INSTITUTE
OF ECONOMICS

Scuola Superiore Sant'Anna

LEM | Laboratory of Economics and Management

Institute of Economics
Scuola Superiore Sant'Anna

Piazza Martiri della Libertà, 33 - 56127 Pisa, Italy
ph. +39 050 88.33.43
institute.economics@sssup.it

LEM
WORKING PAPER SERIES

Synchronization of endogenous business cycles

Marco Pangallo [a]

[a] CENTAI Institute, Turin, Italy.

2023/01 January 2023
ISSN(ONLINE) 2284-0400

Synchronization of endogenous business cycles

Marco Pangallo[1]

[1]CENTAI Institute, Turin 10138, Italy

January 4, 2023

Abstract

Business cycles tend to comove across countries. However, standard models that attribute comovement to propagation of exogenous shocks struggle to generate a level of comovement that is as high as in the data. In this paper, we consider models that produce business cycles endogenously, through some form of non-linear dynamics—limit cycles or chaos. These models generate stronger comovement, because they combine shock propagation with synchronization of endogenous dynamics. In particular, we study a demand-driven model in which business cycles emerge from strategic complementarities within countries, synchronizing their oscillations through international trade linkages. We develop an eigendecomposition that explores the interplay between non-linear dynamics, shock propagation and network structure, and use this theory to understand the mechanisms of synchronization. Next, we calibrate the model to data on 24 countries and show that the empirical level of comovement can only be matched by combining endogenous business cycles with exogenous shocks. Our results lend support to the hypothesis that business cycles are at least in part caused by underlying non-linear dynamics.

Key Words: Synchronization, Business Cycles, Non-linear dynamics, Networks.
JEL Class.: C61 (Dynamic Analysis), E32 (Business Fluctuations, Cycles), F44 (International Business Cycles)

[1]Contact: marco.pangallo@centai.eu. I thank Paul Beaudry, Doyne Farmer, Thomas Peron and Alexander Teytelboym for useful comments at various stages of this work. I also thank seminar participants at the EEA Virtual, NetSci, Conference of Complex Systems, EconophysiX Chair and Sant'Anna. Most of this work is included in two chapters of my DPhil thesis at the Mathematical Institute of the University of Oxford, and was performed while I was also affiliated to the Institute for New Economic Thinking at the Oxford Martin School. I acknowledge funding from INET, Baillie Gifford and from the James S. Mc Donnell Foundation. Replication materials: `https://doi.org/10.5281/zenodo.7503836`

1 Introduction

At one extreme, business cycles have been explained as the consequence of exogenous events that originate outside the economy, such as political decisions, natural catastrophes or wars. According to this view, economies would live in a stable stationary state (up to a growth trend), but their dynamics are perturbed by random shocks that pull them out of their steady state. At the other extreme, business cycles have been explained as the consequence of forces that are endogenous to the economy, such as debt dynamics, overinvestment and exuberant expectations. Under this endogenous view, the steady state of the economy is fundamentally unstable, and macroeconomic fluctuations are mathematically described by some form of non-linear dynamics—limit cycles or chaos. Of course, it is possible to combine these two extreme views by adding exogenous shocks on top of endogenous macroeconomic dynamics. However, whether the exogenous or endogenous source of business cycles dominate has fundamental scientific and policy implications. Unsurprisingly, this question has received enormous empirical and theoretical attention (see the literature review below).

In this paper, we attack this problem from a new angle. We consider a system composed of different countries that produce business cycles independently but are connected through the international trade network. In this situation, the exogenous and endogenous views of business cycles predict different mechanisms driving the comovement, or positive correlation, among the countries' economic activity. According to the exogenous view, comovement originates from the propagation of the shocks. For example, when two countries are in a steady state and country 1 is hit by a positive shock, country 1 increases its demand to country 2 and makes country 2's economy grow above the steady state, too. This produces a positive correlation between the macroeconomic dynamics of the two countries. According to the endogenous view, instead, comovement arises from the *synchronization* of the non-linear dynamics of the two countries. (In this paper, we use the term synchronization in a technical sense, namely as a property of interacting components of a dynamical system to align their non-linear dynamics.) To the extent that shock propagation and synchronization produce different empirical predictions on comovement, we can indirectly test whether the exogenous or endogenous view are better descriptions of business cycles.

It turns out that the propagation of exogenous shocks produces a much lower comovement than the synchronization of endogenous business cycles. In particular, shock propagation falls short of generating a level of comovement that is as high as in the data. By contrast, the empirical level of comovement can be matched by combining the synchronization of endogenous cycles with (relatively small) exogenous shocks. Thus, our paper offers a solution to the so-called trade-comovement puzzle (Kose and Yi, 2006), namely the inability of models based on trade links and exogenous shocks to generate sufficiently high international comovement.

The economic intuition for why synchronization can help generate higher comovement is given in Figure 1. In this illustration, we consider two countries whose economic activity depends much more on their domestic demand than on demand by the other country, as is the case empirically for most of the countries that we consider. Under exogenous business cycles, each country is hit by an idiosyncratic shock process and produces a business cycle independently of the other country. As we can see, the correlation of the macroeconomic dynamics of the two countries is virtually zero. Intuitively, this is because shocks in one country get transmitted to the other country only by the trade linkage which is small relative to domestic demand. On the contrary, if the two countries follow a deterministic endogenous business cycle (in this case, a limit cycle), their non-linear dynamics perfectly synchonize. Combining the exogenous shock processes with the endogenous business cycles gives a Pearson correlation of 0.29 between the dynamics of the two countries, a value that is consistent with empirical observations (Section 4.1.2).

The combination of shocks and deterministic dynamics implies that, at any given time, recessions occur with a certain *probability*. For example, thanks to a series of positive shocks, recessions can be delayed after a peak of the limit cycle has been reached, but recession probability increases as deterministic dynamics move towards the next trough. Because of synchronization, the two countries reach the peak of their deterministic dynamics roughly at

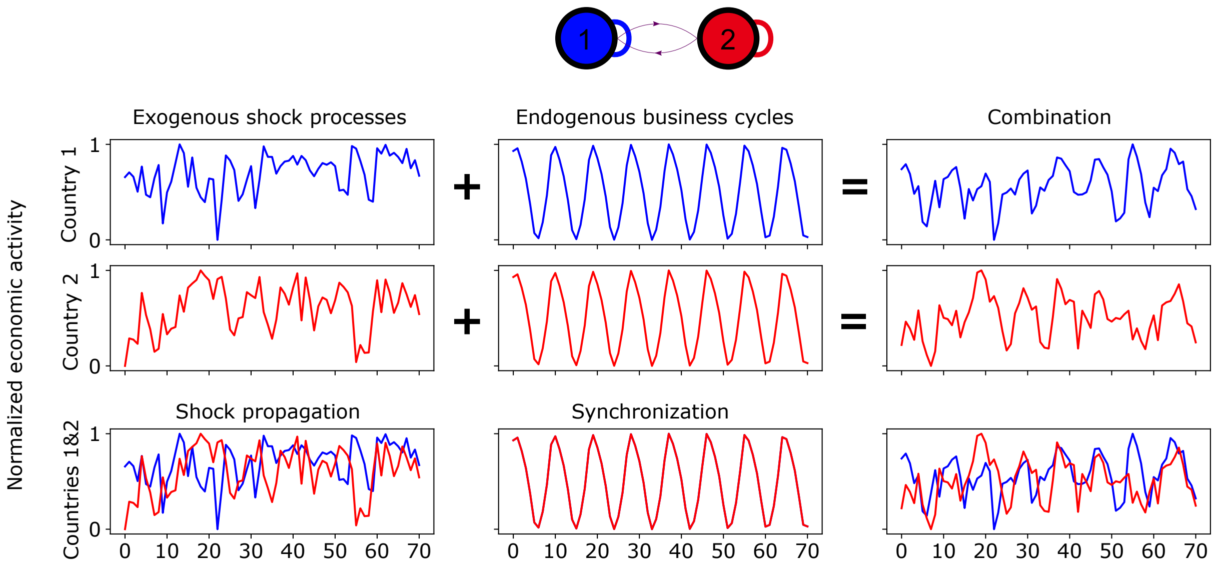

Figure 1: Illustration of the main idea of the paper. Top: We consider two nodes (countries) whose economic activity depends 90% on domestic demand (self-loops in the network representation) and 10% on the demand by the other country (arrows). Bottom: Example dynamics of an index of economic activity for countries 1 and 2, obtained by running representative simulations for 70 time steps. Left: Exogenous shock processes. In this case, comovement originates from the propagation of these exogenous shocks. Center: Endogenous business cycles. Here, comovement originates from the synchronization of the endogenous cycles. Right: Combination of the exogenous shock processes and endogenous business cycles.

the same time. The first country may fall into a recession before the second country because it is hit by a negative shock, while the second country may be enjoying a period of positive shocks. However, by reducing its international demand, the first country would easily drag the second one into a recession, too, because the second country was already predisposed to start a recession anyway.

Our goal in this paper is to illustrate how synchronization theory can be useful to explain business cycle comovement. However, the theory that we develop can be applied to explain comovement of any other economic or financial time series, such as commodity, housing and stock prices. To highlight the generality of our approach, we avoid choosing models that focus on specific economic forces causing endogenous business cycles, such as finance, expectations, overinvestment or technological progress. We instead adapt the reduced-form model in Beaudry et al. (2020), in which endogenous cycles have more abstract origin. In particular, they are caused by strategic complementarities, that is, the tendency of agents to increase their action if other agents increase their action, too. In macroeconomics (Cooper and John, 1988), this can be thought of as the tendency of firms to increase production if other firms increase production themselves, or the tendency of households to increase consumption if other households in their social network increase consumption. This abstract framework can encompass many of the causes for endogenous business cycles listed above.

Our first set of results is theoretical. As we want to understand the interplay between synchronization of deterministic non-liner dynamics, exogenous shocks and network structure, we adapt an approach known as complete synchronization theory (Pecora and Carroll, 1998) to our case. Here, links tend to align the non-linear dynamics of the nodes, but idiosyncratic shock processes tend to pull them apart. Complete synchronization is an elegant mathematical formalism that makes it possible to quantify the relative strength of synchronizing and desynchronizing forces, through an eigenvalue-eigenvector decomposition that takes into account the network connecting the different components. The eigenvalues and eigenvectors give information about how quickly the system may synchronize after some components are hit by idiosyncratic shocks, about which components have more synchronized dynamics depending on the presence of clusters, etc.

Next, we specify the abstract model so that the nodes are 24 countries and a Rest of the World aggregate, and weighted, directed, links represent international demand and domestic demand (self-loops). In our empirical application, we consider different parameterizations under which the deterministic model either produces endogenous business cycles or converges to a stable steady state. We then add exogenous shocks to the dynamics. We show that only a combination of endogenous business cycles and exogenous shocks can match the level of comovement that can be found in the data. Moreover, this combination is also best at reproducing the mean correlation between one country and the other countries in the sample. For instance, both in the data and in the endogenous model China and Brazil are the least correlated to the other countries over the period 1950-2019, whereas Belgium and the Netherlands are the most correlated. Interestingly, these patterns could be largely expected based on the values of the second eigenvector of the Laplacian of the international trade network, in accordance with theory.

Relation to the literature. Our paper relates first of all to the few papers that in the past have seen the appeal of synchronization theory to explain business cycles comovement (Haxholdt et al., 1995; Selover and Jensen, 1999; Brenner et al., 2002; Matsuyama et al., 2014). Our paper is different in several respects. First, we use a different concept of synchronization. These papers use the most well-known concept of *phase synchronization*, that implies that each component of the synchronizing dynamical system has a relatively well-defined natural frequency and phase; under synchronization the frequencies would align to a common frequency and the phases would lock. The problem with this concept is that in economics it is hard to characterize oscillations with a well-defined frequency and phase. In this paper, we use the more general concept of complete synchronization. To our knowledge, this is the first application of this concept to economics. The second important difference is that the papers above are purely theoretical, whereas here we test if indeed synchronization makes it possible to improve over shock propagation in explaining comovement. Other differences are that our paper (i) uses a general model that can be adapted to several economic circumstances, rather than relying on a specific business cycle model; (ii) studies synchronization on a network of an arbitrary number of agents, rather than just two agents or a trivial topology; (iii) produces "realistic" smooth dynamics, as opposed to e.g. limit cycles of period two.

Our paper contributes to the literature that explains international business cycles comovement. While many papers are mostly empirical,[1] some papers attempt to explain comovement by building international real business cycle models (Backus et al., 1992) in which comovement originates from the propagation of exogenous shocks. These works generally struggle to obtain a level of comovement that is as high as in the data (Arkolakis and Ramanarayanan, 2009; Johnson, 2014; Liao and Santacreu, 2015), in line with the already mentioned "trade-comovement puzzle" (Kose and Yi, 2006). As already outlined above, we contribute to these strands of literature by showing how synchronization of non-linear dynamics is a powerful way to generate higher international comovement.

More broadly, we contribute to the debate on the nature of business cycles as exogenous or endogenous. Over the last 80 years, there have been cycles of preferences towards one view or another. Following the early theoretical work on endogenous business cycles by e.g. Kaldor (1940); Hicks (1950); Goodwin (1951), the real business cycles revolution (Kydland and Prescott, 1982) popularized the use of exogenous shocks to characterize business cycles. At the same time, the discovery of chaos theory revived the interest in non-linear dynamic models, and a flurry of papers on general equilibrium models producing endogenous business cycles came out (Boldrin and Woodford, 1990). This interest died out as it was not possible to find conclusive evidence on the existence of non-linear dynamics (limit cycles or chaos) in economic time series (Barnett and Serletis, 2000).[2] Recently, Beaudry et al. (2020) showed

[1]See, for example, Frankel and Rose (1998); Baxter and Kouparitsas (2005); Imbs (2006); Calderon et al. (2007); Di Giovanni and Levchenko (2010); Ng (2010); Hsu et al. (2011); Kalemli-Ozcan et al. (2013); Cravino and Levchenko (2017); Di Giovanni et al. (2018).

[2]See also Brock (1986); Barnett and Chen (1988); Scheinkman and LeBaron (1989); Barnett and Serletis

that some U.S. time series produce a clear spectral peak between 9 and 10 years, a possibility that had been disregarded since the seminal work of Granger (1966). This, together with the general framework based on strategic complementarities that they introduce, aims at reviving the endogenous view of economic fluctuations. Our paper contributes to this revival by showing that allowing for business cycles to be at least in part endogenous makes it possible to produce a level of international comovement that is as high as in the data, thus providing indirect evidence in favor of endogenous business cycles.

Road map. This paper proceeds as follows. In Section 2 we introduce our model of endogenous business cycles, both in its abstract and international specifications, and analyze its dynamics. In Section 3 we obtain theoretical results on complete synchronization. Finally, in Section 4, we compare the predictions of our model to data. Section 5 concludes.

2 A general framework for endogenous business cycles

Over decades, researchers have proposed several economic forces that could generate endogenous business cycles.[3] Some scholars focused on the role of finance, others on real factors such as investment and inventories, still others on bounded rationality and adaptive expectations. Even more forces have been proposed, including overlapping generations effects, preference switching, search and matching, technological progress, and wage bargaining. To discuss synchronization of endogenous business cycles, it would be limiting to focus on a model that picks a particular economic force to produce business cycles in the first place, as results would appear to depend on this choice. It would instead be desirable to use a macroeconomic framework that is as general as possible.

Luckily, Beaudry et al. (2015) have recently proposed such a framework.[4] The authors focus on strategic complementarities (Bulow et al., 1985; Cooper and John, 1988) between agents as the main force leading to endogenous cycles. In their framework, agents' decisions self-reinforce through interactions, leading to an instability of the steady state, but at some point sufficiently far from the steady state, implosive forces set in that prevent dynamics from exploding. Many of the narratives for endogenous business cycles listed above can be cast within this framework, as will be clarified below.

Here we present a slightly modified version of the model in Beaudry et al. (2015), who assume a symmetric solution in which all agents behave alike and face the same homogeneous interaction network. Here, we instead let both decisions and the interaction network be heterogeneous across agents. In this section, we start introducing the model in an abstract interaction network, and then consider a specification of the model in which interactions correspond to international trade. We finally discuss how this model can generate endogenous cycles.

2.1 Abstract formulation

Consider N agents, indexed by $i \in \mathcal{I} = \{1, \ldots, N\}$. Denote by x_i an accumulation variable representing a stock of agent i, and by y_i a decision variable representing a flow. We can think of y as investment and of x as capital; we can also think of y as level of production or consumption of durable goods, and of x as inventories or net worth, respectively.

The time evolution of the accumulation variable x_i is simple. At each time step t, it depreciates by a factor δ, and increases by the decision variable $y_{i,t}$. In formula, $x_{i,t+1} = (1-\delta)x_{i,t} + y_{i,t}$.

(2000); Shintani and Linton (2004); Hommes and Manzan (2006).

[3]See, among many others, the contributions and review articles in: Kaldor (1940); Hicks (1950); Goodwin (1951, 1967); Shleifer (1986); Boldrin and Woodford (1990); Foley (1992); Silverberg and Lehnert (1993); Bullard (1994); Matsuyama (2007); Fazzari et al. (2008); De Grauwe (2011); Nikolaidi and Stockhammer (2017).

[4]Here we reference the initial working paper (Beaudry et al., 2015), rather than the final journal version (Beaudry et al., 2020), because it specifies this general framework in greater detail.

The dynamics of the decision variable y_i is more involved and captures the effect of the interactions with the other agents. The planned level of y for agent i, $y_{i,t+1}$, depends most importantly on a term $\overline{y}_{i,t}$, capturing the interactions among the decision variables y of the agents. The interaction term $\overline{y}_{i,t}$ is defined as $\overline{y}_{i,t} = \sum_{j=1}^{N} \epsilon_{ij} y_{j,t}$, where $\epsilon_{ij} \in [0,1]$, such that $\sum_j \epsilon_{ij} = 1$. Each term ϵ_{ij}, which can be thought of as an *interaction coefficient*, is the weight that the decision variable of agent j has in determining the value of $\overline{y}_{i,t}$ (self-interactions ϵ_{ii} are also considered). The values of ϵ_{ij}, for all i and j, define a weighted, directed, interaction network.

The effect of $\overline{y}_{i,t}$ is mediated by a non-linear function $F(\cdot)$ that determines the effect of interactions. Suppose that the decision variables of the agents with whom agent i interacts increase, so that $\overline{y}_{i,t}$ becomes larger. If, as a consequence, agent i's marginal payoff goes up, one says that there are *strategic complementarities* between agent i and the agents with whom she interacts (Cooper and John, 1988).[5] Intuitively, in case of strategic complementarities, agent i decides to increase her decision variable if the agents with whom she interacts do the same, i.e. $y_{i,t+1} > y_{i,t}$ when $\overline{y}_{i,t+1} > \overline{y}_{i,t}$. If instead, due to an increase of $\overline{y}_{i,t}$, the marginal payoff of agent i decreases, and so agent i reduces her decision variable, one talks about *strategic substitutability.* We choose the function $F(\cdot)$ so that, depending on the value of $\overline{y}_{i,t}$, our model encompasses both the complementarity and substitutability regimes. What generates cyclical behavior in this model is the continuous switching between the two regimes (see Section 2.3).

We now complete the description of the model. We first assume that $y_{i,t+1}$ also linearly depends on one-step lagged values of x and y, with coefficients α_1 and α_2 (further to an intercept α_0). We take α_1 to be negative, and α_2 to be positive and bounded between zero and one (Beaudry et al., 2015, 2020). These assumptions reflect, respectively, decreasing returns to accumulation, i.e. willingness to avoid excessive stocks of inventories, capital or net worth, and sluggishness in the adjustment of the decision variable, i.e. difficulty to quickly modify the level of production, investment, or consumption of durable goods. The assumption on α_1 is instrumental for the switching between regimes of strategic complementarity and substitutability, while the assumption about α_2 introduces realistic smoothness in the dynamics.[6]

We finally add idiosyncratic shock terms $u_{i,t}$ to the evolution of the decision variables. We parameterize this shock process as an AR(1), i.e. $u_{i,t+1} = \rho u_{i,t} + \iota_{i,t}$, where $\rho \in [0,1]$ is a persistence parameter and $\iota_{i,t}$ is white noise normally distributed with mean zero and standard deviation σ.

Our model for endogenous business cycles is thus fully specified by the following equations, for all agents i:

$$
\begin{aligned}
&x_{i,t+1} = (1-\delta)x_{i,t} + y_{i,t}, \\
&y_{i,t+1} = \alpha_0 + \alpha_1 x_{i,t} + \alpha_2 y_{i,t} + F\left(\overline{y}_{i,t}\right) + u_{i,t}, \\
&\overline{y}_{i,t} = \textstyle\sum_{j=1}^{N} \epsilon_{ij} y_{j,t}, \quad \epsilon_{ij} \in [0,1], \quad \sum_j \epsilon_{ij} = 1, \\
&F(y) = \beta_0 + \beta_1 y + \beta_2 y^2 + \beta_3 y^3 + \beta_4 y^4.
\end{aligned} \tag{1}
$$

The last equation represents our arbitrary choice for $F(\cdot)$, that we parameterize as a generic quartic function (Beaudry et al., 2015). The choice of a quartic rather than another functional form is not important, to the extent that $F(\cdot)$ can be parameterized as will be described in Section 2.3.

In the case in which all interaction coefficients ϵ_{ij} are identical and equal to $1/N$ (complete and homogenous interaction network), and all agents behave alike, $x_{i,t} = x_{j,t}$ and $y_{i,t} = y_{j,t}$, for all agents i and j and for all times t (and shocks are identical across all agents), equations (1) exactly recover the model in Beaudry et al. (2015).[7]

[5] Letting $V_{i,t}$ be the payoff function for agent i at time t, strategic complementarities correspond to the condition $\frac{\partial^2 V_{i,t}}{\partial y_{i,t} \partial \overline{y}_{i,t}} > 0$.

[6] A problem of many endogenous business cycle models is that they generate a "sawtooth" dynamics that has no persistence, in stark contrast to empirical time series.

[7] A minor difference with respect to Beaudry et al. (2015) is that $y_{i,t+1}$ depends only on variables at t, while

2.2 International model

We now interpret the agents in the abstract model as countries. What should their interaction coefficients represent? Intuitively, interaction coefficients should represent something that is known to lead to comovement of business cycles across countries. However, several economic forces could potentially fit this purpose; for instance, financial integration could intuitively be assumed to lead to high comovement. Yet, the effect of financial integration on comovement is ambigouous: while Imbs (2006) finds positive effects of financial integration on international comovement, Kalemli-Ozcan et al. (2013), by using a highly disaggregated dataset, conclude that, all else equal, financial integration leads to less comovement. Thus, here instead we follow the large literature that identifies trade linkages as the main factor causing comovement between countries. Indeed, although, theoretically, strong trade ties could lead to specialization and so potentially weaker comovement (Frankel and Rose, 1998), in practice all empirical studies find a statistically significant positive effect of trade on comovement.[8]

A practical difference with respect to the abstract model described in the previous section is that countries have diverse sizes. This would imply that the accumulation and decision variables would be on different scales, making it necessary to consider a different set of parameters for each country. To avoid this impractical situation, we work with "oscillation" variables $x_{i,t}$ and $y_{i,t}$ that define the oscillation around the steady state. In formula, we let total investment (or production, or durable consumption) be given by $Y_{i,t} = \widetilde{Y}_i y_{i,t}$ and total capital (inventories, stock of durables) be $X_{i,t} = \widetilde{X}_i x_{i,t}$. In these expressions, $\widetilde{Y}_i$ and $\widetilde{X}_i$ denote steady-state values of $Y_{i,t}$ and $X_{i,t}$ respectively, and $y_{i,t}$ and $x_{i,t}$ oscillate with mean one (or take value one if the dynamics reaches a steady state).[9] Then

$$\begin{aligned} x_{i,t+1} &:= X_{i,t+1}/\widetilde{X}_i = (1-\delta)x_{i,t} + y_{i,t}, \\ y_{i,t+1} &:= Y_{i,t+1}/\widetilde{Y}_i = \alpha_0 + \alpha_1 x_{i,t} + \alpha_2 y_{i,t} + F\left(\overline{y}_{i,t}\right) + u_{i,t}. \end{aligned} \tag{2}$$

In this application, the interaction term $\overline{y}_{i,t}$ represents the oscillation of the total demand received by country i. Letting total demand be denoted by $D_{i,t}$, we assume that its steady state $\widetilde{D}_i$ is equal to the steady state of the flow variable $Y_{i,t}$, that is $\widetilde{D}_i = \widetilde{Y}_i$. This corresponds to assuming market clearing in the steady state: depending on the interpretation for the decision variable, it could be that total demand equals production, or that total demand of capital goods equals investment, or that total demand of durable goods equals realized consumption of such goods. Under this assumption, we can write $D_{i,t} = \widetilde{Y}_i d_{i,t}$, where again $d_{i,t}$ is the oscillation of demand around its steady-state value. So, in this application $\overline{y}_{i,t} := d_{i,t}$.

The key behavioral assumption to compute total demand is that countries demand fixed proportions of imports from other countries. Letting $\widetilde{Y}_{ij}$ be the steady-state trade flow from country i to country j, and being $a_{ij} = \widetilde{Y}_{ij}/\widetilde{Y}_j$ the fraction of imports from country i over country j's output, we write the total demand that country i receives as $D_{i,t} = \sum_j a_{ij} Y_{j,t}$. This means that, whether country j produces more or less than the steady-state value $\widetilde{Y}_j$, it still demands the fixed proportion a_{ij} of imports from i. Equating the two formulations for $D_{i,t}$, $D_{i,t} = \widetilde{Y}_i d_{i,t}$ and $D_{i,t} = \sum_j a_{ij} Y_{j,t}$, we can express the oscillations of demand $d_{i,t}$

in Beaudry et al. (2015) it depended on both variables at t and (contemporaneous) variables at $t+1$. As there did not seem to be much difference in the implied dynamics, we chose this form as it was computationally simpler. Economically, it corresponds to assuming that the decision variables of the agents with whom agent i interacts have a lagged effect on her decision.

[8] See Frankel and Rose (1998); Clark and Van Wincoop (2001); Imbs (2004); Baxter and Kouparitsas (2005); Kose and Yi (2006); Calderon et al. (2007); Inklaar et al. (2008); Hsu et al. (2011); Kleinert et al. (2015) and Di Giovanni et al. (2018). The magnitude of this effect is more debated. While Frankel and Rose (1998) and Baxter and Kouparitsas (2005) find a very strong effect, Imbs (2004) and Inklaar et al. (2008) find a much smaller effect.

[9] It is $\widetilde{X}_i = \widetilde{Y}_i/\delta$, as we show in Section 2.3.

in terms of the oscillations of the decision variables $y_{j,t}$, as:

$$d_{i,t} = \frac{1}{\widetilde{Y}_i} \sum_j \frac{\widetilde{Y}_{ij}}{\widetilde{Y}_j} \widetilde{Y}_j y_{j,t} = \sum_j \frac{\widetilde{Y}_{ij}}{\widetilde{Y}_i} y_{j,t} = \sum_j \epsilon_{ij} y_{j,t}, \tag{3}$$

where the coefficients $\epsilon_{ij} = \widetilde{Y}_{ij}/\widetilde{Y}_i$ are the fraction of the output of country i that is *exported* to j. Importantly, this includes the fraction of output that remains in country i, ϵ_{ii}, representing domestic demand. These coefficients satisfy the conditions laid out in Eq. 1: $\epsilon_{ij} \in [0, 1]$ and $\sum_j \epsilon_{ij} = 1$.

In sum, the abstract model introduced in Section 2.1 can be interpreted as a model of business cycles of countries coupled by an international trade network. Under strategic complementarities, an increase in domestic and international demand may prompt firms in several countries to increase their decision variables, further increasing domestic and international demand. This could make the steady state unstable, and, as we show in the next section, the model could produce endogenous business cycles.

2.3 Homogeneous dynamics

Here we mathematically describe and build some intuition into how the general framework for endogenous fluctuations laid out in the previous sections may generate cyclical dynamics (see Beaudry et al. 2015 for more details). For simplicity, we focus on a single country i whose dynamics is purely driven by internal demand, namely $\epsilon_{ii} = 1$, and $\epsilon_{ij} = 0$, for all $j \neq i$ (autarchic case). This can be interpreted as the country being composed by a large number of agents with a homogeneous interaction network and we consider the symmetric solution in which all agents behave alike, as in Beaudry et al. (2015). For ease of notation, in this section we drop the subscript i.

We first look for the steady state $(\hat{x}, \hat{y})$. Following the characterization of x and y as oscillations around the steady state, as discussed in Section 2.2, for convenience we constrain parameters so that $\hat{y} = 1$ is the unique steady state. This means that the steady state value of the accumulation variable is $\hat{x} = 1/\delta$ and the following relation must hold: $\alpha_0 = 1 - \alpha_1/\delta - \alpha_2 - F(1)$. In the following, we always select α_0 so that this condition is satisfied. For the steady state to be unique, the slope of the line $-\alpha_0 + y_t(1 - \alpha_1/\delta - \alpha_2)$ must be larger than the derivative of F at $\hat{y} = 1$, namely $1 - \alpha_1/\delta - \alpha_2 > F'(1)$. The Jacobian of this system is

$$\mathcal{J} = \begin{pmatrix} 1-\delta & 1 \\ \alpha_1 & \alpha_2 + F'(1) \end{pmatrix}. \tag{4}$$

The stability of the steady state is completely characterized in terms of the trace $T = 1 - \delta + \alpha_2 + F'(1)$ and determinant $D = (1-\delta)\,(\alpha_2 + F'(1)) - \alpha_1$ of the Jacobian. Following the standard conditions for 2-dimensional maps, stability obtains if $D < 1$, $D > T - 1$ and $D > -T - 1$ (the grey region in Figure 2A, representing the usual diagram for the stability of 2-dimensional maps). If $D > (T)^2/4$, the eigenvalues of the Jacobian are complex, and so the system admits either damped oscillations or sustained cycles. If the line $D = 1$ is crossed above $D = (T)^2/4$, the system undergoes a Hopf (Neimark-Sacker) bifurcation. Beaudry et al. (2015) show that the Hopf bifurcation is supercritical, i.e. the resulting limit cycle is attractive.

To build some intuition into the transition between stability and instability, we fix the parameters α_1, α_2, and δ, and vary the values of the β parameters in Eq. (1) so that $F'(1)$ varies in the interval $F'(1) \in [-1, 1]$. In other words, we consider the level of complementarity or substitutability at the steady state as a bifurcation parameter. Point P corresponds to $F'(1) = -1$. This corresponds to strategic substitutability, and indeed under this parameterization the steady state is a stable node, meaning that eigenvalues are real and of magnitude smaller than one. When the steady state is a node, dynamics converge without oscillations. Next, point Q corresponds to $F'(1) = 0$, i.e. to a transition between strategic complementarity and substitutability. Under these parameters, the system has a stable focus, with complex eigenvalues of magnitude smaller than one. In this setting, dynamics

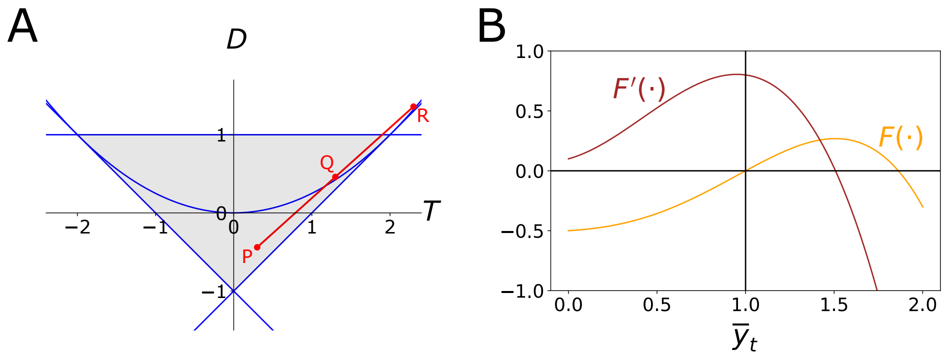

Figure 2: Stability of dynamics for homogeneous agents. (A) Diagram for stability of 2-dimensional maps. The blue lines correspond to $D = 1$, $D = T - 1$, $D = -T - 1$ and $D = (T)^2/4$. The red line corresponds to varying $F'(1)$ with the other parameters fixed. (B) Function $F(\cdot)$ and its first derivative $F'(\cdot)$ for $\overline{y}_t \in [0, 2]$.

converge to the steady state producing damped oscillations. As the level of complementarities increases, the steady state loses stability through a Hopf bifurcation, up to point R, corresponding to $F'(1) = 1$.

The analysis so far has focused on local stability of the steady state. To build some intuition into global dynamics, in Figure 2B we plot the function $F(\cdot)$ and its first derivative for some values of the β parameters. Under these parameters, the function F increases at the steady state, $\overline{y}_t = 1$, but decreases when $\overline{y}_t$ becomes large, guaranteeing the existence of a regime of strategic substitutability that prevents explosive dynamics.

Starting from a situation below the steady state, the regime of strategic complementarity makes it optimal for the agents to increase their decision variables y when the other agents are doing the same, in a way that they overshoot the steady state. Explosive economic forces are at play. However, as the agents start to over-accumulate because the value of y more than offsets the depreciation $(1 - \delta)x$, the rise of y slows down. This is both because the agents dislike accumulation and because, as y increases, the regime of strategic substitutability sets in. Soon, y reverts and starts to decrease. At some point, the decrease in y brings back to a regime of strategic complementarity, making the agents increase their decision variables again, and the process restarts.

Depending on the interpretation for the accumulation variable x, the decision variable y, and the function F, the narrative based on the switching between strategic complementarity and substitutability regimes could correspond to several economic forces causing endogenous business cycles. Beaudry et al. (2020), for example, consider a microfounded model in which x is net worth, y is consumption of durable goods, and F models banks' willingness to give loans. In good times, lending is perceived safe, and so agents can borrow to consume more durable goods. This further strengthens the economic boom, making lending to be perceived even safer (strategic complementarity). When the economy slows down because of overaccumulation, lending instead starts to be perceived less safe, making banks cut back on credit (strategic substitutability). This behavior of banks can easily trigger a recession, which lasts until agents have liquidated assets in excess, at which point the cycle starts again. The narrative in Beaudry et al. (2020) is similar to many of the previously proposed narratives for finance-based endogenous business cycles. However, the strategic complementarity-substitutability framework easily lends itself to other narratives, such as ones based on overinvestment in the Keynesian tradition (Kaldor, 1940; Hicks, 1950; Goodwin, 1951), and ones based on temporal clustering of technological innovations (Judd, 1985; Shleifer, 1986; Matsuyama, 1999).

2.4 A numerical example

We conclude this section by showing representative time series for the various cases discussed above: stable steady state of type node and focus, and limit cycle. We select the α_1, α_2

and δ parameters of the limit cycle so that every cycle lasts about 36 time steps which, if we interpret time steps as quarters, is in line with the empirical evidence found in Beaudry et al. (2020) that the US economy tends to undergo a business cycle every about 9 years. Under the limit cycle parameterization, the β parameters of the function $F(\cdot)$ lead to strong strategic complementarities at the steady state. We then select the parameters of the focus and node by modifying the β parameters so that there are weak complementarities at the steady state. Finally, we also consider a parameterization that leads to chaotic dynamics.[10]

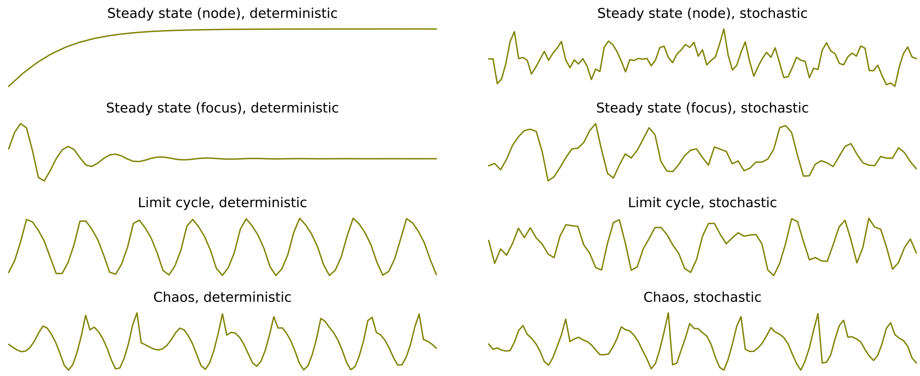

Figure 3: Qualitatively different types of dynamics (node, focus, limit cycle and chaos), depending on the strength of strategic complementarities and other parameters, for both the deterministic and stochastic cases.

We show the time series for each case in Figure 3, showing both the corresponding deterministic and stochastic dynamics. (In all cases, the stochastic terms $u_{i,t}$ follow an AR(1) process with autocorrelation $\rho = 0.3$ and standard deviation $\sigma = 0.1$.) When the steady state is a node, deterministic dynamics converge to the steady state without oscillations, and the corresponding stochastic dynamics are quite irregular. When the steady state is a focus, convergence happens under damped oscillations, and the corresponding stochastic dynamics show more structure. Under limit cycles, deterministic dynamics are perfectly periodic, while stochastic dynamics retain the shape of the cycle with clearer periodicity than under a focus. Finally, under chaotic dynamics, even in the deterministic case there is no perfect periodicity, and it is even more so in the stochastic case.

In conclusion, the dynamical system (1) exhibits a variety of dynamics, some that converge towards stable steady states, in line with the usual narrative of business cycles driven by exogenous shocks, and others that instead follow limit cycles or chaos, resulting in endogenous business cycles. When adding noise, all dynamics could qualitatively resemble business cycle fluctuations, which are clearly not perfectly periodic but do show some structure.

3 Complete synchronization theory

Synchronization theory is about the alignment of dynamics of weakly-interacting self-oscillating units (Pikovsky et al., 2003). It is an important requirement for synchronization theory that different units of a system would still display some form of self-sustaining non-linear dynamics (limit cycles or chaos) if they were uncoupled. Synchronization is not about whether stable units, when coupled, would become unstable (Gualdi et al., 2015); it is about whether

[10] The specific values of the parameters are as follows:

- Node: $\alpha_1 = -0.04$, $\alpha_2 = 0.4$, $\delta = 0.1$, $\beta_0 = -0.19$, $\beta_1 = -0.11$, $\beta_2 = 0.4$, $\beta_3 = 0.2$, $\beta_4 = -0.3$.
- Focus: $\alpha_1 = -0.04$, $\alpha_2 = 0.4$, $\delta = 0.1$, $\beta_0 = -0.2$, $\beta_1 = -0.1$, $\beta_2 = 0.1$, $\beta_3 = 0.3$, $\beta_4 = -0.1$.
- Cycle: $\alpha_1 = -0.04$, $\alpha_2 = 0.4$, $\delta = 0.1$, $\beta_0 = -0.5$, $\beta_1 = 0.1$, $\beta_2 = 0.2$, $\beta_3 = 0.5$, $\beta_4 = -0.3$.
- Chaos: $\alpha_1 = -0.35$, $\alpha_2 = 0.4$, $\delta = 0.1$, $\beta_0 = -0.5$, $\beta_1 = 0.1$, $\beta_2 = 0.2$, $\beta_3 = 0.5$, $\beta_4 = -0.3$.

In all cases, we selected the β parameters so that, conveniently, it is $F(1) = 0$, as in Beaudry et al. (2015).

the dynamics of unstable units, when coupled, would align. The coupling is usually "weak" in the sense that the dynamics of the nodes are only weakly affected by the other nodes at a given time. This characterization of synchronization theory seems ideally suited to our case study, featuring different countries, weakly coupled through international trade (Section 2.2), displaying endogenous business cycles on their own (Section 2.3).

A large part of synchronization theory is about the alignment of the frequency and phase of the oscillations (phase synchronization). To us, this is not so relevant for economics, as economic fluctuations do not generally display clear periodicity. We instead turn to the other part of synchronization theory, that goes under the name of *complete* synchronization (also known as chaotic, or full, or identical synchronization). Under this theory, different units are described by the same parameters and so, under coupling, individual deterministic dynamics would completely align. However, desynchronizing forces such as chaos and noise tend to separate individual dynamics. Complete synchronization theory is an elegant mathematical formalism that makes it possible to quantify the relative strength of synchronizing and desynchronizing forces, through an eigenvalue-eigenvector decomposition that takes into account the network connecting the different units.

3.1 Master stability functions

Here, we build on the master stability approach originally proposed by Pecora and Carroll (1990, 1998). This starts by assuming that all agents are in a *synchronized state* in which they all behave alike. The key idea of the master stability approach is to perform a transformation that uncouples the deviations from the synchronized state into orthogonal components. The approach is semi-analytical, in the sense that the relative strength of synchronizing and desynchronizing forces is obtained by numerically computing Lyapunov exponents, but once these are computed, the key insights are obtained analytically. In the following, we first derive our key equations, and then provide some examples that clarify the working of our approach. To be as general as possible, we mostly focus on the abstract formulation of the model (Section 2.1).

Let $\boldsymbol{s}_t = (x^s_t, y^s_t)$ denote the synchronized state in which all agents behave alike. Let the dynamics of individual agents be given by $x_{i,t} = x^s_t + x^\xi_{i,t}$ and $y_{i,t} = y^s_t + y^\xi_{i,t}$, where $x^\xi_{i,t}$ and $y^\xi_{i,t}$ denote a small deviation from the synchronized state, for example due to idiosyncratic shocks hitting agent i. Let further $\boldsymbol{\xi}_t = \left(x^\xi_{1,t}, y^\xi_{1,t}, x^\xi_{2,t}, y^\xi_{2,t}, \dots x^\xi_{N,t}, y^\xi_{N,t}\right)$ be the $2N$-dimensional vector of deviations.

To show how the evolution of this vector can be obtained, we fully work out a specific example. Consider a star network with three nodes. This is composed by a central node, which we denote as node 1, that is linked to two leaf nodes, which we denote as nodes 2 and 3. The leaf nodes are only connected to the central node. We assume that the dependence on other nodes is $\epsilon_{ij} = \epsilon/k_i$, where k_i is the degree of node i, while the dependency on oneself is $\epsilon_{ii} = 1 - \epsilon$. This defines a weighted interaction network with adjacency matrix

$$\begin{pmatrix} 1-\epsilon & \epsilon/2 & \epsilon/2 \\ \epsilon & 1-\epsilon & 0 \\ 0 & \epsilon & 1-\epsilon \end{pmatrix}. \tag{5}$$

Replacing $x_{i,t} = x^s_t + x^\xi_{i,t}$ and $y_{i,t} = y^s_t + y^\xi_{i,t}$ in Eq. (1) with the interaction network above, we get the following system of equations for the evolution of $\boldsymbol{\xi}_t$:

$$\begin{cases} x^\xi_{1,t+1} = (1-\delta)x^\xi_{1,t} + y^\xi_{1,t}, \\ y^\xi_{1,t+1} = \alpha_1 x^\xi_{1,t} + (\alpha_2 + F'(y^s_t))\, y^\xi_{1,t} - \epsilon F'(y^s_t)\left(y^\xi_{1,t} - \frac{1}{2}y^\xi_{2,t} - \frac{1}{2}y^\xi_{3,t}\right), \\ x^\xi_{2,t+1} = (1-\delta)x^\xi_{2,t} + y^\xi_{2,t}, \\ y^\xi_{2,t+1} = \alpha_1 x^\xi_{2,t} + (\alpha_2 + F'(y^s_t))\, y^\xi_{2,t} - \epsilon F'(y^s_t)\left(y^\xi_{2,t} - y^\xi_{1,t}\right), \\ x^\xi_{3,t+1} = (1-\delta)x^\xi_{3,t} + y^\xi_{3,t}, \\ y^\xi_{3,t+1} = \alpha_1 x^\xi_{3,t} + (\alpha_2 + F'(y^s_t))\, y^\xi_{3,t} - \epsilon F'(y^s_t)\left(y^\xi_{3,t} - y^\xi_{1,t}\right), \end{cases} \tag{6}$$

where we have used the fact that the deviations $\boldsymbol{\xi}_t$ are small to Taylor-expand the function F to first order. This formulation suggests that the evolution of deviations of any single agent is given by the Jacobian of the dynamics corresponding to that agent, plus a term of interaction with other agents. Extrapolating from this example, it is possible to see that the evolution of the whole vector $\boldsymbol{\xi}_t$ is generically given by

$$\boldsymbol{\xi}_{t+1} = \left(\boldsymbol{I}_N \otimes \boldsymbol{J}(\boldsymbol{s}_t) - \epsilon F'(\boldsymbol{s}_t)\boldsymbol{KLK} \otimes \boldsymbol{H}\right)\boldsymbol{\xi}_t, \tag{7}$$

where:

- $\boldsymbol{I}_N$ is the N-dimensional identity matrix.
- $\boldsymbol{J}(\boldsymbol{s}_t)$ is the 2-dimensional Jacobian of the dynamical system.
- $\otimes$ denotes the Kronecker product, i.e. a $2N$-dimensional matrix.
- $F'(\boldsymbol{s}_t)$ is the first derivative of F, due to the fact that we are considering a first-order approximation around the synchronized state.
- $\boldsymbol{K}$ is an N-dimensional square matrix with $1/\sqrt{k_i}$ on the main diagonal and zero everywhere else.
- $\boldsymbol{L}$ is the Laplacian of the network. This is a key mathematical property of the network that is widely used in many applications. $L_{ii} = k_i$ and $L_{ij} = -1$ if i and j are connected. $\boldsymbol{KLK}$ is known as normalized Laplacian, and has similar properties to the Laplacian.[11]
- $\boldsymbol{H}$ is the 2-dimensional square matrix of connectivity in the dynamical system. Since here the only connectivity is through y, it is $H_{22} = 1$ and zero everywhere else.

The problem with Eq. (7) is that the evolution of each component $y^{\xi}_{i,t}$ depends on all other agents j to which i is connected. The key trick of the master stability approach is to diagonalize $\boldsymbol{KLK}$ so as to decompose the deviations $\boldsymbol{\xi}$ into orthogonal, uncoupled, components. Following the terminology in the literature, we call these components *eigenmodes*. Let

$$\boldsymbol{\xi}_t = \left(\boldsymbol{Q} \otimes \boldsymbol{I}_2\right)\boldsymbol{\zeta}_t, \tag{8}$$

where $\boldsymbol{Q}$ is the matrix of eigenvectors of $\boldsymbol{KLK}$. Here, $\boldsymbol{\zeta}_t$ can be interpreted as a projection of $\boldsymbol{\xi}_t$ in the eigenspace. Replacing Eq. (8) in Eq. (7), we see that $\boldsymbol{\zeta}_t$ evolves according to

$$\boldsymbol{\zeta}_{t+1} = \left(\boldsymbol{I}_N \otimes \boldsymbol{J}(\boldsymbol{s}_t) - \epsilon F'(\boldsymbol{s}_t)\boldsymbol{\Lambda} \otimes \boldsymbol{H}\right)\boldsymbol{\zeta}_t, \tag{9}$$

where $\boldsymbol{\Lambda}$ is the matrix with the eigenvalues of $\boldsymbol{KLK}$ on the main diagonal and zero everywhere else. Each of the N eigenmodes of $\boldsymbol{\zeta}_t$ can then be written as

$$\boldsymbol{\zeta}_{i,t+1} = \left(\boldsymbol{J}(\boldsymbol{s}_t) - F'(\boldsymbol{s}_t)\epsilon\lambda_i\boldsymbol{H}\right)\boldsymbol{\zeta}_{i,t}. \tag{10}$$

In this basis, the evolution of each eigenmode i only depends upon itself. We stress that i now is an eigenvector (corresponding to the eigenvalue λ_i of the normalized Laplacian of the network) and *not an agent*. As will be clarified in the examples below, the evolution of each eigenmode corresponds to higher-order properties of the network. Of course one can retrieve the dynamics of the agents by applying the transformation $\boldsymbol{\xi}_t = \left(\boldsymbol{Q} \otimes \boldsymbol{I}_2\right)\boldsymbol{\zeta}_t$.

The eigenmode $i = 1$ corresponds to the eigenvalue $\lambda_i = 0$, and describes dynamics parallel to the synchronization manifold (i.e., "in the same direction as the dynamics"). It

[11] Here we summarize a few properties of $\boldsymbol{KLK}$. Because the rows sum to zero, one eigenvalue is zero. Moreover, it is well known that the other eigenvalues are positive and bounded between 0 and 2. The multiplicity of the 0 eigenvalue reflects the number of disconnected clusters in the network. We sort the eigenvalues in increasing order, so that $\lambda_1 = 0$, $\lambda_2 > 0$ if the network is connected, $\lambda_3 \geq \lambda_2$, ..., $\lambda_N \geq \lambda_{N-1}$, $\lambda_N \leq 2$. The eigenvalue λ_2 is knows as *algebraic connectivity* and the corresponding eigenvector as *Fiedler vector*. The smaller λ_2, the more the network has a modular structure, in which two or more clusters of nodes have strong internal connectivity and weak external connectivity. In the case of two clusters, the components of the Fiedler vector are positive for nodes in a cluster and negative for nodes in the other cluster. Fiedler vectors are commonly used for graph partitioning.

captures the phase shift due to a shock hitting the system at a given time t, as will be clarified below. The eigenmodes $i > 1$ correspond to dynamics orthogonal to the synchronization manifold. If these dynamics always converge to zero after the shock, the synchronized state is stable.

Whether the orthogonal dynamics converge to zero is not obvious from Eq. (10). To know if they do, one must numerically compute the Lyapunov exponents of the system. Letting $K = \epsilon\lambda_i$ denote an effective coupling for eigenmode i, one calculates Lyapunov exponents μ in Eq. (10) for all values of K that are typically obtained. (In our case, because $\epsilon \in [0, 1]$ and $\lambda_i \in [0, 2]$, it is of interest to only consider the interval $K \in [0, 2]$.) When Lyapunov exponents are negative, the corresponding eigenmodes die out over time.

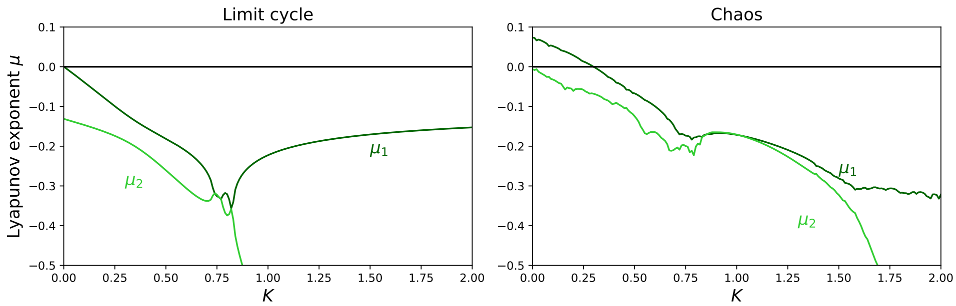

Figure 4: Lyapunov exponents μ_1 and μ_2 as a function of the effective coupling K, both for a parameter combination that leads to a limit cycle and for a parameter combination that leads to chaotic dynamics.

The Lyapunov exponents are shown in Figure 4, for the baseline parameterizations of the model that lead to limit cycles and chaotic dynamics (see footnote 10). There are two Lyapunov exponents because the dynamical system is 2-dimensional. Under limit cycles, both Lyapunov exponents μ_1 and μ_2 are always negative, except for $K = 0$ in which case $\mu_1 = 0$. This means that all eigenmodes $i > 1$ die out, making the synchronized state stable. In other words, if countries followed deterministic limit cycles and were only occasionally hit by idiosyncratic shocks, their business cycles would perfectly synchronize some time after each shock. Under chaotic dynamics, there is a region under which $\mu_1 > 0$, meaning that, even without exogenous shocks, under some conditions countries' business cycles would not synchronize.

3.2 Examples

To clarify the formalism introduced in the previous section, we now consider three examples. The first two examples still use the abstract framework for business cycles with a simple interaction network, while the last example uses real-world countries connected by the international trade network. For simplicity, in all cases we focus on limit cycles rather than chaotic dynamics, but the following analysis is analogous in the chaotic case. Throughout this section, for clarity we study the effect of single idiosyncratic shocks hitting the agents in the network at a certain time τ, and study its propagation. In Section 4, we will let shocks continuously hit agents, as can be expected in the real world.

3.2.1 Two connected agents

We start from the simplest possible system to study the interplay between dynamics and shocks in most detail. We consider two connected agents, agent 1 and agent 2, giving weight $1 - \epsilon$ to their internal dynamics and ϵ to the dynamics of the other agent. We simulate the two agents until they reach the synchronized state (which they always do given that we consider limit cycle dynamics), and then hit agent 1 with a positive shock. We first show the numerical simulations, and then make sense of the results using the formalism developed in Section 3.1.

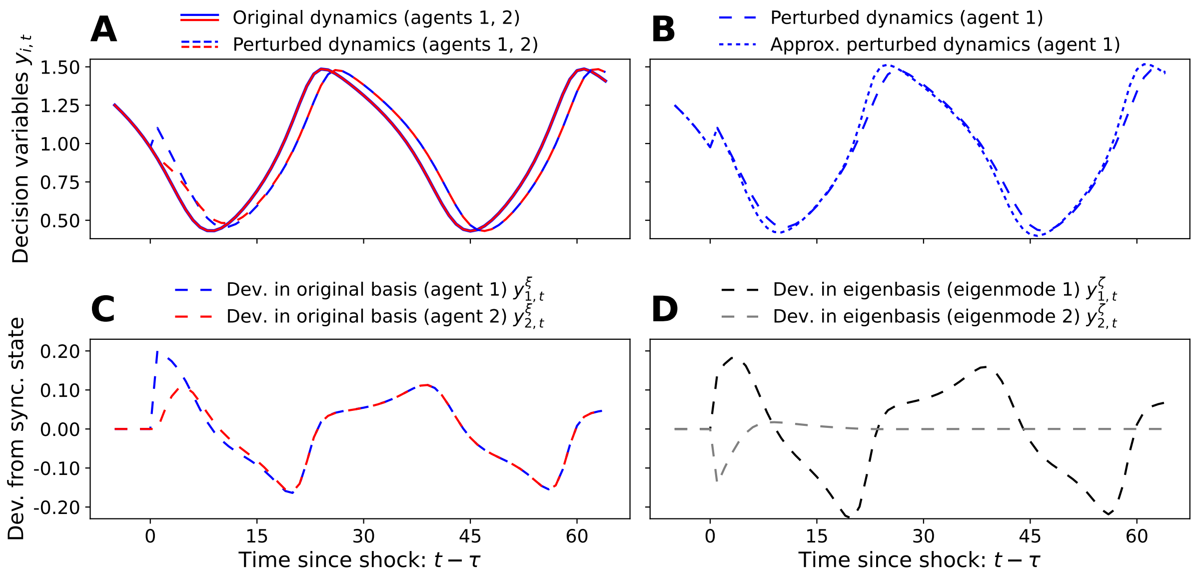

Figure 5: Example dynamics: two connected agents are synchronized in a limit cycle, and at time τ agent 1 gets hit by a positive shock. We show the resulting dynamics at time $t > \tau$. A: Original (unshocked) and perturbed dynamics of agent 1 (blue) and agent 2 (red). B: Perturbed dynamics of agent 1 (as in panel A) compared with the linearized approximation $y^s_t + y^{\xi}_{i,t}$. C: Deviations from the synchronized state in the original basis, for both agent 1 ($y^{\xi}_{1,t}$) and agent 2 ($y^{\xi}_{2,t}$). D: Deviations from the synchronized state in the eigenbasis, for both eigenmode 1 ($y^{\zeta}_{1,t}$) and eigenmode 2 ($y^{\zeta}_{2,t}$).

Figure 5 shows the dynamics in the 60 time steps following the positive shock to agent 1. Panel A compares the original (unshocked) to the perturbed dynamics, for both agents 1 and 2. It is clear that in the perturbed dynamics the two agents quickly reach synchronization again, but their business cycles permanently lead those of the original dynamics by a few time steps. Next, panel B shows how good the approximation of Eq. (7) is. In this panel, the perturbed dynamics of agent 1 (same as previously shown in panel A) is compared to the linearized approximation $y^s_t + y^{\xi}_{i,t}$. We see that that the exact and approximate dynamics are not identical, especially at the peaks and troughs, but overall the approximation is pretty faithful.

Having established this, we focus on the dynamics of the shocks in the original basis and in the eigenbasis. In the original basis (panel C), the terms $y^{\xi}_{1,t}$ and $y^{\xi}_{2,t}$ represent the deviations of agent 1 and agent 2 from the synchronized state. We see that, after agent 1 is hit by a positive shock at τ and deviates positively from the synchronized state, also agent 2 deviates positively from the synchronized state after a few time steps. Then, after a transient, the deviations from the synchronized state become identical across the two agents and start following a limit cycle. In the eigenbasis (panel D), the deviations $y^{\zeta}_{1,t}$ and $y^{\zeta}_{2,t}$ correspond to the first and second eigenmode. While the first eigenmode starts following a limit cycle similar to the one in panel C, the second eigenmode converges to zero after a negative and then slightly positive fluctuation.

To make sense of these results analytically, we now consider the master stability formalism. For this basic network, the eigenvalues of $\boldsymbol{KLK}$ are $\lambda_1 = 0$ and $\lambda_2 = 2$. Looking at Figure 4 (left panel), the first eigenmode, corresponding to $\lambda_1 = 0$ and so $K = 0$, has largest Lyapunov exponent $\mu_1 = 0$. Thus, from Eq. (10), it is $\boldsymbol{\zeta}_{1,t+1} = \boldsymbol{J}(\boldsymbol{s}_t)\boldsymbol{\zeta}_{1,t}$: the dynamics of the first eigenmode are always described by the Jacobian of the synchronized state (black line in Figure 5D). By contrast, the second eigenmode, corresponding to $\lambda_2 = 2$ and so $K = 0.6$,[12], has largest Lyapunov exponent $\mu_1 \approx -0.2$. This means that it exponentially converges to zero, approximately by 20% every time step (grey line in Figure 5D).

[12]In this example, we take $\epsilon = 0.3$.

The eigenvectors corresponding to the two eigenmodes are $\boldsymbol{v}_1 = (1,1)/\sqrt{2}$ and $\boldsymbol{v}_2 = (-1,1)/\sqrt{2}$. Thus, at time t the deviations from the synchronized state in the original basis and in the eigenbasis are related by the following equations:

$$\begin{aligned} y^{\xi}_{1,t} &= \tfrac{1}{\sqrt{2}}\left(y^{\zeta}_{1,t} - y^{\zeta}_{2,t}\right), \qquad & y^{\zeta}_{1,t} &= \tfrac{1}{\sqrt{2}}\left(y^{\xi}_{1,t} + y^{\xi}_{2,t}\right), \\ y^{\xi}_{2,t} &= \tfrac{1}{\sqrt{2}}\left(y^{\zeta}_{1,t} + y^{\zeta}_{2,t}\right), \qquad & y^{\zeta}_{2,t} &= \tfrac{1}{\sqrt{2}}\left(-y^{\xi}_{1,t} + y^{\xi}_{2,t}\right). \end{aligned} \tag{11}$$

The second column in Eq. (11) helps to interpret what the eigenmodes represent. The first eigenmode is proportional to the sum of the deviations of the two agents. As it never converges to zero given the discussion above, it represents the permanent phase shift in the dynamics due to the shock, also described in the literature as a perturbation that is "parallel to the synchronization manifold". The second eigenmode, which is "transverse to the synchronization manifold", is proportional to the difference of the deviations of the two agents from the synchronized state. As it eventually converges to zero, this suggests that the two agents always reach sychronization again after the shock. This can also be seen from the first column in Eq. (11): when $y^{\zeta}_{2,t} = 0$, it is always $y^{\xi}_{1,t} = y^{\xi}_{2,t} = y^{\zeta}_{1,t}/\sqrt{2}$.

We can use this example to study analytically how other shocks would impact the dynamics. That is, we consider generic u_1 and u_2 hitting the agents at time τ, so that $y^{\xi}_{1,\tau} = u_1$ and $y^{\xi}_{2,\tau} = u_2$. From the second column of (11) we can see that if $u_1 = u_2 = u$, $y^{\zeta}_{2,\tau} = 0$, and so for all following time steps $t \geq \tau$, $y^{\zeta}_{2,t} = 0$: The eigenmode transverse to the synchronization manifold never takes positive values. In other words, because the shock is common, it does not cause any relative difference in phase between the two agents. On the contrary, the first eigenmode is maximal, $y^{\zeta}_{1,\tau} = \sqrt{2}u$: the effect of the shock is to only shift the phase, by the maximum amount. If $u > 0$, the phase is shifted forward; if $u < 0$, it is shifted backwards. Consider now the reverse case $u_1 = -u_2 = u$. Now $y^{\zeta}_{1,\tau} = 0$ and $y^{\zeta}_{2,\tau} = \sqrt{2}u$. In this case, the transverse eigenmode is maximal, while the parallel eigenmode is null. So, after a transient in which the phases are different, they synchronize again, with no consequence of the shock in terms of permanent phase shift (this is the only case which does not imply a permanent phase shift).

3.2.2 Six agents with two cliques of three nodes each

To understand the effect of the network structure on shock propagation and synchronizing non-linear dynamics, we consider the network in Figure 6. This network is composed of two cliques of three nodes each, connected by a link between two of the nodes of the cliques. These two nodes are colored lighter to highlight the stronger connectivity to the other clique. This network can be thought of as a stylized model of six countries in two continents.

At time τ, we apply the following shock vector:

$$\left(y^{\xi}_{1,\tau}, y^{\xi}_{2,\tau}, y^{\xi}_{3,\tau}, y^{\xi}_{4,\tau}, y^{\xi}_{5,\tau}, y^{\xi}_{6,\tau}\right) = (0.20, 0.12, 0.16, -0.12, -0.24, 0.00). \tag{12}$$

This shock vector can be thought of as the combination of idiosyncratic shocks hitting all agents differently, combined with a positive shock hitting the clique with agents 1,2 and 3, and a negative shock hitting the other clique. In the eigenbasis, this shock vector corresponds to

$$\left(y^{\zeta}_{1,\tau}, y^{\zeta}_{2,\tau}, y^{\zeta}_{3,\tau}, y^{\zeta}_{4,\tau}, y^{\zeta}_{5,\tau}, y^{\zeta}_{6,\tau}\right) = (0.05, 0.35, 0.00, 0.06, -0.08, -0.17). \tag{13}$$

It should be noted that the shock in the second eigenmode (0.35) is the largest one.

The evolution of the shocks is shown at the bottom of Figure 6. In the panel on the left, we show the shocks in the original basis, i.e. each line corresponds to an agent. It is clear that the dynamics vary across the two cliques, but they are very similar within each of the two cliques. The cyan and orange lines, corresponding to the two "bridge" agents, are the ones that show the closest dynamics to the ones in the other clique. Over time, all dynamics converge again to the synchronized state, with a permanent phase shift (not visible in the figure).

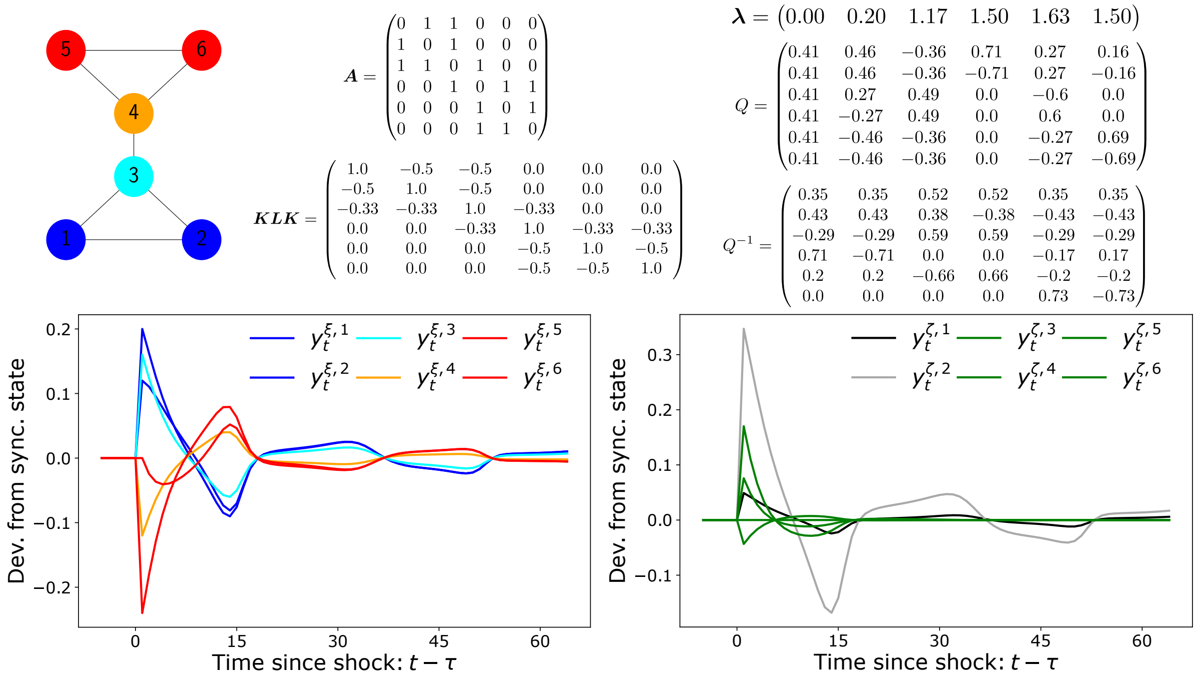

Figure 6: Example dynamics in a network with six agents and two cliques of three nodes each. The top part of the figure shows the network, the adjacency matrix $\boldsymbol{A}$ and the normalized Laplacian $\boldsymbol{KLK}$. It also shows the eigenvalues $\boldsymbol{\lambda}$, the eigenvectors $\boldsymbol{Q}$ of the normalized Laplacian (each column of $\boldsymbol{Q}$ is the eigenvector corresponding to the eigenvalue shown above) and the inverse matrix of eigenvectors $\boldsymbol{Q}^{-1}$. The bottom part of the figure shows the evolution of the deviations in the original basis and in the eigenbasis.

The evolution of the shocks in the eigenbasis is shown in the right panel. The black line is the eigenmode which is parallel to the synchronization manifold. The grey line corresponds to the second eigenmode, the one with largest initial value, and decays very slowly. The green lines correspond to the third to sixth eigenmodes, and decay very quickly.

The interpretation for the eigenmodes is now clear. The second eigenmode corresponds to synchronization across cliques. Indeed, because $y^{\zeta}_{2,t} = 0.43y^{\xi}_{1,t} + 0.43y^{\xi}_{2,t} + 0.38y^{\xi}_{3,t} - 0.38y^{\xi}_{4,t} - 0.43y^{\xi}_{5,t} - 0.43y^{\xi}_{6,t}$, it is maximal in absolute value when shocks are positive in one clique and negative in the other one. This eigenmode takes time to decay, because $K = \epsilon\lambda_2 = 0.3 \cdot 0.2 = 0.06$ is relatively small and so, as shown in Figure 4, the largest Lyapunov exponent is not very negative ($\mu_1 \approx -0.02$). This means that synchronization across cliques takes time, as the eigenmode decays by approximately 2% every time step. The third to sixth eigenmodes correspond to synchronization within cliques, and decay quicker ($K = \epsilon\lambda_i$, $i = 3, \ldots 6$ ranges from 0.35 to 0.49, so the largest Lyapunov exponent μ_1 ranges between -0.14 and -0.18). This indicates that synchronization within cliques happens much quicker.

This phenomenon is general, because the second eigenvector of the Laplacian—the Fiedler vector—takes positive values in a cluster of nodes, and negative values in the other one. If there were M cliques, or more loosely M clusters of highly connected nodes weakly connected across one another, there would be $M-1$ small eigenvalues and then a large gap to the next eigenvalue. The eigenmodes corresponding to the small eigenvalues would indicate, as in the case above, synchronization across clusters. The remaining $N - M$ eigenvalues would be much larger, and the associated eigenmodes would correspond to synchronization within clusters.[13]

[13] There could also be a hierarchy of clusters. Imagine for example that there were two main clusters, and two sub-clusters within each cluster. Here, the Fiedler vector would divide between the two main clusters, but then the following eigenmode (corresponding to a larger eigenvalue) would correspond to synchronization between sub-

3.2.3 Real-world countries linked by trade

To conclude this section, we now discuss a real-world example. We consider the specification of the model in which nodes are countries and links represent international trade. While we defer the reader to Section 4.1.1 for a description of the dataset, here we just discuss the insights that can be gained by applying the theory of complete synchronization to these data.

Due to the way in which we built the networks in the examples above, the analysis was performed for unweighted, undirected networks. However, nothing prevents to apply it to the case of directed, weighted networks described by a generic connectivity matrix $\boldsymbol{W}$ (so that all rows are normalized to one): instead of calculating the eigenvalues of $\boldsymbol{KLK}$ one calculates those of $\boldsymbol{I}_N - \boldsymbol{W}$.[14]

For illustration, we consider year 2011, which is the most recent year available in the data that we use. A visualization of the network is given later (Figure 9). Here, we just note that it is a network that shares certain properties with the one in Section 3.2.2: each node has strong self-loops (representing internal demand) and larger trade links with countries within the same continent than with countries in other continents.

However, this distinction is not as clear-cut as in the example in the previous section, and this is reflected in the spectrum. Indeed, these are the first 10 eigenvalues: 0; 0.072; 0.072; 0.081; 0.098; 0.098; 0.111; 0.116; 0.127; 0.147. We see that there is not a big gap between any eigenvalue, as there was in Section 3.2.2 between the second and the third eigenvalue. Thus, after shocks, no eigenmode converges back to zero much faster than the others. Yet, some community structure clearly exists in the data and is picked by our analysis.

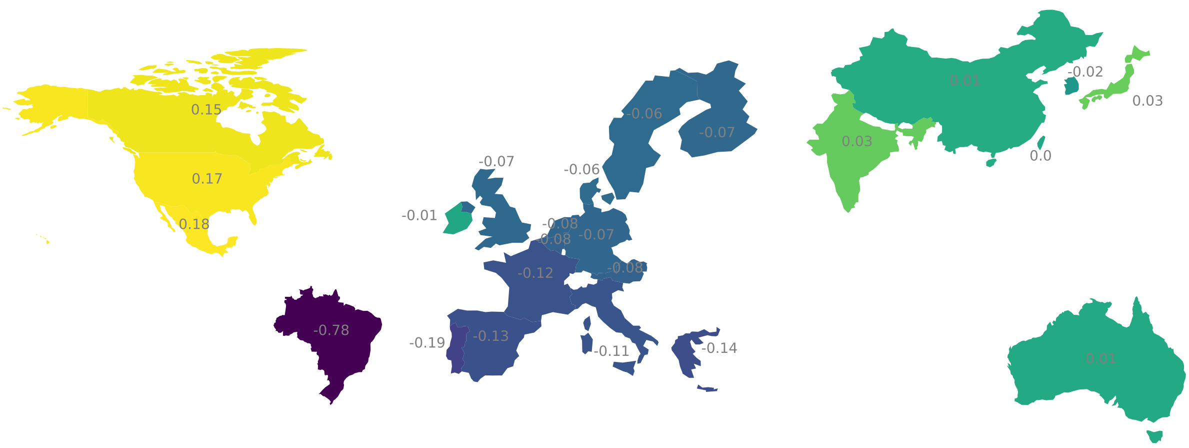

Figure 7: Countries in our sample colored by their component of the Fiedler vector (also annotated).

To show this, Figure 7 represents the countries in our sample colored by their component of the Fiedler vector. This is commonly used in graph partitioning because it is an approximate solution to the problem of dividing a graph into two components such that the sum of the links between the two components has minimal weight. When partitioning a graph using the Fiedler vector, one cuts where the vector is zero, putting all nodes with

clusters. There would be a time-scale separation so that synchronization would first occur within sub-clusters, then across sub-clusters within the same main cluster, and finally across main clusters.

[14] The networks considered in the previous examples have the convenient mathematical property that all eigenvalues of the normalized Laplacian are real, positive and bounded between zero and two. When considering a generic interaction network, the eigenvalues could instead be complex. Indeed, in the example considered in this section we find that the eigenvalues have a very small imaginary part. Applying master stability analysis in the case of complex eigenvalues is almost identical to the case of real eigenvalues, except that one computes complex Floquet exponents (instead of real Lyapunov exponents) and then verifies if the complex values of $K = \epsilon\lambda_i$ are inside the stable region in the complex plane (Pecora and Carroll, 1998). Here, for simplicity we just show the real part of the eigenvalues and eigenvectors.

positive components in one cluster and all nodes with negative components in the other cluster. Components that are close to zero in either the positive or negative direction are less clearly in one cluster than components that are far from zero.

Looking at Figure 7, we see that Brazil's component of the Fiedler vector is strongly negative, and no other country has so strongly negative values.[15] Then, European countries have components ranging between -0.19 and -0.06 (with the exception of Ireland, which has -0.01), with Southern European countries having more negative values than Northern European countries. Next, Asian countries have slightly positive components (with the exception of South Korea which is slightly negative). Finally, the components associated to North American countries range between 0.15 and 0.18. Thus, under graph partitioning one cluster would mostly comprise Brazil and Europe, and the other Asia and North America.

In conclusion, the community structure discovered by our spectral analysis, combined with the theory of complete synchronization that we developed, makes it possible to predict analytically which countries are more likely to experience higher synchronization when each country follows an endogenous business cycle and is hit by idiosyncratic shocks.

3.3 Comovement implications

Further to providing intuition into the mechanisms driving synchronization, the theory of complete synchronization is particularly useful to explain how synchronization of endogenous business cycles could improve on exogenous business cycle models in terms of generating higher comovement between macroeconomic time series.

Consider two countries following two deterministic identical limit cycles in isolation. When coupled, absent shocks, after some time their dynamics perfectly align in the synchronized state $\boldsymbol{s}_t = (x_t^s, y_t^s)$. Suppose now that at time τ, the decision variables $y_{1,\tau}$ and $y_{2,\tau}$ of the countries are hit by idiosyncratic shocks $y_{1,\tau}^{\xi}$ and $y_{2,\tau}^{\xi}$ respectively. The comovement between the time series of the decision variables, calculated for the time steps t immediately following the shocks, is given by the Pearson correlation coefficient

$$\operatorname{cor}\left(y_t^s + y_{1,t}^{\xi}, y_t^s + y_{2,t}^{\xi}\right) = \frac{\operatorname{cov}\left(y_t^s, y_t^s\right) + \operatorname{cov}\left(y_t^s, y_{2,t}^{\xi}\right) + \operatorname{cov}\left(y_{1,t}^{\xi}, y_t^s\right) + \operatorname{cov}\left(y_{1,t}^{\xi}, y_{2,t}^{\xi}\right)}{\operatorname{std}\left(y_t^s + y_{1,t}^{\xi}\right) \cdot \operatorname{std}\left(y_t^s + y_{2,t}^{\xi}\right)}. \tag{14}$$

In the above equation, cov denotes the covariance, std indicates the standard deviation, and we have used linearity of the covariance to decompose it in various terms. The term $\operatorname{cov}\left(y_t^s, y_t^s\right)$ indicates the perfect synchonization of the unperturbed deterministic dynamics; the terms $\operatorname{cov}\left(y_t^s, y_{2,t}^{\xi}\right)$ and $\operatorname{cov}\left(y_{1,t}^{\xi}, y_t^s\right)$ indicate the comovement between the deterministic dynamics and the deviations due to the shocks; and the term $\operatorname{cov}\left(y_{1,t}^{\xi}, y_{2,t}^{\xi}\right)$ indicates the correlation of the deviations (i.e., shock propagation). Compare this equation with the correlation between time series of two countries that are in the same stable steady state and are hit by shocks $(y_{1,t}^{\xi}, y_{2,t}^{\xi})$:

$$\operatorname{cor}\left(y_{1,t}^{\xi}, y_{2,t}^{\xi}\right) = \frac{\operatorname{cov}\left(y_{1,t}^{\xi}, y_{2,t}^{\xi}\right)}{\operatorname{std}\left(y_{1,t}^{\xi}\right) \cdot \operatorname{std}\left(y_{2,t}^{\xi}\right)}. \tag{15}$$

If we assume that the variance of endogenous and exogenous fluctuations is the same, the correlation coefficient is only determined by the numerators of these expressions. Comparing Eq. (14) and Eq. (15), it is clear that the endogenous component of dynamics potentially increases correlation. In particular, when the variance of the shocks is small relative to the variance of y_t^s, the correlation coefficient tends to one when business cycles are endogenous.

[15]This is presumably because Brazil trades a lot with other South American countries, which are included in our sample as "Rest of the World" and not shown here.

4 Empirical application

In the following, we perform simulations to show that, under our model, only comovement produced by endogenous fluctuations is high enough to match the level of comovement that can be found in the data. We first describe the data in Section 4.1, then show our main result in Section 4.2, and finally demonstrate that it is robust to various alternative assumptions in Section 4.3.

4.1 Data

To perform the empirical exercise, we need two main ingredients. First, we need to calibrate the interaction coefficients of the model (Section 4.1.1). Second, we must measure the empirical level of comovement (Section 4.1.2).

4.1.1 Network data

In the international trade network interpretation of the model (Section 2.2), the interaction coefficients between countries represent the share of a country's output that is exported to another country. Importantly, the model also features "self-loops", representing the share of a country's output that remains within the same country. This internal demand, coming from domestic firms and households, is generally much higher than international demand. This makes it possible to set the problem as one of endogenous business cycles that originate within countries and synchronize internationally, as opposed to a problem in which business cycles originate *because of* international linkages (as previously stressed, independent oscillations are a necessary ingredient to apply synchronization theory). Further to the importance of internal demand, another key feature of the model is that countries demand fixed proportions of imports from other countries. These two features of the model are crucial to determine which data should be used to calibrate the interaction network.

To have a consistent framework that includes information on internal demand, exports and imports, we turn to international input-output tables. Although we do not disaggregate our model at the industry-country level and do not distinguish between intermediate and final products, we prefer to use international input-output data over computing internal demand from raw import and export bilateral data and separately collected gross output data. Moreover, to test how good the assumption of fixed proportions of imports is, we need international input-output data that span as many decades as possible.

For these reasons, the ideal dataset is a merge of the recently released Long-Run World Input-Output Database (LR-WIOD), covering the period 1965-2000 (Woltjer et al., 2021), with the established 2013 release of the WIOD, spanning the years 1995-2011 (Timmer et al., 2015). Although the creators of these databases warn against using them together at a detailed industrial level because of differences in their construction (Woltjer et al., 2021), they say that "aggregate levels and trends are likely to be comparable". Selecting the countries that are available in both datasets, we end up with 24 countries and a Rest of the World aggregate.

Figure 8 shows the dynamics of the share of internal demand for all the countries we consider. We show both data from LR-WIOD and data from WIOD 2013, including data from both datasets in the period 1995-2000. This figure has three take-away messages.

First, the trends in the overlapping period do indeed look similar across the two datasets. This justifies our choice to merge the two datasets and supports the hypothesis of Woltjer et al. (2021). To perform the merging, we use LR-WIOD up to 1994 and WIOD 2013 from 1995 on, after adjusting the WIOD 2013 data so that the levels in 1995 match with those of LR-WIOD (note however that the levels were very close even without this transformation). Second, except for some irregular high-frequency movement, the share of internal demand changes slowly, generally decreasing. Similarly, we find that the fraction of bilateral exports over countries' output changes slowly over time (not shown), compared to business cycle frequency. We think that these results support our assumption of assuming fixed import shares as a first approximation. We will show that our results are robust to assuming

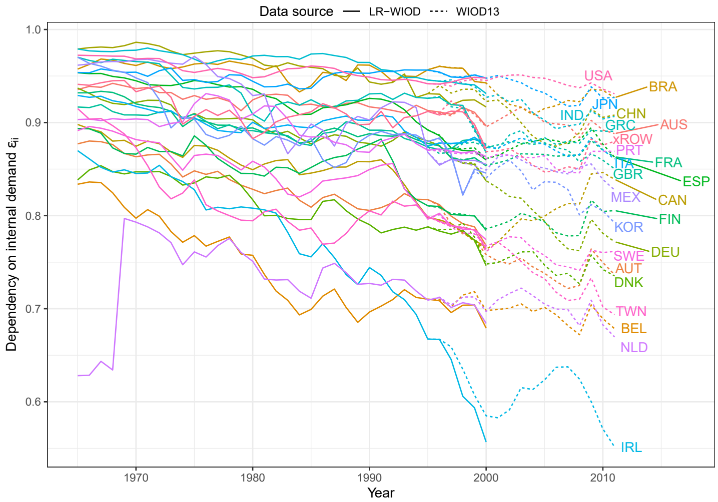

Figure 8: Dependency on internal demand (self-loops), 1965-2011, for all countries in our sample. Solid lines are taken from LR-WIOD, whereas dashed lines are taken from WIOD 2013.

time-varying interaction coefficients (Section 4.3.2). Third, the share of internal demand is almost always above 0.7, except for Ireland in recent years. This supports our view that business cycles originate from strategic complementarities of demand within countries, and then synchronize through international linkages, rather than originating from international demand.

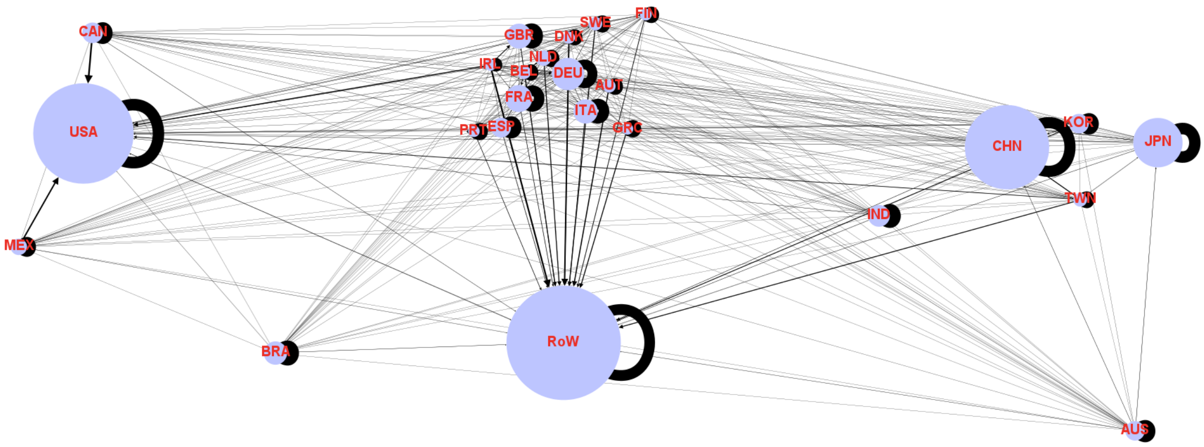

Figure 9: Interaction network in 2011. The position of the node is equal to the countries' geographical centroid. Node size is proportional to countries' gross output, while link size is proportional to interaction coefficients ϵ_{ij} (for visualization purposes, we applied nonlinear transformations to both node and link size). RoW indicates the Rest of the World.

We show an example of the interaction network ϵ_{ij} in Figure 9, for the most recent year available in our dataset.

4.1.2 Time series data

Depending on the dynamics produced by the model (Section 2.3) and the type of idiosyncratic shocks, the simulated level of comovement across countries varies. To check which specification comes closest to the level of comovement that can be observed in the data, we must measure empirical comovement in the first place. This is not a trivial task. Papers measuring business cycle comovement have been using all sorts of variables that measure economic activity and all sorts of filters to detrend the corresponding time series (De Haan et al., 2008).

Here, we consider annual data on employment and value added to measure economic activity. We obtain these data from the Penn World Tables, version 10.0 (Feenstra et al., 2015). More specifically, we use the number of persons engaged as a measure of employment, and real GDP at constant 2017 national prices to measure value added. For the countries considered in this paper, data are available from 1950 to 2019.

We cannot use these raw data to measure comovement, because they show clear common trends that would bias upward the estimate of comovement at business cycle frequencies.[16] To address this issue, we consider various modifications. First, we may divide employment or value added by total or working age population.[17] Second, we filter the resulting data so as to remove long-run fluctuations from the sample. To do so, we use two filters. One is the well-known Hodrick-Prescott filter, with the standard parameter $\lambda = 100$ that is used for annual data. The other is the band-pass filter proposed by Christiano and Fitzgerald (2003), alternatively keeping fluctuations between 2-15 or 2-25 years.[18] For all these options for the filters, we either consider the cyclical component or the ratio between the cyclical and trend components. In conclusion, from all these combinations we get 18 alternative detrending procedures for both employment and value added.

We find that the mean pairwise Pearson correlation among countries' employment is 0.24, with a standard deviation of 0.02 across detrending procedures. When considering value added, the correlation is 0.34, with a 0.05 standard deviation. These results are in line with estimates in the literature (Kose and Yi, 2006).

4.2 Main result

We now check whether endogenous business cycles are indeed necessary to reproduce the empirical level of comovement, or if instead shock propagation in an exogenous business cycles specification is sufficient. To do so, we simulate the model under different parameterizations, and compare the comovement of the decision variables y to empirical comovement as measured in the previous section. We perform this comparison at the aggregate, country and country-pair level.

We simulate the model under the three sets of parameters that, in a deterministic setting, produce limit cycles or a stable steady state (of type focus or node). We choose the parameters in footnote 10 as a baseline, but we find that our results are robust against alternative parameter combinations (Section 4.3.3). Moreover, we combine the deterministic component of the dynamics with an AR(1) shock process with persistence parameter $\rho = 0.3$ and standard deviation σ varying between 0.01 and 0.2. (We confirm the robustness to other values of ρ in Section 4.3.1.) We simulate the model for 280 time steps after removing an initial transient. Following the discussion in Section 2.4, time steps should be interpreted as quarters, so this leads to a 70-years period that is as long as the data we compare against (1950-2019). To calibrate the interaction network, we choose the international trade data in year 1990, which is in the middle of the available data (we show that this assumption almost does not affect the results in Section 4.3.2).

Figure 10 shows our key result. When the model produces limit cycles, the correlation coefficient obtained by averaging across pairwise country correlations and simulation runs ranges from 0.99 in the case of idiosyncratic shocks with standard deviation 0.01 (or 1% of the steady state value of the decision variables) to 0.03 when the standard deviation of idiosyncratic shocks is 0.2 (20% of the steady state). By contrast, when the model produces

[16]The non-stationarity of these time series is confirmed by an augmented Dickey-Fuller test for almost all countries.

[17]Total population is available in the Penn World Tables, while working age population is downloaded from the World Bank.

[18]We prefer the filter by Christiano and Fitzgerald (2003) over the Baxter and King (1999) filter because the latter loses some data points at the extremes of the series.

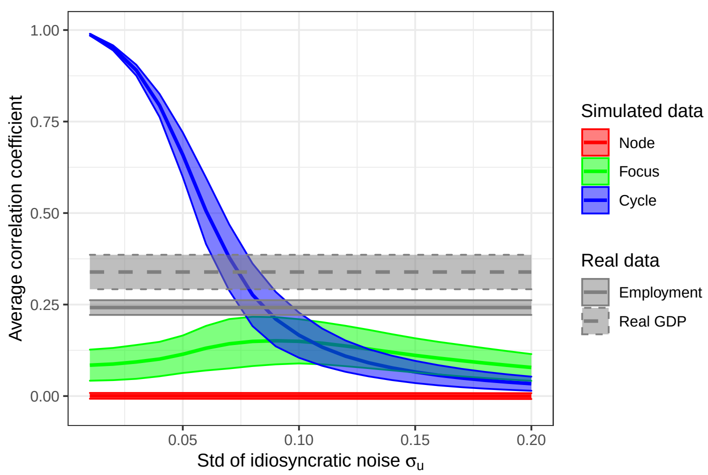

Figure 10: Mean Pearson correlation across the decision variables y of the countries in our sample as a function of the standard deviation of idiosyncratic shocks, in the model and in the data. For the simulations, lines represent the mean across 100 replications, and error bands represent one standard deviation. For the data, lines represent the mean across 18 detrending procedures, and error bands represent one standard deviation (Section 4.1.2).

fluctuations that converge to a stable steady state, the average correlation coefficient is never larger than 0.15, for any value of the standard deviation of idiosyncratic shocks. More specifically, when the steady state is a focus, the average correlation varies between 0.08 and 0.15, whereas when the steady state is a node it is close to zero. As the empirical counterpart of this measures is 0.24 when considering employment and 0.34 when considering value added, only in the limit cycle case the model can reproduce this level of comovement. Doing a t-test for the difference in means between simulated and real data, we reject the hypothesis that the means are equal in all cases in which the steady state is a focus or node, with high statistical significance (t-statistic < -11). In conclusion, the model can only match the empirical level of comovement under endogenous business cycles, for some strength of idiosyncratic shocks, confirming the theoretical intuition in Section 3.3.

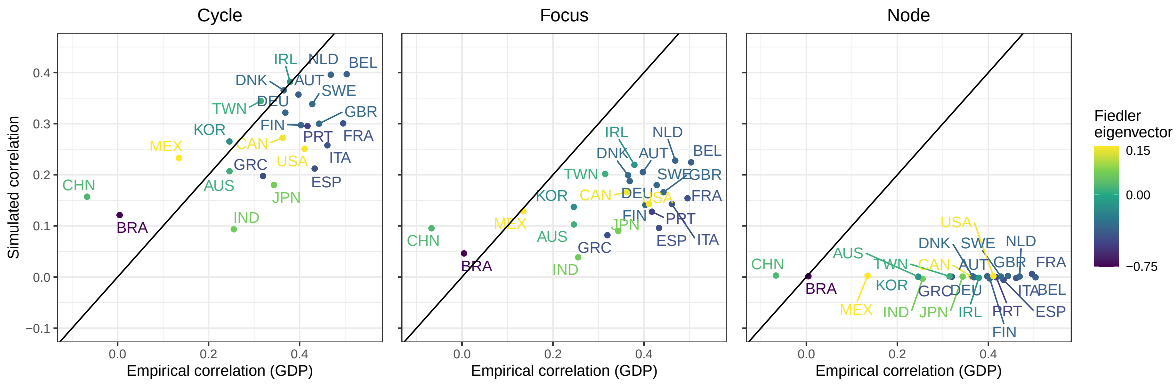

Figure 11: Mean Pearson correlation across the decision variable y of one country and the decision variables of all the other countries in our sample, in the model and in the data, averaged across 100 replications. The color code follows the position of the countries in the interaction network as exemplified by the Fiedler eigenvector, as in Figure 7. The identity line is plotted for reference.

Next, Figure 11 zooms in at the level of countries. Here, we select a standard deviation of idiosyncratic shocks of 0.08, under which the limit cycles model produces a correlation that is close to the data (Figure 10). This figure conveys three results. First, the higher comovement in the limit cycle case than in the focus and node cases is spread among all countries:

Simulated correlations decrease for every country. Second, in the limit cycle case there is a good correlation (Pearson 0.65) between empirical and simulated correlation coefficients. This means that countries whose dynamics are more correlated to other countries in the data are also more correlated to other countries in the model. For instance, the mean of the correlation coefficients of GDP between Belgium and the other countries in the sample is 0.45 in the data and 0.4 in the model, while the same quantity for Brazil is 0.0 in the data and 0.1 in the model. This good correlation between empirical and simulated data extends to the focus case, although it is slightly lower (Pearson 0.57), but it is completely absent when the steady state is a node (Pearson -0.20). The third result is that, in the limit cycle and focus cases, both empirical and simulated correlations are clearly linked to the position of the countries in the international trade network, as exemplified by the Fiedler eigenvector discussed in Section 3.2.3. Countries with high correlations tend to be have similar components in the eigenvector, whereas countries with more distant values such as Brazil and China tend to have lower average correlations.

These results are confirmed in Figure 12. In the data, pairwise correlations are high within Europe, North America and Asia, but they are lower across continents. The simulations, both in the cycle and focus cases, reproduce these patterns, although they fail to predict the few negative correlations that are observed empirically between Brazil and China and the rest of the countries.

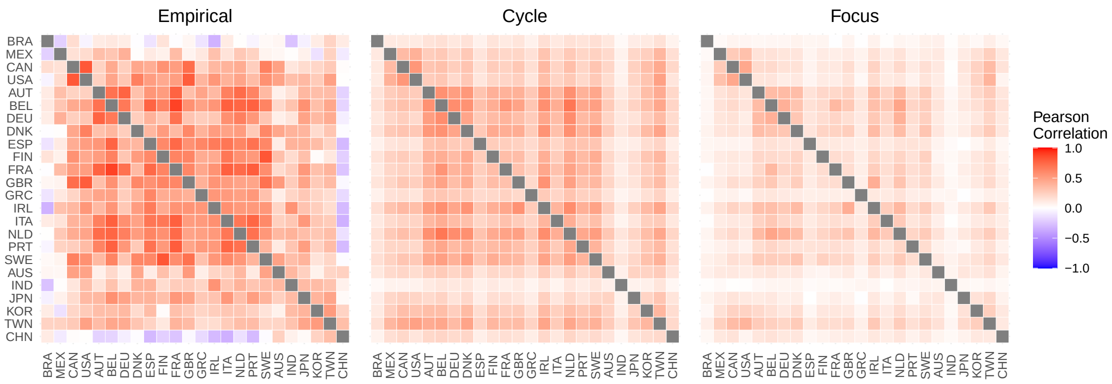

Figure 12: Pairwise Pearson correlation coefficients between the dynamics in different countries, both in the model and in the data, averaged across 100 replications.

In conclusion, the results in this section support the main claim of the paper, namely that endogenous business cycles are needed to explain the level of international comovement that can be found in the data, and show that the theory developed in Section 3 can be fruitful to interpret synchronization.

4.3 Robustness

We now test if the main result is robust to alternative specifications.

4.3.1 Shock persistence parameter

So far, we have used a value $\rho = 0.3$ for the autocorrelation of idiosyncratic shocks. What happens when this persistence parameter varies?

Figure 13 shows that, as we increase the persistence parameter, correlation coefficients become lower in the limit cycle case. This is not surprising, as higher persistence makes the shock processes lead the deterministic dynamics farther away from the values it would take in the absence of shocks. However, for all values of ρ, there is a value of σ that makes the limit cycle model match the empirical level of comovement. By contrast, in the focus and node cases there is not a big effect of the persistence parameter, and in no case the model can produce correlations that are as high as in the data.

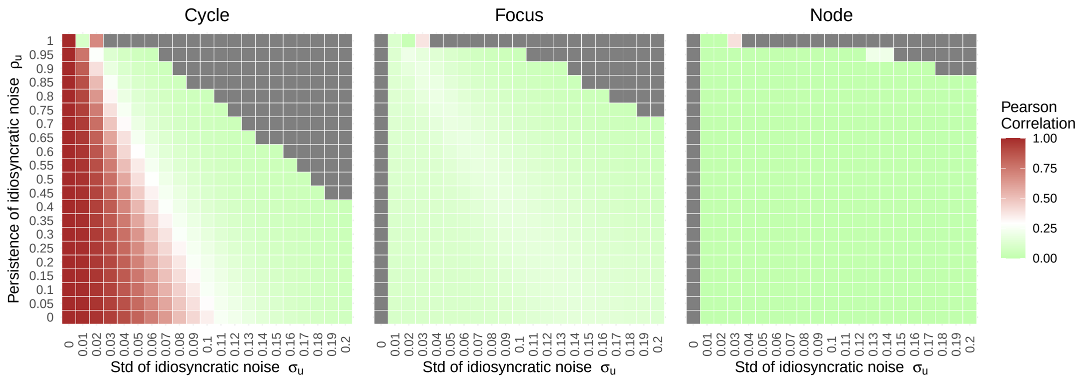

Figure 13: Heatmaps showing the mean Pearson correlation coefficients across all country pairs and 20 simulation runs, for several values of the persistence parameter ρ and the standard deviation of idiosyncratic shocks σ. The color bar is designed so that white corresponds to the mean of the empirical correlation of employment and value added. We show three heatmaps, for each of the main parameterization of the model.

4.3.2 Time-varying network

From the description of the model in Section 2.2, we have been assuming that the interaction coefficients ϵ_{ij} between country i and country j are fixed over time. This was in part supported by the discussion of the data in Section 4.1.1, where we argued that the share of a country's output that is exported to other countries varies slowly over time, at lower frequencies than typical business cycle frequencies. Here, we relax the assumption of fixed interaction coefficients, and directly input to the model time-varying interaction coefficients $\epsilon_{ij,t}$, with $t = 1965, \ldots, 2011$, corresponding to the available data.

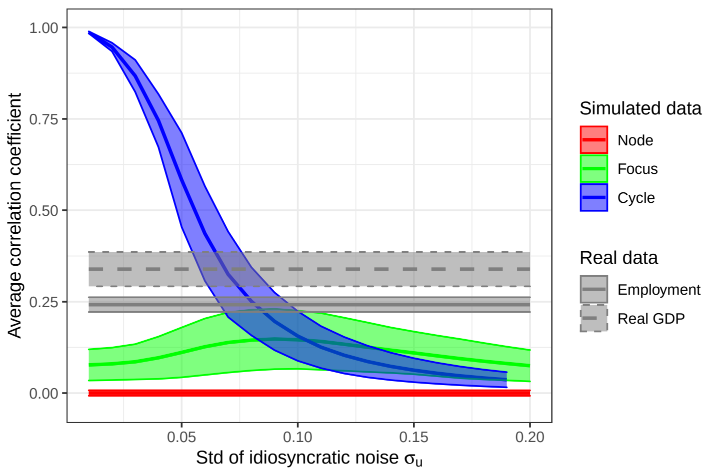

Figure 14: Mean Pearson correlation across the decision variables y of the countries in our sample as a function of the standard deviation of idiosyncratic shocks, in the model and in the data. The interpretation is the same as in Figure 10.

As we can see in Figure 14, the results are virtually unchanged from Figure 10, showing that indeed assuming fixed interaction coefficients is a good approximation for this problem.

4.3.3 Other parameter values

Finally, we test if our results are also valid for parameters that are different from the baseline parameters indicated in footnote 10. To do so, we check for alternative values of strategic complementarities at the steady state by randomly picking the β parameters. We restrict the parameter values so that they do not produce multiple equilibria and do not exit the stable region in Figure 2 from below, as this would lead to a flip bifurcation and a limit cycle of period 2, which is unlikely to be relevant in economic analysis (Beaudry et al., 2015).

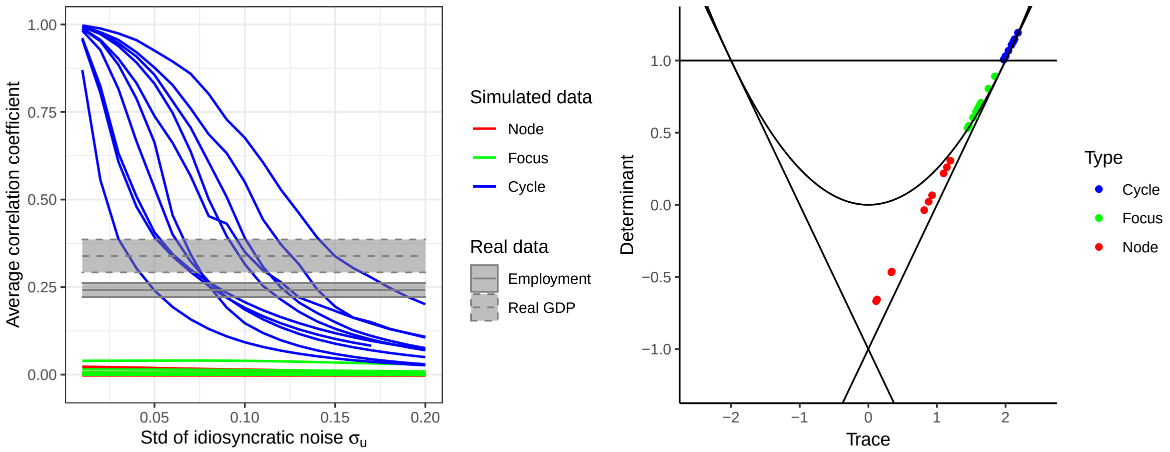

Figure 15: Left: Same as in Figures 10 and 14, except that we only show the mean across 20 simulation runs (and not the error bands) and 10 combinations of the β parameters for each of the cycle, focus and node cases. Right: same diagram as in Figure 2, for the same combinations of the β parameters shown on the left.

As we show in Figure 15 (left panel), the key results are robust to alternative parameter specifications, as only in the limit cycle case the model can produce correlations that are as high as in the data. However, depending on the specific parameterization, the general level of comovement in the limit cycle case can be higher or lower for a given strength of idiosyncratic shocks. This is because, depending on the parameters, limit cycle fluctuations have different amplitude, and so can be more or less perturbed by idiosyncratic shocks with a given standard deviation. It is also possible to note that we do not see a pattern for the focus case that is similar to Figure 10, as in this case the maximal correlation is 0.04, much smaller than 0.15 as in Section 4.2. This is because, in that case, we selected the parameters so that the focus was very close to the boundary of the unstable region, and so oscillations driven by exogenous shocks would die out very slowly. This is not the case for the parameter combinations that we select here in the focus case, which are well within the stable region (Figure 15, right panel).

5 Conclusion

In a popular science book, Krugman (1996) wrote: "One of the luxuries of a format like this one is that I can include the kind of loose speculations that I could never write in a journal and that I can explain, as I am doing now, that I do not necessarily believe in the theory I am advancing. So here is a crazy idea about the global business cycle: it is an example of "phase locking."[...] Like the two back-to-back clocks that started ticking in unison, the two economies would not need to be very strongly linked to develop a synchronized cycle; a modest linkage would do as long as they were predisposed to have cycles in any case and had fairly similar natural periods."

In this paper, we seriously explore this hypothesis, showing that it is highly promising to explain comovement of business cycles across countries. Differently from Krugman, however, we do not put much emphasis on phase locking and periodicity, which are problematic in the analysis of economic time series, but we rather develop a theory that explains synchronization of generic non-linear dynamics of nodes in an interaction network that are hit by exogenous shocks.

Our work can be extended in several ways. First of all, it would be interesting to apply the methods that we introduce in this paper to other models, for example microfounded business cycle models that attribute fluctuations to specific economic forces and microfounded shocks. In parallel, it would be interesting to explore other channels for synchronization, such as financial linkages, multinational corporate control, and animal spirits. Beyond macroeconomics, we hope that the tools that we offer in this paper can find wider applicability, as any disaggregate dynamical model that can be described by some form of non-linear dynamics can be studied with the tools of synchronization theory.

6 Bibliography

Arkolakis, C. and Ramanarayanan, A. (2009) "Vertical specialization and international business cycle synchronization," *Scandinavian Journal of Economics*, Vol. 111, pp. 655–680.

Backus, D. K., Kehoe, P. J., and Kydland, F. E. (1992) "International real business cycles," *Journal of Political Economy*, Vol. 100, pp. 745–775.

Barnett, W. A. and Chen, P. (1988) "The aggregation-theoretic monetary aggregates are chaotic and have strange attractors," in W. A. Barnett, E. R. Berndt, and H. White eds. *Dynamic Econometric Modeling*: Cambridge University Press, pp. 199–246.

Barnett, W. A. and Serletis, A. (2000) "Martingales, nonlinearity, and chaos," *Journal of Economic Dynamics and Control*, Vol. 24, pp. 703–724.

Baxter, M. and King, R. G. (1999) "Measuring business cycles: approximate band-pass filters for economic time series," *Review of economics and statistics*, Vol. 81, pp. 575–593.

Baxter, M. and Kouparitsas, M. A. (2005) "Determinants of business cycle comovement: a robust analysis," *Journal of Monetary Economics*, Vol. 52, pp. 113–157.

Beaudry, P., Galizia, D., and Portier, F. (2015) "Reviving the Limit Cycle View of Macroeconomic Fluctuations,"Technical report, National Bureau of Economic Research.

——— (2020) "Putting the Cycle Back into Business Cycle Analysis," *American Economic Review*, Vol. 110, pp. 1–47.

Boldrin, M. and Woodford, M. (1990) "Equilibrium models displaying endogenous fluctuations and chaos: a survey," *Journal of Monetary Economics*, Vol. 25, pp. 189–222.

Brenner, T., Weidlich, W., and Witt, U. (2002) "International Co-movements of Business Cycles in a "Phase-locking" Model," *Metroeconomica*, Vol. 53, pp. 113–138.

Brock, W. A. (1986) "Distinguishing random and deterministic systems: Abridged version," *Journal of Economic Theory*, Vol. 40, pp. 168–195.

Bullard, J. (1994) "Learning equilibria," *Journal of Economic Theory*, Vol. 64, pp. 468–485.

Bulow, J. I., Geanakoplos, J. D., and Klemperer, P. D. (1985) "Multimarket oligopoly: Strategic substitutes and complements," *Journal of Political Economy*, Vol. 93, pp. 488–511.

Calderon, C., Chong, A., and Stein, E. (2007) "Trade intensity and business cycle synchronization: Are developing countries any different?" *Journal of International Economics*, Vol. 71, pp. 2–21.

Christiano, L. J. and Fitzgerald, T. J. (2003) "The band pass filter," *International Economic Review*, Vol. 44, pp. 435–465.

Clark, T. E. and Van Wincoop, E. (2001) "Borders and business cycles," *Journal of International Economics*, Vol. 55, pp. 59–85.

Cooper, R. and John, A. (1988) "Coordinating coordination failures in Keynesian models," *Quarterly Journal of Economics*, pp. 441–463.

Cravino, J. and Levchenko, A. A. (2017) "Multinational firms and international business cycle transmission," *Quarterly Journal of Economics*, Vol. 132, pp. 921–962.

De Grauwe, P. (2011) "Animal spirits and monetary policy," *Economic Theory*, Vol. 47, pp. 423–457.

De Haan, J., Inklaar, R., and Jong-A-Pin, R. (2008) "Will business cycles in the euro area converge? A critical survey of empirical research," *Journal of Economic Surveys*, Vol. 22, pp. 234–273.

Di Giovanni, J. and Levchenko, A. A. (2010) "Putting the parts together: trade, vertical linkages, and business cycle comovement," *American Economic Journal: Macroeconomics*, Vol. 2, pp. 95–124.

Di Giovanni, J., Levchenko, A. A., and Mejean, I. (2018) "The micro origins of international business-cycle comovement," *American Economic Review*, Vol. 108, pp. 82–108.

Fazzari, S., Ferri, P., and Greenberg, E. (2008) "Cash flow, investment, and Keynes–Minsky cycles," *Journal of Economic Behavior & Organization*, Vol. 65, pp. 555–572.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. (2015) "The next generation of the Penn World Table," *American economic review*, Vol. 105, pp. 3150–82.

Foley, D. K. (1992) "A contribution to the theory of business cycles," *Quarterly Journal of Economics*, Vol. 107, pp. 1071–1088.

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

Goodwin, R. M. (1951) "The nonlinear accelerator and the persistence of business cycles," *Econometrica*, pp. 1–17.

——— (1967) *A Growth Cycle*: Cambridge University Press.

Granger, C. W. (1966) "The typical spectral shape of an economic variable," *Econometrica*, pp. 150–161.

Gualdi, S., Bouchaud, J.-P., Cencetti, G., Tarzia, M., and Zamponi, F. (2015) "Endogenous crisis waves: stochastic model with synchronized collective behavior," *Physical Review Letters*, Vol. 114, p. 088701.

Haxholdt, C., Kampmann, C., Mosekilde, E., and Sterman, J. D. (1995) "Mode-locking and entrainment of endogenous economic cycles," *System Dynamics Review*, Vol. 11, pp. 177–198.

Hicks, J. R. (1950) *A Contribution to the Theory of the Trade Cycle*: Oxford University Press.

Hommes, C. H. and Manzan, S. (2006) "Comments on "Testing for nonlinear structure and chaos in economic time series"," *Journal of Macroeconomics*, Vol. 28, pp. 169–174.

Hsu, C.-C., Wu, J.-Y., and Yau, R. (2011) "Foreign direct investment and business cycle co-movements: The panel data evidence," *Journal of Macroeconomics*, Vol. 33, pp. 770–783.

Imbs, J. (2004) "Trade, finance, specialization, and synchronization," *Review of Economics and Statistics*, Vol. 86, pp. 723–734.

——— (2006) "The real effects of financial integration," *Journal of International Economics*, Vol. 68, pp. 296–324.

Inklaar, R., Jong-A-Pin, R., and De Haan, J. (2008) "Trade and business cycle synchronization in OECD countries—A re-examination," *European Economic Review*, Vol. 52, pp. 646–666.

Johnson, R. C. (2014) "Trade in intermediate inputs and business cycle comovement," *American Economic Journal: Macroeconomics*, Vol. 6, pp. 39–83.

Judd, K. L. (1985) "On the performance of patents," *Econometrica*, pp. 567–585.

Kaldor, N. (1940) "A model of the trade cycle," *Economic Journal*, pp. 78–92.

Kalemli-Ozcan, S., Papaioannou, E., and Peydro, J.-L. (2013) "Financial regulation, financial globalization, and the synchronization of economic activity," *Journal of Finance*, Vol. 68, pp. 1179–1228.

Kleinert, J., Martin, J., and Toubal, F. (2015) "The few leading the many: Foreign affiliates and business cycle comovement," *American Economic Journal: Macroeconomics*, Vol. 7, pp. 134–59.

Kose, M. A. and Yi, K.-M. (2006) "Can the standard international business cycle model explain the relation between trade and comovement?" *Journal of International Economics*, Vol. 68, pp. 267–295.

Krugman, P. (1996) *The self-organizing economy*: Blackwell Publishers.

Kydland, F. E. and Prescott, E. C. (1982) "Time to build and aggregate fluctuations," *Econometrica*, Vol. 50, pp. 1345–1370.

Liao, W. and Santacreu, A. M. (2015) "The trade comovement puzzle and the margins of international trade," *Journal of International Economics*, Vol. 96, pp. 266–288.

Matsuyama, K. (1999) "Growing through cycles," *Econometrica*, Vol. 67, pp. 335–347.

——— (2007) "Credit traps and credit cycles," *American Economic Review*, Vol. 97, pp. 503–516.

Matsuyama, K., Sushko, I., and Gardini, L. (2014) "Globalization and synchronization of innovation cycles."

Ng, E. C. (2010) "Production fragmentation and business-cycle comovement," *Journal of International Economics*, Vol. 82, pp. 1–14.

Nikolaidi, M. and Stockhammer, E. (2017) "Minsky models: A structured survey," *Journal of Economic Surveys*, Vol. 31, pp. 1304–1331.

Pecora, L. M. and Carroll, T. L. (1990) "Synchronization in chaotic systems," *Physical Review Letters*, Vol. 64, p. 821.

——— (1998) "Master stability functions for synchronized coupled systems," *Physical Review Letters*, Vol. 80, p. 2109.

Pikovsky, A., Kurths, J., Rosenblum, M., and Kurths, J. (2003) *Synchronization: a universal concept in nonlinear sciences*, Vol. 12: Cambridge University Press.

Scheinkman, J. A. and LeBaron, B. (1989) "Nonlinear dynamics and stock returns," *The Journal of Business*, Vol. 62, pp. 311–337.

Selover, D. D. and Jensen, R. V. (1999) "Mode-locking and international business cycle transmission," *Journal of Economic Dynamics and Control*, Vol. 23, pp. 591–618.

Shintani, M. and Linton, O. (2004) "Nonparametric neural network estimation of Lyapunov exponents and a direct test for chaos," *Journal of Econometrics*, Vol. 120, pp. 1–33.

Shleifer, A. (1986) "Implementation cycles," *Journal of Political Economy*, Vol. 94, pp. 1163–1190.

Silverberg, G. and Lehnert, D. (1993) "Long waves and "evolutionary chaos" in a simple Schumpeterian model of embodied technical change," *Structural Change and Economic Dynamics*, Vol. 4, pp. 9–37.

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and De Vries, G. J. (2015) "An illustrated user guide to the world input–output database: the case of global automotive production," *Review of International Economics*, Vol. 23, pp. 575–605.

Woltjer, P., Gouma, R., and Timmer, M. P. (2021) "Long-run World Input-Output Database: Version 1.1 Sources and Methods," GGDC Research Memorandum 190.