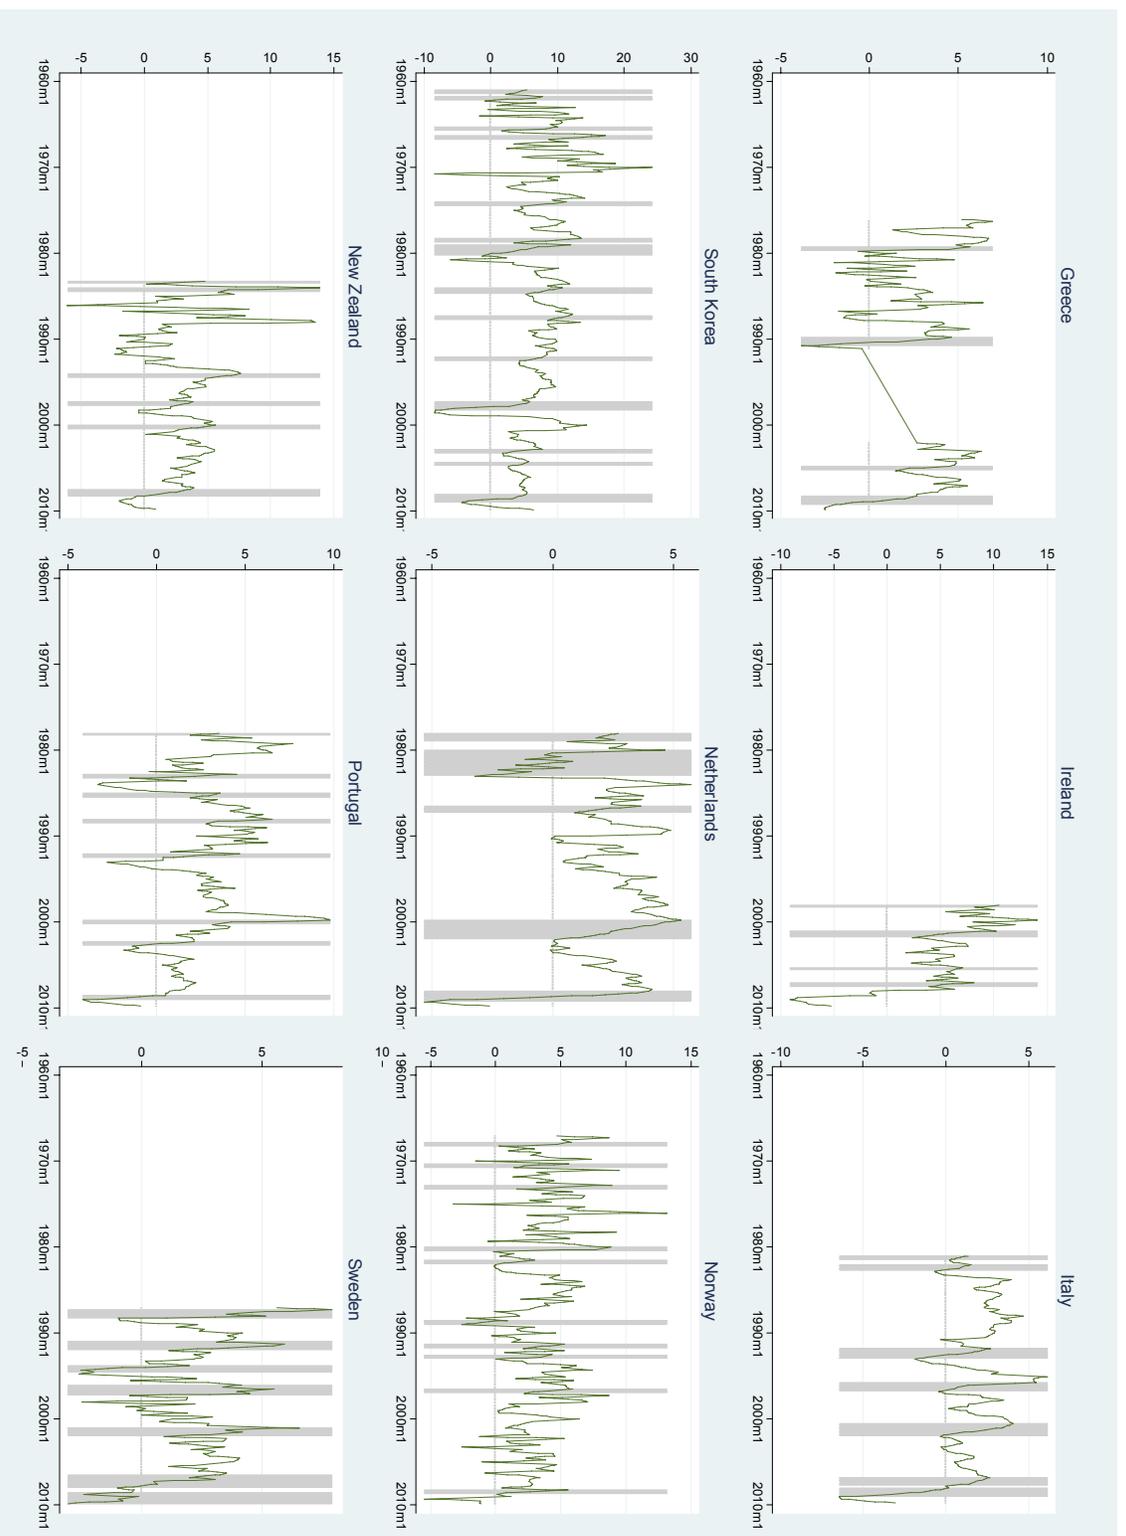


Figure 6 (cont.) – Yearly growth rates of GDP against HFS recessions, all remaining OECD countries

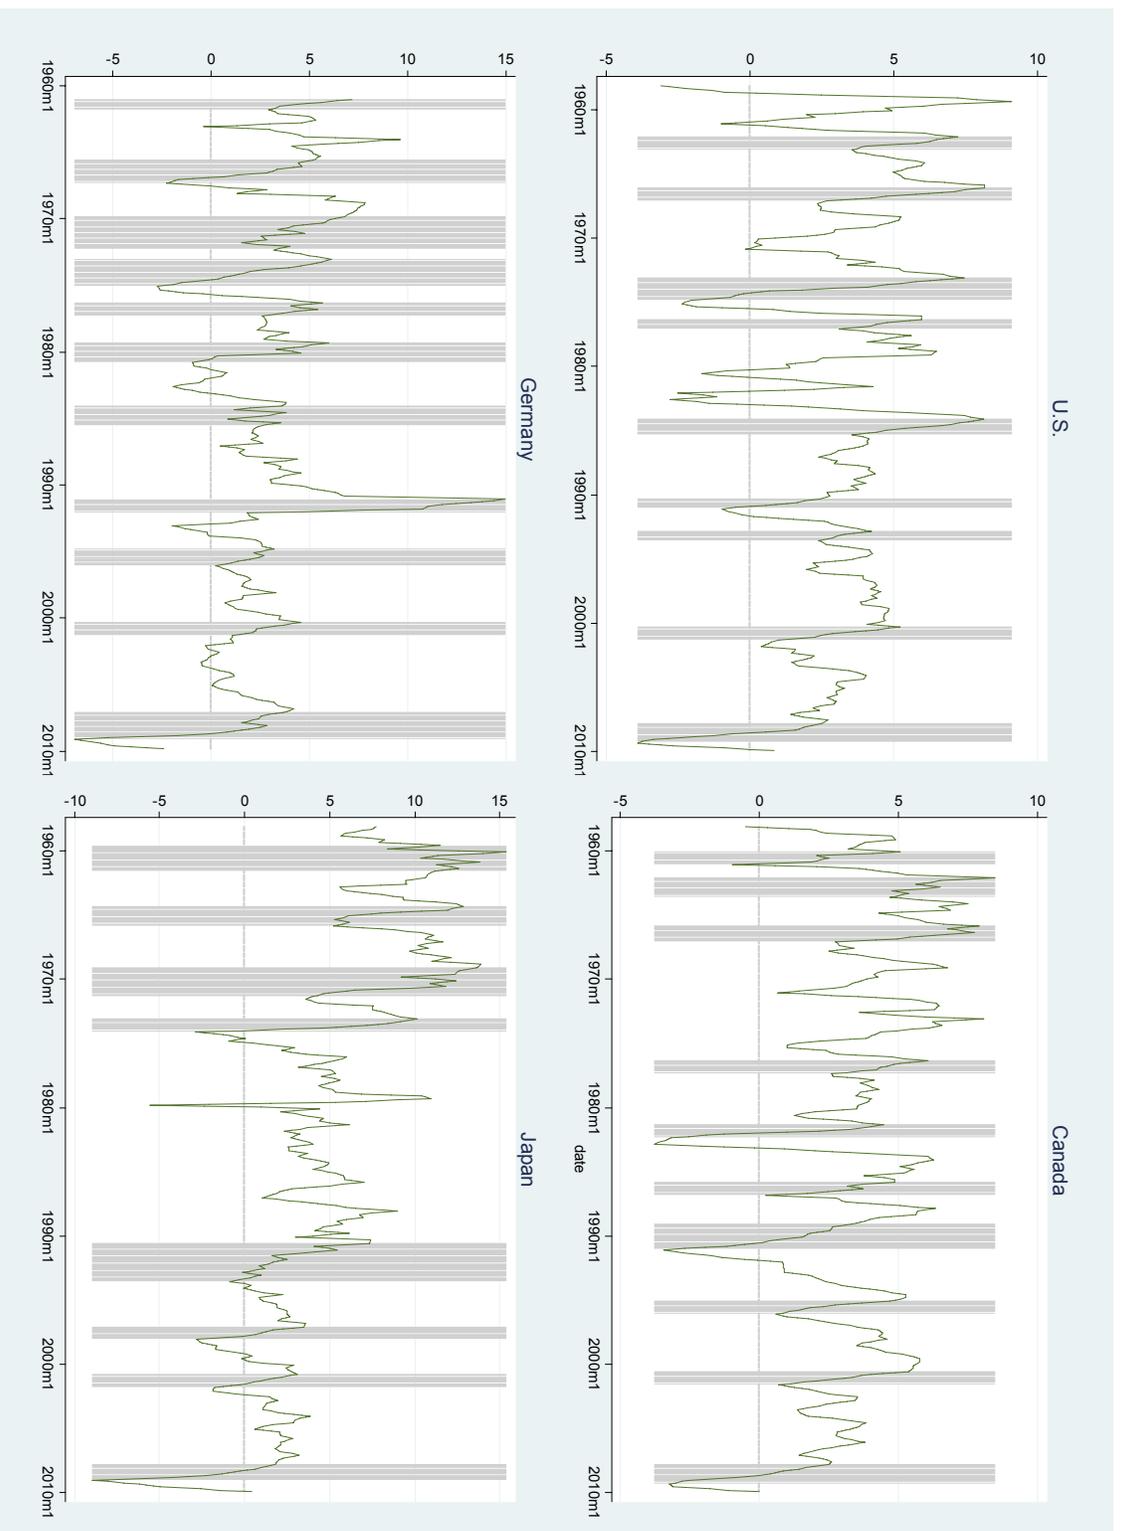
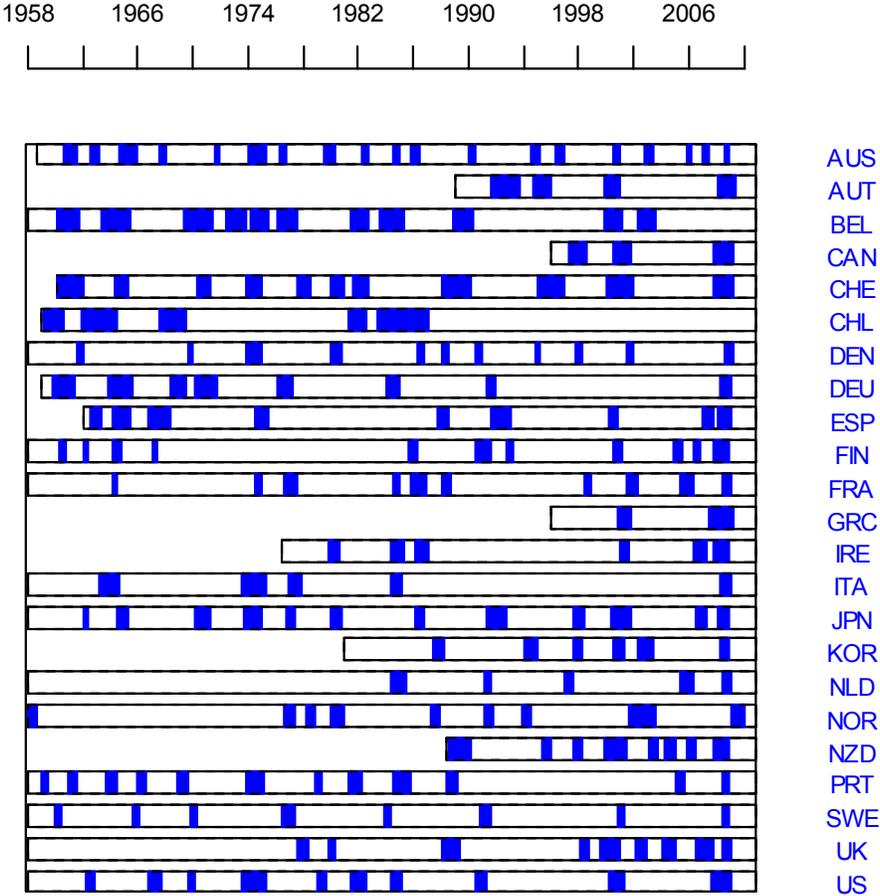
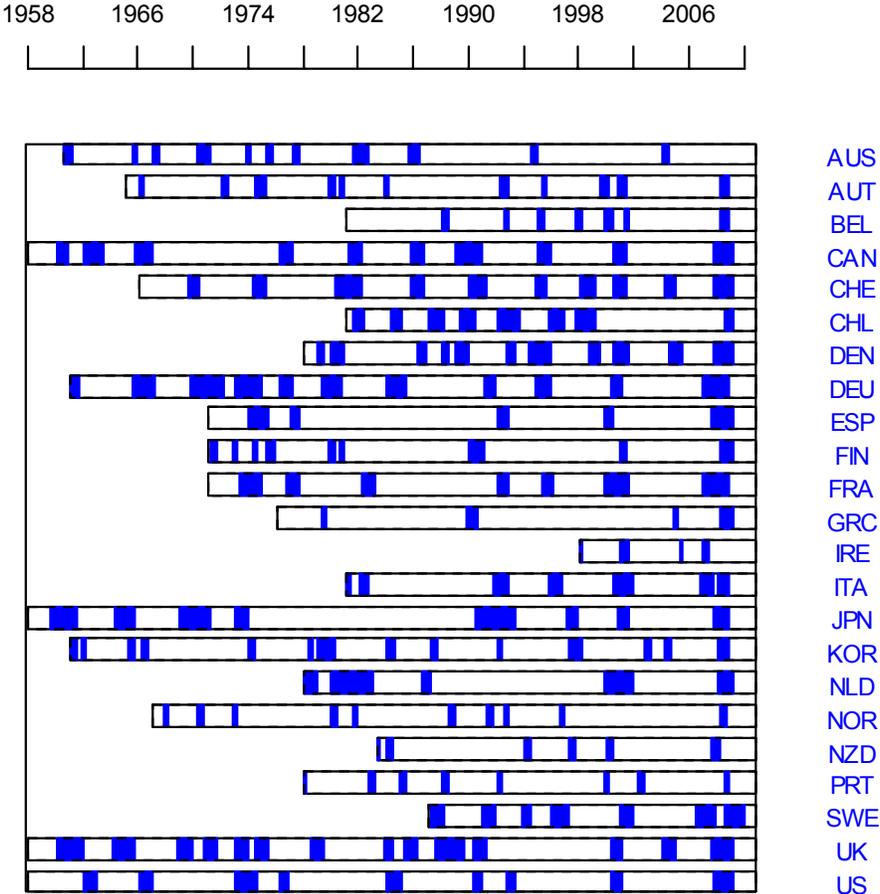


Figure 7 – HFS Chronology per country based on industrial production



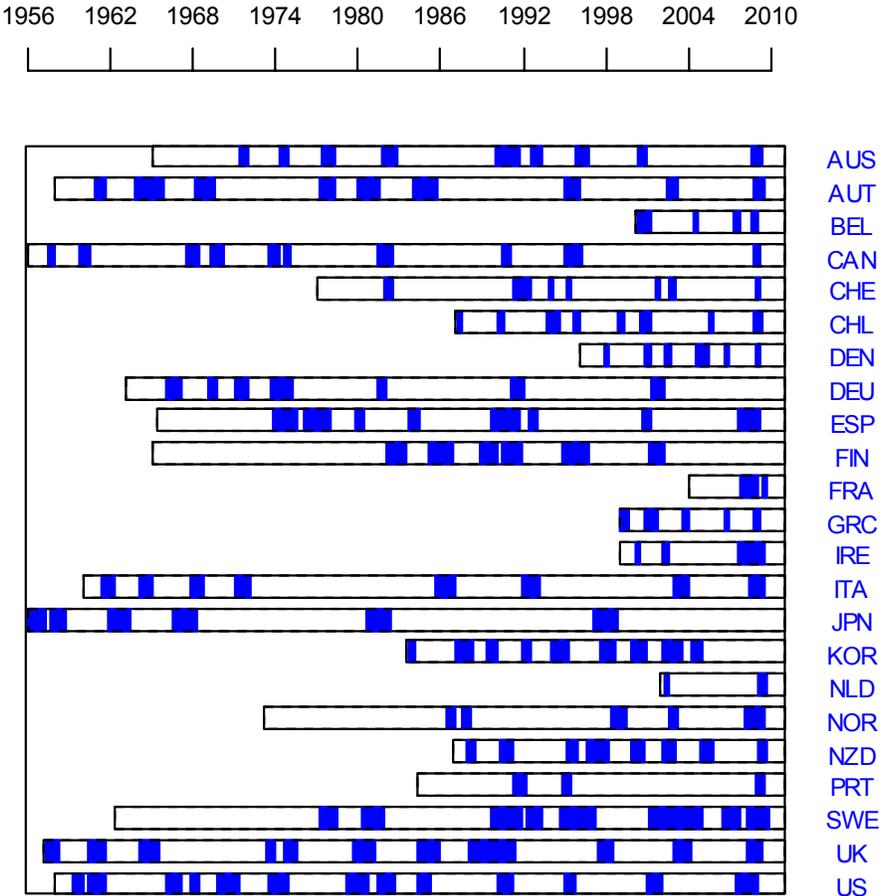
Notes: HFS recession periods denoted as bars for each country listed.

Figure 8 – HFS Chronology per country based on GDP



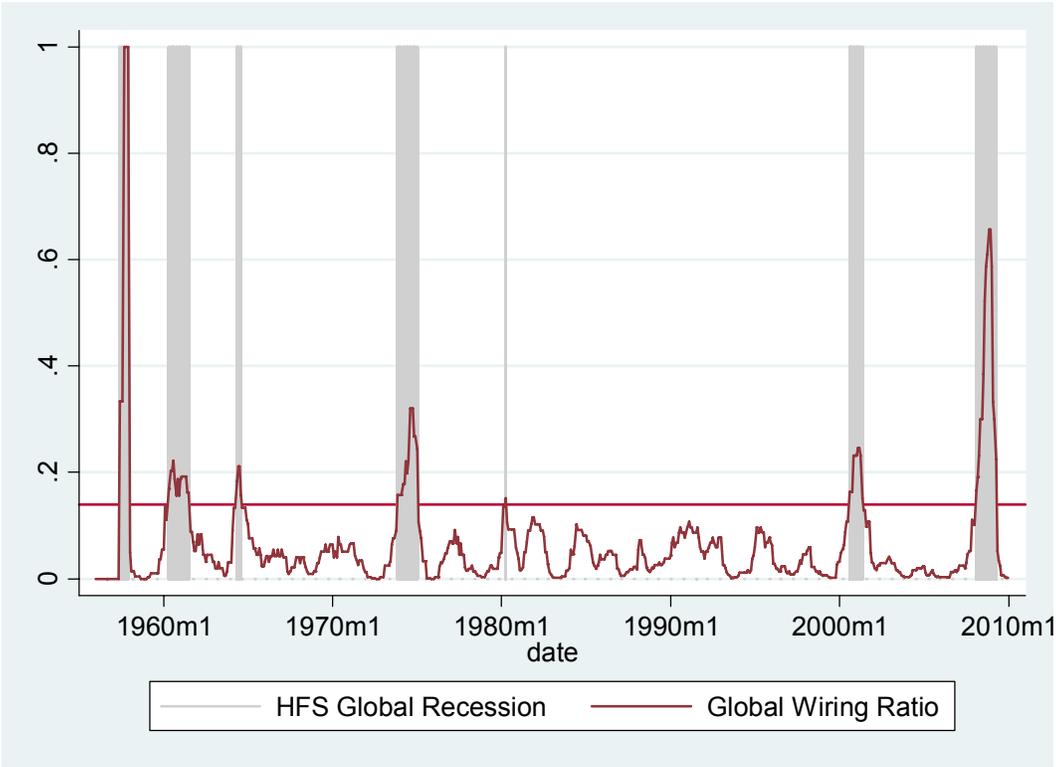
Notes: HFS recession periods denoted as bars for each country listed.

Figure 9 – HFS Chronology per country based on employment



Notes: HFS recession periods denoted as bars for each country listed.

Figure 10 – Wiring ratio of country-specific HFS recession and resulting global recession indicator



Notes: the wiring ratio is constructed as follows. For every period, count the number of pairs of countries whose HFS chronology indicates that period was a recession and normalize by the total number of possible pairs. Using the threshold of 0.14 (which corresponds to 9 or more countries simultaneously experiencing a recession) displayed in the figure as a horizontal line, we then determined the dates of “global” business cycles.