



OPEN

## Forecasting COVID-19 activity in Australia to support pandemic response: May to October 2020

Robert Moss<sup>1✉</sup>, David J. Price<sup>1,2</sup>, Nick Golding<sup>3,4</sup>, Peter Dawson<sup>5</sup>, Jodie McVernon<sup>2,6</sup>, Rob J. Hyndman<sup>7</sup>, Freya M. Shearer<sup>1,3</sup> & James M. McCaw<sup>1,8</sup>


As of January 2021, Australia had effectively controlled local transmission of COVID-19 despite a steady influx of imported cases and several local, but contained, outbreaks in 2020. Throughout 2020, state and territory public health responses were informed by weekly situational reports that included an ensemble forecast of daily COVID-19 cases for each jurisdiction. We present here an analysis of one forecasting model included in this ensemble across the variety of scenarios experienced by each jurisdiction from May to October 2020. We examine how successfully the forecasts characterised future case incidence, subject to variations in data timeliness and completeness, showcase how we adapted these forecasts to support decisions of public health priority in rapidly-evolving situations, evaluate the impact of key model features on forecast skill, and demonstrate how to assess forecast skill in real-time before the ground truth is known. Conditioning the model on the most recent, but incomplete, data improved the forecast skill, emphasising the importance of developing strong quantitative models of surveillance system characteristics, such as ascertainment delay distributions. Forecast skill was highest when there were at least 10 reported cases per day, the circumstances in which authorities were most in need of forecasts to aid in planning and response.

As of January 2021, Australia had effectively controlled COVID-19 transmission, experiencing prolonged intervals of local elimination throughout 2020 despite a steady influx of imported cases, which have been effectively managed by Australia's hotel quarantine system<sup>1</sup>. Scenario modelling conducted in February 2020 informed Australia's rapid and strong initial response<sup>2</sup>, which was sufficient to control the import-driven first wave<sup>3</sup>. In particular, Australia's definitive border measures, including mandatory 14-day hotel quarantine on arrival from overseas, was effective in limiting community exposure from infected international arrivals.

Nonetheless, in May 2020 a breakdown in infection prevention and control measures in the state of Victoria's hotel quarantine system resulted in community exposure that developed into a second wave<sup>4,5</sup>. COVID-19 (ancestral strain) case incidence peaked at 450 locally-acquired infections per day in the last week of July 2020<sup>5,6</sup>. This was substantially higher than the peak of 131 daily cases in the import-driven first wave. Despite the imposition of stringent inter-state movement restrictions, cross-border importation of cases into the state of New South Wales led to a long period of constrained community transmission in that state throughout 2020, kept in check by proactive public health case finding and quarantine without imposition of stringent social measures<sup>7</sup>. In Victoria, a combination of case finding, contact tracing, and prolonged restrictions on movement and gathering sizes were required to bring the second wave to an end in October 2020<sup>5,8</sup>. Subsequent border incursions in various states and territories were associated with small clusters and in some cases localised community transmission, prompting the imposition of strong public health responses supported by variably stringent social restrictions over days and weeks. The daily incidence of locally-acquired COVID-19 cases in each jurisdiction over the study period is shown in Fig. 1.

The public health response to the Victorian second wave and outbreaks in other Australian jurisdictions was informed by the results of real-time analytics and an ensemble forecast of COVID-19 activity for each

<sup>1</sup>Melbourne School of Population and Global Health, The University of Melbourne, Melbourne, VIC, Australia. <sup>2</sup>Department of Infectious Diseases, Melbourne Medical School, at The Peter Doherty Institute for Infection and Immunity, Melbourne, VIC, Australia. <sup>3</sup>Telethon Kids Institute, Perth, WA, Australia. <sup>4</sup>Curtin School of Population Health, Curtin University, Perth, WA, Australia. <sup>5</sup>Defence Science and Technology Group, Melbourne, VIC, Australia. <sup>6</sup>Victorian Infectious Diseases Reference Laboratory Epidemiology Unit, Royal Melbourne Hospital, at The Peter Doherty Institute for Infection and Immunity, Melbourne, VIC, Australia. <sup>7</sup>Department of Econometrics and Business Statistics, Monash University, Melbourne, VIC, Australia. <sup>8</sup>School of Mathematics and Statistics, The University of Melbourne, Melbourne, VIC, Australia. ✉email: rgmoss@unimelb.edu.au



**Figure 1.** Daily incidence of locally-acquired COVID-19 cases for each jurisdiction, as reported after the end of the study period (see "COVID-19 surveillance data" section). Most jurisdictions experienced a first wave in March and April, comprised largely of imported cases. From June to October, Victoria (VIC) experienced a large second wave and New South Wales (NSW) experienced a prolonged low-amplitude small second wave. Other jurisdictions are: Australian Capital Territory (ACT), Northern Territory (NT), Queensland (QLD), South Australia (SA), Tasmania (TAS), and Western Australia (WA). Vertical lines indicate the study period (1 May to 31 October). Note the different vertical scales for New South Wales and Victoria.

Australian state and territory, presented to key government advisory committees in weekly situational reports. Similar analytics and forecasts have supported public health responses around the world<sup>9–13</sup>. Here we detail one of the forecasting models included in the ensemble forecast, which was adapted from Australian seasonal influenza forecasts<sup>14–16</sup> that we have deployed in near-real-time in collaboration with public health colleagues since 2015<sup>17,18</sup>. We evaluate the model's performance by measuring forecast skill relative to an historical benchmark forecast over the study period (May to October 2020) and demonstrate how these outputs supported public health responses during that period. In addition to using the model to predict future rates of locally-acquired COVID-19 cases in each Australian jurisdiction, we demonstrate how our methods were adapted and improved in response to the needs of government, in order to most effectively support the public health decision-making process.

We generated forecasts by combining a stochastic SEIR-type (Susceptible, Exposed, Infectious, Recovered) compartment model with daily COVID-19 case counts through the use of a bootstrap particle filter<sup>16</sup>. The forecasting model incorporated real-time estimates of the effective reproduction number  $R_{\text{eff}}$  for the cohort of active cases (who may not be representative of the whole population), and the transmission potential (TP) averaged over the whole population (characterising the potential for wide-spread transmission)<sup>7</sup>.

In brief,  $R_{\text{eff}}$  is the average number of infections caused by an active case in a population (where some people may be protected against infection) and reflects the *transmission characteristics of active cases* in that population. In contrast, TP characterises the expected number of infections caused by an *average person* in a population (even when there are no active cases in that population) and reflects *whole-population behaviours*. These TP estimates were informed by nationwide behavioural surveys and by mobility data from technology companies<sup>7,19</sup>.

We used the active-case  $R_{\text{eff}}$  to characterise local transmission at the time of forecast, and assumed that transmission from future active cases would gradually become more representative of the whole population. That is, local transmission in the forecasting model gradually shifted from the active-case  $R_{\text{eff}}$  to the whole-population transmission potential (TP) over the forecast horizon. We assess the impact of this assumption about active case

heterogeneity by comparing the forecast skill against that of a separate suite of forecasts where local transmission was maintained at the current active-case  $R_{\text{eff}}$ .

One of the most significant challenges in using predictive models to support public health decision-making in real time is understanding how much confidence to place in the predictions<sup>20–22</sup>. We address the question “When should we trust these forecasts?” by showing how forecast skill can be assessed in real-time *before the ground truth is known*.

## Results

At each week in the study period (May to October 2020) we produced forecasts for each jurisdiction that predicted daily COVID-19 case counts in that jurisdiction over the next four weeks. Note that each forecast was conditioned on the expected case counts by symptom onset date, after accounting for delayed case ascertainment (see “COVID-19 surveillance data” and “Forecasting COVID-19 case numbers in the community” sections for methodological details).

We first examine how well these forecasts characterised future case incidence during the onset of the second wave in Victoria (In “Onset of the second wave in Victoria” section). Over this period, the timeliness and completeness of the near-real-time case data varied substantially, and we show how they impacted the forecast predictions.

We then demonstrate how we adapted the forecasts to support the gradual easing of restrictions in Victoria, as daily case incidence decreased in September and October but continued to persist at low levels (In “Likelihood of a large outbreak” section). Policy changes during this period were subject to reaching specific 14-day moving average case thresholds.

Finally, we present an analysis of the forecast performance across all Australian jurisdictions, in which we measure how accurately our forecasts predicted the future case numbers, relative to an historical benchmark forecast (In “Forecast skill” section). We calculated *skill scores* for our forecasts, where negative values indicate that the forecasts performed worse than the historical benchmark, zero indicates that the forecasts performed as well as the historical benchmark, and positive values indicate that the forecasts performed better than the historical benchmark. The maximum possible skill score is one, which indicates that a forecast yielded perfect predictions with no uncertainty. We identify the circumstances under which forecast skill was highest, evaluate the impact of key model features on forecast skill, and conclude by considering how to assess forecast skill in real-time.

**Onset of the second wave in Victoria.** Throughout May 2020, Victorian case incidence remained low and relatively stable (5.6 cases per day, on average). With such limited chains of local transmission, forecasts generated in May considered local extinction to be the most likely outcome and did not include the possibility of a large outbreak. These predictions were consistent with the findings of a subsequent genomic study, in which the month of May 2020 was characterised by near elimination of COVID-19 in Victoria<sup>23</sup>. However, a new cluster of local cases emerged at the end of May<sup>23</sup>, and by late June more widespread epidemic activity was clearly established, with more than 60 cases per day (Fig. 1).

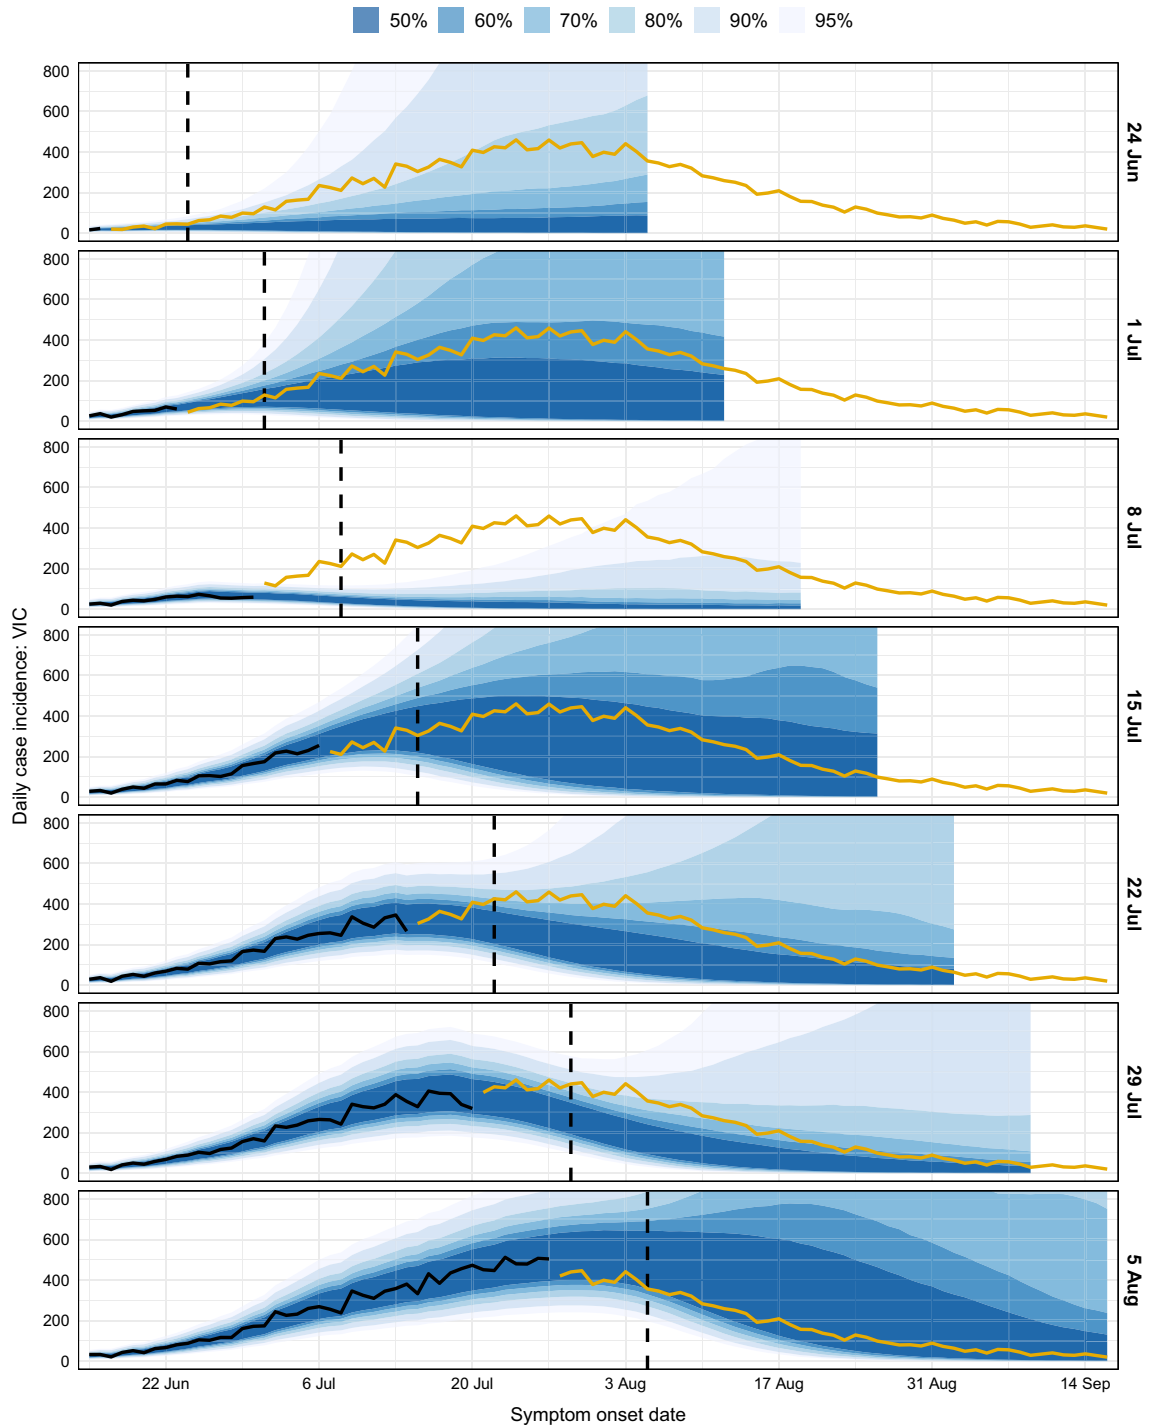
Here we focus on the following key questions: “When did the forecasts first consider a large outbreak *was possible?*”; “When did the forecasts first consider a large outbreak *was likely?*”; and “How did the reported data available at each week influence these predictions?”.

**Likelihood of a large outbreak.** The forecast generated on 24 June was the first to indicate that a large outbreak *might occur*. The 95% credible interval included trajectories with increasing case counts, suggesting that there was at least a 1-in-40 chance of a sustained outbreak, and the lower bounds of each credible interval was greater than zero. Forecasts generated on this date, and for the subsequent six weeks, are shown in Fig. 2, and a qualitative summary of each forecast is provided in Table 1.

The forecast generated on 1 July was the first to demonstrate confidence that a substantial outbreak *would occur* (the 50% credible interval included trajectories with increasing case counts). Its median trajectory exceeded 100 cases per day, the 50% credible interval (CrI) lower bound remained above 10 cases per day until mid-August, and the 50% CrI upper bound was in good agreement with the future case counts over the entire forecast horizon. As local transmission became more established, subsequent forecasts yielded increasingly confident and accurate predictions of the size and timing of this second wave. The 8 July forecast was a marked outlier from the other forecasts over this period, and we address this forecast separately in the next section.

Despite the 50% CrIs being reasonably narrow, the 95% CrIs remained very broad in all of these forecasts. The 95% CrI lower bounds decreased to zero cases per day roughly 2 weeks after the forecast date, and their upper bounds exceeded 1,000 cases per day. This diversity of epidemic trajectories (from the stochastic SEIIR compartment model) illustrates the inherent uncertainty at the time of each forecast about the future course of this second wave, and what impact the imposed public health and social measures might have on local transmission over the forecast horizon.

**The impact of near-real-time data.** A natural consequence of working with near-real-time surveillance data was the challenge of interpreting the reported numbers of locally-acquired cases for recent dates, given the delay between symptom onset and cases being ascertained and reported. Fig. 3 shows the daily case counts reported at each data date, and the expected daily case counts after accounting for delayed ascertainment<sup>7</sup> (see “Forecasting COVID-19 case numbers in the community” section for methodological details). Ascertainment delays are evident in the data reported on and after 24 June (i.e., once there was an indication of established local transmission) and became more pronounced as local case incidence increased. The most recent data in each extract (dashed lines) systematically undershot the corresponding case counts in subsequent data extracts.




**Figure 2.** Forecasts of daily COVID-19 case incidence for Victoria, shown separately for each data date (right-hand axis labels). Each forecast is shown as a spread of credible intervals (blue shaded regions). Vertical dashed lines indicate the data date (i.e., when the forecast was generated). Black lines show the daily case counts at the time of forecasting, after accounting for delayed ascertainment, and dark yellow lines show the future case counts. Note that the daily case counts at the time of forecasting (black lines) end several days prior to the data date; see “Competing models” section for details. The y-axis is truncated at 800 cases; due to the potential for exponential growth in model cases, the upper bounds of the 95% credible intervals reached as high as 45,000 cases per day at the end of the 22 July forecast horizon.

The forecast generated on 8 July is a particularly notable example of how data timeliness and completeness can impact model predictions. Its 50% CrI is much narrower than those of other forecasts and trends downward

Date	Forecast characteristics
24 June	50% CrI is flat (no outbreak), 95% CrI includes possibility of large outbreak
1 July	50% CrI upper bound in good agreement with future case counts
8 July	95% CrI undershoots future case counts
15 July	50% CrI includes all of the future case counts
22 July	50% CrI undershoots future case counts after the peak
29 July	50% CrI undershoots future case counts after the peak
5 August	50% CrI includes all of the future case counts

**Table 1.** Qualitative summaries of the forecast credible intervals (CrIs) during the onset of the second wave, where “Date” is the date on which each forecast was generated.



**Figure 3.** Daily incidence of COVID-19 cases for Victoria from mid-May until the end of June, shown separately for each data date (right-hand axis labels). Each data extract updated the entire time-series. Forecasts were conditioned on the expected case counts after accounting for delayed ascertainment (solid black lines); the reported case counts in the data extracts are also shown (dashed black lines). Updated case counts, as reported on 12 August, are shown as grey columns.

from the most recent observation (30 June), undershooting every future daily case count, even those before the data date itself. This forecast was influenced by an apparent decrease in daily cases reported over 25–30 June (Fig. 3), even after applying the modelled correction for delays between symptom onset and case ascertainment (based on data from earlier data extracts), and might have revealed an opportunity for achieving elimination. However, case counts for these dates were substantially higher in subsequent data extracts. While there were a number of local clusters in Victoria around this time, including in multiple schools and public housing sites<sup>24,25</sup>, to the best of our knowledge this ascertainment delay was not linked to any specific cluster. A plausible explanation is that unprecedented case numbers impacted various aspects of the public health response (e.g., laboratory testing, contact tracing, reporting). The Victorian Government subsequently received additional support from Australian Defence Force personnel<sup>26</sup>.


In contrast, the forecast generated on 5 August is in good agreement with the data, and its 50% CrI includes all of the future case numbers. Note the raw data available at the time of forecast suggested an epidemic in decline, but incidence remained near-maximal for the next two weeks and this was captured by the forecast.

These observations highlight the importance and the challenge of understanding the *meaning* and *limitations* of the available data in near-real-time. While correcting for delays (right censoring) is necessary, such corrections are not static through time. It is fundamentally difficult to estimate the completeness of these data in rapidly-evolving situations, such as during rapid epidemic growth, where bottlenecks and delays are not directly observable and are subject to change. Active communication with those responsible for the data is crucial to ensuring that the data and the resulting forecasts are interpreted appropriately, particularly when there are changes in surveillance, testing, and/or reporting.

**Policy context: easing restrictions in Victoria.** In response to the imposed restrictions (“stage 3” and “stage 4”), daily case numbers in Victoria peaked in August and steadily decreased thereafter<sup>5,7</sup>. By early September the Victorian government had announced a plan to gradually ease these restrictions over a period spanning mid-September to late November, conditional on achieving specific case thresholds over 14-day windows ahead of each planned date for policy changes<sup>27</sup>. However, similar to the months of May and June, local cases remained at low but persistent levels in the second half of September and into October (see Fig. 4). The city of Melbourne had been under intense mobility and gathering restrictions (so-called “lockdown”) since 7 July, and the growing need to ease these restrictions had to be carefully balanced against the risk of resurgence. We used our forecasts to predict the probability that the 14-day moving average for incident cases would achieve these specific thresholds at each day of the forecasting period (see “Meeting case-number thresholds that trigger policy changes” section for methodological details).

Each of the four target case thresholds was achieved in the months of September and October 2020 (see Table 2). For each case threshold, we inspected the forecasts made in the weeks prior to that threshold being achieved, to see how far in advance the forecasts suggested that these thresholds would be achieved. The forecasts predicted a greater than 50% chance of achieving each threshold three to four weeks in advance (see Table 2 and Fig. 5). For example, the  $\leq 50$  cases threshold was achieved on 13 September. Twenty five days earlier, the forecast generated on 19 August predicted a 46.8% chance of achieving this threshold by 13 September. The next forecast (generated on 26 August) predicted a 63.5% chance of achieving this threshold by 13 September.

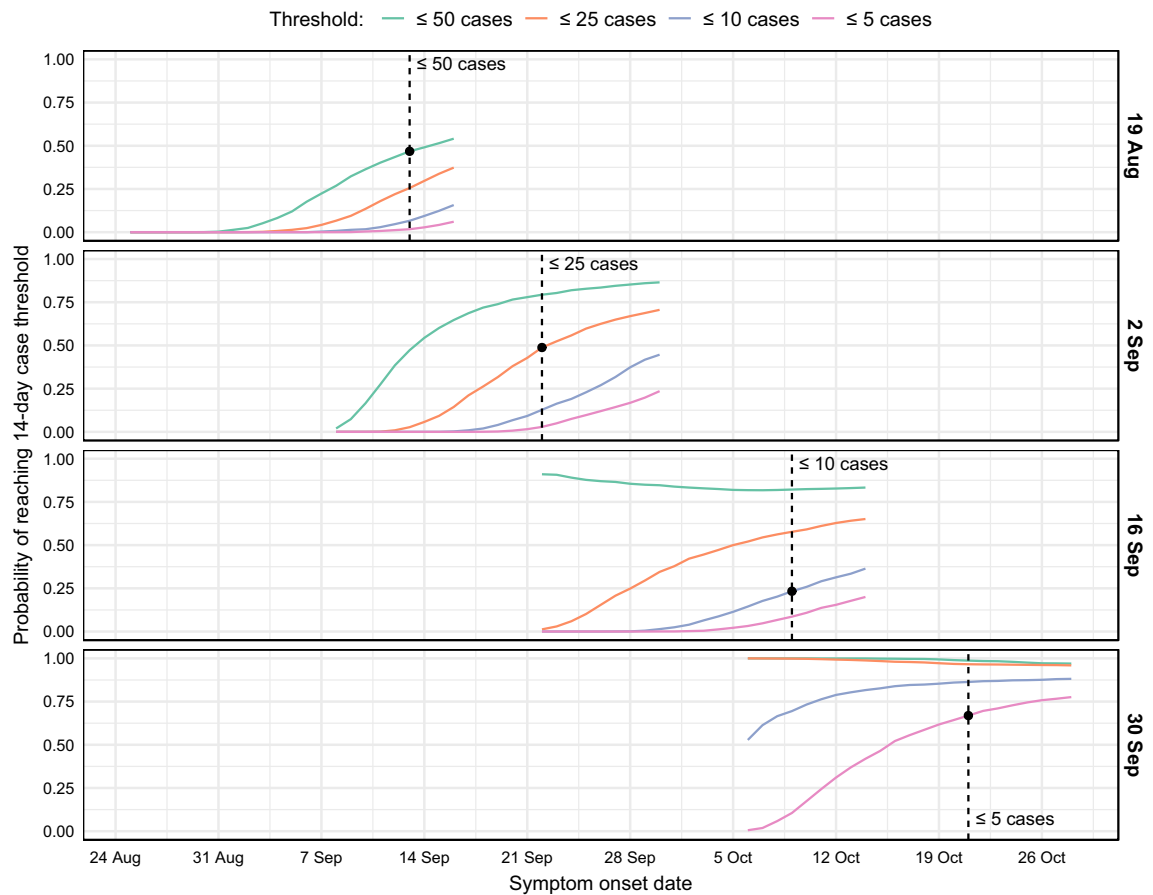
The one exception was the 10-case threshold, which was reached on 9 October. In the weeks prior to this date, daily case counts fluctuated around 10 cases per day (see Fig. 4). Three weeks prior to this threshold being reached, the forecasts predicted only a 23% chance of reaching this threshold by 9 October. By the following week — two weeks in advance — the forecasts predicted a 68% chance of reaching this threshold by 9 October.



**Figure 4.** Daily incidence of COVID-19 cases for Victoria in October, shown separately for each data date (right-hand axis labels). Forecasts were conditioned on the expected case counts after accounting for delayed ascertainment (solid black lines); the reported case counts in the data extracts are also shown (dashed black lines). Updated case counts, as reported after the study period, are shown as grey columns.

Threshold	First reached	Data date	Probability	50% CrI
$\leq 50$ cases	13 Sep	19 Aug	46.8%	7-94
		26 Aug	63.5%	8-44
$\leq 25$ cases	22 Sep	26 Aug	64.1%	1-24
$\leq 10$ cases	9 Oct	23 Sep	68.2%	1-8
$\leq 5$ cases	21 Oct	23 Sep	72.7%	0-4

**Table 2.** Forecast predictions concerning the 14-day moving average case thresholds. For each threshold, this table shows the date when that threshold was first reached, and the following forecast quantities for that date: (a) the forecast data date; (b) the projected probability of reaching the threshold by the date when it was first reached; and (c) the forecast 50% credible interval for daily case incidence on that date.



**Figure 5.** Time series showing the projected probability of reaching each of the 14-day moving average case thresholds, shown separately for each data date (right-hand axis labels). Vertical dashed lines show when each threshold was first reached (with respect to symptom onset date, not reporting date). The black points show the projected probability of having achieved each threshold on the day that it was reached.

**Forecast skill.** Having described the behaviour and qualitative performance of our forecasts during periods of increasing, decreasing, and low levels of local transmission, we now consider their performance more precisely, through skill scores. We measured forecast skill relative to an historical benchmark, described in "Forecast targets and skill" section, which assumed that future case counts would follow the same distribution as the case counts observed to date. This provides a quantitative performance evaluation and helps to address a critical question when using forecasts to support public health decision-making: "When should we trust these forecasts?"

**Forecast uncertainty.** To quantify forecast uncertainty we used *forecast sharpness*, a measure of the breadth of forecast outcomes (e.g., a sharpness of zero means there is a single possible outcome). For jurisdictions that reported substantial COVID-19 activity, as the lead time increased (i.e., looking further into the future) the forecasts exhibited little increase in uncertainty, or even grew more certain (Sect. S1). This trend reflects, in part, that as local epidemic activity increased there was more information available to inform the forecasts and constrain the range of plausible outcomes. During the second wave in Victoria, strong public health measures and

social restrictions markedly reduced the TP, and this caused the forecasts to become sharper at longer lead times (Fig. S2). In contrast, assuming that transmission would be sustained at the current active-case  $R_{\text{eff}}$  caused the forecasts to become more uncertain at longer lead times (Fig. S3). For jurisdictions that primarily reported zero cases on a given day, the forecast trajectories tended to local extinction, and forecast sharpness was close to zero (i.e., minimal uncertainty) for all lead times.

**Forecast bias and case heterogeneity.** The forecasts tended to under-estimate future case counts (negative bias, see Sect. S2). For jurisdictions which had achieved local elimination, such as South Australia in May, this reflects a deliberate model feature: our careful separation of the *potential* for reintroduction from the *consequence* of reintroduction. Without evidence of reintroduction into these jurisdictions, the forecasts did not admit the possibility of subsequent epidemic activity (see "Accounting for observed reintroduction" section) and predicted there would be zero reported cases each day. On the rare occasions where a jurisdiction subsequently reported one or two isolated cases (see Fig. 1) this resulted in a negative bias (under-estimation).

For Victoria, which experienced substantial COVID-19 activity during the study period, model assumptions about local transmission contributed to the underestimation of future case counts. Recall that the model used the active-case  $R_{\text{eff}}$  to characterise local transmission at the time of forecast, with transmission reverting to the transmission potential (TP) over time (see "Forecasting COVID-19 case numbers in the community" section for details; example  $R_{\text{eff}}$  trajectories are shown in Sect. S4). However the second wave in Victoria was dominated by localised outbreaks in health and aged care settings, and other essential services, where public health and social measures to constrain transmission had less impact than in the general population<sup>5</sup>. Over a 5-month period, the  $R_{\text{eff}}$  was systematically greater than the TP, as described and evaluated in Golding et al.<sup>7</sup>. Accordingly, as shown in Fig. 2, the majority of the Victorian forecasts somewhat under-estimated future case counts. Despite this negative bias, the model consistently out-performed forecasts where local transmission was instead maintained at the current active-case  $R_{\text{eff}}$  (see Sect. S3.1). Noting that  $R_{\text{eff}}$  was consistently higher than the TP, this highlights the importance (and challenge) of accounting for heterogeneity of transmission in diverse settings and population sub-groups.

**Forecast skill and the historical benchmark.** Of the five jurisdictions that primarily reported zero cases on a given day, the forecasts consistently out-performed the historical benchmark for all lead times in the Australian Capital Territory and Western Australia, and out-performed the historical benchmark for 1-week lead times in South Australia and Tasmania. The historical benchmark out-performed the forecasts for all lead times in the Northern Territory, where the past case counts were extremely good predictors of the future case counts. For Queensland, which saw a sustained period of 1–5 cases per day (July to September 2020) the forecasts out-performed the historical benchmark for lead times of up to 2 weeks. For longer lead times the forecasts tended to include the possibility of a sustained increase in cases, which was inconsistent with the ground truth and reduced their skill relative to the historical benchmark.

For jurisdictions that experienced substantial COVID-19 activity (New South Wales, Victoria) the forecasts exhibited improved performance relative to the historical benchmark for shorter lead times (Sect. S3.3). For Victoria, which experienced the largest local outbreak, the forecasts substantially out-performed the historical benchmark for lead times of up to 4 weeks (Sect. S3.4).

Forecast skill across all jurisdictions was highest when daily case numbers were zero or at least ten (Sect. S3.1). The forecasts also out-performed the historical benchmark when there were 1–9 daily cases, but to a lesser degree. Across all jurisdictions, 1–9 daily cases comprised 18% of the daily observations; 4% in the state of Victoria, the other 14% in jurisdictions that did not experience sustained epidemic activity. These are the circumstances under which there is great uncertainty as to whether local transmission will become established and drive exponential growth in case numbers. The model forecasts began to include the possibility of a large outbreak in these circumstances (e.g., see the Victorian forecast for 24 June in Fig. 2) and many trajectories substantially exceeded the future case counts, since enacted control measures regularly curtailed transmission.

In contrast, the historical benchmark predictions were defined by the daily case counts reported to date and yielded much narrower predictions than the forecast model in these circumstances. The majority of the 1–9 daily case observations occurred in jurisdictions that did not experience sustained epidemic activity, for which the past observations were good predictors of the future case counts, and the historical benchmark performed well. For Victoria, the historical benchmark predictions substantially underestimated the future case counts in the second wave, and did not include the ground truth as a possibility.

Of course, when considering the model forecasts as a risk evaluation tool, their ability to include the potential for large outbreaks in their projections is a major advantage over the historical benchmark, as is their ability to predict peak timing and size.

**Forecasting in near-real-time.** We highlighted some challenges of using near-real-time data in "Onset of the second wave in Victoria" section. By accounting for delayed ascertainment in the reported case counts (depicted in Figs. 3 and 4) we extracted additional information that, on average, increased the forecast skill (Sect. S3.2), even though it slightly decreased the forecast skill for days where 1–9 cases were reported (Sect. S3.1). Forecast performance was particularly sensitive to the imputed case counts when the reported case counts were low and a difference of, say, 1 case per day represented a large relative change (see Figs. 3 and 4 for examples).

When considering forecast skill as a function of the number of cases ultimately recorded for each day, as above, the results can only be evaluated in retrospect, since these numbers are unknown when the forecast is made. Accordingly, we also considered forecast skill as a function of the number of daily cases *predicted by the forecast model* (Sect. S3.5). In other words, before we know the ground truth, can we *anticipate* what the forecast



skill is likely to be, based only on (a) the current forecast predictions; and (b) the forecast skill of previous forecasts (for which the ground truth is already known). Our intent was to address the question “What does this forecast mean?” before the ground truth is known, with the aim of providing assistance when interpreting the forecast outputs in near-real-time.

The skill scores were similar when aggregated by the number of cases ultimately reported for each day or by the median number of predicted case numbers, particularly when there were at least 5 cases per day (Fig. S17). This means that we can estimate the forecast skill by using the median forecast as a proxy for the ground truth; past performance appears to be indicative of future performance. This indicates that in decision-making contexts (i.e., when future case counts are unknown) we can still obtain a reasonable estimate of forecast skill and assess forecast reliability in real-time. These estimates are particularly reliable when local transmission is established.

## Discussion

**Principal findings.** Forecasting infectious disease activity is a fundamentally difficult problem<sup>28</sup>, and even more so in unprecedented circumstances and in the face of regularly updated public health responses, public perceptions of risk, and changes in behaviour<sup>29–31</sup>. Despite these challenges, we were able to adapt our existing seasonal influenza forecasting methods<sup>14–18</sup> to COVID-19 transmission in Australia, and provide the resulting forecasts to public health authorities to support the Australian COVID-19 response. We tailored our analyses in near-real-time to support public health decision-making in rapidly-evolving situations. This included providing public health authorities with estimated future epidemic activity, timelines for achieving case thresholds that would trigger easing of restrictions, and risk assessments based on repeated incursions of new chains of transmission in primarily infection-free jurisdictions.

Our forecasts exhibited the greatest skill in the state of Victoria, which experienced the overwhelming majority (96%) of locally-acquired cases reported in Australia over the study period, May–October 2020. Forecast skill was highest when there were at least 10 reported cases on a given day, the circumstances in which authorities were most in need of forecasts to aid in planning and response.

In considering several alternate forecast models, forecast skill was higher when we assumed that the active-case  $R_{\text{eff}}$  characterised local transmission at the time of forecast and that, over the forecast horizon, local transmission would trend back towards the whole-population transmission potential (whether  $R_{\text{eff}}$  was lower or higher than the TP). This improvement in skill was driven by characteristics of both the Victorian and New South Wales epidemics. In Victoria, while  $R_{\text{eff}}$  remained consistently higher than the transmission potential, it decreased steadily over time<sup>7</sup>. In New South Wales, the public health response to incursions repeatedly drew  $R_{\text{eff}}$  below the TP<sup>7</sup>. Although our model overestimated the rate at which local transmission would trend back towards the TP, it consistently out-performed forecasts generated under the assumption that local transmission would be sustained at the current active-case  $R_{\text{eff}}$ .

Forecast skill was also higher when the model was conditioned on the most recent data, even though these data were known to be incomplete. By imputing symptom onset dates and estimating the delay between symptom onset and case ascertainment<sup>7</sup>, we obtained additional information to that available in the reported time-series data, improving forecast skill. However, when the reported number of daily cases were low, forecast performance was particularly sensitive to the imputed case counts and there were occasions where forecast skill was higher when the model was not conditioned on the most recent data.

**Study strengths.** Australian jurisdictions maintained high testing levels over the study period. The proportion of infected persons that were identified as cases was likely to be both very high, and to remain relatively constant over the study period. These cases include asymptomatic individuals who were identified through large-scale and systematically exhaustive contact tracing programs<sup>5</sup>. In comparison, in recent influenza seasons the probability that a person with influenza-like illness (ILI) symptoms would seek healthcare and have a specimen collected for testing (as estimated from Flutracking<sup>32,33</sup> survey participants) was around 3–8%<sup>16</sup>. But during the study period, 50–75% of Flutracking participants with ILI symptoms reported having a COVID-19 test (Sect. S5), suggesting high case ascertainment. The end of the second wave in Victoria provides yet further evidence of high case ascertainment. As the 14-day case thresholds for easing restrictions were achieved, there were fewer and fewer unlinked (“mystery”) cases reported. For example, the Victorian 19 October update reported an average of 7.7 cases per day in the past 14 days and a total of just 15 unlinked cases<sup>34</sup>, while the 26 October update reported an average of 3.6 cases per day in the past 14 days and a total of just 7 unlinked cases<sup>35</sup>. Elimination was achieved before the end of the study period, and so unidentified cases, if present, did not cause sustained chains of transmission. This is consistent with the TP, which was estimated to be less than one in Victoria from August to October 2020<sup>7</sup>.

As described above, we had access to detailed line-listed data, and were able to estimate the delay distribution for the time between symptom onset and case ascertainment, and to impute the symptom onset date for cases where this had not yet been reported<sup>7</sup>.

We were able to incorporate the impact of public health interventions via near-real-time estimates of the effective reproduction number ( $R_{\text{eff}}$ ) and the whole-population transmission potential (TP), which were informed by nationwide behavioural surveys, population mobility data, and times from symptom onset to case detection<sup>7,19</sup>. Upon reintroduction of cases in jurisdictions that had achieved local elimination, estimates of the transmission potential allowed us to produce forecasts that were informed by local mobility and behaviour, before reliable estimates of the effective reproduction number were available.

Likewise, we were able to adapt the forecast outputs to support policy decisions in near-real-time for individual jurisdictions without incorporating additional effects into the underlying simulation model. This meant that we did not need to explicitly model control measures that were imposed and relaxed in each jurisdiction

— such an effort would have been extremely complex and it would be impossible to account for decisions that occur within each forecast horizon and their impact on transmission dynamics. The fact that our forecasts exhibited high forecast skill, particularly during periods of sustained local transmission, demonstrates that increasing model complexity is not always necessary or appropriate<sup>12,22,36,37</sup>.

**Study limitations.** Australia had experienced limited COVID-19 activity prior to the study period, and so it was reasonable to assume that the entire population was susceptible to infection. Subsequent waves of local infection, as experienced in Australia after the study period, have resulted in a non-negligible proportion of the population having some degree of natural immunity. In these circumstances, the forecasts would become more sensitive to assumptions about case ascertainment. COVID-19 vaccination in Australia began in February 2021 (i.e., after the study period) and has mitigated the impact of subsequent circulation of the Delta and Omicron variants in 2021 and 2022. By February 2022, more than 94% of people over the age of 16 were fully vaccinated<sup>38</sup>. We have continued to refine and extend the methods presented here to account for changes in case ascertainment, vaccination coverage, booster vaccinations, and waning immunity. Results obtained with these refinements continue to be reported under Australia's National Disease Surveillance Plan for COVID-19<sup>39</sup> and will be described elsewhere.

Our forecast model only incorporated population heterogeneity by allowing local transmission to trend from the active-case  $R_{\text{eff}}$  to the whole-population transmission potential (TP) over the forecast horizon. This limited granularity was appropriate for the purposes to which the forecasts were applied, but meant that the forecasts were not suitable for addressing more nuanced questions about, for example, targeted (geographic or socio-demographic) interventions to mitigate localised clusters<sup>40</sup>.

**Meaning and implications.** The first wave of COVID-19 in Australia was primarily driven by returning travellers who had acquired their infection overseas. In contrast, the second wave in Victoria was driven by local transmission. Cases in the second wave were also over-represented in essential services that involved frequent and intense contact, such as residential aged care facilities, healthcare, manufacturing, and meat processing<sup>5</sup>, and so were less able to reduce their contacts (and so transmission opportunities) than the general population. Accordingly, during the second wave in Victoria the active-case  $R_{\text{eff}}$  was systematically higher than the whole-population transmission potential<sup>7</sup>, even though both decreased over time in response to public health interventions. Forecast bias was compounded by this disparity between  $R_{\text{eff}}$  and transmission potential.

However, since the active-case  $R_{\text{eff}}$  did decrease over time, forecast skill was highest when, over the forecast horizon, local transmission trended back towards the whole-population transmission potential. It is unclear how future trends in local transmission could be better modelled or predicted, highlighting fundamental challenges in regards to (a) understanding the transmission characteristics of persons who are currently infectious; and (b) predicting how the transmission characteristics of infectious persons will change into the future, particularly when case numbers are low. These are intrinsic sources of uncertainty for infectious diseases forecasting, and while they may conceivably be reduced, they cannot be eliminated.

There is an acknowledged need for effective and on-going collaborations between modellers and public health practitioners in order to take full advantage of modelling tools such as epidemic forecasts to support public health policy and decision-making<sup>41–43</sup>. The forecasts presented here were adapted from our Australian seasonal influenza forecasts<sup>14–16</sup>, which we have deployed in near-real-time in collaboration with public health colleagues since 2015<sup>17,18</sup>. A number of modifications were required to adapt these forecasts to COVID-19 transmission in Australia, but developing and applying these methods to seasonal influenza activity in previous years has provided us with valuable experience in adapting to unprecedented circumstances, working with incomplete and evolving data, responding to public health needs, and effectively communicating forecast outputs to public health decision-makers.

## Methods

**Ethics statement.** The study was undertaken as urgent public health action to support Australia's COVID-19 pandemic response. The study used data from the Australian National Notifiable Disease Surveillance System (NNDSS) provided to the Australian Government Department of Health and Aged Care under the National Health Security Agreement for the purposes of national communicable disease surveillance. Data from the NNDSS were supplied after de-identification to the investigator team for the purposes of provision of epidemiological advice to government. Contractual obligations established strict data protection protocols agreed between the University of Melbourne and sub-contractors and the Australian Government Department of Health and Aged Care, with oversight and approval for use in supporting Australia's pandemic response and for publication provided by the data custodians represented by the Communicable Diseases Network of Australia. The ethics of the use of these data for these purposes, including publication, was agreed by the Department of Health and Aged Care with the Communicable Diseases Network of Australia. All methods were carried out in accordance with the relevant guidelines and regulations.

**COVID-19 surveillance data.** We used line-lists of reported cases for each Australian state and territory, which were extracted from the National Notifiable Disease Surveillance System (NNDSS) on a weekly basis. The line-lists contained the date when the individual first reported exhibiting symptoms, the date when the case notification was received by the jurisdictional health department, and whether the infection was acquired locally or overseas.

		Description	Value
(i)	$N$	The population size	Table 4
	$R_{\text{eff}}(t)$	The time-varying effective reproduction number	See text
	$j$	The index of the selected $R_{\text{eff}}(t)$ trajectory	$\sim \mathcal{U}(1, N_{px})$
	$\sigma$	The inverse of the latent period (days <sup>-1</sup> )	See text
	$\gamma$	The inverse of the infectious period (days <sup>-1</sup> )	See text
	$\tau$	The time of the initial exposures (days)	$\sim \mathcal{U}(0, 50)$
	$\hat{e}(t)$	The time-varying number of exogenous exposures	See text
	$\Delta t$	The time-step size (days)	0.01
(ii)	$b_{\text{gobs}}$	The background observation rate	0.05
	$p_{\text{obs}}$	The observation probability	0.8
	$k$	The dispersion parameter	10
	$p_{\text{det}}$	The detection probability (per observation)	See text
(iii)	$N_{px}$	The number of particles	2000
	$N_{\text{min}}$	The minimum number of effective particles	$0.25 \cdot N_{px}$

**Table 3.** Parameter values for (i) the transmission model; (ii) the observation model; and (iii) the bootstrap particle filter.

We used a time-varying delay distribution to characterise the duration between symptom onset and case notification, which was estimated from the reporting delays observed during the study period<sup>7</sup>. Symptom onset dates were imputed for cases where this was not (yet) reported.

**Forecasting COVID-19 case numbers in the community.** We used a discrete-time stochastic SEIIR model with gamma-distributed waiting times to characterise infection in each Australian jurisdiction. Let  $S(t)$  represent the number of *susceptible* individuals,  $E_1(t) + E_2(t)$  represent the number of *exposed* individuals,  $I_1(t) + I_2(t)$  represent the number of *infectious* individuals, and  $R(t)$  the number of *removed* individuals, at time  $t$ . Symptom onset was assumed to coincide with the transition from  $I_1$  to  $I_2$ . We used real-time estimates of the effective reproduction number  $R_{\text{eff}}^7$ , which were informed by nationwide behavioural surveys, mobility data from technology companies, and times from symptom onset to case detection<sup>7,19</sup>, and sampled a unique trajectory  $R_{\text{eff}}^j(t)$  for each particle  $j$ . In order to initialise the first epidemic wave in each jurisdiction, it was assumed that 10 exposures were introduced into the  $E_1$  compartment at time  $\tau$ , to be inferred, giving initial conditions:

$$\begin{aligned}
 S(0) &= N - E_1(0) \\
 E_1(0) &= 10 \\
 E_2(0) &= 0 \\
 I_1(0) &= 0 \\
 I_2(0) &= 0 \\
 R(0) &= 0 \\
 \sigma(t) &= \begin{cases} 0 & \text{if } t < \tau \\ \sigma & \text{if } t \geq \tau \end{cases} \\
 \gamma(t) &= \begin{cases} 0 & \text{if } t < \tau \\ \gamma & \text{if } t \geq \tau \end{cases} \\
 \beta_j(t) &= R_{\text{eff}}^j(t) \cdot \gamma(t).
 \end{aligned}$$

The number of individuals who leave each compartment on each time-step  $\Delta t$  follows a binomial distribution (see Table 3 for parameter definitions):

$$\begin{aligned} \lambda(t) &= \frac{\beta(t) \cdot [I_1(t) + I_2(t)] + \hat{e}(t)}{N} \\ S^{\text{Pr}}(t) &= 1 - \exp(-\Delta t \cdot \lambda(t)) \\ S^{\text{out}}(t) &\sim \text{Bin}(S(t), S^{\text{Pr}}(t)) \\ E_1^{\text{Pr}}(t) &= 1 - \exp(-\Delta t \cdot 2 \cdot \sigma(t)) \\ E_1^{\text{out}}(t) &\sim \text{Bin}(E_1(t), E_1^{\text{Pr}}(t)) \\ E_2^{\text{Pr}}(t) &= 1 - \exp(-\Delta t \cdot 2 \cdot \sigma(t)) \\ E_2^{\text{out}}(t) &\sim \text{Bin}(E_2(t), E_2^{\text{Pr}}(t)) \\ I_1^{\text{Pr}}(t) &= 1 - \exp(-\Delta t \cdot 2 \cdot \gamma(t)) \\ I_1^{\text{out}}(t) &\sim \text{Bin}(I_1(t), I_1^{\text{Pr}}(t)) \\ I_2^{\text{Pr}}(t) &= 1 - \exp(-\Delta t \cdot 2 \cdot \gamma(t)) \\ I_2^{\text{out}}(t) &\sim \text{Bin}(I_2(t), I_2^{\text{Pr}}(t)) \\ S(t + \Delta t) &= S(t) - S^{\text{out}}(t) \\ E_1(t + \Delta t) &= E_1(t) + S^{\text{out}}(t) - E_1^{\text{out}}(t) \\ E_2(t + \Delta t) &= E_2(t) + E_1^{\text{out}}(t) - E_2^{\text{out}}(t) \\ I_1(t + \Delta t) &= I_1(t) + E_2^{\text{out}}(t) - I_1^{\text{out}}(t) \\ I_2(t + \Delta t) &= I_2(t) + I_1^{\text{out}}(t) - I_2^{\text{out}}(t) \\ R(t + \Delta t) &= R(t) + I_2^{\text{out}}(t) \end{aligned}$$

We allowed for the introduction of exogenous exposures, which may be acquired by individuals who, e.g., visit an external location where infectious individuals are present, or encounter an infectious individual from an external location who has travelled to this location. The expected number of exogenous exposures,  $\hat{e}(t)$ , is time-varying and independent trajectories can be sampled for each particle.

We modelled the relationship between model incidence and the observed daily COVID-19 case incidence ( $y_t$ ) using a negative binomial distribution with dispersion parameter  $k$ , since the data were non-negative integer counts and were over-dispersed when compared to a Poisson distribution. Let  $X(t)$  represent the state of the dynamic process and particle filter particles at time  $t$ , and  $x_t$  represent a realisation;  $x_t = (s_t, e_{1t}, e_{2t}, i_{1t}, i_{2t}, r_t, \sigma_t, \gamma_t, j_t, \hat{e}_t)$ . We assumed that cases were observed (i.e., were reported as a notifiable case) with probability  $p_{\text{obs}}$  when they transitioned from  $I_1$  to  $I_2$ . To account for incomplete observations due to, e.g., delays in testing and reporting for the most recent observations, we fitted a time-to-detection distribution and calculated the detection probability ( $p_{\text{det}}$ ) for each observation at each time  $t^7$ . In order to improve the stability of the particle filter for very low (or zero) daily incidence, we also allowed for the possibility of a very small number of observed cases that were *not* directly a result of the community-level epidemic dynamics ( $bg_{\text{obs}}$ ). Accordingly, we defined the likelihood  $\mathcal{L}(y_t | x_t)$  of obtaining the observation  $y_t$  from the particle  $x_t$  as:

$$\begin{aligned} \mathcal{L}(y_t | x_t) &\sim \text{NegBin}(p_{\text{det}} \cdot \mathbb{E}[y_t], k) \\ \mathbb{E}[y_t] &= (1 - n_{\text{inc}}(t)) \cdot bg_{\text{obs}} + n_{\text{inc}}(t) \cdot p_{\text{obs}} \cdot N \\ n_{\text{inc}}(t) &= \frac{I_2(t) + R(t) - I_2(t-1) - R(t-1)}{N} \\ \text{Var}[y_t] &= p_{\text{det}} \cdot \mathbb{E}[y_t] + \frac{(p_{\text{det}} \cdot \mathbb{E}[y_t])^2}{k} \end{aligned}$$

The value of the dispersion parameter  $k$  was chosen so that when zero individuals enter the  $I_2$  compartment over one day, the 50% credible interval has an upper bound of zero observed cases and the 95% credible interval has an upper bound of one observed case.

To generate projected case counts at each day, we used a bootstrap particle filter with post-regularisation<sup>44</sup> and a “deterministic” resampling method<sup>45</sup>, as previously described in the context of our Australian seasonal influenza forecasts<sup>14–18</sup>. We constructed a simulated trajectory  $\tilde{y}^i = \{\tilde{y}_0^i, \tilde{y}_1^i, \dots\}$  from each particle  $i$  in the ensemble by drawing a random sample  $\tilde{y}_t^i$  at each day  $t$ :

$$\tilde{y}_t^i \sim \text{NegBin}(\mu = \mathbb{E}[y_t^i], k).$$

**Parameters and model prior distributions.** Model and particle filter parameters are described in Table 3. Since Australia is one of the most urbanised countries in the world, for each jurisdiction we used capital city residential populations (including the entire metropolitan region, as listed in Table 4) in lieu of the residential population of each jurisdiction as a whole. Parameters  $\sigma$  and  $\gamma$  were sampled from a multivariate log-normal distribution that was defined to be consistent with a generation interval distribution with mean=4.7 and SD=2.9<sup>7,46</sup>, and we sampled independent  $R_{\text{eff}}^j(t)$  trajectories for each particle. The observation probability  $p_{\text{obs}}$  was fixed at 0.8, assuming that 80% of infections would be detected<sup>47</sup>.

Jurisdiction	N
Australian Capital Territory	410,199
New South Wales	5,730,000
Queensland	2,560,000
South Australia	1,408,000
Northern Territory	154,280
Tasmania	240,342
Victoria	5,191,000
Western Australia	2,385,000

**Table 4.** The population sizes used for each jurisdiction. Values were obtained from Population Australia data for 2019 (<https://www.population.net.au/>).

*Competing models.* We evaluated four competing models, by considering all combinations of:

1. Assuming that local transmission would trend back towards the whole-population transmission potential over the forecast horizon (“Trend to TP”), or that local transmission would instead be sustained at the active-case  $R_{\text{eff}}$  distribution over the forecast horizon (“Fixed Reff”); and
2. Conditioning on all daily case counts where the detection probability  $p_{\text{det}} \geq 0.5$ , or only conditioning on all daily case counts where  $p_{\text{det}} \geq 0.95$  (in order to reduce the influence of imputed symptom onset dates, and cases potentially not yet reported).

The labels used to identify each of these four models are listed in Table 5.

**Accounting for observed reintroduction.** We deliberately treated the *consequence* of observed reintroduction separately from the *potential* for reintroduction. Here our interest lies in introducing new exposures to reflect a small number of reported cases in a jurisdiction after a sufficiently long period of no reported cases in that jurisdiction. This captures importation events where the particle ensemble contains few, if any, particles with any exposed or infectious individuals, and is fundamentally different from predicting the probability of a future reintroduction into a jurisdiction that currently has no active COVID-19 cases. We can realise this by defining exogenous exposure trajectories that are non-zero for a short period prior to the observed reintroduction.

For a jurisdiction that reports one or more COVID-19 cases at time  $\tau$  ( $y_\tau$ ) after a prolonged period of no cases (e.g., 28 days or longer), we assume that the initial exposure(s) occurred  $\Delta = 5$  days prior to the case(s) being reported, and we account for the time-to-detection distribution ( $p_{\text{det}}$ ) and the observation probability ( $p_{\text{obs}}$ ):

$$\hat{\epsilon}(t) = \begin{cases} y_\tau \cdot (p_{\text{det}} \cdot p_{\text{obs}})^{-1} & \text{if } \tau = t + \Delta \\ 0 & \text{otherwise.} \end{cases}$$

We deliberately selected this approach instead of keeping  $\hat{\epsilon}(t)$  fixed at some positive value  $\epsilon$ , so that the forecasts for jurisdictions that had achieved local elimination did not suggest there would be subsequent epidemic activity without evidence of reintroduction. This decision was motivated by Australia’s international and domestic movement restrictions, which were effective in limiting community exposure from infected international arrivals and from travel between jurisdictions. The daily probability of an importation event was deemed to be sufficiently low that it would only serve to inflate the widest credible intervals in jurisdictions that had achieved elimination.

**Forecast targets and skill.** The forecast targets were the daily case counts (indexed by symptom onset date) for the 7 days prior to the data date (“back-cast”), for the data date itself (“now-cast”), and for the 28 days after the data date (“forecast”). Forecast performance was measured using Continuous Ranked Probability Scores (CRPS<sup>48</sup>). CRPS is a generalisation of the Mean Absolute Error (MAE) for probabilistic forecasts, which penalises forecasts for placing probability mass at a distance from the ground truth, and is one of the most widely used accuracy metrics for probabilistic forecasts. CRPS values increase as more probability mass is allocated to incorrect values, and as the distance between these incorrect values and the ground truth increases.

We calculated CRPS values for each observation  $y_\tau$  of COVID-19 cases on day  $\tau$ , given observations  $y_{0:t}$  for days  $0 \dots t$  (inclusive):

	Trend to TP	Fixed $R_{\text{eff}}$
$p_{\text{det}} \geq 0.5$	Default	Fixed Reff
$p_{\text{det}} \geq 0.95$	Minimal imputation	Minimal imputation, Fixed Reff

**Table 5.** The four competing forecast models.

$$P(Y_\tau \leq y \mid y_{0:t}) = \sum_{i=1}^{N_{px}} w_i \cdot \mathcal{H}(y - \tilde{y}_\tau^i)$$

$$CRPS(y_\tau \mid y_{0:t}) = \sum_{y=0}^{Y_{max}} [P(Y_\tau \leq y \mid y_{0:t}) - \mathcal{H}(y - y_\tau)]^2$$

$$\mathcal{H}(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{otherwise.} \end{cases}$$

where  $\mathcal{H}$  denotes the Heaviside step function.

We compared the performance of each model  $\mathcal{M}$  to that of an historical benchmark forecast  $\mathcal{N}$ , whose probability mass function  $\mathcal{N}(t)$  for the data date  $t$  was the distribution of the daily case counts reported in the jurisdiction at the time of forecast, accounting for the detection probability  $p_{det}(\tau)$  for each observation  $y_\tau$ :

$$Y_{t+h} \sim \mathcal{N}(t) = \left\{ \left\lfloor \frac{y_\tau}{p_{det}(\tau)} \right\rfloor : \tau \in \{0, \dots, t\} \wedge p_{det}(\tau) \geq 0.5 \right\}$$

$$h \in \{-7, -6, \dots, 27, 28\}.$$

We measured the performance of each model  $\mathcal{M}$  using CRPS skill scores<sup>49</sup>:

$$\text{skill}(\mathcal{M}) = \frac{CRPS_{\mathcal{N}} - CRPS_{\mathcal{M}}}{CRPS_{\mathcal{N}}}.$$

We also measured forecast sharpness  $S(t)$  and forecast bias  $B(t)$  as per Funk et al.<sup>50</sup>:

$$S(t) = \frac{\text{median}(|\tilde{y}_t - \text{median}(\tilde{y}_t)|)}{0.675}$$

$$B(t) = 1 - [P(Y_t \leq y_t) + P(Y_t \leq y_t - 1)].$$

**Meeting case-number thresholds that trigger policy changes.** We have already defined the probability that the case counts on day  $\tau$  will not exceed some threshold  $T$ :  $P(Y_\tau \leq T \mid y_{0:t})$ . We can build on this definition to obtain the probability that the average daily case count over some window  $W$  up to, and including, day  $\tau$  does not exceed this threshold  $T$ :

$$P(\mathbb{E}[Y_{\tau-W+1:\tau}] \leq T \mid y_{0:t}) = \sum_{i=1}^{N_{px}} w_i \cdot \mathcal{H}\left(W \cdot T - \sum_{j=0}^W \tilde{y}_{\tau-j}^i\right).$$

This was used to predict the probability of achieving each of the target case thresholds for relaxing measures in the second wave in Victoria, for each day in the forecast horizon.

## Data availability

All of the simulation code and forecast outputs used to generate the results in this manuscript and the supplementary materials are provided in a public git repository (<https://gitlab.unimelb.edu.au/rgmoss/aus-2020-covid-forecasts>). See the provided README .md file for detailed instructions.

The input data for each data date (i.e., the daily COVID-19 case counts and  $R_{eff}$  trajectories) are provided in a separate dataset (<https://doi.org/10.26188/19315055>). For access to the raw COVID-19 case data, a request must be submitted to [NNDSS.datarequests@health.gov.au](mailto:NNDSS.datarequests@health.gov.au), which will be assessed by a data committee.

Received: 12 September 2022; Accepted: 19 May 2023

Published online: 30 May 2023

## References

1. Australian Government Department of Health. National review of hotel quarantine. <https://www.health.gov.au/resources/publications/national-review-of-hotel-quarantine> (2020). Accessed 12 Feb 2021.
2. Moss, R. *et al.* Coronavirus Disease Model to Inform Transmission Reducing Measures and Health System Preparedness, Australia. *Emerg. Infect. Dis.* **26**, 2844–2853. <https://doi.org/10.3201/eid2612.202530> (2020).
3. Price, D. J. *et al.* Early analysis of the Australian COVID-19 epidemic. *Elife* **9**, e58785. <https://doi.org/10.7554/eLife.58785> (2020).
4. COVID-19 National Incident Room Surveillance Team. COVID-19 Australia: Epidemiology Report 19. *Commun. Dis. Intell.* **44**, 54. <https://doi.org/10.33321/cdi.2020.44.54> (2020).
5. Victorian Department of Health COVID-19 writing group. Population-based analysis of the epidemiological features of COVID-19 epidemics in Victoria, Australia, January 2020 – March 2021, and their suppression through comprehensive control strategies. *The Lancet Regional Health – Western Pacific* **17**, 100297. <https://doi.org/10.1016/j.lanwpc.2021.100297> (2021).
6. COVID-19 National Incident Room Surveillance Team. COVID-19 Australia: Epidemiology Report 29. *Commun. Dis. Intell.* **44**, 87. <https://doi.org/10.33321/cdi.2020.44.87> (2020).
7. Golding, N. *et al.* A modelling approach to estimate the transmissibility of SARS-CoV-2 during periods of high, low, and zero case incidence. *Elife* **12**, e78089. <https://doi.org/10.7554/elife.78089> (2023).
8. Andrews, D. Statement from the Premier – 26 October 2020. <http://web.archive.org/web/20201026231047/https://www.dhhs.vic.gov.au/updates/coronavirus-covid-19/statement-premier-26-october-2020> (2020). Accessed 12 Feb 2021.

9. Salje, H. *et al.* Estimating the burden of SARS-CoV-2 in France. *Science* **369**, 208–211. <https://doi.org/10.1126/science.abc3517> (2020).
10. Keeling, M. J. *et al.* Predictions of COVID-19 dynamics in the UK: Short-term forecasting and analysis of potential exit strategies. *PLoS Comput. Biol.* **17**, 1–20. <https://doi.org/10.1371/journal.pcbi.1008619> (2021).
11. Biggerstaff, M., Slayton, R. B., Johansson, M. A. & Butler, J. C. Improving pandemic response: Employing mathematical modeling to confront Coronavirus Disease 2019. *Clin. Infect. Dis.* **74**, 913–917. <https://doi.org/10.1093/cid/ciab673> (2021).
12. Cramer, E. Y. *et al.* Evaluation of individual and ensemble probabilistic forecasts of COVID-19 mortality in the United States. *Proc. Natl. Acad. Sci.* **119**, e2113561119. <https://doi.org/10.1073/pnas.2113561119> (2022).
13. Paireau, J. *et al.* An ensemble model based on early predictors to forecast COVID-19 health care demand in France. *Proc. Natl. Acad. Sci.* **119**, e2103302119. <https://doi.org/10.1073/pnas.2103302119> (2022).
14. Moss, R., Zarebski, A., Dawson, P. & McCaw, J. M. Forecasting influenza outbreak dynamics in Melbourne from Internet search query surveillance data. *Influ. Other Respir. Viruses* **10**, 314–323. <https://doi.org/10.1111/irv.12376> (2016).
15. Moss, R., Zarebski, A., Dawson, P. & McCaw, J. M. Retrospective forecasting of the 2010–2014 Melbourne influenza seasons using multiple surveillance systems. *Epidemiol. Infect.* **145**, 156–169. <https://doi.org/10.1017/S0950268816002053> (2017).
16. Moss, R., Zarebski, A. E., Carlson, S. J. & McCaw, J. M. Accounting for healthcare-seeking behaviours and testing practices in real-time influenza forecasts. *Trop. Med. Infect. Dis.* **4**, 12. <https://doi.org/10.3390/tropicalmed4010012> (2019).
17. Moss, R. *et al.* Epidemic forecasts as a tool for public health: interpretation and (re)calibration. *Aust. N. Z. J. Public Health* **42**, 69–76. <https://doi.org/10.1111/1753-6405.12750> (2018).
18. Moss, R. *et al.* Anatomy of a seasonal influenza epidemic forecast. *Commun. Dis. Intell.* **43**, 1–14. <https://doi.org/10.33321/cdi.2019.43.7> (2019).
19. Golding, N. *et al.* Estimating temporal variation in transmission of SARS-CoV-2 and physical distancing behaviour in Australia. Tech. Rep., The Peter Doherty Institute for Infection and Immunity (2020). Accessed 12 Feb 2021.
20. Berger, L. *et al.* Rational policymaking during a pandemic. *Proc. Natl. Acad. Sci.* **118**, e2012704118. <https://doi.org/10.1073/pnas.2012704118> (2021).
21. Reich, N. G., Tibshirani, R. J., Ray, E. L. & Rosenfeld, R. On the predictability of COVID-19. *International Institute of Forecasters: Forecasting Blog* (2021). Accessed 04 May 2022.
22. Moss, R. Commentary on “Transparent modeling of influenza incidence”: Because the model said so. *Int. J. Forecast.* **38**, 620–621. <https://doi.org/10.1016/j.ijforecast.2021.01.028> (2022).
23. Lane, C. R. *et al.* Genomics-informed responses in the elimination of COVID-19 in Victoria, Australia: An observational, genomic epidemiological study. *Lancet Public Health* **6**, e547–e556. [https://doi.org/10.1016/S2468-2667\(21\)00133-X](https://doi.org/10.1016/S2468-2667(21)00133-X) (2021).
24. Victorian Department of Health and Human Services. Coronavirus update for Victoria - 3 July 2020. <http://web.archive.org/web/20211010161409/https://www.dhhs.vic.gov.au/coronavirus-update-victoria-03-july-2020> (2020). Accessed 28 Feb 2022.
25. Victorian Department of Health and Human Services. Coronavirus update for Victoria - 9 July 2020. <http://web.archive.org/web/20220306092718/https://www.dhhs.vic.gov.au/coronavirus-update-victoria-9-july-2020> (2020). Accessed 06 Mar 2022.
26. ABC News. 1000 ADF personnel deployed to help coronavirus response in Victoria. *ABC News* (2020). Accessed 04 May 2022.
27. Kinsella, E. & Handley, E. Melbourne’s coronavirus restrictions roadmap: The five steps from stage 4 lockdown to ‘COVID normal’. *ABC News* (2020). Accessed 28 Feb 2022.
28. Wilke, C. O. & Bergstrom, C. T. Predicting an epidemic trajectory is difficult. *Proc. Natl. Acad. Sci.* **117**, 28549–28551. <https://doi.org/10.1073/pnas.2020200117> (2020).
29. Seale, H. *et al.* COVID-19 is rapidly changing: Examining public perceptions and behaviors in response to this evolving pandemic. *PLoS ONE* **15**, e0235112. <https://doi.org/10.1371/journal.pone.0235112> (2020).
30. Pickles, K. *et al.* COVID-19 misinformation trends in Australia: Prospective longitudinal national survey. *J. Med. Int. Res.* **23**, e23805. <https://doi.org/10.2196/23805> (2021).
31. McCaffery, K. J. *et al.* Health literacy and disparities in COVID-19-related knowledge, attitudes, beliefs and behaviours in Australia. *Public Health Res. Pract.* **30**, e30342012. <https://doi.org/10.17061/phrp30342012> (2021).
32. Dalton, C. B., Carlson, S. J., Butler, M. T., Elvidge, E. & Durrheim, D. N. Building influenza surveillance pyramids in near real time, Australia. *Emerg. Infect. Dis.* **19**, 1863–1865. <https://doi.org/10.3201/eid1911.121878> (2013).
33. Dalton, C. B. *et al.* Insights from Flutracking: Thirteen tips to growing a web-based participatory surveillance system. *JMIR Public Health Surveill.* **3**, e48. <https://doi.org/10.2196/publichealth.7333> (2017).
34. Victorian Department of Health and Human Services. Coronavirus update for Victoria - 19 October 2020. <http://web.archive.org/web/20211022132911/https://www.dhhs.vic.gov.au/coronavirus-update-victoria-19-october-2020> (2020). Accessed 28 Feb 2022.
35. Victorian Department of Health and Human Services. Coronavirus update for Victoria - 26 October 2020. <http://web.archive.org/web/20211009202511/https://www.dhhs.vic.gov.au/coronavirus-update-victoria-26-october-2020> (2020). Accessed 28 Feb 2022.
36. Sherden, W. A. *The Fortune Sellers: The Big Business of Buying and Selling Predictions* (Wiley, 1997).
37. Castle, J. L., Doornik, J. A. & Hendry, D. F. Forecasting principles from experience with forecasting competitions. *Forecasting* **3**, 138–165. <https://doi.org/10.3390/forecast3010010> (2021).
38. Australian Government Department of Health. COVID-19 vaccine rollout update – 27 February 2022. <https://www.health.gov.au/resources/publications/covid-19-vaccine-rollout-update-27-february-2022> (2022). Accessed 28 Feb 2022.
39. Australian Government Department of Health. Australian National Disease Surveillance Plan for COVID-19. <https://www.health.gov.au/resources/publications/australian-national-disease-surveillance-plan-for-covid-19> (2022). Accessed 01 Aug 2022.
40. Zachreson, C. *et al.* Risk mapping for COVID-19 outbreaks in Australia using mobility data. *J. R. Soc. Interface* **18**, 20200657. <https://doi.org/10.1098/rsif.2020.0657> (2021).
41. Muscatello, D. J. *et al.* Translation of real-time infectious disease modeling into routine public health practice. *Emerg. Infect. Dis.* **23**. <https://doi.org/10.3201/eid2305.161720> (2017).
42. Doms, C., Kramer, S. C. & Shaman, J. Assessing the use of influenza forecasts and epidemiological modeling in public health decision making in the United States. *Sci. Rep.* **8**, 12406. <https://doi.org/10.1038/s41598-018-30378-w> (2018).
43. Shearer, F., Moss, R., McVernon, J., Ross, J. V. & McCaw, J. M. Incorporating decision analysis in pandemic planning and response. *PLoS Med.* **17**, e1003018. <https://doi.org/10.1371/journal.pmed.1003018> (2020).
44. Musso, C., Oudjane, N. & Le Gland, F. Improving regularised particle filters. In *Sequential Monte Carlo Methods in Practice*, chap., Vol. 12, 247–271 (Springer, 2001). [https://doi.org/10.1007/978-1-4757-3437-9\\_12](https://doi.org/10.1007/978-1-4757-3437-9_12).
45. Kitagawa, G. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *J. Comput. Graph. Stat.* **5**, 1–25. <https://doi.org/10.1080/10618600.1996.10474692> (1996).
46. Golding, N. *et al.* Estimating temporal variation in transmission of COVID-19 and adherence to social distancing measures in Australia. Tech. Rep., The Peter Doherty Institute for Infection and Immunity (2020). Accessed 12 Feb 2021.
47. Mizumoto, K., Kagaya, K., Zarebski, A. E. & Chowell, G. Estimating the asymptomatic proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan, 2020. *Eurosurveillance* **25**, pii=2000180. <https://doi.org/10.2807/1560-7917.ES.2020.25.10.2000180> (2020).
48. Gneiting, T. & Katzfuss, M. Probabilistic forecasting. *Annu. Rev. Stat. Appl.* **1**, 125–151. <https://doi.org/10.1146/annurev-statis-062713-085831> (2014).
49. Hyndman, R. J. & Athanasopoulos, G. *Forecasting: principles and practice* (OTexts, Melbourne, Australia, 2021), 3rd edn. Accessed on 20 Jan 2021.

50. Funk, S. *et al.* Assessing the performance of real-time epidemic forecasts: A case study of Ebola in the Western Area region of Sierra Leone, 2014–15. *PLoS Comput. Biol.* **15**, e1006785. <https://doi.org/10.1371/journal.pcbi.1006785> (2019).

## Acknowledgements

National Notifiable Disease Surveillance System data on COVID-19 were provided by the Office of Health Protection and Response, Australian Government Department of Health and Aged Care, on behalf of the Communicable Diseases Network Australia (CDNA). We thank public health staff from incident emergency operations centres in state and territory health departments, and the Australian Government Department of Health and Aged Care, along with state and territory public health laboratories. We thank members of CDNA for their feedback and perspectives on the study results. We thank Craig Dalton, Sandra Carlson, and the Flutracking team for sharing their insights and Flutracking survey data throughout the study. We thank the Flutracking participants who made the time each week to complete the surveys and contribute to COVID-19 surveillance. This research was supported by use of the Nectar Research Cloud, a collaborative Australian research platform supported by the NCRIS-funded Australian Research Data Commons (ARDC), and by the University of Melbourne's Research Computing Services. This work was directly funded by the Australian Government Department of Health. Additional support was provided by the National Health and Medical Research Council of Australia through its Centres of Research Excellence (SPECTRUM, GNT1170960) and Investigator Grant Schemes (JMcV Principal Research Fellowship, GNT1117140).

## Author contributions

All authors contributed to the study design and interpretation. R.M. conducted the analysis and wrote the initial manuscript text. All authors reviewed and revised the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-35668-6>.

**Correspondence** and requests for materials should be addressed to R.M.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023