Real-time monitoring and short-term forecasting of the COVID-19 pandemic

CSM / CMMID Seminar

Sebastian Funk, 8 October 2020
@sbfnk

Presenting the work of

Sam Abbott
Nikos Bosse
Joel Hellewell
Sophie Meakin
James Munday
Kath Sherratt
Robin Thompson

Table of contents

1. Reproduction numbers for Covid-19
2. Short-term forecasts for Covid-19

Focus on statistical challenges.

Reproduction numbers for Covid-19

See also: Gostic et al. (2020), “Practical considerations for measuring the effective reproductive number, Rt”, https://doi.org/10.1101/2020.06.18.20134858

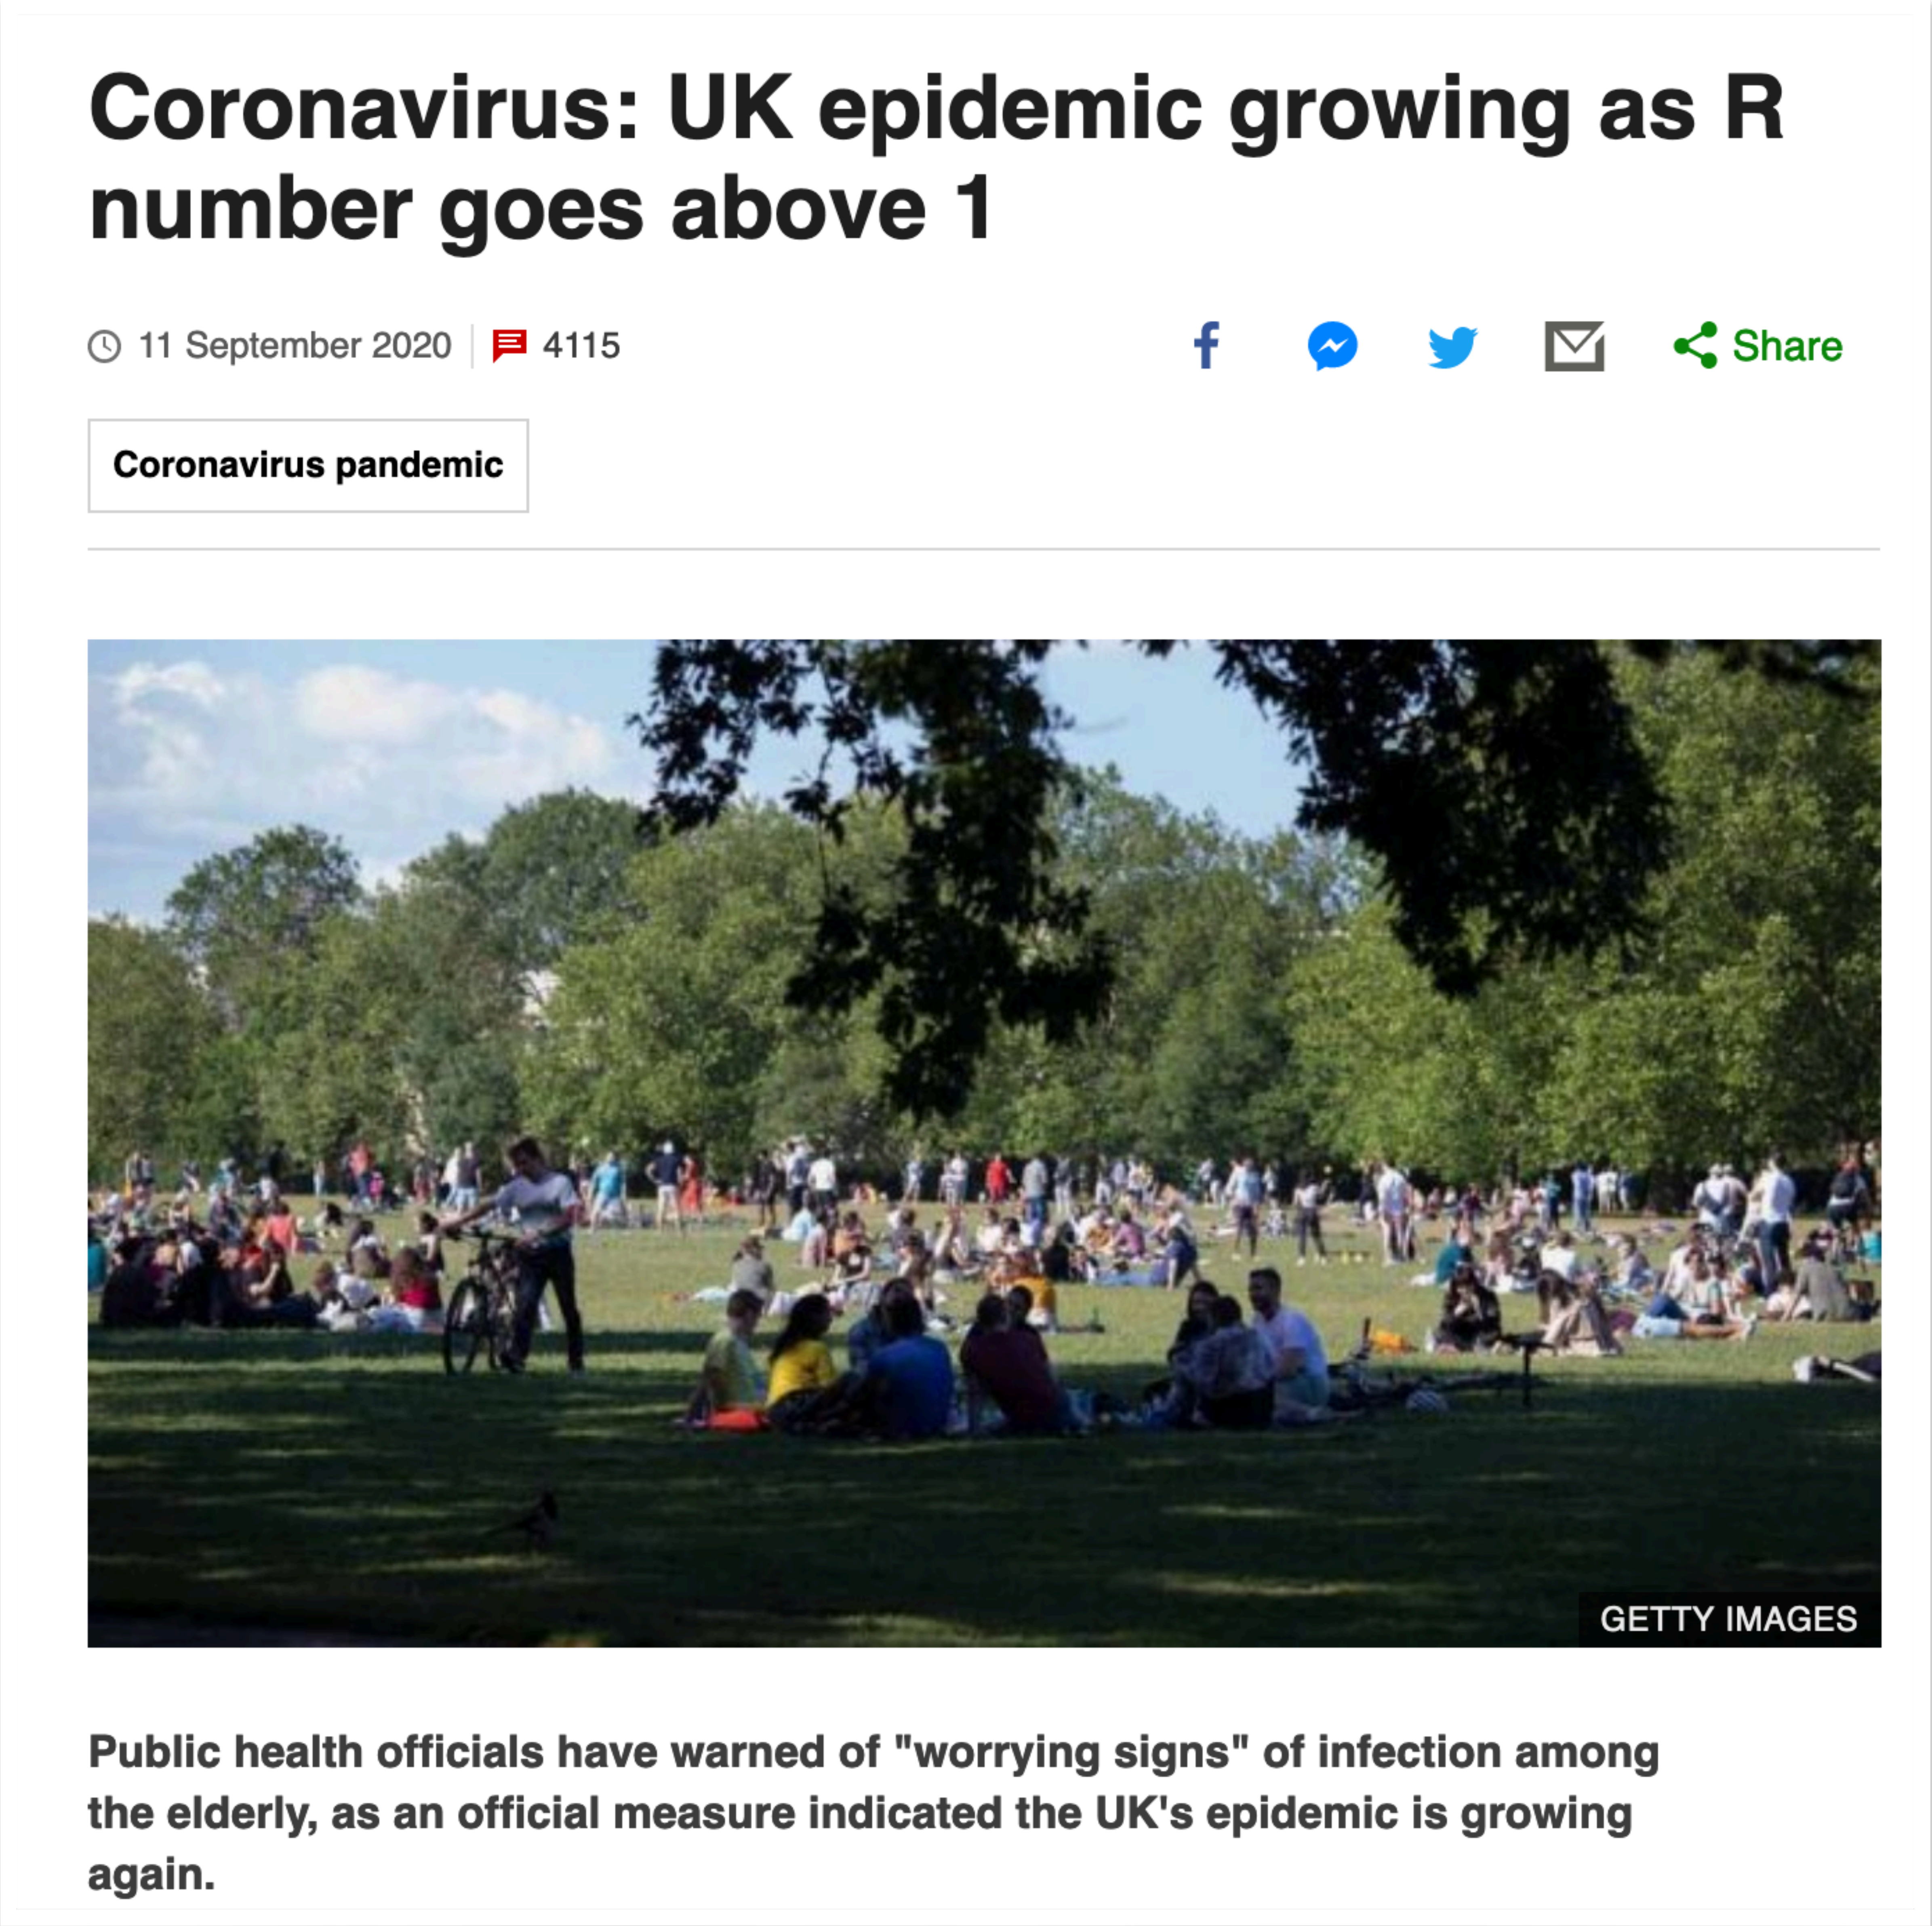

Coronavirus: UK epidemic growing as R number goes above 1

11 September 2020 | 4115

Coronavirus pandemic

GETTY IMAGES

Public health officials have warned of "worrying signs" of infection among the elderly, as an official measure indicated the UK's epidemic is growing again.

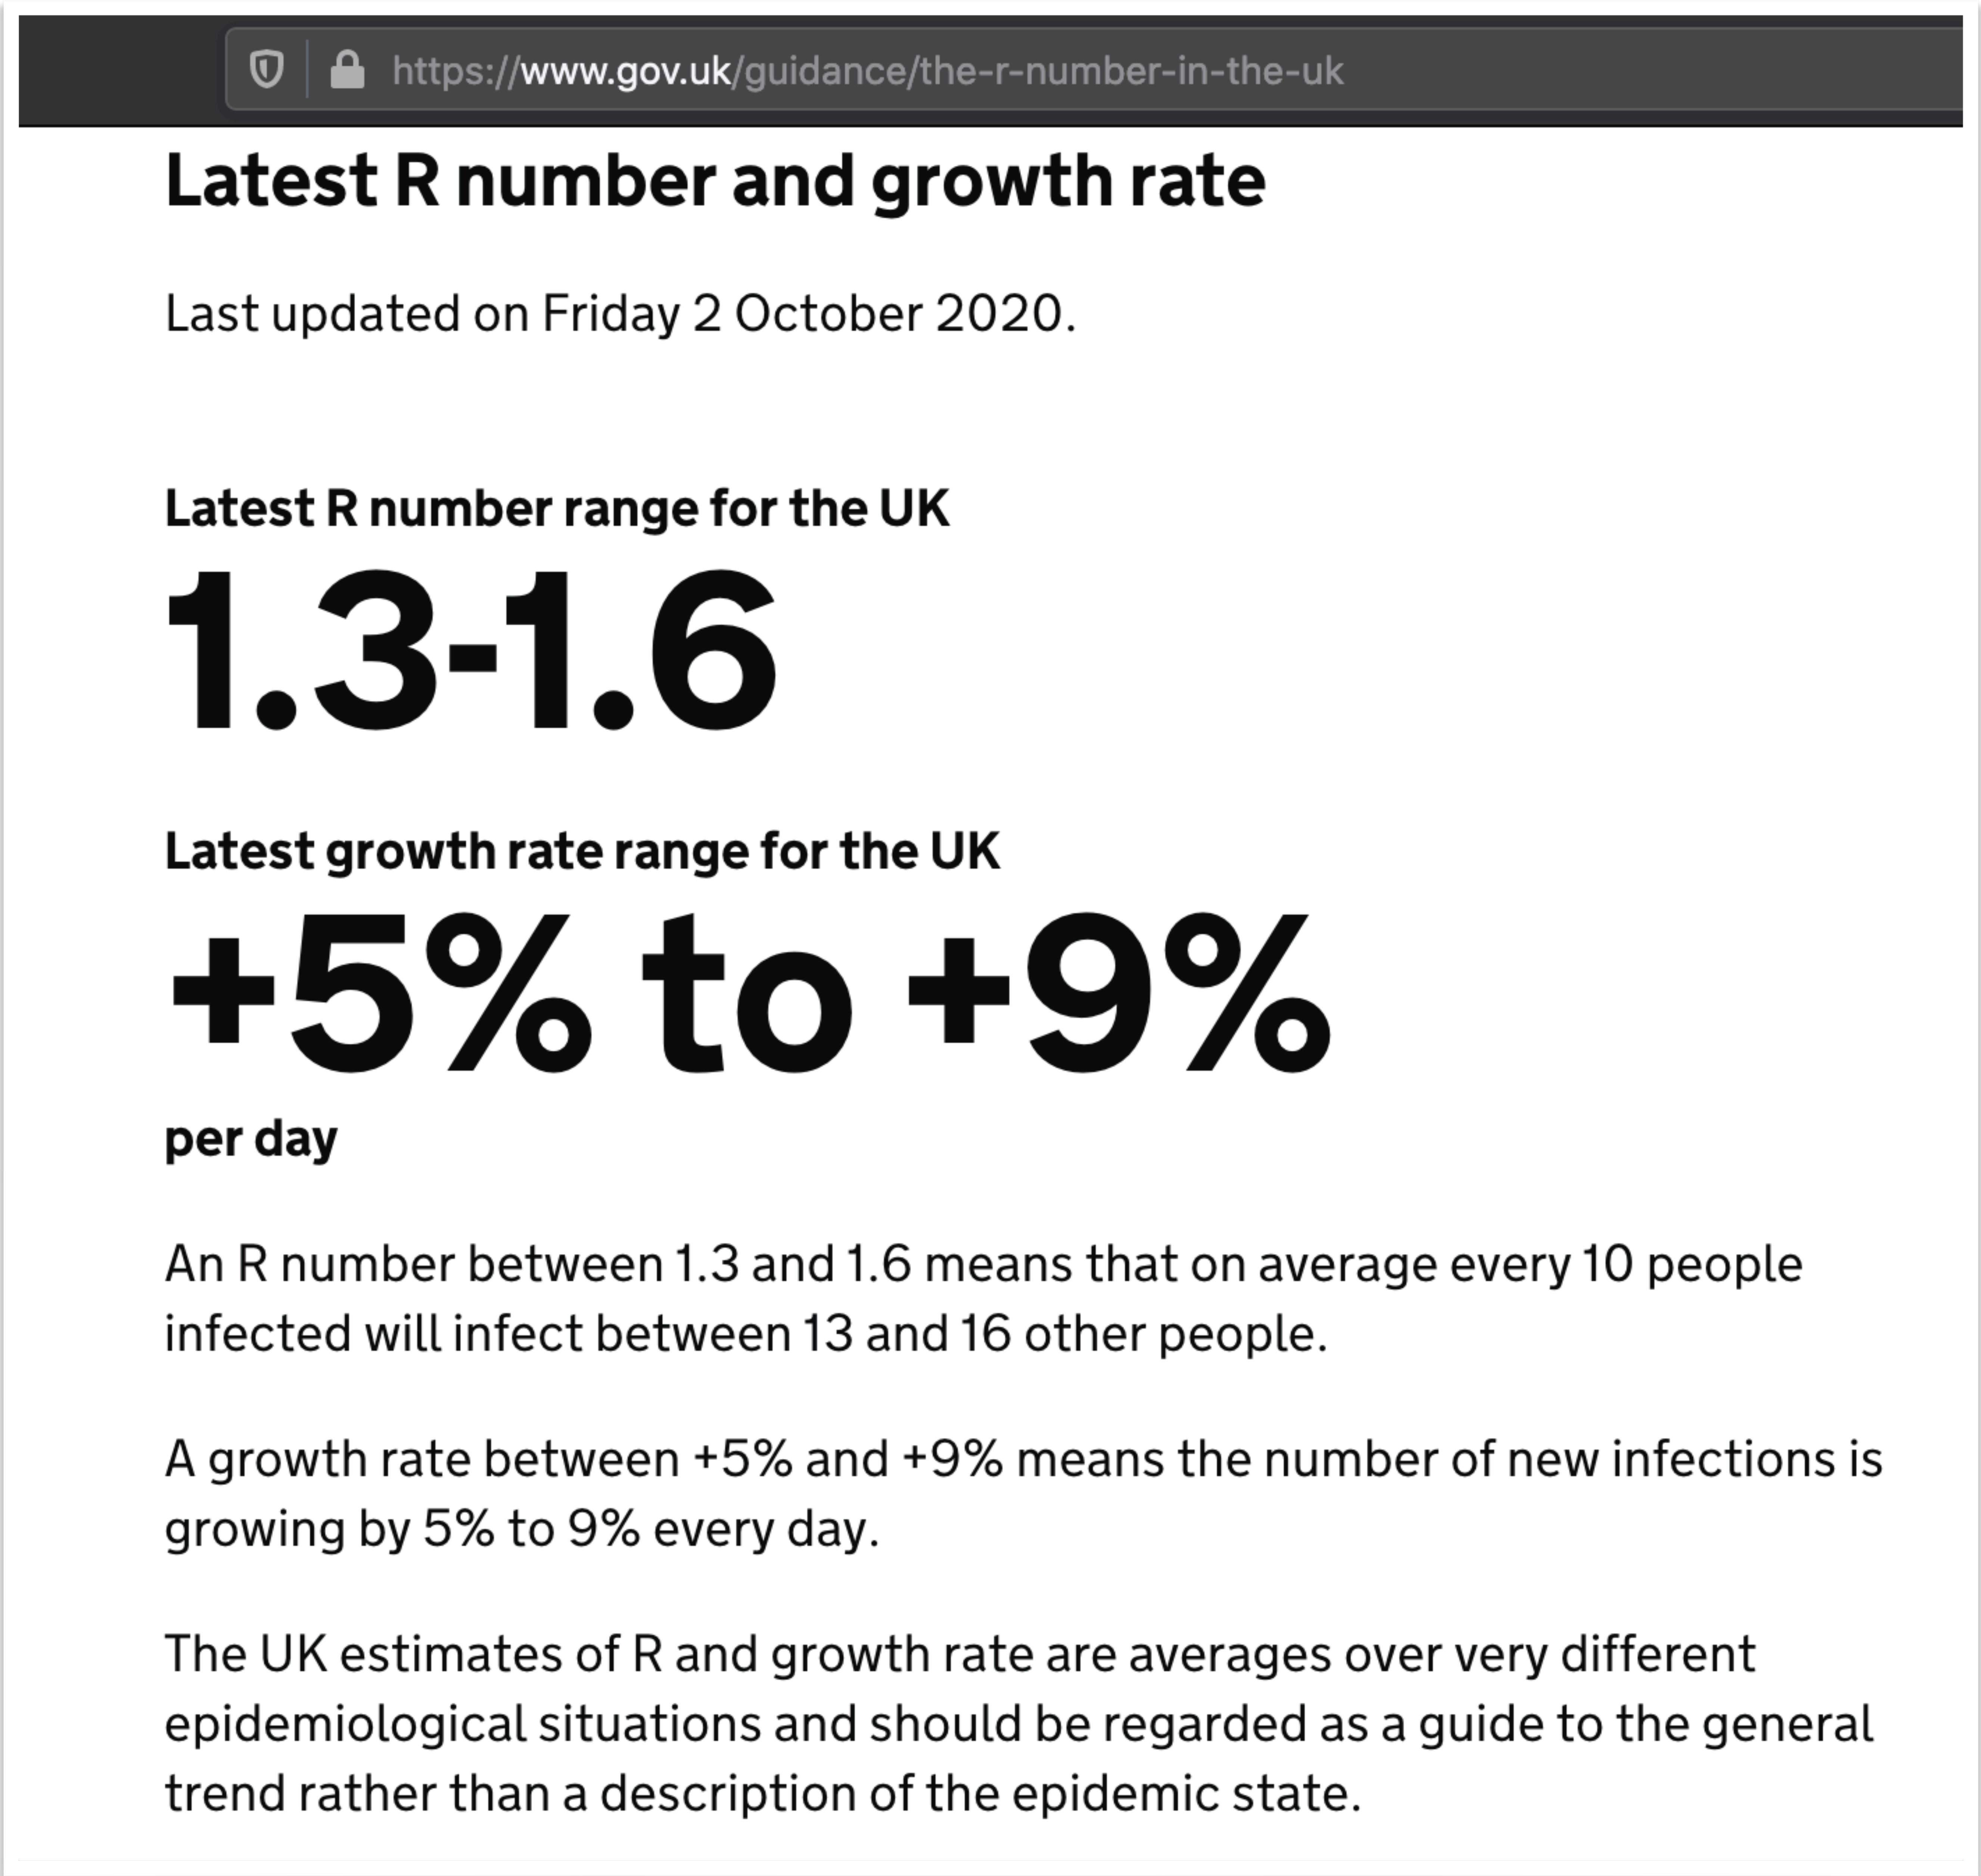

Latest R number and growth rate

Last updated on Friday 2 October 2020.

Latest R number range for the UK

1.3-1.6

Latest growth rate range for the UK

+5% to +9%

per day

An R number between 1.3 and 1.6 means that on average every 10 people infected will infect between 13 and 16 other people.

A growth rate between +5% and +9% means the number of new infections is growing by 5% to 9% every day.

The UK estimates of R and growth rate are averages over very different epidemiological situations and should be regarded as a guide to the general trend rather than a description of the epidemic state.

(Effective) reproduction number

R_t: The average number of secondary infections caused by a typical infectious individual

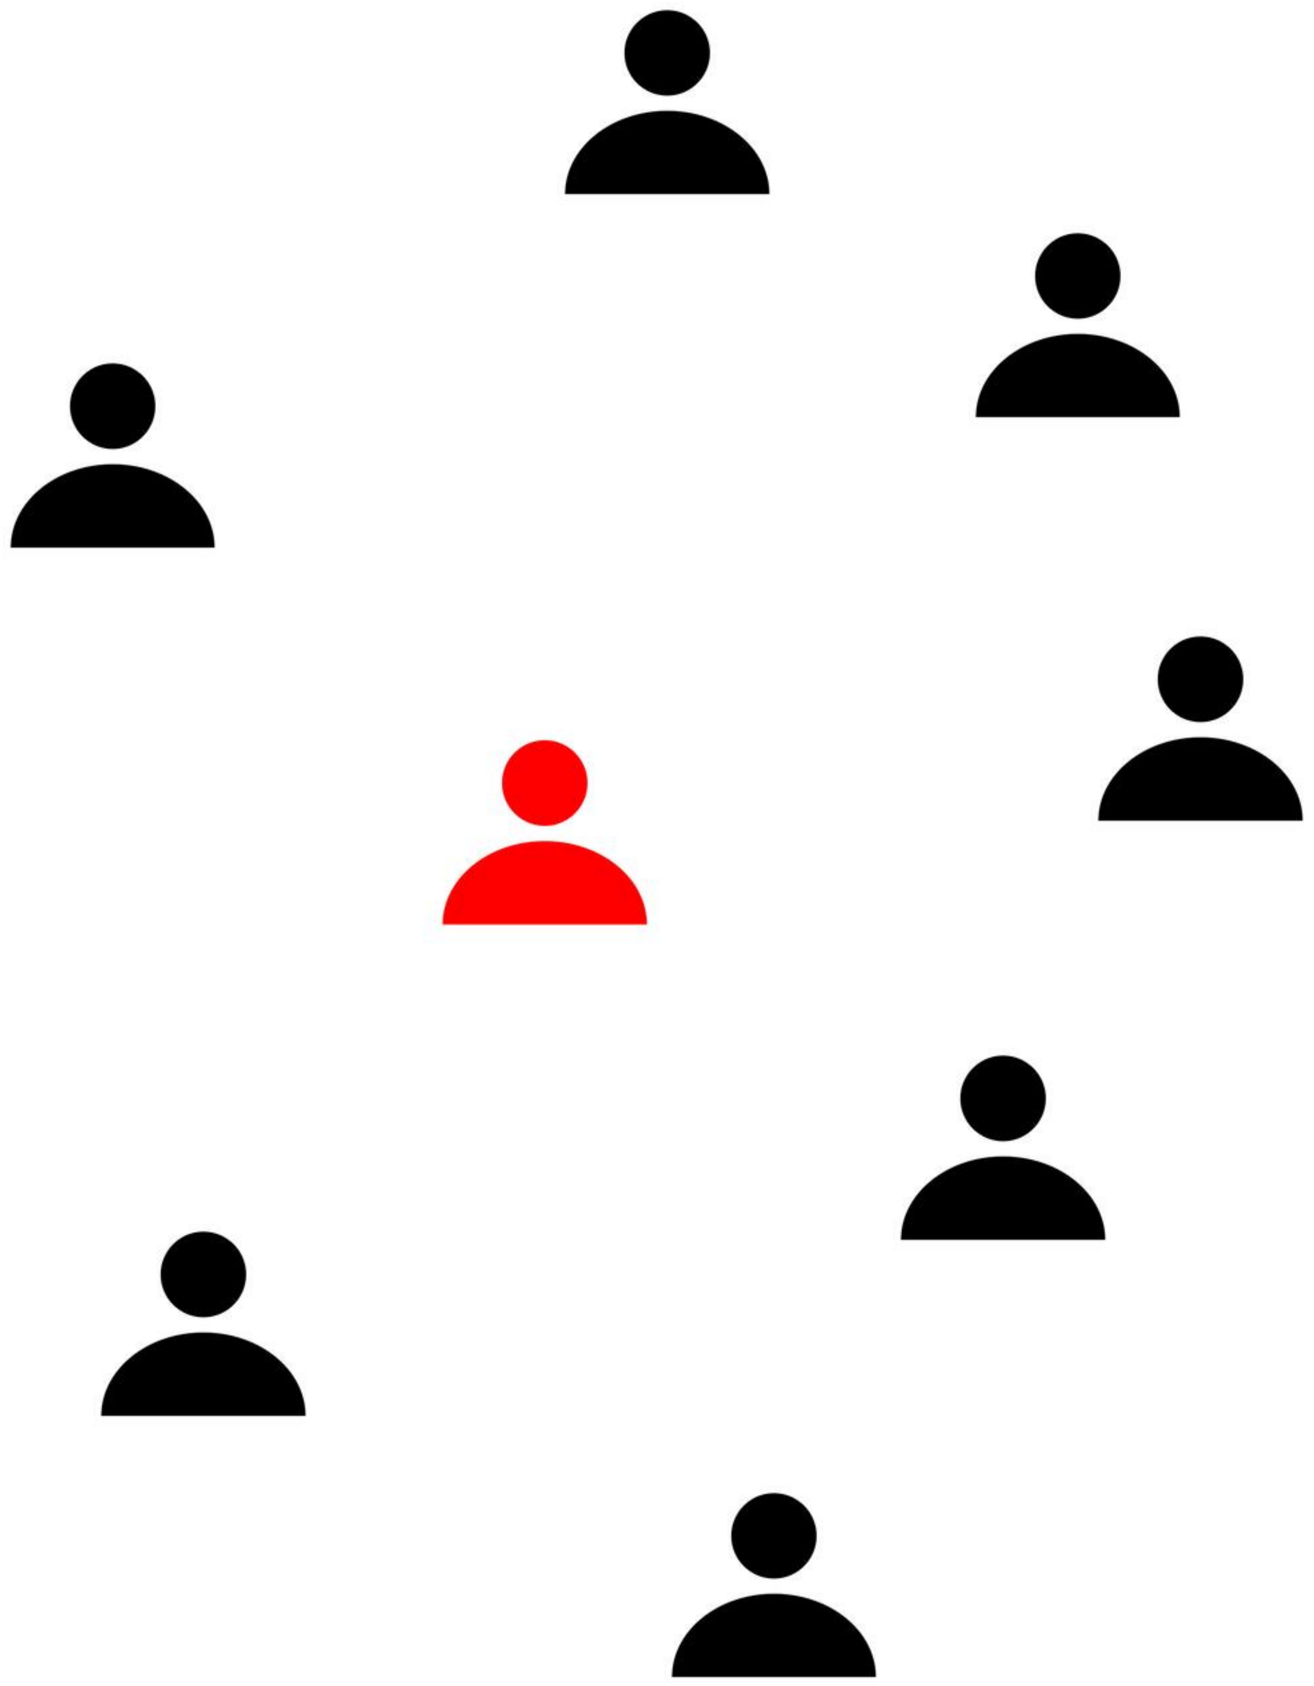

(Effective) reproduction number

R_t: The average number of secondary infections caused by a typical infectious individual

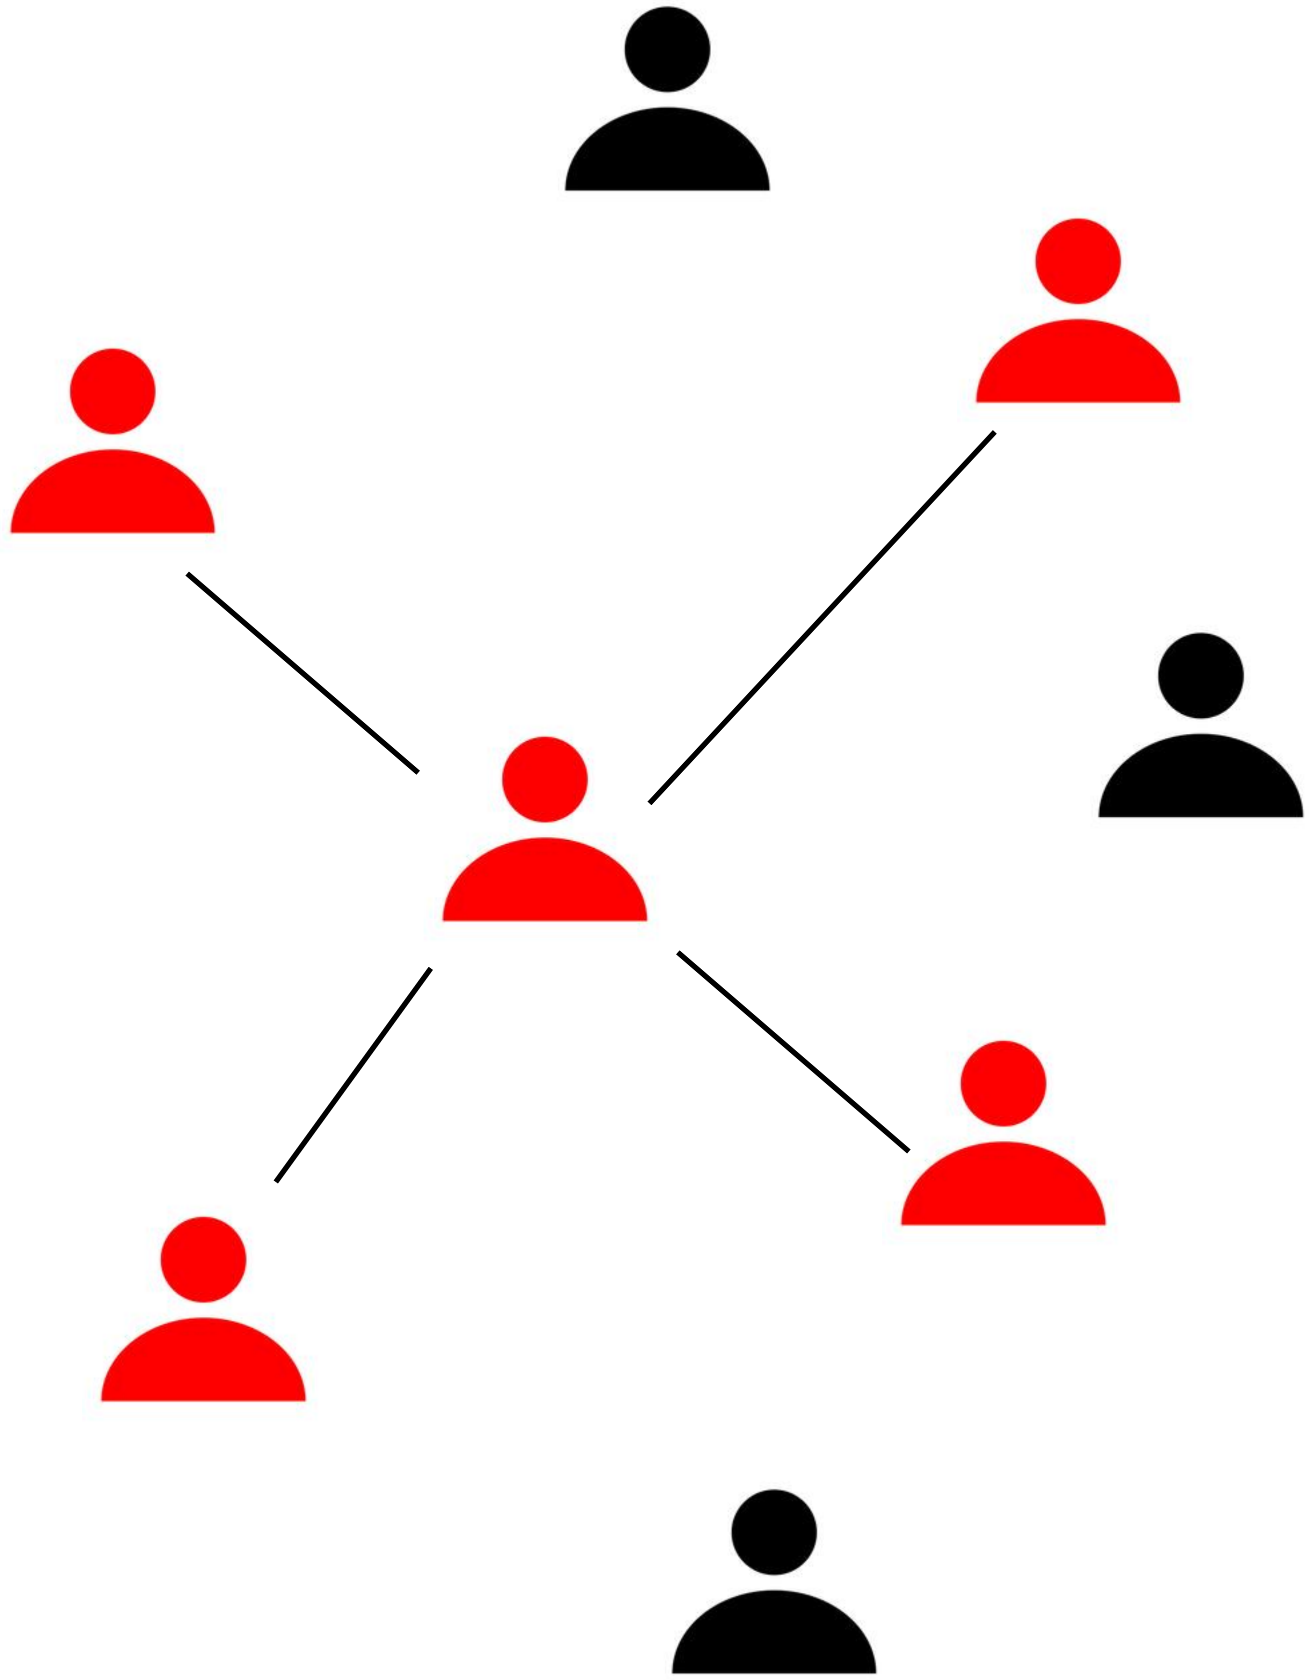

(Effective) reproduction number

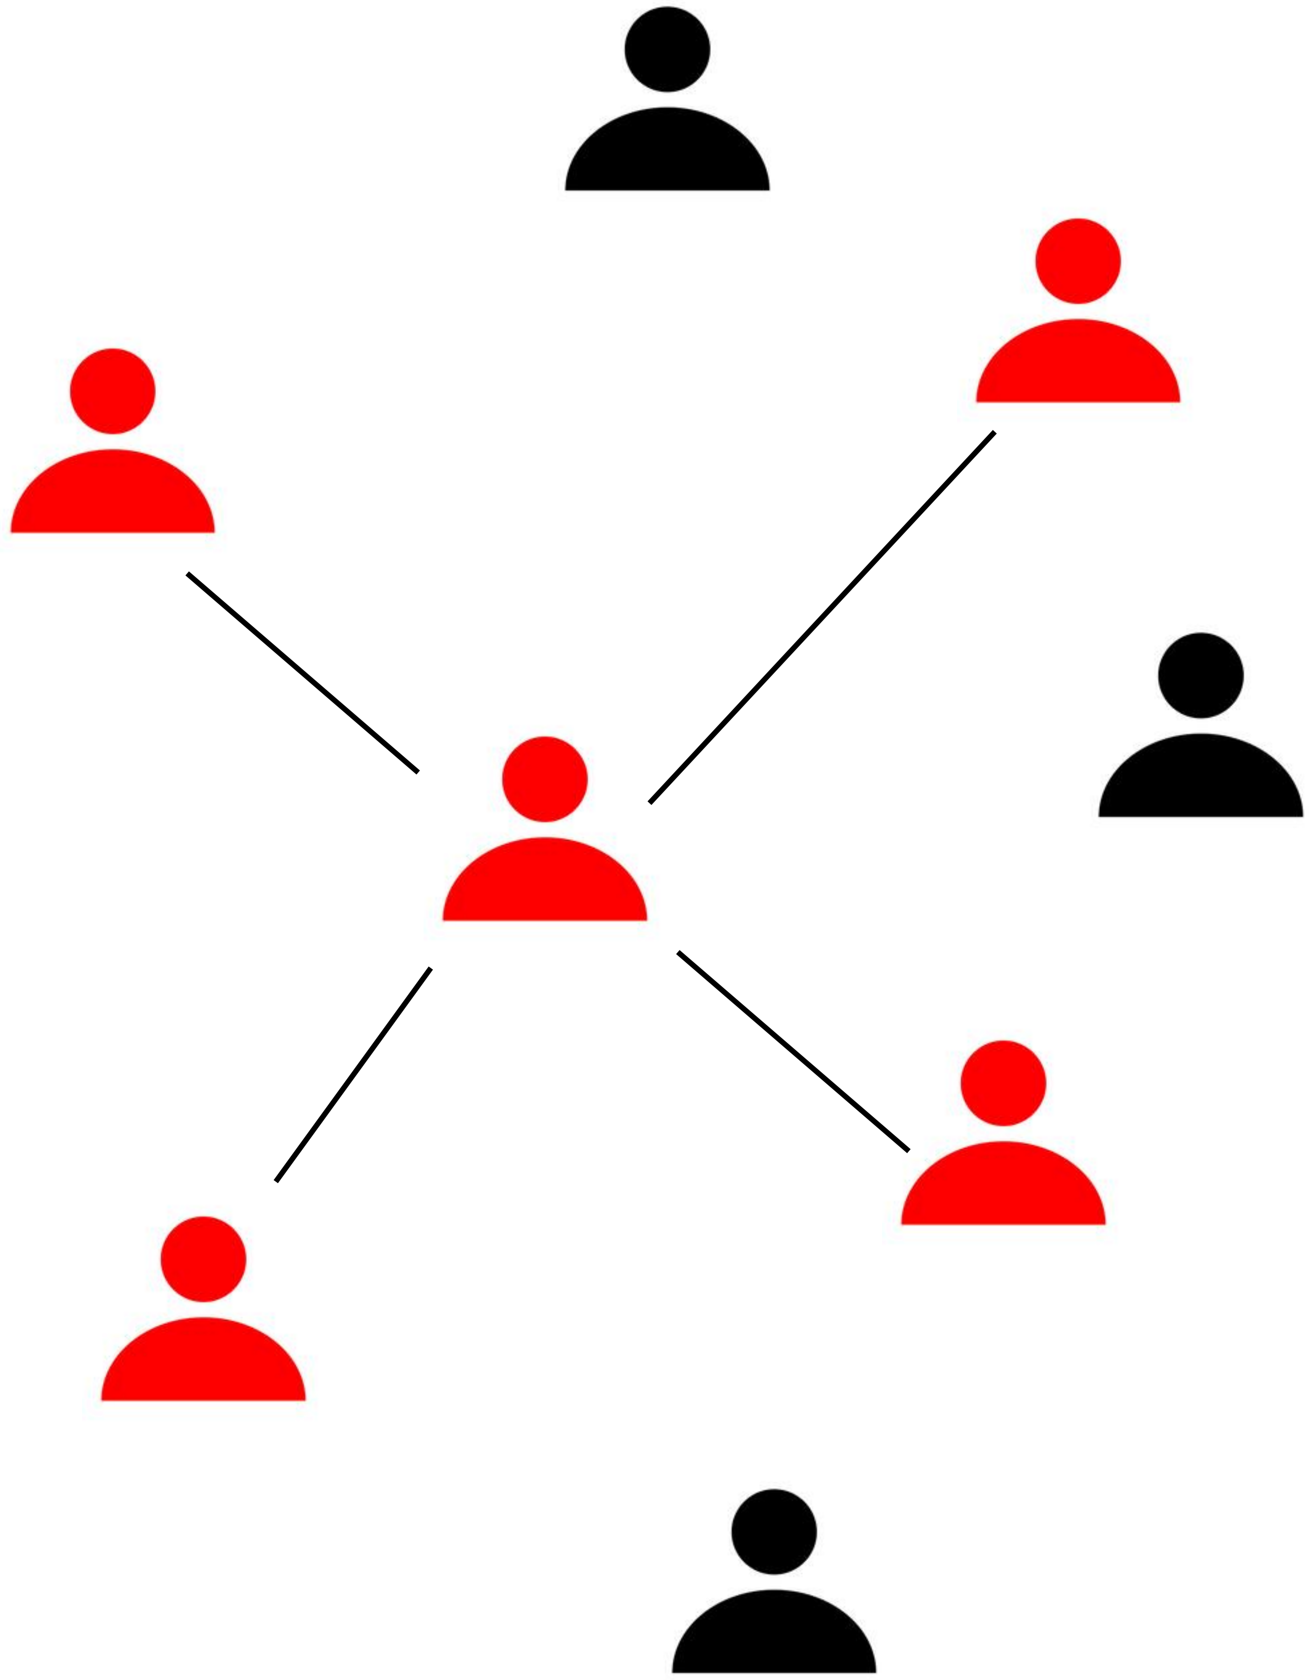

R_t: The average number of secondary infections caused by a typical infectious individual

$R_t > 1$ number of infections increasing
$R_t < 1$ number of infections decreasing

Tells us about
- (historic or current) transmission intensity
- strength of interventions required

Purpose of estimating R_t

Historical
e.g. impact of interventions

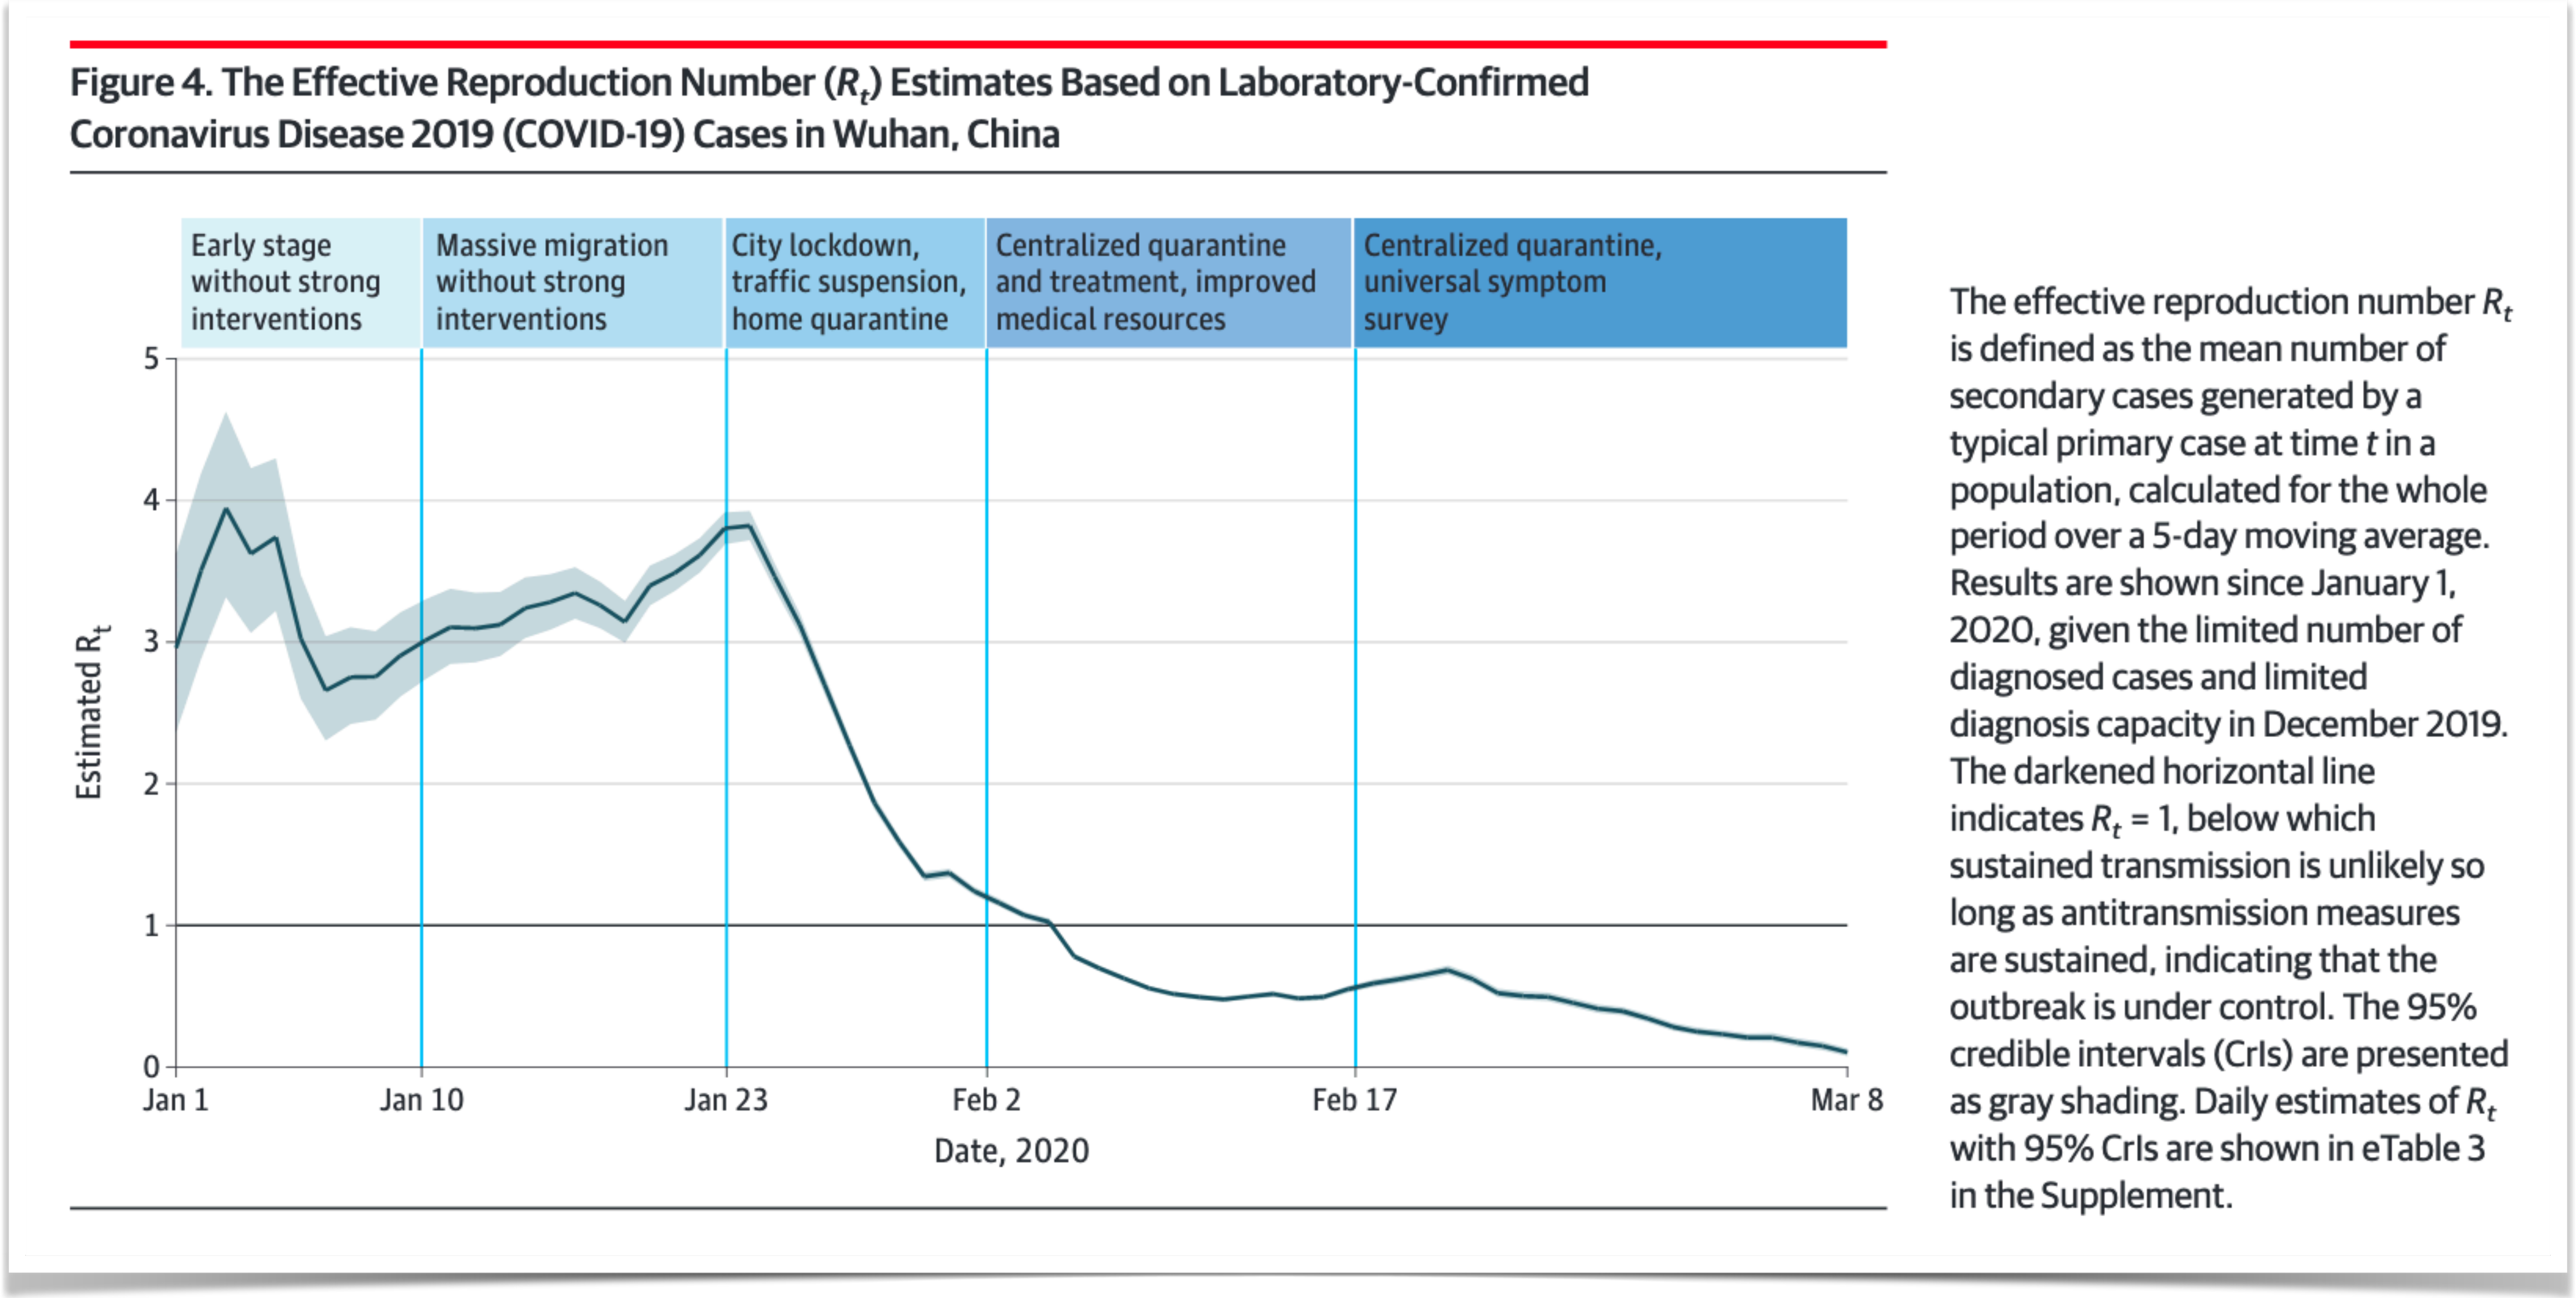

Figure 4. The Effective Reproduction Number (R_t) Estimates Based on Laboratory-Confirmed Coronavirus Disease 2019 (COVID-19) Cases in Wuhan, China

The effective reproduction number R_t is defined as the mean number of secondary cases generated by a typical primary case at time t in a population, calculated for the whole period over a 5-day moving average. Results are shown since January 1, 2020, given the limited number of diagnosed cases and limited diagnosis capacity in December 2019. The darkened horizontal line indicates $R_t = 1$, below which sustained transmission is unlikely so long as antitransmission measures are sustained, indicating that the outbreak is under control. The 95% credible intervals (CrIs) are presented as gray shading. Daily estimates of R_t with 95% CrIs are shown in eTable 3 in the Supplement.

Pan et al., JAMA, 2020

Purpose of estimating R_t

Historical
e.g. impact of interventions

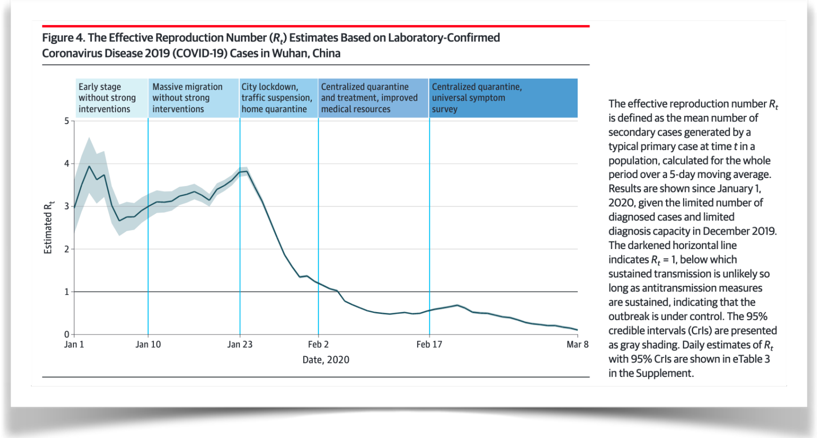

Real-time
for situational awareness

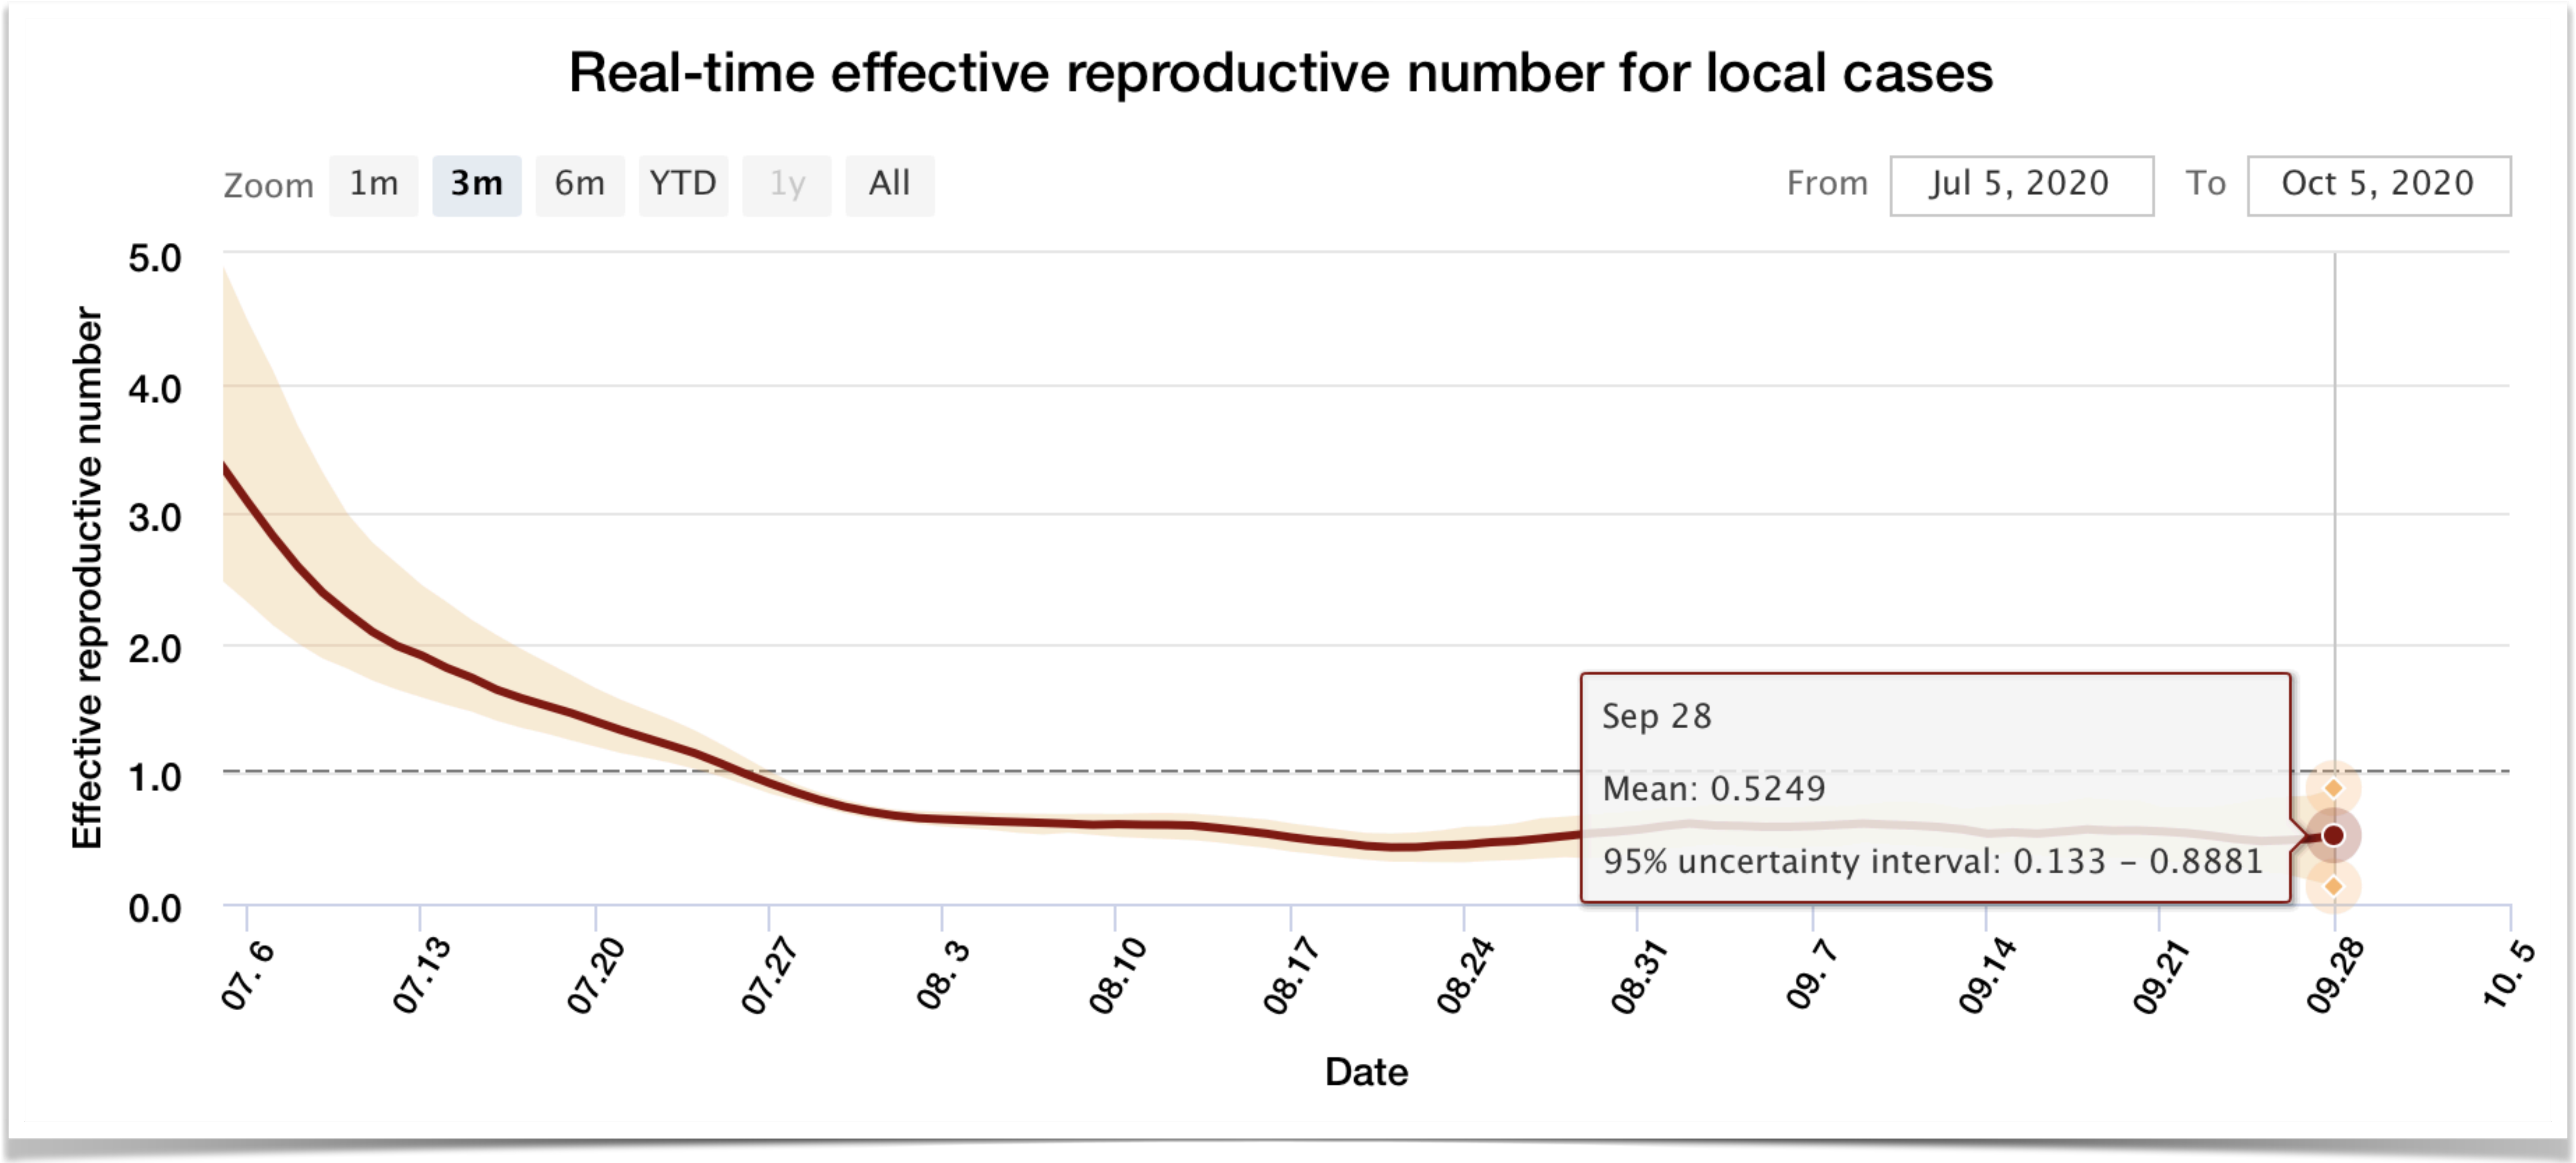

Sep 28
Mean: 0.5249
95% uncertainty interval: 0.133 – 0.8881

https://covid19.sph.hku.hk/

Purpose of estimating R_t

Historical
e.g. impact of interventions

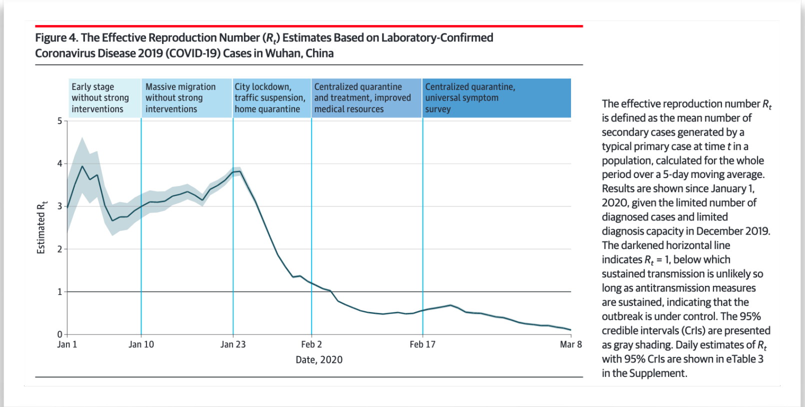

Figure 4. The Effective Reproduction Number (R_t) Estimates Based on Laboratory-Confirmed Coronavirus Disease 2019 (COVID-19) Cases in Wuhan, China

The effective reproduction number R_t is defined as the mean number of secondary cases generated by a typical primary case at time *t* in a population, calculated for the whole period over a 5-day moving average. Results are shown since January 1, 2020, given the limited number of diagnosed cases and limited diagnosis capacity in December 2019. The darkened horizontal line indicates R_t = 1, below which sustained transmission is unlikely so long as antitransmission measures are sustained, indicating that the outbreak is under control. The 95% credible intervals (CrIs) are presented as gray shading. Daily estimates of R_t with 95% CrIs are shown in eTable 3 in the Supplement.

Real-time
for situational awareness

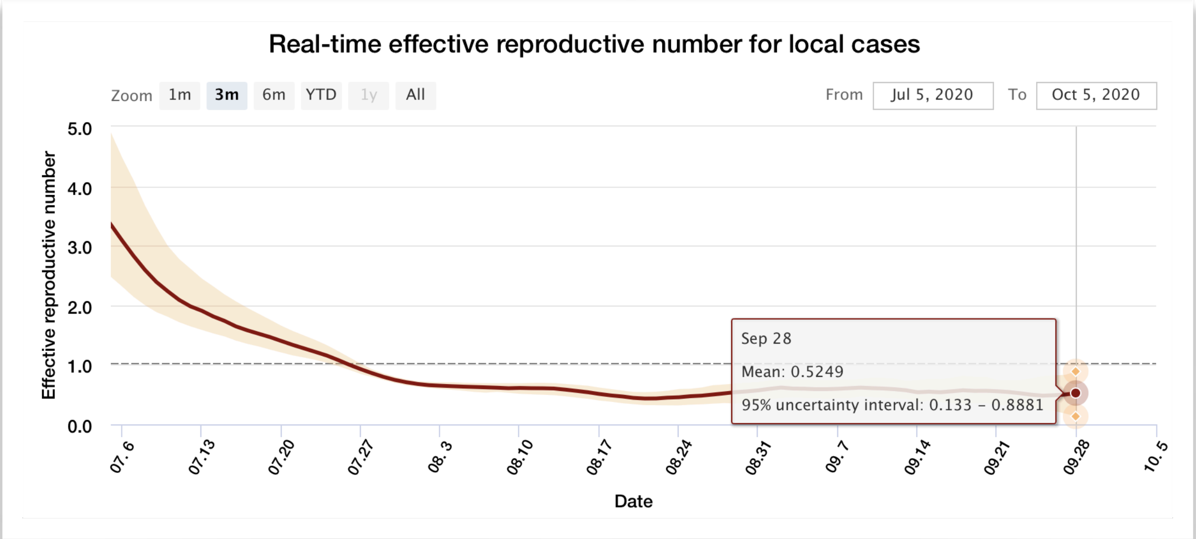

Real-time effective reproductive number for local cases

From Jul 5, 2020 To Oct 5, 2020

Sep 28
Mean: 0.5249
95% uncertainty interval: 0.133 – 0.8881

Purpose of estimating R_t

Historical
e.g. impact of interventions

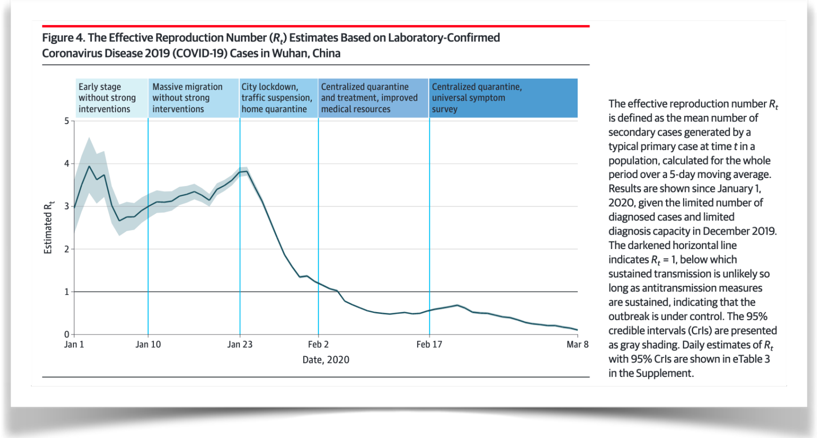

Figure 4. The Effective Reproduction Number (R_t) Estimates Based on Laboratory-Confirmed Coronavirus Disease 2019 (COVID-19) Cases in Wuhan, China

The effective reproduction number R_t is defined as the mean number of secondary cases generated by a typical primary case at time *t* in a population, calculated for the whole period over a 5-day moving average. Results are shown since January 1, 2020, given the limited number of diagnosed cases and limited diagnosis capacity in December 2019. The darkened horizontal line indicates $R_t = 1$, below which sustained transmission is unlikely so long as antitransmission measures are sustained, indicating that the outbreak is under control. The 95% credible intervals (CrIs) are presented as gray shading. Daily estimates of R_t with 95% CrIs are shown in eTable 3 in the Supplement.

Real-time
for situational awareness

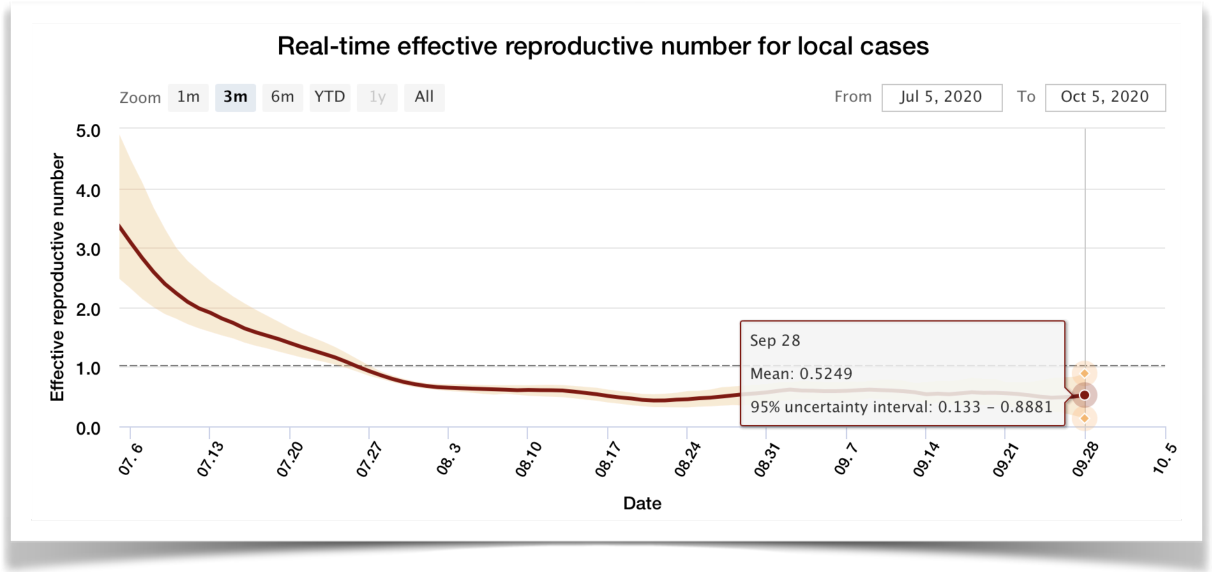

Real-time effective reproductive number for local cases

Sep 28
Mean: 0.5249
95% uncertainty interval: 0.133 – 0.8881

Statistical challenges affect the reliability of estimates differently depending on the given purpose.

How is can R_t be estimated?

- We generally don't observe infection events directly
- Cori et al. (2013); Fraser (2007)

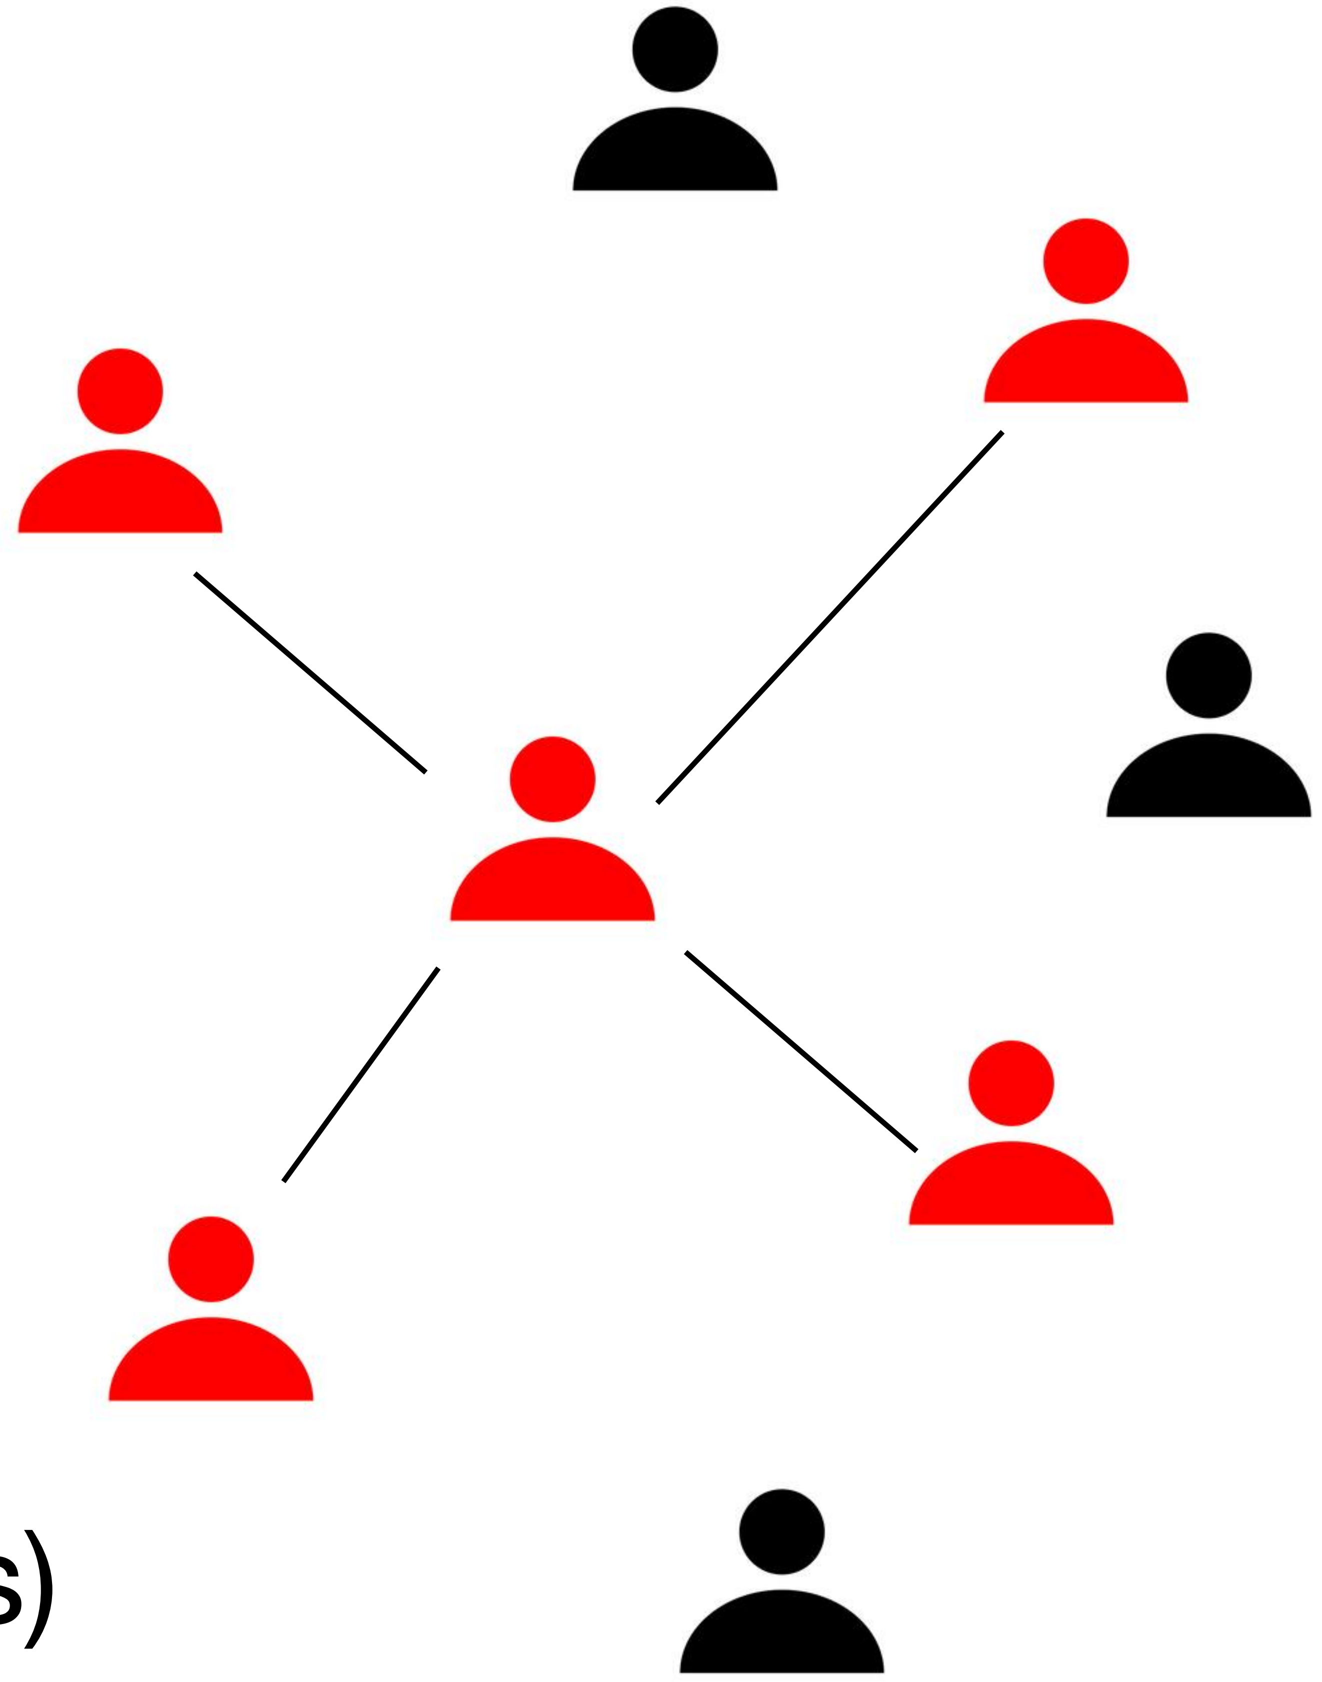

$$R_t = \frac{I_t}{\sum_\tau w_\tau I_{t-\tau}}$$

I_t Number of new infections at time *t* (e.g., days)

w_τ Probability of infecting someone at time τ after becoming infected ("generation interval")

Challenge #1: Generation interval w_j

w_j Probability of infecting someone at time *j* after becoming infected

Difficult to measure:

1) Observe viral load over time; but how does it relate to infectivity?

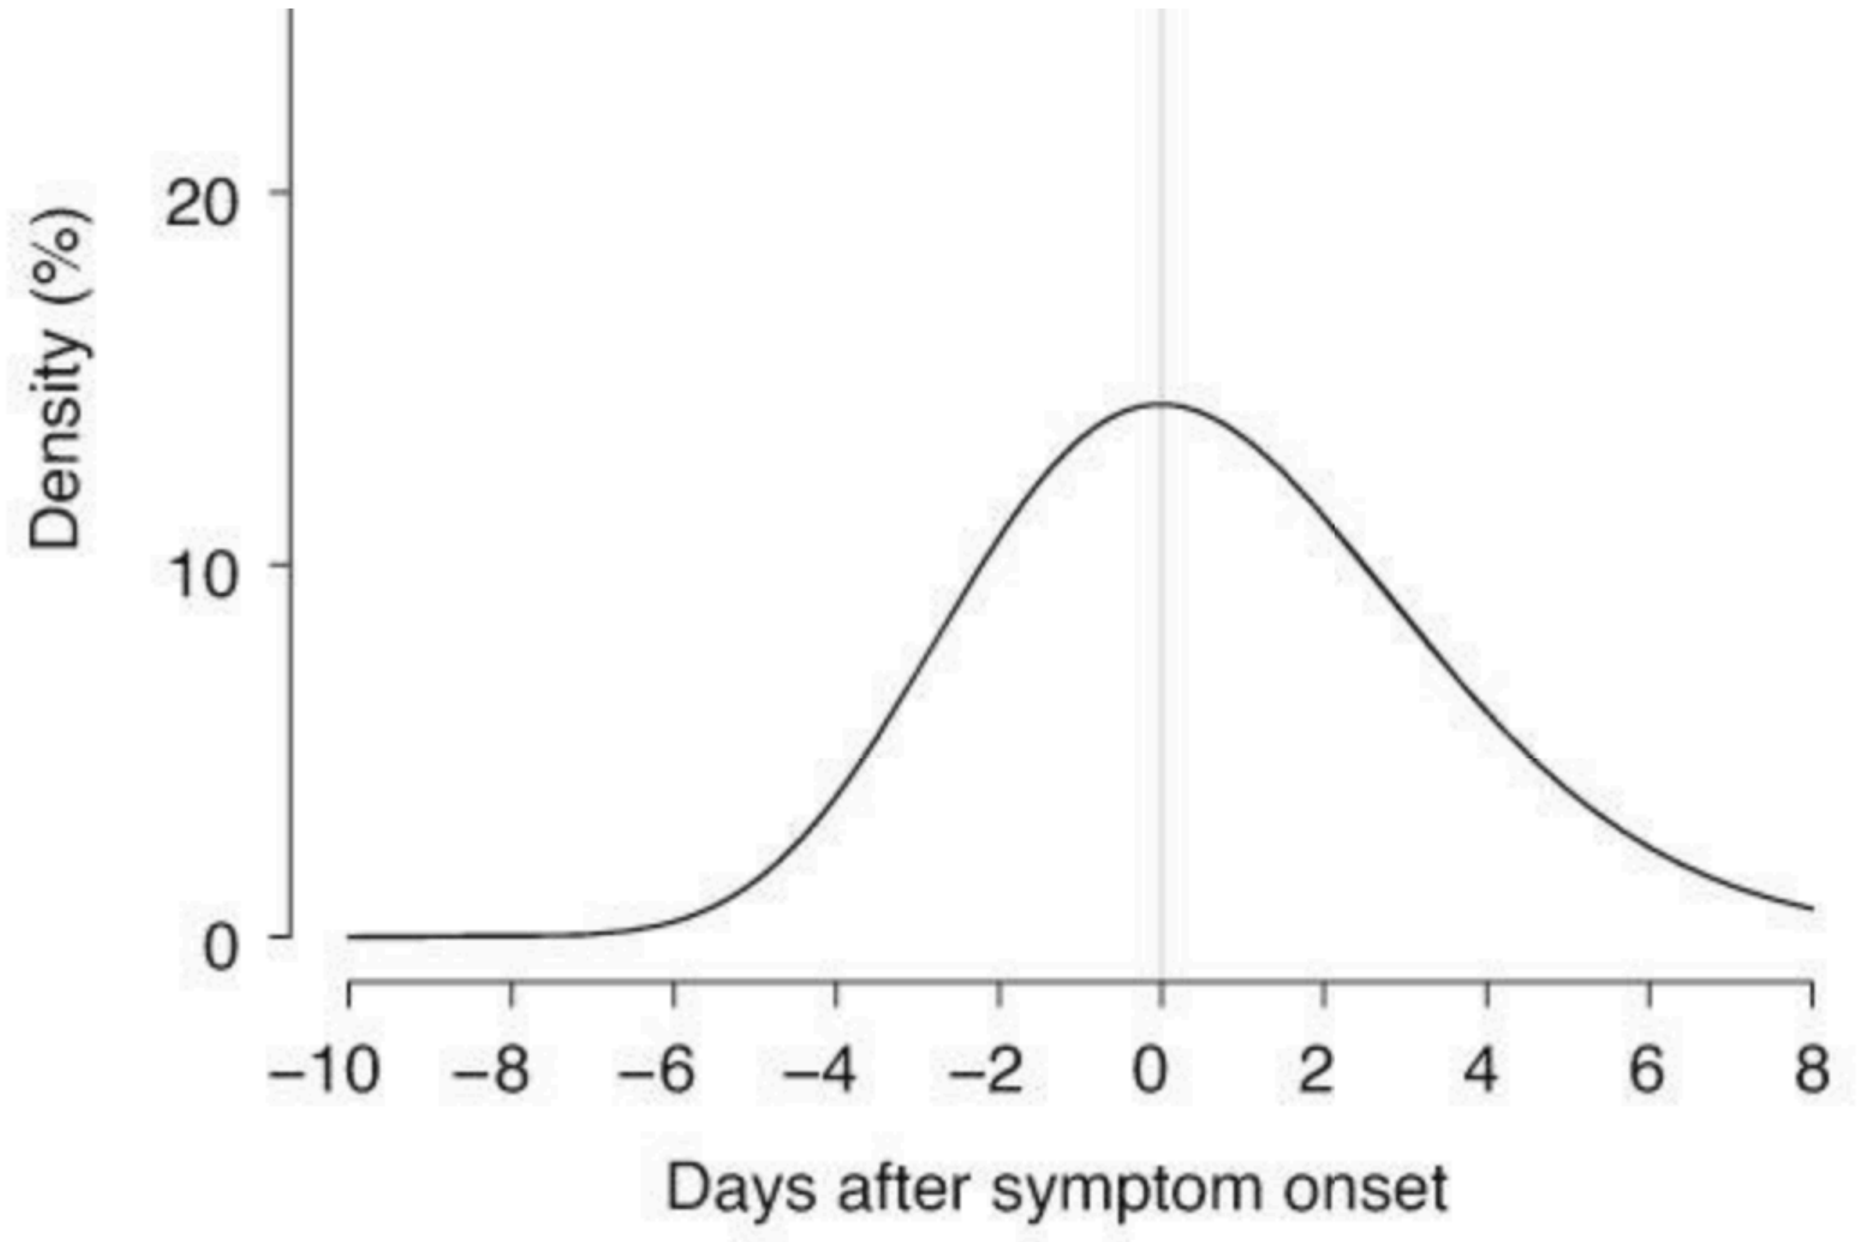

He et al., Nature Medicine, 2020
Wölfel et al., Nature, 2020

Challenge #1: Generation interval w_j

w_j Probability of infecting someone at time *j* after becoming infected

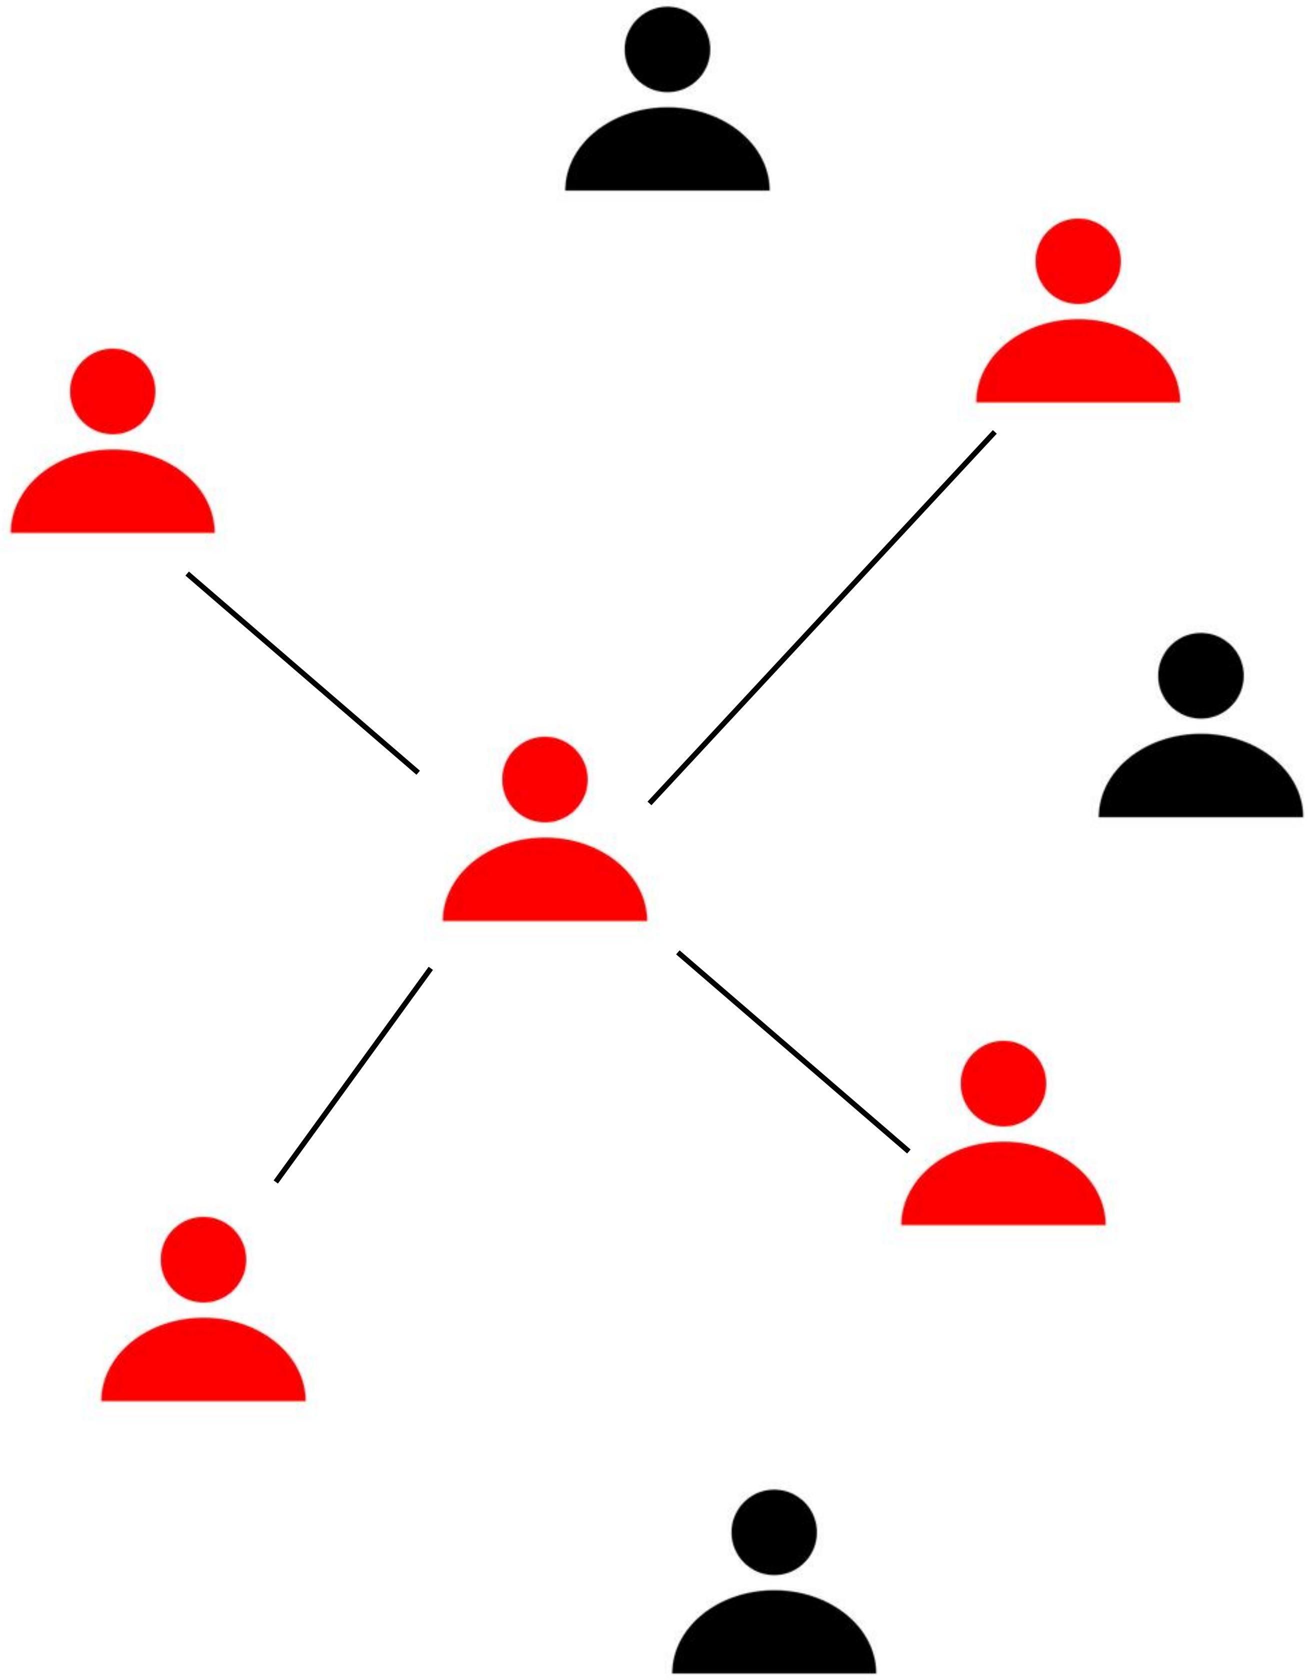

Difficult to measure:

1) Observe viral load over time; but how does it relate to infectivity?
2) Observe time between symptom onsets in transmission pairs (“serial interval”); but:
   - difficult to attribute transmission
   - Serial interval != generation interval (e.g., can be negative)

Challenge #1: Generation interval w_j

w_j Probability of infecting someone at time *j* after becoming infected

Difficult to measure:

1) Observe viral load over time; but how does it relate to infectivity?
2) Observe time between symptom onsets in transmission pairs ("serial interval"); but:
    - difficult to attribute transmission
    - Serial interval != generation interval (e.g., can be negative)

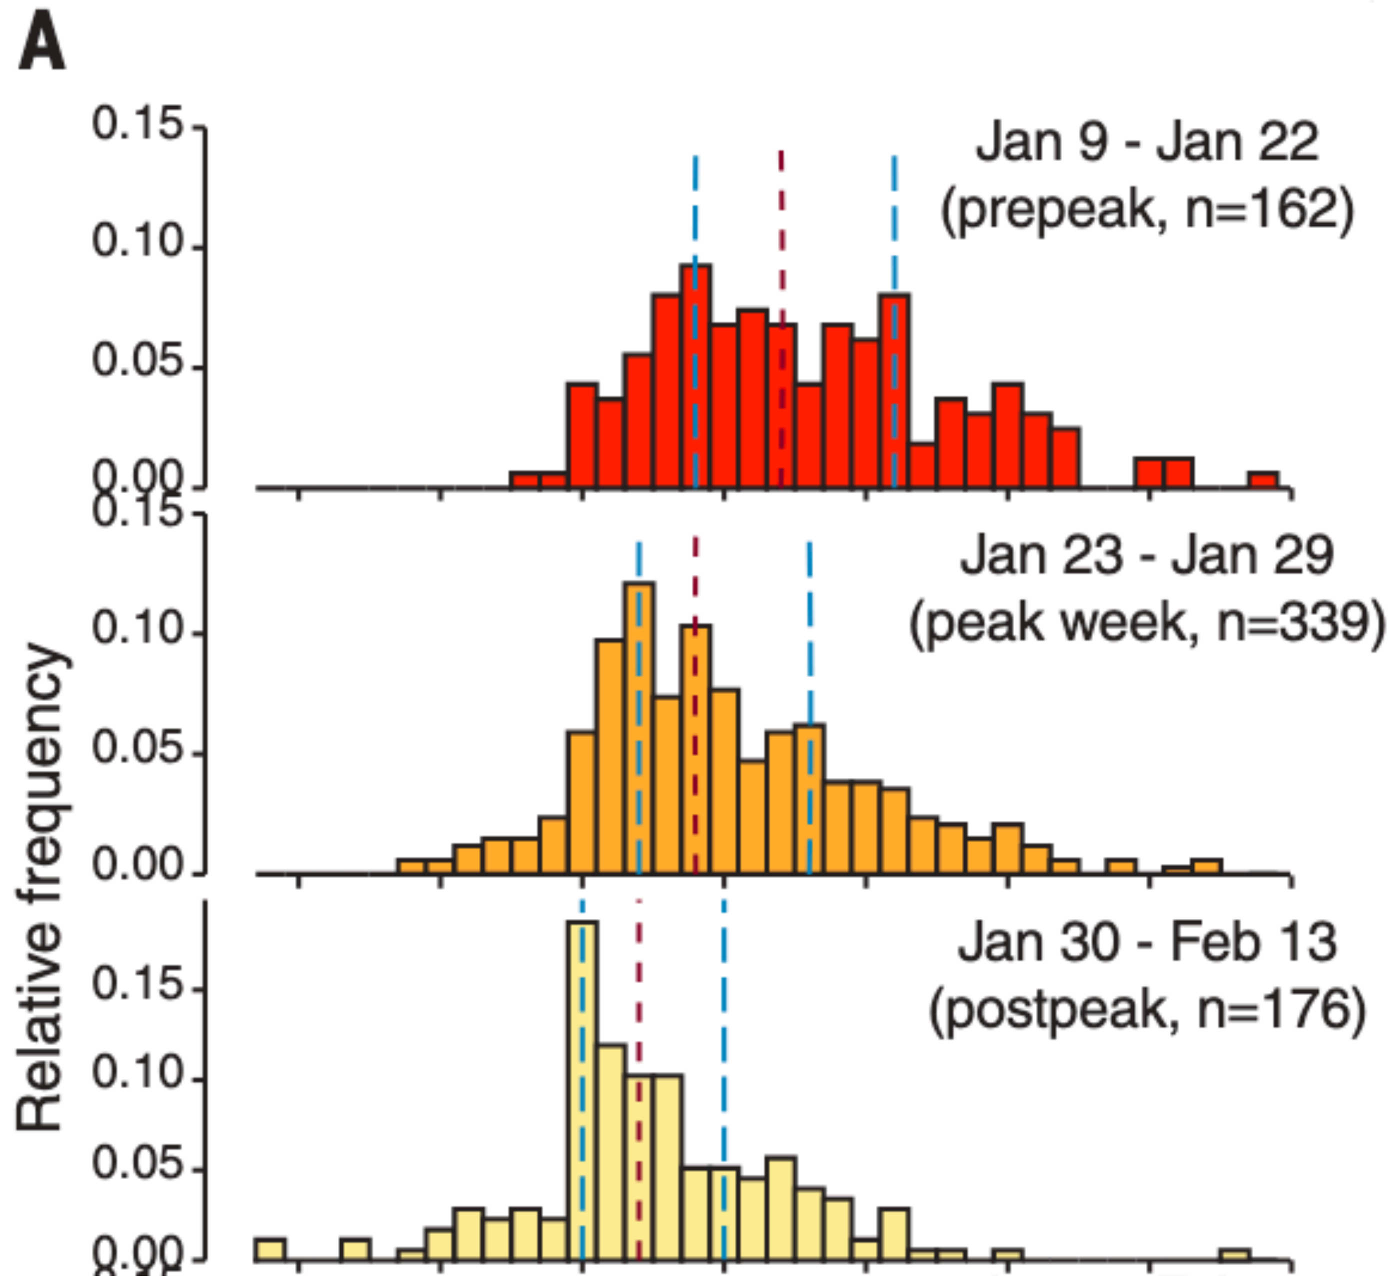

Taslim Ali et al., Science, 2020

Also: interventions can change the generation interval

Challenge #2: Infections are (usually) not observed

I_i Number of new infections at time *i* (e.g., days)

Solution:

- Consider an indirect measure of infections C_i (e.g., test-positive cases, hospital admissions, deaths)
- Problem #1: introduces a delay from infection to observation

If ξ_i is the delay distribution, one could backcalculate:

$$\hat{I}_i = \sum_{j=0}^{n} \xi_j C_{i+j}$$

Better: deconvolution, i.e solve

$$\sum_{j=0}^{n} \xi_j I_{i-j} = C_i$$

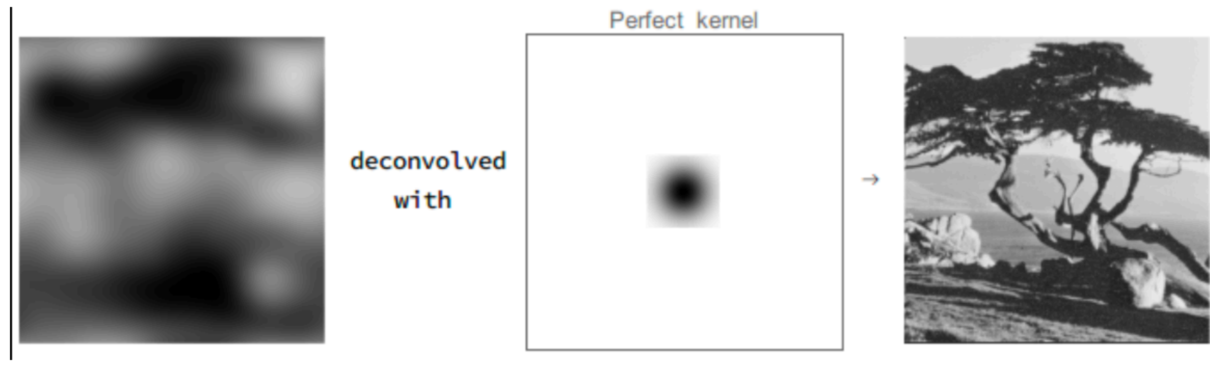

Challenge #2: Infections are (usually) not observed

I_i Number of new infections at time *i* (e.g., days)

Solution:

- Consider an indirect measure of infections C_i (e.g., test-positive cases, hospital admissions, deaths)
- Problem #1: introduces a delay from infection to observation

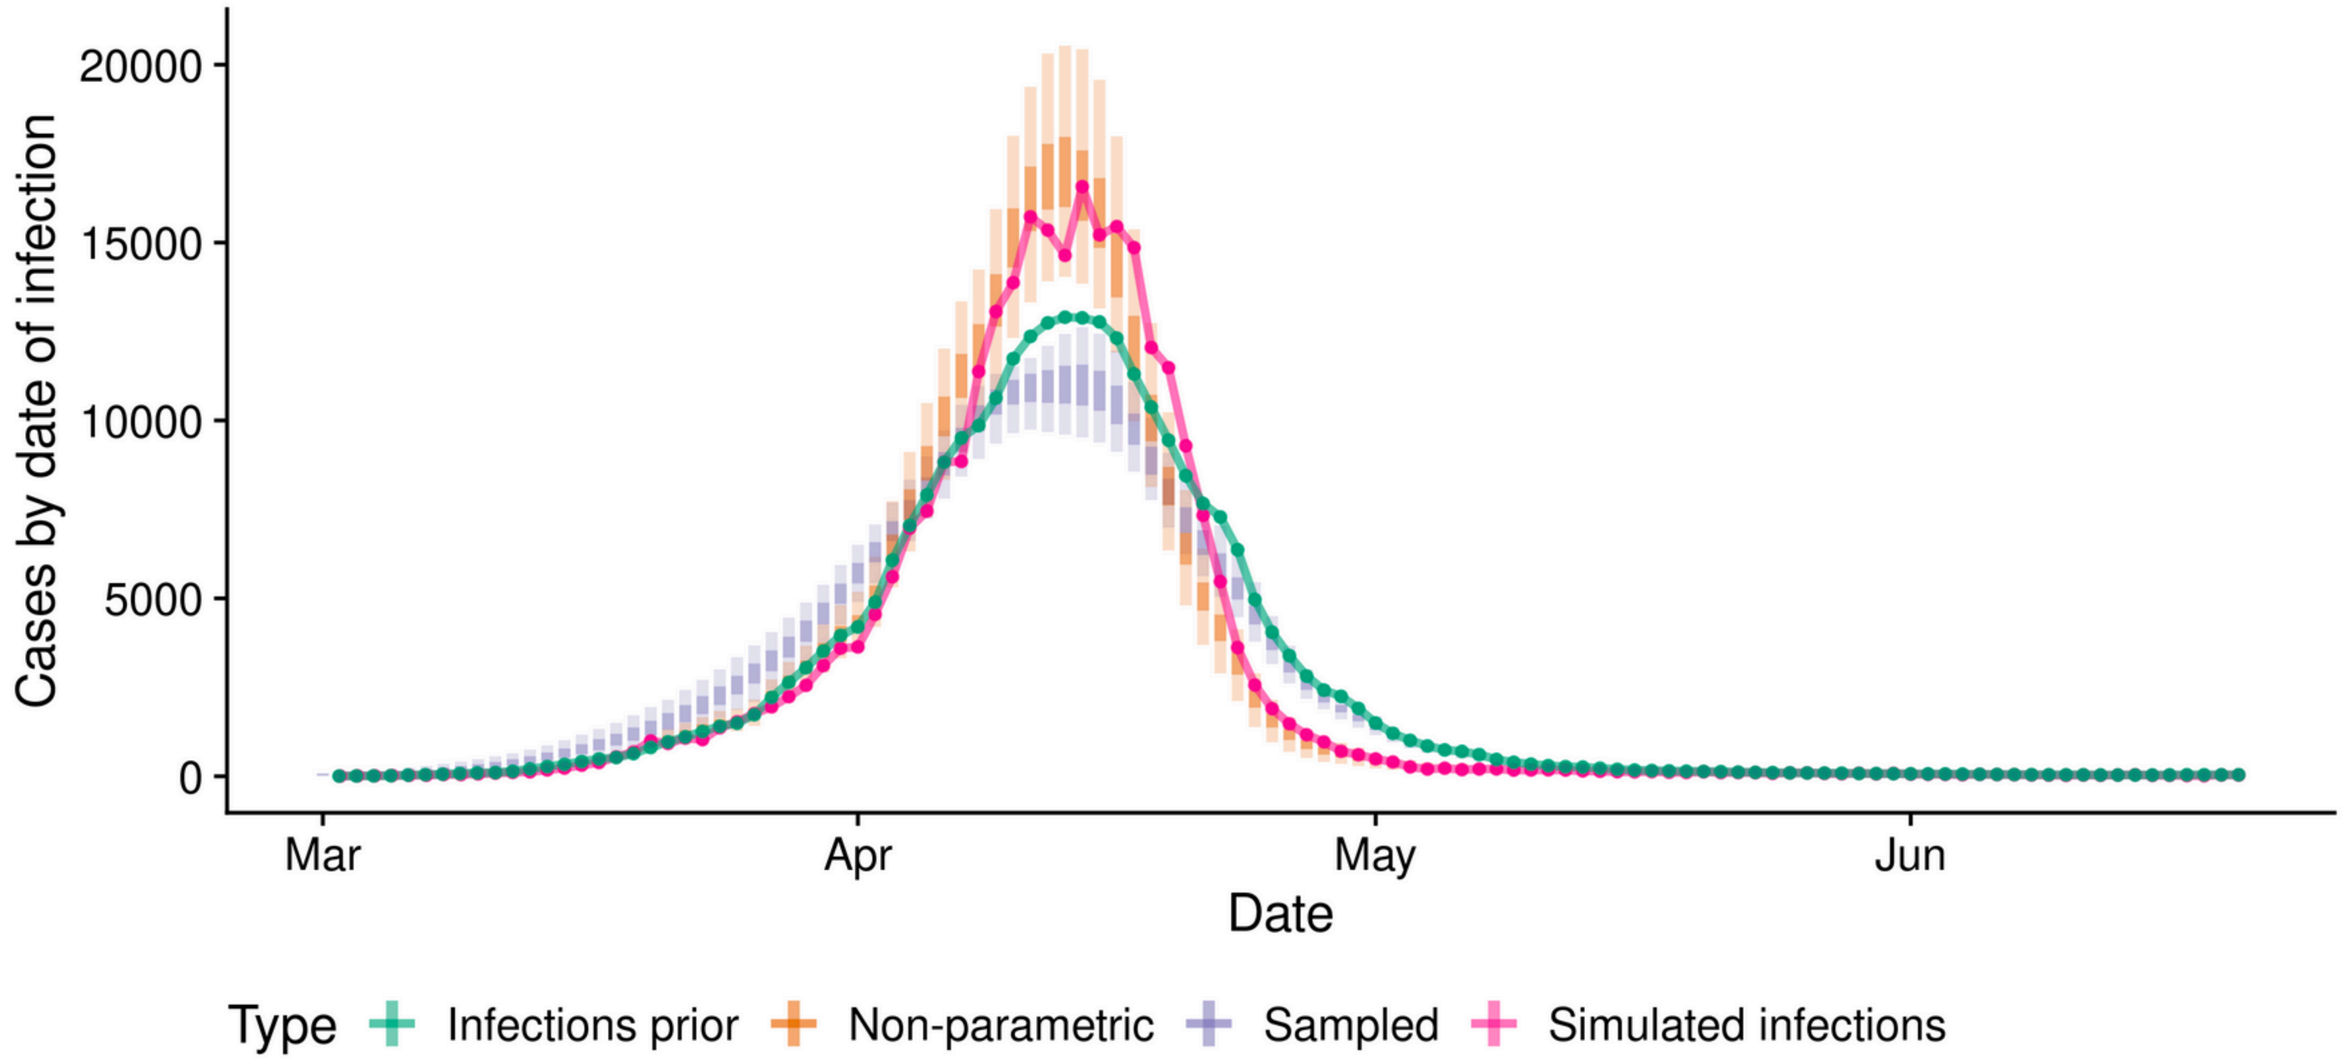

Better: deconvolution, i.e solve

$$\sum_{j=0}^{n} \xi_j I_{i-j} = C_i$$

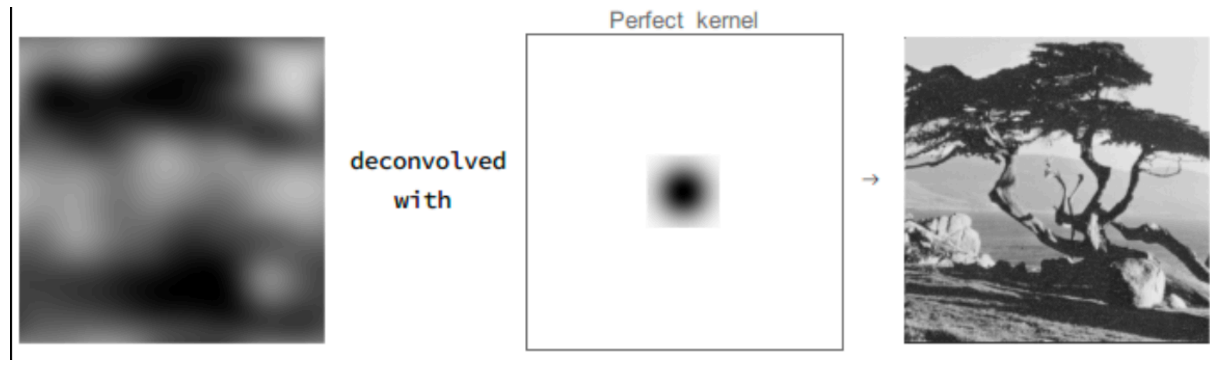

Challenge #2: Infections are (usually) not observed

I_i Number of new infections at time *i* (e.g., days)

Solution:

- Consider an indirect measure of infections C_i (e.g., test-positive cases, hospital admissions, deaths)
- If delays and are not accounted for, R_t will reference the wrong date

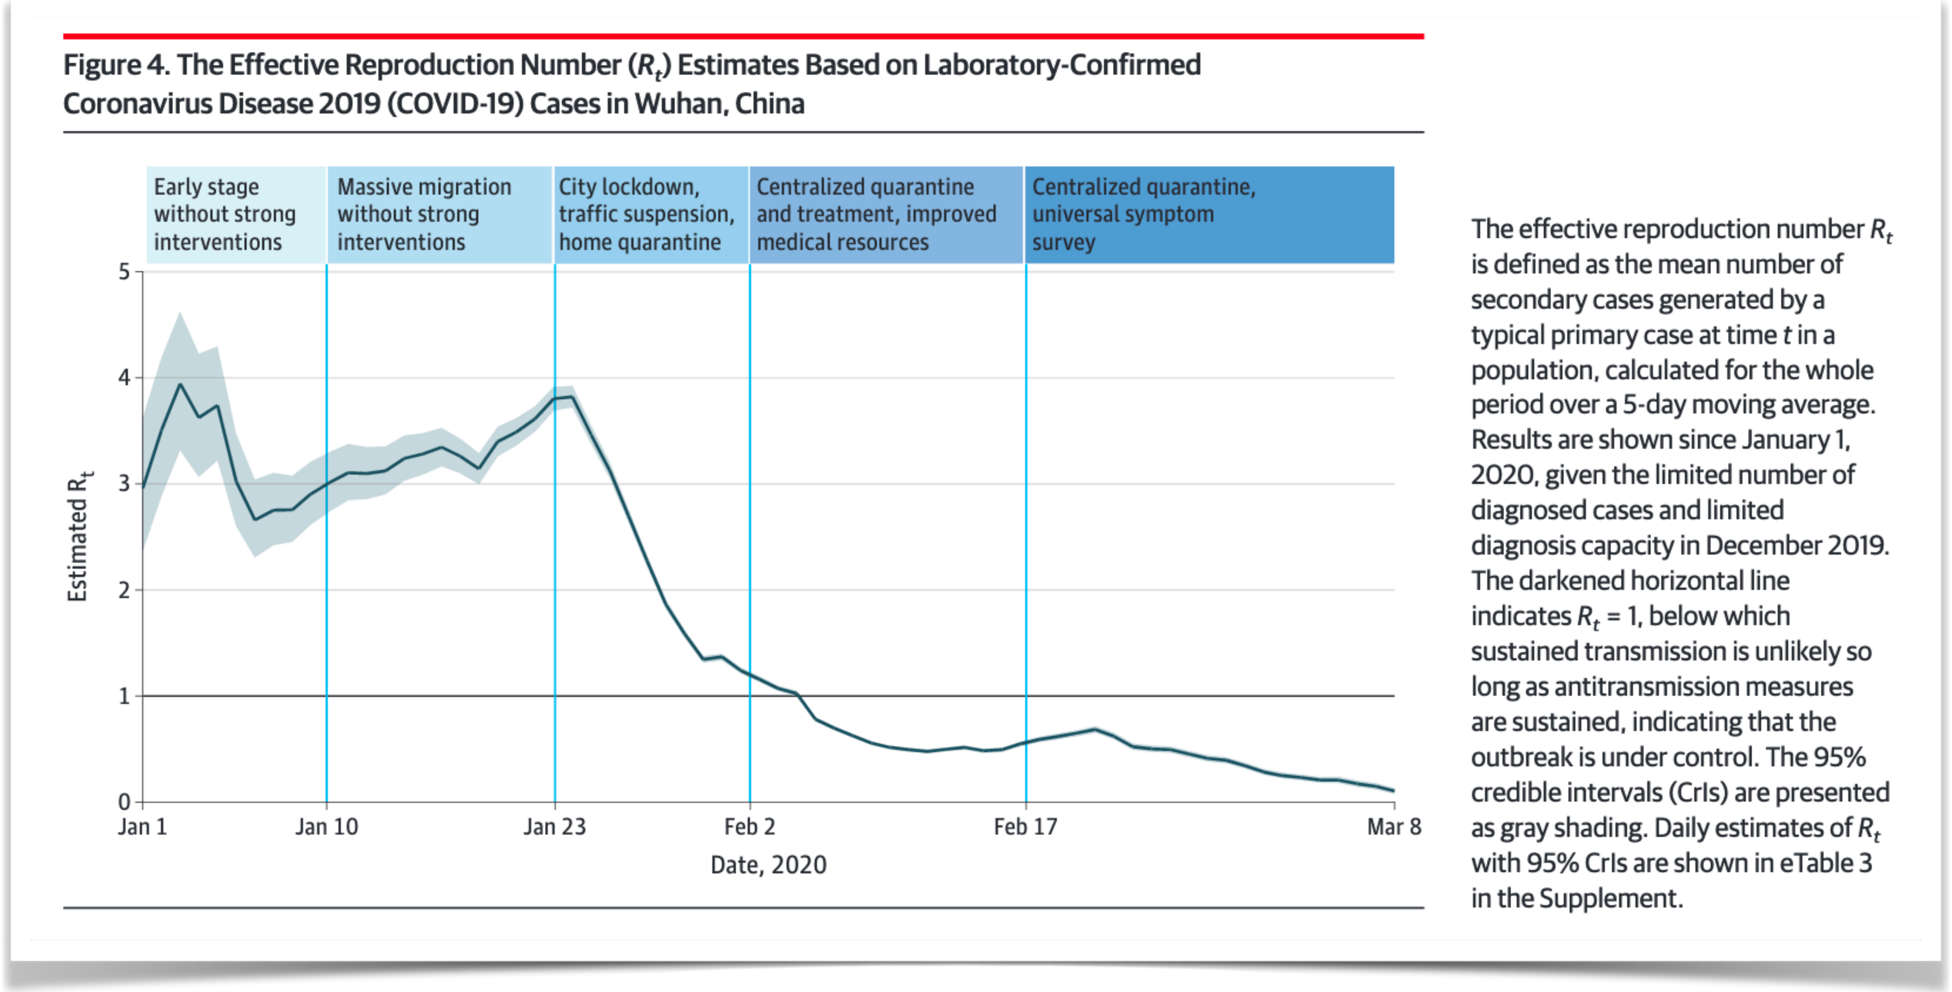

Figure 4. The Effective Reproduction Number (R_t) Estimates Based on Laboratory-Confirmed Coronavirus Disease 2019 (COVID-19) Cases in Wuhan, China

The effective reproduction number R_t is defined as the mean number of secondary cases generated by a typical primary case at time *t* in a population, calculated for the whole period over a 5-day moving average. Results are shown since January 1, 2020, given the limited number of diagnosed cases and limited diagnosis capacity in December 2019. The darkened horizontal line indicates R_t = 1, below which sustained transmission is unlikely so long as antitransmission measures are sustained, indicating that the outbreak is under control. The 95% credible intervals (CrIs) are presented as gray shading. Daily estimates of R_t with 95% CrIs are shown in eTable 3 in the Supplement.

Pan et al., JAMA, 2020

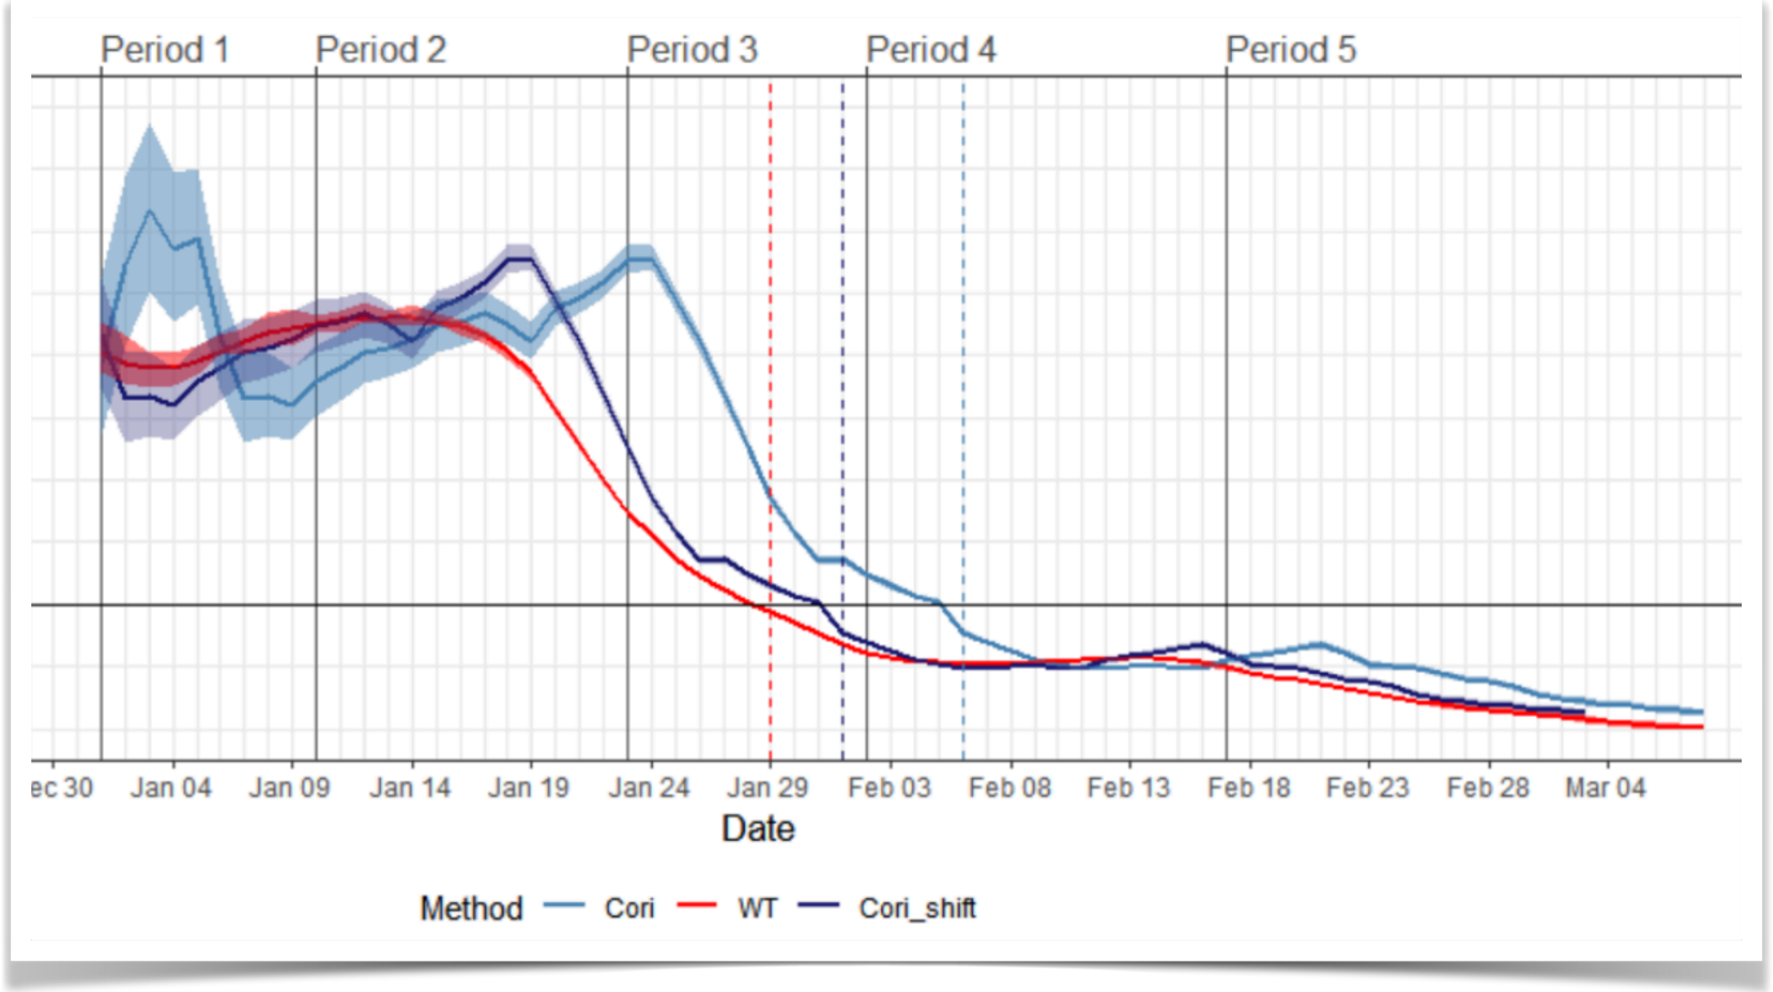

Lipsitch et al., GitHub, 2020

Challenge #2: Infections are (usually) not observed

I_i Number of new infections at time *i* (e.g., days)

Solution:

- Consider an indirect measure of infections C_i (e.g., test-positive cases, hospital admissions, deaths)
- Problem #2: introduces right truncation -> necessitates a “nowcast”

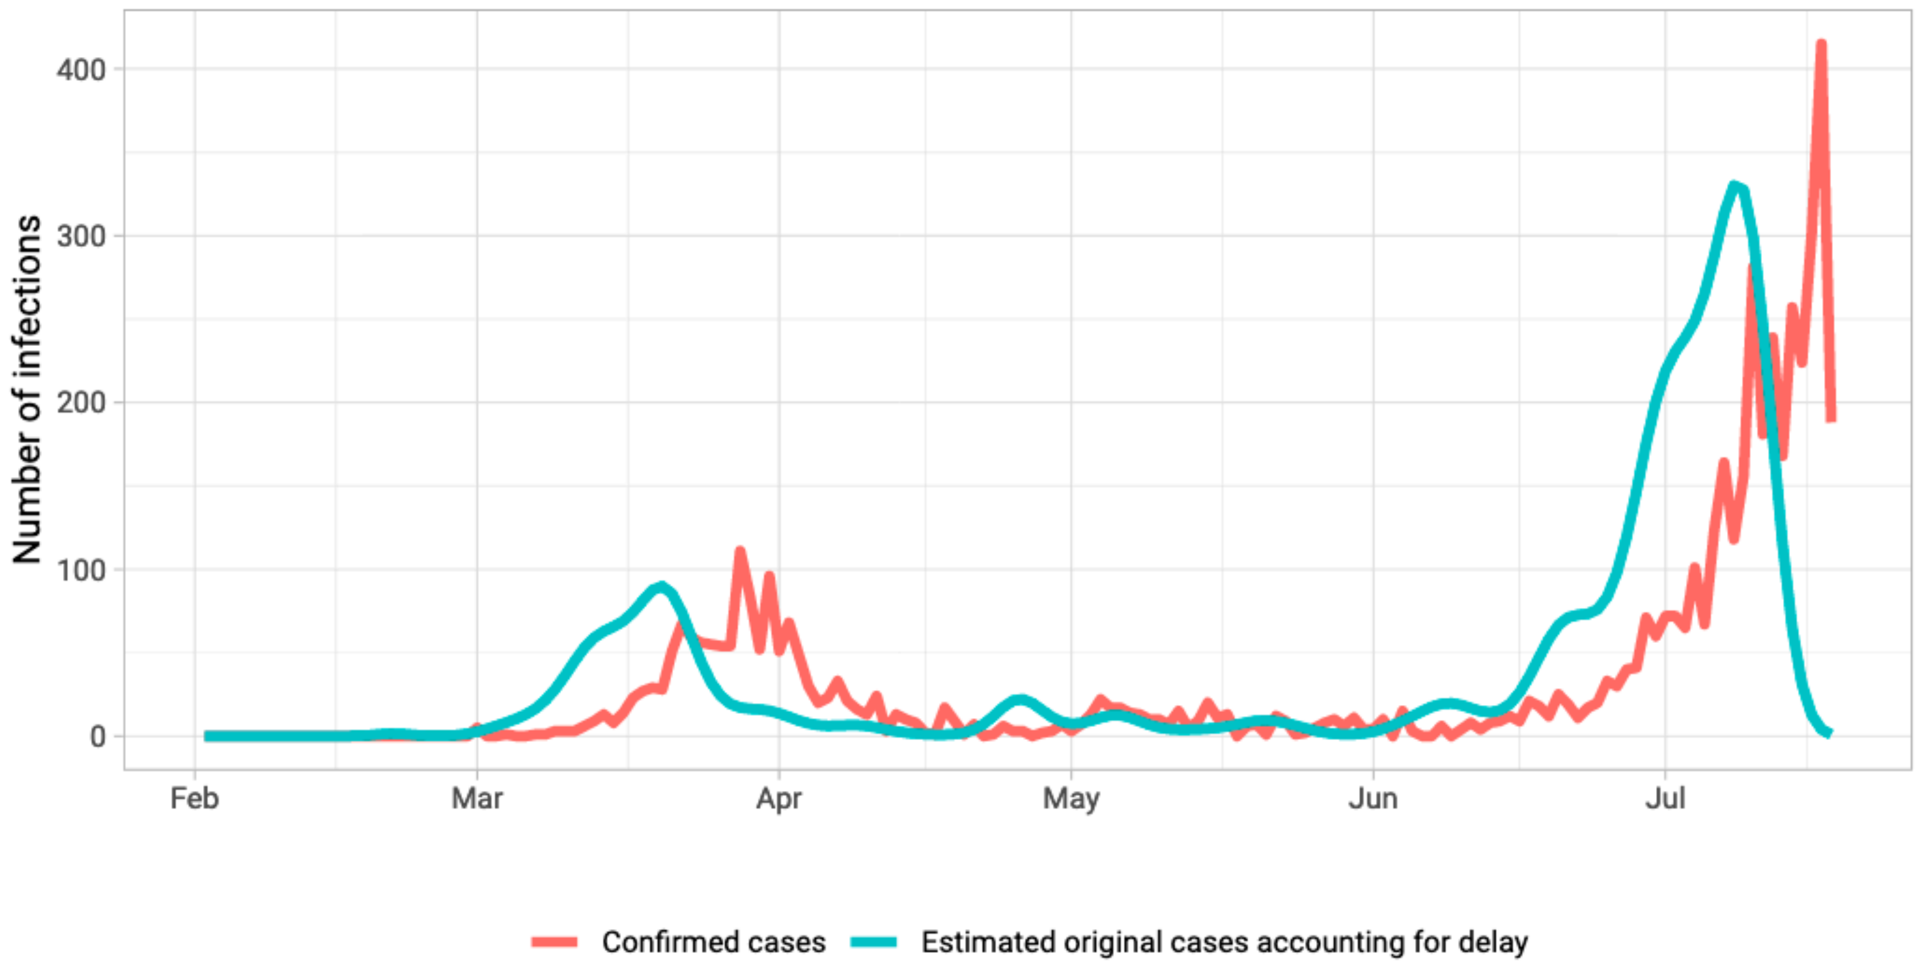

Back-projection of Victorian Covid-19 infections

Non-parametric back-projection of incidence cases assuming average delay of 10 days between infection and observation, using methods in Becker et al (1991). No correction for right truncation of data, so the last 15 days will be badly biased downwards.

http://freerangestats.info/blog/

Challenge #2: Infections are (usually) not observed

I_i Number of new infections at time *i* (e.g., days)

Solution:

- Consider an indirect measure of infections C_i (e.g., test-positive cases, hospital admissions, deaths)
- Problem #3: Period effects and other artefacts

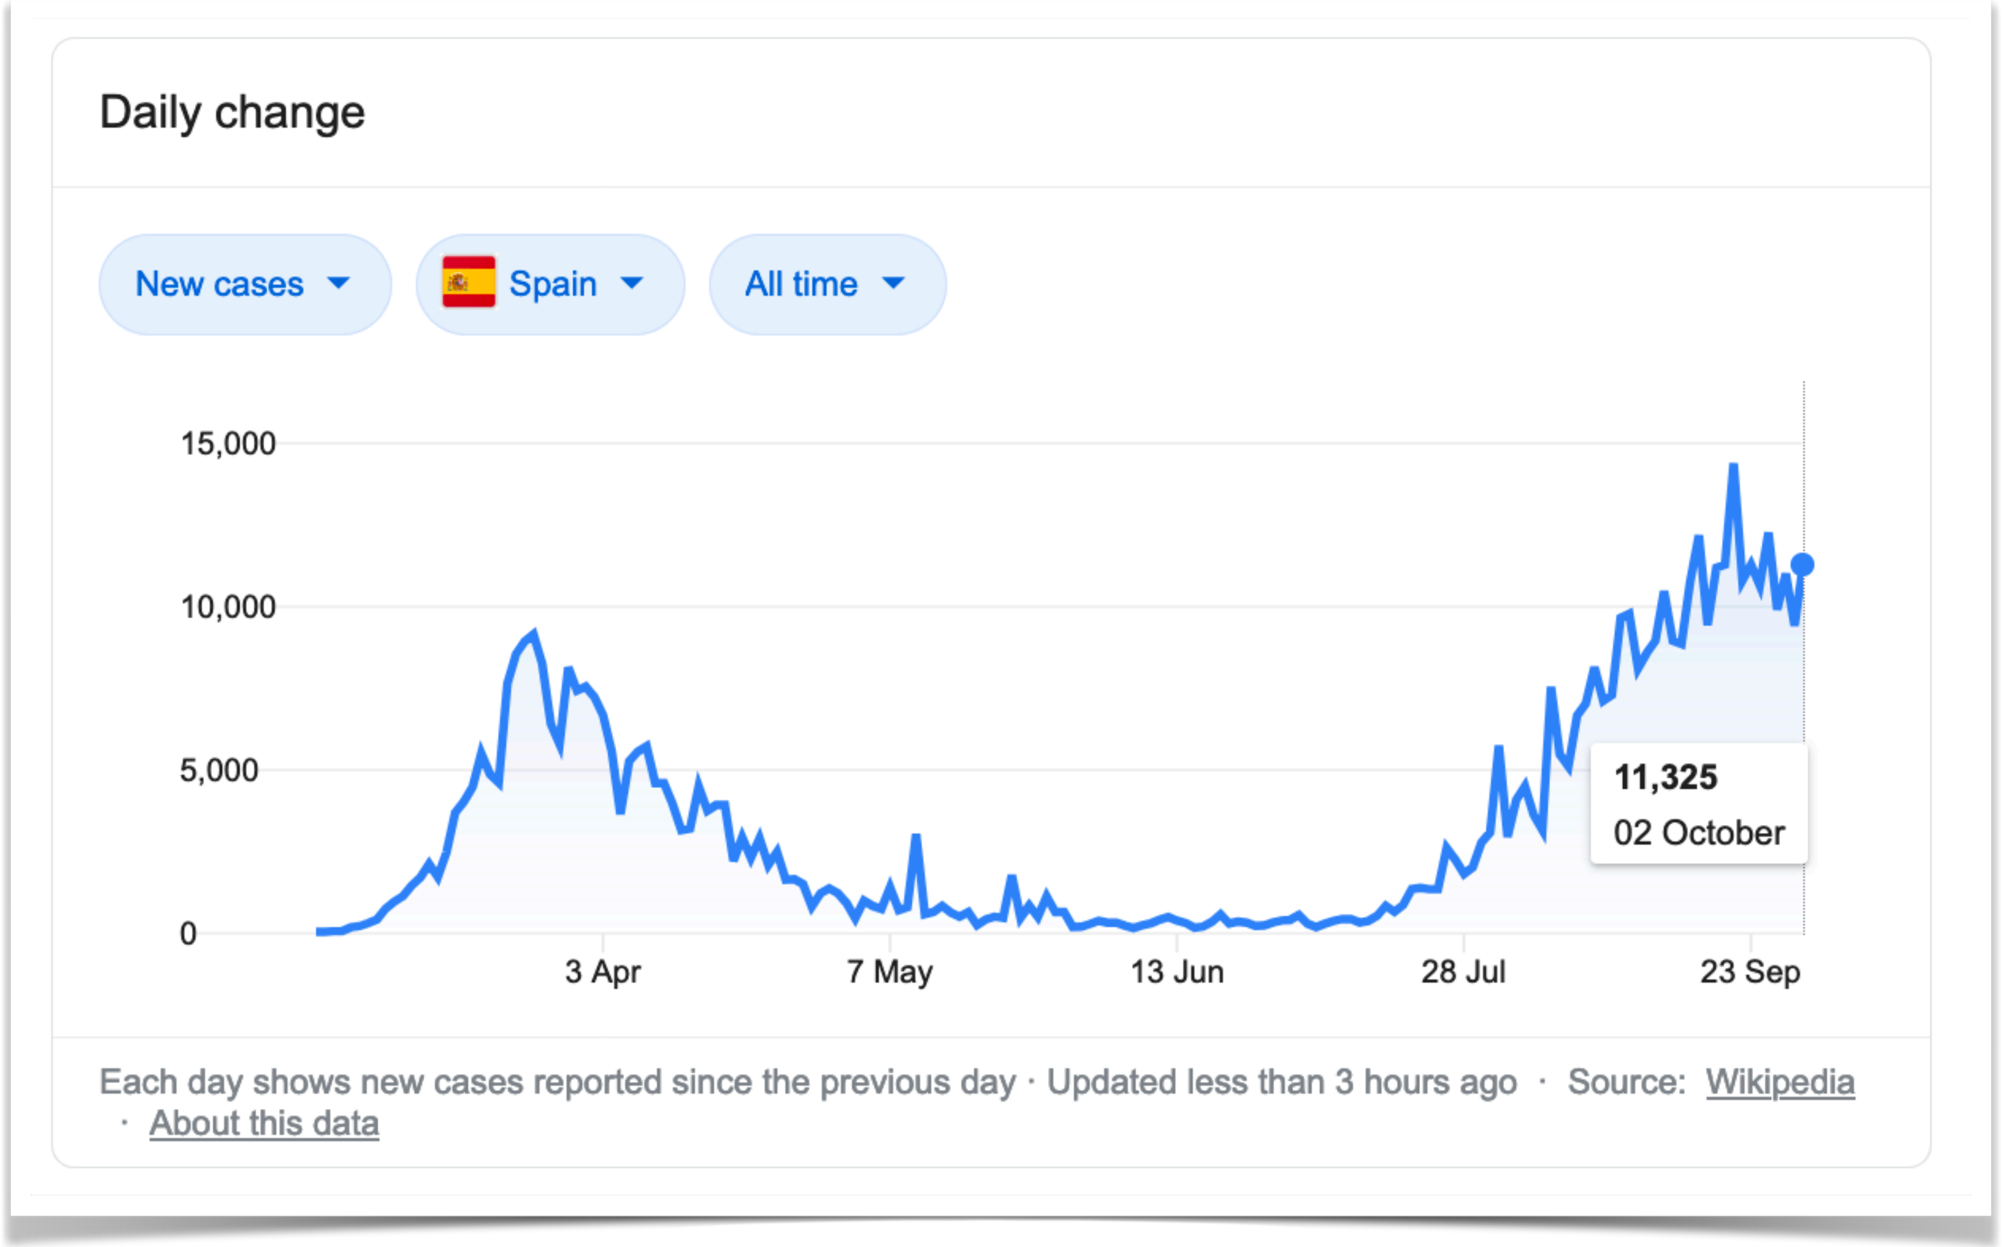

Each day shows new cases reported since the previous day · Updated less than 3 hours ago · Source: Wikipedia · About this data

Challenge #3: Choice of smoothing window

The time period over which R_t is assumed to be smooth or constant affects the estimate.

Solution:

- Optimise one-step ahead predictions.

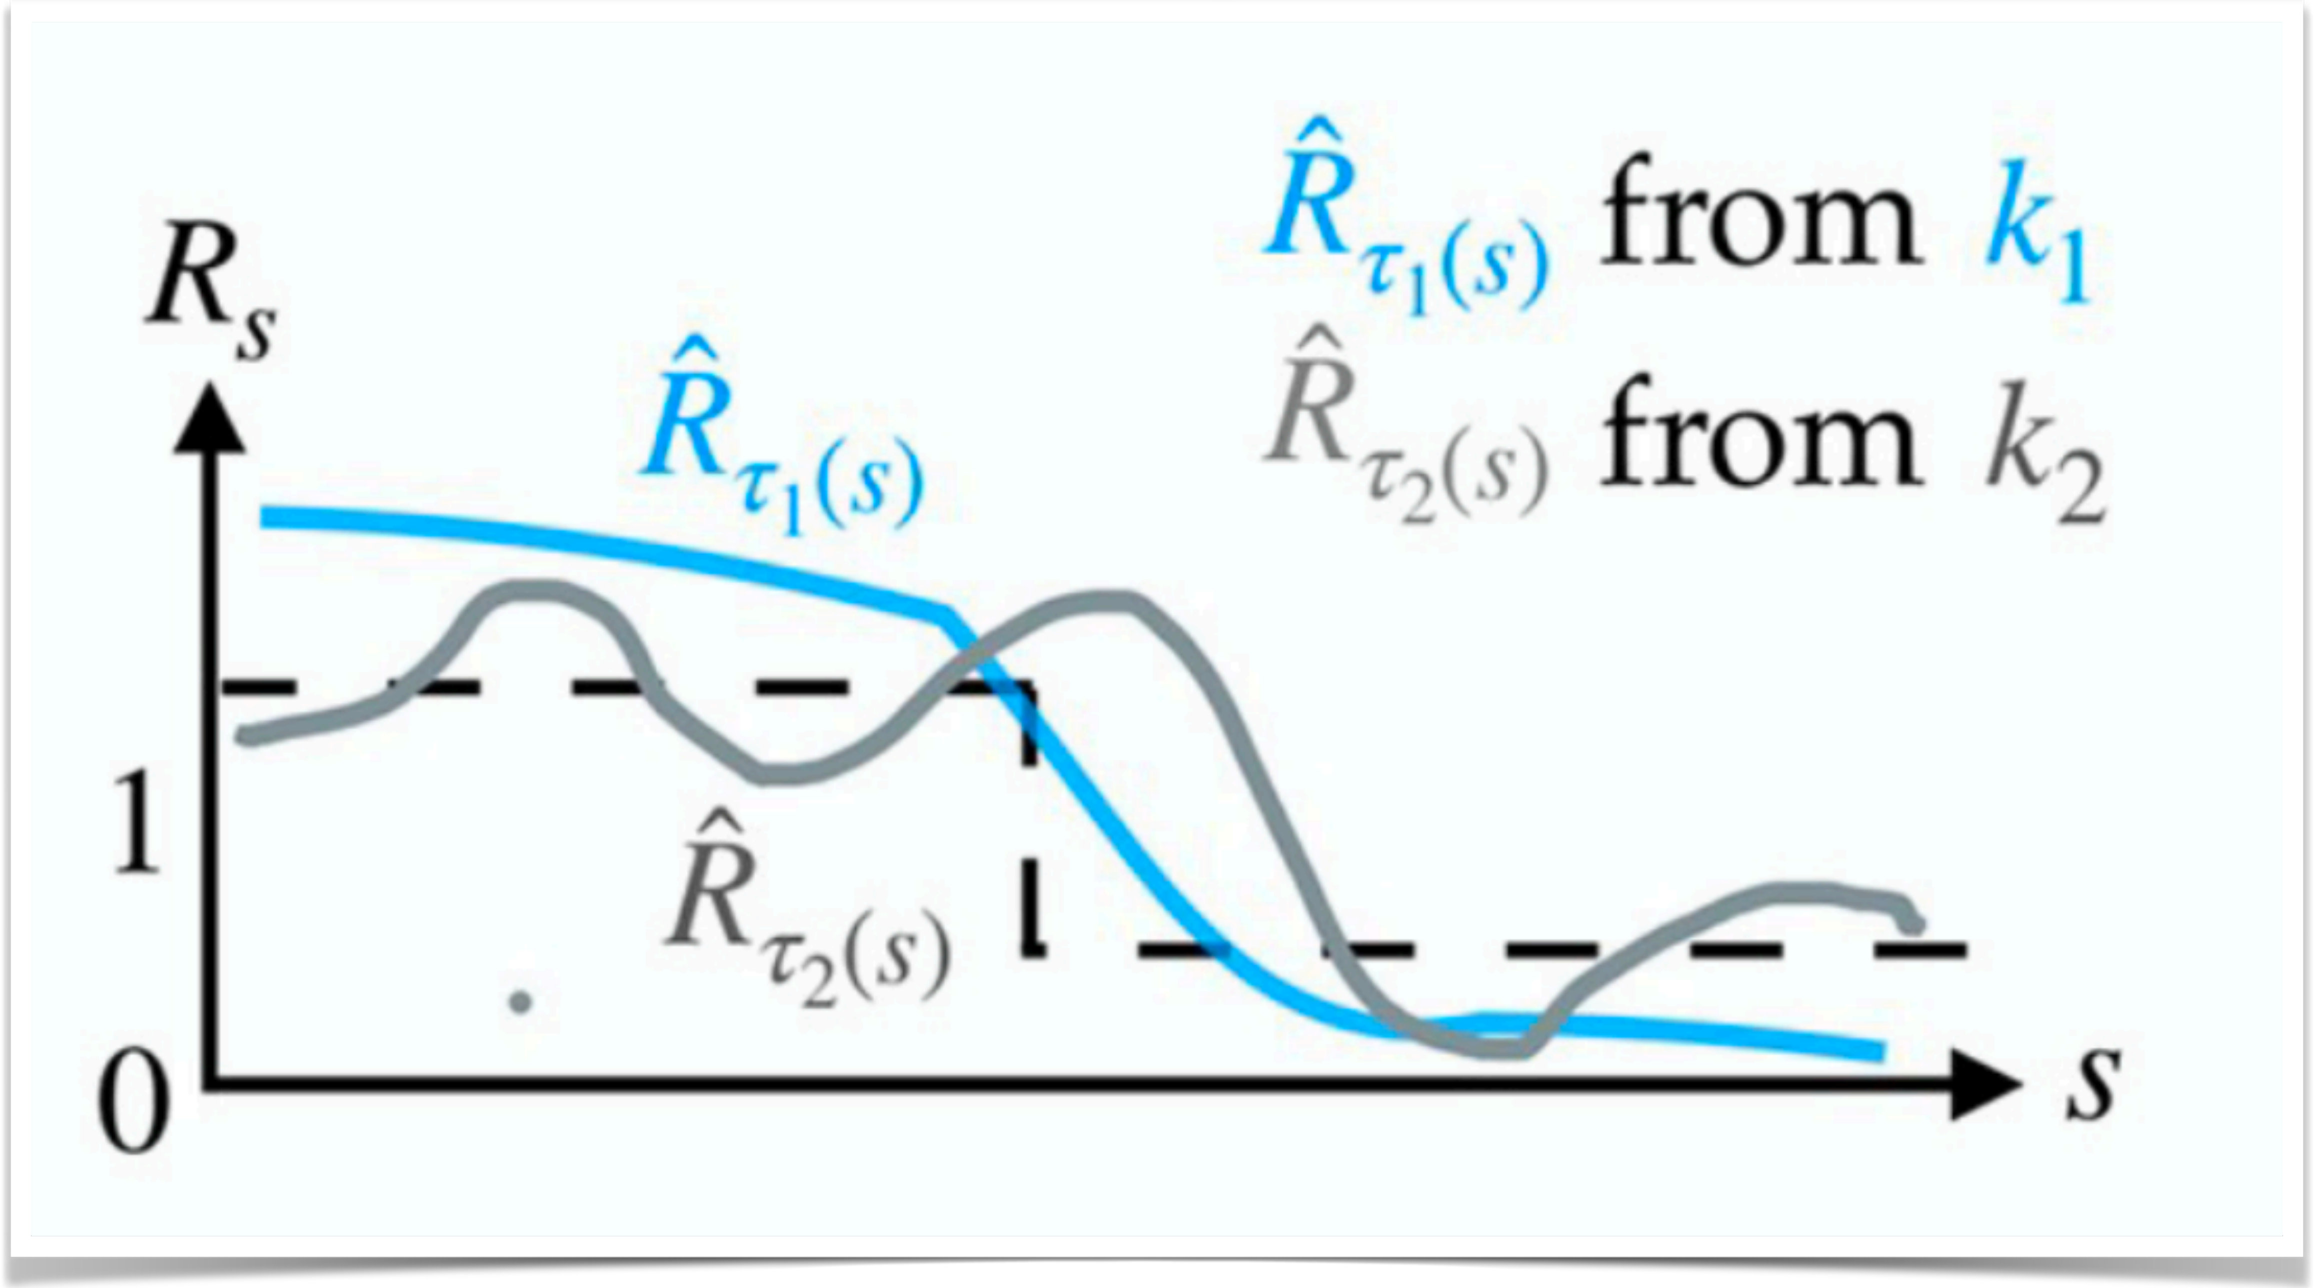

Parag & Donnelly, PLOS Comp Biol, 2020

Our approach

- Bayesian approach combining nowcasting and Rt estimation
- Generation interval estimates from Singapore allowing for negative serial intervals
- Latent process for estimating I_t
- Negative binomial reporting with multiplicative day-of-the-week effect
- R_t estimates with correlation between R_{t+1} and R_t based on Gaussian Process prior (squared exponential kernel)
- All implemented in Stan and as open-source R package https://epiforecasts.io/EpiNow2/
- Abbot et al., Wellcome Open Res, 2020

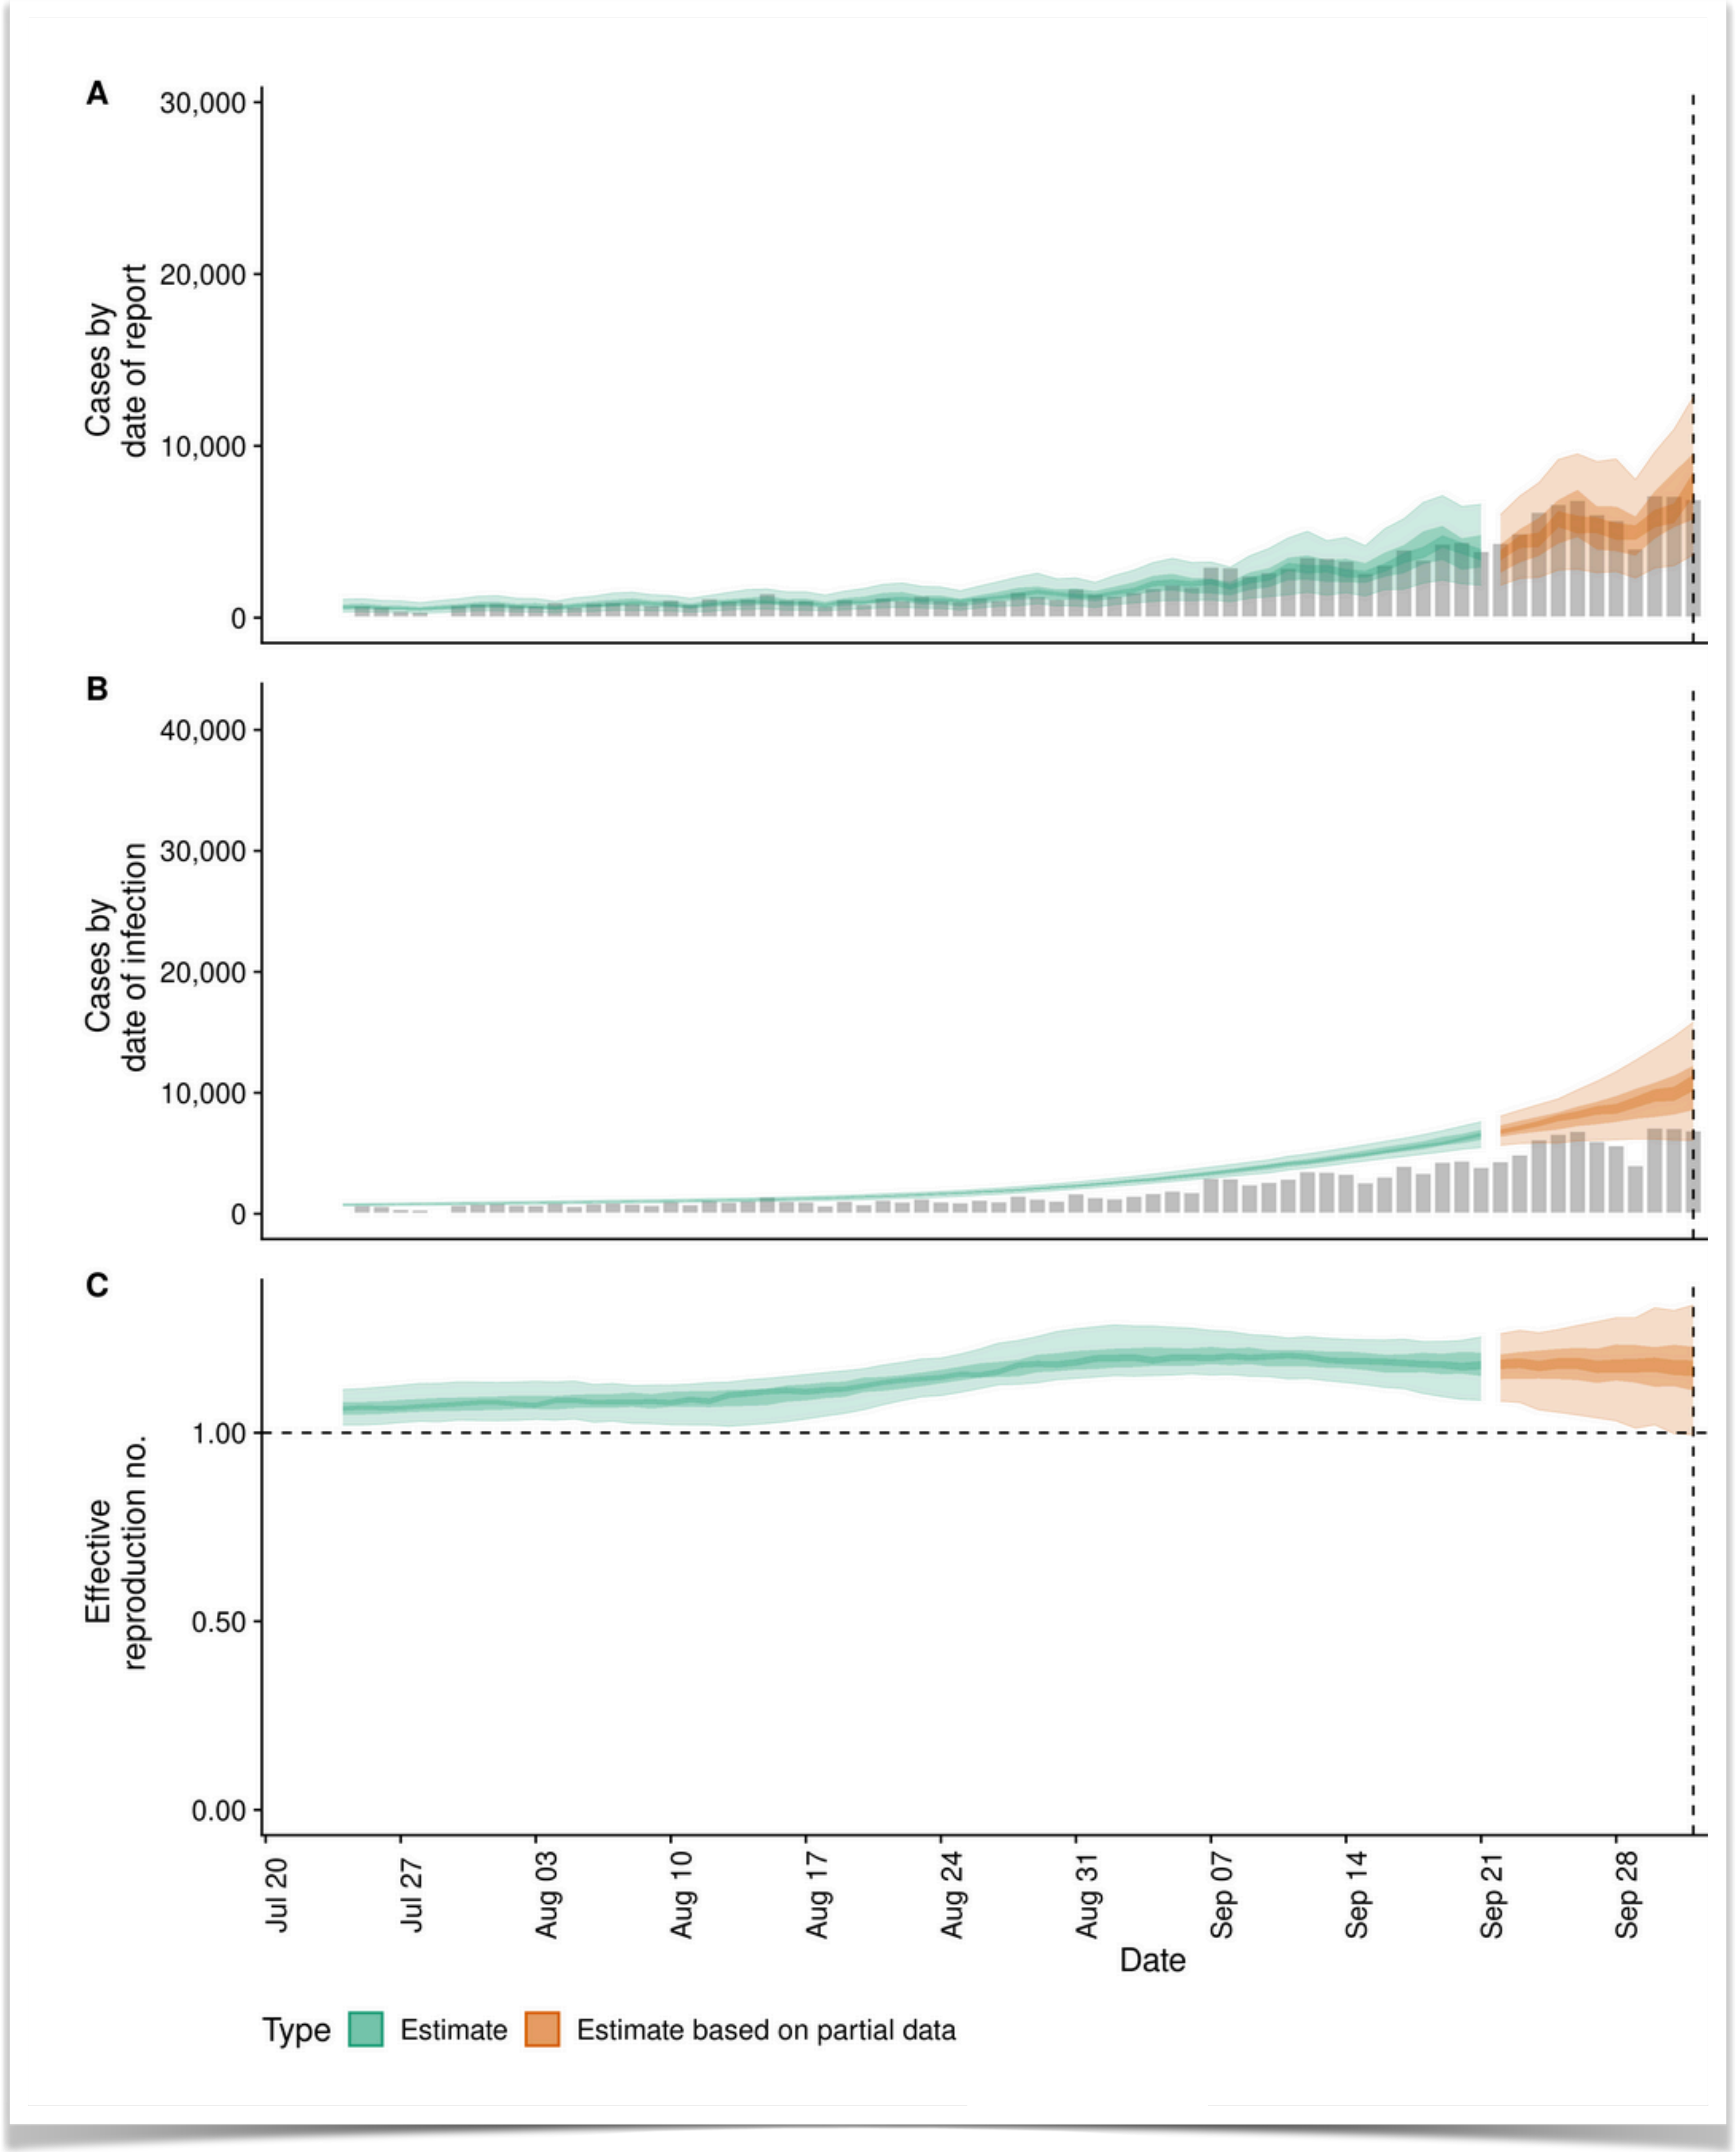

https://epiforecasts.io/covid/posts/national/united-kingdom/

Our approach

$$R_t \sim R_{t-1} \times \exp\left(\mathrm{GP}\right)$$

$$I_t = R_t \sum_{\tau} w_\tau I_{t-\tau}$$

$$D_t = \sum_{\tau} \xi_\tau I_{t-\tau}$$

$$C_t \sim \mathrm{NB}(D_t \omega_{(t \mod 7)}, \phi)$$

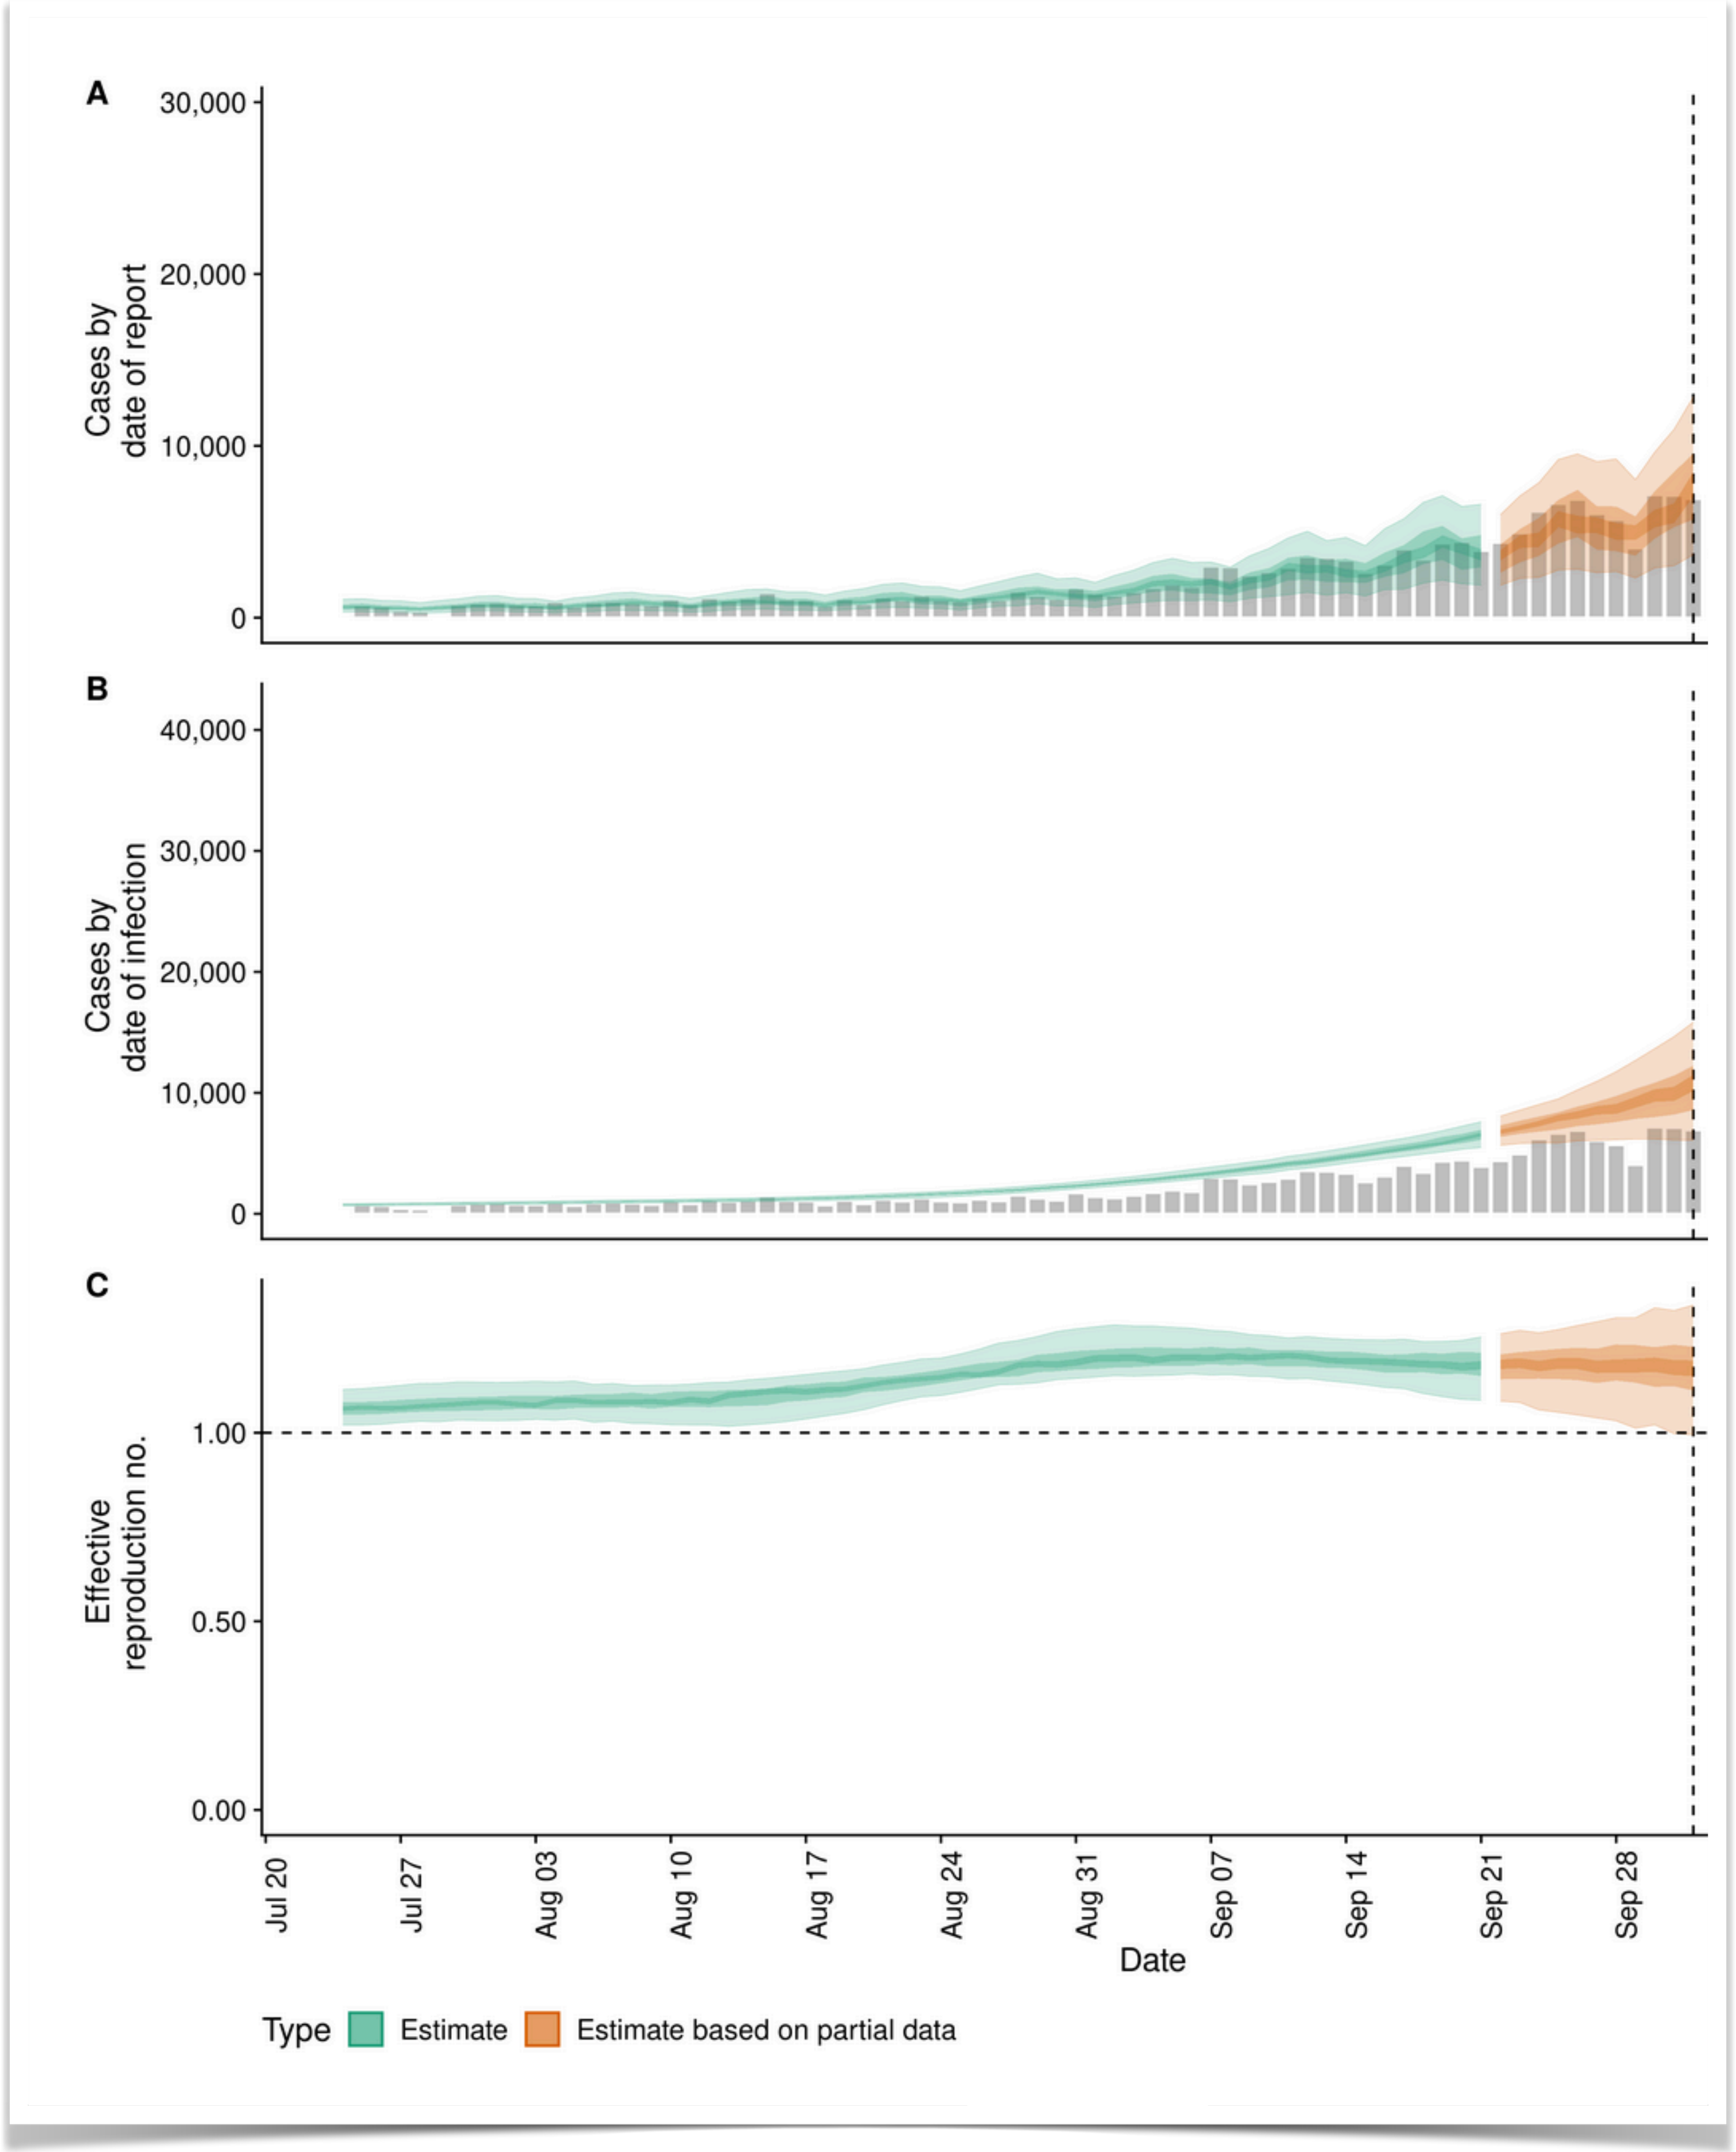

Stan code: https://git.io/JUxRt

https://epiforecasts.io/covid/posts/national/united-kingdom/

Our approach: public-facing work

- Subsampling from global linelist to estimate delay distributions (Xu et al., Scientific Data, 2020)
- Use daily global cases and deaths based on public data provided by ECDC
- Daily estimates are downloaded and used by several organisations (e.g. WHO Euro)
- Several other groups around the world have used our methods and tools to inform policy makers independently (Price et al., eLife, 2020).

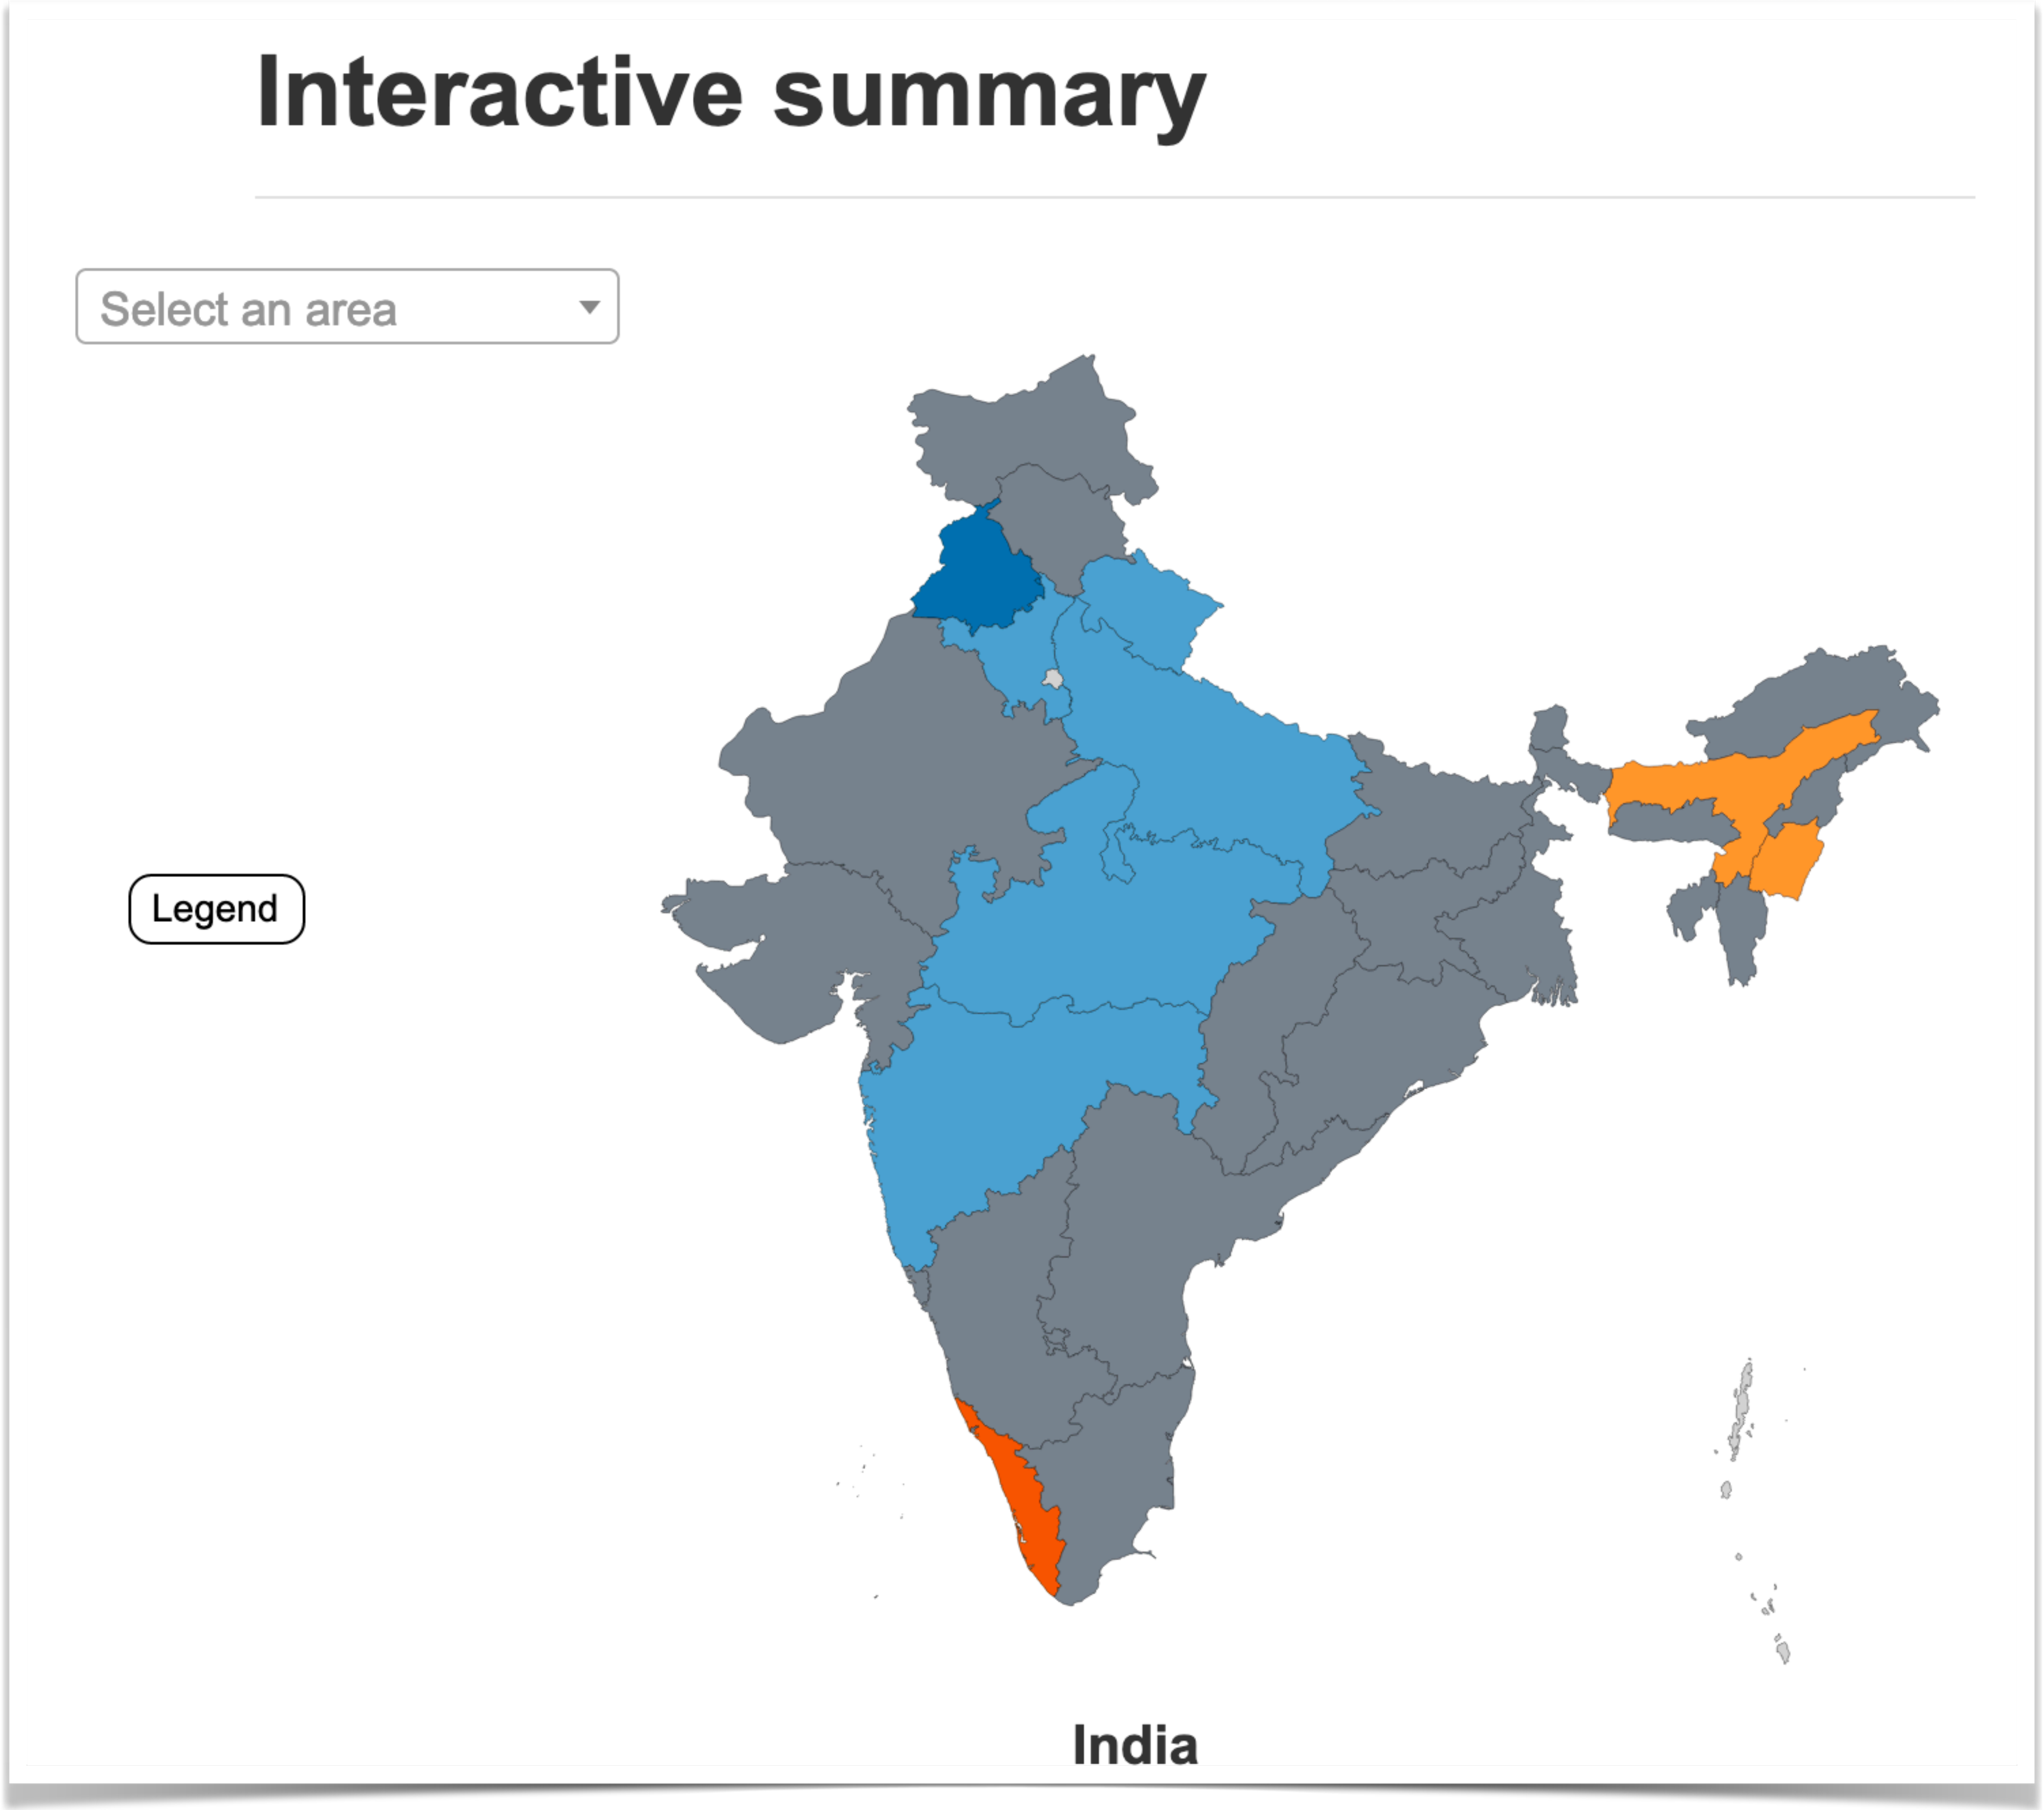

https://epiforecasts.io/covid/posts/national/india/

Our approach: UK work

- We have been producing regional and national level estimates of Rt in the UK based on hospital admissions data and deaths since the end of March.
- These have been used by policy-makers and scientific advisors to assess changes in transmission.
- Our estimates have been combined with those from other groups to produce an ensemble estimate.
- We have also used our methods to assess outbreaks in care homes, hospitals and in the community.

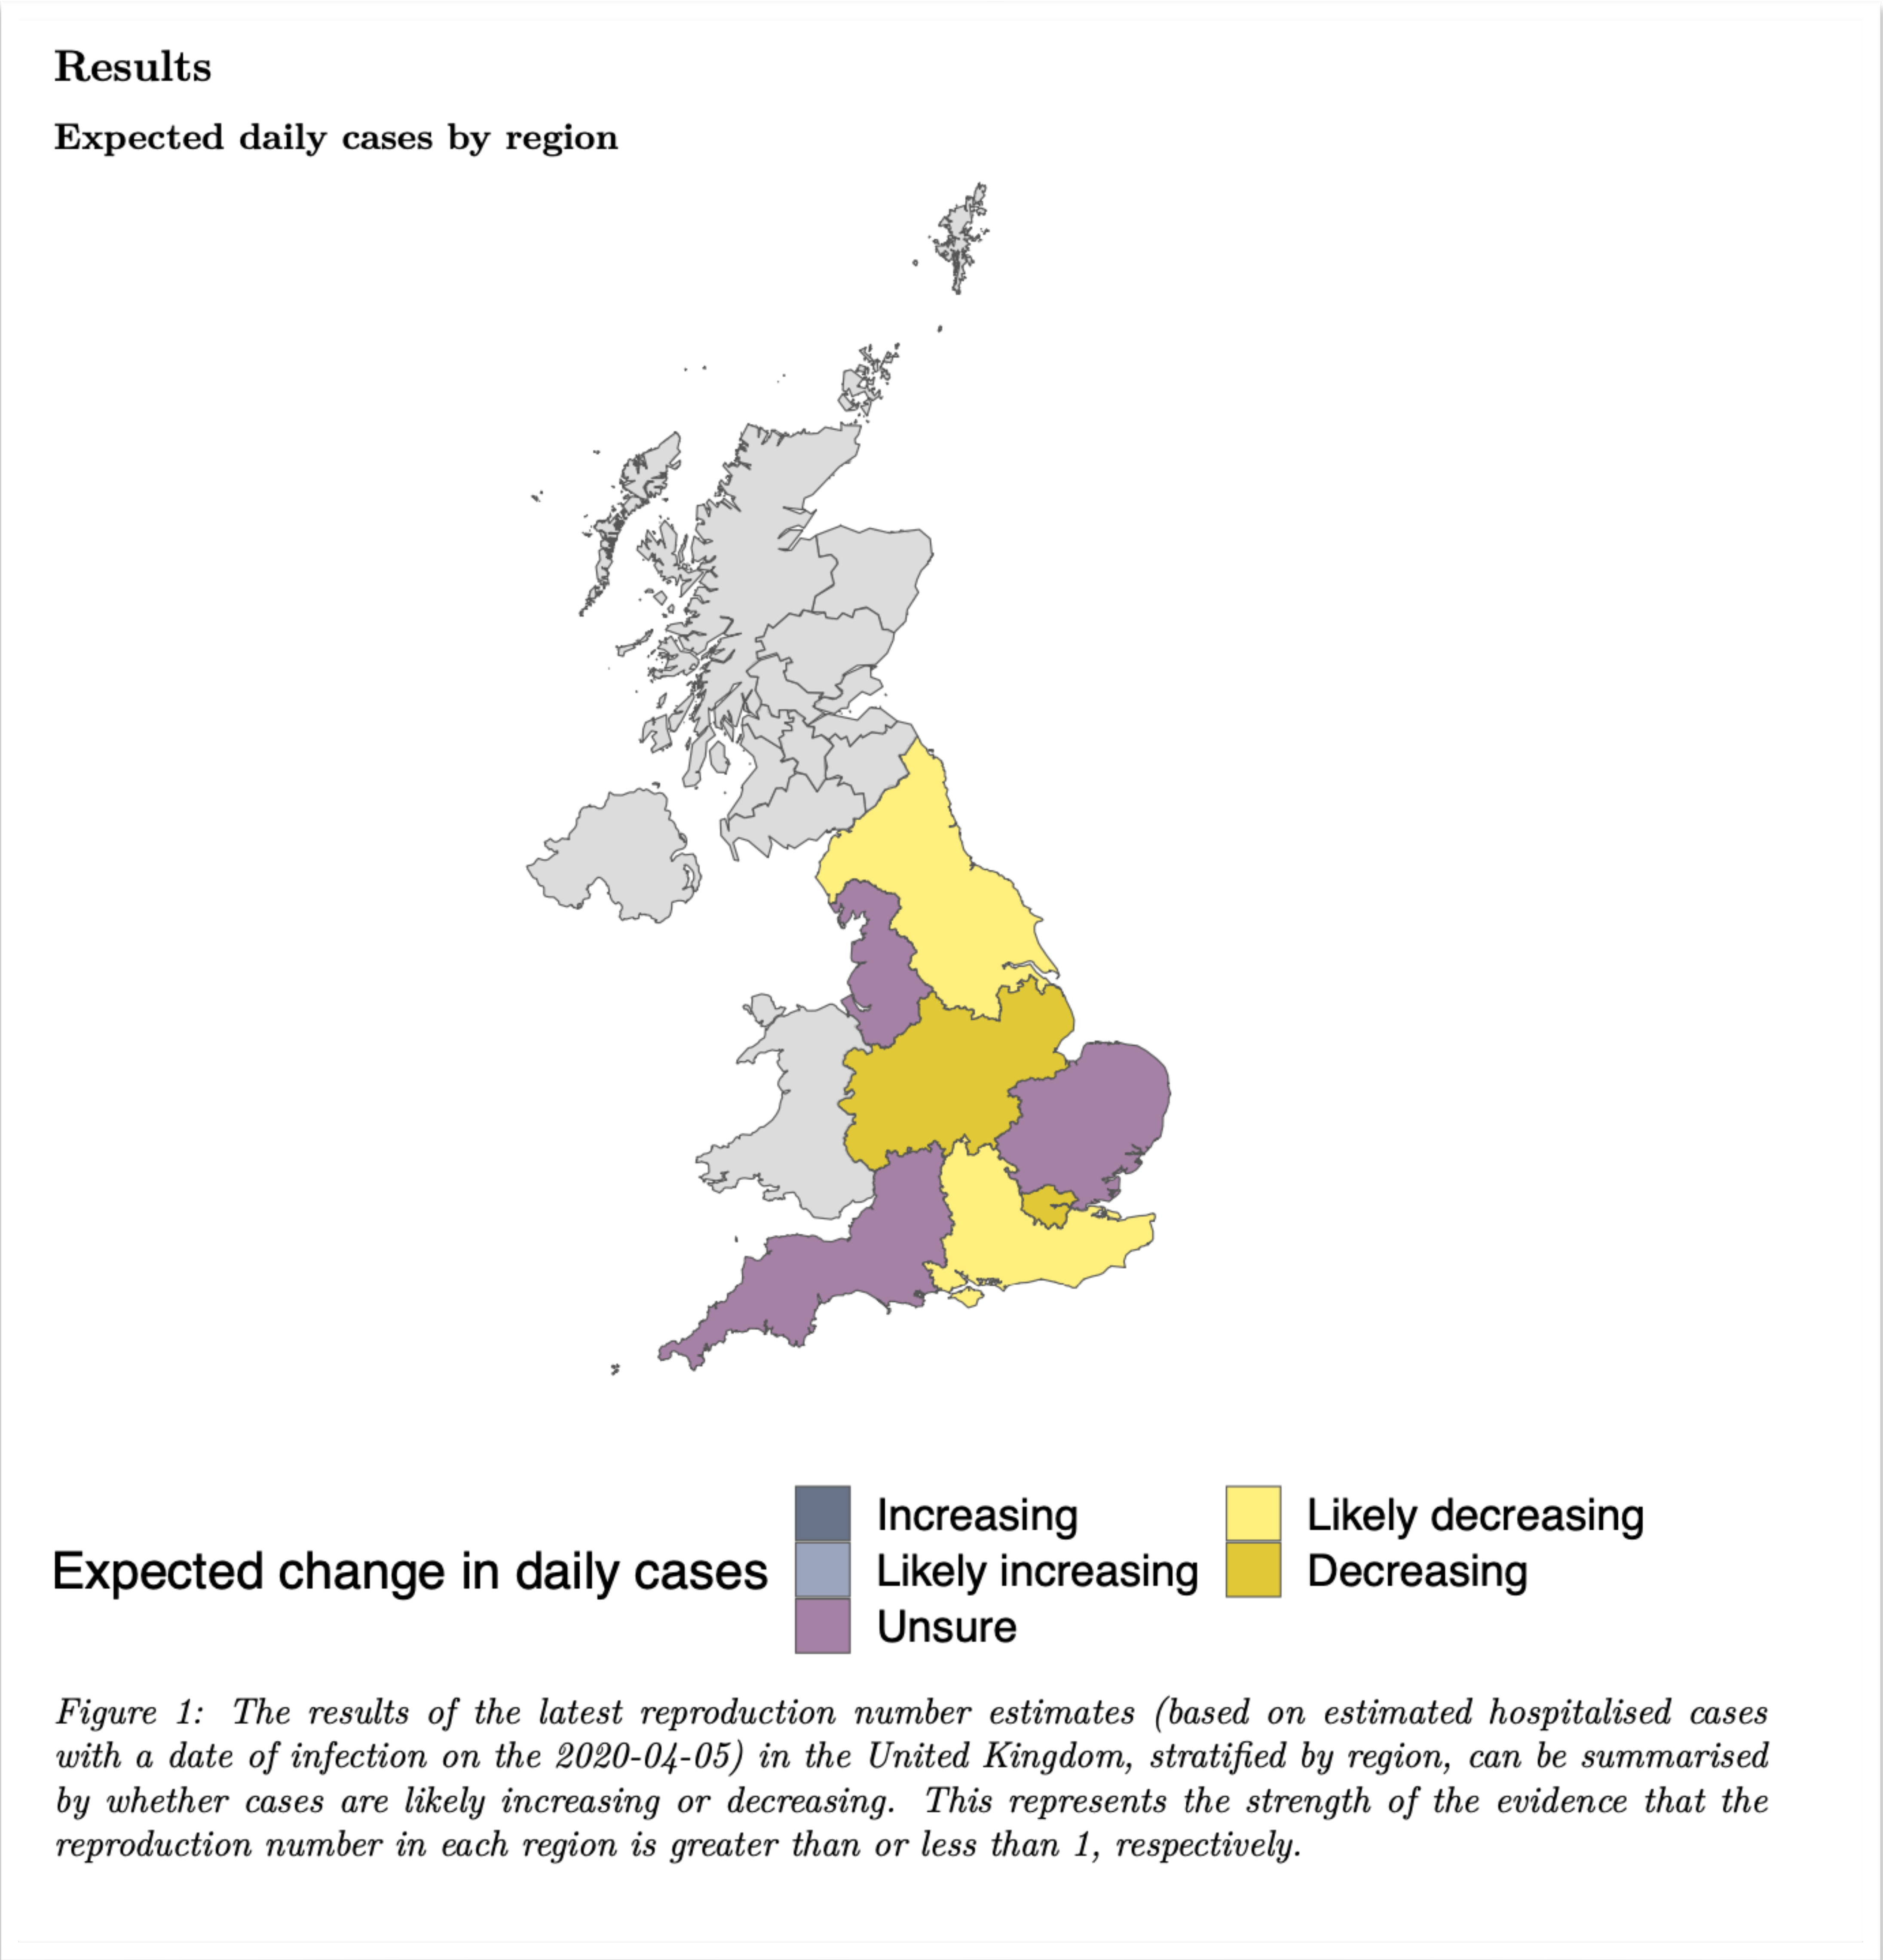

Results

Expected daily cases by region

Figure 1: The results of the latest reproduction number estimates (based on estimated hospitalised cases with a date of infection on the 2020-04-05) in the United Kingdom, stratified by region, can be summarised by whether cases are likely increasing or decreasing. This represents the strength of the evidence that the reproduction number in each region is greater than or less than 1, respectively.

Multiple data sources

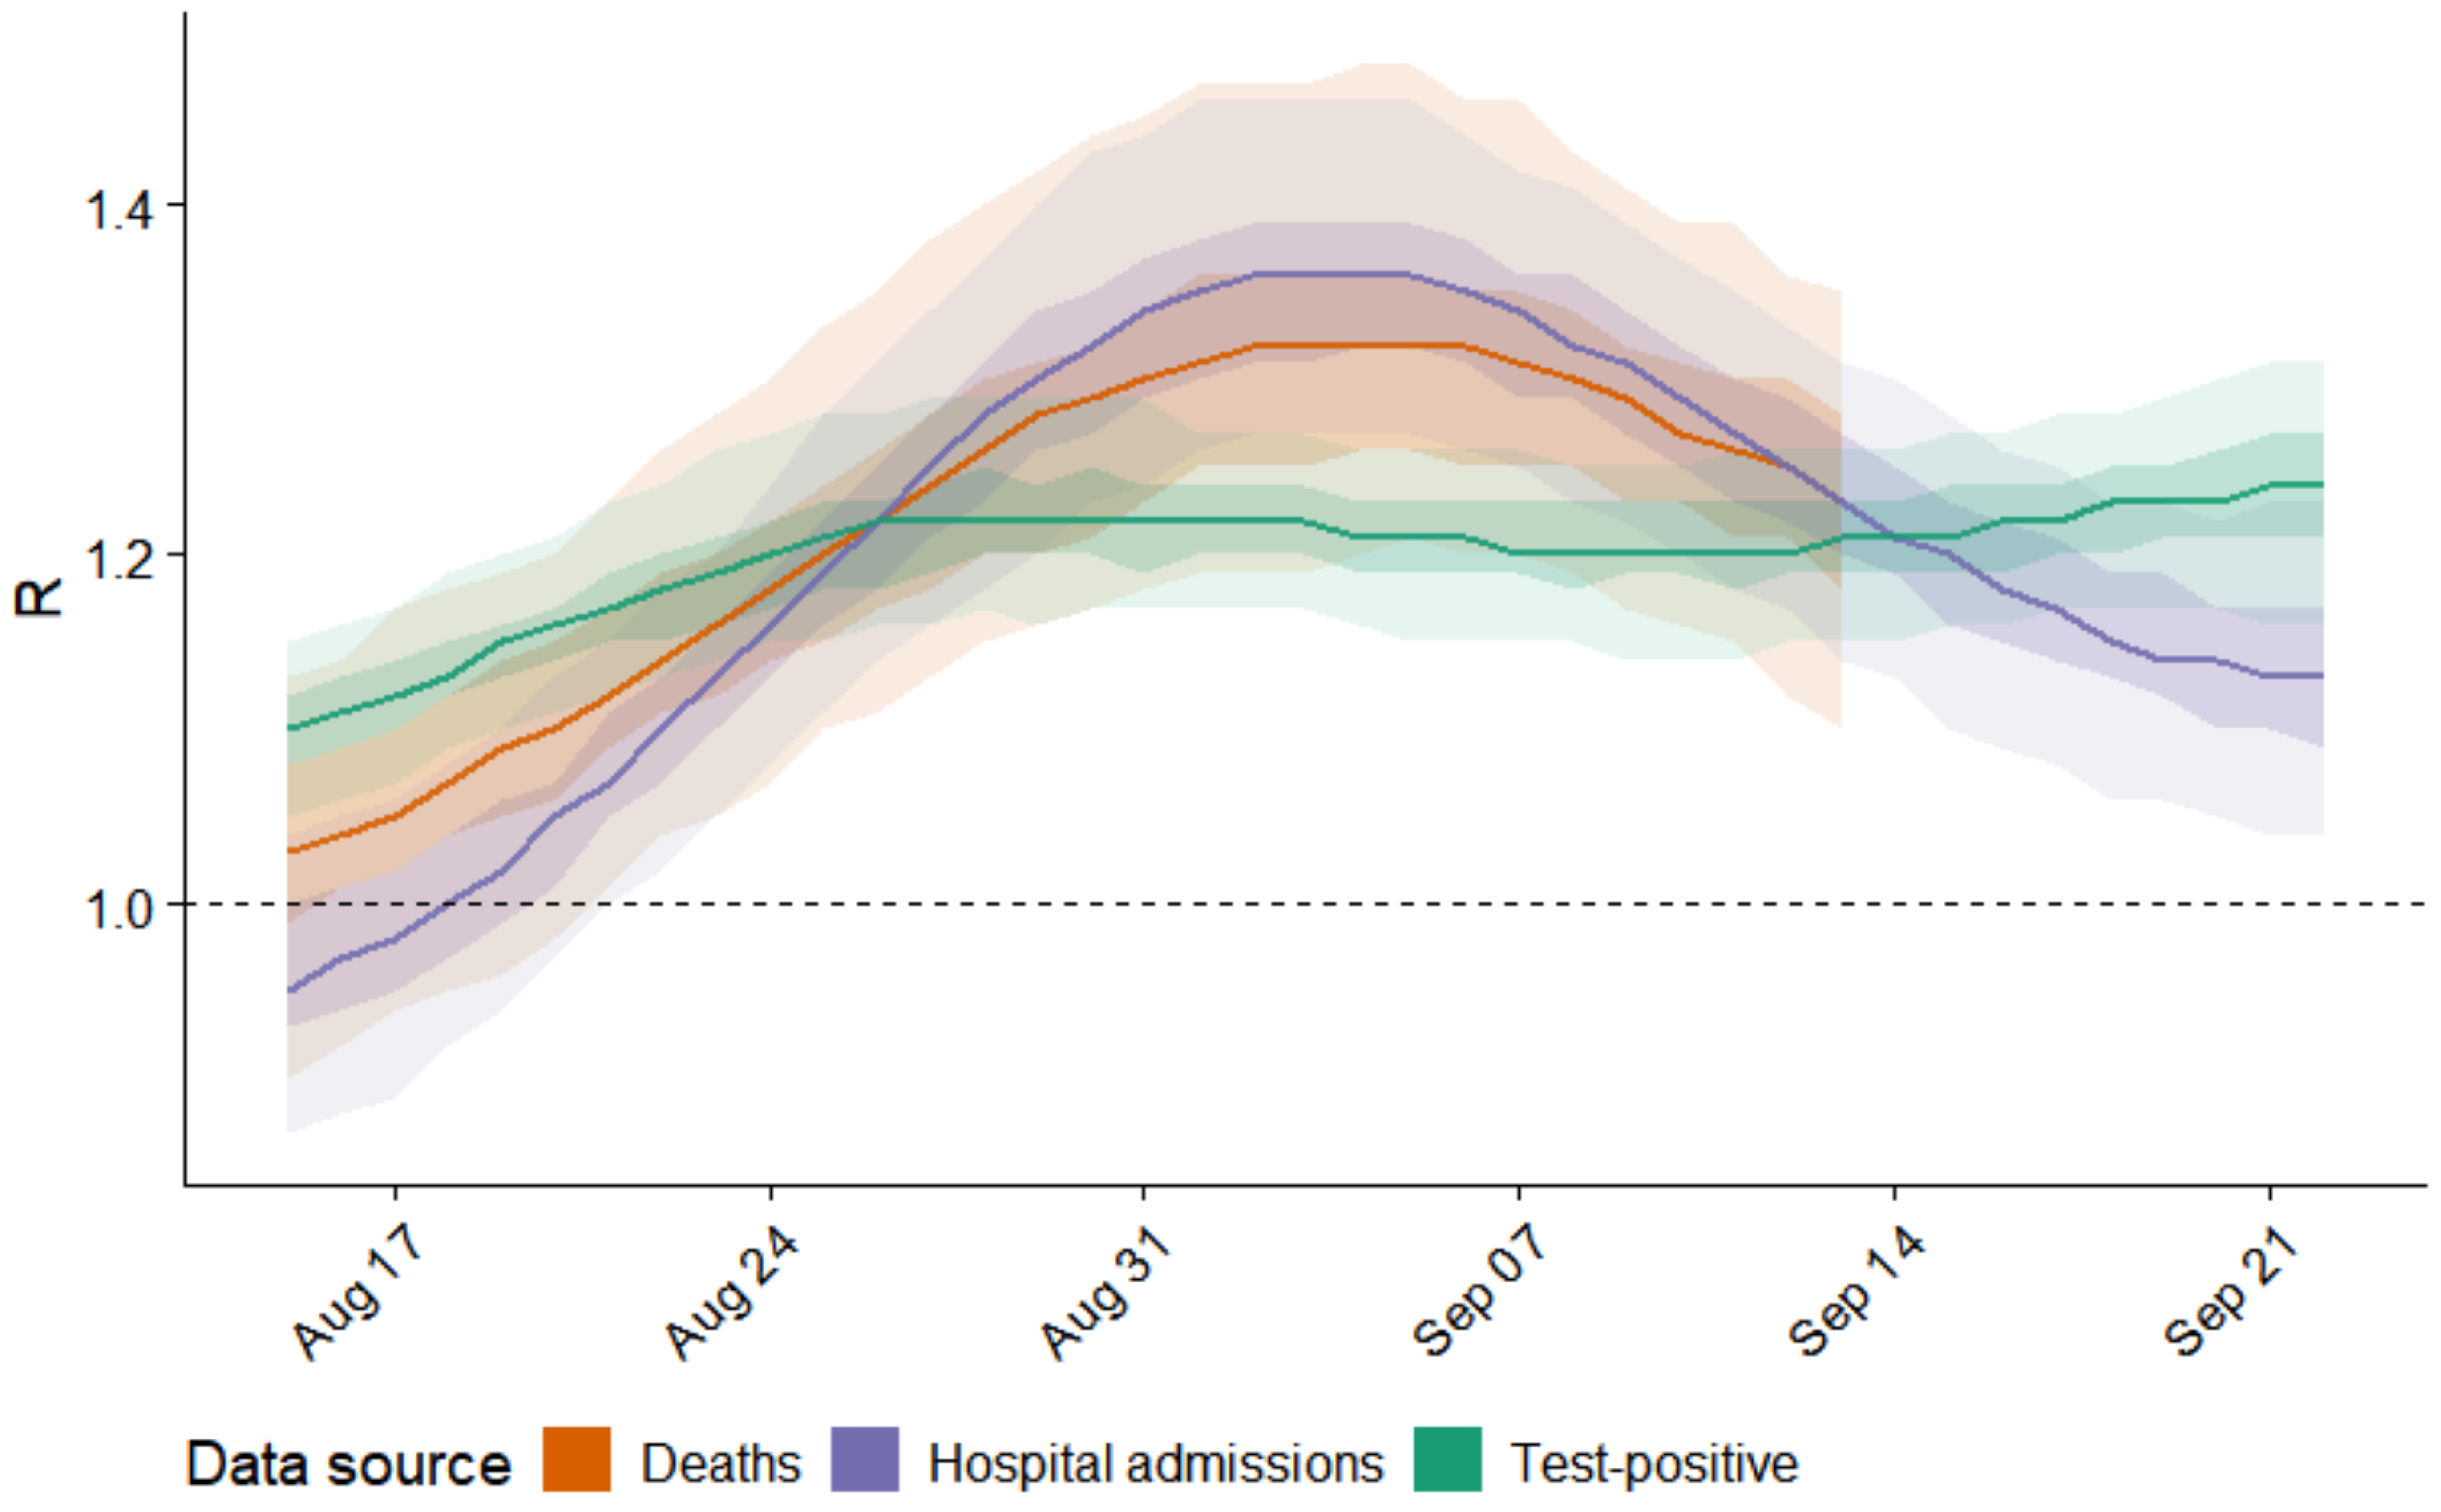

Sherratt et al., in prep

Many other ways of estimating R_t exist

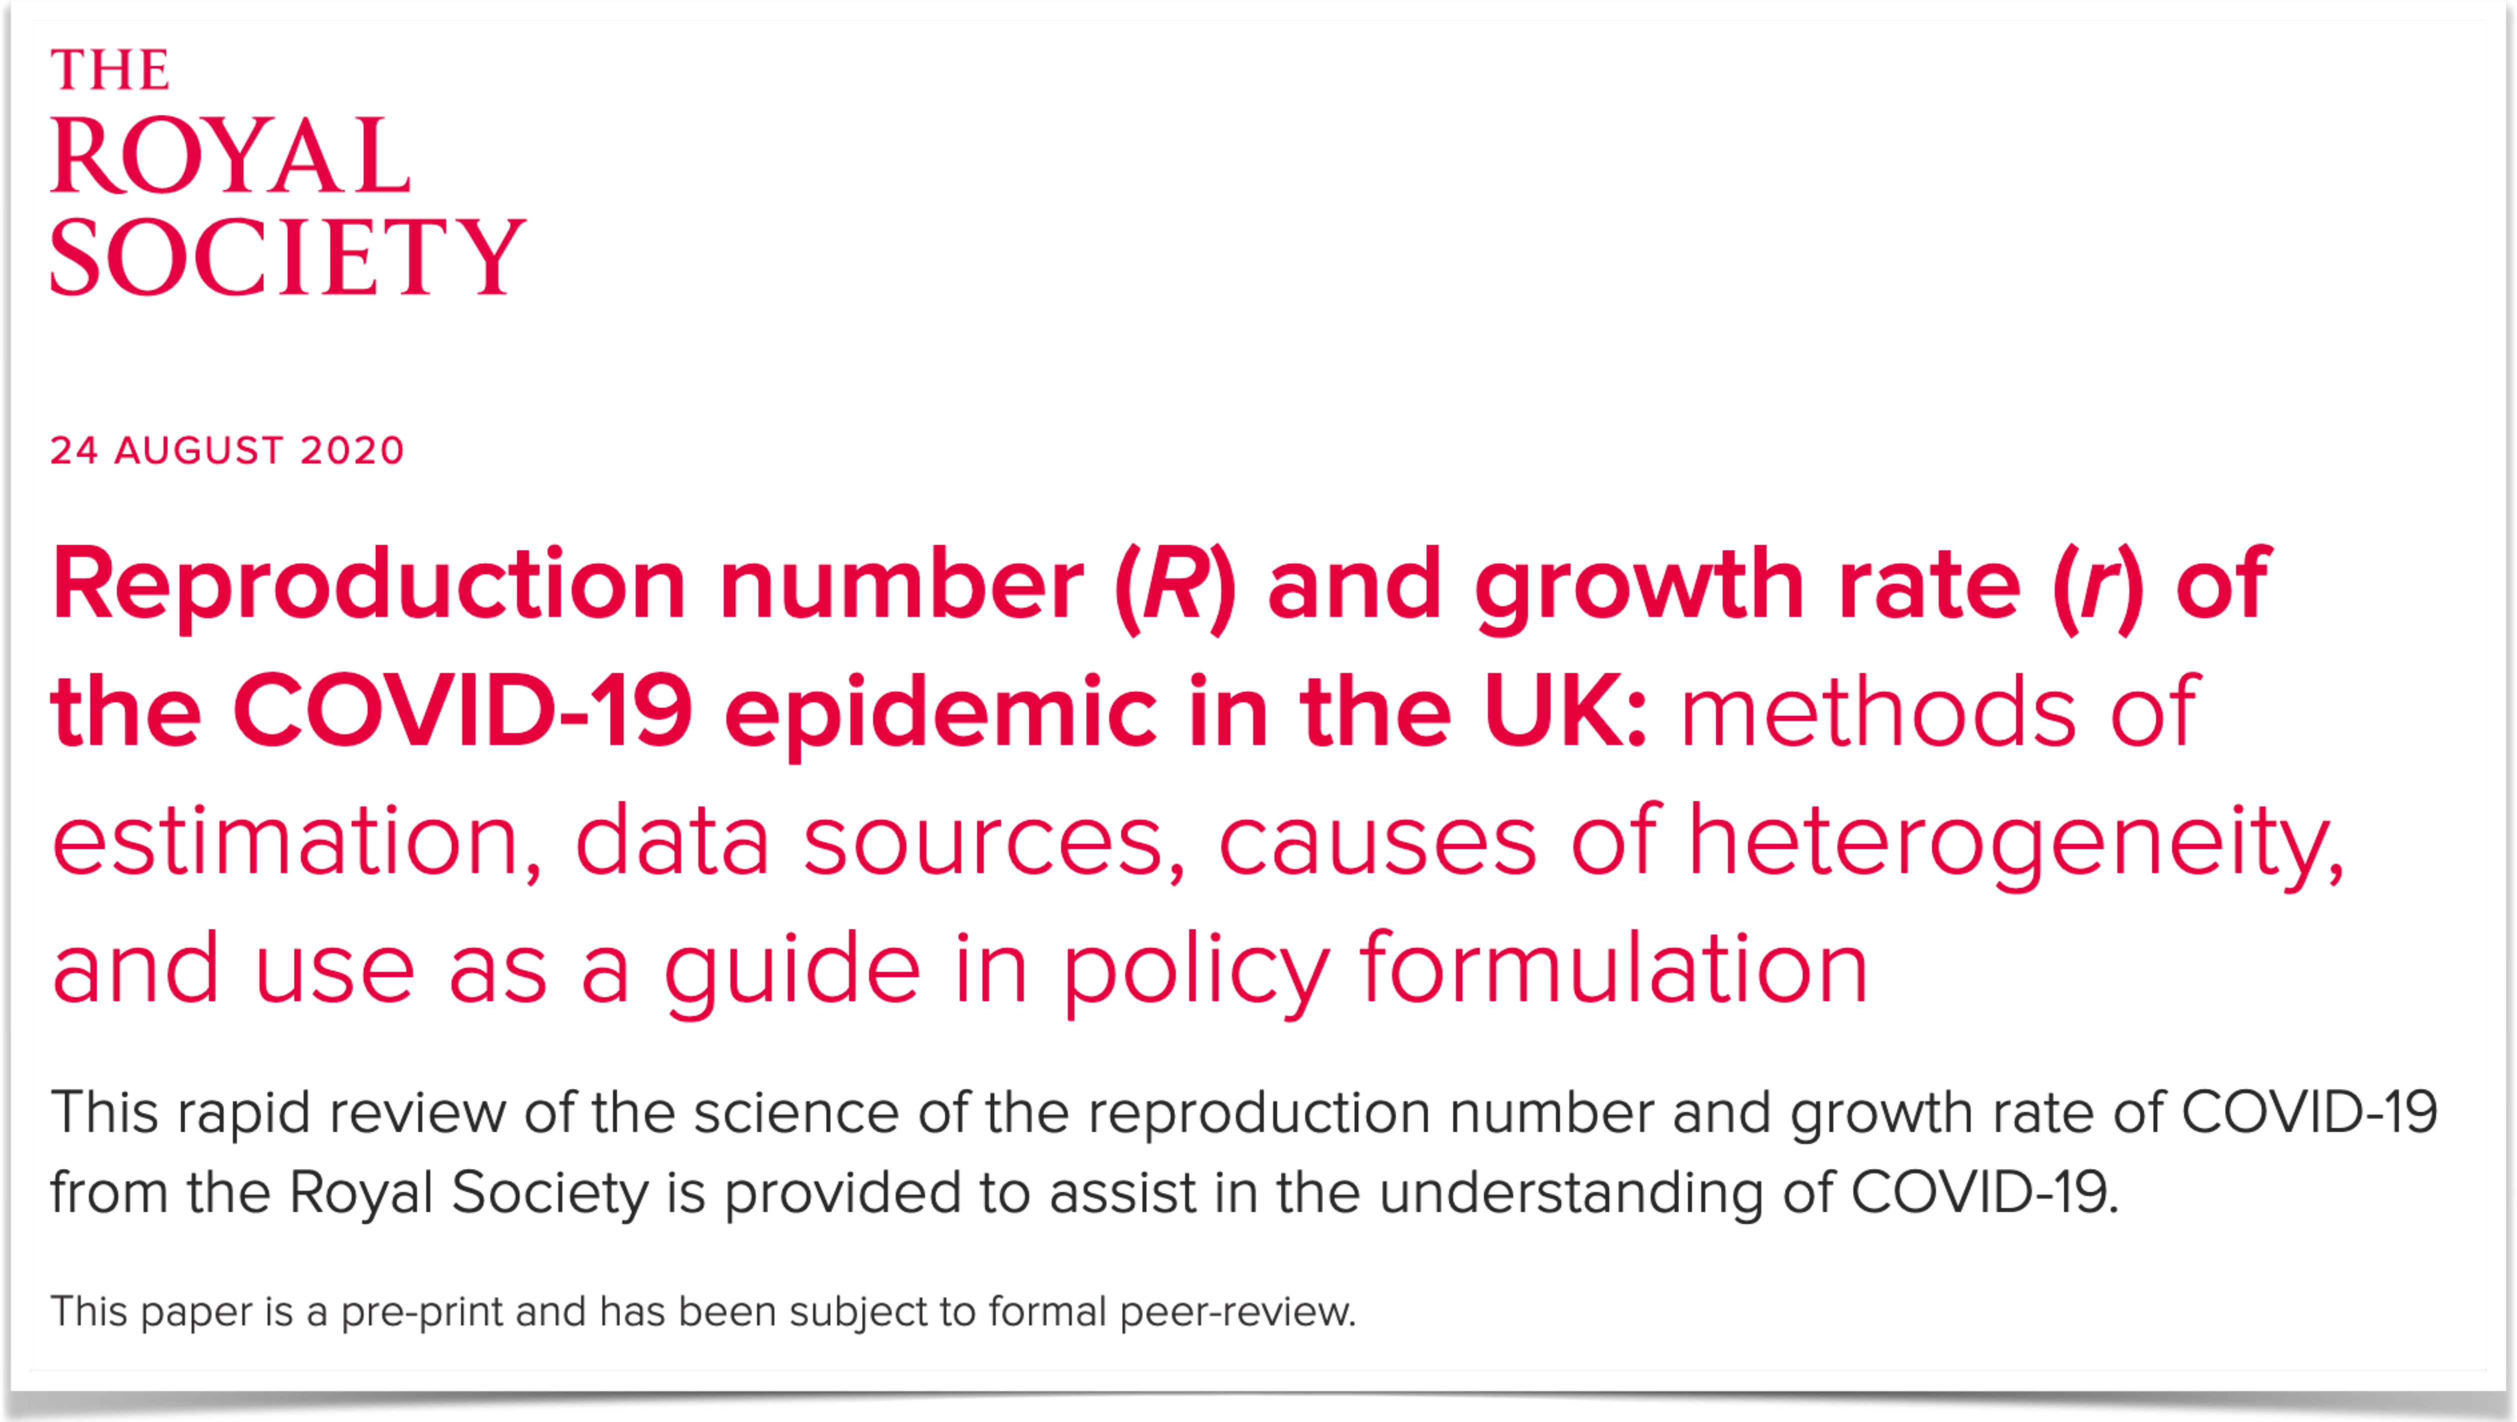

THE
ROYAL
SOCIETY

24 AUGUST 2020

Reproduction number (*R*) and growth rate (*r*) of the COVID-19 epidemic in the UK: methods of estimation, data sources, causes of heterogeneity, and use as a guide in policy formulation

This rapid review of the science of the reproduction number and growth rate of COVID-19 from the Royal Society is provided to assist in the understanding of COVID-19.

This paper is a pre-print and has been subject to formal peer-review.

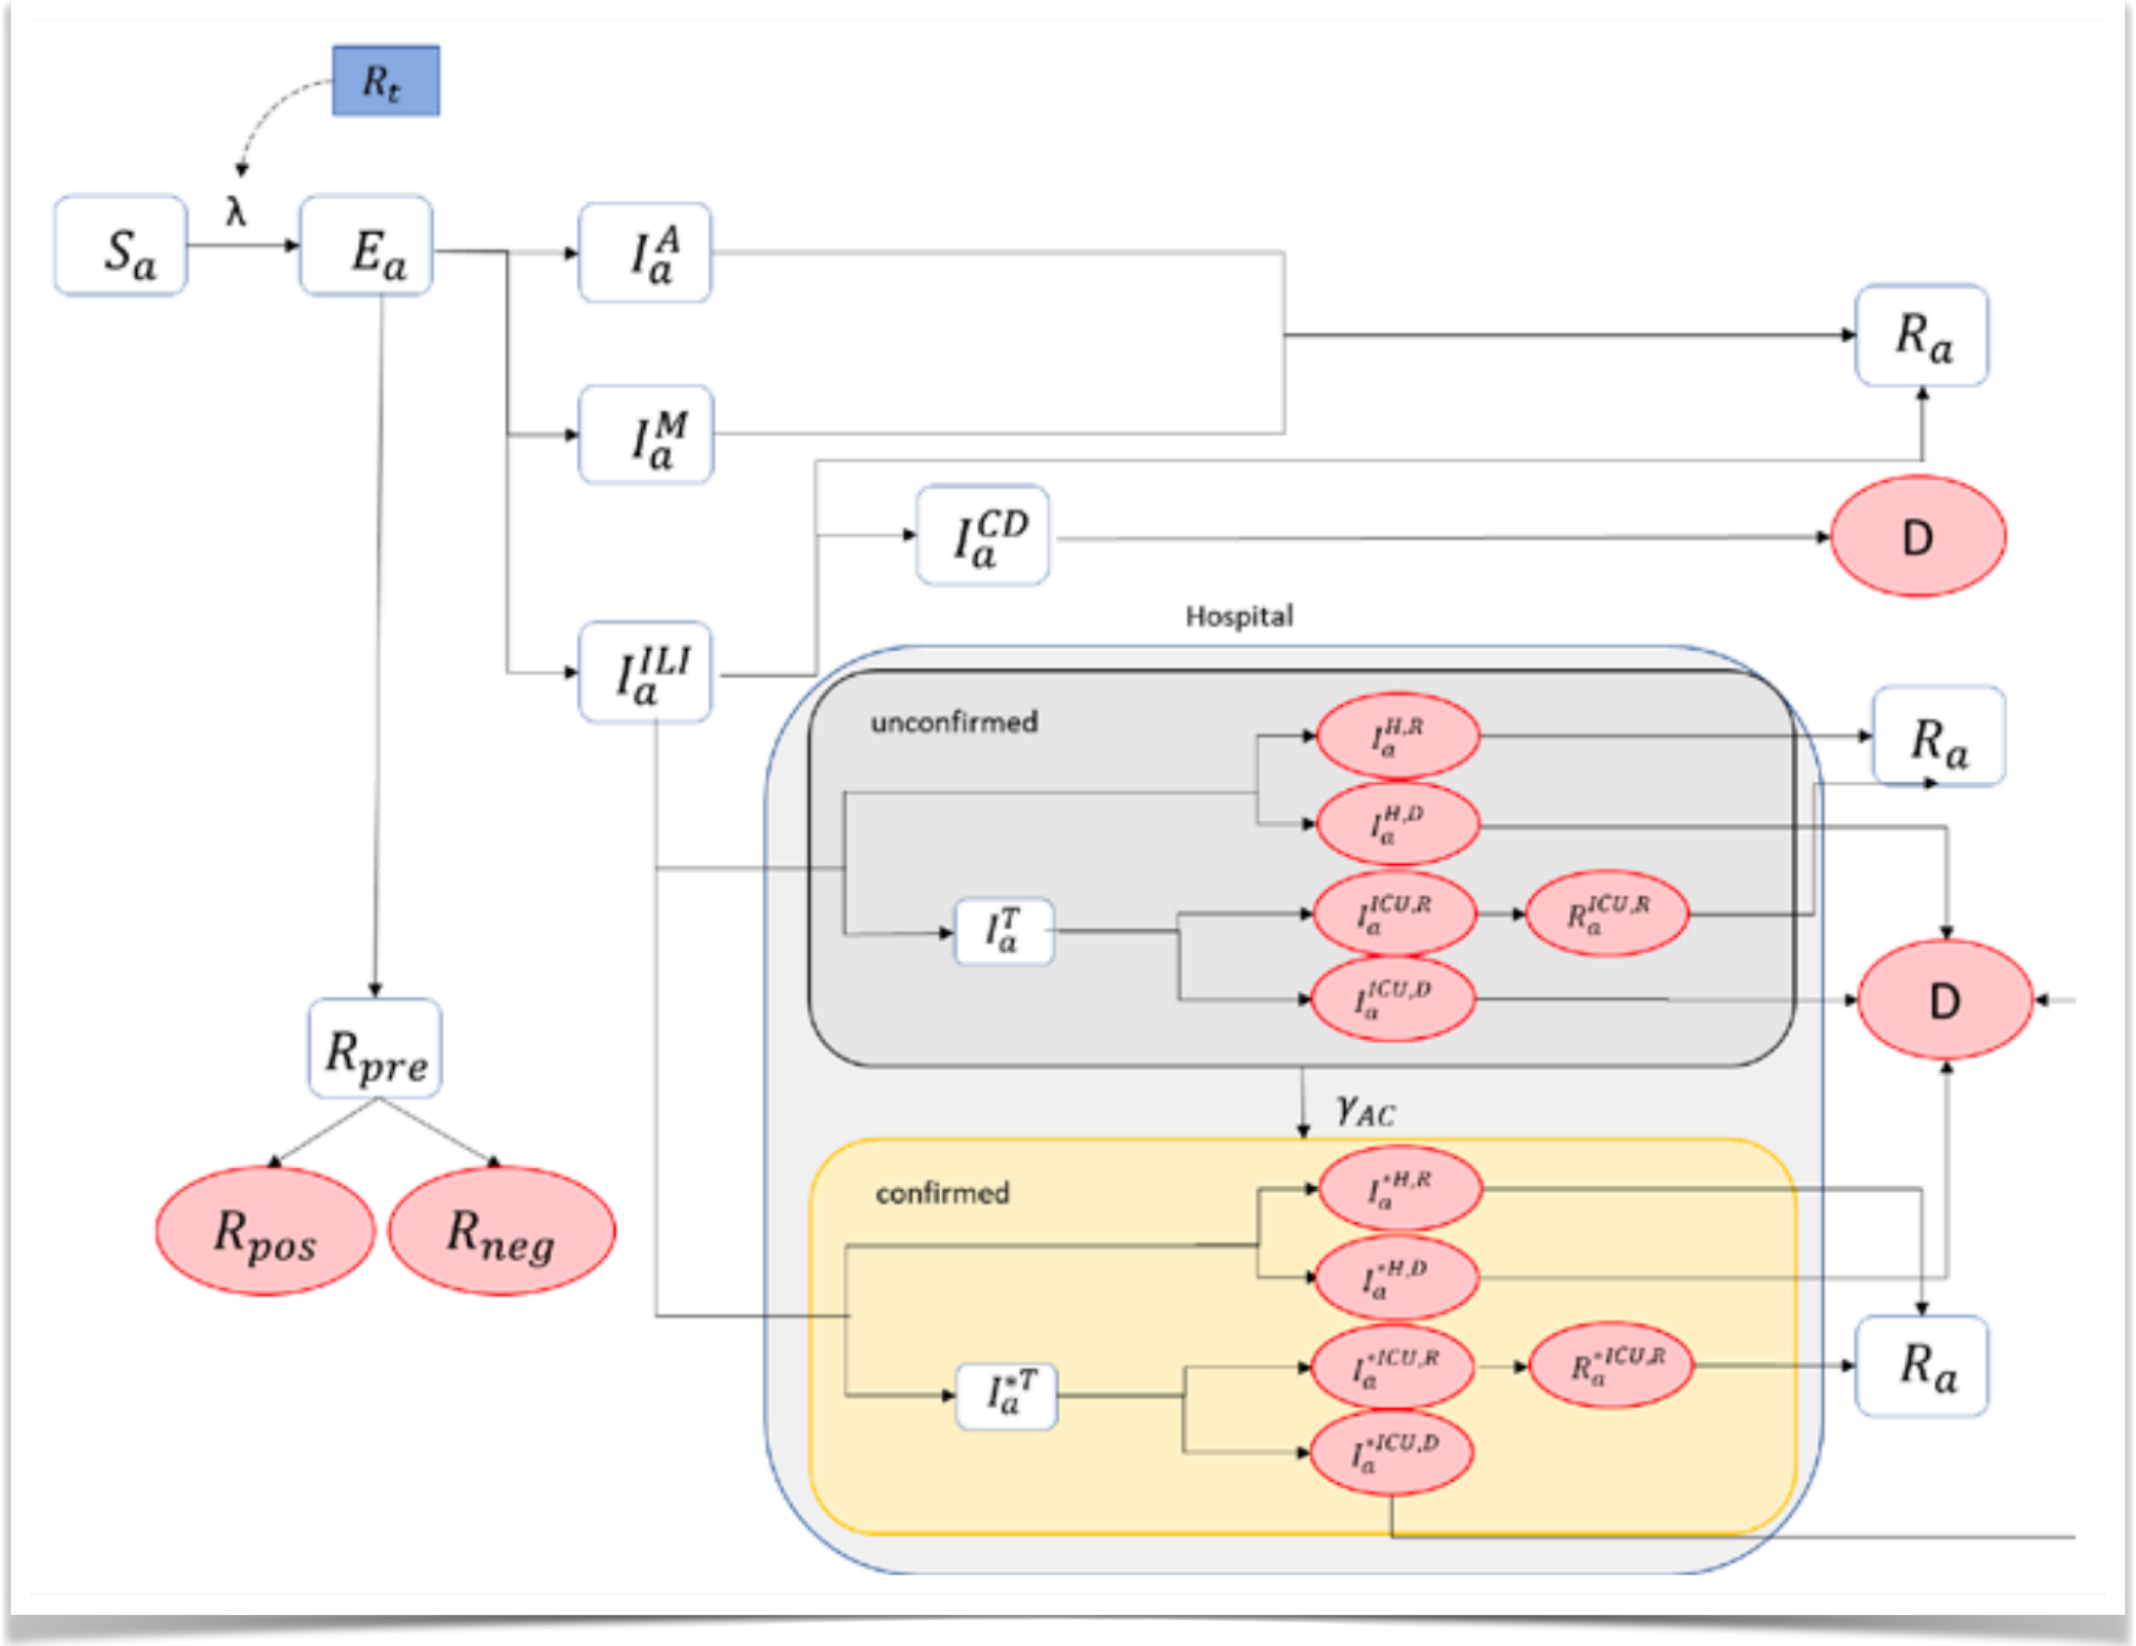

How can we tell which approach works best?

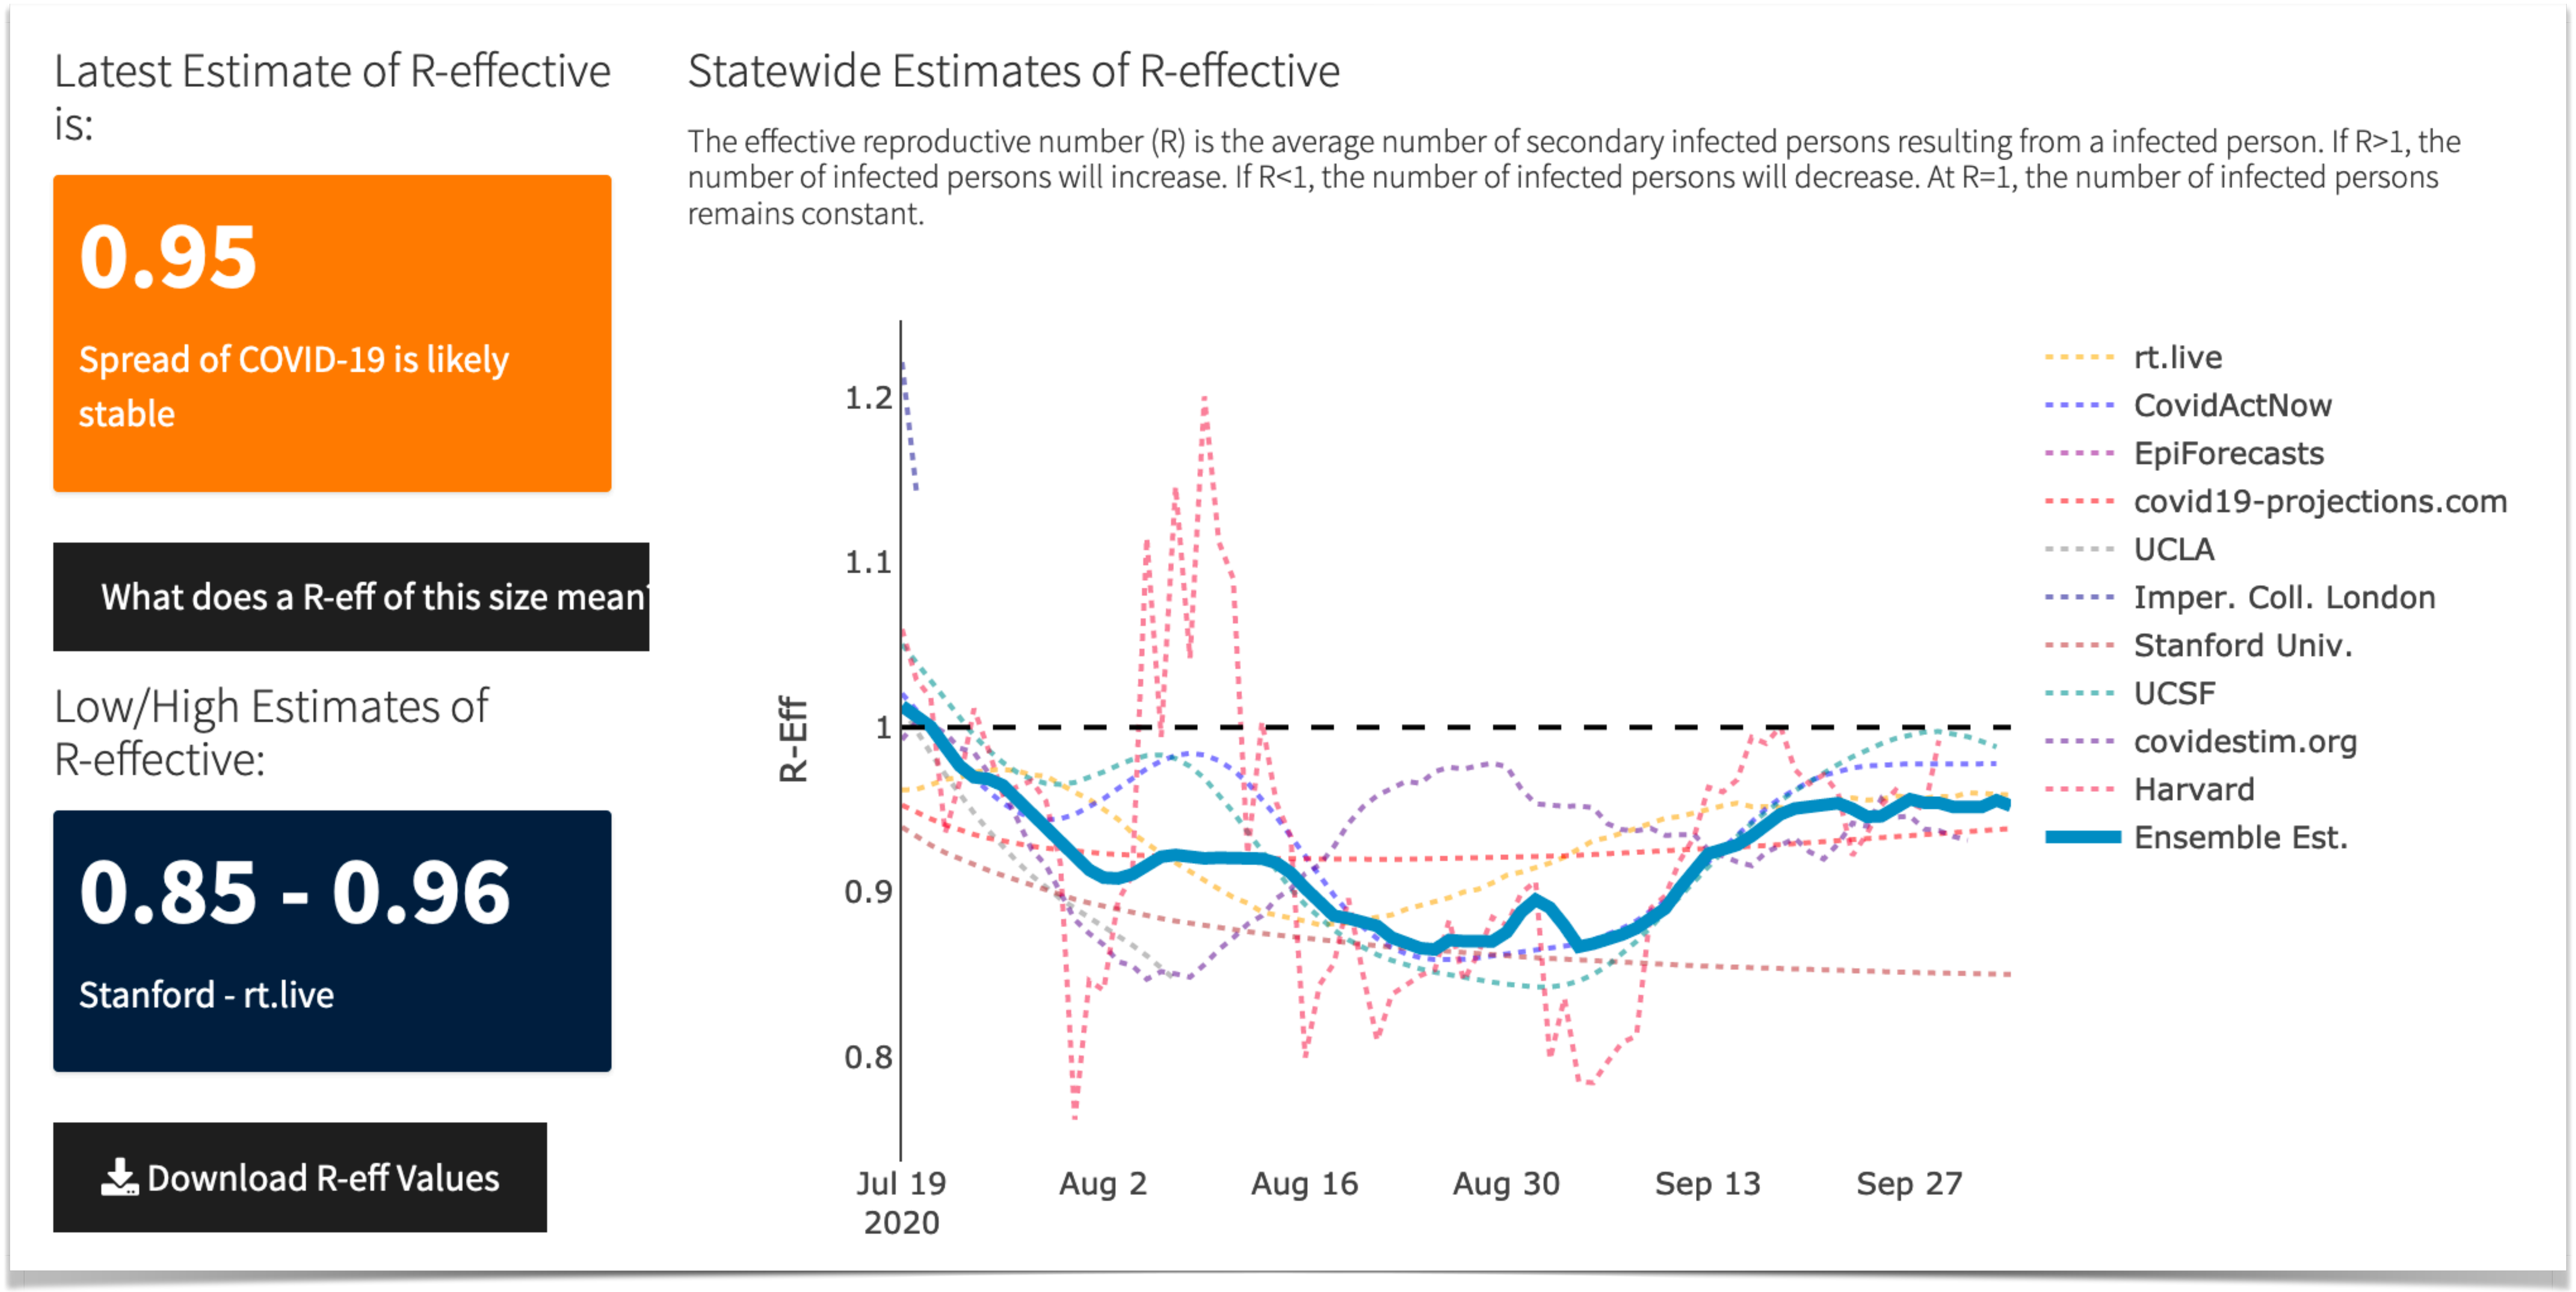

Latest Estimate of R-effective is:

0.95

Spread of COVID-19 is likely stable

What does a R-eff of this size mean

Low/High Estimates of R-effective:

0.85 - 0.96

Stanford - rt.live

Download R-eff Values

Statewide Estimates of R-effective

The effective reproductive number (R) is the average number of secondary infected persons resulting from a infected person. If R>1, the number of infected persons will increase. If R<1, the number of infected persons will decrease. At R=1, the number of infected persons remains constant.

https://calcat.covid19.ca.gov/cacovidmodels/

Short-term forecasts for Covid-19

Short-term forecasts based on R_t estimation

Ebola in Western Area, Sierra Leone (2015)

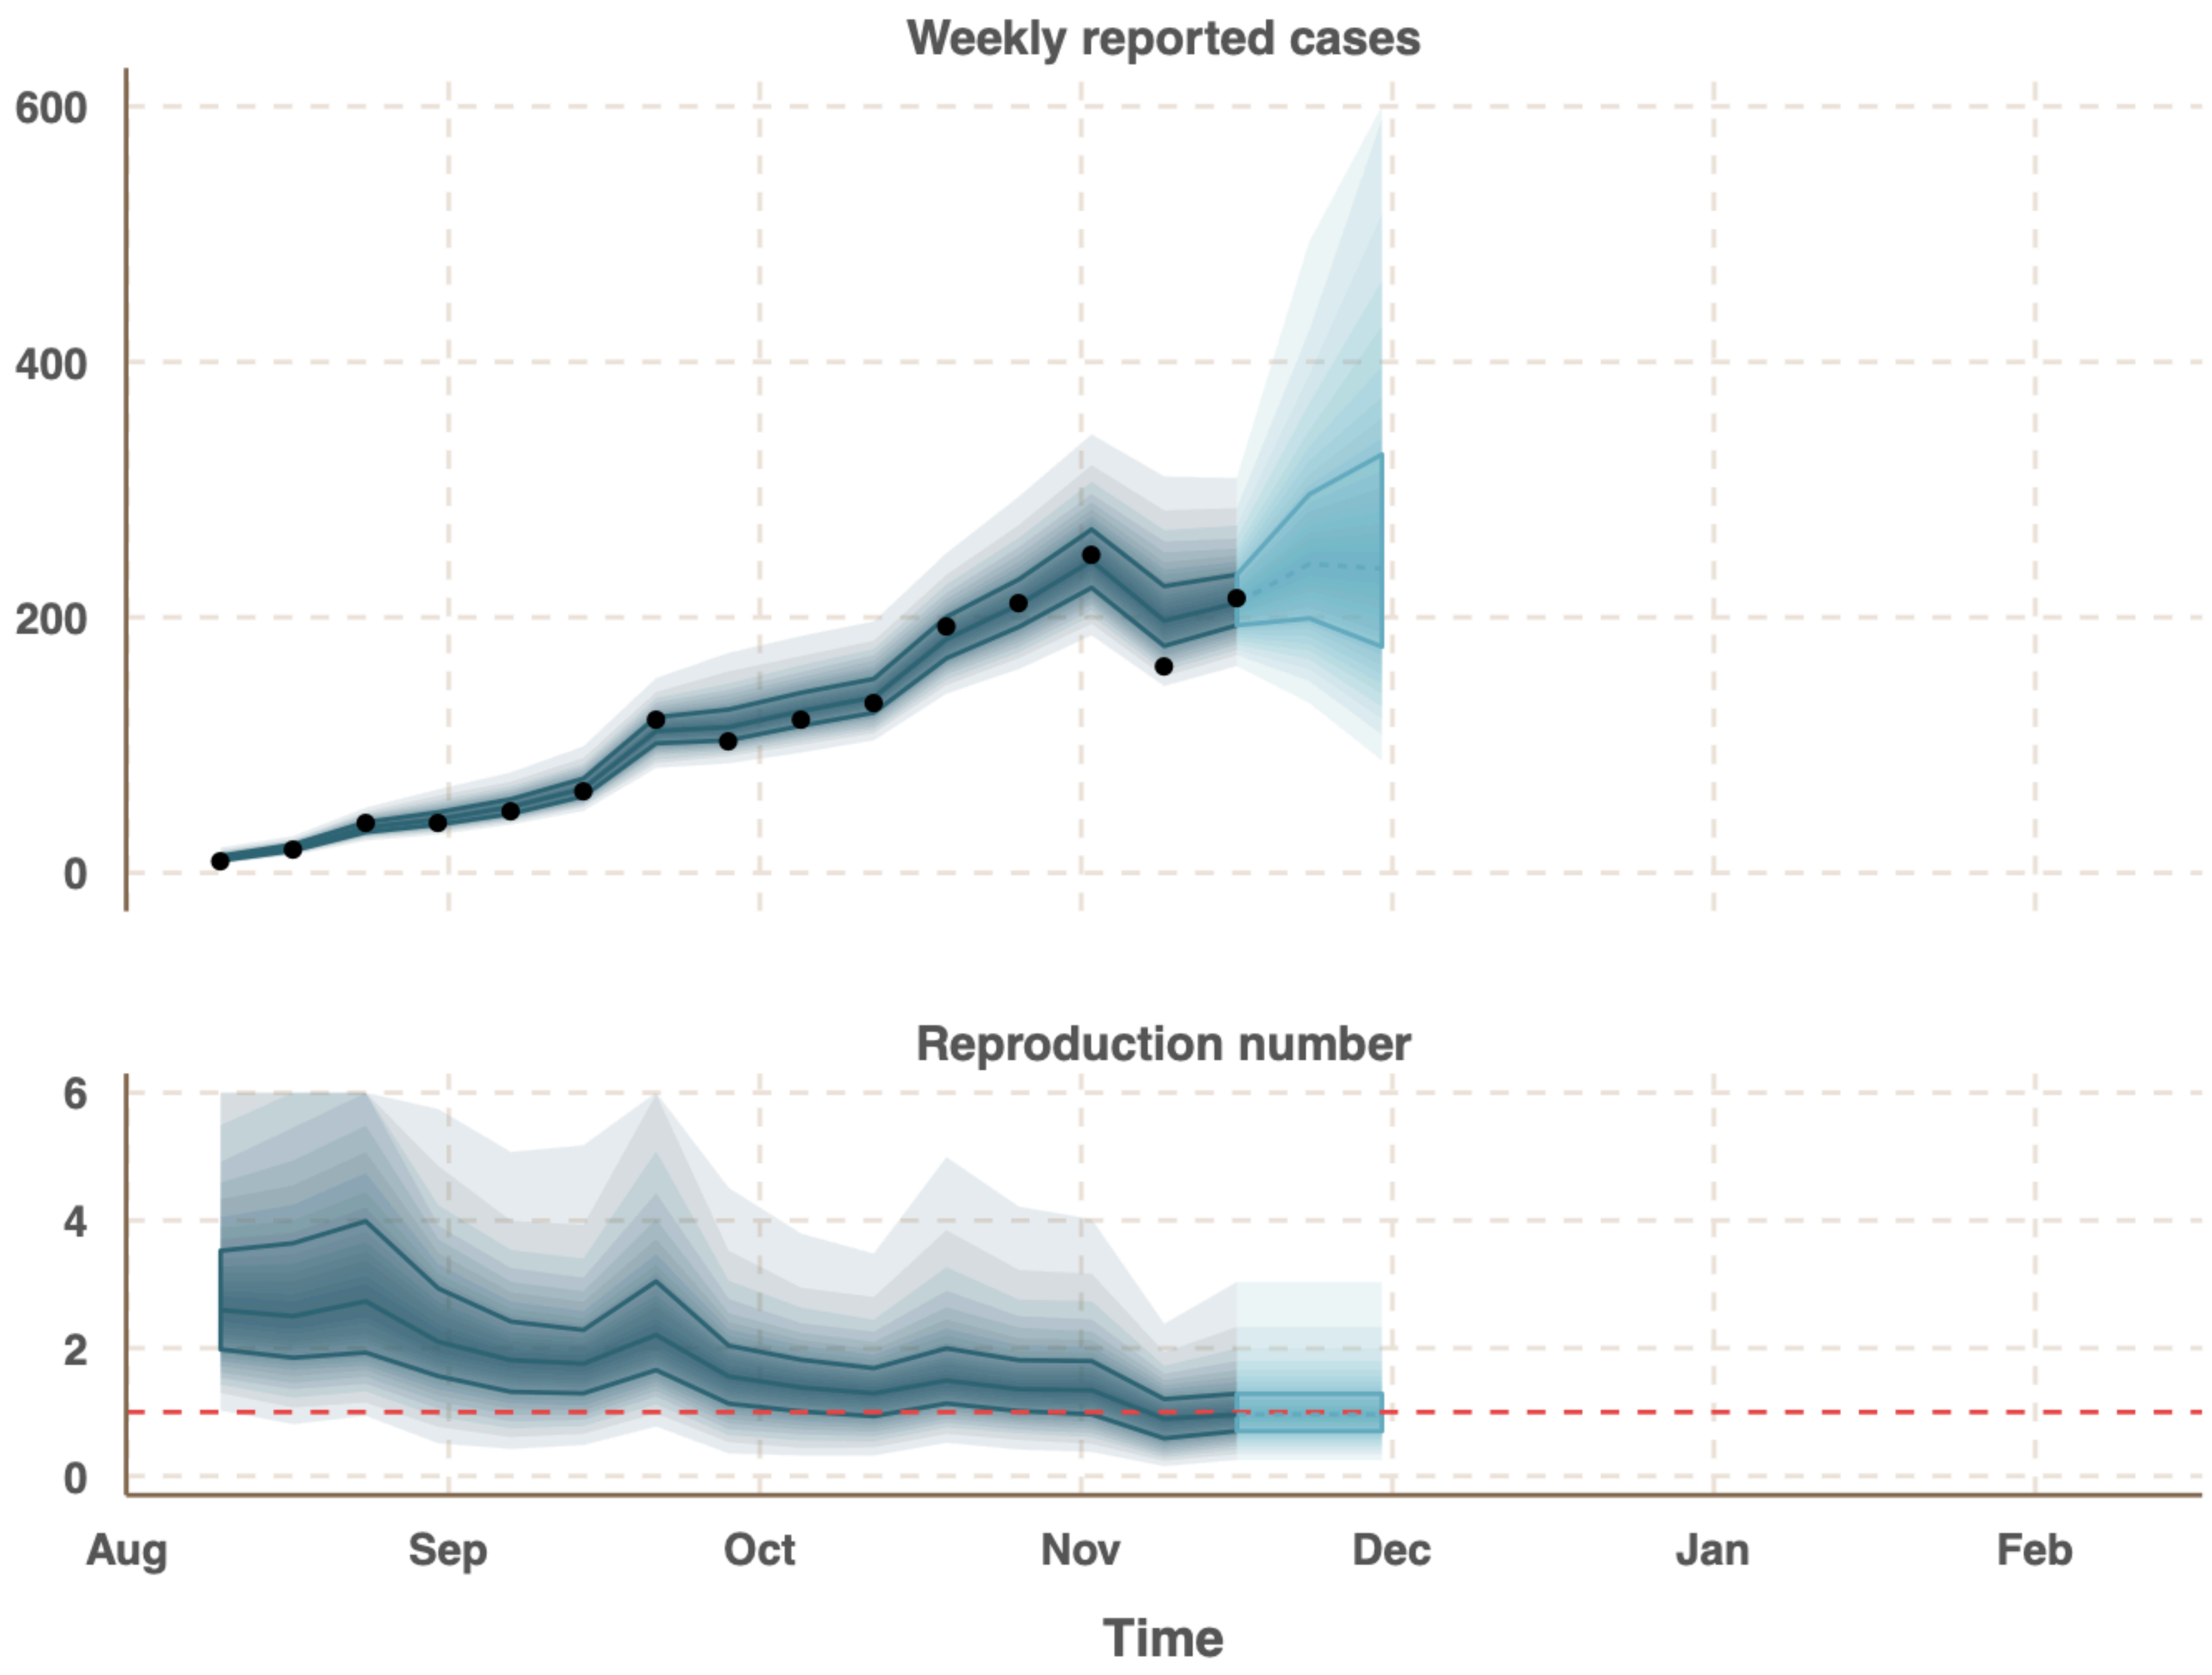

Camacho et al., PLOS Currents, 2015

Short-term forecasts based on R_t estimation

Colombia

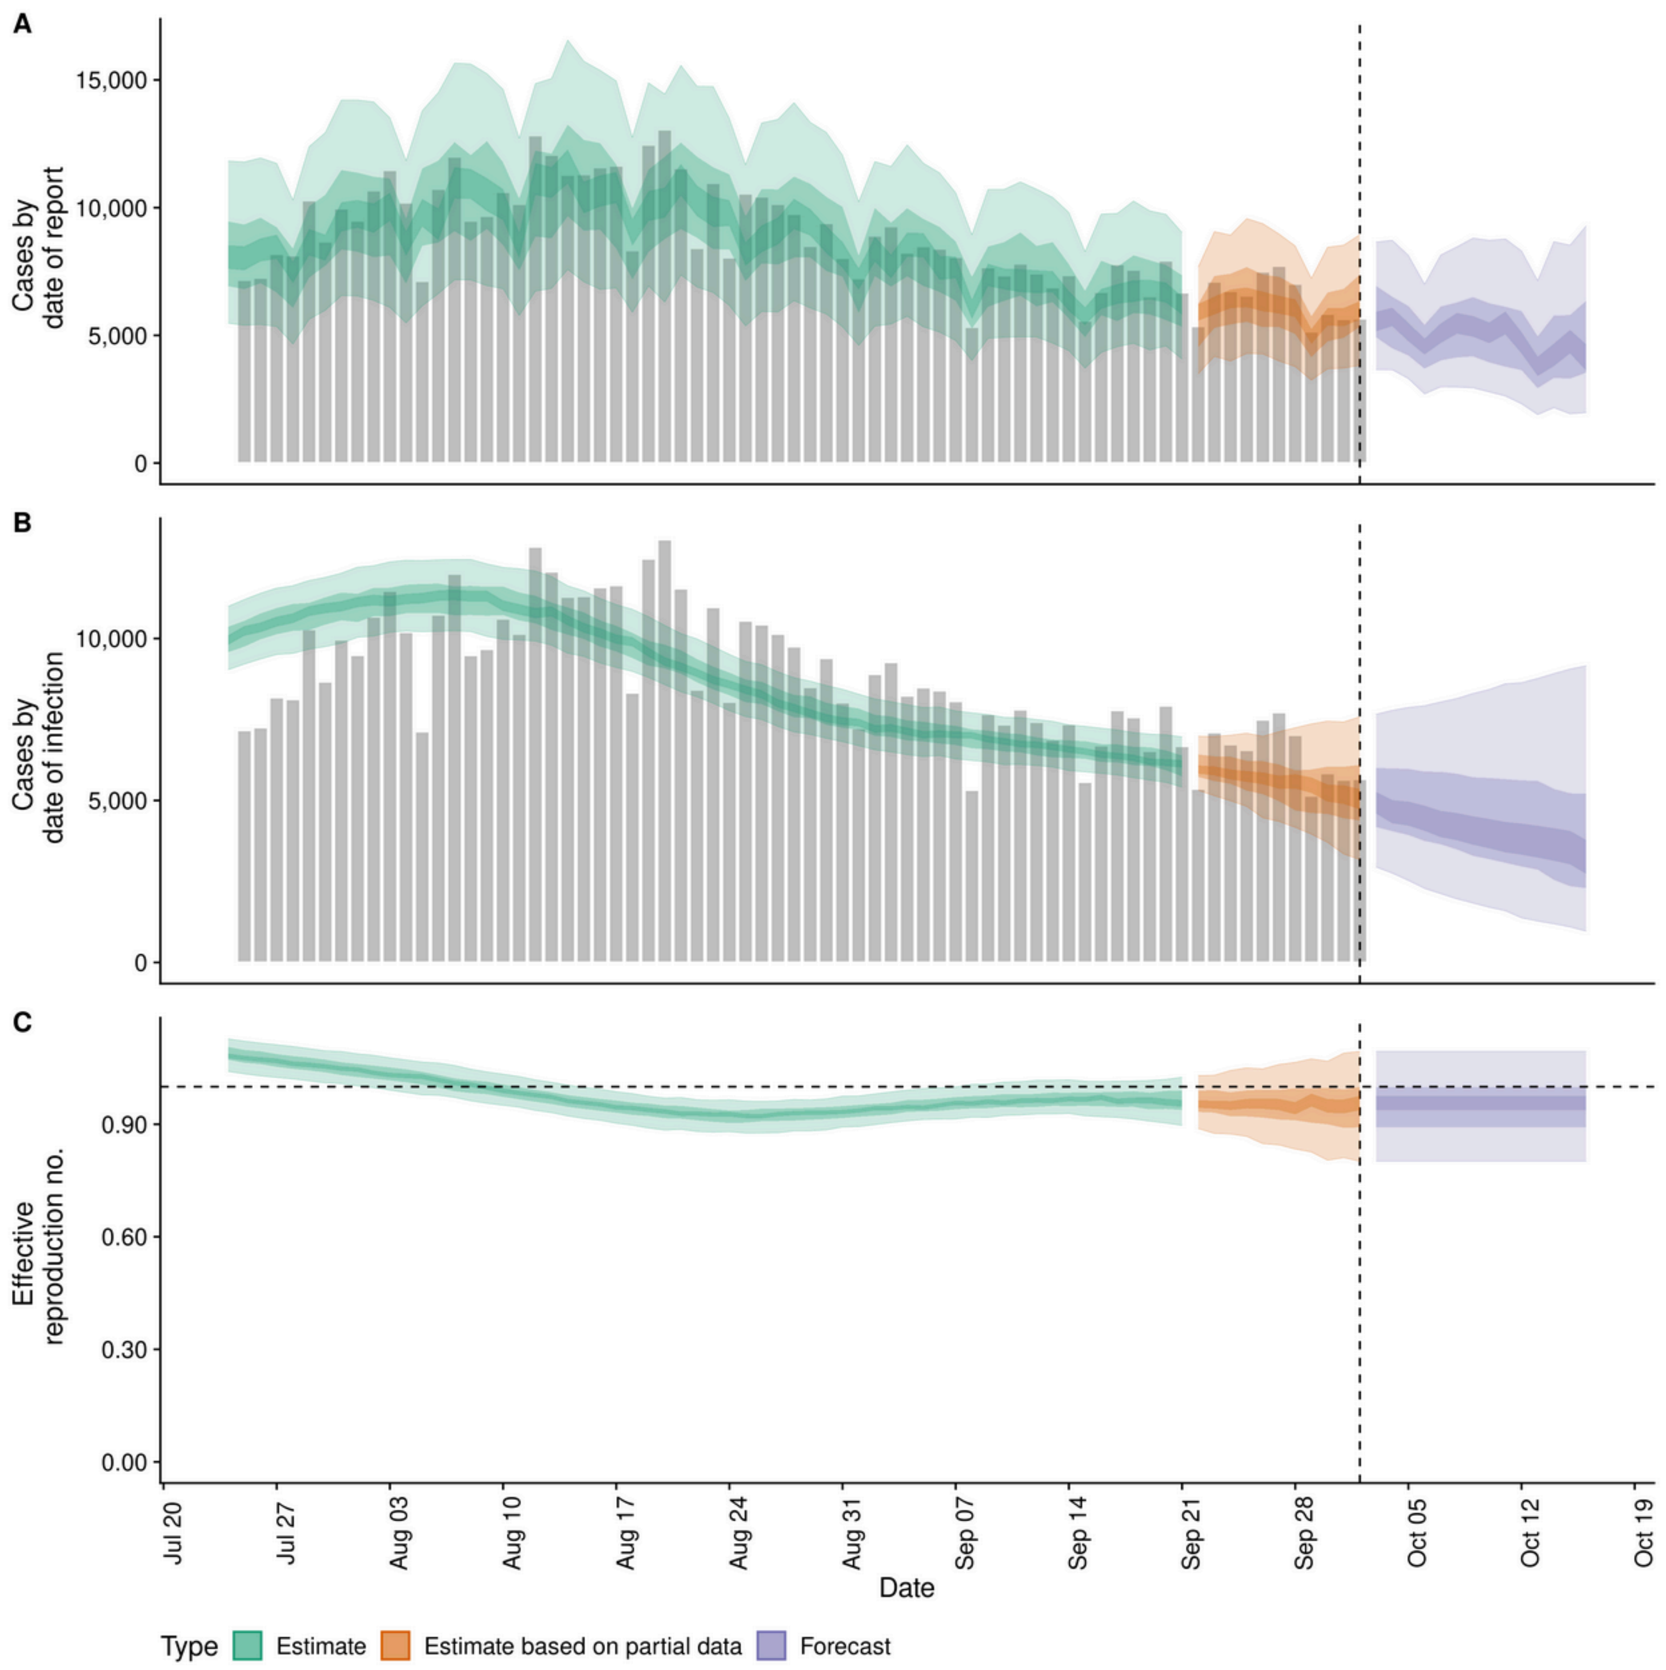

Confirmed cases, their estimated date of report, date of infection, and time-varying reproduction number estimates

https://epiforecasts.io/covid/posts/national/colombia/

Italy

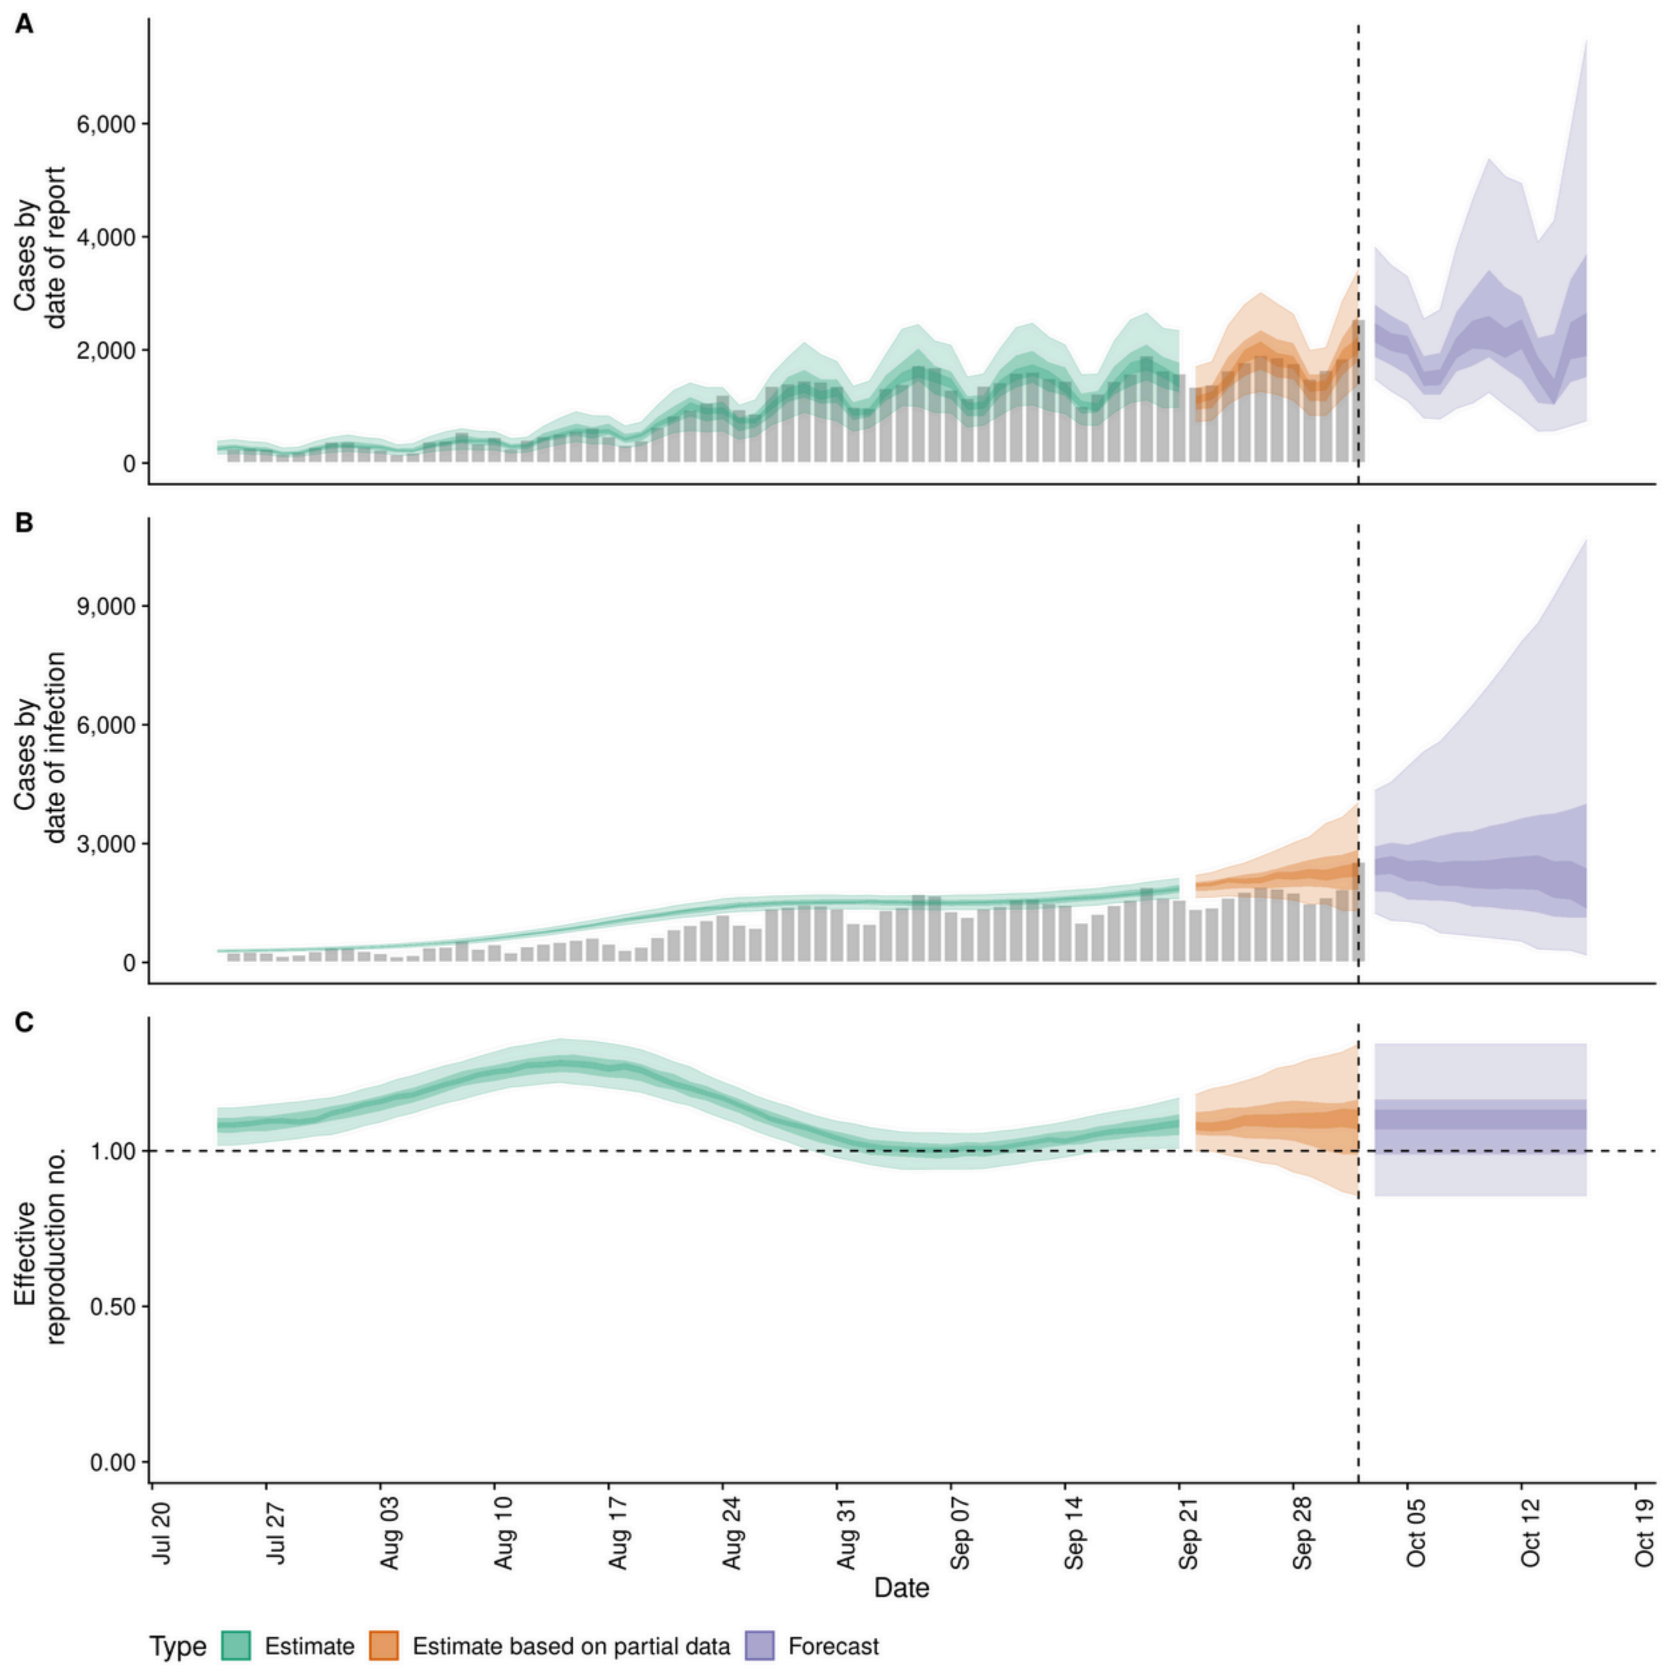

Confirmed cases, their estimated date of report, date of infection, and time-varying reproduction number estimates

https://epiforecasts.io/covid/posts/national/italy/

Forecasts can be assessed / validated

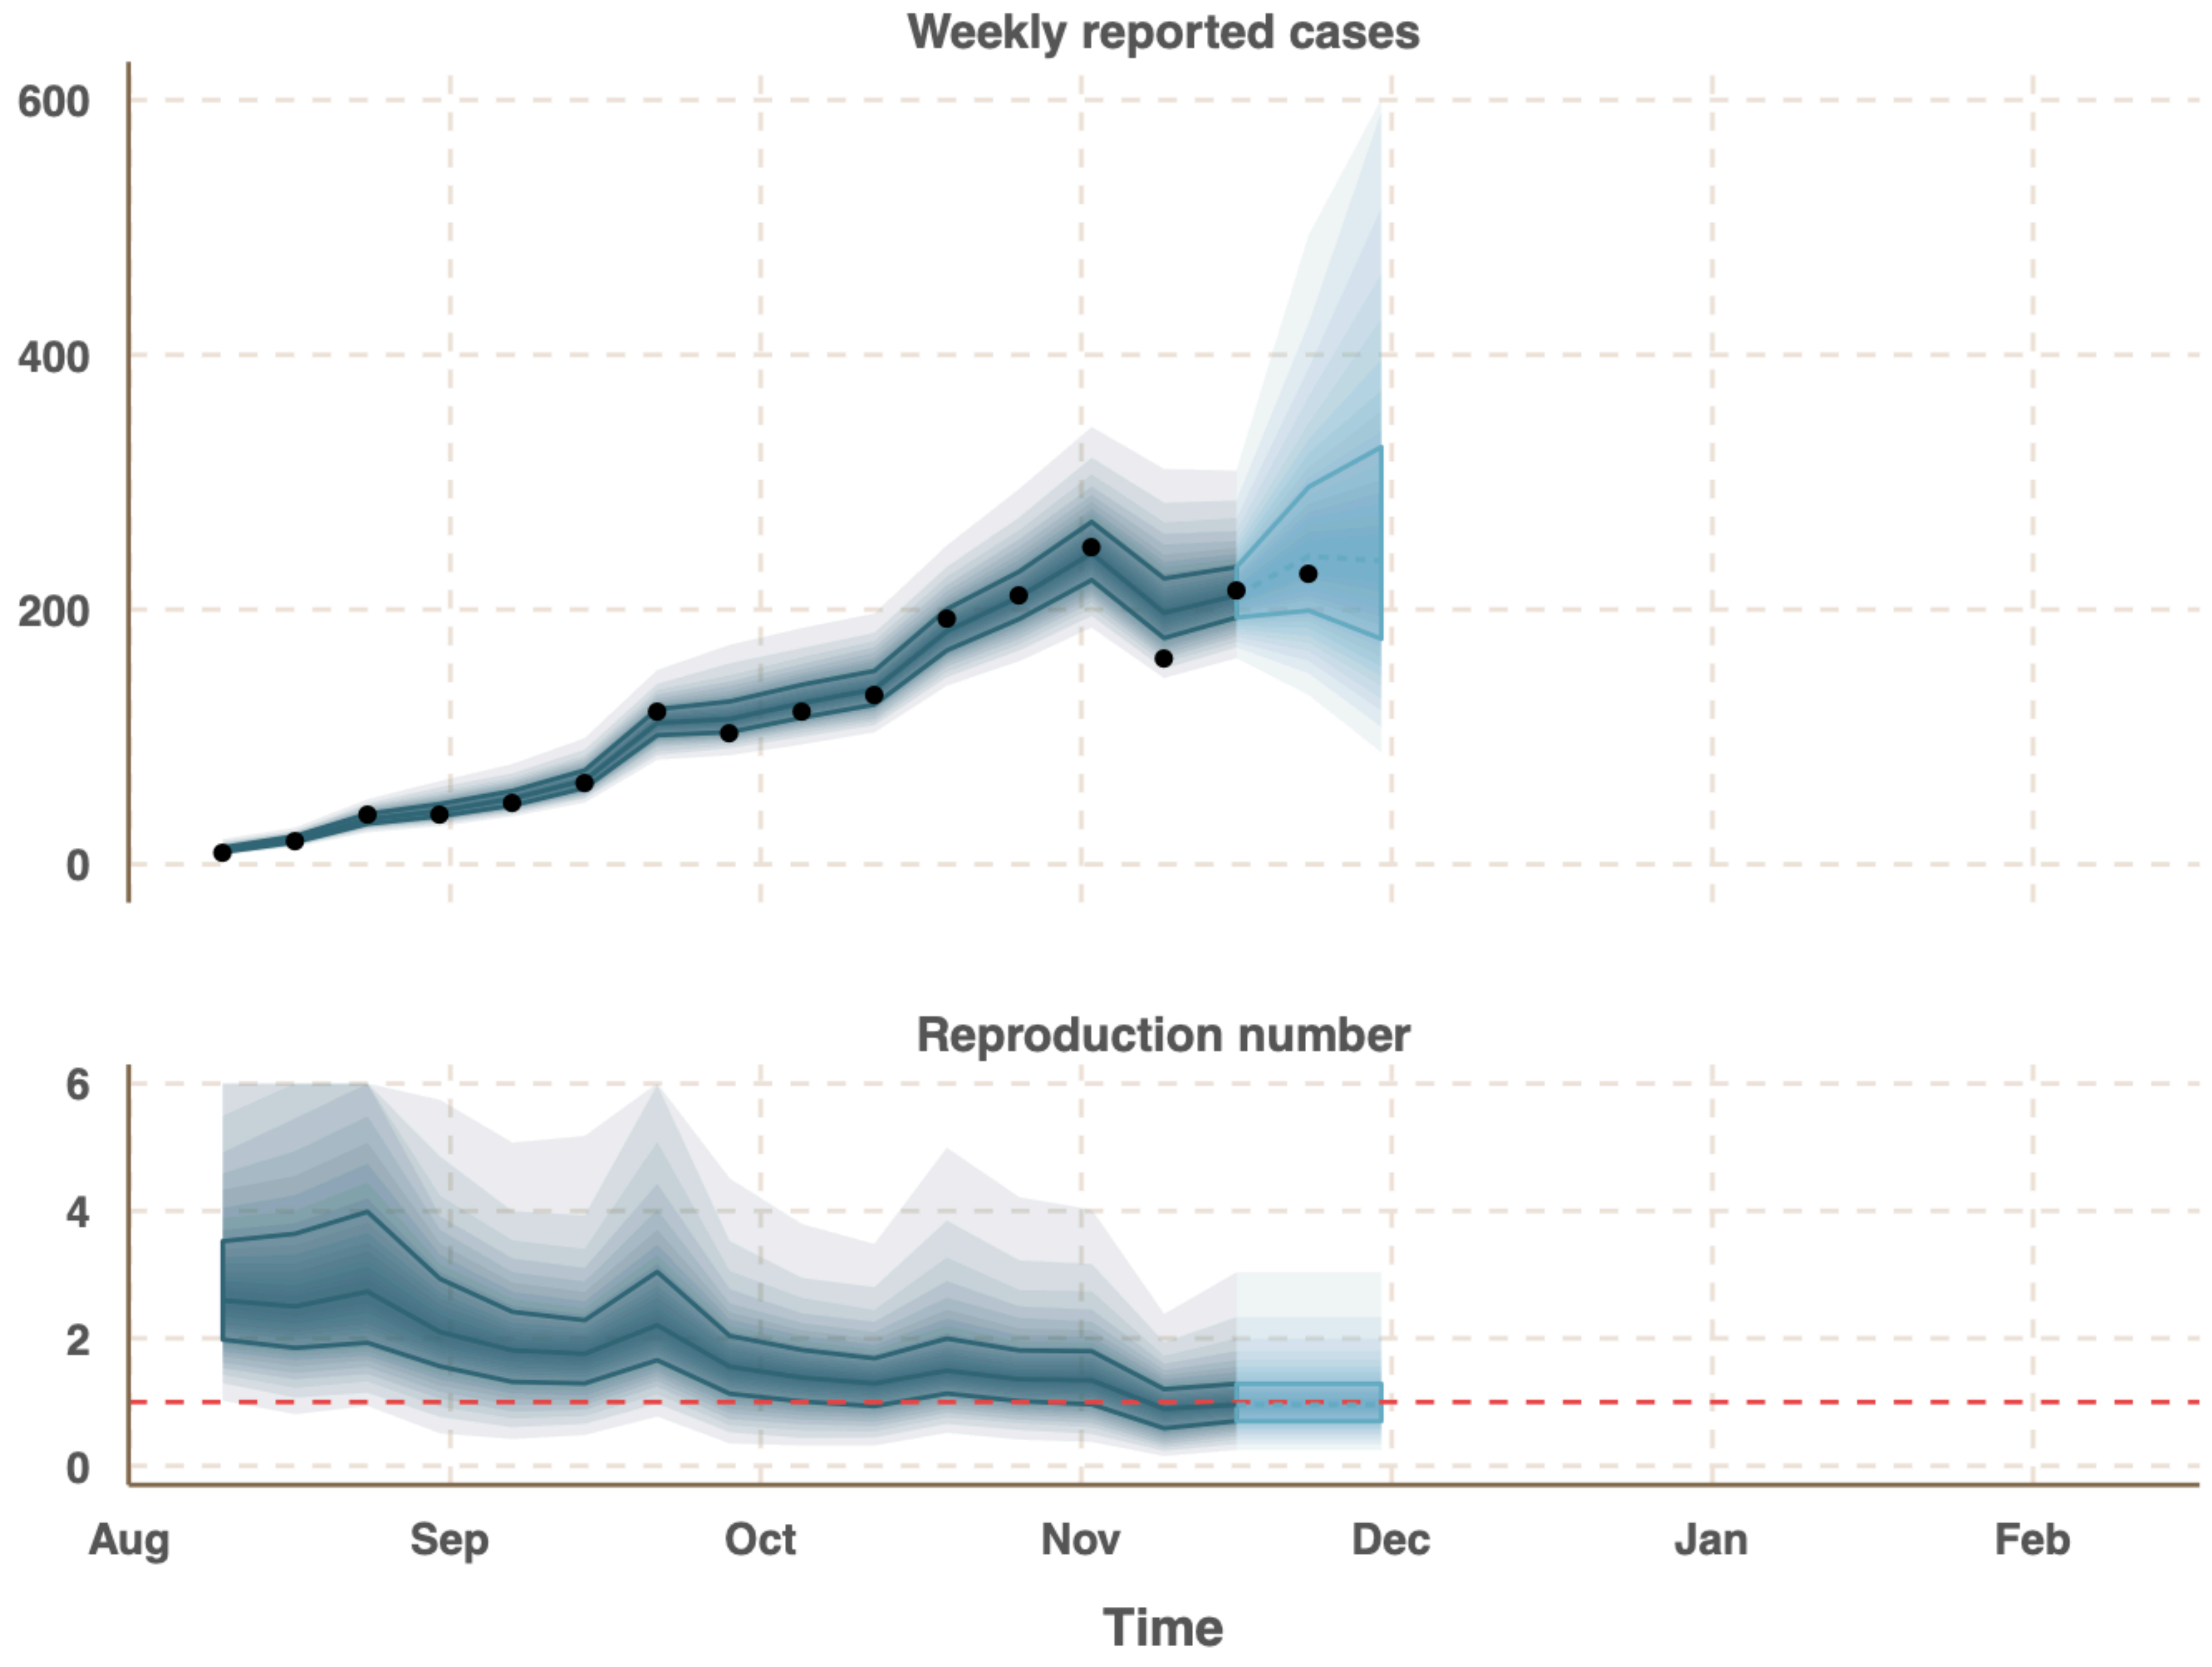

Funk et al., PLOS Comp Biol, 2019

Forecasts from multiple models

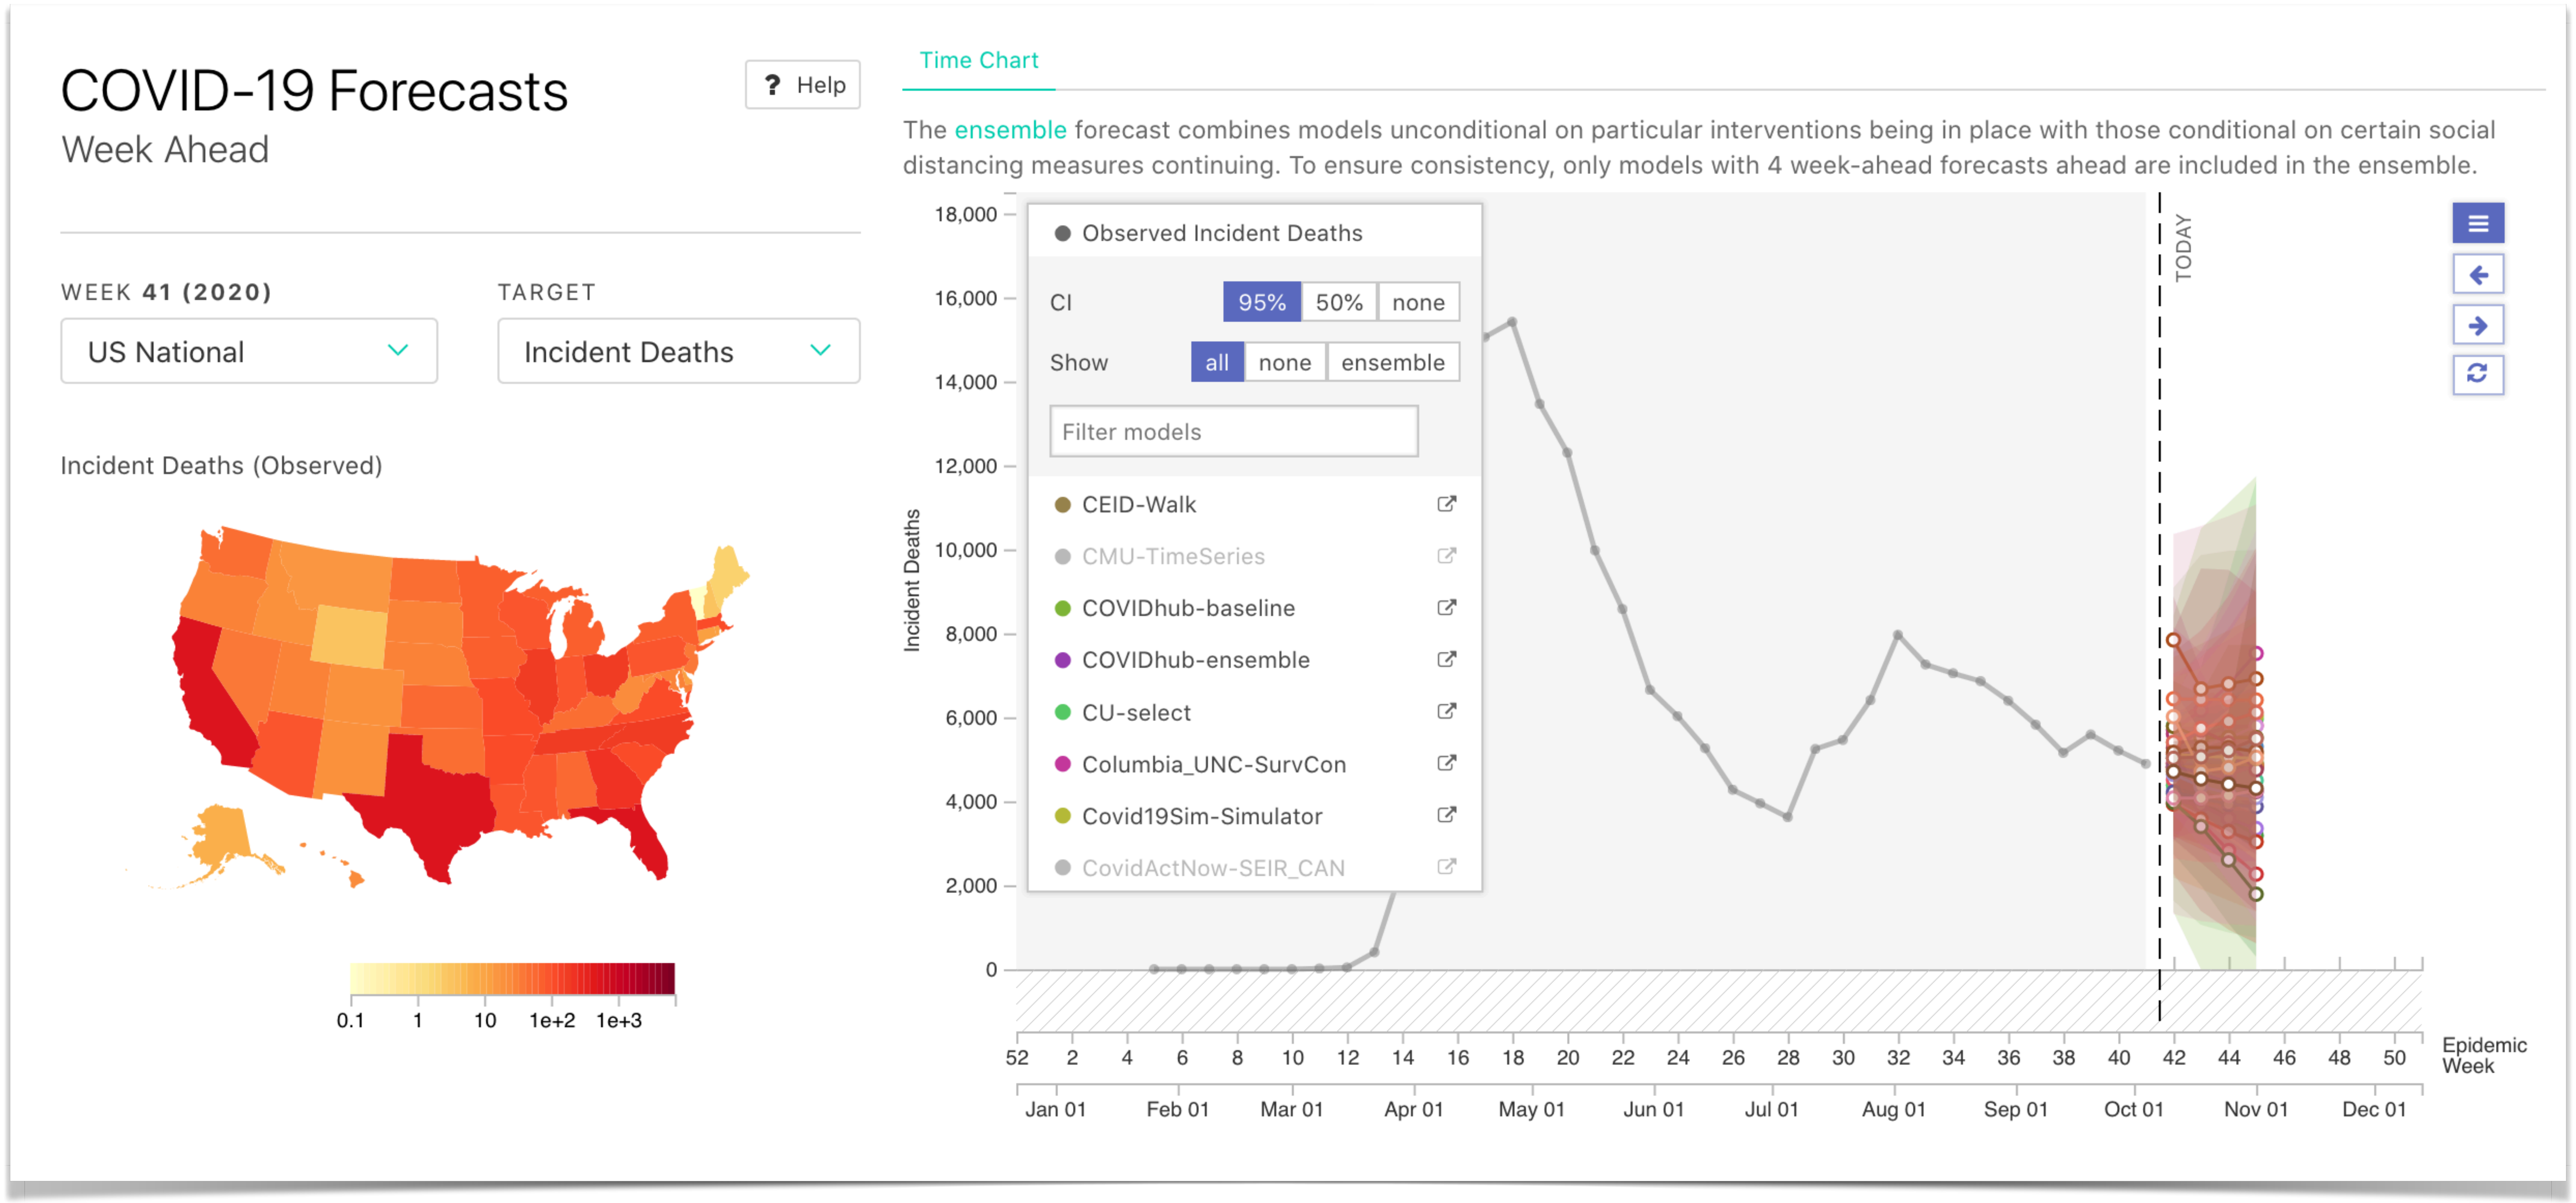

COVID-19 Forecasts

Week Ahead

WEEK 41 (2020): US National

TARGET: Incident Deaths

Incident Deaths (Observed)

Time Chart

The ensemble forecast combines models unconditional on particular interventions being in place with those conditional on certain social distancing measures continuing. To ensure consistency, only models with 4 week-ahead forecasts ahead are included in the ensemble.

http://viz.covid19forecasthub.org

Assessing quantile forecasts

- From each forecaster, collect median and K central prediction intervals $\{l_k, u_k\}$ with k=1…K
- Calculate “weighted interval score” (WIS, Bracher et al, arXiv, 2020). The WIS is negatively oriented, i.e. better forecasts have lower scores.

$$\mathrm{WIS}_{\alpha_{0:K}}(y, F) = \frac{1}{K+1}\left(|y - m| + \sum_{k=1}^{K} \left(\frac{\alpha_k}{2}(u_k - l_k) + \underbrace{(l_k - y)\mathbf{1}(y < l_k) + (y - u_k)\mathbf{1}(y > u_k)}\right)\right)$$

Data

Interval width

Distance median ↔ data

“Sharpness”, i.e.
penalises wide
prediction intervals

“Calibration”, i.e.
penalises data ou
prediction intervals

- Strictly proper forecast score, i.e. optimised for forecasts from the true distribution

Comparing forecast performance

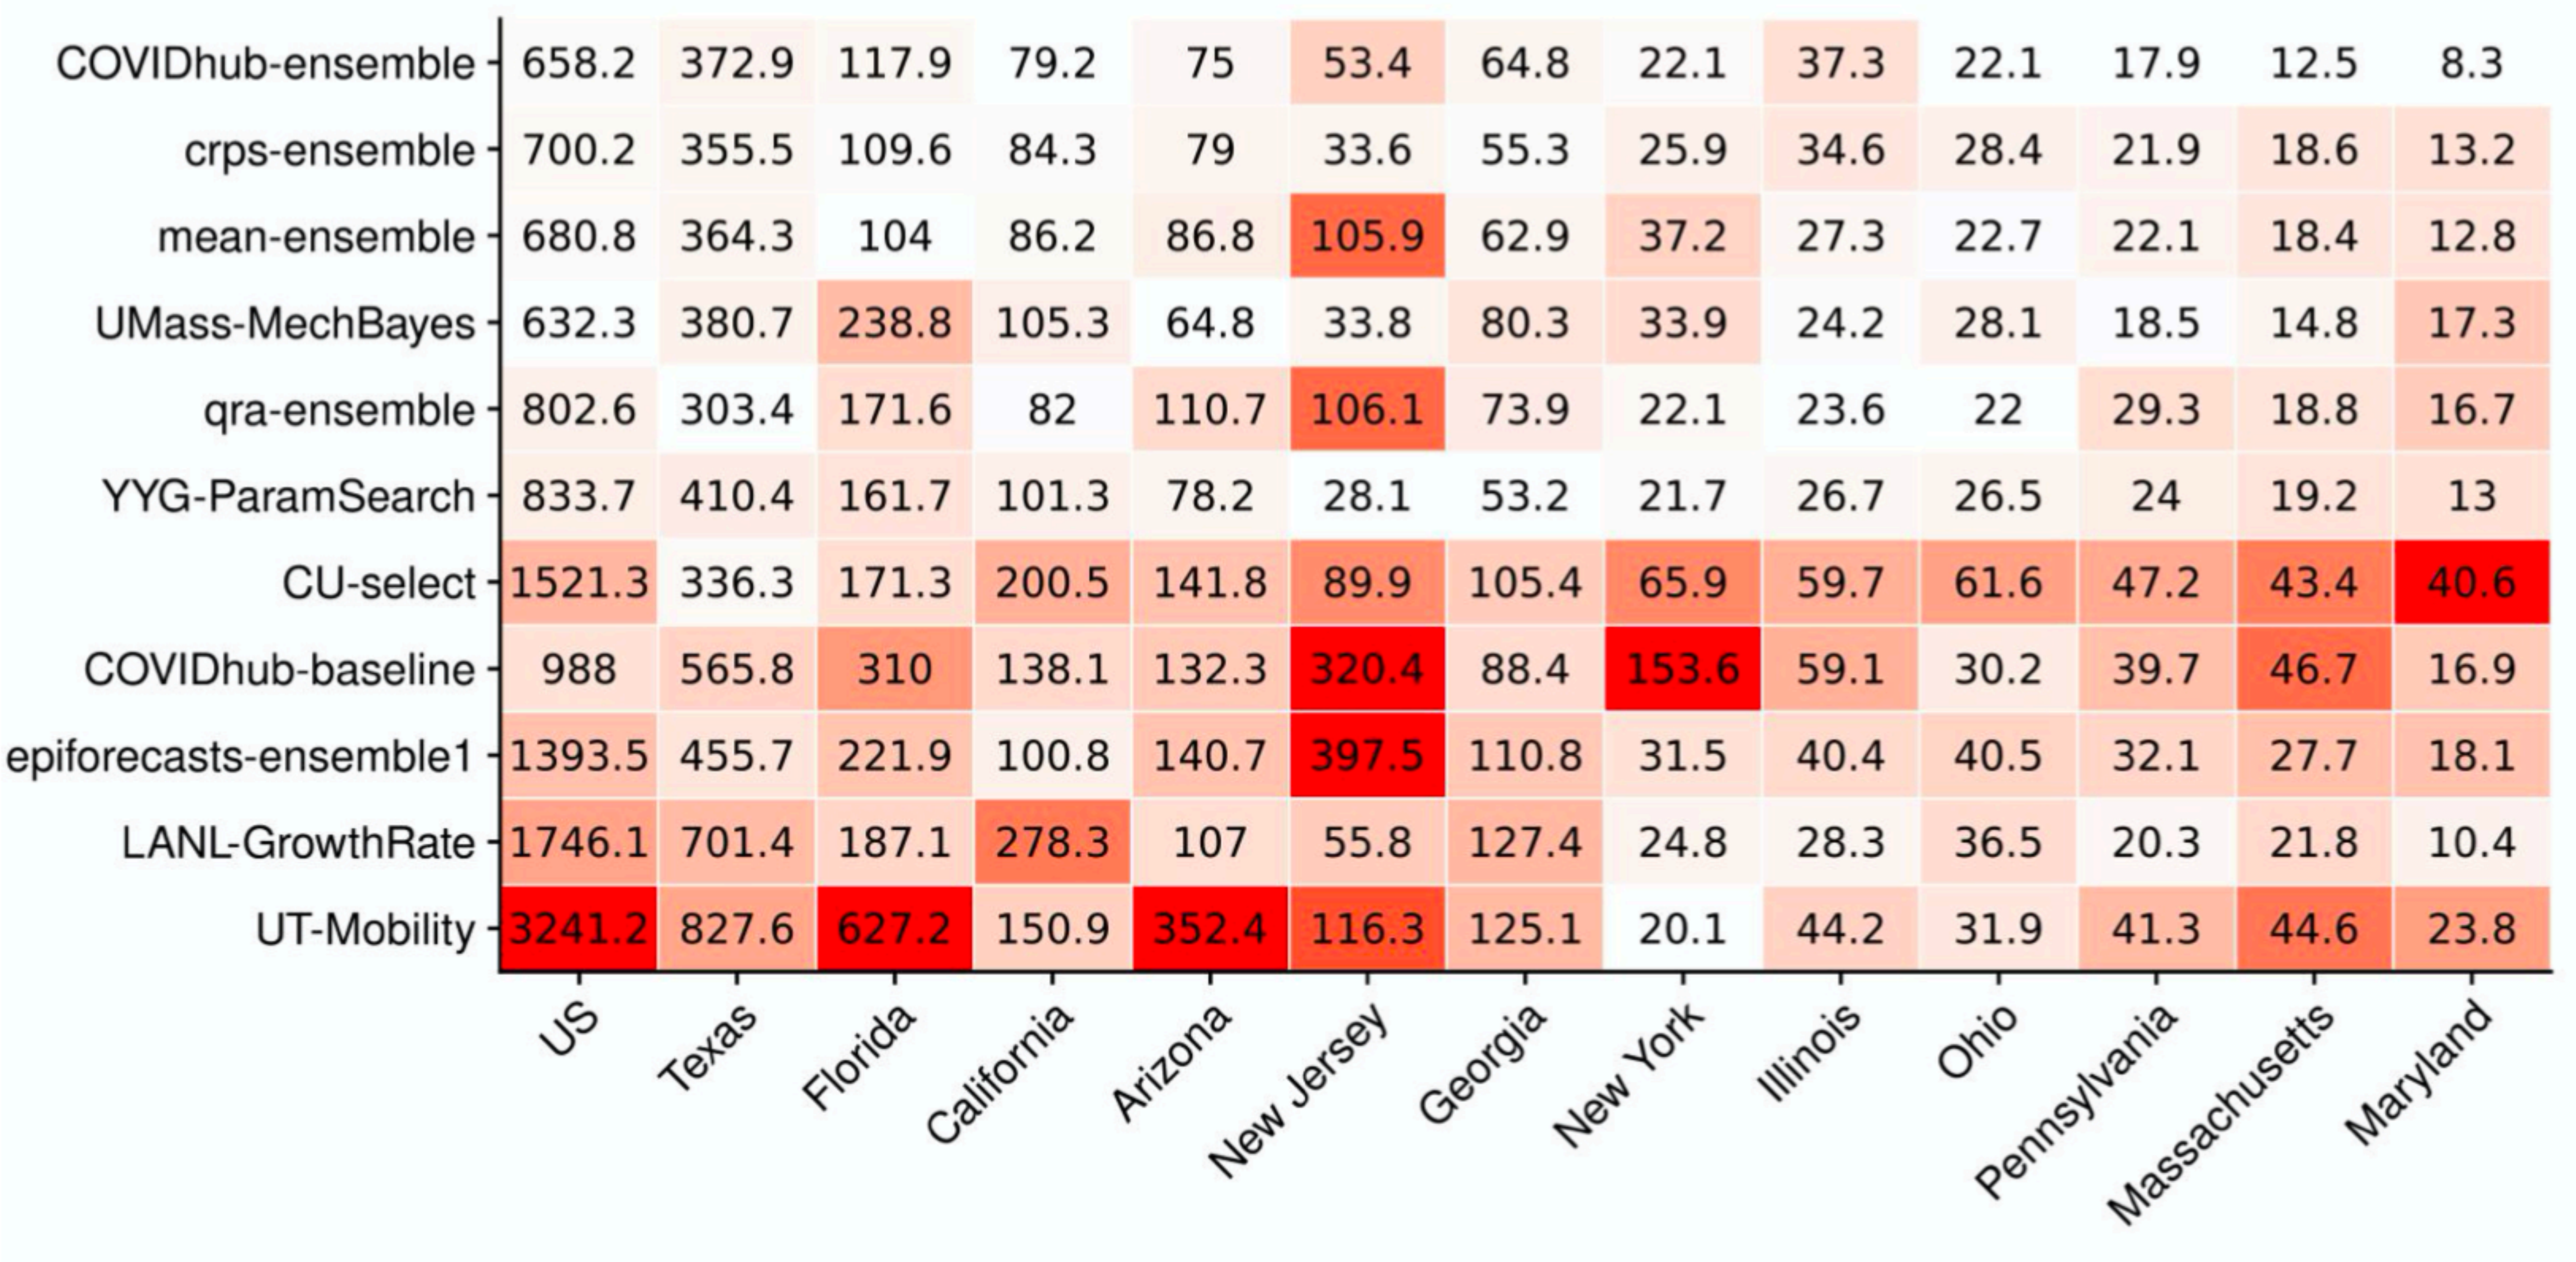

	US	Texas	Florida	California	Arizona	New Jersey	Georgia	New York	Illinois	Ohio	Pennsylvania	Massachusetts	Maryland
COVIDhub-ensemble	658.2	372.9	117.9	79.2	75	53.4	64.8	22.1	37.3	22.1	17.9	12.5	8.3
crps-ensemble	700.2	355.5	109.6	84.3	79	33.6	55.3	25.9	34.6	28.4	21.9	18.6	13.2
mean-ensemble	680.8	364.3	104	86.2	86.8	105.9	62.9	37.2	27.3	22.7	22.1	18.4	12.8
UMass-MechBayes	632.3	380.7	238.8	105.3	64.8	33.8	80.3	33.9	24.2	28.1	18.5	14.8	17.3
qra-ensemble	802.6	303.4	171.6	82	110.7	106.1	73.9	22.1	23.6	22	29.3	18.8	16.7
YYG-ParamSearch	833.7	410.4	161.7	101.3	78.2	28.1	53.2	21.7	26.7	26.5	24	19.2	13
CU-select	1521.3	336.3	171.3	200.5	141.8	89.9	105.4	65.9	59.7	61.6	47.2	43.4	40.6
COVIDhub-baseline	988	565.8	310	138.1	132.3	320.4	88.4	153.6	59.1	30.2	39.7	46.7	16.9
epiforecasts-ensemble1	1393.5	455.7	221.9	100.8	140.7	397.5	110.8	31.5	40.4	40.5	32.1	27.7	18.1
LANL-GrowthRate	1746.1	701.4	187.1	278.3	107	55.8	127.4	24.8	28.3	36.5	20.3	21.8	10.4
UT-Mobility	3241.2	827.6	627.2	150.9	352.4	116.3	125.1	20.1	44.2	31.9	41.3	44.6	23.8

Bosse et al., in prep

Forecast ensembles

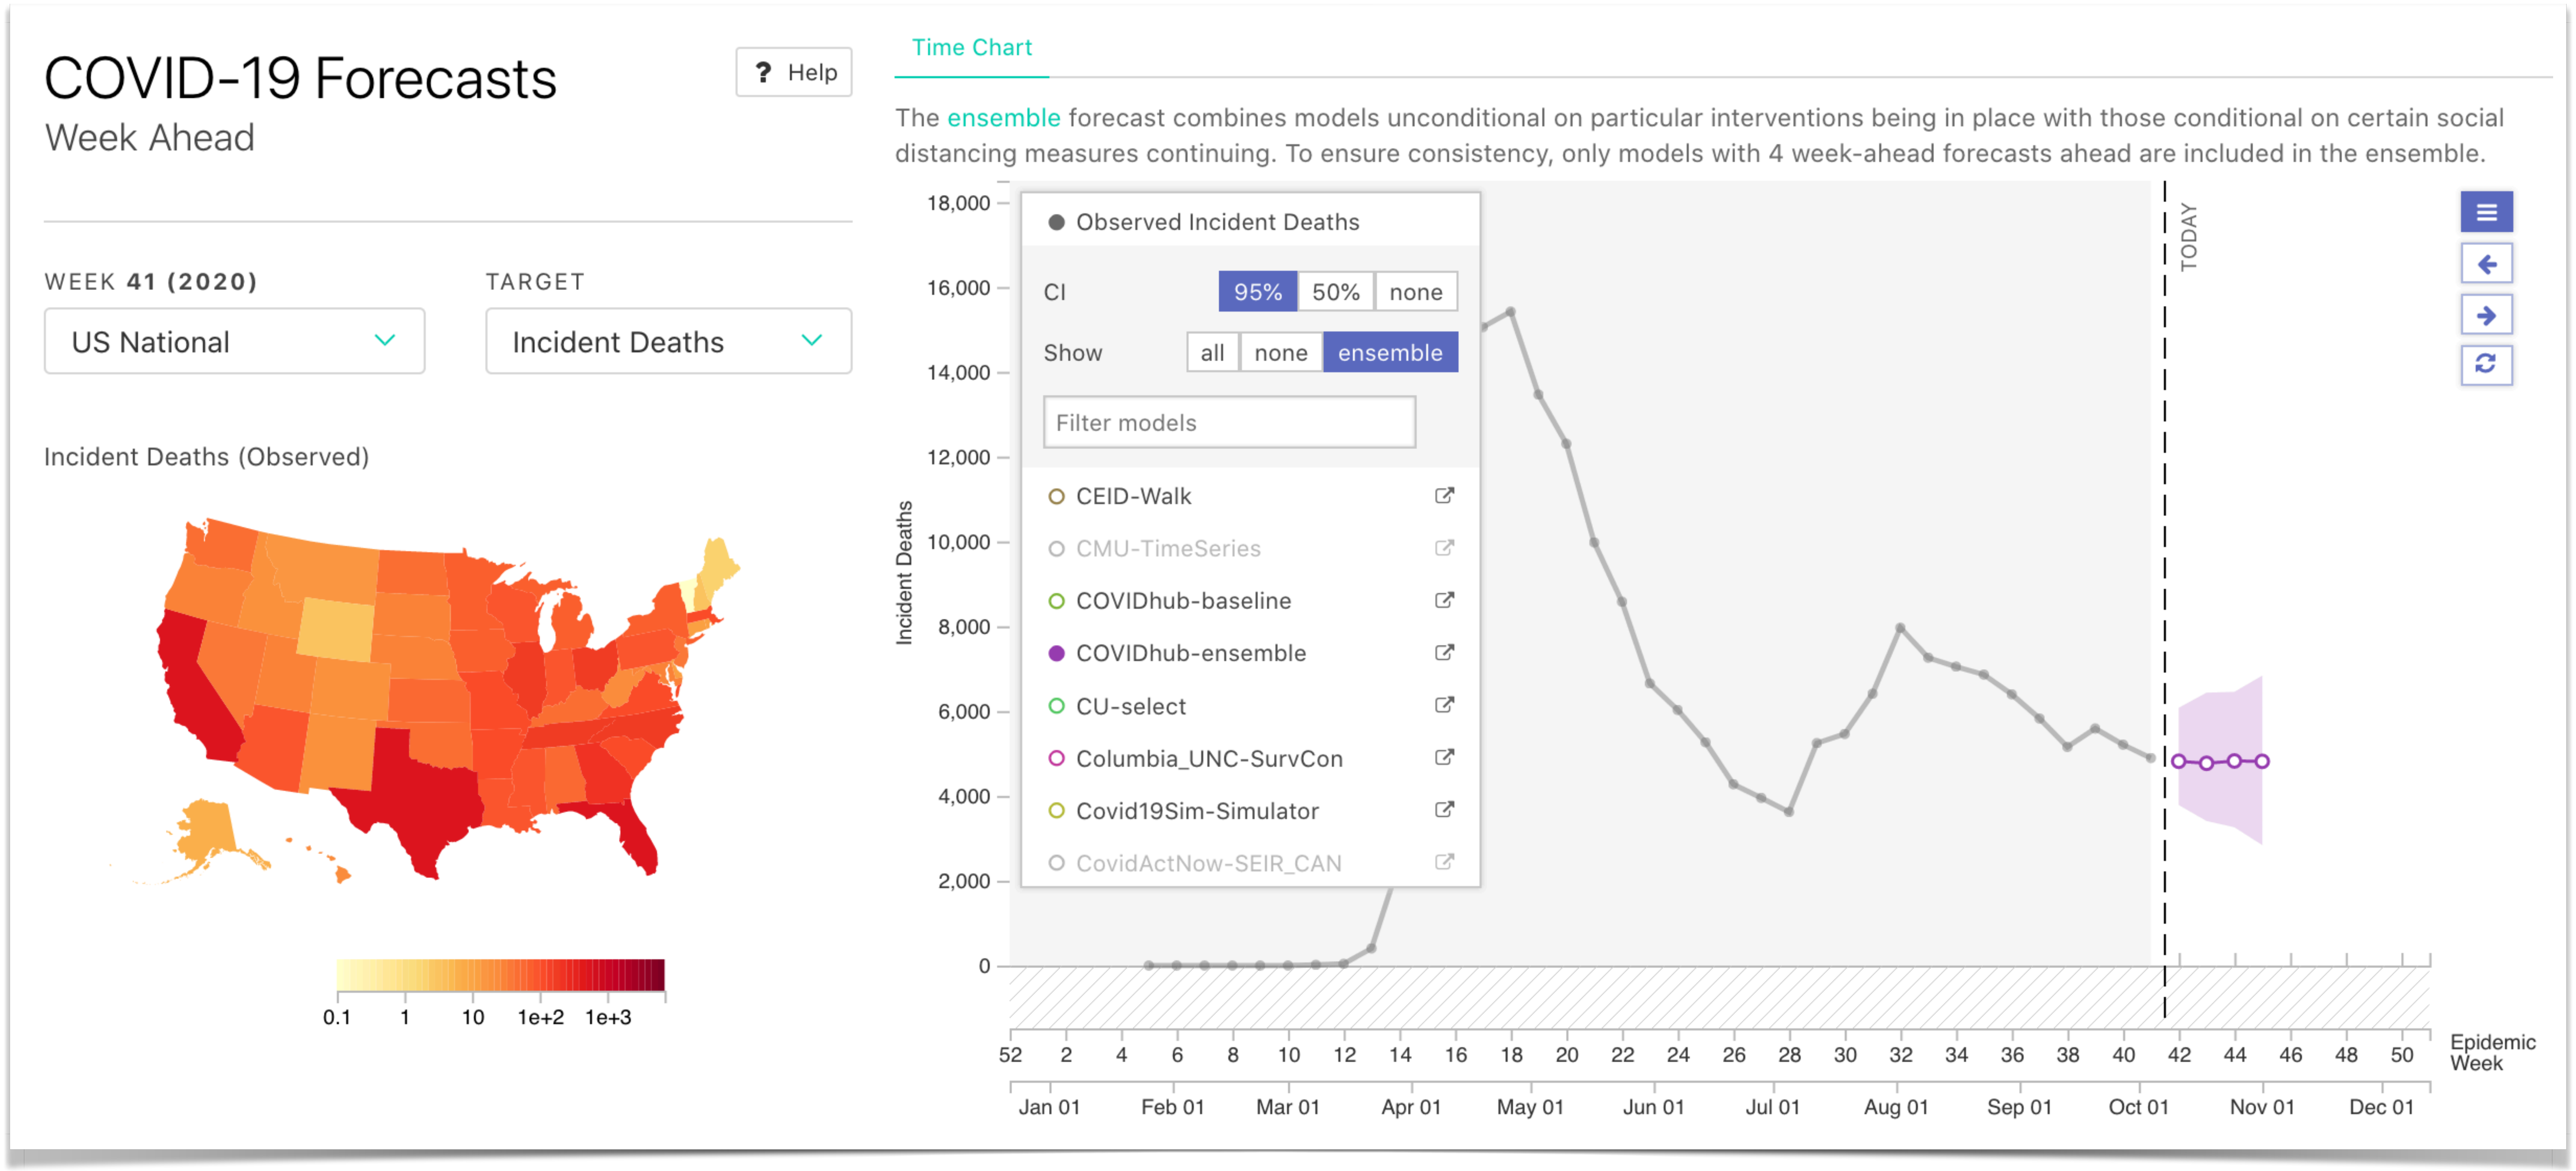

COVID-19 Forecasts

Week Ahead

WEEK 41 (2020): US National

TARGET: Incident Deaths

Incident Deaths (Observed)

Time Chart

The ensemble forecast combines models unconditional on particular interventions being in place with those conditional on certain social distancing measures continuing. To ensure consistency, only models with 4 week-ahead forecasts ahead are included in the ensemble.

● Observed Incident Deaths

CI: 95%

Show: ensemble

- ○ CEID-Walk
- ○ CMU-TimeSeries
- ○ COVIDhub-baseline
- ● COVIDhub-ensemble
- ○ CU-select
- ○ Columbia_UNC-SurvCon
- ○ Covid19Sim-Simulator
- ○ CovidActNow-SEIR_CAN

http://viz.covid19forecasthub.org

Forecast ensembles from predictive quantiles

Quantiles of combined distribution as weighted average of the *n* individual model quantiles:

$$F^{-1}(\alpha) = \sum_{i=1}^{n} w_i F_i^{-1}(\alpha), \qquad 0 < \alpha \leq 1$$

(Vincent, 1912, Behav Monogr; Genest, 1992, Ann Stat)

1) Simplest option: $w_i = 1/n$

2) Quantile regression average (QRA). Learn w_i from past performance, e.g. via linear regression to optimise past WIS

Comparing forecast performance

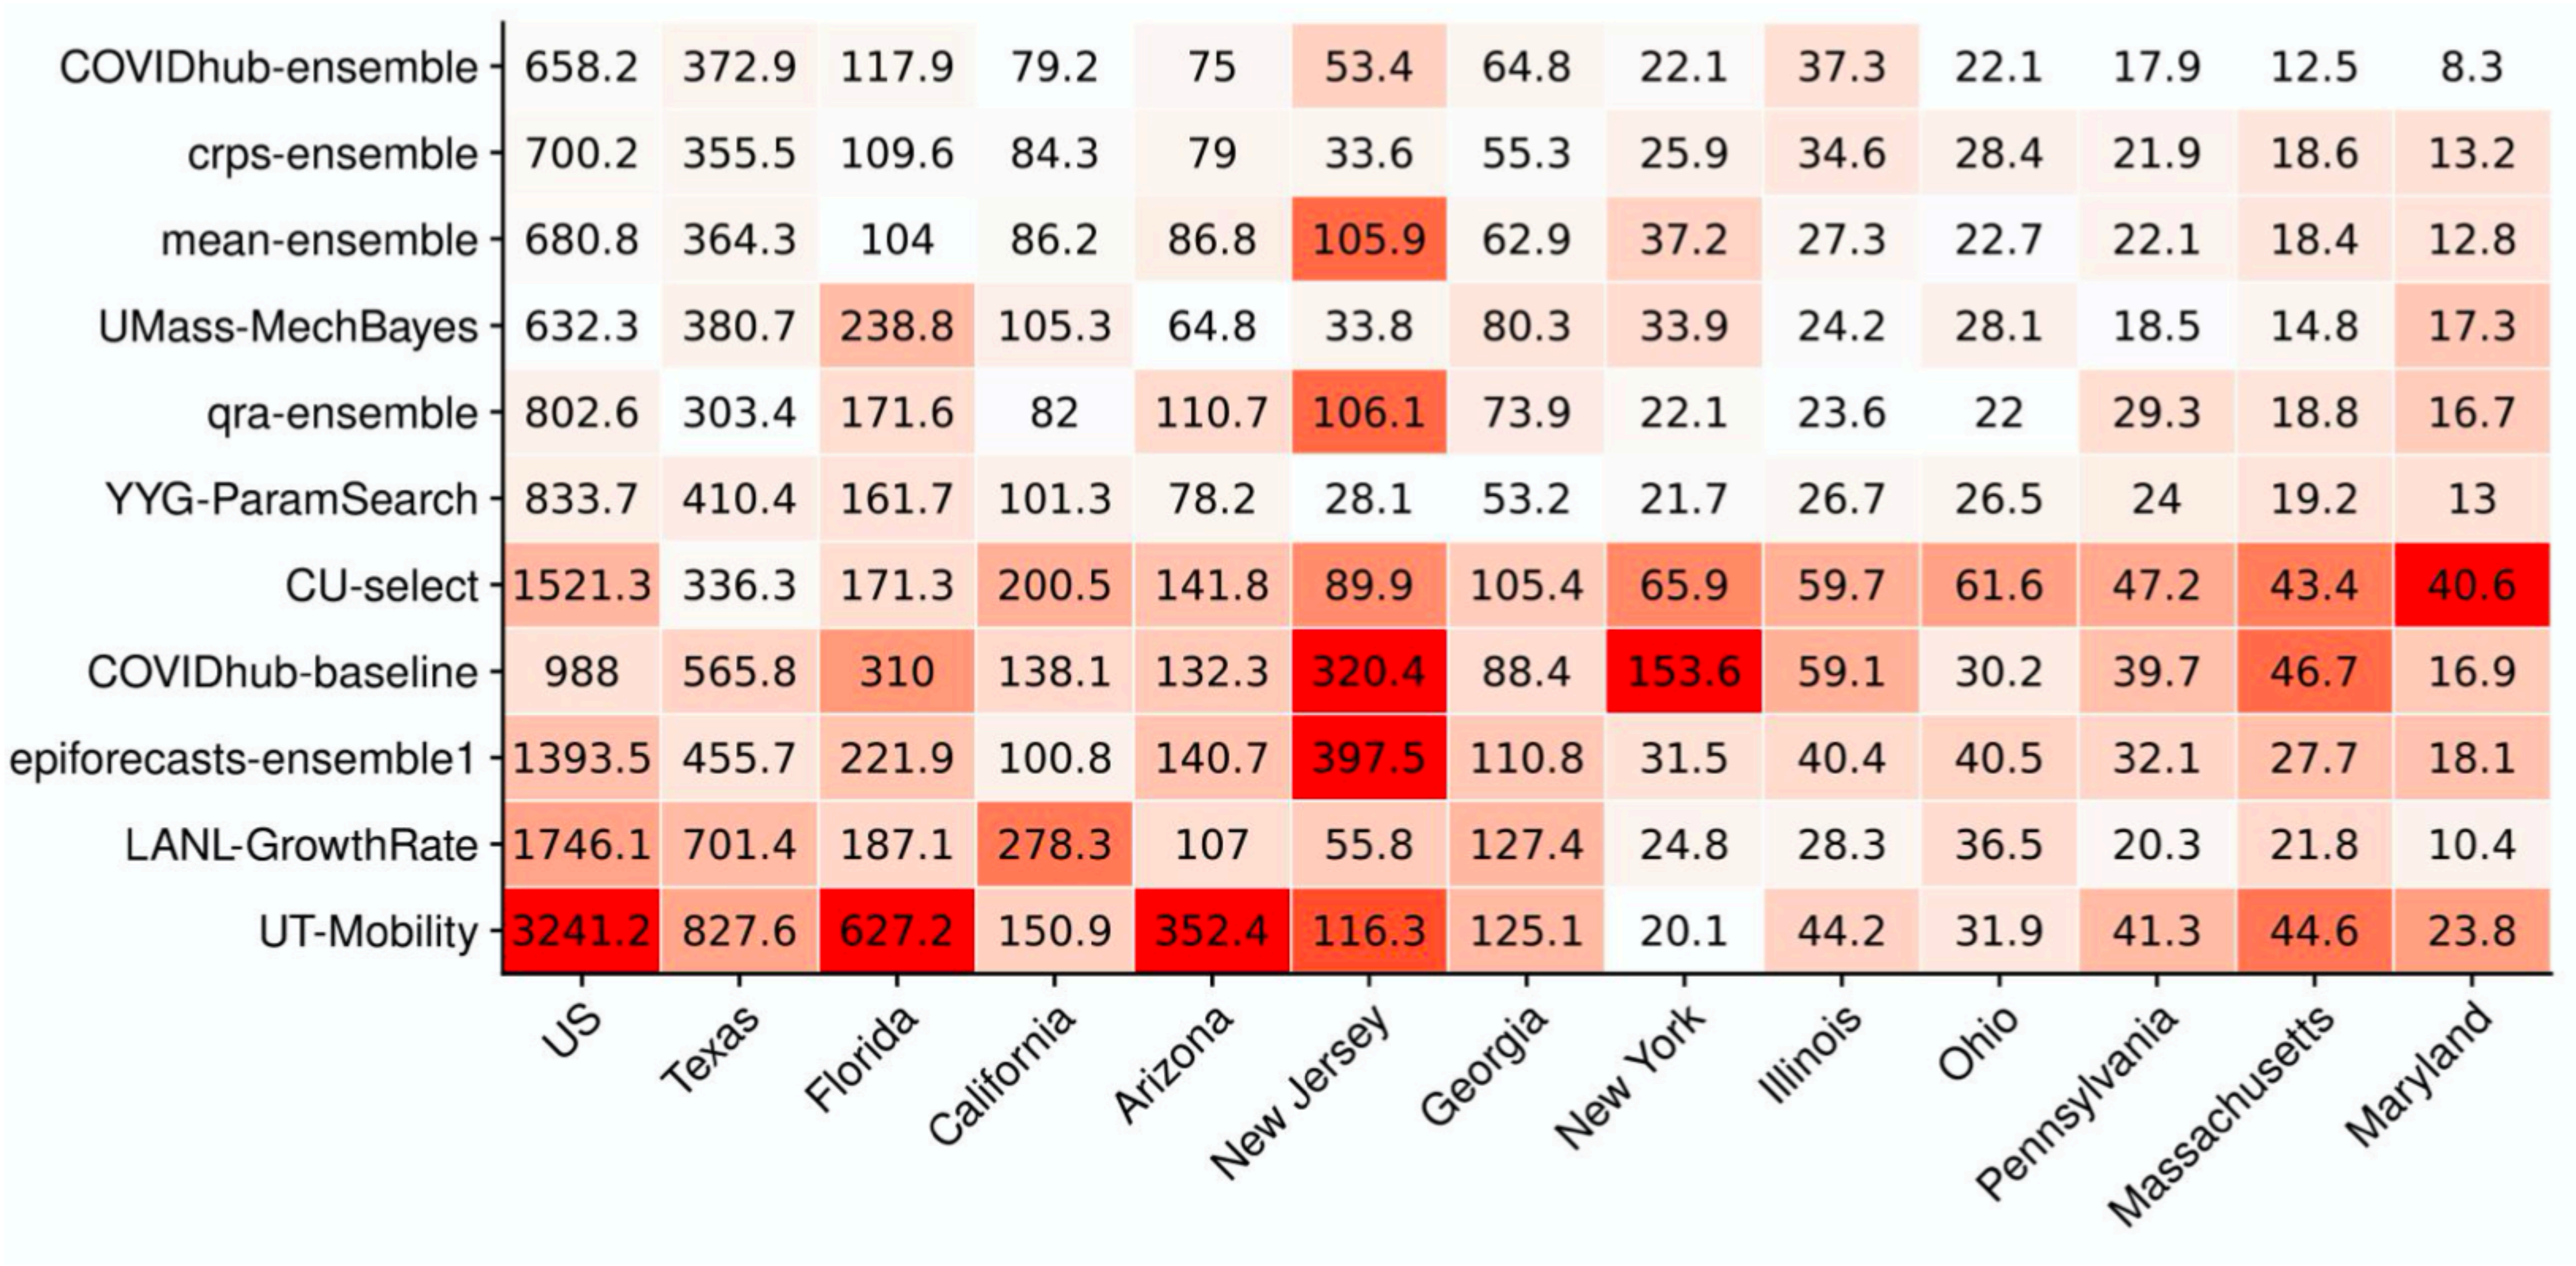

	US	Texas	Florida	California	Arizona	New Jersey	Georgia	New York	Illinois	Ohio	Pennsylvania	Massachusetts	Maryland
COVIDhub-ensemble	658.2	372.9	117.9	79.2	75	53.4	64.8	22.1	37.3	22.1	17.9	12.5	8.3
crps-ensemble	700.2	355.5	109.6	84.3	79	33.6	55.3	25.9	34.6	28.4	21.9	18.6	13.2
mean-ensemble	680.8	364.3	104	86.2	86.8	105.9	62.9	37.2	27.3	22.7	22.1	18.4	12.8
UMass-MechBayes	632.3	380.7	238.8	105.3	64.8	33.8	80.3	33.9	24.2	28.1	18.5	14.8	17.3
qra-ensemble	802.6	303.4	171.6	82	110.7	106.1	73.9	22.1	23.6	22	29.3	18.8	16.7
YYG-ParamSearch	833.7	410.4	161.7	101.3	78.2	28.1	53.2	21.7	26.7	26.5	24	19.2	13
CU-select	1521.3	336.3	171.3	200.5	141.8	89.9	105.4	65.9	59.7	61.6	47.2	43.4	40.6
COVIDhub-baseline	988	565.8	310	138.1	132.3	320.4	88.4	153.6	59.1	30.2	39.7	46.7	16.9
epiforecasts-ensemble1	1393.5	455.7	221.9	100.8	140.7	397.5	110.8	31.5	40.4	40.5	32.1	27.7	18.1
LANL-GrowthRate	1746.1	701.4	187.1	278.3	107	55.8	127.4	24.8	28.3	36.5	20.3	21.8	10.4
UT-Mobility	3241.2	827.6	627.2	150.9	352.4	116.3	125.1	20.1	44.2	31.9	41.3	44.6	23.8

Bosse et al., in prep

Short-term forecasts in the UK

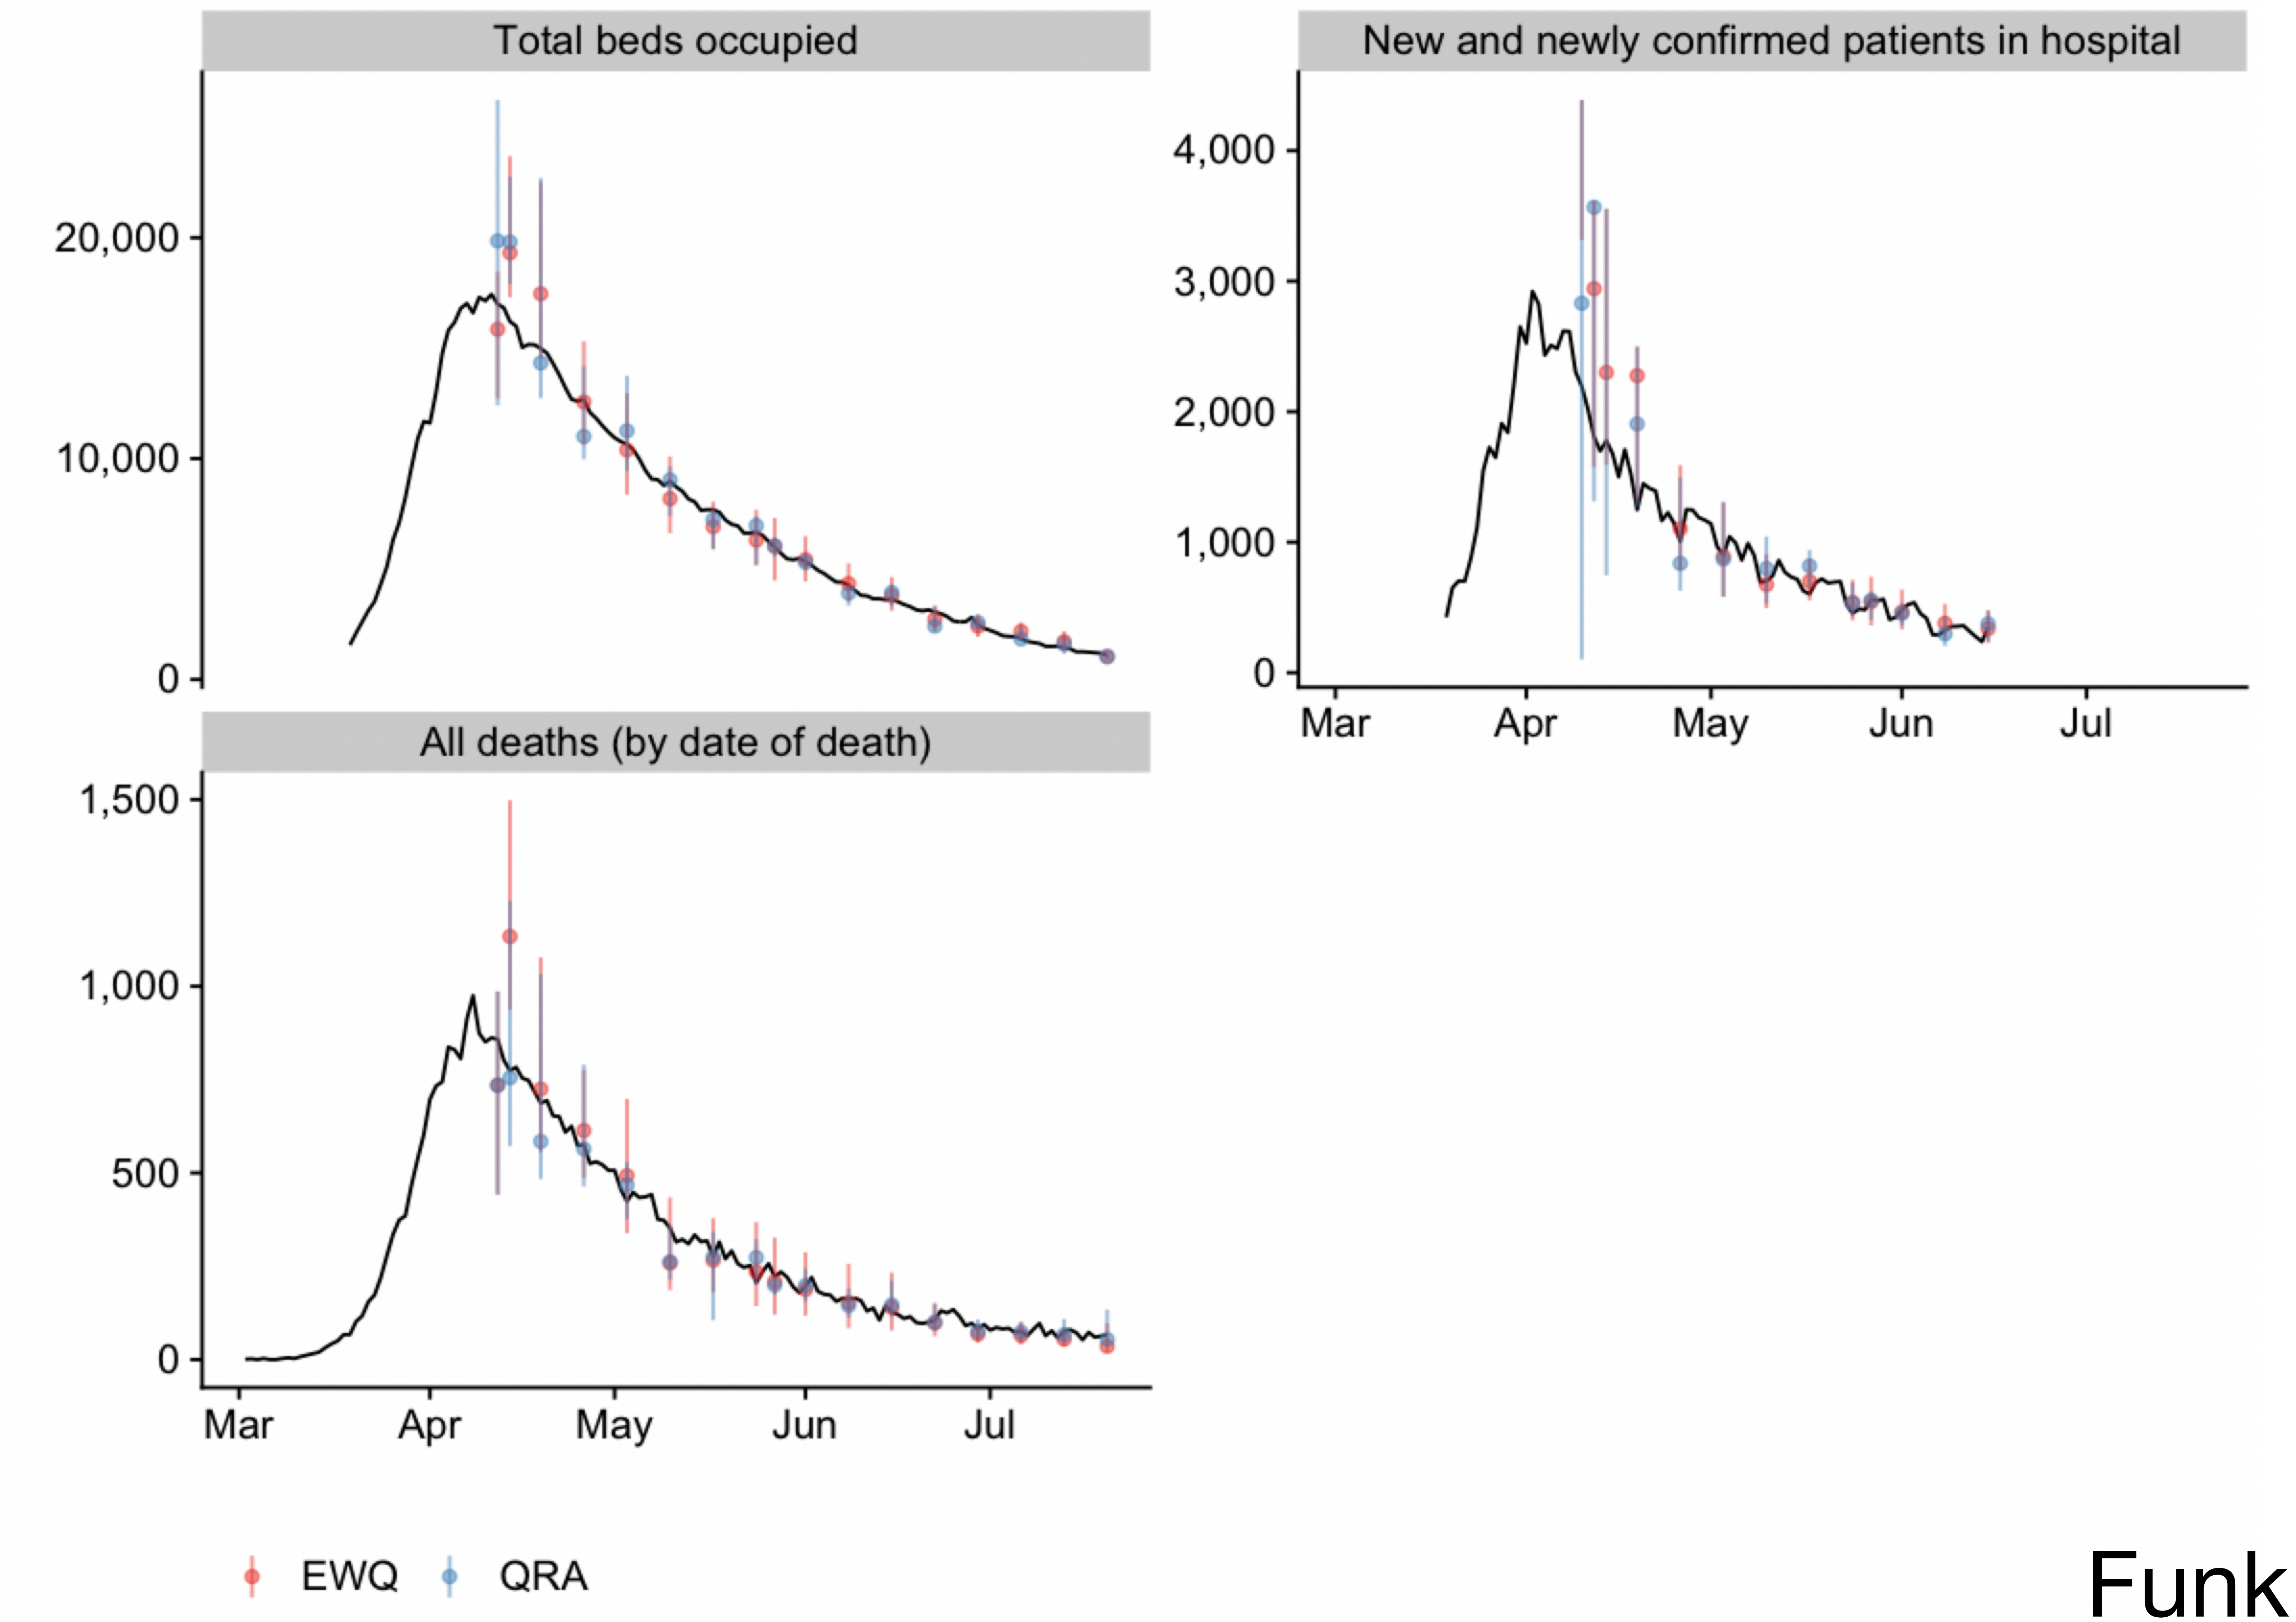

Funk et al., in prep

Acknowledgements

CMMID Covid-19 working group

Team

Sam Abbott
Nikos Bosse
Joel Hellewell
Sophie Meakin
James Munday
Kath Sherratt
Robin Thompson

Collaborators

Johannes Bracher (Karlsruhe)
Katie Gostic (Chicago)
Tilmann Gneiting (Heidelberg)
Evan Ray (Massachusetts)
Nick Reich (Massachusetts)
Ryan Tibshirani (Carnegie Mellon)

Met Office

Google

Microsoft | AI

wellcome

Summary

- Challenges in estimating the reproduction number R_t
    1. Generation interval
    2. Delayed and noisy observations
    3. Smoothing
    4. Validation
- R_t is closely linked to short-term forecasts
- Short-term forecasts can be evaluated and combined

https://epiforecasts.io/covid/