Research article

Epidemiological waves - Types, drivers and modulators in the COVID-19 pandemic

John Harvey [a,b,1,*], Bryan Chan [c,1], Tarun Srivastava [d,1], Alexander E. Zarebski [e], Paweł Dłotko [f], Piotr Błaszczyk [g,h], Rachel H. Parkinson [e], Lisa J. White [i], Ricardo Aguas [j], Adam Mahdi [k]

[a] Department of Mathematics, Swansea University, Swansea, UK
[b] School of Mathematics, Cardiff University, UK
[c] Department of Economics, London School of Economics and Political Science, London, UK
[d] Department of Engineering Science, University of Oxford, Oxford, UK
[e] Department of Biology, University of Oxford, Oxford, UK
[f] Dioscuri Centre in Topological Data Analysis, Mathematical Institute, Polish Academy of Sciences, Warsaw, Poland
[g] Faculty of Computer Science, Electronics and Telecommunications, AGH University of Science and Technology, Krakow, Poland
[h] Oxford Internet Institute, University of Oxford, Oxford, UK
[i] Li Ka Shing Centre for Health Information and Discovery, Big Data Institute, University of Oxford, Oxford, UK
[j] Nuffield Department of Medicine, Mahidol-Oxford Tropical Medicine Research Unit, University of Oxford, Oxford, UK
[k] Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford, UK

ARTICLE INFO

ABSTRACT

Introduction: A discussion of 'waves' of the COVID-19 epidemic in different countries is a part of the national conversation for many, but there is no hard and fast means of delineating these waves in the available data and their connection to waves in the sense of mathematical epidemiology is only tenuous.

Methods: We present an algorithm which processes a general time series to identify substantial, significant and sustained periods of increase in the value of the time series, which could reasonably be described as 'observed waves'. This provides an objective means of describing observed waves in time series. We use this method to synthesize evidence across different countries to study types, drivers and modulators of waves.

Results: The output of the algorithm as applied to epidemiological time series related to COVID-19 corresponds to visual intuition and expert opinion. Inspecting the results of individual countries shows how consecutive observed waves can differ greatly with respect to the case fatality ratio. Furthermore, in large countries, a more detailed analysis shows that consecutive observed waves have different geographical ranges. We also show how waves can be modulated by government interventions and find that early implementation of NPIs correlates with a reduced number of observed waves and reduced mortality burden in those waves.

Conclusion: It is possible to identify observed waves of disease by algorithmic methods and the results can be fruitfully used to analyse the progression of the epidemic.

* Corresponding author. School of Mathematics, Cardiff University, UK.
  E-mail address: harveyj13@cardiff.ac.uk (J. Harvey).
[1] JH, BC and TS contributed equally to the authorship of this manuscript.

1. Introduction

The COVID-19 pandemic has thrust epidemiology into the public consciousness. Outbreaks, epidemic peaks and waves of transmission are widely discussed. However, we lack a consensus on the definition of many of these terms [27]. Depending upon context, the meaning of the term 'epidemic wave' can vary from a well-defined property of a mathematical object through to a loosely defined portion of a time series. Despite definitional problems, using these descriptive terms has value for planning and public health [26]. With Severe Acute Respiratory Syndrome Coronavirus 2 (SARS-CoV-2) believed likely to become endemic [21] and join the ranks of other pathogens which regularly cause epidemics, a reflection on what constitutes an epidemic wave can help to clarify our thoughts and discussion [22].

SARS-CoV-2 has spread across the world after emerging in early December 2019 in Wuhan, China (Fig. 1) [17]. Governments across the world implemented non-pharmaceutical interventions (NPIs) with different levels of stringency and speed in an attempt to prevent and/or minimise the importation and local transmission of the virus [23]. Unfortunately, these NPIs often have a substantial associated cost and so it is important to understand how best to reduce transmission in a cost efficient manner. Understanding the epidemic in any single country is a challenge given the myriad potential sources of regional heterogeneities; making meaningful comparisons between countries is even harder. Before we can critically engage in comparative exercises that explore how the size of an epidemic wave relates to factors such as socio-economic status, the stringency of interventions, and other health indicators, we must clearly define what a wave is.

In this paper we make three contributions aimed at resolving this issue. First, to facilitate more nuanced discussion, we make explicit the different ways in which researchers use the term 'epidemic wave'. Second, we provide an algorithm, with a public domain implementation, to partition epidemic time series (of confirmed cases and deaths) into non-overlapping 'observed waves', as defined below. We stress that this is not the provision of yet another definition of an epidemic wave; rather it is an exercise in emphasising some of the features any plausible definition should have, making explicit the ones we have chosen and making it possible to implement that definition consistently and at scale. Third, we consider different circumstances that may lead to the formation of observed waves, such as changes in transmission or case ascertainment, and geographic aggregation and in doing so we provide a more nuanced interpretation of the data.

The importance of distinguishing different waves and their types in epidemiology is critical for understanding the long term expected dynamics which in itself determines what health policies are optimal.

The sequence of waves provides critical information on the underlying characteristics of transmission of the pathogen in human populations. An analysis of the drivers of successive waves can unveil how the pathogen is evolving, how that evolution enables immune escape, and whether transmission is seasonal or how transmission is affected by interventions, for example.

The implications for policy are immediate: if a virus is non-seasonal, evolution is slow and interventions are impactful, then policies should make a strong push for early elimination; if transmission is highly seasonal, evolution is fast and interventions have limited impact, then transmission is likely to become endemic and policies should focus on mitigation and mortality management.

By considering case studies, we outline how both transmission and testing act as drivers of 'observed waves', using the time series of cases and the time series of deaths to understand the relative importance of these drivers. We also observe how a wave structure can be driven by the aggregation of data across a large geographical area.

We posit that, when comparing epidemic curves across countries, it is beneficial to do so at the level of observed waves (as
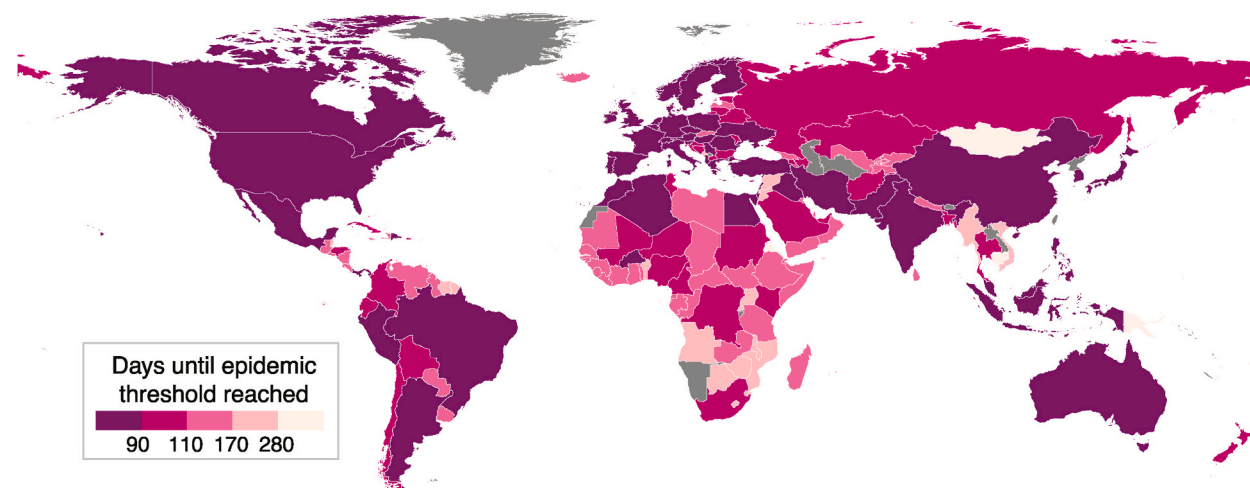

Fig. 1. Choropleth showing the number of days since the emergence of the first cases in China on the 31st of December 2019, until the cumulative number of deaths in each country surpassed 10. Countries with darker colours passed the threshold earlier than the lighter coloured countries. After starting in China, epidemics occurred in Europe, the Middle East and North America before moving south to South America, Africa and the Pacific. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

identified by our algorithm), in order to remove biases inherent to changing practices over time, such as testing effort and intervention stringency. This may also reduce bias due to between-country differences in measurement, as the detection of the presence of a wave is less sensitive to measurement practices than the use of absolute numbers.

We provide an example of how to use annotated time series to study the modulators of waves. Possible modulators could include the effects of seasonality, shifts in population behavior in response to infection levels, and the implementation by governments of NPIs. In this work we focus on the latter and carry out a global comparative analysis of the influence of NPIs and testing on the size of observed waves. This analysis aims to determine whether stringent and/or swift interventions reduced the burden of the epidemic.

2. Methods

As a comparative study of the effects of COVID-19 internationally, using novel methods to synthesize evidence from each country's data, this study has some similarities to a systematic review with meta-analysis. For this reason, the reporting of this study is guided by the standards of the Preferred Reporting Items for Systematic Review and Meta-Analysis (PRISMA) Statement [20], but given the clear differences from a standard systematic review, it is frequently necessary to deviate from the statement.

2.1. Data collection

We extracted data from the OxCOVID19 Database [18], which contains geographically unified information on epidemiology, government response, demographics, mobility and weather at a national and sub-national level collected from various sources, as well as from the Our World in Data COVID-19 dataset [24]. Data were last extracted on 23 June 2021, with a cutoff for time series data on or before 31 May 2021 to account for incomplete data due to delays in reporting.

From the OxCOVID19 database we selected epidemiological data from the European Centre for Disease Prevention and Control [1], obtaining the number of cases and deaths reported on each day.

To track measures imposed by governments we also drew from the OxCOVID19 database data curated by researchers from the Blavatnik School of Government, University of Oxford [11], obtaining the values of their 'stringency index' for each country as well as whether any of the following were in force: school closure, workplace closure, cancellation of public events, restrictions on gathering, public transport closure, stay-at-home requirements, restrictions on internal movement, international travel controls, testing polices and contact tracing systems.

Information on testing was obtained from the Our World in Data COVID-19 dataset [24]. Note that the statistics provided by each country may differ in whether they refer to the number of tests conducted or the number of individuals tested. This provided the number of tests performed per day, with a seven-day smoothing applied.

2.2. Synthesis methods

Missing daily values for cumulative confirmed cases and deaths [1] were filled in by linear interpolation. The number of new cases and deaths per day was then computed as the difference between the cumulative values for successive days. Negative values were replaced with the last non-negative observation. Such values commonly arise when reporting authorities subsequently correct their figures for total cases or total deaths. After this initial cleaning substantial stochasticity is still present in the time series, due to factors such as backlogs in the number of cases over weekends and errors in consolidating municipal sources. To better understand the underlying trend and identify significant waves, we computed a fourteen-day moving average to smooth the data.

To account for differences in epidemic start date, measures of time are calculated relative to the date at which each country first reached 10 cumulative deaths (referred to in this paper as start date T_0). Reducing this threshold would give undue weight to importation events at the expense of locally acquired infections. Smaller countries (population <2.5 million) were excluded from the analysis as T_0 was often very late in these countries.

The data series for the number of performed tests is provided with this seven-day smoothing already performed and we use the smoothed value directly [24]. We computed the positivity rate as a ratio of new cases per day to new tests per day (using raw data in both cases), and subsequently smoothed this rate with a seven-day moving average. Earlier in the time series positivity rates greater than 1 occasionally arise; these were removed.

No pre-processing was required for the Blavatnik School data [11].

The times series for cases and deaths were then analyzed using our novel wave detection method. This allowed us to separate the time series into 'waves', by which we mean significant, substantial and sustained increases. The methodology and a more in-depth discussion of the problem of wave identification are given in Section 2.3.

Some epidemiological and socioeconomic characteristics were analyzed according to the number of waves a country has experienced. Mann–Whitney tests were used to determine statistical significance of these characteristics.

A small number of in-depth case studies were carried out to better illustrate the types and drivers of waves in different countries. Two of these which demonstrate the main features are reported here.

The stringency and speed of the introduction of NPIs and large-scale testing were compared to the severity of epidemic waves (as measured by the per capita death rate during a wave identified by our algorithm). This analysis aims to determine whether stringent and/or swift interventions reduced the burden of the epidemic by using Kendall's rank correlation to measure the association between the measures taken and the severity of the burden.

2.3. Wave detection method

A wave in a time series is informally understood to be a significant, substantial and sustained increase in values followed by a significant and substantial decrease - a temporary increase of at least a certain minimal height which lasts for a certain minimal duration. Both those parameters depend on the context of the problem and, even in a fixed context, practitioners will have legitimate disagreements about whether a given period of a time series is indeed a 'wave'.

In the context of the epidemiology of infectious diseases, an epidemic wave is a clearly understood feature of mathematical models of transmission dynamics. A wave is characterized by early exponential growth in the number of infected agents, during which the effective reproduction number $R_t > 1$, followed by a saturation phase as R_t approaches 1. In an unmitigated epidemic $R_t = 1$ is a threshold condition that corresponds to the moment in which the population reaches herd immunity, after which R_t becomes lower than 1 and cases decrease over time. In systems with mechanisms of fast susceptibility replenishment (loss of immunity) and time dependent transmission modulators (seasonality and epochal evolution), a population can fluctuate around the herd immunity threshold, leading to several epidemic waves over time (e.g. flu epidemics) [10,16]. In such cases, an epidemic wave can be defined as the period between two local minima. If the model results in the infection rate converging to a stable background level, we can consider the period after the final local minimum to be the final wave of the model.

In the real world, where health surveillance systems are struggling to keep up with an unfolding pandemic, case and death notification definitions are volatile and/or dependent on testing capacity, making it hard to detect when an epidemic wave begins or ends [6]. Furthermore, during the current COVID-19 pandemic, unlike in unmitigated epidemics, changes in social behaviour in response to the development of the pandemic, both as a result of government restrictions and individual choice, have drastically changed the
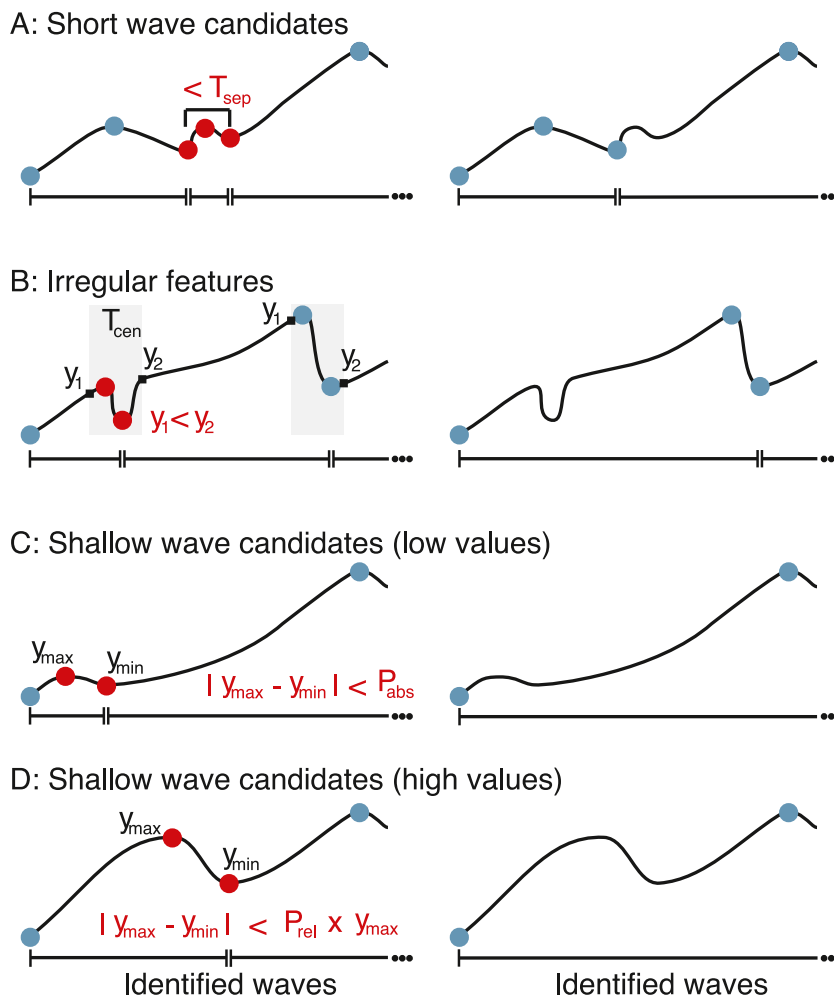

Fig. 2. The wave identification algorithm. Time series showing local maxima and minima (solid dots) before (left panels) and after (right panels) filtering with each wave identification algorithm (**A**: Sub-algorithm A; **B**: Sub-algorithm B; **C**: Sub-algorithm C; **D**: Sub-algorithm D). Red dots represent local minima and maxima that are removed with each algorithm, while those marked with blue dots are retained. Broken horizontal lines beneath each time series show the identified wave durations. Further detail on how each sub-algorithm operates is given in the Supplementary Material. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

intensity of transmission, R_t, and re-shaped the underlying transmission networks in affected countries. These changes are exogenous to standard epidemiological models, rendering the concept of a wave more ambiguous and raising the need for a clearer discourse around what a wave is, especially given how epidemic waves have been brought to the forefront of public discussion when justifying different levels of policy restriction.

In order to remove the ambiguity surrounding time series of cases and deaths, we propose a robust algorithm for identifying waves (understood as significant, substantial and sustained increases) in time series and apply it to time series of cases and deaths. We do not test for the presence of a significant decrease after the peak, assuming that if a country is currently experiencing an elevated level of transmission it is undergoing a wave which will eventually terminate. As a consequence, this means ongoing waves are still identified.

We describe these as 'observed waves', since they differ from the theoretical waves of epidemiological models and, due to measurement error, may also not accurately reflect the underlying SARS-CoV-2 transmission intensity. We consider the essential features of a 'wave' and use them to define algorithms which process the many brief rises and falls in case numbers over time and translate them into meaningful 'observed waves'. For the remainder of the paper, the term 'waves' will refer to 'observed waves'.

Since both cases and deaths can inform the decision about when a wave of COVID-19 begins and ends, we use both series together. Suppose that, during the time period covered by some wave in the deaths time series, no peak occurs in the cases time series. On the assumption that the recording of deaths is less subject to change, we assume that a wave of cases should in fact have been recorded. We therefore identify an appropriate peak which was filtered out earlier by the algorithm and restore it to the list of peaks in order to identify a wave of cases matching the wave of deaths. This mitigates concerns about changes in testing over time.

2.3.1. Prominence

Many algorithms exist to identify 'spikes' in time series, as this is a key challenge in signal processing [30]. These algorithms identify peaks of brief pulse-like signals. In contrast, epidemic waves are likely to be broad, so that algorithms to identify spikes are not applicable in this context. Instead we aim to identify waves using a prominence-based approach. Prominence can be informally defined, geometrically, as a measurement of how far a peak extends in the vertical direction above the surrounding regions of the time series. The advantage of a prominence-based approach is that it is agnostic as to both the width (i.e., duration) of the peak and the presence of additional local maxima on the shoulder of the peak. This methodology, which stems from recent developments in topological data analysis, has previously been used to analyse medical time series [5,9].

Prominence may be understood as follows. Consider a fully submerged mountain range, and imagine that the water level is gradually dropping. As the water level falls, 'islands' become visible. These are equivalent to portions of the time series that have values above a certain threshold. The global maximum will be seen first, followed by the neighboring values in the time series. As the water level drops further, a second island becomes visible, around the second highest local maximum. As we continue this process, islands will eventually become connected. This occurs as the threshold passes local minima. The 'prominence' of a peak is its height above the water level at the point when it merges with an island which originated at a higher peak. (The global maximum could be considered to have infinite prominence, but in this work it is convenient to use the value of the global maximum as the prominence). By considering prominence, we pay less attention to the absolute height of a peak and more attention to its height relative to adjacent peaks. This visual definition is complete, but a definition using mathematical notation is also available in the Supplementary Material.

2.3.2. Algorithm

The wave detection algorithm consists of four parts (A-D, see Fig. 2]). It processes a time series in order to identify significant, substantial and sustained periods of increased values. These periods are described as waves and the output of the algorithm is a labelling of the peaks of each wave and of the troughs between them.

After initial preprocessing (e.g. smoothing, handling of missing values), a list of all local minima and maxima (critical points) present in the time series is generated. This list corresponds to a set of wave candidates, which are the periods between consecutive local minima. The algorithm then processes this list of minima and maxima, reducing it by removing neighboring pairs of critical points from the list. This is equivalent to merging two neighboring wave candidates into one. The conditions on a pair of consecutive critical points to be removed are designed to identify those wave candidates that are insufficiently sustained or substantial and therefore suitable for merging. After the list is reduced, the surviving wave candidates will be of sufficient duration and prominence to be described as waves.

Four sub-algorithms operate to remove pairs from the list of critical points. The operation of each sub-algorithm is described briefly here and illustrated in Fig. 2, with more details available in the Supplementary Material and an implementation available at https://github.com/covid19db/epidemiological-waves/. Sub-algorithms A and B ensure that only 'sustained' periods of increase are identified as waves. Sub-algorithms C and D ensure that waves are 'significant and substantial' by merging waves of insufficient prominence.

Sub-algorithm A: Merging short wave candidates. For a feature to be described as a wave it must last for a certain minimal period. Wave candidates with duration less than T_{sep} are merged into neighbours, in order of increasing prominence. In Fig. 2 a pair of local minima is separated by less than T_{sep}. The pair of critical points to be removed is the intervening maximum and the greater of the two neighboring minima. The result is that the second wave is merged into the third. Where local maxima are separated by less than T_{sep}, the gap between waves is short, and in this case we also merge the wave candidates, considering them to be insufficiently separated from each other.

Sub-algorithm B: Merging irregular features. Where a local minimum and local maximum are separated by a short amount of time, this might indicate either an artifact in the data (for example, it can occur that a large number of historical cases are announced on a single day, creating a spike of cases) or a genuine shift. When two consecutive critical points are separated by less than $T_{cen}/2$, we centre a time window (shaded) of duration T_{cen} on the pair. The data from this period is censored; it is essential that $T_{cen} < T_{sep}$ so that

the window contains only a minimum and a maximum. If the feature is transient (as in the first censored window in Fig. 2, where the feature is a decrease in values but $y_1 < y_2$), the minimum and maximum are removed. The first wave is merged into the second. However, in the second window the decrease is still visible after censoring ($y_1 > y_2$) so this pair is retained. The third wave survives this filter.

Sub-algorithm C: Merging shallow wave candidates for low values. Waves with prominence less than P_{abs} are merged into neighbours. The prominence of the maximum in red (which is $y_{max} - y_{min}$), does not exceed P_{abs}. The maximum is therefore removed along with the greater of the neighboring minima, so that the first wave is merged into the second. When values are higher, this sub-algorithm is redundant, as sub-algorithm D would also merge all low-prominence waves.

Sub-algorithm D: Merging shallow wave candidates for high values. Case numbers may not fall sufficiently after a peak for a wave to be clearly separated from a neighbour. Waves with prominence less than $h \times P_{rel}$, where h is the value at the peak, are merged into a neighbour. This assumes a background level of 0 for the time series, which may not be suitable for other applications. In the illustration we have set $P_{rel} = 0.61$, which is the value used in this work. The prominence of the maximum in red (which is $y_{max} - y_{min}$), does not exceed $0.61 \times y_{max}$. The maximum is therefore removed along with the greater of the neighboring minima, so that the first wave is merged into the second.

2.3.3. Parameters

Certain parameters must be set by the user in order to determine how substantial or sustained a feature must be before it is identified as a wave. These are T_{sep}, T_{cen}, P_{abs} and P_{rel}. The choice of T_{sep} can be informed by biological information on the transmission of the disease. For example, epidemics that occur over a longer period of time such as tuberculosis and HIV will require a longer value for T_{sep} than more dynamic epidemics such as measles. The motivation for sub-algorithms A and B is to remove local fluctuations which are due to changes in things beyond the transmission process and hence this parameter must be informed by the natural timing of that transmission process. P_{abs} and P_{rel} should be set large enough to exceed the level of noise in the data and also meet the user's subjective opinion of what level of cases should be considered substantial and how well separated two waves must be.

The parameters used in this paper are set out in Table 1. At a minimum, the values T_{sep} and T_{cen} must be large compared to the serial interval of COVID-19, which is less than five days [19], and the mean incubation period, which is also approximately five days [17]. However, we find setting $T_{sep} = 35$ d and $T_{cen} = 35$ d to be a reasonable parameter choice as it matches well to the duration of what is generally understood to be a wave. P_{abs} and P_{rel} were chosen as the output provided a suitable match to expert opinion. P_{abs} is set to vary with the population of the country, but with a floor to account for noise and a ceiling for sufficiently large countries.

The final output is a list of local maxima and minima dividing the total period of the time series into waves. All remaining waves have prominence exceeding P_{abs}, length at least T_{sep}, and there is a long ($>T_{sep}$) and substantial ($P_{rel} = 61\%$) drop between any pair of peaks. The algorithm will indicate that no wave was detected if the time series never exceeds P_{abs}.

3. Results

3.1. Data collection

Epidemiological data from Ref. [1] was available for 192 countries. Of these, 148 also had data from Ref. [11] and 103 had data from Ref. [24] covers 103 countries.

Once attention was restricted to countries with a population >2.5 million, only 132 countries remained, of which 121 had data from Ref. [11] and 93 had data from Ref. [24].

3.2. Identification of epidemic waves of COVID-19

The algorithm was applied in the context of COVID-19 for every country for which data was available. By applying the algorithm to both the cases time series and the deaths time series, we could use cross-validation of the results to address the confounding effect of changing case ascertainment and better identify the waves of cases.

Fig. 3: (A–D) gives a sample of how the different layers of the algorithm merge wave candidates together.

Fig. 4 shows four characteristic examples of time series (Panels A–D) studied along with the division of those series into waves. It is

Table 1
Parameters used in the Wave Identification Algorithm. For cases $P_{abs} = \min (\max (45, \text{Population} \times 3.3/1{,}000{,}000), 500)$ and for deaths, $P_{abs} = \min (\max (7, \text{Population} \times 1/1{,}000{,}000), 70)$.

Parameter	Definition	Value (Cases)	Value (Deaths)
T_{sep}	Minimal wave-length, i.e. separation between consecutive minima or maxima.	35 days	35 day
T_{cen}	Each minimum and maximum separated by less than $T_{cen}/2$ are tested in Sub-algorithm B to determine if they persist over a period of duration T_{cen}. It is required that $T_{cen} < T_{sep}$.	35 days	35 days
P_{abs}	Absolute prominence threshold, defined as the minimal prominence for local maxima of waves.	See caption.	See caption.
P_{rel}	Relative prominence threshold, defined as the minimal prominence for local maxima of waves as a proportion of the value at the local maximum.	0.61	0.65

clear how decisions in the setting of parameters control which features are removed as noise and which structures are combined to be considered as one wave. The utility of cross-validating between the cases and deaths time series is also demonstrated.

There are a small number of outliers (Botswana, Cameroon, Equatorial Guinea, Republic of the Congo, and Western Sahara) for which a much larger number of waves are detected by the algorithm. On inspection, the time series for these countries are extremely noisy. This might be resolved by extending the smoothing window, but this would come at the cost of reduced accuracy for other countries. For this reason, we excluded these countries from the analysis. Countries with time series that never reached P_{abs} are also excluded from the analysis.

Table 2 reports certain epidemiological and socioeconomic characteristics according to the number of waves a country has experienced. Of the trends which can be observed there, only two are significant at the 5% level. A greater number of waves is associated with a longer stringency response time (a one-tailed Mann-Whitney test indicates that countries with more than one wave had significantly slower response time than those with only one wave, $p = 0.0002$) and with a higher GNI ($p < 0.0001$). The link to population density is not statistically significant.

3.3. Case studies

The two main drivers of waves in the case incidence times series are transmission and testing. A wave may be driven by an increase in transmission, an increase in testing, or, indeed, by a combination of these two drivers if the testing regime changes during a transmission wave. As a result, case incidence data from two subsequent waves can often not be usefully compared.

To illustrate some circumstances which may generate a time-series wave, we consider the observations from Italy and the United States as case studies. Fig. 5 shows the structure of the waves of the epidemic for both countries.

Case Study: Italy. In Italy there are two waves of confirmed cases and two waves of mortality, which occur at closely matched periods in time. However, as is clear from Fig. 5, Panel A, the relative number of cases to deaths about each peak differs greatly from the first to the second wave, translating into a decreasing CFR (case fatality ratio) trend which needs to be critically engaged with.

If we assume that the time series of deaths provides a more reliable and consistent indicator of trends in viral activity than the time series of cases then the relative difference in drivers can be inferred, at least after the fact, from the presence or absence of an accompanying mortality incidence peak.

In the case of Italy, then, we may conclude that the first wave, from Spring 2020, is an example of a 'transmission-driven wave'. These waves cannot be directly observed but may be inferred through the consequent demand for health care or active surveillance. The shape of these waves is a direct reflection of how the intensity of transmission changes over time and is best approximated by test positivity ratio curves (provided enough tests per day) and mortality data (in countries with close to no under-reporting of deaths).

On the other hand, from Autumn 2020 onwards, the wave of confirmed cases is observed principally as a result of increased case ascertainment and so is a 'testing-enhanced wave'. Unlike transmission waves, those do not cause increased hospitalisation and mortality, though the mere fact of observing an increasing number of cases may cause burdens. A true wave of infections may well underlie the testing-enhanced wave but its scale cannot be determined without strong assumptions regarding unknown transmission dynamics.

Case Study: United States. In Fig. 5, Panel B, we visually perceive three waves of cases and deaths in the United States. The
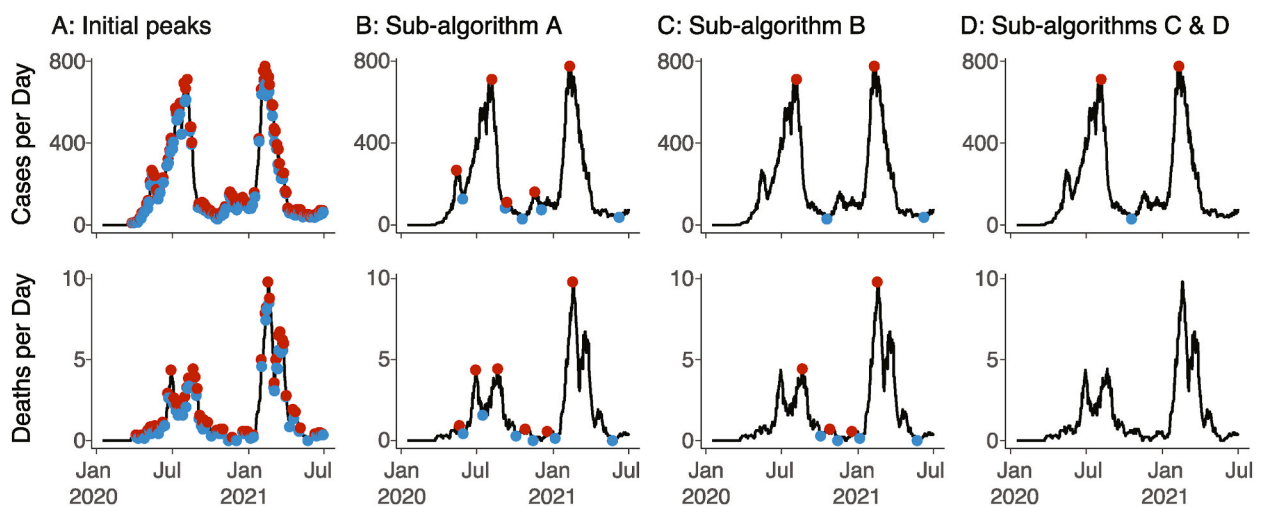

Fig. 3. Illustration of how each Sub-Algorithm operates to reduce the list of identified minima (solid blue circles) and maxima (solid red circles) in the time series of cases and deaths for Ghana. In black, the smoothed data for each time series is shown. Solid circles indicate which minima and maxima have survived the pruning process at each stage. **A:** The initial set of minima and maxima. **B:** After Sub-Algorithm A. **C:** After Sub-Algorithm B. **D:** After Sub-Algorithms C and D. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)
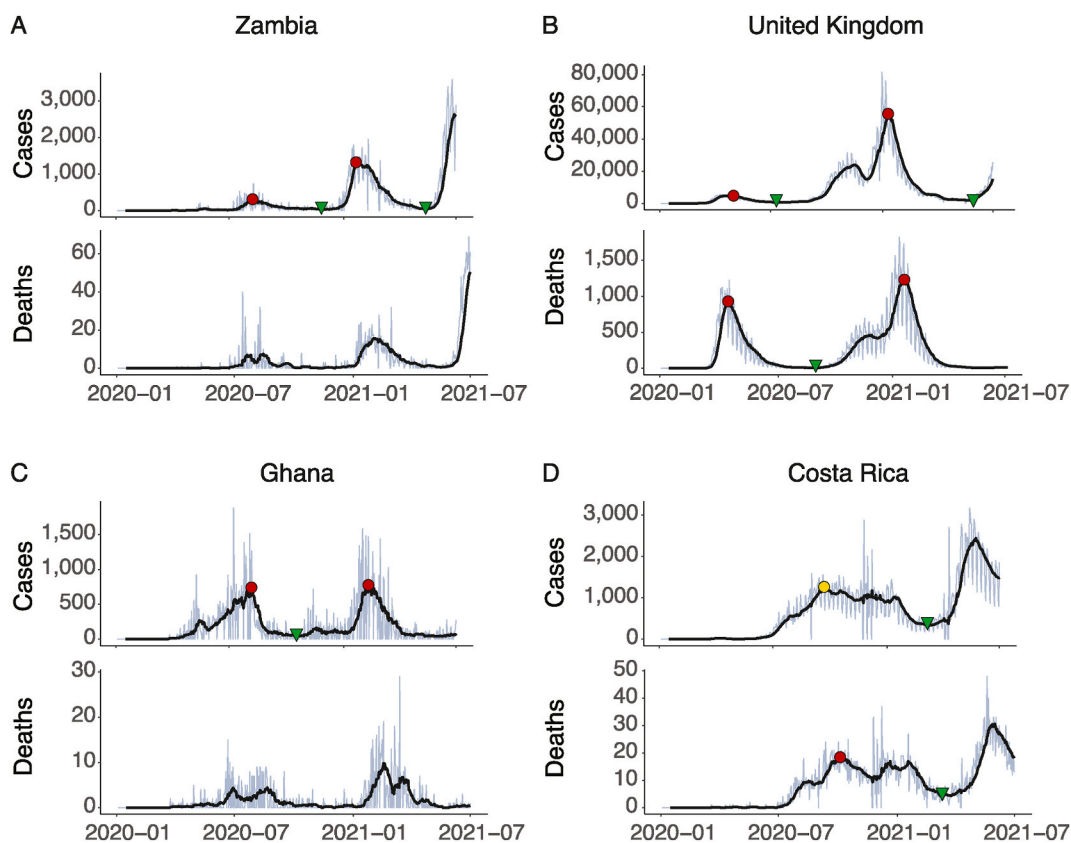
Fig. 4. Identification of epidemic waves of COVID-19. **A:** Zambia shows a clear structure with two waves (peaks marked by red circles, troughs by green triangles) in the cases data, while no waves are identified in the deaths data. **B:** The UK shows a structure that might arguably have two or three waves, but sub-algorithm D combines the final two. **C:** In Ghana sub-algorithm B filters out an early spike in cases. It is not clear visually whether this is noise or a meaningful epidemiological event; the algorithm cannot do better than the reader in determining this from simply inspecting a graph. No waves in deaths are identified due to low absolute counts. **D:** The number of cases in Costa Rica does not fall by 70% after the first wave, so it is not identified by the algorithm as a wave. This shows how important the parameter P_{rel} can be. However, cross-validating against the time series of deaths allows the wave to be identified (yellow circle). (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

Table 2
Basic characteristics of countries stratified by how many waves of COVID-19 they have experienced. The 'Count' variable indicates the number of countries in the cohort; not all characteristics are available for all countries in the cohort. The time T_0 is the date of the tenth cumulative death.

	Countries with 1 wave		Countries with 2 waves		Countries with 3 or more waves	
Count	57		84		51	
	Median	IQR	Median	IQR	Median	IQR
Mortality (cumulative per 10,000)	2.81	8.72	3.65	9.52	5.46	11.40
Cumulative case rate (total confirmed per 10,000)	196.99	609.05	233.18	698.56	353.62	675.67
Peak case rate (peak new cases per day per 10,000)	2.52	6.43	2.02	5.07	2.87	6.51
Stringency response time (days from T_0 to cancellation of public events)	−53.5	255	−33.5	51.25	−20	27.75
Total stringency (integral under stringency index curve)	26,601	6867	30,062	8186	31,413	5,6498
Testing response time (days from T_0 to 1 test per 1000 cumulative)	5	64.5	34	46.75	14	34.50
Population density (people per square km)	84.63	263.55	82.35	98.63	105.32	188.61
GNI per capita (PPP adj 2017 US$)	12,191	17.806	13,459	24,079	20,083	34,832

algorithm combines the first two wave structures into one wave due to the lack of a sufficiently deep fall in numbers between the two peaks. Again, there is a clear difference in the relative number of cases to deaths.

In this particular instance, we observe a striking geographical variation between the waves, with the outbreak being concentrated in different regions of the country at different times (Fig. 6, Panels A–B).

The aggregation of the regions into a single time series results in the appearance of multiple waves at the aggregate level. In general, this might be the result of a lag caused by geographical spread or of multiple index cases. We refer to these as 'spatially asynchronous
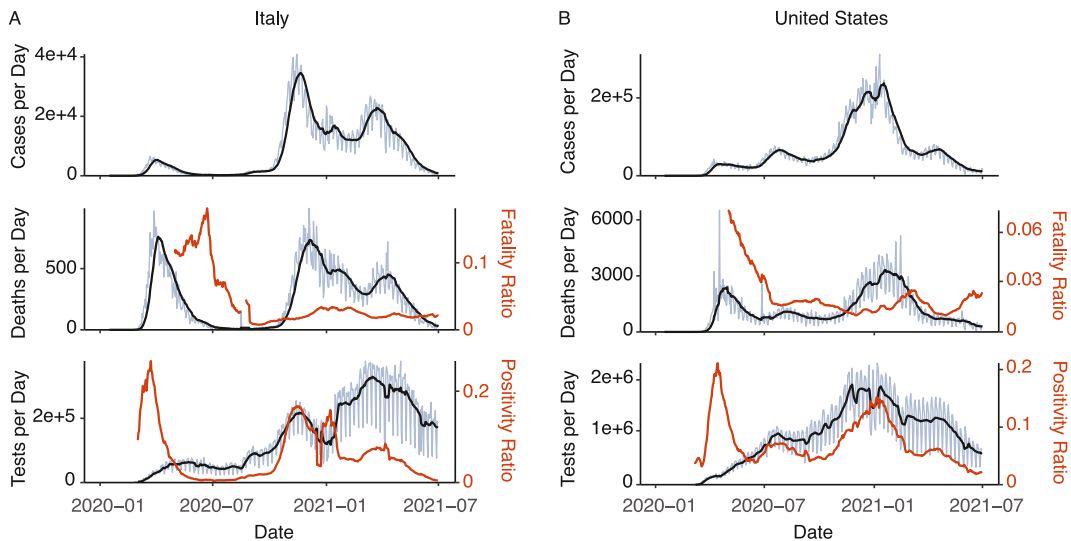
Fig. 5. The structure of the waves of the epidemic for two example countries: **A:** Italy and **B:** the United States. In Italy there are two waves of confirmed cases and two waves of mortality, which occur at closely matched periods in time. However, the relative number of cases to deaths about each peak differs greatly from the first to the second wave, translating into a decreasing CFR (case fatality ratio) trend which needs to be critically engaged with. In the United States, we visually perceive three waves of cases and deaths, with the algorithm combining the first two wave structures into one wave. Again, there is a clear difference in the relative number of cases to deaths. In this particular instance, we observe geographical variation between the waves, with the outbreak being concentrated in different regions at different times. This is an example of spatially asynchronous waves. Daily counts are shown in light blue, with a 7-day average in black. Identified wave durations for daily cases are shown in green. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)
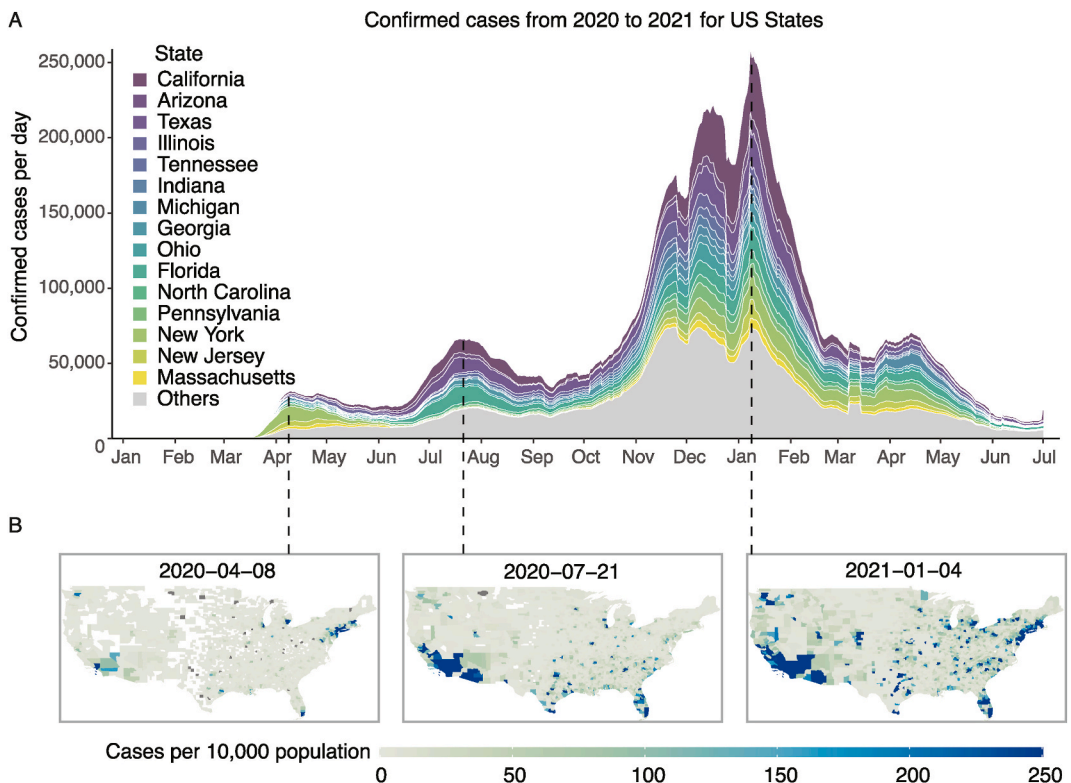

Fig. 6. A visualization of how the geographical concentration of COVID-19 has varied over time across the United States. **A:** Absolute numbers of cases by individual state. **B:** Relative numbers by individual county. Individual states show wave structures but the result of aggregating data at the national level is that local waves overlap and are not so distinct from each other.

waves'. Countries with this wave typology might benefit from evaluating the local epidemic curves in isolation and designing local intervention policies.

The concept of spatially asynchronous waves is more likely to be relevant in large countries. However, even in a small country, the progression of case numbers can usefully be thought of as the result of several superimposed outbreaks separated in both time and space, creating wave patterns in case numbers which may not fit our definition of a wave. Kaxiras and Neofotistos [15] found that a classification of countries according to the number of subepidemics and whether their intensity increases or decreases was well explained by government action, with declining subepidemics associated with stringent action and increasing subepidemics associated with less stringent action.

3.4. Modulation of waves by NPIs

Table 3 shows the relationship between the features of a government's response and the per capita death rate in that country. Table 4 shows the relationship between testing levels and the per capita death rate in that country. Both tables show positive correlations between the mortality burden and the overall scale of interventions or testing, which are strengthened when the relationships are studied by waves, an analysis made possible by the wave identification algorithm.

There is also a positive correlation between the time taken to implement the interventions or testing and the mortality burden. In the case of NPIs, this correlation is stronger for the overall burden than it is for the individual waves, which can be explained by the positive correlation between the response time and the number of waves experienced. No significant relationship between testing response time and the burden of the second wave was found.

4. Discussion

4.1. Identification of epidemic waves of COVID-19

Public discussion of waves of COVID-19 is common, with epidemiologists acknowledging the ambiguities around the term. Marc Lipsitch of Harvard has described a wave as "a useful metaphor" rather than a precise term, distinguishing waves from 'momentary' spikes [29] and Stephen Morse of Columbia has highlighted how epidemic waves in the traditional sense are "biological" phenomena, ending due to reduction in the susceptible population, while COVID-19 waves so far have ended due to measures "artificially designed to slow it down" [28].

This is not the first paper to seriously investigate the concept of observed epidemic waves. Taubenberger and Morens [25] note the three 'explosive pandemic waves' of 1918 influenza as a particularly unusual aspect of the pandemic. In the current context of COVID-19, seeking a rigorous means of identifying waves of COVID-19, Zhang et al. [31] propose an identification method which locates sustained periods of time when R_t is significantly larger than or less than 1. Values for R_t are obtained from the publicly available website [2] which uses a Kalman filter to estimate growth rate. Our method attends less to rates of increase and decrease as the defining aspect of a wave, focussing instead on identifying periods of high case numbers.

The shifts in sampling effort and sampling framework over the course of the COVID-19 pandemic greatly complicate the task of estimating the true level of infection. This paper contributes a semi-automated means of decomposing epidemiological time series into shorter periods, the waves. Without an algorithm of this type it is impossible to carry out the analysis above of how characteristics differ depending on the number of waves experienced. We also expect that analysis by wave period can reduce the impact of these shifts and provide customised time windows for each country studied. The method tallies well with natural intuition and allows for cross-validation between time series which the user expects to be correlated (cases and deaths). The parameters can be set by the user according to biological knowledge of the disease and the user's own determination of how a wave should be defined in a particular context.

This method is grounded in a translation of visual intuition into mathematical rules and so it cannot perform better than a visual

Table 3
Kendall's rank correlation between the mortality burden of the pandemic in different waves and aspects of government response. Stringency Response Time is defined to be the time from T_0 until public events were cancelled. The analysis is carried out for alternative definitions of Stringency Response Time in the Supplementary Material. Total Stringency and Wave Stringency are the integrals of the stringency index over the relevant period. Epidemic phase is an ordinal variable indicating which wave the country is currently in, and whether it has passed the peak of its current wave. The analyses for Wave 2 Deaths were carried out only for countries which experienced a second wave.

Burden metric	Stringency metric	Kendall's rank correlation	p-value
Wave 1 Deaths	Stringency Response Time	0.24	<0.0001
Wave 2 Deaths	Stringency Response Time	0.23	0.0019
All Deaths	Stringency Response Time	0.34	<0.0001
Epidemic Phase	Stringency Response Time	0.24	0.0003
Wave 1 Deaths	Total Stringency	0.25	<0.0001
Wave 2 Deaths	Total Stringency	0.13	0.082
All Deaths	Total Stringency	0.2	0.0005
Wave 1 Deaths	Wave 1 Stringency	0.32	<0.0001
Wave 2 Deaths	Wave 2 Stringency	0.38	<0.0001

Table 4

Kendall's rank correlation between the mortality burden of the pandemic in different waves and aspects of testing response. Testing Response Time is the time from T_0 until reaching a threshold for the level of testing (1 test per 1000 population). The analysis is carried out for alternative definitions of Testing Response Time in the Supplementary Material. Epidemic phase is an ordinal variable indicating which wave the country is currently in, and whether it has passed the peak of its current wave. The analyses for Wave 2 Deaths were carried out only for countries which experienced a second wave.

Burden metric	Testing metric	Kendall's rank correlation	p-value
Wave 1 Deaths	Testing Response Time	0.18	0.013
Wave 2 Deaths	Testing Response Time	−0.022	0.79
All Deaths	Testing Response Time	0.14	0.047
Epidemic Phase	Testing Response Time	0.021	0.78
Wave 1 Deaths	Total Tests	0.26	0.0005
Wave 2 Deaths	Total Tests	0.41	<0.0001
All Deaths	Total Tests	0.42	<0.0001
Wave 1 Deaths	Wave 1 Tests	0.42	<0.0001
Wave 2 Deaths	Wave 2 Tests	0.37	0.0033

inspection of the time series. However, it does provide the advantage of setting an objective framework for identification of waves and the ability to process a large number of time series swiftly and consistently. Another aspect of this method is that every date in the time series is allocated to a wave, meaning every case and every death is allocated to a particular wave. This may not always be desirable. If there is a clear plateau between two waves then the location of the minimum which marks the end of one wave and the beginning of its successor will not be robust. A threshold test could easily be implemented to identify gaps between waves.

4.2. Modulation of waves by NPIs

We must urge caution in the interpretation of the positive correlations observed in Tables 3 and 4

The feedback process between the use of NPIs and mortality is complex and we have not carried out a causal analysis of this relationship. Most likely, our results suggest that countries that have introduced more stringent and/or prolonged measures did so in response to higher mortality figures, not that higher stringency measures caused more deaths.

There is also a positive correlation between the time taken to implement the interventions or testing and the mortality burden. In the case of NPIs, this correlation is stronger for the overall burden than it is for the individual waves, which can be explained by the positive correlation between the response time and the number of waves experienced.

The finding of a positive correlation between time to implement NPIs and mortality in the second wave of an epidemic is somewhat surprising, since the second wave is distant in time. In fact, the correlation is similar to that with the burden of the first wave. This may be because response time influenced conditions at the beginning of the second wave, or it may indicate a country's general policy approach.

Although some early NPIs such as implementation of testing strategies, isolation of confirmed cases and tracing of their contacts slowed down early transmission, in most places these measures were not sufficient to contain the outbreak [8,13]. Studies have found that while moderate measures are capable of reducing the size of the ensuing epidemic, severe measures are needed to suppress an epidemic so as to avoid overwhelming healthcare systems [7].

The positive correlation between the strength of interventions and the burden of the pandemic has been observed before. This relationship has been the subject of lively debate centered on the difficulty in collecting reliable data, the importance of timing, and the fact that interventions are an endogenous variable, with governments implementing interventions in response to current or anticipated case load [3]. Indeed, attending more closely to the timing of interventions and effects on the transmission rate rather than on absolute case numbers demonstrates that a wide range of restrictions are associated with a reduction in R_t [4,12,14], with early implementation being crucial.

Our findings provide further confirmation of the importance of early introduction of NPIs. Countries which introduced early NPIs tended to experience a lower mortality burden in the first wave, were less likely to experience a second wave and, if a second wave was experienced, that wave came with a lower mortality burden. These three effects combine to produce a very strong relationship between early introduction and low overall mortality.

5. Conclusion

It is possible to distill the natural visual understanding of 'waves' in a time series into simple mathematical rules which can be used to objectively annotate a large number of time series, identifying their component waves. In the context of COVID-19, these waves may result from increased transmission, increased testing, or some combination of the two. Not only this, waves can occur due to the aggregation of time series from a large geographical area, so that the second wave is really a first wave, but for a different region of the country. When engaging in comparative study of the relationships between interventions taken and the mortality burden of the disease, using the wave as the time unit for analysis can result in clearer conclusions. The speed with which interventions are implemented is strongly correlated with the wave structure of the subsequent epidemic.

Author contribution statement

Funding statement

Data availability statement

No new data were created during this study. All analysis and code, including the code used to access the pre-existing data from other repositories, are available at https://github.com/covid19db/epidemiological-waves/.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.heliyon.2023.e16015.

References

[1] European Centre for Disease Prevention and Control, Covid-19 Pandemic, 2020. https://www.ecdc.europa.eu/en/covid-19-pandemic.

[2] F. Arroyo-Marioli, F. Bullano, S. Kucinskas, C. Rondón-Moreno, Tracking R of COVID-19: a new real-time estimation using the Kalman filter, PLoS One 16 (1) (2021), e0244474.

[3] L. Besançon, G. Meyerowitz-Katz, A. Flahault, Sample size, timing, and other confounding factors: toward a fair assessment of stay-at-home orders, Eur. J. Clin. Invest. 51 (6) (2021).

[4] J.M. Brauner, S. Mindermann, M. Sharma, D. Johnston, J. Salvatier, T. Gavenčiak, A.B. Stephenson, G. Leech, G. Altman, V. Mikulik, A.J. Norman, J.T. Monrad, T. Besiroglu, H. Ge, Hartwick Ma, Y.W. Teh, L. Chindelevitch, Y. Gal, J. Kulveit, Inferring the effectiveness of government interventions against COVID-19, Science 371 (6531) (2021).

[5] Y.M. Chung, Chuan S. Hu, Y.L. Lo, H.T. Wu, A persistent homology approach to heart rate variability analysis with an application to sleep-wake classification, Front. Physiol. 12 (2021).

[6] S. Conti, P. Ferrara, G. Mazzaglia, M.I. D'Orso, R. Ciampichini, C. Fornari, F. Madotto, M. Magoni, G. Sampietro, A. Silenzi, C.V. Sileo, A. Zucchi, G. Cesana, L. Manzoli, L.G. Mantovani, Magnitude and time-course of excess mortality during COVID-19 outbreak: population-based empirical evidence from highly impacted provinces in northern Italy, ERJ Open Res. 6 (3) (2020) 458–2020.

[7] N.G. Davies, A.J. Kucharski, R.M. Eggo, A. Gimma, W.J. Edmunds, T. Jombart, K. O'Reilly, A. Endo, J. Hellewell, E.S. Nightingale, Billy J. Quilty, C.I. Jarvis, T. W. Russell, P. Klepac, Bosse Ni, S. Funk, S. Abbott, G.F. Medley, H. Gibbs, C.A.B. Pearson, S. Flasche, M. Jit, S. Clifford, K. Prem, C. Diamond, J. Emery, A. K. Deol, S.R. Procter, K. van Zandvoort, Y.F. Sun, J.D. Munday, A. Rosello, M. Auzenbergs, G. Knight, R.M.G.J. Houben, Y. Liu, Effects of non-pharmaceutical interventions on COVID-19 cases, deaths, and demand for hospital services in the UK: a modelling study, Lancet Public Health 5 (7) (2020) e375–e385.

[8] LF Buss WM de Souza, D.S. Candido, J.-P. Carrera, S. Li, A.E. Zarebski, R.H.M. Pereira, C.A. Prete, A.A. de Souza-Santos, Kris V. Parag, M.C.T.D. Belotti, M. F. Vincenti-Gonzalez, Janey Messina, F.C. da Silva Sales, P.S. Andrade, V.H. Nascimento, Fabio Ghilardi, Leandro Abade, Bernardo Gutierrez, U. Moritz, G. Kraemer, K. Carlos, V. Braga, R.S. Aguiar, N. Alexander, P. Mayaud, O.J. Brady, I. Marcilio, Nelson Gouveia, Guangdi Li, A. Tami, S.B. de Oliveira, V.B. G. Porto, Fabiana Ganem, W.A.F. de Almeida, F.F.S.T. Fantinato, E.M. Macário, W.K. de Oliveira, M.L. Nogueira, O.G. Pybus, C.-H. Wu, Julio Croda, E.C. Sabino, N.R. Faria, Epidemiological and clinical characteristics of the COVID-19 epidemic in Brazil, 2020 4:8, Nat. Hum. Behav. 4 (8) (2020) 856–865.

[9] M. Dindin, Y. Umeda, F. Chazal, Topological data analysis for arrhythmia detection through modular neural networks, Adv. Artif. Intell.: 33rd Can. Conf. Artif. Intell. (2020) 177–188.

[10] J. Dushoff, J.B. Plotkin, S.A. Levin, D.J.D. Earn, Dynamical resonance can account for seasonality of influenza epidemics, Proc. Natl. Acad. Sci. U.S.A. 101 (48) (2004) 16915–16916.

[11] T. Hale, N. Angrist, R. Goldszmidt, B. Kira, A. Petherick, T. Phillips, S. Webster, E. Cameron-Blake, L. Hallas, S. Majumdar, H. Tatlow, A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker), Nat. Hum. Behav. 5 (4) (2021) 529–538.

[12] N. Haug, L. Geyrhofer, A. Londei, E. Dervic, A. Desvars-Larrive, V. Loreto, B. Pinior, S. Thurner, P. Klimek, Ranking the effectiveness of worldwide COVID-19 government interventions, 2020 4:12, Nat. Hum. Behav. 4 (12) (2020) 1303–1312.

[13] J. Hellewell, S. Abbott, Phd Jarvis, T.W. Russell, J.D. Munday, A.J. Kucharski, W.J. Edmunds, S. Funk, R.M. Eggo, J. Hellewell, S. Abbott, A. Gimma, Bosse Ni, C. I. Jarvis, T.W. Russell, J.D. Munday, A.J. Kucharski, J. Edmunds, S. Funk, Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts, Lancet Global Health 8 (2020) e488–e496.

[14] S. Hsiang, D. Allen, S. Annan-Phan, K. Bell, I. Bolliger, T. Chong, H Druckenmiller, L.Y. Huang, A. Hultgren, E. Krasovich, P. Lau, J. Lee, E. Rolf, J. Tseng, T. Wu, The effect of large-scale anti-contagion policies on the COVID-19 pandemic, 2020 584:7820, Nature 584 (7820) (2020) 262–267.

[15] E. Kaxiras, G. Neofotistos, Multiple epidemic wave model of the COVID-19 pandemic: modeling study, J. Med. Internet Res. 22 (7) (2020), e20912.

[16] Katia Koelle, Sarah Cobey, Bryan Grenfell, Mercedes Pascual, Epochal evolution shapes the phylodynamics of interpandemic influenza a (h3n2) in humans, Science 314 (5807) (2006) 1898–1903.

[17] Q. Li, X. Guan, P. Wu, X. Wang, L. Zhou, Y. Tong, R. Ren, K.S.M. Leung, E.H.Y. Lau, J.Y. Wong, X. Xing, N. Xiang, Y. Wu, C. Li, Q. Chen, D. Li, T. Liu, J. Zhao, M. Liu, W. Tu, C. Chen, L. Jin, R. Yang, Q. Wang, S. Zhou, R. Wang, H. Liu, Y. Luo, Y. Liu, G. Shao, H. Li, Z. Tao, Y. Yang, Z. Deng, B. Liu, Z. Ma, Y. Zhang, G. Shi, T.T.Y. Lam, J.T. Wu, G.F. Gao, B.J. Cowling, B. Yang, G.M. Leung, Z. Feng, Early transmission dynamics in Wuhan, China, of novel coronavirus–infected pneumonia, N. Engl. J. Med. 382 (13) (2020) 1199–1207.

[18] A. Mahdi, P. Błaszczyk, P. Dłotko, D. Salvi, T.-S. Chan, J. Harvey, D. Gurnari, Y. Wu, A. Farhat, N. Hellmer, A. Zarebski, B. Hogan, L. Tarassenko, OxCOVID19 Database, a multimodal data repository for better understanding the global impact of COVID-19, 2021 11:1, Sci. Rep. 11 (1) (2021) 1–11.

[19] H. Nishiura, N.M. Linton, A.R. Akhmetzhanov, Serial interval of novel coronavirus (COVID-19) infections, Int. J. Infect. Dis. 93 (2020) 284–286.

[20] Matthew J. Page, Joanne E McKenzie, Patrick M. Bossuyt, Isabelle Boutron, Tammy C. Hoffmann, Cynthia D. Mulrow, Larissa Shamseer, Elie A Akl Jennifer M Tetzlaff, Sue E. Brennan, Roger Chou, Julie Glanville, Jeremy M Grimshaw, Asbjørn Hróbjartsson, Manoj M. Lalu, Tianjing Li, Elizabeth W. Loder, Evan Mayo-Wilson, Steve McDonald, Luke A. McGuinness, Lesley A. Stewart, Thomas James, Andrea C. Tricco, Vivian A. Welch, Penny Whiting, David Moher, The PRISMA 2020 statement: an updated guideline for reporting systematic reviews, BMJ 372 (2021).

[21] N. Phillips, The coronavirus is here to stay - here's what that means, Nature 590 (7846) (2021) 382–384.

[22] M. Pigliucci, Species as family resemblance concepts: the (dis-)solution of the species problem? Bioessays 25 (6) (2003) 596–602.

[23] B.J. Quilty, S. Clifford, S. Flasche, R.M. Eggo, Effectiveness of airport screening at detecting travellers infected with novel coronavirus (2019-nCoV), Euro Surveill. 25 (5) (2020).

[24] M. Roser, H. Ritchie, E. Ortiz-Ospina, J. Hasell, Coronavirus Pandemic (COVID-19), Our World in Data, 2020.

[25] J.K. Taubenberger, D.M. Morens, 1918 influenza: the mother of all pandemics, Emerg. Infect. Dis. 12 (1) (2006) 15.

[26] T. Vega, J.E. Lozano, T. Meerhoff, T. Snacken, J. Mott, R. Ortiz de Lejarazu, B. Nunes, Influenza surveillance in Europe: establishing epidemic thresholds by the moving epidemic method, Influenza Other Respir Viruses 7 (4) (2013) 546–558.

[27] A.L. Wagner, What Makes a 'wave' of Disease? An Epidemiologist Explains, The Conversation, 2020. https://theconversation.com/what-makes-a-wave-of-disease-an-epidemiologist-explains-141573.

[28] M. Walker, COVID-19: what makes 'waves' during a pandemic? MedPage Today (2020). https://www.medpagetoday.com/infectiousdisease/covid19/89599.

[29] T. Wibawa, M. Walsh, South Korea Says it Has a Second Wave of Coronavirus Infections — but what Does that Really Mean? Australian Broadcasting Corporation News, 2020. https://www.abc.net.au/news/2020-06-24/coronavirus-covid-19-countries-second-wave-explainer-south-korea/12385882.

[30] C. Yang, Z. He, W. Yu, Comparison of public peak detection algorithms for MALDI mass spectrometry data analysis, BMC Bioinf. 10 (2009).

[31] S.X. Zhang, F.A. Marioli, R. Gao, A second wave? What do people mean by COVID waves? – a working definition of epidemic waves, Risk Manag. Healthc. Pol. 14 (2021) 3775–3782.