Performance of early warning signals for disease emergence: a case study on COVID-19 data

Daniele Proverbio[1,2], Françoise Kemp[1], Stefano Magni[1], and Jorge Gonçalves[1,3]

[1]Luxembourg Centre for Systems Biomedicine, University of Luxembourg, Campus Belval, 6 avenue du Swing, L-4367 Belvaux, LU
[2]College of Engineering, Mathematics and Physical Sciences, University of Exeter, EX4 4QL, Exeter, UK
[3]Department of Plant Sciences, University of Cambridge, CB2 3EA, Cambridge, UK

March 30, 2021

Abstract

The sudden emergence of infectious diseases pose threats to societies worldwide and it is notably difficult to detect. In the past few years, several early warning signals (EWS) were introduced, to alert to impending critical transitions and extend the set of indicators for risk assessment. While they were originally thought to be generic, recent works demonstrated their sensitivity to some dynamical characteristics such as system noise and rates of approach to critical parameter values. Moreover, testing on empirical data is so far limited. Hence, validating their performance remains a challenge. In this study, we analyse the performance of common EWS such as increasing variance and autocorrelation in detecting the emergence of COVID-19 outbreaks in various countries, based on prevalence data. We show that EWS are successful in detecting disease emergence provided that some basic assumptions are satisfied: a slow forcing through the transitions and not fat-tailed noise. We also show cases where EWS fail, thus providing a verification analysis of their potential and limitations. Overall, this suggests that EWS can be useful for active monitoring of epidemic dynamics, but that their performance is sensitive to surveillance procedures. Our results thus represent a further step towards the application of EWS indicators for informing public health policies.

1 Introduction

Among the natural phenomena that undergo critical transitions, epidemics pose important and long-lasting threats to human societies. The current COVID-19 pandemic is an extreme and vivid example of how harmful epidemics could be [1]. Hence, developing tools for rapid and early detection of disease emergence is important to perform science-based risk assessment [2]. In principle, detailed mechanistic understanding could help formulating predictive models. However, combinations of non-linearity, noise and lack of curated datasets for validation hamper the development of complete models. Therefore, numerous recent studies have suggested using different methods, agnostic of detailed mechanistic models, that could alert for shifts in epidemic dynamics [3, 4, 5, 6, 7]. Said methods imply the calculation of summary statistics from observed data, and are rooted in the theory of critical transitions in dynamical systems [8].

Critical transitions are a broad class of phenomena characterised by sudden shifts of the system dynamics. Their main drivers are (but not limited to) bifurcations (qualitative changes of equilibria due to leading eigenvalues crossing a threshold value) and random perturbations [9]. As dynamical bifurcations are properties that are typical of different systems, statistical indicators that could detect them have been suggested to be generically applicable to a broad variety of phenomena [10, 11]. The emergence of epidemics is characterised by a transcritical bifurcation [7] when the control parameter R, the average

number of secondary infections from a single contagious case in a susceptible population, crosses its threshold value 1 [12]. Hence, epidemics are suitable candidates for the application of statistical early warning signals (EWS) such as increasing variance or autocorrelation prior to the transition. Theoretical and computational studies already investigated their performance on abstract epidemiological models [7, 4, 13, 14, 15], but so far only few testings on empirical data were performed [16].

In this study, we follow the strategy of "natural experiments" [17], that is, to test theoretical predictions about complex systems on wide and variegate sets of data, provided that we account for possible confounders. To do so, we screen COVID-19 epidemic curves from numerous countries, to build a testable dataset for EWS predictions and to assess their performance, according to some basic assumptions from critical transitions theory [18]. In particular, we consider the re-emergence of the disease (the so called "second wave") [19, 20] and its underlying characteristics, such as the quality of prevalence data, the rate of approach to R's critical value, and the system noise. We notice that, if basic theoretical assumptions are satisfied, EWS from the critical transitions framework are able to detect the transition to disease emergence. However, they are sensitive to rapid and strong fluctuations in epidemic patterns. As a result, we suggest that they are suitable candidates for epidemic monitoring and deserve further attention, to expand the current toolbox of indicators.

2 Theoretical background on EWS

Sudden dynamical changes happen when a system is pushed over a critical point and are common in many complex systems [21]. Recent studies have shown that the trend of some statistical indicators may signal the approach to a critical bifurcation in a slowly forced dynamical system [9, 22]. In general, a slowly forced system with variables $\mathbf{x}$ and control parameter q is:

$$\begin{cases} \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, q) \\ \dot{q} = \epsilon \mathrm{g}(\mathbf{x}, q) \end{cases} \tag{1}$$

with $0 < \epsilon \ll 1$. When describing epidemics dynamics, a SIR-like model with slowly evolving control parameter R can be expressed as a slowly forced system (*cf.* "Mathematical models and assumptions", Eq. 4). At first approximation, R can be modeled with a linear trend, *cf.* Eq. 3 and 5. If its regression coefficient is small, then the modelling assumption Eq. 1 is satisfied.

In addition, random fluctuations can push the system state around the deterministic trend of Eq. 1. The statistical properties of such fluctuations change as the system approaches the transition [23]. Hence, the evolution trends of statistical indicators of fluctuations around the equilibria have been proposed as early warning signals (EWS) of impending transitions [24, 25]. In "Materials and Methods", we outline how summary statistics and early warning signals are derived for epidemiological models. If fluctuations are normally distributed, the trend of the most common summary statistics (variance, autocorrelation, skewness, coefficient of variation) is expected to increase next to the transition [11, 26, 3]. In addition to qualitative evaluation, recent studies have suggested to quantify the increasing trend with either deviations from the mean or with Kendall's τ coefficient of monotonicity [27, 18, 13].

In epidemiology studies, several theoretical results suggest that variance increases are the best performing indicators in terms of signal-to-noise ratio and of detection Area Under the Cure (AUC) [3, 4]. However, as noticed in follow-up studies [25, 27, 28], such indicators perform best when some modelling assumptions are satisfied. Said assumptions are linked to how EWS are theoretically derived, *cf.* above and in "Materials and Methods". First, critical transitions are local phenomena. Hence, EWS are not global measures, but are expected to work in the vicinity of the regime shift. Second, it should be possible to express the epidemic dynamics in terms of a fast-slow system like Eq. 1. If this is not satisfied, literature results suggest that the expected patterns will be either distorted or will not occur [7, 3]. Third, the closer random fluctuations are to being white noise, the more robust the performance of EWS is. If there is deviation from white noise, their trends can be modified or disrupted. For instance, decreasing variance was observed in case of intrinsic multiplicative noise [28].

The results described above were obtained by means of theoretical analysis and computer simulations, where each assumption was satisfied, or singularly varied. It is then compelling to validate such modelling predictions against empirical data, and to test their performance in different contexts.

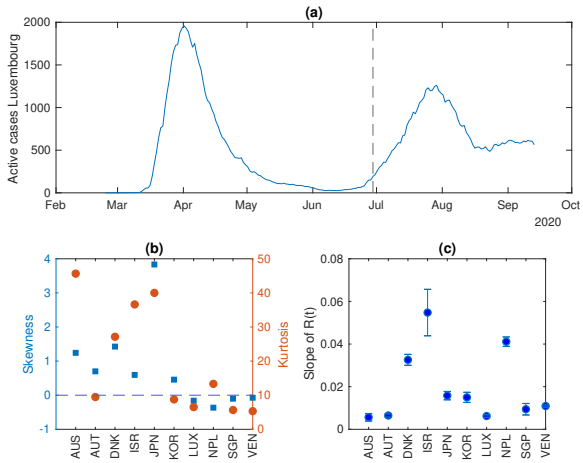

Fig. 1 Analysis of the dynamical characteristics of the countries included in the test set $\mathcal{Y}$. a) An example of an epidemiological curve of active cases, not detrended. The dashed line indicates the transition, measured by $R > 1$. To identify it objectively, we identify the date at which the probability of $R(t)$ to be greater than 1 is at its maximum (see "Materials and Methods"); b) Measures of the distribution of data fluctuations. Skewness indicates the symmetry of the distribution, whereas kurtosis indicates the relevance of its peak with respect to the tails. Large deviations from $\mu = 0$ (dashed line) and $\gamma = 3$ are associated with non-normal distributions. c) The regression coefficient of $R(t)$ and its associated uncertainty, as obtained from the linear fit Eq. 13.

3 Analysis of country-wise dynamical characteristics associated to the spread of COVID-19

Usually, early warning signals are tested in controlled experiments that satisfy the assumptions listed above [29]. However, those are often not feasible in complex phenomena like epidemics, which would also require frequent monitoring of prevalence data [13, 15]. Instead, we use worldwide data, collected during the COVID-19 pandemic, to construct a dataset that displays the dynamical characteristics of interest. The COVID-19 disease, caused by severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) [30], rapidly diffused in the whole world [31]. After a first wave, characterised by an initial R_0 [32], the effective time-dependent reproduction number $R(t)$ [33] decreased due to a number of non-pharmaceutical interventions. Later on, many countries experienced an epidemic emergence (often called "second waves") associated to $R(t)$ re-crossing 1 from below. By analysing dynamical characteristics (noise, rate of evolution of $R(t)$) of each country, it is then possible to construct a meaningful test set that reflects the modelling assumptions, and to compare the performance of early warning signals in real world situations.

Before analysing the early warning signals, we screened worldwide time series to improve the quality of the dataset and to subdivide it according to various dynamical characteristics. For details about data collection and curation, refer to "Materials and Methods". For a list of selected countries that satisfy the requirements explained in "Data Collection and Curation", refer to Table 1. The table also reports the second wave dates identified by the analysis of $R(t)$, which serve as ground truth for detection. Upon the selected time series, we assess if fluctuations of the data are close to be normally distributed, and we study the speed of approach of the control parameter $R(t)$ to its critical value (*cf.* "Constructing the test set"). Examples of data selection and subsequent analysis are reported in SI Appendix, Figs S1 and S2.

Fig. 1 reports the initial analysis of the epidemic curves, for those countries that displayed a second wave and had curated prevalence data. First, we estimate the "ground truth" day of disease emergence. In order to have objective and repeatable results, we used automatic, Bayesian-based algorithms [34, 35] to identify when the empirical $R(t)$ and its lower confidence interval crosses 1. Sometimes, the methods obtained $R > 1$ after some small increase in prevalence, that was nonetheless associated to random

Table 1: Selected countries for the dataset, abbreviations and date of second epidemic insurgence.

Country	Abbr.	Date
State of Victoria (Australia)	AUS	27/06/2020
Austria	AUT	01/07/2020
Denmark	DNK	03/08/2020
Israel	ISR	01/06/2020
Japan	JPN	28/06/2020
Korea, South	KOR	13/08/2020
Luxembourg	LUX	29/06/2020
Nepal	NPL	29/07/2020
Singapore	SGP	25/07/2020
Veneto (Italy)	VEN	29/07/2020

Table 2: Refer to "Data Collection and Curation" for how the date marking the second wave was obtained. Abbreviations correspond to ISO codes.

events rather than to deterministic emergence; an example is show in Fig. 1a. The selected countries exhibit different ranges of noise distribution (Fig. 1b). Austria, Luxembourg, Nepal, Singapore and Veneto are close to being Gaussian distributed, while others are more divergent. This could be associated to social dynamics or imperfect data reporting [36]. The approach to the transition is also different, as indicated in Fig. 1c by the regression coefficient of a linear fit for $R(t)$ (*cf.* Eq. 13). State of Victoria, Austria, Luxembourg, Singapore and Veneto display a slow approach to the critical value and can thus be better suited to be appropriately described as slow-fast systems like Eq. 1. Japan and South Korea show intermediate values, while other countries (Denmark, Israel and Nepal) have a faster evolution of the control parameter and are thus associated to violation of this assumption.

By this analysis, we notice that Luxembourg satisfies the modelling assumptions and is thus the closest to being a controlled experiment w.r.t. the current criteria. In fact, not only the country is small and was under a homogeneous lockdown during the first wave, but a Large Scale Testing (LST) strategy was implemented to monitor the virus diffusion[1]. Additional countries (State of Victoria, Austria, South Korea, Singapore, Veneto) meet some of the modelling assumptions and will thus be included in a test set $\mathcal{Y}$. In fact, they either display a similar rate of approach to the transition as Luxembourg (within the error bars), or a similar noise distribution. The remaining ones (DNK, ISR, JPN, NPL) do not satisfy either of the modelling assumptions and are grouped in a set $\mathcal{N}$.

4 Local performance of EWS on controlled data

In accordance to the above analysis, we first focus on Luxembourg to test the theoretical predictions about the local behavior of common EWS. Summary statistic indicators are estimated from the detrended fluctuations around prevalence data as per standard methods [18]. We first check that the detrending method is not inducing severe bias, by comparing the deviations from the deterministic trend resulting from a Gaussian kernel smoothing [37] or a moving average filtering [11]. As displayed in the inset of Fig. 2, the Pearson correlation coefficient between the two filtering methods is $\rho = 0.95$. This excludes explicit bias in the analysis due to the filtering process, and allows us to use the Gaussian one from here on, without loss of generality.

Then, we study the behavior of the variance next to the transition point, identified as the day when the estimated $R(t)$ crosses 1 (dashed line in Fig. 2). The increase in variance prior to the transition, as expected from theoretical studies, is evident in Fig. 2. We also confirm that the increasing trend is not much sensitive to the size of the moving window. Although the lead time is slightly advanced for shorter window sizes, the corresponding Kendall's τ measure of monotonous increase [38] is similar: $\tau = 0.94$ for a window size of 7 days, and $\tau = 0.81$ for a window size of 14 and $\tau = 0.92$ for a window

[1]Refer to the ACAPS database (`https://bit.ly/3nFFqUS`) for policies and to the Luxembourg Health Ministry (`https://bit.ly/36SSPCn`) for information about the LST. This country-wide testing strategy reached more than 70.000 tests per week over a population of about 600.000, thus allowing extensive and frequent monitoring.

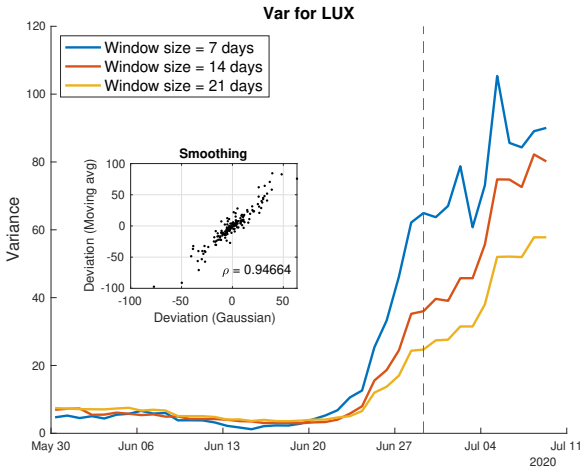

Fig. 2 Analysis of the variance, the most robust EWS [4], in the Luxembourg controlled setting. Its increase is evident prior to the transition (dashed line). The increasing trend is little sensitive to the sliding window size, as displayed by the three curves and by τ values reported in the text. Inset: correlation between deviations from the mean trend, obtained with a moving average or a Gaussian filtering.

size of 21 days. The increasing trend for each window size can be appreciated in Fig. 2. On the contrary, smaller window size is associated with lower smoothing. Hence, from here on, we will use a window of 14 days, which groups enough data to be robust without being over dependent on past history and represents a reasonable trade-off between smoothing and detection power.

These findings confirm that, in controlled settings that satisfy the modelling assumptions, the theoretical predictions are valid. Indeed, an increasing trend in variance in the vicinity of the transition point could serve as early warning to detect the transition to disease emergence.

5 Global performance of EWS

After confirming the local behavior of the variance in a tightly controlled setting, we widen the analysis to the global performance of other EWS, *i.e.* far from the bifurcation, and for different countries from the pre-defined dataset, *cf.* Table 1. This way, we further test the theoretical predictions and the EWS potential use in more general contexts.

Among the indicators, we estimate lag-1 autocorrelation (AC(1)), skewness and coefficient of variation (CV), which are often proposed as alternatives to the variance. The size of the moving window is set to 14 days as discussed above. To compare the trend of EWS with the approach to the bifurcation, the probability of $R(t)$ to be greater than zero, Eq. 15, is also calculated and reported.

Fig. 3 shows the results, for Luxembourg, Austria, State of Victoria (from the test set). In addition, Israel, which does not satisfy the EWS assumptions *cf.* Fig. 1, is reported to allow inspecting a deviant case. The figure focuses on EWS trends after the first wave, up to more than a month after the second epidemic insurgence. The graphs for other countries from Table 1 are reported in SI Appendix Fig S4, along with their associated prevalence data and estimated $R(t)$.

Let us first focus on Luxembourg and Austria. The indicator $\mathcal{P}(R > 1)$ identifies the approach to the transitions, marked by the straight line at its maximum. The variance follows its theoretically predicted behavior closely (*cf.*, e.g., Fig. 9b in [7]), with a small but visible increase prior to the transition and a subsequent monotonous trend along the second wave. In Austria, though, it still displays some fluctuations after the relaxation of the first wave. The same happens for the coefficient of variation CV; this is expected, as it depends on the variance and on a stable equilibrium in infectious numbers. On the other hand, the lag-1 autocorrelation shows an increasing trend very close to the transition point, but returns possibly spurious signals during the global time series. Finally, the skewness does not display immediately detectable relevant trends, as anticipated by computational studies [13]. This might be

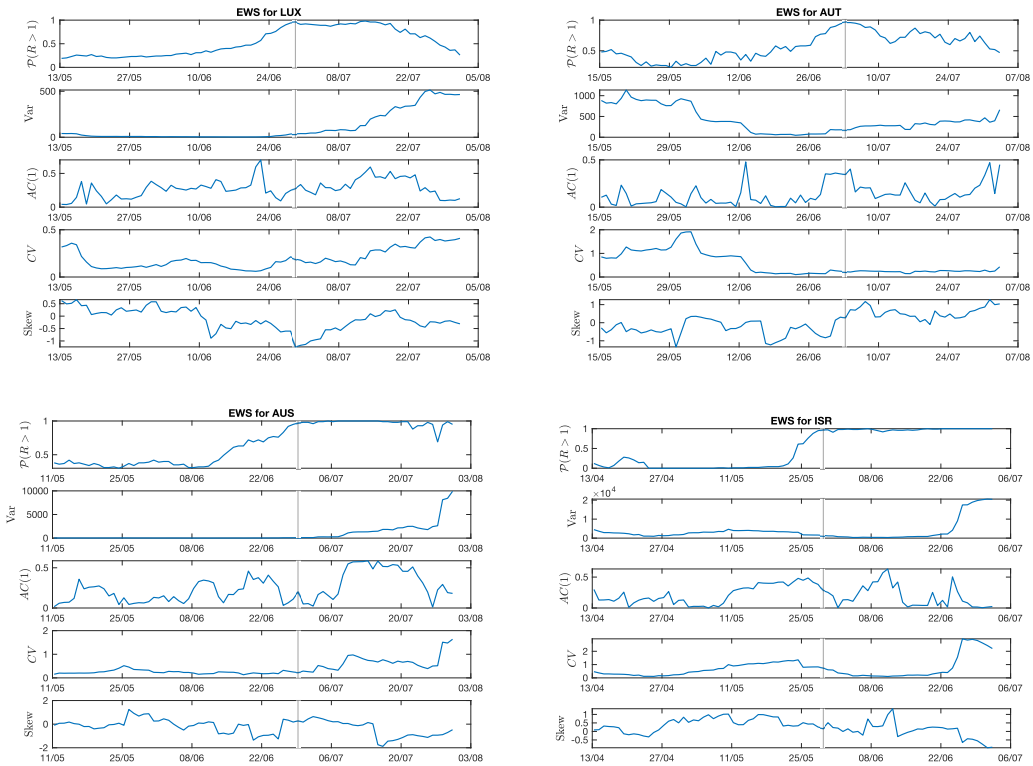

Fig. 3 Evolution of EWS far from the transition point, for four example countries (Luxembourg and Austria, with controlled features), State of Victoria (Australia) with small deviations from controlled features, and Israel that does not satisfy theoretical conditions. Considered EWS are the most common ones (variance, lag-1 autocorrelation, coefficient of variation, skewness). In addition, to mark the approach to the transition, $\mathcal{P}(R > 1)$ from the Bayesian estimation (see Eq. 15) is displayed. The vertical line reports the transition date. Other countries are reported in SI Appendix.

related to noise properties, as suggested by [39].

Variance and CV on Australian data, when processed by eye, show a delay in their increasing trend, which begins around the 7th of July (here, we consider the expected *small* increase like the one in LUX and AUT). This is likely a feature of the so-called "bifurcation delay" (see "Mathematical models and assumptions"), which is associated with deviations from Gaussian noise [11, 28] as those reported in Fig. 1. In this case, the delay is of about two weeks, which could impair online detection of epidemic insurgence.

Finally, Israel is an interesting case study as it diverges from the theoretical assumptions, see Fig. 1. In fact, its transition to epidemic insurgence is rapid, and the noise distribution is far from being Gaussian. These disrupt the EWS trends as predicted by the theory. In fact, the variance remains flat around the transition, CV and skewness slightly decrease, while lag-1 autocorrelation does not display informative patterns. A bifurcation delay is reported more than 20 days after the transition, but is as abrupt as the exponential increase in infectious data (*cf.* SI Appendix). This shows that the application of early warning signals indicators on correct contexts is crucial to obtain reliable signals to develop risk assessment analysis.

To go beyond qualitative visual inspection of time series, we provide a quantitative estimation of their performance in detecting the transition. The Kendall's τ score is used to evaluate if a certain indicator corresponds to an increasing or decreasing trend and compare this for different data types [27, 13]. Hence, we evaluate τ for each indicator, over the same 14 days window, and we assess which values are associated with a passage through the transition point. The increase in τ is reflected in the Receiver Operator Characteristics (ROC) curve and quantified by Area Under the Curve (AUC) scores (see "Ma-

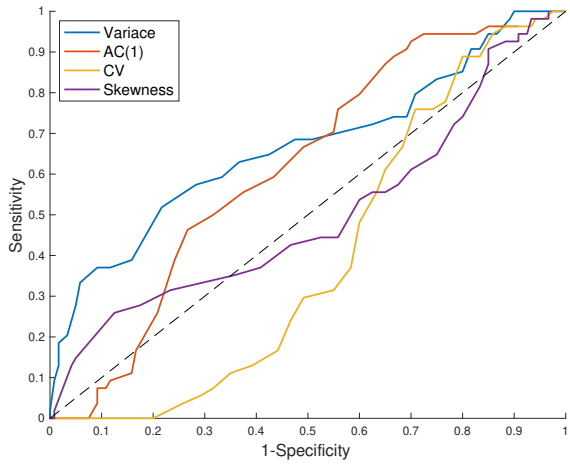

Fig. 4 ROC curves for each considered indicator, estimated on the test set. Each point corresponds to a null value to define if the detection if positive. The diagonal line corresponds to the ROC of a random classifier. Curves above it imply better performance.

Table 3: AUC scores for different indicators, over $\mathcal{Y}$ and $\mathcal{N}$ datasets.

Indicator	over $\mathcal{Y}$	over $\mathcal{N}$
Variance	0.6671	0.1981
AC(1)	0.5258	0.4995
CV	0.3968	0.0976
Skewness	0.4664	0.2925

terials and Methods"). Fig. 4 shows the ROC curves for the considered indicators among all countries in $\mathcal{Y}$, while Table 3 reports the corresponding AUC values. The variance is the only indicator that consistently performs better than a random classifier, while the lag-1 autocorrelation performs slightly better for a number of null values $\tau_0 > 0.5$. This validates aforementioned theoretical results from literature. Instead, the skewness does not improve the detection performance. This is probably due to its fluctuations around the 0 value, as noticed in Fig. 3, which is in turn associated with noise distribution of original data. Interestingly, the coefficient of variation is overall the worst performer. We speculate that this is due to its sensitivity to data fluctuations, which are often non negligible even in countries that belong to the test set $\mathcal{Y}$ (*cf.* Fig. 1). We remark that our findings are sensitive to the estimated time of emergence, which also complicates the estimation of the lead time. Currently, the best lead time is of 5 days for Luxembourg, which remarks the performance of EWS in settings that are close to the analytical assumptions.

The same analysis was performed over the set $\mathcal{N}$ of countries that do not satisfy the theoretical assumptions. Their AUC values are reported in Table 3. Such values clearly show that the considered indicators are not able to detect the transitions, overall performing worse than random classifiers. This supports what already observed for Israel in Fig. 3, where a slight *decreasing* trend was observed. It contradicts what was expected and thus returns a false negative. Time series for indicators of other countries in $\mathcal{N}$ are reported in SI Appendix Fig. S4. Hence, if a system is not known or there is difficulty in determining the type of data, incorrect conclusions could be drawn when interpreting the time series trend.

6 Discussion

Research on early warning signals from the theory of dynamical systems has greatly progressed in the last years. Building on dynamical commonalities across diverse systems, it allowed to formulate generic

indicators for the detection of transition points. An important application is the detection of epidemic emergence, for which analytical and computational studies have been thoroughly performed. In fact, when facing epidemics, theory-based indicators can provide useful evidence, particularly when scarce data and few prior information are a constraint for using large scale statistics or Machine Learning (ML) methods. In addition, they can complement the existing toolbox of indicators to improve epidemic risk assessment, and serve as basis for theoretically informed feature selection for ML predictive algorithms like in [16].

To support the growing corpus of theoretical studies, validating their predictions on empirical data is necessary to guarantee the adequacy of proposed indicators on real world applications. When randomised experimental studies are not possible, observational studies provide stronger evidence if consistent patterns are seen in multiple locations and at multiple times, after checking for possible confounders.

In this study, we used world-wide available data about the ongoing COVID-19 pandemic, concentrating on the re-emergence of the disease after a first wave in Spring 2020. To limit biases associated with country-specific testing and reporting capacities, we constructed a limited, but curated set of time series. On the one hand, data quality could be a limiting factor, on the other it is representative of real world monitoring capacities, that could be further improved. In fact, the dataset selection highlights the importance of quality prevalence data [13] and monitoring.

Like $R(t)$ and other indicators [40], EWS are estimates that rest on assumptions; hence, we screened the dataset to assess the matching of empirical features and theoretical assumptions. Then, we tested the best performing EWS, as suggested in the literature, in detecting the epidemic re-emergence. We acknowledge that our definition of "ground truth" transition date is somewhat conservative, as we requested to have maximum probability of $R(t)$, the control parameter, to be greater than 1. In the real world, the appropriate detection threshold is conditional on the various costs of a late outbreak alert, and requires an assessment by public health authorities which could modify the estimated lead time. Nonetheless, our results confirm the general validity of the indicator system proposed and confirm the expected trends in EWS indicators. In particular, we showed that dynamical EWS are likely to operate successfully in contexts where the approach to the transition is gradual and not subject to high fluctuations. Further studies could associate these features to social behaviours and political strategies. Finally, we analysed the limitations of the indicator system in other contexts, characterised by different dynamical features such as rapid increases of $R(t)$. This remark that, for EWS to properly work, the real system must fulfill the discussed conditions underlying the EWS modelling. These highlighted open problems remark that knowledge of the type of collected data is imperative to avoid misleading judgements in response to time series trends: EWS as epidemiological constructs will only remain valuable and relevant when used and interpreted correctly [41]. Our results thus constitute a further step towards the validation of literature findings and call for future studies, which will contribute to the exciting field of EWS in epidemic control and will likely have a tremendous impact in informing public health policies.

Materials and Methods

We describe the theoretical framework leading to the study of early warning signals as well as the steps to reproduce the analysis.

Mathematical models and assumptions

As mentioned in the Main Text, early warning signals are measures that rest upon modelling assumptions. To clarify which are the most relevant ones, which are tested in this study, we recall how EWS are theoretically derived.

To illustrate how early warning systems can be derived from SIR models, we here recall the process described in [7]. When consistent with a mean-field approximation, the dynamics of the infectiousness of COVID-19 is well described by SIR-like models [42, 43]. It describes disease processes in homogeneous populations of susceptible individuals (X), which can get infectious (Y) and eventually removed (Z) by death or recovery. In the deterministic SIR, transitions among states are governed by the infection

rate β and the removing rate γ which lumps recovery and death rate. Empirical values for COVID-19 can be traced in the abundant literature, e.g. [44]. In addition, systems are often open with influx rate μ' and outflux rate μ'' of people. However, as many travelling restrictions[2] were in place during the year 2020, we can assume that such fluxes are small and balanced: $\mu' = \mu'' = \mu$. Along with that, we model an influx rate of new cases η that can trigger subsequent disease outbreaks. Finally, intervention measures introduce a probability p that some susceptible individuals are isolated and protected, either physically or by vaccination [45]. The model reads:

$$\begin{aligned} \dot{X} &= \mu(1-p) - \beta XY - (\eta+\mu)X \\ \dot{Y} &= \beta XY + \eta X - (\gamma+\mu)Y \\ \dot{Z} &= \mu p + \gamma Y - \mu Z \end{aligned} \tag{2}$$

It is assumed that the population size N is constant, that is $X + Y + Z = 1$. In this case, the control parameter R is given by [7]:

$$R = \frac{\beta}{\gamma+\mu}(1-p) \tag{3}$$

and reaches its critical value 1 for $p^* = 1 - (\gamma+\mu)/\beta$, at which the dynamics undergo a transcritical bifurcation on the (Y, p) bifurcation diagram [7]. We remark that we consider here an example from literature: R might change in time after being driven by other evolving parameters such as one that tunes the contact rate β [43]. Without protection and extra fluxes, the basic reproduction number for COVID-19 was estimated at the beginning of the pandemic in the range $2 < R < 4$ (*cf., e.g.*, [32]).

If $p(t)$ changes slowly over time, we can mathematically express the SIR model 2 approaching the transition as a slow-fast system:

$$\begin{cases} \text{Eq.2} \\ \dot{p} = \epsilon f(X, Y, p) \end{cases} \tag{4}$$

where $0 < \epsilon \ll 1$ and f is a function that describes its change. A limit case is $\epsilon \to 0$. This condition is necessary to interpret the dynamical shift as a slow crossing through a bifurcation point, and to compute its associated summary statistics [23]. Often [9, 7], it is assumed a constant $f = \tilde{p}$. This way, the protection probability is, at first approximation, a linear function of time:

$$p = p_0 + \tilde{p}t \tag{5}$$

and R as well, following Eq. 3. If $\tilde{p} > 0$, then R gets reduced and, if it was above 1, the bifurcation is crossed from above towards elimination. If $\tilde{p} < 0$, R increases and, if it was below 1, the bifurcation is crossed from below, towards a new emergence. Only the second case can be investigated with COVID-19 data, as most countries implemented suppression measures very rapidly [46] and thus Eq. 4 is not satisfied.

When the transition is approached from below, and if there are few cases, stochastic fluctuations are not negligible. Hence, we need to consider the transitions described by a stochastic master equation. Reducing it to Eq. 4 and a Fokker-Plank equation for the fluctuations was already performed in [7, 47]. Hence, we briefly recall them to illustrate the assumptions underlying the behavior of early warning signals prior to the transition.

First, note that system 2, along with condition $\dot{N} = 0$, can be reduced to its first two equations. Hence, we just need to consider the transitions in and out X and Y. Second, for each small time step dt, the quasi-steady state p is constant.

Transitions in and out states are described as random jump processes. Such states are (X, Y), $(X-1, Y)$, $(X, Y+1)$ and so on. Using $\alpha = (X, Y)$ to describe a state, $\bar{\alpha}$ is any other state and $P(X, Y, t) = \text{Prob}(X(t), Y(t) = (x, y))$ is the probability that the state vector is equal to some non negative integer number (x, y). Finally, $T_i(\alpha, \bar{\alpha})$ is the transition probability between states, and is a function of transition rates. The subscript i denotes all possible jumps in and out of the states. Examples of $T_i(\alpha, \bar{\alpha})$ are found

[2]Refer to the ACAPS website (`https://bit.ly/3nFFqUS`) for a dataset of government measures.

in [7, 13], depending on the system of interest. Consequently, the master equation for the stochastic process is:

$$\frac{dP(\alpha,t)}{dt} = \sum_{\alpha\neq\bar{\alpha}} T_i(\alpha|\bar{\alpha})P(\bar{\alpha},t) - \sum_{\alpha\neq\bar{\alpha}} T_i(\bar{\alpha}|\alpha)P(\alpha,t) \tag{6}$$

In general, Eq. 6 is nonlinear. To have analytical results about its average behavior and the fluctuations around it, the van Kampen expansion can be used [48] to approximate the discrete random variables with continuous random variables. This depends on having large N, which holds in our case when we consider the population of medium to big countries. The approximation of Eq. 6 is a system of normal random variables. To leading order, the expansion is equivalent to Eq. 4. To next-to-leading order, the obtained Fokker-Plank equation is equivalent to the following system of stochastic differential equations [47]:

$$\begin{aligned} \frac{d\sigma}{dt} &= b_{11}\sigma(t) + b_{12}\zeta(t) + \Gamma_1(t) \\ \frac{d\zeta}{dt} &= b_{21}\sigma(t) + b_{22}\zeta(t) + \Gamma_2(t) \end{aligned} \tag{7}$$

where σ and ζ are continuous random variables from the van Kampen expansion which represents the fluctuations from the susceptible and infectious states, respectively. Γ_j are white noise processes and the elements of the matrix $B = \{b_{kl}\}$ are functions of transition rates. Eq. 7 connects the stochastic description of the epidemic with SIR-like models like Eq. 2 with a noise term.

Eq. 7 can be analysed by its Fourier transform. By considering the fluctuations around the infectious state, we can derive the power spectrum and, through integration, the variance, the autocorrelation and other statistical moments. Their specific values depend on the eigenvalues of matrix B and of the covariance matrices of Γ_j. The trend of these summary statistic indicators on the fluctuations, prior to the transition, constitutes the set of signals that could detect the transition itself; for instance, the increase in variance, often measured in terms of Kendall's τ. The Kendall's τ score is a non-parametric measure of ranks correlation, which is usually used to identify increasing trends [22, 27]. Such increasing trends are known in the literature as early warning signals (EWS) [24].

Finally, let us recall the general theory of EWS on bifurcations. Any system approaching a transcritical bifurcation is, in its vicinity, topologically equivalent to a transcritical normal form [49, 50]. This is a minimal model that retains the systems' behavior and resilience properties in the vicinity of a bifurcation. Models can be transformed to a normal form after an appropriate change of variables [21]. The one associated to a transcritical bifurcation has the form:

$$\dot{\theta} = q\theta - \theta^2 \tag{8}$$

where q is the bifurcation parameter and θ the variable of interest (Y from eq. 2, in this specific case). This form represents a system whose extinction state and positive steady state coalesce and exchange stability when q reaches its critical value. If there are statistical fluctuations ξ, we can write their evolution as a linearization around the equilibrium $\tilde{\theta}$, thus resulting in a Langevin equation [23, 25]:

$$d\xi = -\frac{\partial f}{\partial\theta}|_{\tilde{\theta}}\xi + \sqrt{\sigma^2 g^2(\tilde{\theta})}dW \tag{9}$$

where dW is a Wiener process, σ^2 models the noise level and g is the diffusion coefficient of the associated Fokker-Planck equation. With a linear $|\partial f/\partial\theta|_{\tilde{\theta}}| = k$, Eq. 9 is an Ornstein–Uhlenbeck process with known statistical moments [51]. For instance, the theoretical variance is:

$$\mathrm{Var}_{\text{theo}} = \frac{\sigma^2 g^2(\tilde{x})}{2k} \tag{10}$$

and the theoretical autocorrelation is:

$$\mathrm{AC}(1)_{\text{theo}} = e^{-k}. \tag{11}$$

Hence, we can complement the theoretical results described above with those from the theory of statistical indicators of normal forms, e.g. [23, 25, 28]. Important remarks are that: a) multiplicative noise

can modify the indicators trend, e.g. by making it decrease; b) EWS are expected to work best in the vicinity of the transition; c) there can be bifurcation delays associated to out-of-equilibrium phenomena, i.e. changes of the system state (and of its indicators) that lag behind the bifurcation of the limit case.

Data collection and curation

In order to validate the use of EWS to detect the emergence of epidemics, and to assess their performance, we constructed a dataset that satisfies the assumptions explained above. We analysed all countries that faced a resurgence of positive COVID-19 cases between beginning of March to mid September, when many countries issued new social measures that rapidly impacted the epidemic trends. When possible, prevalence data (*i.e.* active cases, in accordance to Eq. 2 and results in [13]) were obtained directly. For instance, active cases from Luxembourg were obtained from the government website (`COVID19.public.lu/fr/graph`). Otherwise, active cases A were estimated by the proxy [52]:

$$A = C - D - \tilde{R} \tag{12}$$

where C indicates the cumulative positive cases, D the number of registered deaths and $\tilde{R}$ the number of recovered patients. Country data were obtained from the John Hopkins University collection [31] and the European Centre for Disease Prevention and Control database (`https://www.ecdc.europa.eu/en/COVID-19/data`). For Italy, we concentrated on the Veneto region since it had an identifiable second wave during our considered time interval. Data were retrieved from the Github repository of the Italian "Dipartimento della Protezione Civile - Emergenza Coronavirus" (`https://github.com/pcm-dpc/COVID-19`). Databases were accessed up to 15/09/2020. Information about the share of positive tests was obtained from the OurWorldInData curated dashboard [53].
Pre-screening on data quality was performed. We rejected time series with very few active cases, as mostly driven by stochastic processes that do not satisfy the van Kampen expansion. We also did not include time series for which the share of positive cases was daily new cases/performed tests $> 5\%$ next to the transition, as WHO guidelines suggest possible undertesting (we refer to WHO reports such as `https://bit.ly/3dARcy1`). As EWS from critical transitions are based on mean-field homogeneous SIR-like models, we didn't consider countries with clear spatial heterogeneity [54]. When possible, we used regional data for those countries, e.g. Italy. Finally, we rejected some public time series that were behaving drastically differently from the common "bell-shaped" epidemiological curve [52] (see Appendix SI Fig. S1).
To identify *a posteriori* the transition, we computed the statistical time-dependent $R(t)$ as suggested in [34, 35, 4] using a Bayesian MCMC estimator. The most likely day t_{em} in which the transition happened, assumed as our ground truth, was obtained by identifying the maximum of $\mathcal{P}(R > 1)$, from Eq. 15 (see also SI Appendix).

Constructing the test set

To analyse the global distribution of fluctuations, we detrended the time series with a 7-days moving Gaussian kernel as suggested in [37, 55]. The window size reflects typical cycles of data reporting and of COVID-19 fluctuations [56]. Results from Gaussian kernel show a Pearson correlation coefficient $\rho = 0.95$ with the common moving average.
The distribution of fluctuations over the complete time series (from the beginning of the pandemic as reported in [31] to mid September) is indicative of the average noise. As indicators of deviations from Gaussian noise, we compute skewness and kurtosis [57].
To measure the speed of transition of the control parameter to its critical value, we computed the statistical time-dependent $R(t)$ and its confidence intervals with Bayesian Markov Chain Monte Carlo methods like in [34, 4] (see SI Appendix). The notation "$R(t)$" identifies values for R estimated from empirical data. Then, consistently with the linear approximation Eq. 5, we fit a linear function:

$$R(t) = a + b \cdot t \tag{13}$$

in the interval $t \in [\tilde{t}, t_{em}]$. Here, t_{em} corresponds to the day associated with novel disease emergence as explained above; $\tilde{t}$ is the day associated with the minimum of $R(t)$ after the first wave and possible

fluctuations. For the fitting, we used the *scipy* Python library. Refer to Appendix SI for details. The regression coefficient parameter $b \pm \sigma_b$ measures the evolution speed of the control parameter. As an indication, the control parameter is said "slowly evolving" if it goes from its minimum value to 1 in a period of time that is much longer than the COVID-19 serial interval (around 4 days [58]), which is a proxy of the disease time scale.

Estimation of EWS

Estimation of early warning signals from time series data was performed as per standard methods [18]. First, as EWS are expressed as summary statistics over fluctuations around an equilibrium, we detrended the time series to concentrate on the deviations from the deterministic trend. We verified that smoothing with moving averages and Gaussian kernels [37] are highly correlated and do not induce explicit bias. Then, we computed the statistical indicators associated to each point with a backward sliding window, i.e. the associated time point is the rightmost one, to be agnostic of future values. As mentioned in the main text, we initially focused on the variance, which is identified as the most robust indicator [3, 4, 13, 7] for epidemic emergence:

$$\mathrm{Var}_{i,t} = \sum_{s=t_0}^{t} \frac{(A_{i,s} - \hat{A}_{i,s})^2}{M} \tag{14}$$

for any point i with active cases A, over a sliding window with size $t - t_0$ including M time points. $\hat{A}$ is the estimator for the mean, from the Gaussian kernel. We also calculated other common EWS such as lag-1 autocorrelation AC(1), coefficient of variation (CV) and skewness, which were constructed similarly. The sampling frequency was not enough to estimate the reddening of the power spectrum [59] or the sample entropy [3]. All indicators were estimate with their corresponding MATLAB functions.
Since EWS are local measures around the transition point, we especially focused on their behaviour in its vicinity, thus looking for trends analogous to those reported in [4] and, in particular, in [7], Fig. 9 d–f and insets therein.

Estimation of the probability of exponential increase

From empirical data, the estimation of the critical point $R = 1$ is associated with a probability distribution. Bayesian frameworks like in [34, 4] allow estimation of posterior density functions $p(R|\text{data})$, which can be used to assess the probability of the control parameter to be greater than 1. Following the fact that $R > 1$ is associated with an exponential increase of infectious cases after a transcritial bifurcation, $\mathcal{P}(R > 1)$ can be interpreted as the probability of seeing an epidemic outbreak. Such probability is defined as:

$$\mathcal{P}(R > 1) = \int_1^{\infty} p(R'|\text{data}) dR' \tag{15}$$

It is compared with other EWS in Fig. **3**.

ROC curves

The Receiver Operator Characteristics (ROC) curve is a parametric plot of the sensitivity and specificity of a classification method as a function of the detection threshold [4, 60]. It allows scoring how successfully a particular indicator classifies whether or not an event is happening; in our case, if the trends of several EWS are able to detect an epidemic insurgence. As anticipated, the trend is quantified by the Kendall's τ score. As null (non emergent) data, we considered values of τ from 0, which corresponds to a flat trend, to 1, which corresponds to monotonous increase with constant regression coefficient. The test data were then τ values for $-30 < t < 5$ days before the transition. $t > -30$ was chosen to avoid significant overlaps with the first epidemic wave, $t < 5$ to account for possible small bifurcation delays. Then, we calculated the ROC using data associated to each time point separately, considering every country in the test set, to show how the detectability of emergence changes with time and threshold values.

Finally, we calculated the Area Under the Curve (AUC), which quantifies the overall performance of an indicator. A value of AUC$= 0.5$ means that the EWS is no better than a random guess in detecting the transition, while AUC> 0.5 implies better performance. This is used to compare various indicators, as well as the performance of the same indicators over the two different datasets $\mathcal{Y}$ and $\mathcal{N}$.

Author contribution

DP conceived and designed the study, collected the data, performed the analysis and interpreted the results. SM and FK performed the analysis and interpreted the results. JG supervised the project and interpreted the results. All authors wrote and approved the manuscript.

Declaration of interest

The authors declare no competing interests.

Acknowledgments

The authors thank the Research Luxembourg - COVID-19 Taskforce for mutual collaborations. DP's and SM's work is supported by the FNR PRIDE DTU CriTiCS, ref 10907093. FK's work is supported by the Luxembourg National Research Fund PRIDE17/12244779/PARK-QC. JG is partly supported by the 111 Project on Computational Intelligence and Intelligent Control, ref B18024.

Code availability

The analysis was performed in MATLAB and Python. They require the *Statistics Toolbox* and the *scipy* library, respectively. Code and curated data are accessible on:
`https://github.com/daniele-proverbio/EWS_epidemic`.

References

[1] Solomon Hsiang, Daniel Allen, Sébastien Annan-Phan, Kendon Bell, Ian Bolliger, Trinetta Chong, Hannah Druckenmiller, Luna Yue Huang, Andrew Hultgren, Emma Krasovich, et al. The effect of large-scale anti-contagion policies on the covid-19 pandemic. *Nature*, 584(7820):262–267, 2020.

[2] Leonhard Horstmeyer, Christian Kuehn, and Stefan Thurner. Network topology near criticality in adaptive epidemics. *Physical Review E*, 98(4):042313, 2018.

[3] Tobias S Brett, John M Drake, and Pejman Rohani. Anticipating the emergence of infectious diseases. *Journal of The Royal Society Interface*, 14(132):20170115, 2017.

[4] Tobias Brett, Marco Ajelli, Quan-Hui Liu, Mary G Krauland, John J Grefenstette, Willem G van Panhuis, Alessandro Vespignani, John M Drake, and Pejman Rohani. Detecting critical slowing down in high-dimensional epidemiological systems. *PLoS computational biology*, 16(3):e1007679, 2020.

[5] Paige B Miller, Eamon B O'Dea, Pejman Rohani, and John M Drake. Forecasting infectious disease emergence subject to seasonal forcing. *Theoretical Biology and Medical Modelling*, 14(1):1–14, 2017.

[6] Brendon Phillips, Madhur Anand, and Chris T Bauch. Spatial early warning signals of social and epidemiological tipping points in a coupled behaviour-disease network. *Scientific reports*, 10(1): 1–12, 2020.

[7] Suzanne M O'Regan and John M Drake. Theory of early warning signals of disease emergenceand leading indicators of elimination. *Theoretical Ecology*, 6(3):333–357, 2013.

[8] Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H Van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.

[9] Peter Ashwin, Sebastian Wieczorek, Renato Vitolo, and Peter Cox. Tipping points in open systems: bifurcation, noise-induced and rate-dependent examples in the climate system. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1962):1166–1184, 2012.

[10] Marten Scheffer and Stephen R Carpenter. Catastrophic regime shifts in ecosystems: linking theory to observation. *Trends in ecology & evolution*, 18(12):648–656, 2003.

[11] Christian Kuehn. A mathematical framework for critical transitions: normal forms, variance and applications. *Journal of Nonlinear Science*, 23(3):457–510, 2013.

[12] Judith Legrand, Rebecca Freeman Grais, Pierre-Yves Boelle, Alain-Jacques Valleron, and Antoine Flahault. Understanding the dynamics of Ebola epidemics. *Epidemiology & Infection*, 135(4):610–621, 2007.

[13] Emma Southall, Michael Tildesley, and Louise Dyson. Prospects for detecting early warning signals in discrete event sequence data: Application to epidemiological incidence data. *PLoS Comput Biol*, 16(9):e1007836, 2020.

[14] Tobias S Brett, Eamon B O'Dea, Éric Marty, Paige B Miller, Andrew W Park, John M Drake, and Pejman Rohani. Anticipating epidemic transitions with imperfect data. *PLoS computational biology*, 14(6):e1006204, 2018.

[15] Eamon B O'Dea and John M Drake. Disentangling reporting and disease transmission. *Theoretical Ecology*, 12(1):89–98, 2019.

[16] Tobias S Brett and Pejman Rohani. Dynamical footprints enable detection of disease emergence. *PLoS biology*, 18(5):e3000697, 2020.

[17] Jared Diamond, James A Robinson, et al. *Natural experiments of history*. Harvard University Press, 2010.

[18] Vasilis Dakos, Stephen R Carpenter, William A Brock, Aaron M Ellison, Vishwesha Guttal, Anthony R Ives, Sonia Kefi, Valerie Livina, David A Seekell, Egbert H van Nes, et al. Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data. *PloS one*, 7(7):e41010, 2012.

[19] Shunqing Xu and Yuanyuan Li. Beware of the second wave of covid-19. *The Lancet*, 395(10233): 1321–1322, 2020.

[20] Giacomo Cacciapaglia, Corentin Cot, and Francesco Sannino. Second wave covid-19 pandemics in europe: a temporal playbook. *Scientific reports*, 10(1):1–8, 2020.

[21] Steven H Strogatz. *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering*. CRC press, 2018.

[22] Vasilis Dakos, Marten Scheffer, Egbert H van Nes, Victor Brovkin, Vladimir Petoukhov, and Hermann Held. Slowing down as an early warning signal for abrupt climate change. *Proceedings of the National Academy of Sciences*, 105(38):14308–14312, 2008.

[23] Nils Berglund and Barbara Gentz. *Noise-induced phenomena in slow-fast dynamical systems: a sample-paths approach*. Springer Science & Business Media, 2006.

[24] Marten Scheffer, Stephen R Carpenter, Timothy M Lenton, Jordi Bascompte, William Brock, Vasilis Dakos, Johan Van de Koppel, Ingrid A Van de Leemput, Simon A Levin, Egbert H Van Nes, et al. Anticipating critical transitions. *science*, 338(6105):344–348, 2012.

[25] Christian Kuehn. A mathematical framework for critical transitions: Bifurcations, fast–slow systems and stochastic dynamics. *Physica D: Nonlinear Phenomena*, 240(12):1020–1035, 2011.

[26] Vasilis Dakos, Egbert H Van Nes, Paolo d'Odorico, and Marten Scheffer. Robustness of variance and autocorrelation as indicators of critical slowing down. *Ecology*, 93(2):264–271, 2012.

[27] Carl Boettiger and Alan Hastings. Early warning signals and the prosecutor's fallacy. *Proceedings of the Royal Society B: Biological Sciences*, 279(1748):4734–4739, 2012.

[28] Suzanne M O'Regan and Danielle L Burton. How stochasticity influences leading indicators of critical transitions. *Bulletin of mathematical biology*, 80(6):1630–1654, 2018.

[29] Lei Dai, Daan Vorselen, Kirill S Korolev, and Jeff Gore. Generic indicators for loss of resilience before a tipping point leading to population collapse. *Science*, 336(6085):1175–1177, 2012.

[30] Savannah F Pedersen, Ya-Chi Ho, et al. Sars-cov-2: a storm is raging. *The Journal of clinical investigation*, 130(5), 2020.

[31] Ensheng Dong, Hongru Du, and Lauren Gardner. An interactive web-based dashboard to track COVID-19 in real time. *The Lancet infectious diseases*, 20(5):533–534, 2020. doi: https://doi.org/10.1016/S1473-3099(20)30120-1.

[32] Ying Liu, Albert A Gayle, Annelies Wilder-Smith, and Joacim Rocklöv. The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine*, 27(2), 02 2020. ISSN 1708-8305. doi: 10.1093/jtm/taaa021. URL `https://doi.org/10.1093/jtm/taaa021`. taaa021.

[33] Christian L Althaus. Estimating the reproduction number of Ebola virus (EBOV) during the 2014 outbreak in West Africa. *PLoS Currents*, 6, 2014. doi: 10.1371/currents.outbreaks.91afb5e0f279e7f29e7056095255b288.

[34] Kevin Systrom, Thomas Vladek, and Mike Krieger. Rt.live. `https://github.com/rtcovidlive/covid-model`, 2020.

[35] S Abbott, J Hellewell, JD Munday, JY Chun, RN Thompson, NI Bosse, YWD Chan, TW Russell, CI Jarvis, CMMID nCov working group, et al. Temporal variation in transmission during the covid-19 outbreak [version 1; peer review: awaiting peer review]. *Wellcome Open Res*, 5(112), 2020. doi.org/10.12688/wellcomeopenres.16006.1.

[36] Howard Bauchner, Robert M Golub, and Jody Zylke. Editorial concern—possible reporting of the same patients with covid-19 in different reports. *Jama*, 323(13):1256–1256, 2020.

[37] Adjani Gama Dessavre, Emma Southall, Michael J Tildesley, and Louise Dyson. The problem of detrending when analysing potential indicators of disease elimination. *Journal of theoretical biology*, 481:183–193, 2019.

[38] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.

[39] Vishwesha Guttal and Ciriyam Jayaprakash. Changing skewness: an early warning signal of regime shifts in ecosystems. *Ecology letters*, 11(5):450–460, 2008.

[40] David Adam. A guide to R—the pandemic's misunderstood metric. *Nature*, 583(7816):346–348, 2020.

[41] Paul L Delamater, Erica J Street, Timothy F Leslie, Y Tony Yang, and Kathryn H Jacobsen. Complexity of the basic reproduction number (R0). *Emerging infectious diseases*, 25(1):1, 2019.

[42] Roy M Anderson and Robert M May. Population biology of infectious diseases: Part I. *Nature*, 280 (5721):361–367, 1979. doi: https://doi.org/10.1038/280361a0.

[43] Daniele Proverbio, Francoise Kemp, Stefano Magni, Andreas Dominik Husch, Atte Aalto, Laurent Mombaerts, Alexander Skupin, Jorge Goncalves, Jose Ameijeiras-Alonso, and Christophe Ley. Dynamical SPQEIR model assesses the effectiveness of non-pharmaceutical interventions against COVID-19 epidemic outbreaks. *medRxiv*, 10.1101/2020.04.22.20075804, 2020.

[44] Joseph T Wu, Kathy Leung, and Gabriel M Leung. Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study. *The Lancet*, 395(10225):689–697, 2020. doi: https://doi.org/10.1016/S0140-6736(20)30260-9.

[45] Liangrong Peng, Wuyue Yang, Dongyan Zhang, Changjing Zhuge, and Liu Hong. Epidemic analysis of covid-19 in china by dynamical modeling. *arXiv*, preprint arXiv:2002.06563, 2020.

[46] Seth Flaxman, Swapnil Mishra, Axel Gandy, H Juliette T Unwin, Thomas A Mellan, Helen Coupland, Charles Whittaker, Harrison Zhu, Tresnia Berah, Jeffrey W Eaton, et al. Estimating the effects of non-pharmaceutical interventions on covid-19 in europe. *Nature*, 584(7820):257–261, 2020.

[47] David Alonso, Alan J McKane, and Mercedes Pascual. Stochastic amplification in epidemics. *Journal of the Royal Society Interface*, 4(14):575–582, 2007.

[48] Nicolaas Godfried Van Kampen. *Stochastic processes in physics and chemistry*, volume 1. Elsevier, 1992.

[49] Yuri A Kuznetsov. *Elements of applied bifurcation theory*, volume 112. Springer Science & Business Media, 2013.

[50] John Lee. *Introduction to topological manifolds*, volume 202. Springer Science & Business Media, 2010.

[51] Linda JS Allen. An introduction to stochastic epidemic models. In *Mathematical epidemiology*, pages 81–130. Springer, 2008.

[52] Giulia Giordano, Franco Blanchini, Raffaele Bruno, Patrizio Colaneri, Alessandro Di Filippo, Angela Di Matteo, and Marta Colaneri. Modelling the covid-19 epidemic and implementation of population-wide interventions in italy. *Nature Medicine*, 26:1–6, 2020.

[53] Max Roser, Hannah Ritchie, Esteban Ortiz-Ospina, and Joe Hasell. Coronavirus Pandemic (COVID-19). *Our World in Data*, https://ourworldindata.org/coronavirus, 2020.

[54] Bruno Alessandro Rivieccio, Ester Luconi, Patrizia Boracchi, Elena Pariani, Luisa Romanò, Silvia Salini, Silvana Castaldi, Elia Biganzoli, and Massimo Galli. Heterogeneity of covid-19 outbreak in italy. *Acta Bio Medica: Atenei Parmensis*, 91(2):31, 2020.

[55] TM Lenton, VN Livina, V Dakos, EH Van Nes, and M Scheffer. Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1962):1185–1204, 2012.

[56] Itay Ricon-Becker, Ricardo Tarrasch, Pablo Blinder, and Shamgar Ben-Eliyahu. A seven-day cycle in covid-19 infection, hospitalization, and mortality rates: Do weekend social interactions kill susceptible people? *medRxiv*, 10.1101/2020.05.03.20089508, 2020. URL `https://www.medrxiv.org/content/early/2020/10/01/2020.05.03.20089508`.

[57] JR Taylor. *An Introduction to Error Analysis - University Science Books*. Mill Valley, California, 1997.

[58] Hiroshi Nishiura, Natalie M. Linton, and Andrei R. Akhmetzhanov. Serial interval of novel coronavirus (covid-19) infections. *International Journal of Infectious Diseases*, 93:284 – 286, 2020. ISSN 1201-9712. doi: https://doi.org/10.1016/j.ijid.2020.02.060. URL `http://www.sciencedirect.com/science/article/pii/S1201971220301193`.

[59] Reinette Biggs, Stephen R Carpenter, and William A Brock. Turning back from the brink: detecting an impending regime shift in time to avert it. *Proceedings of the National academy of Sciences*, 106(3): 826–831, 2009.

[60] Tom Fawcett. An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874, 2006.

[61] C. Reno, J. Lenzi, A. Navarra, E. Barelli, D. Gori, A. Lanza, R. Valentini, B. Tang, and M. P. Fantini. Forecasting COVID-19-Associated Hospitalizations under Different Levels of Social Distancing in Lombardy and Emilia-Romagna, Northern Italy: Results from an Extended SEIR Compartmental Model. *ournal of clinical medicine*, 9:1492, 2020. doi: 10.3390/jcm9051492.

[62] Robert C. Reiner, Ryan M. Barber, James K. Collins, Peng Zheng, Christopher Adolph, James Albright, Catherine M. Antony, Aleksandr Y. Aravkin, Steven D. Bachmeier, Bree Bang-Jensen, Marlena S. Bannick, Sabina Bloom, Austin Carter, Emma Castro, Kate Causey, Suman Chakrabarti, Fiona J. Charlson, Rebecca M. Cogen, Emily Combs, Xiaochen Dai, William James Dangel, Lucas Earl, Samuel B. Ewald, Maha Ezalarab, Alize J. Ferrari, Abraham Flaxman, Joseph Jon Frostad, Nancy Fullman, Emmanuela Gakidou, John Gallagher, Scott D. Glenn, Erik A. Goosmann, Jiawei He, Nathaniel J. Henry, Erin N. Hulland, Benjamin Hurst, Casey Johanns, Parkes J. Kendrick, Apurva Khemani, Samantha Leigh Larson, Alice Lazzar-Atwood, Kate E. LeGrand, Haley Lescinsky, Akiaja Lindstrom, Emily Linebarger, Rafael Lozano, Rui Ma, Johan Månsson, Beatrice Magistro, Ana M. Mantilla Herrera, Laurie B. Marczak, Molly K. Miller-Petrie, Ali H. Mokdad, Julia Deryn Morgan, Paulami Naik, Christopher M. Odell, James K. O'Halloran, Aaron E. Osgood-Zimmerman, Samuel M. Ostroff, Maja Pasovic, Louise Penberthy, Geoffrey Phipps, David M. Pigott, Ian Pollock, Rebecca E. Ramshaw, Sofia Boston Redford, Grace Reinke, Sam Rolfe, Damian Francesco Santomauro, John R. Shackleton, David H. Shaw, Brittney S. Sheena, Aleksei Sholokhov, Reed J. D. Sorensen, Gianna Sparks, Emma Elizabeth Spurlock, Michelle L. Subart, Ruri Syailendrawati, Anna E. Torre, Christopher E. Troeger, Theo Vos, Alexandrea Watson, Stefanie Watson, Kirsten E. Wiens, Lauren Woyczynski, Liming Xu, Jize Zhang, Simon I. Hay, Stephen S. Lim, Christopher J. L. Murray, and IHME COVID-19 Forecasting Team. Modeling COVID-19 scenarios for the United States. *Nature Medicine*, 27:94–105, 2020. doi: 10.1038/s41591-020-1132-9. URL `https://doi.org/10.1038/s41591-020-1132-9`.

[63] Marco Ajelli, Stefano Merler, Laura Fumanelli, Ana Pastore y Piontti, Natalie E Dean, Ira M Longini, M Elizabeth Halloran, and Alessandro Vespignani. Spatiotemporal dynamics of the ebola epidemic in guinea and implications for vaccination and disease elimination: a computational modeling analysis. *BMC medicine*, 14(1):1–10, 2016.

[64] Quan-Hui Liu, Marco Ajelli, Alberto Aleta, Stefano Merler, Yamir Moreno, and Alessandro Vespignani. Measurability of the epidemic reproduction number in data-driven contact networks. *Proceedings of the National Academy of Sciences*, 115(50):12680–12685, 2018.

[65] Luis MA Bettencourt and Ruy M Ribeiro. Real time bayesian estimation of the epidemic potential of emerging infectious diseases. *PloS one*, 3(5):e2185, 2008.

[66] Zhanwei Du, Xiaoke Xu, Ye Wu, Lin Wang, Benjamin J Cowling, and Lauren Ancel Meyers. Serial interval of covid-19 among publicly reported confirmed cases. *Emerging infectious diseases*, 26(6): 1341, 2020.

[67] Minah Park, Alex R Cook, Jue Tao Lim, Yinxiaohe Sun, and Borame L Dickens. A systematic review of covid-19 epidemiology based on current evidence. *Journal of clinical medicine*, 9(4):967, 2020.

A Data collection and curation: examples

Among all the countries that registered a resurgence of COVID-19 epidemic between April and September 2020, we first selected those for which meaningful prevalence data could be directly obtained from official sources or reconstructed with Eq. 12 in Main Text. Examples of discarded data series are reported in Fig. 5. We recall the underlying hypothesis of this work: that dynamical early warning signals are expected to work when the investigated system can be described by a proper dynamical model. Hence, we did not consider time series for which active cases do not display the typical SEIR-like behavior *e.g.* [61, 62, 43] (mostly due to data management and reporting), or for which recovered cases are reported with different frequencies, resulting in sawtooth curves for active cases. In the later case, the detrended fluctuations would be associated to reporting standards and would not be representative of dynamical fluctuations associated to EWS. Other selecting criteria to increase the quality of the dataset are discussed in "Materials and Methods" of the Main Text.

Curves of active cases for all countries considered in Main Text, their smoothing and the associated $R(t)$ are displayed in Fig. 6. The smoothed curves serve as basis for the detrending, the analysis of the noise distribution of each country and the subsequent investigation of early warning signals, as explained in Main Text.

B Estimating $\mathrm{R(t)}$ with Bayesian MCMC

Following standard methodologies [35, 4], we reconstruct the day-by-day evolution of the reproduction number $R(t)$ by fitting a Poisson transmission model with Markov Chain Monte Carlo methods.

When modelling "arrivals" of discrete-state stochastic processes (*cf.* "Materials and Methods" in Main Text), Poisson processes are widely employed. For instance, it effectively modeled the transmission of Ebola [63] and Influenza [64]. It states that, given an average rate of λ new cases per day, the probability of seeing k new cases is distributed according to the Poisson distribution:

$$P(k|\lambda) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{16}$$

In turn, λ depends on R as [65]:

$$\lambda_t = k_{t-1} e^{\gamma(R_t - 1)} \tag{17}$$

where γ is the reciprocal of the serial interval, which is around 4 days for COVID-19 [66, 67]. To account for such uncertainty, we treat γ as a random sample from a Gaussian distribution centered in 4 days with standard deviation 0.2. Hence, the probability of observing a time series $x = \{x_t\}$ between t_0 and T is given by:

$$P(x|R) = \prod_{t=t_0}^{T-\delta} P(k_{t+\delta}|\lambda_t) \tag{18}$$

for each time step δ. Before t_0, no cases were reported. Following Bayes' rule, the posterior distribution of R, for each time point, is given by (up to a normalization constant):

$$P(R|x) \propto P(x|R)q(R), \tag{19}$$

where $q(R)$ is a prior distribution. For each time point after $t_0 + 1$, the prior equals the preceding posterior. We follow he implementation of [34] to generate thousands of MCMC samples with Metropolis Hastings, starting with a Gaussian prior $\mathcal{N}(R, \sigma)$. We used $\sigma = 0.15$ as it maximizes the log-likelihood over every state in the current implementation. From the posterior distribution, we also estimated the probability that $R(t)$ is greater than 1, which was in turn used to define the "ground truth" date of regime transition for the epidemic trend (see also "Materials and Methods" of Main Text for further discussion). Fig. 6 shows the results of the Bayesian $R(t)$ estimation (median values and 50% Credible Intervals) for the considered countries.

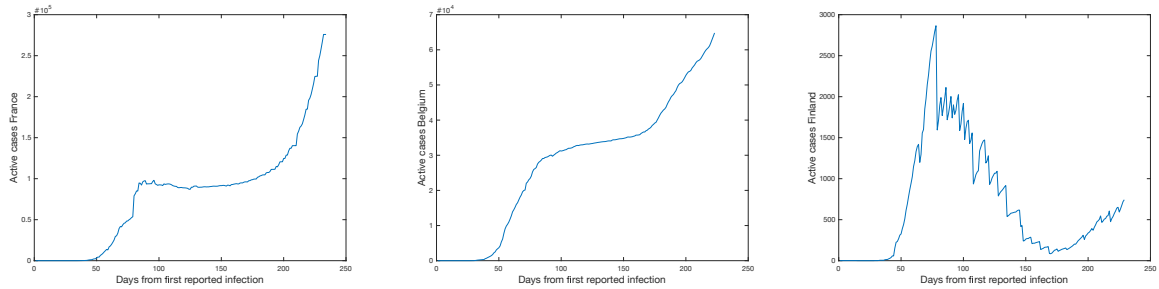

Fig. 5 Examples of discarded time series, following some of the criteria explained in "Material and Methods" of Main Text. Among others, France and Belgium had active cases curves that differ from the typical bell-shaped SEIR-like behavior. This is related to reporting of recovered and dead patients. On the other hand, Finland is an example of sawtooth evolution, due to recovered cases being reported with different frequencies than daily cases. Data from [31].

C Determining the rate of approach for $\mathbf{R(t) \to 1}$

As explained in the Main Text, we are interested in evaluating how rapidly the transition point is approached, for each country. To do so, the simplest linear trend

$$y = a + b \cdot x, \tag{20}$$

corresponding to Eq. 13 in Main Text, is assumed and fitted to $R(t)$ time series, obtained as described above and reported in Fig. 6. The estimated regression coefficient is informative about the rate of approach of $R(t) \to 1$.

The fitting was performed with *scipy.optimize* routine, considering as uncertainty the 50% credible interval from the distribution of $R(t)$, which is representative of one standard deviation. The goodness of fit was evaluated with the reduced χ^2 score. Results are displayed in Fig. 7. The $\chi^2_{\text{red}} < 1$ guarantees the goodness of the linear fit, which allows to extrapolate the regression coefficient as a measure of the rate of $R(t) \to 1$. The values of $b \pm \sigma_b$ are then reported in Fig. 1 of Main Text.

D Evolution of EWS for all countries

In this section, we show and discuss the evolution of the considered early warning indicators for all countries, including those that are not shown in Fig. 3 of Main Text. In Fig. 8 we observe the evolution of the indicators, either globally (8a) or locally, just prior to the bifurcation (8b for the variance, the most robust one as discussed in Main Text). We can observe the patterns discussed in the Main Text, associated to the different countries belonging to the test set $\mathcal{Y}$ or not ($\mathcal{N}$). Within $\mathcal{Y}$, the variance increases prior to the transition and gives very few spurious signals before, whereas other indicators can be more misleading when the transition still did not happen. Overall, predicted trends of EWS prior to the bifurcation are associated to satisfying theoretical assumptions such as gradual approach of $R(t) \to 1$ and white noise (*cf.* Fig. 1 in Main Text and discussion thereafter). Not satisfying these requirements might disrupt the expected increasing trend and results in misleading signals, see in Fig. 8 the countries listed in $\mathcal{N}$. Hence, if a system is not known or there is difficulty in determining the type of data, incorrect conclusions could be drawn when interpreting the time series trend. Fig. 4 and Table 2 of Main Text then quantify the performance of EWS for all countries.

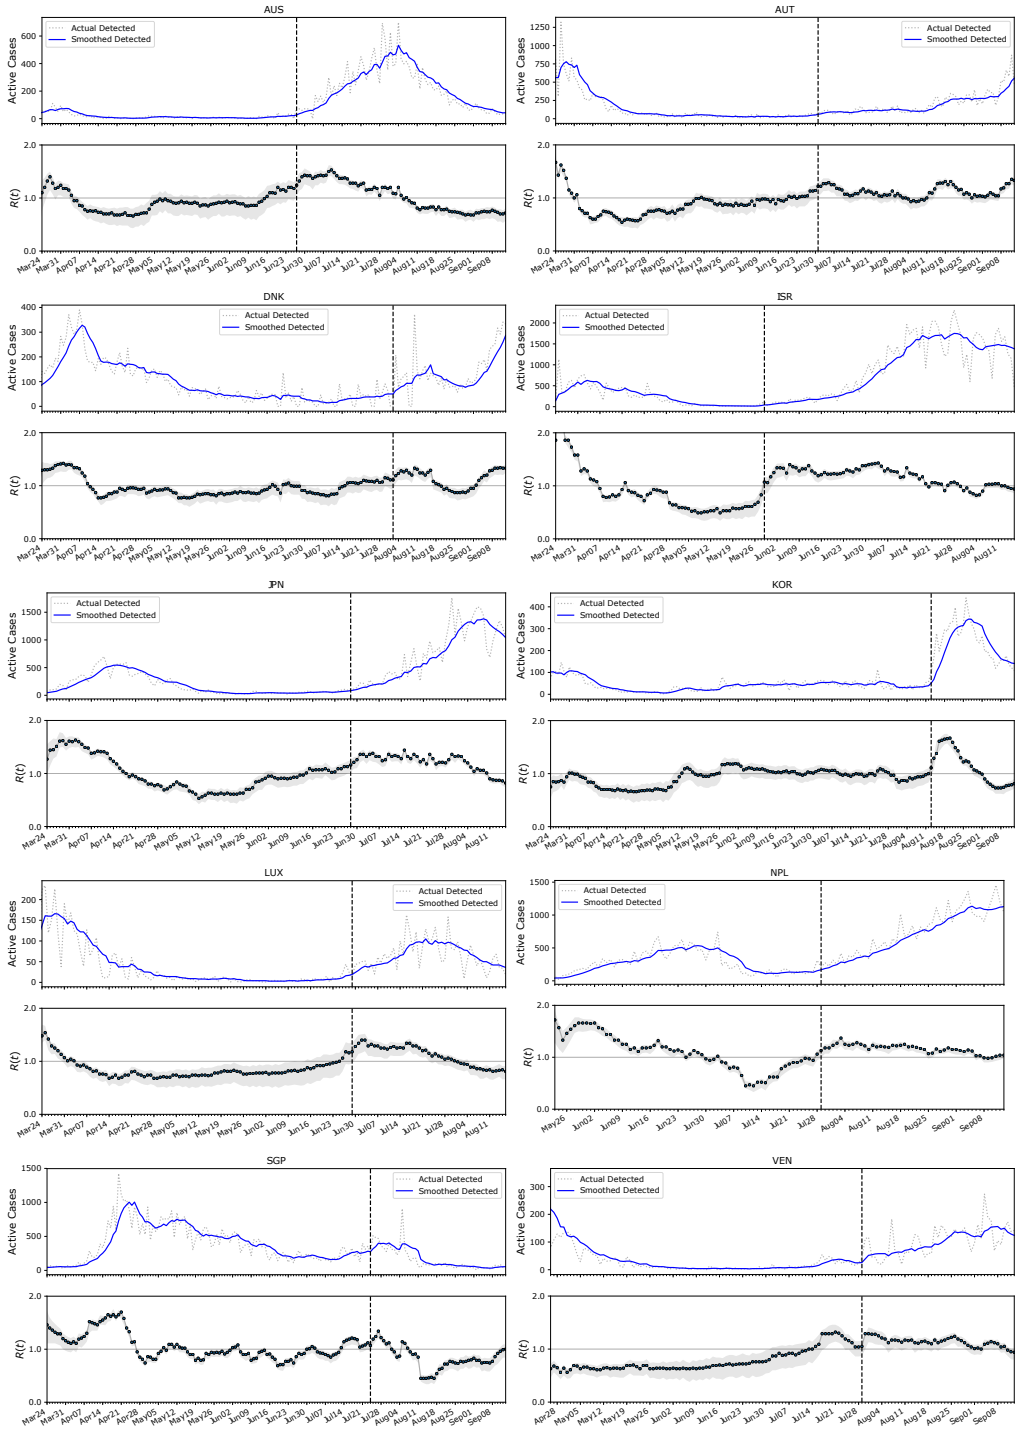

Fig. 6 Curves of active cases for the considered countries, along with their associated $R(t)$ (median values and 50% Credible Intervals). The vertical dashed line identifies the day marked for the transition an reported in Table 1 of Main Text.

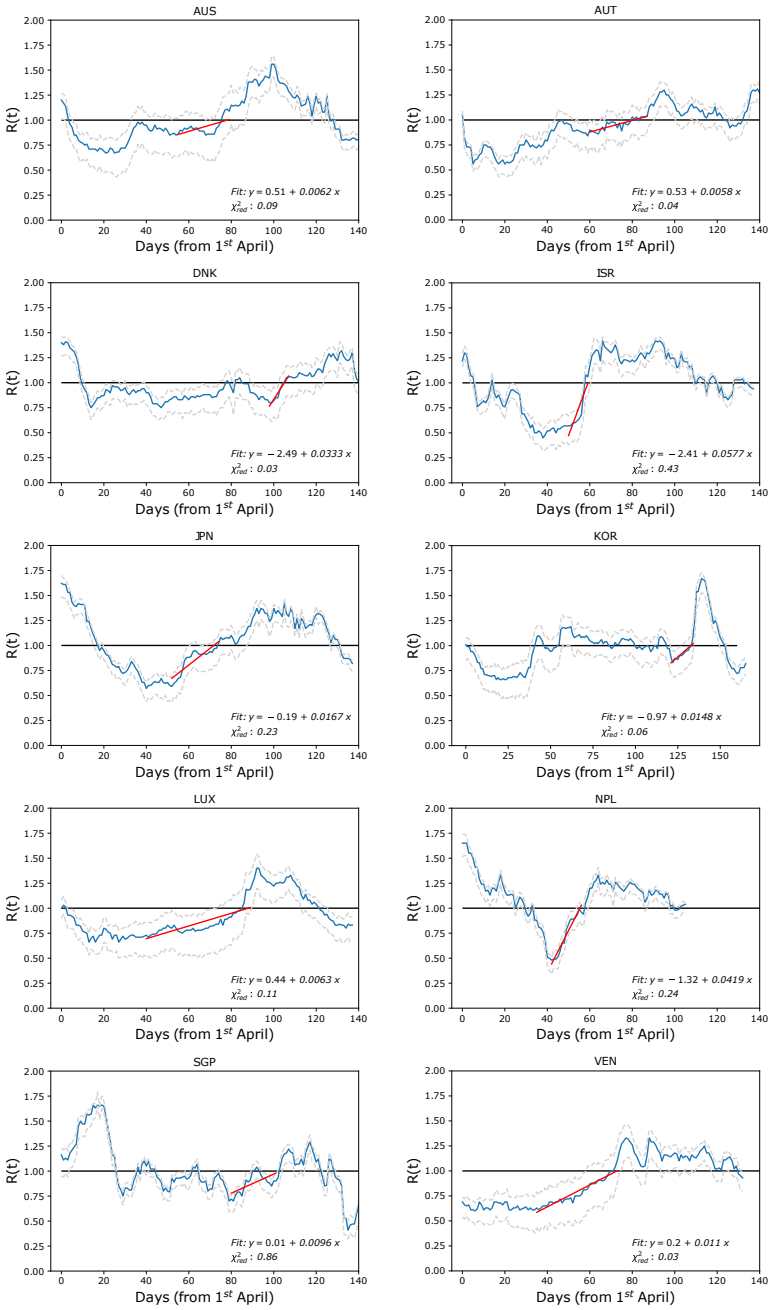

Fig. 7 Fitting $R(t)$ evolution (blue line; dashed lines are $\pm 50\%$ CI) with a linear trend prior to the transition, to estimate the rate of approach to the threshold value 1. The fit begins around the minimum of $R(t)$ (excluding small fluctuations) until when the median value crosses 1 (horizontal line). The best fit curve is in red. $\chi^2_{\text{red}} < 1$ guarantees that a good fit is achieved for the simple linear function Eq. 20. The regression coefficient b is a proxy for the rate of approach to 1. Its associated uncertainty is not reported here but is shown in Fig. 1 of Main Text.

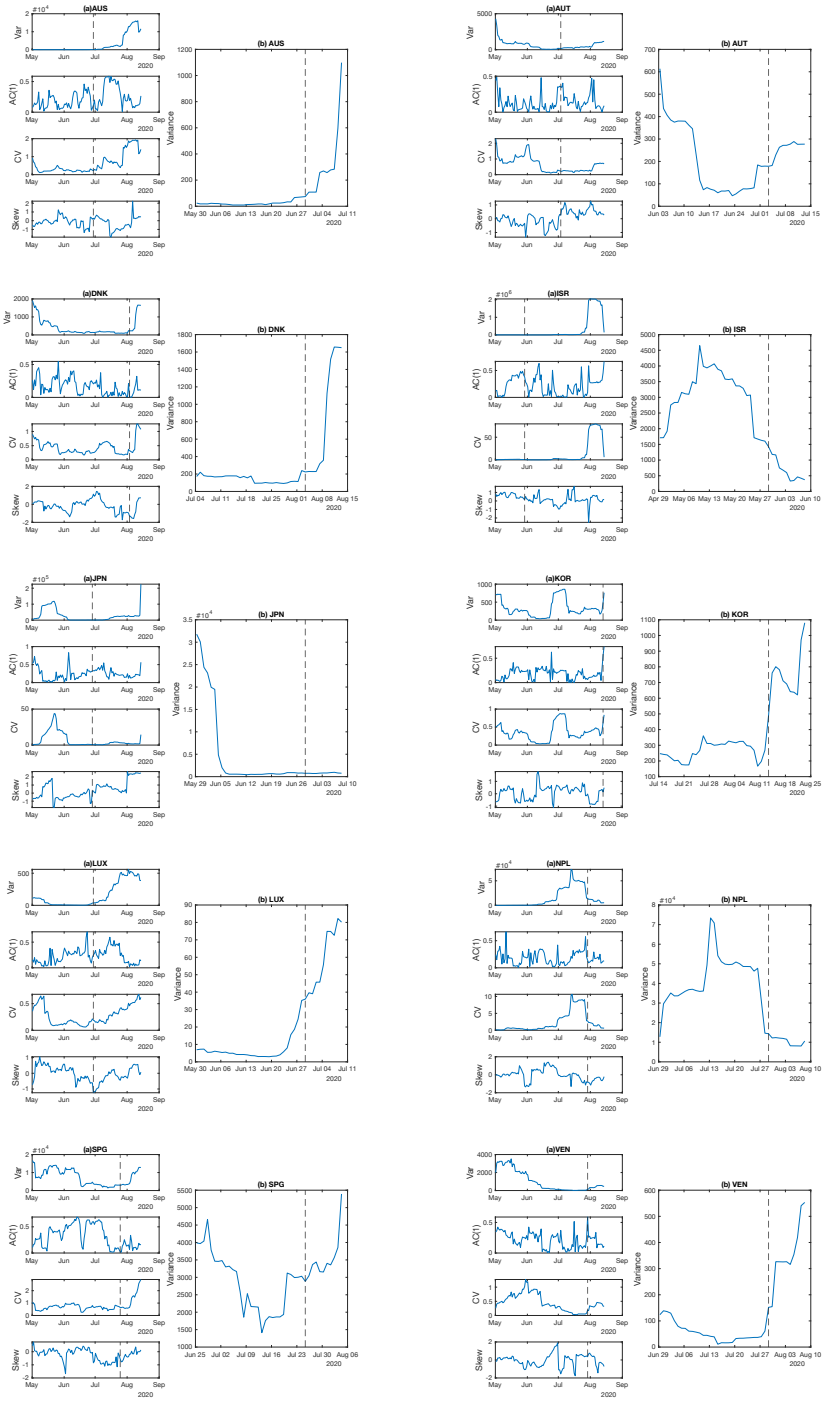

Fig. 8 Evolution of the considered indicators for all countries. Figs (a) report their global evolution, from the end of the first wave to the second. Figs (b) focus on the behavior of the variance close to the transition (local behavior). As discussed in the Main Text, we report the values of indicators after a moving window of size 14 days, associated to the rightmost data point to avoid "looking into the future". The vertical dashed line marks the transition point.