



Menu

HL

Graphic detail | Covid-19 death tallies

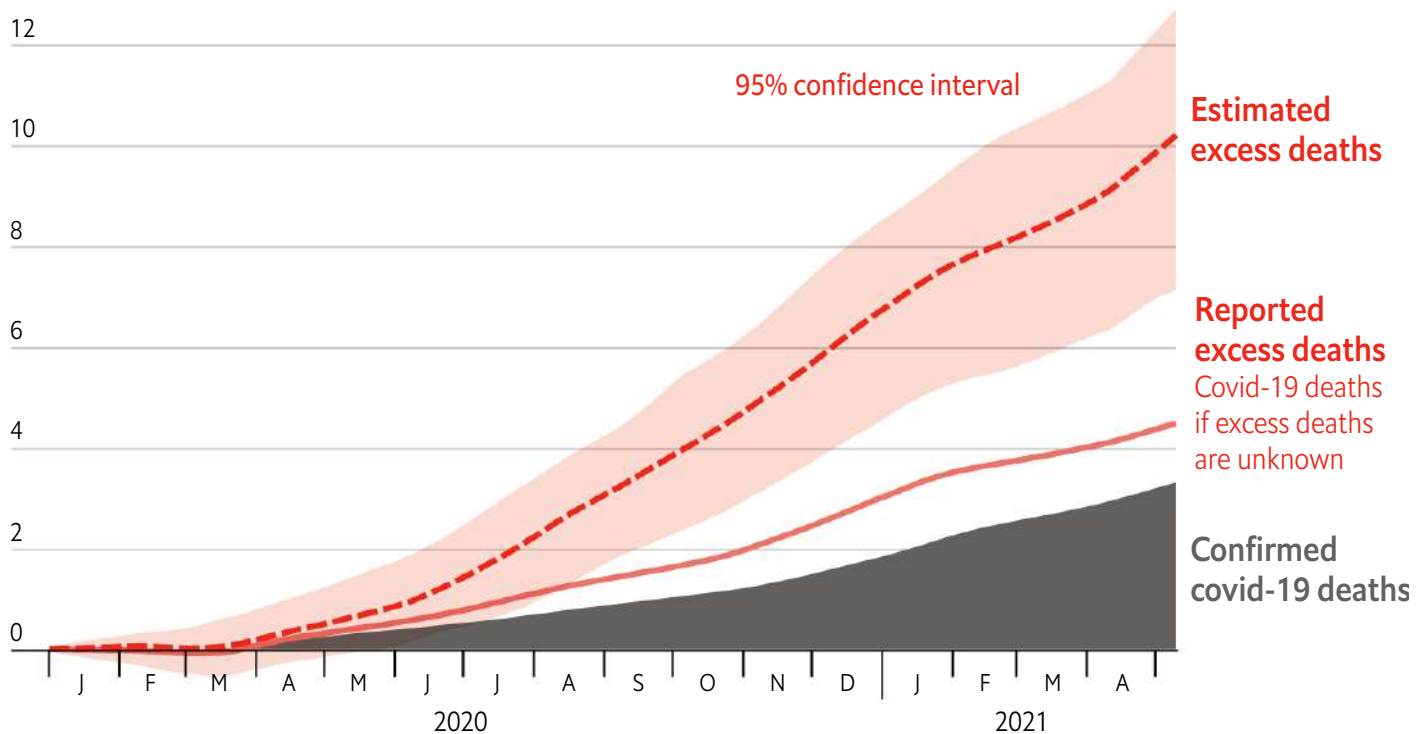
# How we estimated the true death toll of the pandemic

## Dealing with potential outcomes, known unknowns, and uncertainty

May 13th 2021

Save Share Give

Global estimated excess deaths v confirmed covid-19 deaths, m



Sources: Johns Hopkins University, CSSE; The Economist excess deaths model

*You can see our main findings, and country-by-country estimates, [here](#). Replication code and data is available [here](#).*

**A**CADEMIC RESEARCH often seeks to answer counter-factual questions. What would have happened if a drug had been administered earlier to someone suffering from a disease? What if wealth were distributed more equally than it is? What would have occurred if a revolution that failed had instead succeeded—or if a successful one had failed? Even in quantum physics, we cannot observe something both happening and not: as soon as we open the box containing Schrödinger’s cat, the animal is either alive or dead. Yet any time scholars make causal claims, they are implicitly estimating the results of a comparison between what actually occurred and what might have, but didn’t.

The covid-19 pandemic has changed the world in innumerable ways, but its most immediate and tragic consequence has been the loss of life. Just as with countless other research topics, calculating a full account of the disease’s human toll requires estimating the answer to a counter-factual question. How many more people have died than the number who would still have perished if SARS-CoV-2 had never emerged?

### **Why not just count them?**

The answer to this question is measured using “excess mortality”, which is the gap between the total number of deaths that occur for any reason and the amount that would be expected under normal circumstances. Excess mortality caused by the pandemic

normal circumstances. Excess mortality caused by the pandemic is probably smaller than covid-19's total effect on the world's population: the reduction in the number of babies born during the pandemic could well exceed the net increase in deaths.

However, it is much more comprehensive than official statistics on how many people have been killed by the disease.

Public-health authorities classify anyone who dies within weeks of a positive covid-19 test as a covid-19 death. In some cases, this method can lead to over-counting: some people, particularly ones who are extremely old or suffer from other severe ailments, would have died anyway between the time they were infected and the present. More often, however, this approach leads to underestimates. Lots of people who would still be alive today had they not been infected with covid-19 do not show up in official figures, because they succumbed to the disease without being tested for it.

In addition, the pandemic has had indirect effects on death rates. Many people with treatable medical conditions have died because they did not receive appropriate health care. Conversely, other people have survived who, were it not for the behavioural changes induced by covid-19, would have died from causes like influenza, air pollution or car accidents.

These effects can be quite large, and vary from country to country in both size and direction. According to the most recent data, excess deaths in America are 7% higher than its official covid-19 deaths. In contrast, excess mortality is 21% below the publicly

reported toll from covid-19 in Britain, and 20% lower in France. Places that have had extremely mild pandemics, such as Australia, Norway, and Taiwan, actually have deficit rather than

excess deaths, meaning that fewer people have died overall than the average in previous years.

In middle-income and especially poor countries, which conduct less testing for covid-19 than rich ones do, excess mortality almost always exceeds the official death toll, often by large margins. For example, in Romania, Iran, South Africa and Mexico, excess deaths are more than double those officially attributed to covid-19. Moreover, these multipliers vary over time. During the course of the pandemic, some countries have improved their ability to count deaths accurately. Others, particularly those facing severe outbreaks, got worse.

### **Known unknowns**

Data on excess deaths are available for dozens of countries, and many places update them frequently. In these mostly rich parts of the world, we have a reliable measure of the pandemic's true toll. However, other countries report total death numbers only with long lags, and a surprisingly large number do not publish them at all. Such places also tend to do relatively little testing for covid-19, meaning that their official death counts capture just a small share of the disease's true impact. A large majority of the world's population lives in regions, including India, China and most of Africa, where data on excess deaths are heavily delayed or entirely unavailable

UNAVAILABLE.

The only way to fill in these continent-sized gaps in the data is to estimate them. Fortunately, lots of other statistics about these countries are available, which can be used to make highly educated guesses. Virtually all countries publish official counts of covid-19 cases and deaths, which can serve as a starting point. Most countries also report the share of their covid-19 tests that are positive, which can be used to estimate the magnitude of under-counting of cases. In some places, seroprevalence surveys offer estimates of the share of the population that has detectable antibodies to SARS-CoV-2, the virus that causes the disease, which is a measure of past infection. Other indicators may matter as well, such as the steps governments have taken to curb the pandemic; the extent to which people move around physically; systems of government; media freedom; demography; and geographic location. In cases where countries do not publish some of these data, their absence is useful information in and of itself.

## **Collect and connect**

To estimate excess deaths wherever and whenever data are unavailable—which in some places is the entirety of the pandemic, and in others only particular time periods—*The Economist* gathered all of the relevant information we could find about each country (a full list of sources and categories is provided below). In total, we used 121 different indicators, and produced estimates for over 200 distinct countries and territories. Of these statistics, about half were constant over time

(an example is demography, where we used population estimates from 2019) and half varied by day (the number of diagnosed covid-19 cases, or the extent to which people moved around according to their mobile phones). These numbers were unavailable for many countries. To improve predictions in such places, we calculated weighted averages obtained from neighbouring, nearby or similar countries for each variable that was missing. And to avoid distortions caused by the specific day of the week on which data are reported, we aggregated daily numbers into weekly averages.

After amassing all of these predictors, we lined them up with data on excess deaths, which we have been collecting for more than a year. [Our tracker](#) has numbers for 78 countries, drawing heavily from the [Human Mortality Database](#) and [World Mortality Dataset](#). In addition, for this exercise, we made an extra effort to obtain at least partial data about countries with especially large populations, which weigh heavily in any estimate of the global total. Using the supplemental appendix of an [article](#) in the British Journal of Medicine, we were able to calculate 12 weeks of data on excess deaths in China from the Chinese Centres for Disease Control (which did not respond to our requests). For India, we could only identify reliable data for the city of Mumbai. Rather than assuming that this one region is representative of all of India, we had our models treat Mumbai as a separate country, from which it could learn patterns to apply elsewhere, in India and beyond. In order to strike a balance between information obtained from big countries and small ones, we weighted data

from each country by the logarithm of its population. This meant that our eventual model would assign much more importance to Mexico than to Monaco, without ignoring the latter entirely.

## **Estimation station**

Our next step was to feed all of these numbers into a machine-learning algorithm called a “gradient booster”, in order to determine which combinations of our variables yielded the most accurate estimates of excess deaths per 100,000 people. This method makes predictions based on a series of “decision trees”. For instance, one tree might first check if test positivity rates are above or below a given threshold, and then ask whether the number of tests per person exceeds a different cut-off. It might find that if both figures are greater than these benchmarks, excess deaths tend to be higher than the model expected before asking those questions. It thus adjusts its prediction incrementally upwards. The algorithm constructs thousands of these trees, each one detecting new patterns that improve the forecasts generated by the “forest” of all previous ones.


To assess whether our model’s estimates were accurate, we split up our data into ten chunks. One by one, we set aside a single chunk, trained our algorithm on the remaining nine-tenths of the data, and asked the model to make a prediction of excess deaths for the countries we did not allow it to see. If the model did a good job of estimating mortality in places for which we knew the real number but it did not, we reasoned that it would fare similarly well at making predictions for countries where no one



knows the true value.

### Daily excess deaths per 100,000 people, actual vs test model predictions\*

Jan 1st 2020-May 10th 2021



\*Models created without data on the countries for which they were asked to predict.  
The final model uses all data available, for even higher accuracy.

Our model was far from perfect. In some cases, its best guess of the number of excess deaths in a given country in a given week was far too high; others were far too low. Overall, however, the model was reasonably well-calibrated. Given 100 different examples of combinations of countries and weeks where its



predictions were between, say, 0.45 and 0.55 excess deaths per 100,000 people, some might actually have been as low as -0.1 and others might in fact have been as high as 1.5. However, the average of all 100 cases was generally close to 0.5. The same was true for other levels of predicted excess deaths, such as those between -0.2 and -0.1 per 100,000, or between 0.75 and 0.85.

Our algorithm struggled most at the very peak of the direst outbreaks. In the small handful of cases where it estimated horrifying weekly tolls above three excess deaths per 100,000 people, the real numbers clustered between two and three. (You can see this pattern at the top right of the chart above, where the red line, representing the best-fitting trend between predicted and actual data, dips below the black 45-degree line that represents a perfectly calibrated model.) However, such severe cases are so rare that they barely influence our total global estimate. For example, the model's current estimate for the out-of-control outbreak ravaging India is still below 1.5 excess deaths per 100,000 people per week, putting that country comfortably within the range in which its forecasts are generally on target.

### **Believe it or not**

Our algorithm produces only a single best guess, or “point estimate”, for each country in each week. However, such central estimates may be of little value if they sit near the middle of vast ranges of uncertainty. As one example commonly cited in statistics classes illustrates, picture a poor student given two buckets of water, one near boiling, the other near freezing. His

best guess for the temperature was correctly a comfortable 50° C

Best guess for the temperature was, correctly, a comfortable 50 °C. However, failing to appreciate the importance of the variation surrounding that estimate, he put one foot in each.

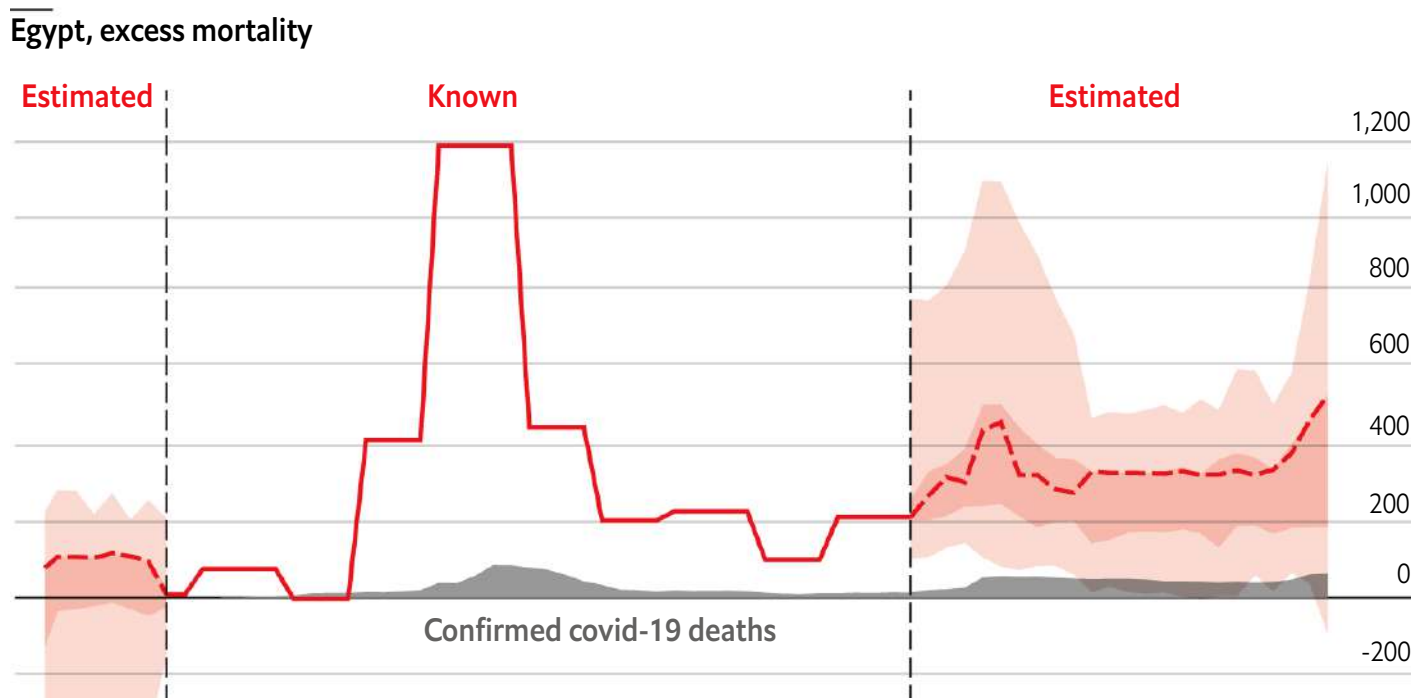
To avoid such pitfalls, our final step was to calculate the ranges within which the model suggested the true answer would lie, with varying degrees of certainty. We used a method called “bootstrapping”. Rather than merely using our full dataset to produce a single model, we created 100 artificial datasets of the same size, first by picking countries at random, and then similarly by selecting weeks of data at random for each chosen country. After a country or week was chosen, we did not remove it but instead left it in place. As a result, each individual artificial dataset contains many copies of certain country-week pairs, whereas others are absent entirely.


Next, we trained a different gradient-boosting model on each of these simulated datasets. Each model contains a different set of decision trees that reflect different patterns and relationships between variables, reflecting the idiosyncrasies of the specific artificial dataset used to create it. In turn, each model yields a slightly different prediction for every single country-week combination. For each excess-death total we publish, we use the model trained on the actual, complete dataset as our central estimate, and then the middle 95 or 50 of the predictions generated by the 100 other models as our confidence intervals. These bounds reflect the ranges within which the models suggest that the true value should lie, with 95% or 50% certainty.

## Combining it all

Our final step was to combine these estimates with published data on excess deaths. For any week in any country where a true number of excess deaths is reported, we simply use the known value; for those where one is missing, we fill in our estimate instead. For some countries, such as India, we had to use the model's prediction for every week; in others, such as America, we only needed to estimate the most recent week.

This process is illustrated in the chart below, showing data from Egypt. To an extent, our estimates move up and down over time in the same direction as officially recorded covid-19 deaths. However, the relationship between official and estimated deaths itself varies. During periods when recorded outbreaks were worse, our model reckoned that undercounting was more severe as well (as suggested by available data).





Sources: Johns Hopkins University CSSE; *The Economist* excess-deaths model

By adding up known and estimated totals for every country in every week, we can compute an estimated global total of excess deaths due to the pandemic, and uncertainty ranges that reflect the precision of those estimates. To read the model's conclusions, and what the results imply, please see our [Briefing](#). Interested readers can also explore the model themselves: we provide complete replication code with annotations and data at our dedicated [GitHub repository](#).

## Sources:

### Excess deaths

*The Economist*; Human Mortality Database; [World Mortality Dataset](#); Registro Civil (Bolivia); Vital Strategies; Office for National Statistics; Northern Ireland Statistics and Research Agency; National Records of Scotland; Registro Civil (Chile); Registro Civil (Ecuador); Institut National de la Statistique et des Études Économiques; Santé Publique France; Istituto Nazionale di Statistica; Dipartimento della Protezione Civile; Secretaría de Salud (Mexico); Ministerio de Salud (Peru); Data Science Research Peru; Departamento Administrativo Nacional de Estadística (Colombia); South African Medical Research Council; Instituto de Salud Carlos III; Ministerio de Sanidad (Spain); Datadista; Liu et al (2021)

### Covid-19 data (deaths cases testing and vaccinations)

**COVID-19 data (deaths, cases, testing, and vaccinations)**

Our World In Data; Johns Hopkins University, CSSE

## **Prevalence of covid-19 antibodies**

[SeroTracker.com](https://serotracker.com)

## **Demography and urbanization rates**

Our World in Data, World Bank, United Nations

## **Demography-adjusted infection fatality rate**

*The Economist*, based on Brazeau et al. (2020) and UN population figures

## **Healthcare quality**

Our World in Data, World Bank

## **Political regime and media freedom data**

V-Dem Institute, PolityIV Project, Freedom House, Boix et al (2015)

## **Economy and connectivity**

World Bank, Our World in Data, World Tourism Organization

## **Mobility**

COVID-19 Community Mobility Reports (Google)

## **Geography**

Decker et al (“maps” R package), Mayer T et al (2011)

## **Government policy responses to Covid-19**

OxCGRT (University of Oxford)




**SUBSCRIBER ONLY** | OFF THE CHARTS

# Taking you behind the scenes of our data journalism

Directly to your inbox every week

Sign up


## More from Graphic detail





## Satellite data show Ukraine's forces are testing Russia's defences

The attacks may be part of the highly anticipated counter-offensive



## How well does your country provide for its citizens?

A global index that attempts to compare people's well-being across 170 countries





Our model suggests that global deaths remain 5% above pre-covid forecasts

Attributing this increase to covid would make it the fourth-leading cause of death

**Subscribe**

**The Trust Project**

**Group subscriptions**

**Help and contact us**

**Reuse our content**

**Keep updated**



Published since September 1843 to take part in “*a severe contest between intelligence, which presses forward, and an unworthy, timid ignorance obstructing our progress.*”

---

## The Economist

About

Advertise

Press centre

## The Economist Group

The Economist Group

Economist Intelligence

Economist Impact

Economist Events

Working Here

Economist Education Courses

Which MBA?

Executive Jobs

Executive Education Navigator

---

[Terms of Use](#) [Privacy](#) [Cookie Policy](#) [Manage Cookies](#) [Accessibility](#)

[Modern Slavery Statement](#) [Sitemap](#) [California: Do Not Sell My Personal Information](#)

Copyright © The Economist Newspaper Limited 2023. All rights reserved.