RESEARCH ARTICLE SUMMARY

CORONAVIRUS

Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions

Jonas Dehning*, Johannes Zierenberg*, F. Paul Spitzner*, Michael Wibral, Joao Pinheiro Neto, Michael Wilczek*, Viola Priesemann*†

INTRODUCTION: When faced with the outbreak of a novel epidemic such as coronavirus disease 2019 (COVID-19), rapid response measures are required by individuals, as well as by society as a whole, to mitigate the spread of the virus. During this initial, time-critical period, neither the central epidemiological parameters nor the effectiveness of interventions such as cancellation of public events, school closings, or social distancing is known.

RATIONALE: As one of the key epidemiological parameters, we inferred the spreading rate λ from confirmed SARS-CoV-2 infections using the example of Germany. We apply Bayesian inference based on Markov chain Monte Carlo sampling to a class of compartmental models [susceptible-infected-recovered (SIR)]. Our analysis characterizes the temporal change of the spreading rate and allows us to identify potential change points. Furthermore, it enables short-term forecast scenarios that assume various degrees of social distancing. A detailed description is provided in the accompanying paper, and the models, inference, and forecasts are available on GitHub (https://github.com/Priesemann-Group/covid19_inference_forecast). Although we apply the model to Germany, our approach can be readily adapted to other countries or regions.

RESULTS: In Germany, interventions to contain the COVID-19 outbreak were implemented in three steps over 3 weeks: (i) Around 9 March 2020, large public events such as soccer matches were canceled; (ii) around 16 March 2020, schools, childcare facilities, and many stores were closed; and (iii) on 23 March 2020, a far-reaching contact ban (*Kontaktsperre*) was imposed by government authorities; this included the prohibition of even small public gatherings as well as the closing of restaurants and all nonessential stores.

From the observed case numbers of COVID-19, we can quantify the impact of these measures on the disease spread using change point analysis. Essentially, we find that at each change point the spreading rate λ decreased by ~40%. At the first change point, assumed around 9 March 2020, λ decreased from 0.43 to 0.25, with 95% credible intervals (CIs) of [0.35, 0.51] and [0.20, 0.30], respectively. At the second change point, assumed around 16 March 2020, λ decreased to 0.15 (CI [0.12, 0.20]). Both changes in λ slowed the spread of the virus but still implied exponential growth (see red and orange traces in the figure).

To contain the disease spread, i.e., to turn exponential growth into a decline of new cases, the spreading rate has to be smaller than the recovery rate μ = 0.13 (CI [0.09, 0.18]). This critical transition was reached with the third change point, which resulted in λ = 0.09 (CI [0.06, 0.13]; see blue trace in the figure), assumed around 23 March 2020.

From the peak position of daily new cases, one could conclude that the transition from growth to decline was already reached at the end of March. However, the observed transient decline can be explained by a short-term effect that originates from a sudden change in the spreading rate (see Fig. 2C in the main text).

As long as interventions and the concurrent individual behavior frequently change the spreading rate, reliable short- and long-term forecasts are very difficult. As the figure shows, the three example scenarios (representing the effects up to the first, second, and third change point) quickly diverge from each other and, consequently, span a considerable range of future case numbers.

ON OUR WEBSITE

Read the full article at https://dx.doi.org/10.1126/science.abb9789

Inference and subsequent forecasts are further complicated by the delay of ~2 weeks between an intervention and the first useful estimates of the new λ (which are derived from the reported case numbers). Because of this delay, any uncertainty in the magnitude of social distancing in the previous 2 weeks can have a major impact on the case numbers in the subsequent 2 weeks. Beyond 2 weeks, the case numbers depend on our future behavior, for which we must make explicit assumptions. In sum, future interventions (such as lifting restrictions) should be implemented cautiously to respect the delayed visibility of their effects.

CONCLUSION: We developed a Bayesian framework for the spread of COVID-19 to infer central epidemiological parameters and the timing and magnitude of intervention effects. With such an approach, the effects of interventions can be assessed in a timely manner. Future interventions and lifting of restrictions can be modeled as additional change points, enabling short-term forecasts for case numbers. In general, our approach may help to infer the efficiency of measures taken in other countries and inform policy-makers about tightening, loosening, and selecting appropriate measures for containment of COVID-19. ■
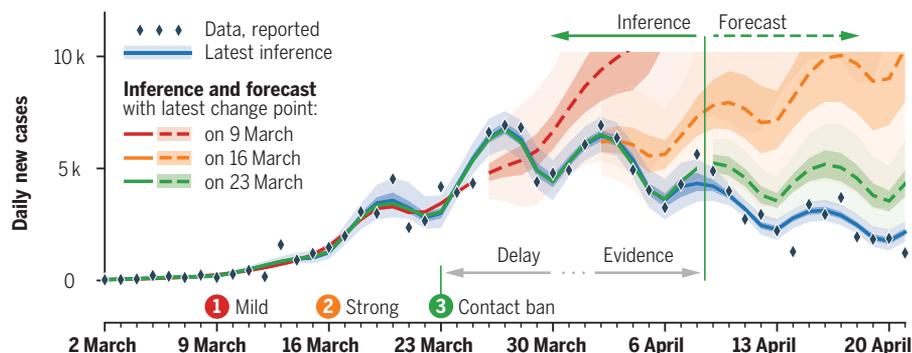
Bayesian inference of SIR model parameters from daily new cases of COVID-19 enables us to assess the impact of interventions. In Germany, three interventions (mild social distancing, strong social distancing, and contact ban) were enacted consecutively (circles). Colored lines depict the inferred models that include the impact of one, two, or three interventions (red, orange, or green, respectively, with individual data cutoff) or all available data until 21 April 2020 (blue). Forecasts (dashed lines) show how case numbers would have developed without the effects of the subsequent change points. Note the delay between intervention and first possible inference of parameters caused by the reporting delay and the necessary accumulation of evidence (gray arrows). Shaded areas indicate 50% and 95% Bayesian credible intervals.

CORONAVIRUS

Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions

Jonas Dehning[1]*, Johannes Zierenberg[1]*, F. Paul Spitzner[1]*, Michael Wibral[2], Joao Pinheiro Neto[1], Michael Wilczek[1,3]*, Viola Priesemann[1,3,4]*†

As coronavirus disease 2019 (COVID-19) is rapidly spreading across the globe, short-term modeling forecasts provide time-critical information for decisions on containment and mitigation strategies. A major challenge for short-term forecasts is the assessment of key epidemiological parameters and how they change when first interventions show an effect. By combining an established epidemiological model with Bayesian inference, we analyzed the time dependence of the effective growth rate of new infections. Focusing on COVID-19 spread in Germany, we detected change points in the effective growth rate that correlate well with the times of publicly announced interventions. Thereby, we could quantify the effect of interventions and incorporate the corresponding change points into forecasts of future scenarios and case numbers. Our code is freely available and can be readily adapted to any country or region.

During the initial outbreak of an epidemic, reliable short-term forecasts are key to estimate medical requirements and capacities and to inform and advise the public and decision makers (*1*). During this initial phase, three tasks are important to provide time-critical information for crisis mitigation: (i) establishing central epidemiological parameters, such as the basic reproduction number, that can be used for short-term forecasting; (ii) simulating the effects of different possible interventions aimed at the mitigation of the outbreak; and (iii) estimating the actual effects of the measures taken not only to make rapid adjustments but also to adapt short-term forecasts. Addressing these tasks is challenging because of the large statistical and systematic errors that occur during the initial stages of an epidemic, when case numbers are low. This is further complicated by the fact that mitigation measures are taken rapidly while the outbreak unfolds, but they take effect only after an unknown delay. To obtain reasonable parameter estimates for short-term forecasting and policy evaluation despite these complications, any prior knowledge available needs to be integrated into modeling efforts to reduce uncertainties. This includes knowledge about basic mechanisms of disease transmission, recovery, and preliminary estimates of epidemiological parameters from other countries or from closely related pathogens. The integration of prior knowledge, the quantitative assessment of the remaining uncertainties about epidemiological parameters, and the principled propagation of these uncertainties into forecasts is the domain of Bayesian modeling and inference (*2*, *3*).

We draw on an established class of models for epidemic outbreaks: The susceptible-infected-recovered (SIR) model (*4*–*7*) specifies population compartments and the rates at which they change (susceptible people becoming infectious and infectious people recovering). This simple model can be formulated in terms of coupled ordinary differential equations (in mean field), which enable analytical treatment (*8*, *9*) or fast evaluation (ideally suited for Bayesian inference). Accordingly, SIR-like models have been used to model epidemic spreads, from Bayesian Markov chain Monte Carlo (MCMC) parameter estimation (*10*–*12*) to detailed scenario discussions (*13*–*16*). This family of models has played a dominant role in the analyses of the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) pandemic, from inference (*17*–*19*) through scenario forecasting (*20*–*27*) to control strategies (*28*, *29*).

Here, we combined the SIR model (and generalizations thereof) with Bayesian parameter inference and augmented the model with a time-dependent spreading rate. The time dependence was implemented as potential change points in the spreading rate, which we assume to be driven by governmental interventions and the associated change of individual behavior (nonpharmaceutical interventions). On the basis of three distinct measures taken in Germany, we detected three corresponding change points in the reported COVID-19 case numbers. By 1 April 2020, we had reported evidence for the first two change points and predicted the third (*30*). Now, with data until 21 April 2020, we have evidence for all three

change points. First, the spreading rate decreased from 0.43 (with 95% credible interval, CI [0.35,0.51]) to 0.25 (CI [0.20,0.30]). This inferred decrease was initiated around 7 March 2020 (CI [3,10]) and matches the timing of cancellation of large public events such as trade fairs and soccer matches. Second, the spreading rate decreased further to 0.15 (CI [0.12,0.20]). This decrease was initiated around 16 March 2020 (CI [14,18]) and matches closure of schools, childcare facilities, and nonessential stores. Third, the spreading rate decreased further to 0.09 (CI [0.06,0.13]). This decrease was initiated around 24 March 2020 (CI [21,26]) and matches the strict contact ban (including the closing of all restaurants and nonessential stores), which was announced on 22 March 2020. Whereas already the first two change points strongly slowed the spread of the virus, the third change point can be associated with the start of a sustained decline in daily new cases.

Our framework is designed to infer the effectiveness of past measures and to explore potential future scenarios, along with propagating the respective uncertainties. In the following, we demonstrate the potential impact of timing and magnitude of change points and report our inference about the three past governmental interventions in Germany. Our framework can be readily adapted to any other country or region. The code (already including data sources from many other countries) and figures are available on GitHub (*31*).

Basic inference of central epidemiological parameters during the initial phase of the COVID-19 outbreak in Germany

To assess the general effect of different possible interventions on the spread of COVID-19 in Germany, we first focused on the initial phase of the outbreak, when no serious mitigation measures had been implemented. In the absence of interventions, an epidemic outbreak can be described by SIR models with a constant spreading rate (see the materials and methods). In Germany, the first serious interventions occurred around 9 March 2020 and affected the case reports after an observation delay [a combination of incubation period with a median of 5 to 6 days (*32*)] and a test delay (the time until a doctor is visited plus test evaluation time, which we assumed to be ~2 to 3 days, combined). Thus, to infer central epidemiological parameters, we considered the initial phase to be 2 March 2020 to 15 March 2020. The central epidemiological parameters estimated here will also be estimated under the full model with change points on the data records up to 21 April 2020, allowing for a consistency check.

We performed Bayesian inference for the central epidemiological parameters of an SIR model using MCMC sampling (Fig. 1). The central parameters are the spreading rate λ, the recovery rate μ, the reporting delay D, and

[1]Max Planck Institute for Dynamics and Self-Organization, 37077 Göttingen, Germany. [2]Campus Institute for Dynamics of Biological Networks, University of Göttingen, 37075 Göttingen, Germany. [3]Institute for the Dynamics of Complex Systems, University of Göttingen, 37077 Göttingen, Germany. [4]Bernstein Center for Computational Neuroscience, 37075 Göttingen, Germany.
*These authors contributed equally to this work.
†Corresponding author. Email: viola.priesemann@ds.mpg.de

the number of initially infected people I_0. We chose informative priors based on available knowledge for λ, μ, and D, and we chose uninformative priors for the remaining parameters (see the materials and methods). Also, we intentionally kept the informative priors as broad as possible so that the data would constrain the parameters (Fig. 1).

As median estimates, we obtained λ = 0.41, μ = 0.12, D = 8.6, and I_0 = 19 (see Fig. 1, C to H, for the posterior distributions and the 95% CIs). Converting to the basic reproduction number R_0 = λ/μ, we found a median R_0 = 3.4 (CI [2.4,4.7]), which is consistent with previous reports finding median values between 2.3 and 3.3 (18, 33, 34). Overall, the model exhibited good agreement with new cases (Fig. 1A) and cumulative cases (Fig. 1B), both of which show the expected exponential growth (linear in log-lin plot). The observed data were clearly informative about λ, I_0, and σ, as indicated by the difference between the priors (gray line) and posteriors (histograms) in Fig. 1, D to F. However, μ and D were largely determined by our prior choice of parameters (histograms match the gray line in Fig. 1, C and H). This

is to be expected for the initial phase of an epidemic outbreak, which is dominated by exponential growth.

To quantify the impact of possible interventions, we concentrated on the effective growth of active infections before and after the intervention. As long as the number of infections and recoveries are small compared with the population size, the number of active infections can be approximated by an exponential growth with effective growth rate λ* = λ – μ (see the materials and methods). As a consequence, λ and μ cannot be estimated independently. This was further supported by a systematic scan of the model's log-likelihood in the (λ – μ) space, which showed an equipotential line for the maximum likelihood (Fig. 1I). This strongly suggests that the growth rate λ* is the relevant free parameter with a median λ* = 28% (Fig. 1G). The control parameter of the dynamics in the exponential phase is thus the (effective) growth rate: If the growth rate is larger than zero (λ > μ), then case numbers grow exponentially; if, however, the growth rate is smaller than zero (λ < μ), then recovery dominates and new cases decrease. The two different dy-

namics (supercritical and subcritical, respectively) are separated by a critical point at λ* = 0 (λ = μ) (35).

Magnitude and timing of interventions matter for the mitigation of the outbreak

To simulate the effect of different possible interventions, we modeled the effects of interventions as change points in the spreading rate (see the materials and methods). We considered different, hypothetical interventions after the initial phase to show that both the amount of change in behavior (leading to a change in spreading rate λ; Fig. 2A) and the exact timing of the change (Fig. 2B) determine future development. Hypothetical interventions build on the inferred parameters from the initial phase (in particular median $λ_0$ = 0.41 and median μ = 0.12; Fig. 1) and were implemented as change points in the spreading rate from the inferred $λ_0$ to a new value $λ_1$. With such a change point, we modeled three potential scenarios of public behavior.

(i) No social distancing

Public behavior is unaltered and spread continues with the inferred rate ($λ_1 = λ_0$ with median $λ_1$ = 0.41 > μ).

(ii) Mild social distancing

The spreading rate decreases to 50% ($λ_1 = λ_0/2$ with median $λ_1$ = 0.21 > μ). Although people effectively reduce the number of contacts by a factor of two in this second scenario, the total number of reported cases continues to grow alongside scenario (i) for the time period of the reporting delay D (median D = 8.6 from initial phase; see below for a more constrained estimation). Also, we still observe an exponential increase of new infections after the intervention becomes effective, because the growth rate remains positive, $λ_1$* = $λ_1$ –μ > 0.
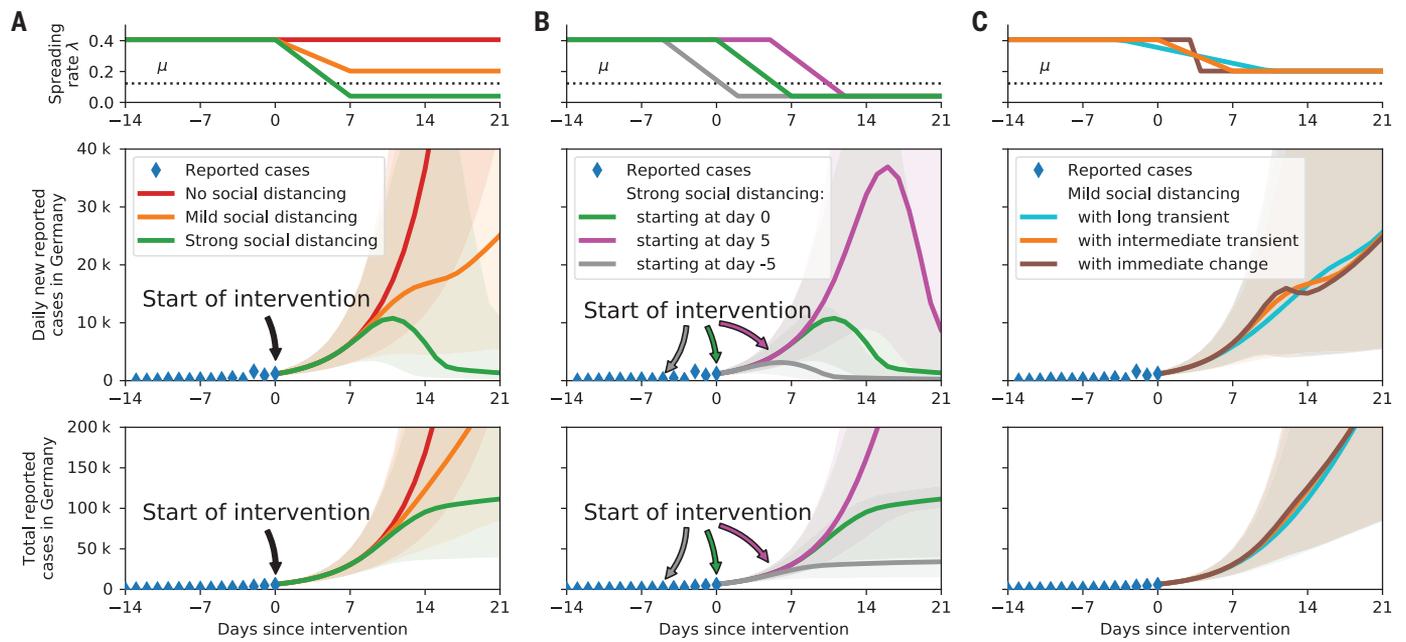
(iii) Strong social distancing

Here, the spreading rate decreases to 10% ($λ_1$ = $λ_0$/10 with median $λ_1$ = 0.04 < μ). The assumptions here are that contacts are severely limited, but even when people stay at home as much as possible, some contacts are still unavoidable. Even under such drastic policy changes, no effect is visible until the reporting delay D is over. Thereafter, a quick decrease in daily new infections manifests within 2 weeks (delay plus change point duration), and the total number of cases reaches a stable plateau. Only in this scenario (iii) a plateau is reached, because here the growth rate becomes negative, $λ_1$* < 0, which leads to decreasing numbers of new infections.

Furthermore, the timing of an intervention matters: Apart from the strength of an intervention, its onset time has great impact on the total case number (Fig. 2B). For example,
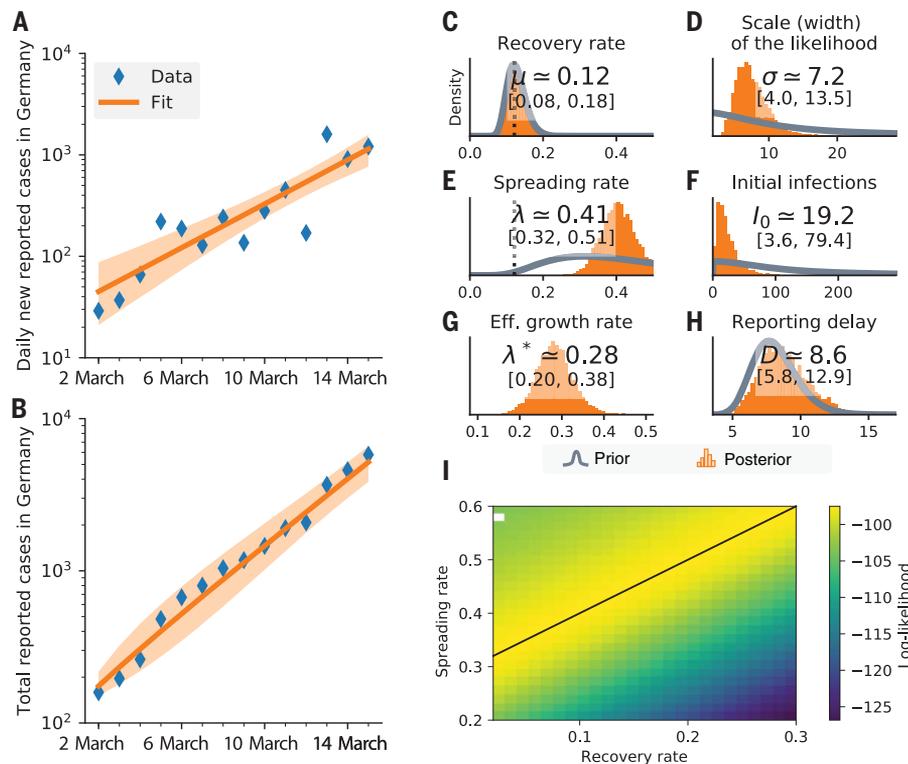
Fig. 1. Inference of central epidemiological parameters of the SIR model during the initial onset period, 2 to 15 March 2020. (**A**) The number of new cases and (**B**) the total (cumulative) number of cases increase exponentially over time. (**C** to **H**) Prior (gray) and posterior (orange) distributions for all model parameters: estimated spreading rate λ, recovery rate μ, and reporting delay D between infection date and reporting date; number of cases I_0 at the start of the simulation; scale-factor σ of the width of the likelihood distribution; and the effective growth rate λ* = λ – μ. (**I**) Log-likelihood distribution for different combinations of λ and μ. A linear combination of λ and μ yields the same maximal likelihood (black line). White rectangle indicates that the inference did not converge.

Fig. 2. The timing and effectiveness of interventions strongly affect future COVID-19 cases. (**A**) We assume three different scenarios for interventions starting on 16 March 2020: (i, red) no social distancing, (ii, orange) mild social distancing, or (iii, green) strict social distancing. (**B**) Delaying the restrictions has a major impact on case numbers: strict restrictions starting on 16 March 2020 (green), 5 days later (magenta), or 5 days earlier (gray). (**C**) Comparison of the time span over which interventions ramp up to full effect. For all ramps that are centered around the same day, the resulting case numbers are fairly similar; however, a sudden change of the spreading rate can cause a temporary decrease of daily new cases (although $\lambda > \mu$ at all times; brown).

focusing on the strong intervention (iii), by which a stable plateau is reached, the effect of advancing or delaying the change point by just five days leads to more than a threefold difference in cumulative cases.

Whereas we found that the timing of an intervention has a great effect on case numbers, the duration over which the change takes place has only minor effect if the intervals of change are centered around the same date. In Fig. 2C, we illustrate the adjustment of $\lambda_0 \rightarrow \lambda_1$ for mild social distancing with durations of 14, 7, and 1 day(s). The change point duration is a simple way to incorporate variability in individual behavior and is not linked to the reporting delay D. As an interesting effect, a sudden change in the spreading rate can lead to a temporary decrease of new case numbers despite the fact that the effective growth rate remains positive at all times.

Change point detection for the spread of COVID-19 in Germany

To model real-world data, we further refined the SIR model. We accounted for systematic variations of case reports throughout the week (lower case numbers toward the weekend) by explicitly modeling a weekly reporting modulation (see the materials and methods). Comparisons confirmed that models with this correction outperformed those without (see table S2). In the supplemental material, we further generalize our model to include an

explicit incubation period [as in susceptible-exposed-infectious-recovered (SEIR) models; fig. S3] that yields results consistent with our main model.

We incorporated the effect of nonpharmaceutical interventions into our models by introducing flexible change points in the spreading rate (see the materials and methods). During the COVID-19 outbreak in Germany, governmental interventions occurred in three stages from (i) the cancellation of large events with more than 1000 participants (around 9 March 2020), through (ii) closing of schools, childcare facilities, and most stores (in effect 16 March 2020), to (iii) the contact ban and closing of all nonessential stores (in effect 23 March 2020). The aim of all these interventions was to reduce the (effective) growth rate $\lambda^* = \lambda - \mu$. As soon as the growth rate becomes notably negative ($\lambda^* < 0$), the number of daily new cases decreases after the respective reporting delay.

Detecting change points in the spreading rate and quantifying the amount of change as quickly as possible becomes a central modeling challenge when short-term forecasts are required. To address this challenge, we assumed an initial spreading rate λ_0 (the exponential growth phase; Fig. 1) and up to three potential change points motivated by German governmental interventions: In our modeling, the first change point ($\lambda_0 \rightarrow \lambda_1$) is expected around 9 March 2020 (t_1) as a result of the official recommendations to cancel large events. A sec-

ond change point ($\lambda_1 \rightarrow \lambda_2$) is expected around 16 March 2020 (t_2), when schools and many stores were closed. A third change point ($\lambda_2 \rightarrow \lambda_3$) is expected around 23 March 2020 (t_3), when all nonessential stores were closed and a contact ban was enacted. We modeled the behavioral changes that were introduced at these change points to unfold over a few days (Δt_i), but the changes in duration can be partly compensated by changes in the onset time (t_i) (see Fig. 2C, scenarios). We chose priors for all parameters based on the information available to us up to 28 March 2020 (see the materials and methods). In addition, we performed a sensitivity analysis by using wider priors in the supplemental material (figs. S5 to S7 and table S2), which yielded consistent results. On 28 March 2020, the data were already informative about the first change point, and thereby helped to inform our forecast scenarios.

The inferred parameters for the models with change points are consistent with the inferred parameters from the exponential onset phase (Figs. 1 and 3 and figs. S1 and S2). In particular, all estimated λ_0 values from models with multiple change points are compatible with the value of the model without change points (during the exponential onset phase, $\lambda_0 = 0.41$, CI [0.32,0.51], assuming a stationary λ until 15 March 2020; Fig. 1E). Also, the scale factor σ and the number of initial infections I_0 for the models with change points are consistent with the initial model inference

Fig. 3. Bayesian analysis of the German COVID-19 data until 21 April 2020 reveals three change points that are consistent with three major governmental interventions.
(**A**) Time-dependent model estimate of the effective spreading rate $\lambda^*(t)$.
(**B**) Comparison of daily new reported cases and the model (green solid line for median fit with 95% CIs, dashed line for median forecast with 95% CI). Inset, same data in log-lin scale.
(**C**) Comparison of total reported cases and the model [same representation as in (B)]. (**D** to **F**) Priors (gray lines) and posteriors (green histograms) of all model parameters; inset values indicate the median and 95% CIs of the posteriors. For the same model with one or two change points, please see the corresponding figures in the supplementary materials (figs. S1 and S2 and table S2).
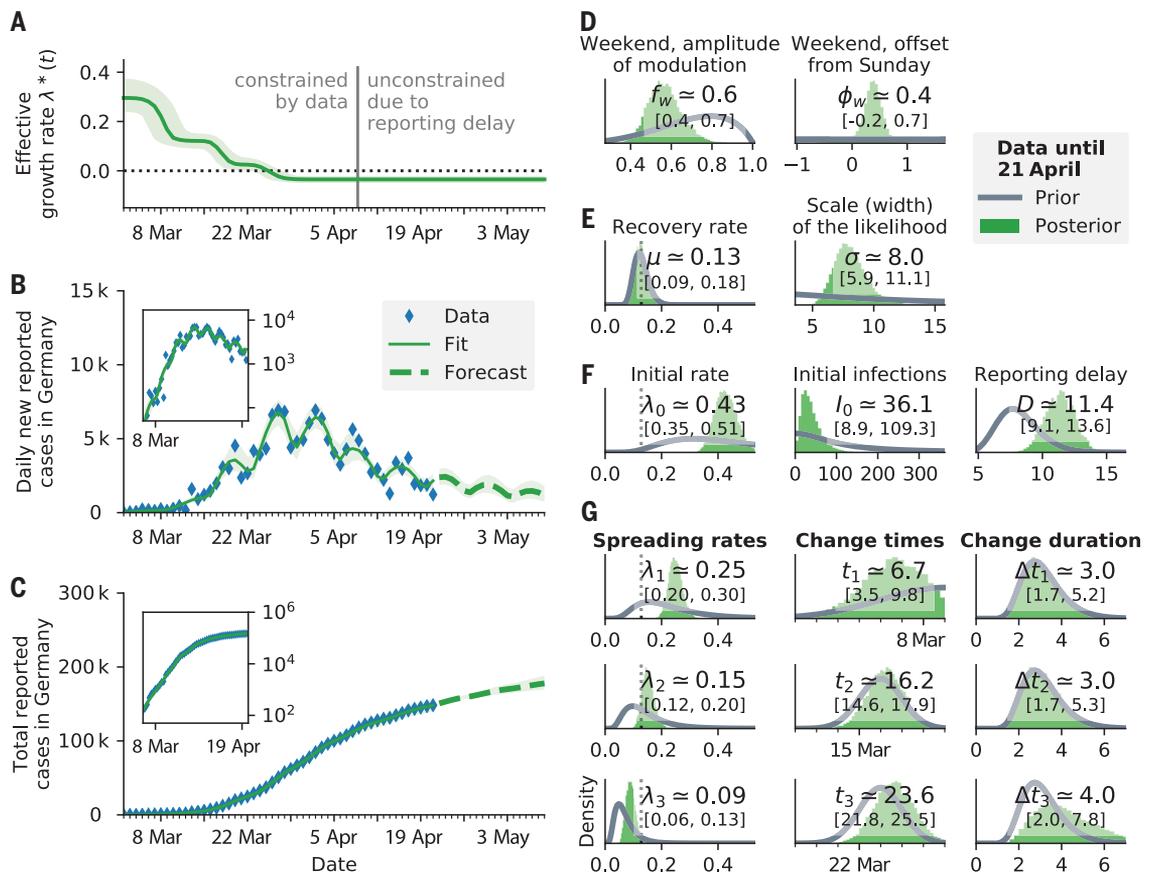
during the exponential onset phase (Fig. 1, D to F).

Models with two or three change points fit the observed data better

The models with three change points describe the data better than models with fewer change points, as indicated by the leave-one-out (LOO) cross-validation–based Bayesian model comparison (*36*) (lowest LOO score in Table 1). However, the LOO scores of the model with two and three change points differ by <1 SE. This originates from an extended duration of the second change point in the two-change-point model, which partially captures the effect of the third intervention. As expected, the models with none or a single change point had LOO scores that were at least 1 SE higher (worse) than those of the best models, and we did not consider them further.

When comparing our inference based on three change points with the number of confirmed cases, we found them to largely match (Fig. 3, B and C). The dominant periodic change in the daily new reported cases (Fig. 3B) was well described by the weekday modulation. In addition to the periodic change, the daily new case numbers also reflect the fairly sudden change of the spreading rate at the change points (compare Fig. 2 and fig. S4 for the effect

of change points without the modulation). The cumulative effect of change points manifested in an overarching decay in new case numbers that was visible after 5 April 2020 and followed the third change point (with reporting delay).

Change point detection quantifies the effect of nonpharmaceutical interventions

Ideally, detected changes can be related to specific mitigation measures so that one gains insights into the effectiveness of different measures (Fig. 3). Our model comparison favored the model with three change points with the following posteriors of the parameters: First, $\lambda(t)$ decreased from $\lambda_0 = 0.43$ (CI [0.35,0.51]) to $\lambda_1 = 0.25$ (CI [0.20,0.30]). The date of the change point was inferred to be 7 March 2020 (CI [3,10]), and this inferred date matches the timing of the first governmental intervention, which included cancellations of large events, as well as increased awareness. After this first intervention, the (effective) growth rate $\lambda^*(t) = \lambda(t) - \mu$ decreased by more than a factor of 2, from median $\lambda_0 - \mu = 0.3$ to median $\lambda_1 - \mu = 0.12$, given that the recovery rate was inferred as $\mu = 0.13$ (CI [0.09,0.18]). At the second change point, $\lambda(t)$ decreased from $\lambda_1 = 0.25$ to $\lambda_2 = 0.15$ (CI [0.12,0.20]), which is larger than our prior assumption. The date of this change point was inferred to be 16 March 2020 (CI [14,18]), and

this inferred date matches the timing of the second governmental intervention including closing schools and some stores. After this second intervention, the median growth rate became $\lambda^*(t) = \lambda_2 - \mu = 0.02 \approx 0$ (CI [0.00, 0.06]) and is thus in the vicinity of the critical point yet still slightly positive. The first two interventions thereby mitigated the spread of COVID-19 in Germany by drastically reducing the growth rate, but did most likely not lead to a sustained decline of new infections. A third change point, when $\lambda(t)$ decreased from $\lambda_2 = 0.15$ to $\lambda_3 = 0.09$ (CI [0.06,0.13]), was inferred on 24 March 2020 (CI [21,26]), and this inferred date matches the timing of the third governmental intervention including contact ban and closing of all non-essential stores. Only after this third intervention did the median (effective) growth rate, $\lambda^*(t) = \lambda_3 - \mu = -0.03 < 0$ (CI [−0.05,−0.02]), become negative, enabling a sustained decrease in the number of new infections. In summary, we have related the change points to the nonpharmaceutical interventions to quantify their effect.

Discussion

We have presented a Bayesian approach for monitoring of the effect of nonpharmaceutical interventions on the epidemic outbreak of an infectious disease. Using the example of the COVID-19 outbreak in Germany, we applied

this approach to infer the central epidemiological parameters and three change points in the spreading rate from the number of reported cases. We showed that change points in the spreading rate affect the confirmed case numbers with a delay of ~2 weeks (median reporting delay of D = 11.4 days plus a median change point duration of 3 days). Thereby, we were able to relate the inferred change points to the three major governmental interventions in Germany: In our model with three change points, we found a clear reduction of the spreading rate related to each governmental intervention and the concurring adaptation of individual behavior (Fig. 3), (i) the cancellation of large events with >1000 participants (around 9 March 2020), (ii) the closing of schools, childcare centers, and most stores (in effect 16 March 2020), and (iii) the contact ban and closing of all nonessential stores (in effect 23 March 2020).

Our results suggest that the full extent of interventions and the concurring adaptation of behavior lead to a swift and sustained decrease of daily new cases. The first two interventions brought a reduction of the growth rate λ* from 30% to 12% and down to 2% (CI [0%, 6%]), respectively. These numbers still implied the possibility of exponential growth. Only with the third intervention did we find that the epidemic changed from growth to decay. However, the decay rate of ~–3% (CI [–5%,–2%]) remained close to zero. Therefore, even a minor increase in the spreading rate may again switch the dynamics to the unstable regime with exponential growth.

We used a formal Bayesian model comparison to validate the presence of change points. Our model comparison ruled out models with fewer than two change points (Table 1 and table S2). Although this may seem trivial, it has important consequences for making short-term forecasts that decision makers rely on. Demonstrating and quantifying the effect of past change points can be used to formulate priors for the effects of future, similar change points. These priors help to project the effects of more recent change points into future forecasts, even when these change points are not yet apparent in the reported case numbers.

Consequently, it is important to look out for and identify potential change points as early as possible to incorporate them into forecasts.

The detection of change points and their interpretation depend crucially on an accurate estimate of the reporting delay D. Therefore, the validity of its estimate should be evaluated. In our model, D contains at least three distinct factors: the biological incubation period (median 5 to 6 days) (32), an additional delay from first symptoms to symptoms motivating a test (1 to 3 days), and a possible delay before testing results come in (1 to 4 days). The sum of these delays seems compatible with our inferred median delay of D = 11.4 days, especially given the wide range of reported incubation periods.

We chose to keep our main model comparatively simple because of the small number of data points initially available during an epidemic outbreak. With few data points, only a limited number of parameters can be effectively constrained. Therefore, we chose to approximate a time-dependent spreading rate $\lambda(t)$ by using episodes of constant spreading rates λ_i that are separated by change points where a transition occurs. Although we introduced fairly broad priors on the spreading rates, we obtained comparably narrow posterior distributions for each spreading rate λ_i (Fig. 3). We additionally evaluated extensions of our main model with three change points, e.g., by explicitly taking into account the incubation period (fig. S3). These models yield consistent results for the three change points, and all have LOO scores within 1 SE of each other. Thus, we consider our main model to be sufficient to describe case numbers in Germany at present.

Our framework can be easily adapted to other countries and enables incorporation of future developments. For other countries, or for forecasts within smaller communities (e.g., federal states or cities), additional details may become important, such as explicit modeling of incubation time distributions (17, 37) (i.e., as in fig. S3), spatial heterogeneity (17, 21), isolation effects (20, 37), subsampling effects hiding undetected cases even beyond the re-

porting delay (38, 39), or the age and contact structure of the population (26). In countries where major changes in test coverage are expected, this will have to be included as well. The methodology presented here is in principle capable of incorporating such details. It also lends itself to modeling of continuous drifts in the spreading rate, e.g., reflecting reactions of the public to news coverage of a catastrophic situation or people growing tired of mitigation measures. Such further adaptations, however, can only be performed on a per-country basis by experts with an intimate knowledge of the local situation. Our approach provides a solid and extensible base for this. For Germany, several developments that occurred after the time span of the presented analysis should be included in the model. First, people may have transiently changed their behavior over the Easter holidays; second, we expect a series of change points, as well as continuous drifts, with governments trying to ease and calibrate mitigation measures. Third, extensions to hierarchical models will enable regional assessments, e.g., on the level of federal states.

In Germany, following the three major governmental interventions in March, effective growth rates remained close to zero and warranted careful consideration of future measures. With the data available until 21 April 2020, we estimated an effective growth rate of ~–3% for the beginning of April; the growth rate remained close to zero—the threshold between exponential growth or decay. Together with the delay of ~2 weeks between infection and case report, a growth rate close to zero warrants caution in lifting restrictions for two reasons. First, lifting restrictions too much will quickly lead to renewed exponential growth, and second, we would be effectively blind to this worsened situation for nearly 2 weeks, during which time transmission would be uninhibited. This may result in a growth in case numbers beyond the level that the health system can cope with, especially if active cases are not close to zero before lifting restrictions.

In conclusion, our Bayesian approach allows the detection and quantification of the effect of nonpharmaceutical interventions and, combined with potential subsequent interventions, the forecasting of future case number scenarios. Our analysis highlights the importance of precise timing and magnitude of interventions for future case numbers. It also stresses the importance of including the reporting delay D between the date of infection and the date of the confirmed case in the model. The reporting delay D, together with the time required to implement interventions, means that changes in our behavior today can only be detected in confirmed cases in 2 weeks' time. This delay, combined with a spreading rate close to zero, indicates that careful planning of future measures is essential.

Table 1. Model comparison showing that the three-change-point model describes the data best. Shown is the LOO cross-validation for the main models (SIR with weekend correction) and a different number of change points. Lower LOO scores represent a better match between model and data.

Model	LOO score	Effective number of parameters (pLOO)
Zero change points	927 ± 9	8.31
One change point	819 ± 16	13.46
Two change points	796 ± 17	12.53
Three change points	787 ± 17	13.42

Materials and methods

As a basis for our Bayesian inference and the forecast scenarios, we used the differential equations of the well-established SIR model. We also tested the robustness of our results by means of more sophisticated models, in particular an SEIR-like model that explicitly incorporates an incubation period (fig. S3). Although SIR model dynamics are well understood in general, here, our main challenge was to estimate model parameters specifically for the COVID-19 outbreak and to use them for forecasting. To that end, we combined a Bayesian approach to incorporate prior knowledge with MCMC sampling to compute the posterior distribution of the parameters and to sample from it for forecasting. Put simply, we first estimated the parameter distribution that best described the observed situation, and then we used many samples from this parameter distribution to evolve the model equations and thus forecast future developments.

The data used come from the Johns Hopkins University Center for Systems Science and Engineering (JHU CSSE) dashboard (*40*). The JHU CSSE provides up-to-date data on COVID-19 infections, usually a few days ahead of official German sources. The exact versions of the data and code are available at (*31*). Data were incorporated until 21 April 2020. Note that after this cutoff date, additional modeling of the effects of behavioral changes over the Easter holidays became necessary.

Simple model: SIR model with stationary spreading rate

We consider a time-discrete version of the standard SIR model. In short, we assume that the disease spreads at rate λ from the infected population compartment (I) to the susceptible compartment (S), and that the infected population compartment recovers (R) at rate μ. This well-established model for disease spreading can be described by the following set of (deterministic) ordinary differential equations [see, e.g., (*5*), (*6*), (*20*)]. Within a population of size N,

$$\begin{aligned}
\frac{dS}{dt} &= -\lambda \frac{SI}{N} \\
\frac{dI}{dt} &= \lambda \frac{SI}{N} - \mu I \\
\frac{dR}{dt} &= \mu I
\end{aligned} \qquad (1)$$

During the onset phase of an epidemic, only a very small fraction of the population is infected (I) or recovered (R), and thus $S \approx N \gg I$ such that $S/N \approx 1$. Therefore, the differential equation for the infected reduces to a simple linear equation, exhibiting an exponential growth

$$\frac{dI}{dt} = (\lambda - \mu)I, \text{ solved by } I(t) = I(0)e^{(\lambda-\mu)t} \qquad (2)$$

Because our dataset is discrete in time ($\Delta t = 1$ day), we solve the above differential equations with a discrete time step ($dI/dt \approx \Delta I/\Delta t$), such that

$$\begin{aligned}
S_t - S_{t-1} &= -\lambda \Delta t \frac{S_{t-1}}{N} I_{t-1} &&=: -I_t^{\text{new}} \\
R_t - R_{t-1} &= \mu \Delta t I_{t-1} &&=: R_t^{\text{new}} \\
I_t - I_{t-1} &= \left(\lambda \frac{S_{t-1}}{N} - \mu\right)\Delta t I_{t-1} &&= I_t^{\text{new}} - R_t^{\text{new}}
\end{aligned} \qquad (3)$$

I_t models the number of all (currently) active infected people, whereas I_t^{new} is the number of new infections that will eventually be reported according to standard World Health Organization (WHO) convention. We explicitly include a reporting delay D between new infections I_t^{new} and newly reported cases (C_t) as

$$C_t = I_{t-D}^{\text{new}} \qquad (4)$$

We begin our simulations at time $t = 0$ with I_0 infected cases and start including real-word data of reported cases $\hat{C}_t$ from day $t > D$ (see below for a parameterization).

In our model, we do not explicitly incorporate the inflow of additional infected people by travel for two reasons. First, we implicitly model the initial surge of infections with I_0. Second, previous work showed that travel during the outbreak has only modest effects on the dynamics, e.g., travel restrictions in China merely delayed the exponential spread if not combined with reductions of spreading (*41*).

Full model: SIR model with weekly reporting modulation and change points in spreading rate

Our change point detection builds on a generalization of the simple SIR model with stationary spreading rate. We now assume that the spreading rate λ_i, $i = 1, ..., n$, may change at certain time points t_i from λ_{i-1} to λ_i, linearly over a time window of Δt_i days. Thereby, we account for mitigation measures, which were implemented in Germany step by step. Thus, the parameters t_i, Δt_i, and λ_i are added to the parameter set of the simple model above, and the differential equations are augmented by the time-varying λ_i.

In addition, we include a weekly modulation to account for lower case reports around the weekend, which subsequently accumulate during the week. To model the systematic variation of case reports during the week, we adapt the newly reported cases by a reporting fraction

$$\begin{aligned}
C_t &= I_{t-D}^{\text{new}}\left(1 - f(t)\right), &&\text{with} \\
f(t) &= (1 - f_w) \cdot \left(1 - \left|\sin\left(\frac{\pi}{7}t - \frac{1}{2}\Phi_w\right)\right|\right)
\end{aligned} \qquad (5)$$

Table 2. Priors on the model parameters for the SIR model with stationary spreading rate.

Parameter	Variable	Prior distribution
Spreading rate	λ	LogNormal[log(0.4), 0.5]
Recovery rate	μ	LogNormal[log(1/8), 0.2]
Reporting delay	D	LogNormal[log(8), 0.2]
Initially infected	I_0	HalfCauchy(100)
Scale factor	σ	HalfCauchy(10)

Table 3. Priors on the model parameters for the SIR model with change points and weekly reporting modulation.

Parameter	Variable	Prior distribution
Change points	t_1	Normal(2020/03/09, 3)
	t_2	Normal(2020/03/16, 1)
	t_3	Normal(2020/03/23, 1)
Change duration	Δt_i	LogNormal[log(3), 0.3]
Spreading rates	λ_0	LogNormal[log(0.4), 0.5]
	λ_1	LogNormal[log(0.2), 0.5]
	λ_2	LogNormal[log(1/8), 0.5]
	λ_3	LogNormal[log(1/16), 0.5]
Recovery rate	μ	LogNormal[log(1/8), 0.2]
Reporting delay	D	LogNormal[log(8), 0.2]
Weekly modulation amplitude	f_w	Beta(mean = 0.7, std = 0.17)
Weekly modulation phase	Φ_w	VonMises(mean = 0, κ = 0.01) (nearly flat)
Initially infected	I_0	HalfCauchy(100)
Scale factor	σ	HalfCauchy(10)

where f_w and Φ_w will also be constrained by the data.

Estimating model parameters with Bayesian MCMC

We estimate the set of model parameters $\theta = \{\lambda_i, t_i, \mu, D, \sigma, I_0, f_w, \Phi_w\}$ using Bayesian inference with MCMC. The parameter σ is the scale factor for the width of the likelihood $P(\hat{C}_t|\theta)$ between observed data and model (see below). Our implementation relies on the Python package PyMC3 (42) with NUTS (No-U-Turn Sampling) (43) using multiple, independent Markov chains. The structure of our approach is as follows.

Initialization of the Markov chains through variational inference

The posterior is approximated by Gaussian distributions ignoring correlations between parameters through automatic differentiation variational inference (ADVI) (44), which is implemented in PyMC3. From this distribution, four starting points for four chains are sampled.

Burn-in phase

Each chain performs 1000 burn-in (tuning) steps using NUTS, which are not recorded. This serves as equilibration to sample from an equilibrium distribution in the next step.

Sampling phase

Each chain performs 4000 steps, which are used to approximate the posterior distribution. To ensure that the chains are equilibrated and sampled from the whole posterior distribution (ergodicity), we verified that the R-hat statistic is below 1.05, which is implemented in PyMC3. The rank-normalized R-hat diagnostic tests for lack of convergence by comparing the variances within chains and between chains: For identical within-chain and between-chain variances, R-hat becomes 1, indicating convergence. For well-converged chains, the resulting samples will describe the real-world data well, so that consistent forecasts are possible in the forecast phase.

Forecast using MCMC samples

For the forecast, we take all samples from the MCMC step and continue time integration according to different forecast scenarios. Note that the overall procedure yields an ensemble of forecasts as opposed to a single forecast that would be solely based on one set of (previously optimized) parameters.

MCMC sampling details

Each MCMC step requires us to propose a new set of parameters θ to generate a (fully deterministic) time series of new infected cases $C(\theta) = \{C_t(\theta)\}$ of the same length as the observed real-world data $\hat{C} = \{\hat{C}_t\}$ and to accept or reject θ. In our case, the NUTS implementation (in PyMC3) first proposes a new set of parameters θ based on an advanced gradient-based algorithm and subsequently accepts or rejects it such that the resulting samples reflect the posterior distribution $p(\theta|\hat{C}) \propto p(\hat{C}|\theta)p(\theta)$ where $p(\hat{C}|\theta)$ is the likelihood for the data given the parameters and $p(\theta)$ is the prior distribution of the parameters (see below). The likelihood quantifies the similarity between model outcome and the available real-world time series. Here, the likelihood is the product over local likelihoods

$$p(\hat{C}_t|\theta) \sim \text{StudentT}_{\nu=4}\left(\text{mean} = C_t(\theta), \text{width} = \sigma\sqrt{C_t(\theta)}\right),$$

quantifying the similarity between the model outcome for one time point t, $C_t(\theta)$, and the corresponding real-world data point $\hat{C}_t$. We chose the Student's t distribution because it resembles a Gaussian distribution around the mean but features heavy tails, which make the MCMC more robust with respect to outliers (45) and thus reporting noise. The case-number–dependent width is motivated by observation noise through random subsampling (38), resulting in a variance proportional to the mean. Our likelihood neglects any noise in the dynamic process, because the SIR model is deterministic, but could be in principle extended to incorporate typical demographic noise from stochastic spreading dynamics (35, 46).

Priors that constrain model parameters

Because short-term forecasts are time critical at the onset of an epidemic, the available real-world data are typically not informative enough to identify all free parameters or to empirically find their underlying distributions. We therefore chose informative priors on initial model parameters where possible and complemented them with uninformative priors otherwise. Our choices are summarized in Table 2 for the simple model, i.e., the SIR model with stationary spreading rate for the exponential onset phase, and in Table 3 for the full model with change points and are discussed in the following.

Priors for the simple model (Table 2)

To constrain our simple model, an SIR model with stationary spreading rate for the exponential onset phase, we chose the following informative priors. Because of the ambiguity between the spreading and recovery rate in the exponential onset phase (see description of the simple model), we chose a narrow log-normal prior for the recovery rate $\mu \sim \text{LogNormal}[\log(1/8), 0.2]$ with a median recovery time of 8 days (20). Note that our implementation of μ accounts for the recovery of infected people and isolation measures because it describes the duration during which a person can infect others. For the spreading rate, we

assume a broad log-normal prior distribution $\lambda \sim \text{LogNormal}[\log(0.4), 0.5]$ with a median of 0.4. This way, the prior for $\lambda - \mu$ has a median of 0.275 and the prior for the base reproduction number ($R_0 = \lambda/\mu$) has a median of 3.2, consistent with the broad range of previous estimates (18, 33, 34). In addition, we chose a log-normal prior for the reporting delay $D \sim \text{LogNormal}[\log(8), 0.2]$ to incorporate both the incubation time between 1 and 14 days with median 5 days (32) plus a delay from infected people waiting to contact the doctor and get tested (assumed as 3 days).

The remaining model parameters are constrained by uninformative priors, in practice the half-Cauchy distribution (47). The half-Cauchy distribution, $\text{HalfCauchy}(x,\beta) = 2/\pi\beta[1 + (x/\beta)^2]$, is essentially a flat prior from zero to $O(\beta)$ with heavy tails beyond. Thereby, β merely sets the order of magnitude that should not be exceeded for a given parameter. We chose for the number of initially infected people in the model (16 days before first data point) $I_0 \sim \text{HalfCauchy}(100)$ assuming an order of magnitude $O(100)$ and below. In addition, we chose the scale factor of the width of the likelihood function as $\sigma \sim \text{HalfCauchy}(10)$; this choice means that the variance in reported numbers may be up to a factor of 100 larger than the actual reported number.

Priors for the full model (Table 3)

To constrain our full model, an SIR model with weekly reporting modulation and change points in the spreading rate, we chose the same priors as for the simple model but added the required priors associated with the change points. In general, we assume that each set of interventions (together with the increasing awareness) leads to a reduction (and not an increase) of λ at a change point. Because we cannot know yet the precise reduction factor, we adhere to assuming a reduction by ~50%, but always with a fairly wide uncertainty so that, in principle, even an increase at the change point would be possible. We model the time dependence of λ as change points, not as continuous changes, because the policy changes were implemented in three discrete steps, which were presumably followed by the public in a timely fashion.

For the spreading rates, we chose log-normal distributed priors as in the simple model. In particular, we chose for the initial spreading rate the same prior as in the simple model, $\lambda_0 \sim \text{LogNormal}[\log(0.4), 0.5]$; after the first change point $\lambda_1 \sim \text{LogNormal}[\log(0.2), 0.5]$, assuming the first intervention to reduce the spreading rate by 50% from our initial estimates ($\lambda_0 \approx 0.4$) with a broad prior distribution; after the second change point $\lambda_2 \sim \text{LogNormal}[\log(1/8), 0.5]$, assuming the second intervention to reduce the spreading rate to the level of the recovery rate, which would lead

to a stationary number of new infections and corresponds approximately to a reduction of λ at the change point by 50%; and after the third change point $\lambda_3 \sim$ LogNormal[log(1/16), 0.5], assuming the third intervention to reduce the spreading rate again by 50%. With that intervention, we expect λ_3 to be smaller than the recovery rate μ, causing a decrease in new case numbers and a saturation of the cumulative number of infections.

For the timing of change points, we chose normally distributed priors. In particular, we chose $t_1 \sim$ Normal(2020/03/09,3) for the first change point because on the weekend of 8 March 2020, large public events such as soccer matches or fairs were canceled. For the second change point, we chose $t_2 \sim$ Normal(2020/03/16,1), because on 15 March 2020, the closing of schools and other educational institutions, along with the closing of nonessential stores, was announced and implemented on the following day. Restaurants were allowed to stay open until 6:00 p.m. For the third change point, we chose $t_3 \sim$ Normal(2020/03/23,1), because on 23 March 2020, a far-reaching contact ban (*Kontaktsperre*), which includes the prohibition of even small public gatherings as well as complete closing of restaurants and nonessential stores, was imposed by the government. Further policy changes may occur in future; however, for now, we do not include more change points.

The change points take effect over a certain time period Δt_i for which we chose $\Delta t_i \sim$ LogNormal[log(3), 0.3] with a median of 3 days over which the spreading rate changes continuously as interventions have to become effective. The precise duration of the transition has hardly any effect on the cumulative number of cases (Fig. 2C). We assumed a duration of 3 days because some policies were not announced at the same day for all states within Germany; moreover, the smooth transition also can absorb continuous changes in behavior.

The number of tests that are performed and reported vary regularly over the course of a week and are especially low during weekends. To account for this periodic variation, we modulated the number of inferred cases by the absolute value of a sine function with, in total, a period of 7 days. We chose this function because it is a nonsymmetric oscillation, fitting the weekly variation of cases on a phenomenological level. For the amplitude of the modulation, we chose a weakly informative Beta prior: $f_w \sim$ Beta(mean = 0.7, std = 0.17) and for the phase a nearly flat circular distribution: $\Phi_w \sim$ VonMises(mean = 0, κ = 0.01).

Model comparison

Because change-point detection entails evaluating models with different numbers of parameters, some form of fair model comparison is needed. This is necessary to compensate for the higher flexibility of more complex models, because this flexibility carries the risk of overfitting and overconfident forecasts. The standard approach to avoid overfitting in machine learning is cross-validation, which has recently also been advocated for Bayesian model comparison [e.g., (*3*, *36*)], especially for models used for predictions and forecasts. Thus, one would ideally like to compare the models with different numbers of change points by the probability they assign to previously unobserved data points. Technically this is measured by their out-of-sample prediction accuracy, i.e., their log pointwise predictive density (lppd):

$$\text{lppd} = \sum_{i=1}^{N} \log\left(\int p(y_i^{\text{os}}|\theta) p_{\text{post}}(\theta) d\theta \right) \quad (6)$$

where the vector $[y_1^{\text{os}}, ..., y_N^{\text{os}}]$ is a an out-of-sample dataset of N new data points and $p_{\text{post}}(\theta) = p_{\text{post}}(\theta \mid y, M_j)$ is the posterior distribution of the parameters given the in-sample data y and the model M_j. In practice, the integral is approximated using a sufficient number of samples from $p_{\text{post}}(\theta)$. However, this approach is only reasonable if a sufficient amount of out-of-sample data are available, which is not the case in the early stages of a disease outbreak. Therefore, the pointwise out-of-sample prediction accuracy was approximated using LOO in PyMC3 to compute Eq. 6 individually for each left-out data point based on the model fit to the other data points. The sum of these values, multiplied by a factor of –2, then yields the LOO cross-validation (LOO-CV) score. Lower LOO-CV scores imply better models.

REFERENCES AND NOTES

1. M. Enserink, K. Kupferschmidt, With COVID-19, modeling takes on life and death importance. *Science* 367, 1414–1415 (2020). doi: 10.1126/science.367.6485.1414-b; pmid: 32217707
2. E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge Univ. Press, 2003).
3. A. Gelman et al., *Bayesian Data Analysis* (CRC Press, ed. 3, 2013).
4. W. O. Kermack, A. G. McKendrick, G. T. Walker, A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. A* 115, 700–721 (1927). doi: 10.1098/rspa.1927.0118
5. H. Hethcote, The mathematics of infectious diseases. *SIAM Rev.* 42, 599–653 (2000). doi: 10.1137/S0036144500371907
6. H. Andersson, T. Britton, *Stochastic Epidemic Models and Their Statistical Analysis* (Springer, 2000). https://doi.org/10.1007/978-1-4612-1158-7.
7. N. C. Grassly, C. Fraser, Mathematical models of infectious disease transmission. *Nat. Rev. Microbiol.* 6, 477–487 (2008). doi: 10.1038/nrmicro1845; pmid: 18533288
8. R. Parshani, S. Carmi, S. Havlin, Epidemic threshold for the susceptible-infectious-susceptible model on random networks. *Phys. Rev. Lett.* 104, 258701 (2010). doi: 10.1103/PhysRevLett.104.258701; pmid: 20867419
9. T. Harko, F. S. N. Lobo, M. K. Mak, Exact analytical solutions of the susceptible-infected-recovered (SIR) epidemic model and of the SIR model with equal death and birth rates. *Appl. Math. Comput.* 236, 184–194 (2014). doi: 10.1016/j.amc.2014.03.030
10. T. Britton, P. D. O'Neill, Bayesian inference for stochastic epidemics in populations with random social structure. *Scand. J. Stat.* 29, 375–390 (2002). doi: 10.1111/1467-9469.00296
11. J. Lourenço et al., Epidemiological and ecological determinants of Zika virus transmission in an urban setting. *eLife* 6, e29820 (2017). doi: 10.7554/eLife.29820; pmid: 28887877
12. N. R. Faria et al., NHLBI Recipient Epidemiology and Donor Evaluation Study-III (REDS-III), Genomic and epidemiological characterisation of a dengue virus outbreak among blood donors in Brazil. *Sci. Rep.* 7, 15216 (2017). doi: 10.1038/s41598-017-15152-8; pmid: 29123142
13. B. Shulgin, L. Stone, Z. Agur, Pulse vaccination strategy in the SIR epidemic model. *Bull. Math. Biol.* 60, 1123–1148 (1998). doi: 10.1016/S0092-8240(98)90005-2; pmid: 9866452
14. O. N. Bjørnstad, B. F. Finkenstadt, B. T. Grenfell, Dynamics of measles epidemics: Estimating scaling of transmission rates using a time series SIR model. *Ecol. Monogr.* 72, 169–184 (2002). doi: 10.1890/0012-9615(2002)072[0169:DOMEES]2.0.CO;2
15. L. Hufnagel, D. Brockmann, T. Geisel, Forecast and control of epidemics in a globalized world. *Proc. Natl. Acad. Sci. U.S.A.* 101, 15124–15129 (2004). doi: 10.1073/pnas.0308344101; pmid: 15477600
16. A. Pandey et al., Strategies for containing Ebola in West Africa. *Science* 346, 991–995 (2014). doi: 10.1126/science.1260612; pmid: 25414312
17. R. Li et al., Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science* 368, 489–493 (2020). doi: 10.1126/science.abb3221; pmid: 32179701
18. A. J. Kucharski et al., Early dynamics of transmission and control of COVID-19: A mathematical modelling study. *Lancet Infect. Dis.* 20, 553–558 (2020). doi: 10.1016/S1473-3099(20)30144-4; pmid: 32171059
19. J. Lourenço, R. Paton, M. Ghafari, M. Kraemer, C. Thompson, P. Simmonds, P. Klenerman, S. Gupta, Fundamental principles of epidemic spread highlight the immediate need for large-scale serological surveys to assess the stage of the SARS-CoV-2 epidemic. medRxiv 2020.03.24.20042291 [Preprint]. 26 March 2020; https://doi.org/10.1101/2020.03.24.20042291.
20. B. F. Maier, D. Brockmann, Effective containment explains subexponential growth in recent confirmed COVID-19 cases in China. *Science* 368, 742–746 (2020). doi: 10.1126/science.abb4557; pmid: 32269067
21. P. Bittihn, R. Golestanian, Containment strategy for an epidemic based on fluctuations in the SIR model. arXiv:2003.08784 [q-bio.PE] (19 March 2020).
22. R. M. Anderson, H. Heesterbeek, D. Klinkenberg, T. D. Hollingsworth, How will country-based mitigation measures influence the course of the COVID-19 epidemic? *Lancet* 395, 931–934 (2020). doi: 10.1016/S0140-6736(20)30567-5; pmid: 32164834
23. J. R. Fauver et al., Coast-to-coast spread of SARS-CoV-2 in the United States revealed by genomic epidemiology. medRxiv 2020.03.25.20043828 [Preprint]. 26 March 2020; https://doi.org/10.1101/2020.03.25.20043828.
24. A. Arenas, et al., A mathematical model for the spatiotemporal epidemic spreading of COVID19. medRxiv 2020.03.21.20040022 [Preprint]. 23 March 2020; https://doi.org/2020.03.21.20040022.
25. O. Mitjà et al., 62 signatories, Experts' request to the Spanish Government: Move Spain towards complete lockdown. *Lancet* 395, 1193–1194 (2020). doi: 10.1016/S0140-6736(20)30753-4; pmid: 32224297
26. S. L. Chang, N. Harding, C. Zachreson, O. M. Cliff, M. Prokopenko, Modelling transmission and control of the COVID-19 pandemic in Australia. arXiv:2003.10218 [q-bio.PE] (23 March 2020).
27. W. Bock et al., Mitigation and herd immunity strategy for COVID-19 is likely to fail. medRxiv 2020.03.25.20043109 [Preprint]. 5 May 2020; https://doi.org/10.1101/2020.03.25.20043109.
28. C. Gros, R. Valenti, L. Schneider, K. Valenti, D. Gros, Containment efficiency and control strategies for the Corona pandemic costs. arXiv:2004.00493 [physics.soc-ph] (1 April 2020).
29. V. Zlatić, I. Barjašić, A. Kadović, H. Štefančić, A. Gabrielli, Bi-stability of SUDR+K model of epidemics and test kits applied to COVID-19. arXiv:2003.08479 [q-bio.PE] (18 March 2020).
30. J. Dehning et al., Inferring COVID-19 spreading rates and potential change points for case number forecasts. arXiv:2004.01105 [q-bio.PE] (2 April 2020).
31. J. Dehning et al., Analysis code for: Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions, Zenodo (2020); https://doi.org/10.5281/zenodo.3823382.
32. S. A. Lauer et al., The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: Estimation and application. *Ann. Intern. Med.* 172, 577–582 (2020). doi: 10.7326/M20-0504; pmid: 32150748
33. J. Zhang et al., Age profile of susceptibility, mixing, and social distancing shape the dynamics of the novel coronavirus disease 2019 outbreak in China. medRxiv 2020.03.19.20039107 [Preprint]. 20 March 2020; https://doi.org/10.1101/2020.03.19.20039107.
34. Y. Liu, A. A. Gayle, A. Wilder-Smith, J. Rocklöv, The reproductive number of COVID-19 is higher compared to SARS coronavirus.

J. Travel Med. **27**, taaa021 (2020). doi: 10.1093/jtm/taaa021; pmid: 32052846

35. M. A. Muñoz, Colloquium: Criticality and dynamical scaling in living systems. *Rev. Mod. Phys.* **90**, 031001 (2018). doi: 10.1103/RevModPhys.90.031001

36. A. Vehtari, A. Gelman, J. Gabry, Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Stat. Comput.* **27**, 1413–1432 (2017). doi: 10.1007/s11222-016-9696-4

37. L. Peng, W. Yang, D. Zhang, C. Zhuge, L. Hong, Epidemic analysis of COVID-19 in China by dynamical modeling. arXiv:2002.06563 [q-bio.PE] (16 February 2020).

38. J. Wilting, V. Priesemann, Inferring collective dynamical states from widely unobserved systems. *Nat. Commun.* **9**, 2325 (2018). doi: 10.1038/s41467-018-04725-4; pmid: 29899335

39. Y.-C. Chen, P.-E. Lu, C.-S. Chang, T.-H. Liu, A Time-dependent SIR model for COVID-19 with undetectable infected persons. arXiv:2003.00122 [q-bio.PE] (28 February 2020).

40. E. Dong, H. Du, L. Gardner, An interactive web-based dashboard to track COVID-19 in real time. *Lancet Infect. Dis.* **20**, 533–534 (2020). doi: 10.1016/S1473-3099(20)30120-1; pmid: 32087114

41. M. Chinazzi *et al.*, The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science* **368**, 395–400 (2020). doi: 10.1126/science.aba9757; pmid: 32144116

42. J. Salvatier, T. V. Wiecki, C. Fonnesbeck, Probabilistic programming in Python using PyMC3. *PeerJ Comput. Sci.* **2**, e55 (2016). doi: 10.7717/peerj-cs.55

43. M. D. Hoffman, A. Gelman, The No-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**, 1593–1623 (2014).

44. A. Kucukelbir, D. Tran, R. Ranganath, A. Gelman, D. M. Blei, Automatic differentiation variational inference. *J. Mach. Learn. Res.* **18**, 1–45 (2017).

45. K. L. Lange, R. J. A. Little, J. M. G. Taylor, Robust statistical modeling using the t Distribution. *J. Am. Stat. Assoc.* **84**, 881–896 (1989). doi: 10.2307/2290063

46. S. di Santo, P. Villegas, R. Burioni, M. A. Muñoz, Simple unified view of branching process statistics: Random walks in balanced logarithmic potentials. *Phys. Rev. E* **95**, 032115 (2017). doi: 10.1103/PhysRevE.95.032115; pmid: 28415350

47. A. Gelman, Prior distributions for variance parameters in hierarchical models. *Bayesian Anal.* **1**, 515–534 (2006). doi: 10.1214/06-BA117A

48. Johns Hopkins University, COVID-19 Data Repository by the Center for Systems Science and Engineering (CSSE) (2020); https://raw.githubusercontent.com/CSSEGISandData/COVID-19/bda67e3db0e8dca4540297633d431a8021c035c8/csse_covid_19_data/csse_covid_19_time_series/time_series_covid19_confirmed_global.csv.

ACKNOWLEDGMENTS

SUPPLEMENTARY MATERIALS

Science

Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions

Jonas Dehning, Johannes Zierenberg, F. Paul Spitzner, Michael Wibral, Joao Pinheiro Neto, Michael Wilczek and Viola Priesemann

Keeping the lid on infection spread

From February to April 2020, many countries introduced variations on social distancing measures to slow the ravages of severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2). Publicly available data show that Germany has been particularly successful in minimizing death rates. Dehning *et al.* quantified three governmental interventions introduced to control the outbreak. The authors predicted that the third governmental intervention—a strict contact ban since 22 March—switched incidence from growth to decay. They emphasize that relaxation of controls must be done carefully, not only because there is a 2-week lag between a measure being enacted and the effect on case reports but also because the three measures used in Germany only just kept virus spread below the growth threshold.

Science, this issue p. eabb9789

ARTICLE TOOLS	http://science.sciencemag.org/content/369/6500/eabb9789
SUPPLEMENTARY MATERIALS	http://science.sciencemag.org/content/suppl/2020/05/14/science.abb9789.DC1
RELATED CONTENT	http://stm.sciencemag.org/content/scitransmed/12/541/eabb5883.full http://stm.sciencemag.org/content/scitransmed/11/499/eaat0360.full http://stm.sciencemag.org/content/scitransmed/12/534/eabb1469.full http://stm.sciencemag.org/content/scitransmed/9/396/eaal3653.full
REFERENCES	This article cites 31 articles, 6 of which you can access for free http://science.sciencemag.org/content/369/6500/eabb9789#BIBL
PERMISSIONS	http://www.sciencemag.org/help/reprints-and-permissions

Use of this article is subject to the Terms of Service