Using routinely collected data to inform infection-prevention policy decisions

Mark Geoffrey Pritchard

Worcester College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy in Clinical Medicine

Michaelmas Term, 2024

The diseases of the epidemic class follow laws of their own; they remain nearly stationary during months, years, and, as we learn from medical history, centuries; then suddenly rise, like a mist from the earth, and shed desolation on nations—to disappear as rapidly or insensibly as they came. The pestilences of ancient history, the plagues of England, cholera, influenza, small-pox, and typhus, are examples of this peculiar tribe of diseases. Epidemics have furnished much matter for discussion, and still offer large scope for inquiry.

William Farr, Causes of death in England and Wales (1840)

Contents

List of Figures

List of Tables

Abstract

Measures to reduce transmission are a vital response to infectious disease epidemics. Collectively such measures are effective in reducing the burden of infectious disease but effectiveness of individual interventions is less certain. Methodologies for causal inference from observational data are well developed, but many methods have requirements that are not met by epidemic data. They may require an individual's outcome to be independent of anyone else's treatment, but the very purpose of infection-prevention measures is to break chains of transmission, benefiting both treated and untreated individuals.

I combine causal inference methods, mechanistic models, and observational data to estimate effects of interventions that were used to reduce the spread of severe acute respiratory syndrome coronavirus 2 in the United Kingdom. I combine difference-in-differences methodology with a renewal-equation model. If its assumptions are met, this can detect effects of interventions on transmissibility, but if assumptions are violated, erroneous results can arise with no indication that an error is occurring. I apply the method to mass testing and mandatory use of face masks. Difference-in-differences results suggest that interventions increased incidence of detected infections.

I investigate optimal timing of vaccination against respiratory viral infections with models incorporating immune boosting from re-exposure to the virus. Boosting can lead to synchrony in susceptibility and cause periodic outbreaks even without seasonal variation in infectiousness. In scenarios with more immune boosting, vaccinating sooner tends to lead to fewer infections, while in scenarios with less boosting, later vaccination is beneficial.

Analyses in this thesis highlight potential problems with causal analyses that disregard mechanisms of disease transmission, and with models that oversimplify immunity. These analyses suggest that greater understanding of changing immunity over time is necessary to determine optimal approaches to reducing transmission of these respiratory viral infections.

Acknowledgements

I am extremely grateful to my supervisors, Ben Cooper, Peter Horby and Christl Donnelly for their insight and support throughout my DPhil studies. On the day I met Ben, he told me that the best thing about a research degree was the opportunity to explore rabbit holes and see where they took you. I have thoroughly enjoyed the opportunity to spend three years following my curiosity, and am very thankful for the times he helped me find my way back out of the warren.

I have been fortunate to work with amazing people across the university's Centre for Tropical Medicine and Global Health, the Pandemic Sciences Institute and the Department of Statistics. Special thanks go to Sean Cavany, Cherry Lim and Myo Maung Maung Swe for helpful discussions on causal inference, Bayesian statistics and mathematical modelling. The New Richards Building may be no more, but I will always have fond memories of the time we shared there.

I thank David Eyre and Koen Pouwels for their help throughout this process, advising on my thesis committee then acting as examiners for my transfer and confirmation of status milestones. I thank Susanna Dunachie and Graham Medley for helping me to make my thoughts on immune boosting more coherent.

I am grateful to the National Institute for Health Research Health Protection Research Unit in Emerging and Zoonotic Infections, not only for the funding that made this work possible, but for the excellent educational and networking programme that supported me throughout my studies. Lance Turtle and Andrew Carey truly made the Health Protection Research Unit feel like a community. I am also grateful to the Oxford School of Public Health, and particularly to Jill Morris and Jane Bishop, for their support to take time out of public health training to pursue this DPhil.

Finally, my undying thanks to Suzanne. As I follow my meandering path through life, you make sure I never come adrift. It would be clichéd to say I couldn't have done this without you. But then, some clichés are true.

This study is funded by the National Institute for Health Research Health Protection Research Unit in Emerging and Zoonotic Infections (200907), a partnership between the United Kingdom Health Security Agency, the Uni-

versity of Liverpool, the University of Oxford and the Liverpool School of Tropical Medicine. The views expressed are those of the authors and not necessarily those of the National Institute for Health Research, the United Kingdom Health Security Agency or the Department of Health and Social Care.

I acknowledge the use of the University of Oxford Advanced Research Computing facility in carrying out this work, `http://dx.doi.org/10.5281/zenodo.22558`.

Acronyms and notation

ATT	average treatment effect in the treated
CI	confidence interval
CrI	credible interval
MCMC	Markov chain Monte Carlo
NHS	National Health Service
ODE	ordinary differential equation
RCT	randomized controlled trial
RSV	respiratory syncytial virus
SARS-COV-2	severe acute respiratory syndrome coronavirus 2
SEIR	susceptible–exposed–infectious–resistant
SIR	susceptible–infectious–resistant
SIRS	susceptible–infectious–resistant–susceptible
SIS	susceptible–infectious–susceptible
UK	United Kingdom

Only one genus name is repeated in abbreviated form in this thesis:
B. Bordetella

Symbols for parameters and model compartments are used consistently where possible but are defined for each chapter and may differ between chapters. The following notation is used consistently:

$\bar{x}$	mean value of x
$\hat{x}$	estimated or fitted value of x
$x \in (a, b)$	$a < x < b$
$x \in [a, b]$	$a \leq x \leq b$
$\mathbb{E}(\cdot)$	expected value or mean
e	Euler's number, $\mathrm{e} = \sum_{x=0}^{\infty}(1/x!)$
i	imaginary unit, $\mathrm{i}^2 = -1$
$\mathbb{N}$	the natural numbers, $\{1, 2, 3, \dots\}$
$\mathcal{N}(\mu, \sigma^2)$	normally distributed with mean μ and variance σ^2
$\Pr(\cdot)$	Probability
$\mathcal{R}_0$	basic reproduction ratio
$\mathcal{R}_e(t)$	effective reproduction ratio at time t

Note that some models have compartments labelled e or i. Context should

prevent confusion with Euler's number or the imaginary unit, and compartment labels are always italicized.

When describing the value that a variable takes, upper case letters are used to name the variable and lower case letters for the value the variable takes. For instance, $X = x$ should be read as the variable named X has a value x. Superscripts are used to define potential outcomes (described on page 17), so that $Y^{A=a}$ represents the potential outcome for Y in the scenario where variable A has the value a.

In discrete-time models, subscript values index time rather than indicating a measure of time, *i.e.* x_1 should be interpreted as the discrete-time equivalent of $x(t_1)$, the value of x when $t = t_1$, and the next value of x is x_2 regardless of the magnitude of the interval between t_1 and t_2.

Chapter One

Introduction

If the purpose of biomedical research is to effect improvements in health then it is vital to have good understanding of causal relationships between interventions and outcomes. Ideally we wish to be able to tell decision makers – be they government officials or private individuals – what outcome to expect without an intervention and what to expect if they intervene. Any intervention has opportunity costs (resources could have been spent elsewhere) and risks of harm. Decision makers therefore need reliable estimates of the effectiveness of interventions.

Interventions explicitly intended to reduce the spread of infectious disease have been used since at least the seventeenth century when isolation, quarantine and beak-shaped face masks were adopted to reduce the spread of *Yersinia pestis*/plague (Coleman, 1986; Pan *et al.*, 2020; Shaw-Taylor, 2020). Together with improvements in nutrition, housing and healthcare, infection-prevention measures have played a vital role in improving human health (Mackenbach, 2021). The burden of disease has shifted from a situation in which most people died of infectious diseases, often in childhood, to one where most people can expect to live to old age and only a minority of deaths are attributed to infections (GBD 2013 Mortality and Causes of Death Collaborators, 2015; Shaw-Taylor, 2020).

Despite evidence for their collective importance, individual interventions' effects are less certain (Aiello and Larson, 2002; Dancer, 2016). Some interventions clearly offer large benefits for small costs and a low risk of harm. Estimated effects of handwashing vary between studies and are subject to possible biases, but there is general consensus that improving hand hygiene produces net health benefits (Freeman *et al.*, 2014; Mo *et al.*, 2022). If an intervention's benefits are smaller, they might not justify the costs, or may even fail to outweigh its harms. If public-health practice is to be evidence-based, we need to understand the magnitude of proposed interventions' effects.

Routine practice in community communicable-disease control includes advising people with infections such as measles, whooping cough and infectious diarrhoea, that they must stay away from work or school for a period of time (Hawker *et al.*, 2019). Local authorities in England can apply to have someone involuntarily detained to protect public health but these powers are very rarely used, almost exclusively to reduce spread of multidrug-resistant tuberculosis (Arnold *et al.*, 2019). Most routine isolation therefore takes place with affected individuals' consent. If people are to be compelled to comply with public-health policy, the Nuffield Council on Bioethics (2007) argues that there is a stronger requirement for evidence of effectiveness. Such compulsion has been legislated during epidemics and pandemics such as the influenza pandemic of 1919 (Bootsma and Ferguson, 2007), the Ebola epidemic in Guinea, Liberia and Sierra Leone between 2013 and 2016 (Farmer, 2020), and the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2)/covid-19 pandemic since 2020 (Hale *et al.*, 2021). Policy makers need reliable evidence to guide decisions about such interventions but evaluating measures during a public-health emergency may be even more challenging than evaluating routine measures.

Randomized controlled trials (RCTs) are experimental studies designed to answer causal questions. Well-conducted RCTs can provide unbiased estimates of causal effects (Rubin, 1974). Collins *et al.* (2020) refer to this as 'the magic of randomization'. They argue that the risk of drawing biased conclusions from observational studies means that observational data cannot be substituted for RCT results. Frameworks to assess quality of evidence typically rank results from RCTs higher than those from observational studies (*e.g.* Guyatt *et al.*, 1995, 2008), and new medicinal products must generally be assessed in RCTs before their use is approved (Stratford *et al.*, 2020). Many other – non-pharmaceutical, or physical – interventions have never been tested in trials. Ogilvie *et al.* (2005) found evidence for an 'inverse evidence law', in which interventions with the greatest effect on population health may be least likely to have high-quality evidence of effectiveness.

Transmission of microparasites[1] – viruses, bacteria, fungi and protozoa – requires them to move from an infectious host into tissues of a susceptible individual. Infection-prevention measures may affect numbers of infectious and susceptible individuals, or disrupt microparasites' routes between them.

Some infection-prevention measures use medicinal products to reduce spread of infection. Effective treatment can reduce the number of infectious individuals. Usually prevention of further cases is secondary to the objective of reducing symptoms, severity or sequelae for the person receiving treatment, but some treatments – such as antibiotics given to people with *Bordetella pertussis*/whooping cough – have a greater role in reducing transmission than in benefiting the recipient (Kilgore *et al.*, 2016). Vaccines given on a population-wide scale can protect recipients from infection and reduce population-wide transmission (Anderson and May, 1982; Halloran,

[1] In contrast to macroparasites (helminths and arthropods); Anderson and May (1979) and May and Anderson (1979) introduced the term 'microparasite' to emphasize parallels between infectious-disease dynamics and parasite ecology.

2019). And post-exposure medications can prevent microparasites from invading tissues, such as secondary prevention with rabies vaccine and immunoglobulin to prevent rabies virus from entering motor neurones (Fooks *et al.*, 2017).

Many measures to block routes of transmission are physical interventions. These are diverse. Some can be used by individuals to protect themselves, such as hand washing, insecticide-treated nets and barrier contraceptives. Others require changes in behaviour for the purpose of protecting others, such as requirements to isolate or stay away from school or work with infectious diarrhoea. Still others can only be conducted at a region-wide scale, such as provision of sewerage, and microbiological testing and contact-tracing processes. Measures may be used across communities, or may be limited to specific settings, such as hospitals, nurseries or care homes.

For most physical interventions, there is no evidence of effectiveness from RCTs (Dancer, 2016), mostly because they have never been tested in RCTs. Trials that have been conducted may have poor adherence to allocated treatment (Jefferson *et al.*, 2023) or imprecise descriptions of interventions (Hoffmann *et al.*, 2013). But an absence of evidence from trials does not remove the need to make decisions about potential interventions.

'Causal-inference' methods and mathematical models trade the *magic of randomization* for assumptions about unobserved underlying processes that connect observations. Inferring causal effects from observational data can be seen as controversial (Dowd, 2011). Pearl (2000) writes that there was a 'century of neglect within the study of causal inference' following Karl Pearson's objections to cause-and-effect conclusions (also see, Pearl and Mackenzie, 2018). The journal *JAMA* prohibits causal language outside reports of RCT results (Livingston, 2020). But causal conclusions will be drawn by read-

ers regardless of whether they are explicitly stated in a manuscript. Hernán (2018) argues that it is much more useful to describe causal assumptions and conclusions explicitly, allowing the audience to assess their validity, rather than ambiguously report associations and hope readers interpret them appropriately.

1.1 Motivation and outline of the thesis

Memories of the public-health response to the covid-19 pandemic remain fresh. The conditions of 2020 and 2021 are unlikely to repeat exactly, but it remains useful to assess the effects of interventions used in response to SARS-CoV-2. Firstly, those interventions may be suggested for future outbreaks of respiratory-virus infection and policy makers will need high-quality evidence. Secondly, retrospective analyses – removed from pressures of attempting to inform urgent decisions – can explore the usefulness of tools available to assess effectiveness of interventions. This can help to maximize researchers' ability to respond efficiently and reliably to future outbreaks.

In this thesis, I use causal-inference methods, mathematical models and publicly available data from the United Kingdom (UK) to explore how we can understand effects of interventions used to reduce the spread of covid-19.

For the remainder of this chapter, I introduce the background that this thesis aims to build upon. I summarize infection-prevention measures used to reduce transmission of respiratory-tract viruses and steps taken in the UK to reduce the spread of covid-19 (Section 1.2). In Section 1.3 I discuss the role of 'potential outcomes' in causal inference, and introduce notation and assumptions that I use to describe causal effects. Then I introduce concepts

related to mechanistic modelling of infectious disease dynamics that are used for analyses and simulations in later chapters (Section 1.4).

In chapter 2, I consider the role of the difference-in-differences causal inference method in infectious-disease epidemiology. I combine this method with a renewal-equation model to produce causal estimates of interventions' effects that are linked to our understanding of microparasitic-disease dynamics. I explore how violation of the method's assumptions in simulated datasets affects results. Then I apply the method to two publicly available datasets, estimating the effect of changes in policy regarding use of face coverings in the UK in 2020, and the effect of mass testing of asymptomatic people for SARS-CoV-2 in the Liverpool City region in late 2020.

In Chapter 3 I discuss 'natural immune boosting',[2] a hypothesized effect of exposure to a microparasite that prolongs immunity to reinfection. I use a simple compartmental epidemic model to explore the role that immune boosting could have for a microparasite that induces only brief immunity, as is common for viral infections of the respiratory tract. I apply the model to respiratory syncytial virus (RSV) diagnoses in Scotland before and during the period of social distancing measures introduced to reduce the spread of covid-19, and consider how our assumptions regarding immune boosting affect conclusions about the effectiveness of interventions.

In Chapter 4 I consider the implications of natural immune boosting on healthcare workers' immunity and the optimal timing of repeat vaccination. I use data from acute National Health Service (NHS) hospitals in the Midlands of England to inform a model of the effects of counterfactual vaccination programmes in the second half of 2021.

[2]The phrase 'natural immune boosting' is often used to refer to interventions in complementary and alternative medicine. My use of the phrase in this thesis refers to boosting by (unintended) exposure to the pathogen, in contrast to intentional immune boosting by repeat vaccination.

Table 1.1: Common respiratory viruses

		Syndromes†				
Virus	Classification*	Coryza	Pharyngitis	Croup	Brochiolitis	Pneumonia
Rhinovirus	A–C	+++	++	+	+	
Coronavirus	6 species‡	++	+	+		+
Respiratory syncytial virus	A–B	++	+	++	+++	+++
Human metapneumovirus	A–B	+	+	+	++	++
Parainfluenza	1–3, 4a, 4b	+	++	+++	+	++
Influenza	A–C	+	++	++	+	++
Adenovirus	A–G	+	++	++	++	++
Bocavirus	1–4	++	++	++	++	++

* Including species, groups, subgroups, types and serotypes.
† +++, represents the major cause; ++, a common cause; +, an occasional cause; and blank is a rare cause or not reported.
‡ Including severe acute respiratory syndrome coronavirus 2.
Data from Peiris, 2020, and for bocavirus from Chakraborty and Bhukya, 2023.

The final chapter contains conclusions from my work and suggestions for future areas of study.

1.2 Preventing viral infections of the respiratory tract

Respiratory-tract infections are ubiquitous. Adults typically experience two to three 'common colds' per year, and ten per cent of children experience twelve or more infections per year (Turner, 1997). Respiratory infections typically cause non-specific symptoms with overlapping syndromes of disease (Table 1.1). Clinically, it may not be possible to distinguish bacterial from viral infections, let alone identify specific pathogens (Lydon *et al.*, 2019).

Many infections of the upper respiratory tract are mild and pass unnoticed. Up to seventy per cent of viral respiratory infections in healthy adults cause no symptoms (Galanti *et al.*, 2019). But infections of the lower respiratory tract (pneumonia and bronchiolitis) are the main infectious cause

of disability-adjusted life years.[3] Their burden is particularly high in the first year of life and for adults over seventy years of age (GBD 2016 Lower Respiratory Infections Collaborators, 2018). RSV caused an estimated 100 thousand childhood deaths in 2019 (Li *et al.*, 2022) and may cause up to a quarter of excess winter deaths among adults (Thompson *et al.*, 2003). The influenza pandemic of 1918–1920 and, one hundred years later, the SARS-CoV-2 pandemic, each caused more deaths than the First World War,[4] demonstrating ongoing vulnerability to newly emergent strains of virus.

Many systemic infections (including measles and varicella zoster, which are transmitted via aerosols to the respiratory mucosa before replicating in other tissues; Gershon *et al.*, 2015; Moss and Griffin, 2012) cause long-term immunity, which may be lifelong. In unvaccinated populations, most susceptible individuals are children. Consequently these are often childhood infections. In contrast, viral infections of mucosa lead to transient immunity (Slifka and Ahmed, 1996). Reinfection is possible at all ages, and severe disease is common among the elderly (Thompson *et al.*, 2003). Vaccine-derived immunity to respiratory viruses is similarly temporary and frequent repeat vaccination may be recommended for people vulnerable to severe disease (UK Health Security Agency, 2022d, 2023).

Transmission routes

Transmission requires transfer of virions from one individual's respiratory mucosa to another's. Four main routes of transmission are classically described (Kutter *et al.*, 2018; Leung, 2021):

[3] Lower respiratory tract infections collectively before 2020, and SARS-CoV-2 specifically in 2020 and 2021; GBD 2021 Diseases and Injuries Collaborators, 2024.

[4] Approximately 8.5 million military deaths and 13 million civilian deaths caused by the First World War (Showalter and Royde-Smith, 2024); 50–100 million deaths attributed to the influenza pandemic, and 27 million to SARS-CoV-2 by the end of 2023 (Dattani, 2023)

1. direct contact,
2. indirect contact via fomites (objects contaminated with virions),
3. via droplets, sometimes termed 'droplet contact', and
4. airborne via aerosols.

Influenza virus, coronaviruses, rhinoviruses and bocavirus are also found in stools and may be transmitted via a faecal–oral route or may be aerosolized by flushing lavatories (Chakraborty and Bhukya, 2023; Leung, 2021). The relative importance of each route of transmission is uncertain, but infection-prevention guidance tends to recommend considering all routes as important for all respiratory viruses (Kutter *et al.*, 2018).

The first three routes listed above each involve a susceptible individual's hands being contaminated, followed by self-inoculation of their respiratory mucosa. The difference between these routes is whether their hands were contaminated from (1) an infectious person's hands, (2) an object that an infectious person had touched, or (3) an object that had been contaminated by respiratory droplets. Good hand hygiene is expected to interrupt all three routes (Mo *et al.*, 2022). Physical separation of people, either by distance or barriers, should also reduce direct contact and droplet transmission (Chu *et al.*, 2020).

Aerosols are small particles that can stay in the air for a long time and travel a long distance, being moved more by air flow than by gravity. Traditionally, an aerodynamic diameter of 5 µm was considered the threshold distinguishing aerosols from droplets. Behaviour of small particles depends on environmental conditions such as air movement and humidity, and particles up to 100 µm can behave as aerosols (Tellier *et al.*, 2019; Wang *et al.*, 2021). Aerosols can be inhaled, depositing virions into the nasopharynx or, for particles smaller than 10 µm, directly into alveoli (termed 'respirable' *vs.* 'in-

spirable' by Larson and Liverman, 2011). Aerosol transmission is possible via ventilation so does not require infectious and susceptible individuals to share a space (Hwang *et al.*, 2021). Use of face masks by both infectious and susceptible individuals is proposed to reduce aerosol transmission (Wang *et al.*, 2021).

Evidence of effectiveness of physical interventions

As with infection-prevention measures in general, there is substantial evidence that measures taken in response to the covid-19 pandemic were effective in reducing transmission of infectious disease. Hatoun *et al.* (2020) found that implementation of physical interventions in Massachusetts in March 2020 was associated with reductions in a range of viral and bacterial infections among children. The much smaller observed reduction in diagnoses of urinary-tract infections, which are generally not transmitted from person to person, suggested that the effect was due to changes in transmission and not solely due to difficulties accessing healthcare (*ibid.*). Further evidence of an effect came after the first year of interventions when reduced numbers of RSV infections were reported from multiple countries (Friedrich *et al.*, 2021; van Summeren *et al.*, 2021; Yeoh *et al.*, 2021).

Evidence for individual interventions is less certain. Of six systematic reviews in the Cochrane Library that are particularly relevant to this question (Table 1.2), only one included evidence from RCTs. Jefferson *et al.* (2023) including 78 RCTs of masks and hand hygiene. They found considerable uncertainty in results and a high risk of bias (especially due to a lack of blinding of participants to interventions). They found no evidence of a benefit from face masks. Hand hygiene was associated with a reduction in acute respiratory illness but not with a reduction in influenza-like illness or laboratory-confirmed influenza (*ibid.*). Viswanathan *et al.* (2020) identified

Table 1.2: Cochrane reviews relevant to non-pharmaceutical interventions to reduce the spread of viral infections of the respiratory tract

Citation	Interventions	Disease	Studies
Anglemyer *et al.*, 2020	Digital contact tracing	Epidemics	6 observational, 6 modelling
Nussbaumer-Streit *et al.*, 2020	Quarantine / isolation	Covid-19	8 observational, 43 modelling
Viswanathan *et al.*, 2020	Universal screening	Covid-19	17 observational, 5 modelling
Burnsa *et al.*, 2021	Restrictions on international travel	Covid-19	13 observational, 49 modelling
Kratzer *et al.*, 2022	School-based measures	Covid-19	3 laboratory, 5 observational, 5 modelling
Jefferson *et al.*, 2023	Masks and hand hygiene	ARI	78 RCT

ARI, acute respiratory infections; RCT, randomized controlled trial

two modelling studies of effectiveness of universal screening for covid-19 to reduce spread of transmission. All other studies in their review considered accuracy of screening tests but not their effect on disease transmission. All observational studies of restrictions on international travel identified by Burnsa *et al.* (2021) were assessments of numbers of infections detected by screening tests, with no discussion of onward transmission. Most observational studies of quarantine measures had no control group, and many were reports of situations in which multiple measures were being introduced, so these studies were unable to quantify the effect of quarantine (Nussbaumer-Streit *et al.*, 2020). In May 2020, Anglemyer *et al.* (2020) found few studies using digital contact tracing during outbreaks or epidemics. During the covid-19 pandemic, smartphone applications that identified proximity to other smartphones through Bluetooth connections were widely used. Kendall *et al.* (2023) estimate that in its first year of use in England and Wales, this application averted one million covid-19 infections and almost ten thousand deaths.

Jefferson *et al.* (2023) found that few trials assessed harms from interventions. Kratzer *et al.* (2022) were particularly interested in unintended effects of measures used in schools. They found the evidence base to be very patchy, concluding that, 'There is no deliberate, systematic research effort to identify, quantify, and/or theorise the unintended consequences of school-based SARS-CoV-2/covid-19 prevention measures' (*ibid.*).

The covid-19 pandemic in the United Kingdom

The first cases of pneumonia caused by the novel coronavirus that is now known as SARS-CoV-2, emerged in Wuhan, China, in December 2019 (Huang *et al.*, 2020). The first cases in the UK were diagnosed in England at the end of January 2020 (Wright, 2021). Infections were diagnosed in Northern Ireland and Wales in late February 2020, and in Scotland on 2 March (BBC News, 2020a,b,d).

Figure 1.1*A* shows recorded incidence from January 2020 to the end of December 2022. Availability of testing for SARS-CoV-2 in the UK increased dramatically during the pandemic. Initially tests were only available to individuals who met clinical and epidemiological criteria. By April 2021, rapid testing for SARS-CoV-2 antigen with lateral flow devices was freely available for anyone who wanted a test in England, Scotland or Wales. This free testing ended in 2023, as did the facility to self-report results (Hale *et al.*, 2021). Recorded incidence is therefore not expected to be a constant proportion of infections.

The Office for National Statistics (2023a,b) covid-19 infection survey invited participants randomly sampled from addresses in the UK, and the 'real-time assessment of community transmission (REACT)' study sampled individuals registered with NHS general practitioners (Ward *et al.*, 2023). Response rates were sometimes low (only twelve per cent of households

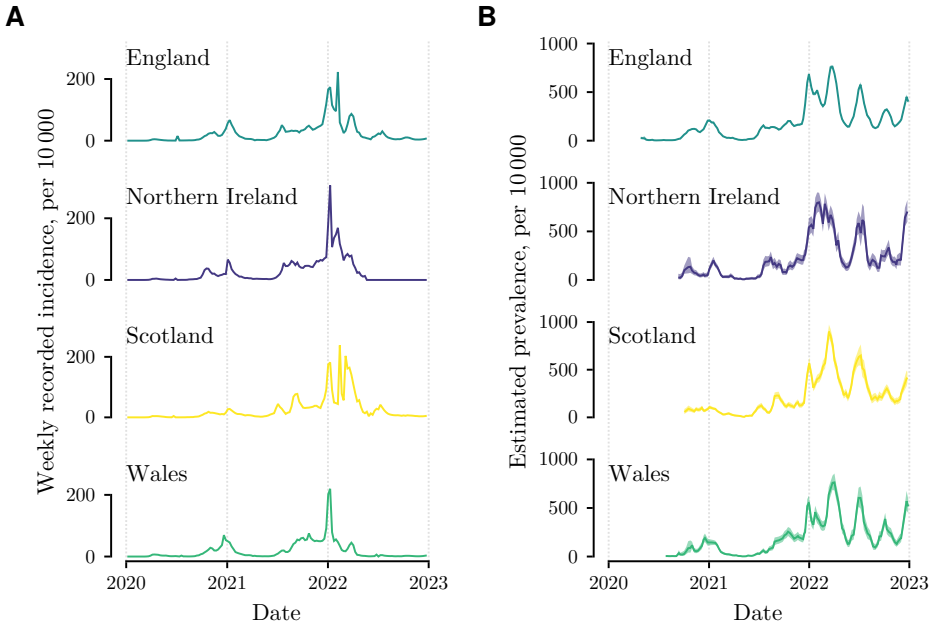

Figure 1.1: Covid-19 cases in the United Kingdom, 2020–2022. *A*. Weekly reported incidence per 10 000 people; calculated from cumulative confirmed case numbers in Hale *et al.*, 2021. *B*. Estimated prevalence, with 95% confidence intervals shaded; data from Office for National Statistics, 2023a.

in some rounds; Office for National Statistics, 2023c) so results needed to be adjusted to match the target population (Office for National Statistics, 2023b), but this should give a less-biased estimate of SARS-CoV-2 prevalence. Reported incidence is qualitatively similar to results from the covid-19 infection survey for 2020 and 2021 (Figure 1.1*B*), but in 2022 the survey revealed several peaks in prevalence that were not apparent in numbers of reported cases.

Measures taken against covid-19 in the United Kingdom

Health policy is devolved in the UK so each constituent nation could recommend or mandate different measures. All four nations introduced equivalent emergency regulations and support measures in March 2020. Their policies then diverged during late spring and summer 2020 as restrictions

were variously introduced and relaxed in response to changes in prevalence of infection.

From June 2020, districts within England had different levels of restriction, with 'local lockdowns' introduced where the highest levels of transmission were identified. 'Tiered' restrictions were introduced in October and again in December 2020, mandating different levels of restriction for different local authority areas. These were replaced by national (England-wide) restrictions in November 2020 and January 2021 (Brown and Kirk-Wade, 2021).

In December 2020, the first mass-vaccination programme against covid-19 started in the UK (Baraniuk, 2021). AstraZeneca COVID-19 ChAdOx1-S and Pfizer BioNTech COVID-19 BNT162b2 were the first vaccines in use (Baraniuk, 2021; UK Health Security Agency, 2023). Moderna mRNA-1273 was added to the programme in 2021 (UK Health Security Agency, 2023). By October 2021, forty million people, or ninety per cent of the population over eighteen years of age in the England, had received at least one vaccine dose (NHS England, 2021a). A third, booster, dose was offered to healthcare workers and people who were elderly or clinically vulnerable from September 2021, and to everyone over sixteen years of age from November 2021. Since then, spring and autumn vaccination campaigns have invited those who are susceptible to severe disease to have further doses of vaccine against covid-19 (UK Health Security Agency, 2023). The success of this vaccination campaign contributed to the 'Roadmap out of lockdown' (Cabinet Office, 2021), as mandatory restrictions were gradually lifted from early 2021 onwards.

Hale *et al.* (2021) compiled an extensive dataset of governments' policies in response to covid-19 from the start of 2020. They collected data on twenty-four policy measures and reported them with daily resolution. For each

measure, a code of 0 indicates that the measure was not in use. Higher values are given for greater levels of restriction. The number of levels varies between measures. For example, measure H6, facial coverings, has five levels (Hale *et al.*, 2021):

0 No policy;

1 Recommend;

2 Required in some specified public spaces or some situations when social distancing is not possible;

3 Required in all public spaces or all situations when social distancing is not possible; and

4 Required at all times outside the home.

The codes for different measures are combined into policy indices, that give a score of 0 if all components are coded as 0, and a maximum score of 100 if all components are coded at their highest level (*ibid.*). The *Government Response Index* used sixteen measures (those displayed in Figure 1.2) and the *Stringency Index* used a subset of nine of these (C1–C8 and H1). Figure 1.3 shows the Government response index and Stringency index for each constituent nation of the UK between January 2020 and December 2022.

1.3 Causal inference for infectious disease

Understanding causal relationships between co-occurring events has been vital for people making decisions since pre-history (Woodward, 2003). To adopt a behaviour that correlates with favourable outcomes but does not cause them is essentially a form of superstition (Scheibe and Sarbin, 1965). This leads to wasted efforts and possible harms. Non-human animals also appear to base decisions on causal beliefs (Woodward, 2003).

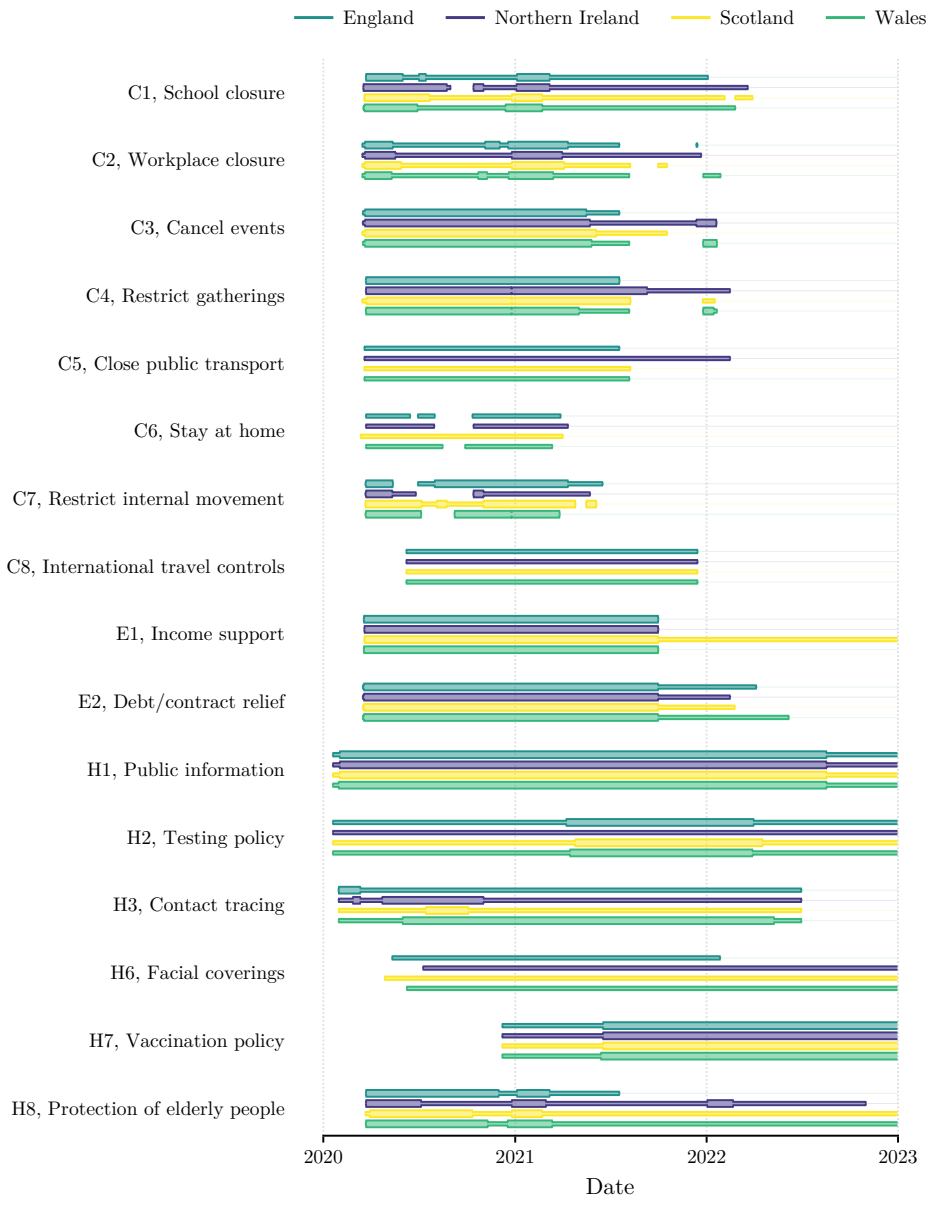

Figure 1.2: Use of interventions to reduce spread of severe acute respiratory syndrome coronavirus 2 in the United Kingdom. Thick bars represent periods when interventions were used at the maximal level, thin bars when interventions were present at a lower level, and no bars when interventions were not in use or no data are available. Dates marked on the horizontal axis represent 1 January of each year. Data from Hale *et al.* (2021).

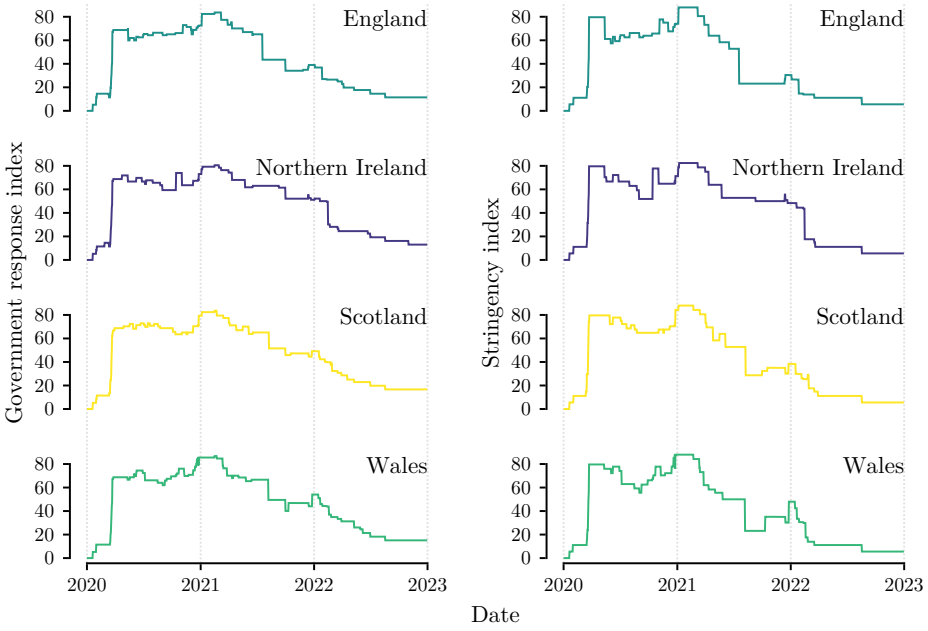

Figure 1.3: Levels of government response to covid-19 in the United Kingdom between 2020 and 2023. Left-hand plots show Government response index, and right-hand plots show Stringency index. Dates marked on the horizontal axis represent 1 January of each year. Data from Hale *et al.* (2021).

Philosophers and statisticians have sometimes struggled to define a *cause* (see, Holland, 1986; Woodward, 2003). In a legal context, the question of whether an act caused an outcome (such as an injury) depends on an assessment of whether the outcome would still have occurred without the act, all else being equal (Steel, 2012). This is similar to Woodward's (2003) 'manipulationist' approach to causation, or Schulz's (2011) title: 'If you'd wiggled A, then B would've changed.' Rubin (2005) describes this in terms of 'potential outcomes'.

Potential outcomes

Consider an individual deciding between two treatment options, $A = 0$ *vs.* $A = 1$, intended to affect an outcome, Y. At the time of making such decisions we might consider possible outcomes in terms of probabilities, such

as a 10 per cent chance of cardiovascular disease in the next five years *vs.* a 15 per cent chance. But retrospectively, having taken, or not taken, the treatment for five years, we would either observe a cardiovascular event or we would not. In hindsight, the outcome is deterministic. Our consideration of what might have been under different circumstances might still be probabilistic but if the individual had in fact made a different choice five years earlier, there would still have been exactly one observed outcome. Potential outcomes are outcomes that could have been observed. They are therefore events or non-events, not probabilities.

Let $Y_j^{A=0}$ represent the outcome individual j would experience if they chose treatment $A_j = 0$, and $Y_j^{A=1}$ their outcome if they chose $A_j = 1$. I refer to $Y_j^{A=0}$ as $Y_j^{A=1}$ 'potential outcomes' because at the time the decision is being made, they are possibilities in the individual's future. Some authors (*e.g.* Hernán and Robins, 2020) refer to them as 'counterfactual outcomes.' I reserve the word 'counterfactual' for outcomes that do not occur. In particular, I wish to emphasize that exactly one potential outcome will occur,

$$A_j = a_j \;\Rightarrow\; Y_j = Y_j^{A=a_j} = y_j. \tag{1.1}$$

This relationship – that the outcome we observe is the potential outcome corresponding to the treatment option that was followed – is referred to as 'consistency' (Hernán and Taubman, 2008).

Causal estimands and the fundamental problem of causal inference

The causal estimand is the causal effect to be estimated (Rubin, 2005). With data on potential outcomes for a series of n participant,such as in Table 1.3, we could calculate individual-level causal effects (more generally, these are termed 'unit-level effects' to include situations in which proposed interventions are for populations rather than individuals). The causal effect of the

Table 1.3: Data table for causal estimands

Units	Covariates	Potential outcomes: Control	Potential outcomes: Treatment	Unit-level causal effects
1	$\boldsymbol{x_1}$	$Y_1^{A=0}$	$Y_1^{A=1}$	$\delta_1 = Y_1^{A=1} - Y_1^{A=0}$
2	$\boldsymbol{x_2}$	$Y_2^{A=0}$	$Y_2^{A=1}$	$\delta_2 = Y_2^{A=1} - Y_2^{A=0}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
j	$\boldsymbol{x_j}$	$Y_j^{A=0}$	$Y_j^{A=1}$	$\delta_j = Y_j^{A=1} - Y_j^{A=0}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
n	$\boldsymbol{x_n}$	$Y_n^{A=0}$	$Y_n^{A=1}$	$\delta_n = Y_n^{A=1} - Y_n^{A=0}$

Based on Rubin, 2005

intervention for individual j is,

$$\tau_j := Y_j^{A=1} - Y_j^{A=0}. \tag{1.2}$$

We could equivalently calculate causal effects in terms of a ratio of outcomes, $Y_j^{A=1}/Y_j^{A=0}$. The sharp null hypothesis is that the treatment has no effect on any individual's outcome, $\tau_1 = \tau_2 = \cdots = \tau_n = 0$. We could calculate population-level effects, such as the mean population causal effect,

$$\bar{\tau} := \mathbb{E}\left(Y^{A=1} - Y^{A=0}\right) = \sum_{i=1}^{n} \frac{Y_i^{A=1} - Y_i^{A=0}}{n}, \tag{1.3}$$

where $\mathbb{E}(\cdot)$ represents the expected, or mean, value. This can trivially, but usefully (Holland, 1986), be rearranged as,

$$\bar{\tau} = \mathbb{E}\left(Y^{A=1}\right) - \mathbb{E}\left(Y^{A=0}\right) = \sum_{i=1}^{n} \frac{Y_i^{A=1}}{n} - \sum_{i=1}^{n} \frac{Y_i^{A=0}}{n}. \tag{1.4}$$

We can summarize causal effects for any subset of the cohort, such as a specified age group or those with a particular comorbidity, provided we compare control and treatment potential outcomes for the same individuals (Rubin, 2005),

$$\begin{aligned} \bar{\tau}_c &:= \mathbb{E}\left(Y^{A=1} - Y^{A=0} \mid C = c\right) \\ &= \sum_{i \in \mathcal{S}_c} \frac{Y_i^{A=1} - Y_i^{A=0}}{|\mathcal{S}_c|}, \quad \mathcal{S}_c = \{i \mid C_i = c\}, \end{aligned} \tag{1.5}$$

Table 1.4: A more realistic data table

		Potential outcomes		
Units	Covariates	Control	Treatment	Unit-level causal effects
1	$\boldsymbol{x_1}$	$Y_1^{A=0}$	missing	unknown
2	$\boldsymbol{x_2}$	missing	$Y_2^{A=1}$	unknown
3	$\boldsymbol{x_3}$	missing	$Y_3^{A=1}$	unknown
4	$\boldsymbol{x_4}$	$Y_4^{A=0}$	missing	unknown
⋮	⋮	⋮	⋮	⋮

where C represents a covariate of interest, and $|\mathcal{S}_c|$ is the cardinality (number of members) of $\mathcal{S}_c$.

Since we can only ever observe one potential outcome for each individual, a more realistic dataset might look like Table 1.4. With this dataset, we cannot observe unit-level causal effects, or test the sharp null hypothesis, or calculate population causal effects as described in equations 1.3–1.5. Holland (1986) describes this as the 'fundamental problem of causal inference'. To estimate causal effects we must make inferences about missing potential outcomes.

Causal inference in randomized controlled trials

In an RCT, individuals, or clusters of individuals, are assigned to treatment and control groups at random. Assuming all participants adhere to their allocated treatment, observed outcomes for individuals who received the treatment, $\{y_i \mid A_i = 1\}$, are a random sample of the whole cohort's set of treated potential outcomes. Greenland and Robins (1986) describe this in terms of exchangeability between the groups, $\mathbb{E}\big(Y^{A=a} \mid A = a\big) = \mathbb{E}\big(Y^{A=a} \mid A \neq a\big)$. The mean of the observed outcomes is an unbiased estimator of the mean of the whole cohort's potential outcomes, $\mathbb{E}(Y \mid A = a) = \mathbb{E}\big(Y^{A=a}\big)$. Substituting these sample means into the equation for the mean

population causal estimand (equation 1.4),

$$\begin{aligned}\hat{\tau}_{\text{RCT}} &= \mathbb{E}\Big(Y^{A=1} \mid A=1\Big) - \mathbb{E}\Big(Y^{A=0} \mid A=0\Big) \\ &= \sum_{i \in \mathcal{S}_1} \frac{y_i}{|\mathcal{S}_1|} - \sum_{i \in \mathcal{S}_0} \frac{y_i}{|\mathcal{S}_0|}, \quad \mathcal{S}_a = \{i \mid A_i = a\}\,. \end{aligned} \tag{1.6}$$

Note that this equation differs from the analysis of potential outcomes in a subset of the population (equation 1.5), as the values being compared describe different individuals. Jerzy Neyman proved that this is an unbiased estimator of the causal effect (discussed by Rubin, 2005). Although the estimate is unbiased, it could differ substantially from the true causal effect, depending on the play of chance (Rubin, 1974).

Causal inference in observational studies

When using observational data, we expect those who receive an intervention to differ systematically from those who do not receive it. For instance, in a study of therapeutic treatments, those who are treated might have more severe disease than those who are not treated, or they might have a different set of comorbidities. Outcomes for those who do not receive treatment are not a random sample of the cohort's untreated potential outcomes. Even if the sharp null hypothesis is true we could, wrongly, infer evidence for a causal effect, $\mathbb{E}\left(Y \mid A=1\right) - \mathbb{E}\left(Y \mid A=0\right) \neq 0$. An objective to infer causal effects can therefore be rephrased as an objective to find unbiased estimates of the potential outcomes.

Causal diagrams and adjustment for covariates

We can consider differences between exposed and unexposed cohorts in terms of 'confounding variables'. A causal diagram can help to identify confounders. Causal diagrams are graphical models of relationships between variables. Often these are represented as directed acyclic graphs (Greenland

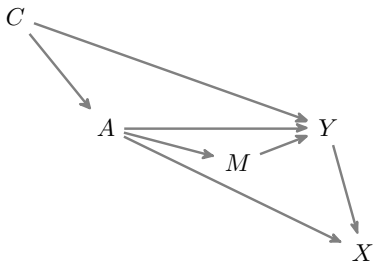

Figure 1.4: Example causal diagram. For the $A \rightarrow Y$ causal effect, C is a confounder, M is a mediator and X is a collider.

et al., 1999; Pearl, 1995), as in Figure 1.4. In these graphs, all connections between nodes are directed (arrows) and no directed paths (following the direction of the arrows) pass through the same node more than once. When two nodes are connected by an arrow, such as $A \rightarrow M$, we can describe them as a parent (in this case, A) and child (M). More generally, any nodes on a directed path leading into a node are described as ancestors, and any on a path away are described as descendants (*e.g.* node M has ancestors C and A, and descendants Y and X; Greenland *et al.*, 1999). Note that the absence of an arrow is more informative than its presence as an arrow does not rule out a relationship with magnitude zero (*i.e.* no effect) but absence of an arrow always rules out any causal effect (Pearl, 1995). In Figure 1.4, if Y has no effect on X, the magnitude of the effect represented by the arrow $Y \rightarrow X$ is 0, but this is still consistent with the assumptions encoded in the causal diagram.

A confounder is a common cause of both the putative cause and the outcome. In a causal diagram, this is represented as a variable that is an ancestor of both the exposure and the outcome (Greenland *et al.*, 1999). A descendant of both the putative cause and the outcome variables is termed a collider. They are commonly identified by two arrows meeting at a node,

such as X in the path, $A \rightarrow X \leftarrow Y$. Any descendant of X would also be a collider. A path between two nodes is considered 'open' (regardless of the directions of the arrows) unless there is a collider on it. Variables represented on an open path are expected to be correlated. Controlling for a variable on an open path, for example by stratification, closes the path and removes the correlation. Controlling for a collider opens a path. If we assume the causal diagram is correct, it provides the necessary information to determine the minimum set of variables that must be controlled for an unbiased analysis of causal effects. We can also test the implications of the causal graph by assessing whether variables are conditionally independent when not connected by any open paths (Pearl *et al.*, 2016). Note that the causal diagram does not inform the shape of any relationship. For instance, knowing that C causes A does not imply a linear relationship between them.

Assume that Figure 1.4 is an accurate causal representation of our research question, and we are interested in the total causal effect of A on Y. The diagram shows three open paths between A and Y,

1. $A \rightarrow Y$, the direct causal effect;
2. $A \rightarrow M \rightarrow Y$, an indirect effect, where M is a mediator; and
3. $A \leftarrow C \rightarrow Y$, which Greenland *et al.* (1999) term a 'back-door path', where C is a confounder.

Assuming C is a binary variable and the only confounder then if we stratify the data by C, any remaining variation in the outcome must be due to chance, so

$$\mathbb{E}\left(Y^{A=a} \mid A = a, C = c\right) = \mathbb{E}\left(Y^{A=a} \mid C = c\right). \tag{1.7}$$

This condition is variously described as 'exchangeability' (Greenland and Robins, 1986), 'ignorable treatment assignment' (Rubin, 1978), or simply 'no unmeasured confounding' (VanderWeele, 2008). If this condition holds then

we can analyse observational data in a similar way to the analysis of an RCT. We could report a result for a subgroup of interest only, such as,

$$\hat{\tau}_{\text{obs},c} = \sum_{i \in \mathcal{S}_{1c}} \frac{y_i}{|\mathcal{S}_{1c}|} - \sum_{i \in \mathcal{S}_{0c}} \frac{y_i}{|\mathcal{S}_{0c}|}, \quad \mathcal{S}_{ac} = \{i \mid A_i = a,\, C_i = c\}\,, \tag{1.8}$$

or we could produce a weighted mean across all values of C,

$$\hat{\tau}_{\text{obs}} = \sum_{c \in \mathcal{C}} \frac{|\mathcal{S}_c|\, \hat{\tau}_{\text{obs},c}}{n}. \tag{1.9}$$

Note that if there is any value of c for which there is not at least one treated individual and one untreated individual, this calculation includes a term $0/0$, and cannot be calculated. This leads to the requirement of 'positivity', that all treatments must be possible under all combinations of confounding covariates (Hernán and Robins, 2020).

We could also use a statistical model to calculate causal effects. Still assuming that C is binary, we could use linear regression to solve,

$$\mathbb{E}\left(Y \mid A = a,\, C = c\right) = \phi_0 + \phi_1\, a + \phi_2\, c. \tag{1.10}$$

Given the assumption of consistency, so $y_j = Y_j^{A=a_j}$, and equality 1.7, we find,

$$\mathbb{E}\left(Y^{A=a} \mid C = c\right) = \hat{\phi}_0 + \hat{\phi}_1\, a + \hat{\phi}_2\, c \;\Rightarrow\; \hat{\tau}_{\text{obs}} = \hat{\phi}_1. \tag{1.11}$$

We can take the same approach with a continuous variable if we know how it is related to the outcome. If it has a linear relationship then the same regression equation (1.11) could be used. This can be expanded to account for multiple confounding variables. Further complexity is added if there in an interaction, for example if the intervention's effectiveness is dependent on values of other covariates (see, VanderWeele, 2015).

A similar approach is possible if we know the relationship between C and the intervention, A. For a binary intervention, if we knew that the log

odds of $A = 1$ were linearly related to C, we could use logistic regression to solve,

$$\text{logit}\,\Pr(A = 1 \mid C = c) = \theta_0 + \theta_1\, c. \tag{1.12}$$

Then let,

$$w_j := \frac{1}{\Pr(A = a_j \mid C = c_j)} = \begin{cases} 1 + \exp\big(\hat{\theta}_0 + \hat{\theta}_1\, c_j\big), & a_j = 0, \\ 1 + \exp\big(-\hat{\theta}_0 - \hat{\theta}_1\, c_j\big), & a_j = 1, \end{cases} \tag{1.13}$$

be 'inverse-probability of treatment weights'. Weighting each observation by w_j generates a pseudopopulation in which A and C are independent (Robins *et al.*, 2000). We can then calculate an unconfounded estimate of the causal effect,

$$\hat{\tau}_{\text{iptw}} = \frac{\sum_{i \in \mathcal{S}_1} w_i\, y_i}{\sum_{i \in \mathcal{S}_1} w_i} - \frac{\sum_{i \in \mathcal{S}_0} w_i\, y_i}{\sum_{i \in \mathcal{S}_0} w_i}. \tag{1.14}$$

These methods can each be expanded for greater numbers of measured confounders, provided that the positivity assumption is met, *i.e.* $\Pr(A = a \mid \boldsymbol{C} = \boldsymbol{c}) \neq 0$ for all a and $\boldsymbol{c}$. Doubly-robust methods that model the relationship of confounding variables with both the intervention and the outcome are accurate if either of the modelled relationships is correct (see Hernán and Robins, 2020 for more detailed description, including calculation of confidence intervals).

The feature shared by all of these methods is that confounding variables must be identified and measured. It is impossible to prove that all relevant confounders have been identified, and the no-unmeasured-confounders assumption is often violated (VanderWeele, 2008). If a known confounding variable has not been measured, this strongly suggests that these methods will not be appropriate.

Challenges for infectious disease: The stable-unit–treatment-value assumption

Comparison of potential outcomes, $Y_j^{A=0}$ and $Y_j^{A=1}$, relies on there being only one value for each unit–treatment combination. If the value of $Y_j^{A=1}$ depends on other variables, we cannot assert that the magnitude of any difference between $Y_j^{A=0}$ and $Y_j^{A=1}$ is the causal effect of the treatment. Rubin (1980, 1990) named this requirement the 'stable unit–treatment value assumption', or simply the 'stability assumption'. The requirement of 'consistency' (see page 18) necessarily follows from the stability assumption. Given a requirement that there is only one value for the potential outcome $Y_j^{A=a}$, then this must be the value observed when $A_j = a$. Note that the requirement of only one potential outcome per unit per intervention does not require that there is only one intervention option. There can be many possible treatments, each with a corresponding potential outcome. For example, the HIV-CAUSAL Collaboration (2011) estimated a series of causal effects of using different thresholds of CD4-positive T cell–counts to guide decisions to start antiretroviral therapy for human immunodeficiency virus infection. This assumption also does not require that multiple units will have the same outcome.

For the stability assumption to hold, two subsidiary assumptions must be true:

(1) There must be only one form of the intervention

Hernán and Taubman (2008) describe this as a 'well-defined intervention.' How well it needs to be defined depends on the research question. If our only objective is to discard the sharp null hypothesis, all that is needed is to demonstrate that some difference in potential outcomes exists, $Y_j^{A\neq 0} \neq Y_j^{A=0}$ (Rubin, 1986). To estimate the magnitude of interventions' effects, we need

more precise understanding of the intervention.

'Weight loss' would not be a well-defined intervention (Hernán and Taubman, 2008). We do not expect an individual who lost weight through increased physical activity to have the same outcome as if they had lost weight through cachexia associated with severe disease. 'Running for 60 minutes, three times per week', could be a well-defined intervention. It is always possible to add to the definition – running at what speed? on what surface? at what elevation? – but what is necessary is that it is defined sufficiently well that we can expect the intervention to be associated with only one outcome per individual (Hernán, 2016).

There may be several causal questions for a single intervention. Are we interested in the effect of an individual taking a treatment or the effect of them being prescribed it? The effect of being allocated to a treatment in an RCT may differ from the effect of it being prescribed once trial results are known (Hernán and Hernández-Díaz, 2012). For policies at a population-level, there could be different outcomes for personal use of an intervention *vs.* official recommendations to use it *vs.* legal mandates. Analyses must be clear which potential outcomes they are comparing.

The 'well-defined intervention' must also incorporate an understanding of what else changes, similar to the legal definition's requirement that 'all else is equal'. The treatment may affect other variables. Cox (1986) gives an example of an agricultural treatment that was so successful that more birds came and consumed the crop. Should the causal effect be that the treatment increased the yield or reduced it? An epidemiological example is provided in Saad-Roy and Traulsen's (2023) model of use of infection-prevention measures which suggested that making one intervention mandatory could lead to reductions in use of other measures, resulting in a constant overall risk of infection. If the treatment causes a change in other

variables then these are mediators of the treatment's effect. How these are handled depends on whether we are interested in total effects of the intervention, or in separating direct effects from indirect effects (VanderWeele, 2015).

(2) There must be no interference from treatment of other units

If the effect of giving a treatment to individual j depends on whether the treatment was also given to other individuals, then there are multiple possible values of the potential outcome. This is a particular issue for studies of infection-prevention measures. Not only do we expect one individual's treatment to affect other people's outcomes, that may be the very effect we are attempting to estimate.

One approach is to consider interventions at the level of clusters of individuals. Cluster-RCTs randomly allocate collections (clusters) of people to treatments, rather than individuals. To analyse these studies, the cluster is the *unit* to consider for the stable unit–treatment value, so that the stability assumption is met so long as there is no interference between clusters, regardless of interference within clusters.

Consider a household with two individuals, $k = 1, 2$. Without a vaccine, individual 1 gets infected then infects individual 2. If either individual were vaccinated, individual 2 would not be infected. Let $A_k = 1$ if individual k is vaccinated and $A_k = 0$ otherwise. What is the causal effect of vaccinating individual 2?

$$\tau_2 = Y_2^{A=1} - Y_2^{A=0} = \begin{cases} 0 - 1 = -1, & \text{if } A_1 = 0, \\ 0 - 0 = 0, & \text{if } A_1 = 1. \end{cases}.$$

Our estimate of the causal effect for individual 2 changes with the treatment of individual 1.

VanderWeele and Tchetgen Tchetgen (2011) considered this scenario, allowing interference within households of two people but no interference between households. In each of j households, the kth individual, $k \in \{1, 2\}$, had four potential outcomes, $Y_{jk}^{A_1=a_1, A_2=a_2}$, $a_1, a_2 \in \{0, 1\}$, depending on whether the first and second members of the household were vaccinated. Halloran and Struchiner (1995) considered interference among three individuals, giving eight potential outcomes per person. As the number of individuals whose treatment may cause interference increases, n, this approach rapidly becomes impractical as the number of potential outcomes increases exponentially, 2^n.

Another approach described by Halloran and Struchiner (*ibid.*) is to consider transmission as a two-part intervention:

1. Is this individual exposed?
2. If exposed, is this individual infected?

An individual's potential outcomes conditional on having been exposed are then independent of other individuals' treatments. For unexposed individuals, neither of the post-exposure potential outcomes is observed.

Halloran (2019) describes multiple measures of vaccine efficacy, including effects for the recipient of the vaccine and indirect effects for people who are not vaccinated.[5] Vaccinated individuals may also benefit from indirect effects if the vaccine does not offer complete protection. Hudgens and Halloran (2008) proposed a two-stage RCT. In that trial, clusters are allocated to having a high or low proportion vaccinated. Then within each cluster, individuals are allocated to receive the vaccine or not at a ratio that results in the assigned proportion. This allows estimates of the direct effect of vaccin-

[5] 'Indirect effect' here refers to effects for other community members and the 'direct effect' is the effect of receiving the vaccine. This is distinct from the definitions of direct and indirect used in mediation analyses, as discussed on page 23.

ation on the recipients of the vaccine, and community-level indirect effects on transmission.

Difference-in-difference methods

In Chapter 2 I combine the difference-in-differences method with a mathematical model of infection dynamics to estimate the effects of interventions during the covid-19 pandemic in the UK. Difference-in-differences methods attempt to estimate causal effects of exposures applied to groups or populations. The causal effect being estimated is the effect of the intervention on the group. This has a benefit for infectious-disease transmission as the *unit* is the group so interference of an individual's treatment on other individuals does not affect the causal question. The stability assumption requires only that outcomes in each group are independent of other groups' treatments.

Difference-in-differences analyses rely on an assumption that, without treatment, groups' potential outcomes would have moved in parallel. The causal effect of the intervention is therefore the difference between the observed outcomes for the treated groups and a counterfactual that moved in parallel with the untreated groups' outcomes. I discuss this method in more detail at the start of Chapter 2.

1.4 Mechanistic modelling of infectious disease

Models are simplified representations of the world. The most frequently-quoted statement about models is: 'All models are wrong' (Box, 1976). Taking a globe as a somewhat literal model of the world (inspired by McElreath, 2020, ch. 2), we find many features of the real world are missing. We could describe the globe as *wrong*. But we can use the globe to infer things about the real world (McElreath, 2020; Thompson, 2022). Indeed, inference from

the model may be much easier than inference from the real world. With a globe, we can very quickly determine whether London is further north or south than Ottawa. Determining this from the real world would take a lot more time and effort. It will always be possible to make a model more realistic by making it more complicated. But if we tried to *correct* our globe, until we had eventually created a full-scale replica of the world, it would provide no additional insight beyond direct observation of the real world. Box recognised that it was not necessary or desirable to change every aspect of a model that differed from reality when he wrote: 'The scientist must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad.' (Box, 1976)

Thompson (2022) uses the metaphor of 'Model Land' to describe the difference between models and the real world. In Model Land, we can make any assumptions we desire. These may be intentionally contrary to knowledge about the real world. Nothing is inherently wrong in Model Land. The challenge lies in determining whether results in Model Land can be used to draw conclusions about the real world. She concludes that the accurate interpretation of a model takes the form: 'X will happen, conditional on this model being adequate for the purpose of predicting X' (*ibid.*).

Phenomenological and mechanistic models

One of the earliest mathematical representations of infectious disease outbreaks was made by Farr (1840). Studying epidemic smallpox between 1837 and 1839, he found that the rate at which numbers of deaths from smallpox increased, was decreasing over time (*i.e.* the second derivative was negative). In 1866 he applied this finding to an epizootic of rinderpest and predicted the timing of its peak to within two weeks of its subsequently observed true

date.[6] Farr's model, often referred to as 'Farr's law', is,

$$\frac{x(t_1)/x(t_2)}{x(t_3)/x(t_4)} = \frac{x(t_2)/x(t_3)}{x(t_4)/x(t_5)} = K < 1, \tag{1.15}$$

where $x(t)$ is disease-specific mortality at time t; t_1, t_2, ..., are equally-spaced time points; and K is constant (Fine, 1979). An alternative form is, $x(t) = \exp\left(-at^2 + bt + c\right)$, $a > 0$, which describes the epidemic as following a *normal* curve (*ibid.*).

Farr's model can be described as 'phenomenological' as it describes a relationship between observations but does not incorporate biologically-meaningful understanding of why this relationship might be seen (Santillana *et al.*, 2018). It is worth recalling that Farr was trying to understand epidemics without the germ theory of infectious disease.[7] Santillana *et al.* (*ibid.*) argue that phenomenological models may be sufficient for generating forecasts or for recognising when underlying disease dynamics have changed. Farr's method remained in use into the twenty-first century. It has been used to predict trends in acquired immunodeficiency syndrome (Bregman and Langmuir, 1990) and covid-19 diagnoses (Pacheco-Barrios *et al.*, 2020; Xu *et al.*, 2020), and has been applied to the non-infectious epidemic of deaths from drug overdoses in the United States (Darakjy *et al.*, 2014). Its application to acquired immunodeficiency syndrome diagnoses was particularly unsuccessful (Nishiura, 2007). Changes in characteristics of the disease, such as the incubation period (time from exposure to disease), contributed to substantial errors in the prediction (Brookmeyer, 1996; Gail and Brookmeyer, 1990). Even in the original report of their predictions,

[6] Farr's prediction in a letter to *Daily News*, 17 February 1866, and follow-up with information about how the epizootic progressed in the *Journal of Social Science*, 20 March 1866, are reprinted in full in Brownlee, 1915.

[7] Farr (1840) wrote a lengthy footnote about the hypothesis that epidemics could be caused by transmissible animalcules, before concluding that this hypothesis 'does not satisfactorily explain the cause of epidemics.'

Bregman and Langmuir recognised that changes in diagnostic criteria would affect numbers of diagnoses, but they were unable to adjust their projection to account for this. More generally, given that phenomenological models do not explain why disease numbers progress as they do, they cannot predict or explain situations when the pattern changes. They have no parameters to manipulate so cannot be used to answer 'what if?' questions. In the language of causal inference, they cannot encompass counterfactual scenarios.

In contrast, mechanistic models attempt to incorporate a description of processes that cause disease transmission. Limitless variation in such models is possible. I limit my introduction to types of model that I use in this thesis.

Characteristics of models used in this thesis

Before discussing details of specific models, I summarize characteristics shared by all models used in this thesis:

(1) All are population-level models

None of the models in this thesis attempt to trace individuals' outcomes. Instead they simplify populations into groups that share some characteristic – such as those who are susceptible or infectious or immune – and describe how numbers in these groups are expected to change over time.

(2) All assume a constant population size

This may often be an unrealistic assumption but it greatly simplifies models. It is achieved by forcing birth and mortality rates to equal each other. This has two important consequences:

1. there can be no mortality due to the infection (unless infection also caused an increase in birth rates); and

2. attempts to use the model to explore changes in birth rates simultaneously change modelled life expectancy, not through any biological process but as an inevitable consequence of the model maintaining the constant population.

(3) Transmission is assumed to be prevalence-dependent

Prevalence-dependent transmission (which is also referred to as frequency-dependent or mass-action transmission)[8] assumes that infectiousness is not dependent on population size. Let $i(t)$ be the number of infectious individuals at time t, and n be the (constant) population size, so prevalence is $i(t)/n$. We can describe infectiousness in terms of a per-capita force of infection,

$$\lambda(t) \propto \frac{i(t)}{n} \;\Rightarrow\; \lambda(t) = \beta \frac{i(t)}{n}, \tag{1.16}$$

where β is a transmission coefficient. If transmissibility changes over time, either due to seasonal changes or in response to an intervention, we can substitute a time-varying β,

$$\lambda(t) = \beta(t) \frac{i(t)}{n}. \tag{1.17}$$

For the remainder of this section, I describe types of model used in this thesis.

Deterministic compartmental models in continuous time (ordinary differential equation models)

Compartmental models are widely used to describe physical and biological processes (Godfrey, 1983). For chemical reactions, compartments can repres-

[8] Names for different forms of transmission have not been used consistently (Keeling and Rohani, 2008). Begon *et al.* (2002) inadvertently coined my preferred term when they wrote of their 'wish not to introduce new terms – such as the logical "prevalence-dependent transmission" [...] – if this can be avoided.' I support the principle of trying to avoid ballooning numbers of terms, but the meaning of 'prevalence-dependent' seems so immediately clear that I feel it cannot help but aid understanding.

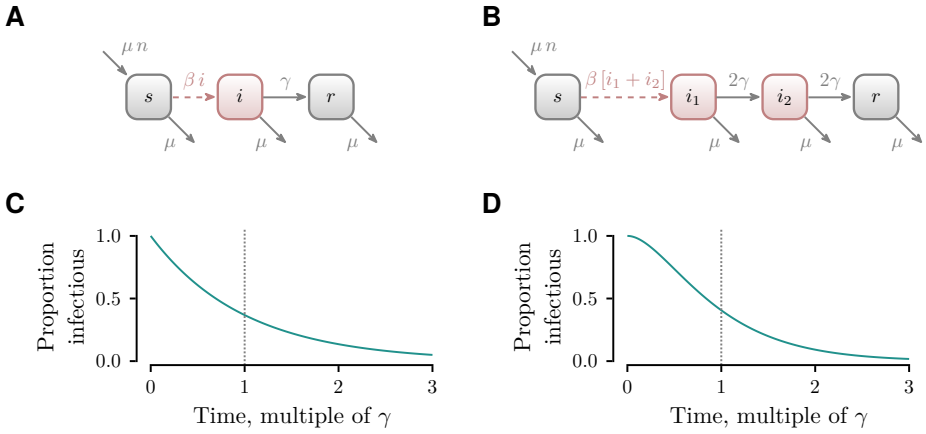

Figure 1.5: Example flow diagrams for compartmental models with susceptible, s, infectious, i, and resistant, r, categories. Model A has a single infectious compartment, giving an exponential duration of infectiousness (C); B has two sequential infectious compartments, giving an Erlang-distributed duration of infectiousness (D). Compartments that contribute to the force of infection are shaded in red; events that occur in proportional to the force of infection are represented by dashed red lines. Parameters are described in the text.

ent concentrations, in pharmacokinetics they can represent amounts of drug in different tissues, and in ecology they can represent numbers of predators and prey. In epidemic models, compartments often represent numbers of individuals who share a state with respect to history of infection by a pathogen, such as being susceptible, infectious or immune.

Figure 1.5A represents a susceptible–infectious–resistant (SIR) compartmental model. This is based on models initially described by Kermack and McKendrick (1927, 1932, 1933). Let $s(t)$ represent the number of individuals who are susceptible to infection at time t, where $s(t) \geq 0$ for all $t \geq t_0$. Equivalently, let $i(t)$ represent the number who are infectious, and $r(t)$ the number who are resistant (immune) to infection. Each compartment is assumed to be homogeneous so for any given number of individuals who are infectious, disease dynamics do not depend on which individuals these are.

Rates of change of each compartment size are described by ordinary

differential equations (ODEs). For example, let,

$$\left.\begin{aligned}\frac{\mathrm{d}s(t)}{\mathrm{d}t} &= \mu\,[n - s(t)] - \beta\,\frac{s(t)\,i(t)}{n},\\ \frac{\mathrm{d}i(t)}{\mathrm{d}t} &= \beta\,\frac{s(t)\,i(t)}{n} - [\gamma + \mu]\,i(t),\\ \frac{\mathrm{d}r(t)}{\mathrm{d}t} &= \gamma\,i(t) - \mu\,r(t),\end{aligned}\right\} \tag{1.18}$$

where μ is both the birth and mortality rate, γ is the recovery rate, and $t \geq t_0$. In contrast to the force of infection, which varies with the number of infectious individuals, the recovery and birth rates are unaffected by compartment sizes. I refer to these rates that are independent of other compartments as 'transitions'.

We can demonstrate that the model has a constant population size by calculating the rate of population change,

$$\frac{\mathrm{d}n}{\mathrm{d}t} = \frac{\mathrm{d}s(t)}{\mathrm{d}t} + \frac{\mathrm{d}i(t)}{\mathrm{d}t} + \frac{\mathrm{d}r(t)}{\mathrm{d}t} = 0.$$

We can simplify the model by removing one differential equation, since we know, $r(t) = n - [s(t) + i(t)]$. This makes the problem two-dimensional, rather than three-dimensional, simplifying analysis of the model (Strogatz, 2018).

The model has a Markov property. This means that, given conditions at time t, progression of the model is independent of conditions at any previous time, $\tau < t$. Transitions out of any compartment follow an exponential distribution (equivalent to radioactive decay). For instance, the mean duration of infectiousness is $1/[\gamma + \mu]$ (approximately $1/\gamma$ when $\gamma \gg \mu$), with an exponential distribution (Figure 1.5C). 'Memory' can be added to the model by dividing durations in one state between a series of compartments. For example, we could represent infectiousness with two serial infectious

compartments (the SIIR, or SI^{2}R model; Figure 1.5*B*),

$$\left.\begin{aligned} \frac{\mathrm{d}s(t)}{\mathrm{d}t} &= \mu\,[n - s(t)] - \beta\, s(t)\,\frac{i_1(t) + i_2(t)}{n}, \\ \frac{\mathrm{d}i_1(t)}{\mathrm{d}t} &= \beta\, s(t)\,\frac{i_1(t) + i_2(t)}{n} - [2\gamma + \mu]\, i_1(t), \\ \frac{\mathrm{d}i_2(t)}{\mathrm{d}t} &= 2\gamma\, i_1(t) - [2\gamma + \mu]\, i_2(t), \\ r(t) &= n - [s(t) + i_1(t) + i_2(t)]\,. \end{aligned}\right\} \tag{1.19}$$

Assuming μ is much smaller than γ, the mean duration of infectiousness in this model is $1/\gamma$, as in the SIR model. In the SIIR model, this duration follows an Erlang distribution (a Gamma distribution with an integer shape parameter; Figure 1.5*D*). This is more biologically realistic for an infectious period than the exponential distribution (Lloyd, 2001; Wearing *et al.*, 2005).

Further compartments can be added as needed. A common variation is the susceptible–exposed–infectious–resistant (SEIR) model, in which a latent (non-infectious) period is modelled in the *exposed* compartment,

$$\left.\begin{aligned} \frac{\mathrm{d}s(t)}{\mathrm{d}t} &= \mu\,[n - s(t)] - \beta\, s(t)\,\frac{i(t)}{n}, \\ \frac{\mathrm{d}e(t)}{\mathrm{d}t} &= \beta\, s(t)\,\frac{i(t)}{n} - [\varsigma + \mu]\, e(t), \\ \frac{\mathrm{d}i(t)}{\mathrm{d}t} &= \varsigma\, e(t) - [\gamma + \mu]\, i(t), \\ r(t) &= n - [s(t) + e(t) + i(t)]\,, \end{aligned}\right\} \tag{1.20}$$

where $1/\varsigma$ is the mean duration of the latent period. Latent periods can also be given an Erlang distribution by use of serial exposed compartments. A general form of these models with m exposed compartments and k infectious compartments can be described with the acronym SE^{m}I^{k}R.

Infectiousness of respiratory viruses varies seasonally (Moriyama *et al.*, 2020; Shek and Lee, 2003; Tamerius *et al.*, 2011). One way to represent this is to replace the transmission coefficient, β, with a time-varying version, such as,

$$\beta(t) = \beta_0[1 + \beta_1\,\cos(2\pi t - \varphi)], \tag{1.21}$$

where β_0 is the mean transmission coefficient, $\beta_1 \in [0, 1]$ is the proportional amplitude of seasonal forcing, $\varphi \in [0, 2\pi)$ is a phase offset, and t is time in years. Strogatz (2018) suggests that analysis is simplified if each differential equation is time-invariant, which can be achieved by adding a compartment, $x(t)$,

$$\left.\begin{aligned} \beta(t) &= \beta_0[1 + \beta_1 \cos x(t)], \\ \frac{\mathrm{d}x(t)}{\mathrm{d}t} &= 2\pi, \quad x(0) = -\varphi. \end{aligned}\right\} \tag{1.22}$$

A more computationally-efficient version is achieved by substituting,

$$\frac{\mathrm{d}(\cos(2\pi\, t))}{\mathrm{d}t} = -2\pi \sin(2\pi\, t),$$

giving a model with two additional equations and additional parameters $x_1(t)$ and $x_2(t)$,[9]

$$\left.\begin{aligned} \beta(t) &= \beta_0[1 + \beta_1\, x_1(t)], & \\ \frac{\mathrm{d}x_1(t)}{\mathrm{d}t} &= -2\pi\, x_2(t), & x_1(0) = \cos\varphi, \\ \frac{\mathrm{d}x_2(t)}{\mathrm{d}t} &= 2\pi\, x_1(t), & x_2(0) = -\sin\varphi. \end{aligned}\right\} \tag{1.23}$$

Note that this approach to seasonal variation is essentially phenomenological, as we have not ascribed any meaning to the parameters β_1 or φ. Alternative approaches to seasonal forcing attempt to incorporate reasons for changes in transmission. Schenzle's (1984) model of measles transmission assumed that seasonal changes were due to changes in contact rates among children when schools were open *vs.* closed. In his model, age-specific transmission rates among children switch between two levels depending on whether or not schools are open (switching to a lower level during holidays and on each Sunday). Shaman *et al.* (2010) included specific humidity in a transmission model to incorporate measurable environmental effects on transmissibility.

[9] Described in documentation for the `DifferentialEquations.jl` package; Rackauckas and Nie, 2017.

Generation interval

The generation interval is the time that passes from when one individual is infected, to the time they infect a secondary case (labelled transmission interval by Fine, 2003). The generation interval is not constant throughout an epidemic (Svensson, 2007). Kenah *et al.* (2008) described generation interval contraction in terms of a 'race', such that when there are many infectious individuals, an initially susceptible individual may have many contacts with infectious individuals, only the first of which leads to infection.

Svensson (2007) investigated and formally defined generation intervals. Assuming constant infectiousness throughout the infectious period, for an infectious individual who causes at least one secondary infection, the mean time to a secondary infection is,

$$\mathbb{E}(T_p) = \mathbb{E}(X) + \mathbb{E}(Y)/2, \tag{1.24}$$

where T_p stands for 'primary' generation interval, X is the duration of the latent period and Y is the duration of the infectious period. As not all infectious individuals will cause a secondary infection, infectious individuals must be weighted by number of secondary cases, or equivalently we choose a secondary infection event at random and define the generation interval T_s ('secondary' generation interval) as the time since the infector was themselves infected. This weighted mean generation time is given by,

$$\mathbb{E}(T_s) = \mathbb{E}(X) + \frac{\mathbb{E}(Y^2)}{2\,\mathbb{E}(Y)}. \tag{1.25}$$

For a Gamma-distributed infectious period with shape parameter k, this equates to,

$$\mathbb{E}(T_s) = \mathbb{E}(X) + \mathbb{E}(Y)\,\frac{k+1}{2k}. \tag{1.26}$$

For $\mathrm{SE}^m\mathrm{I}^k\mathrm{R}$ models, assuming μ is much smaller than γ, this gives,

$$\mathbb{E}(T_s) = \frac{1}{\varsigma} + \frac{k+1}{2k\gamma}. \tag{1.27}$$

This is not affected by the number of exposed compartments, m. For models with exponential-distributed infectious periods, $k = 1$, the expected generation interval equals the mean latent period plus the mean duration of infectiousness. For models with serial infectious compartments, $k > 1$, the mean generation time is less than the sum of the times in the latent and infectious compartments. Krylova and Earn (2013) found that if SIR models were parameterized so that $1/\gamma$ equals the microparasite's infectious period then the ability of the model to reproduce real-world–like dynamics depends on the number of infectious compartments. However, if $1/\gamma$ is calibrated to the mean generation interval then models with any number of infectious compartments, including those with a single infectious compartment, perform equally well.

Basic and effective reproduction ratios

The basic reproduction ratio (or number), $\mathcal{R}_0$, is the expected number of secondary infections that would arise from each infectious individual in a fully-susceptible population (Dietz, 1993; Heesterbeek, 2002). In general form, it can be described as,

$$\mathcal{R}_0 = \int_0^\infty p(\tau)\, f(\tau)\, \mathrm{d}\tau, \tag{1.28}$$

where $p(\tau)$ is an individual's probability of surviving, and $f(\tau)$ is expected infectiousness, at time τ since infection. For models discussed so far, this is simply the product of the rate of secondary infections and the expected duration of infectiousness. In a fully-susceptible population, $s(t) = n$, so the rate of new infections from each infectious individual is β. Given the duration of infectiousness is $1/[\gamma + \mu]$, $\mathcal{R}_0 = \beta/[\gamma + \mu]$. Using the SIR model, we can show that this intuitive result matches the general form in

equation 1.28, with $p(\tau) = \exp(-\mu\,\tau)$, and $f(\tau) = \beta\,\exp(-\gamma\,\tau)$,

$$\mathcal{R}_0 = \int_0^\infty \beta\,\exp(-[\gamma+\mu]\,\tau)\;\mathrm{d}\tau = \left[-\frac{\beta\,\exp(-[\gamma+\mu]\,\tau)}{\gamma+\mu}\right]_0^\infty = 0 + \frac{\beta}{\gamma+\mu}.$$

Diekmann *et al.* (2010) describe how to calculate the basic reproduction ratio from more complicated compartmental models using next-generation matrices. It can also be estimated from the epidemic growth rate or the final size of an epidemic (Dietz, 1993). For a pathogen causing lifelong immunity and $\mathcal{R}_0 > 1$, it can be estimated as $\mathcal{R}_0 = 1 + L/A$, where L is life expectancy and A is mean age at infection (Anderson and May, 1982).

Microparasites do not have consistent basic reproduction ratios. It can differ for urban *vs.* rural settings, summer *vs.* winter, schools *vs.* healthcare settings *vs.* the wider community (*ibid.*). Assuming SIR-like dynamics, any disturbance to β, γ or μ would change the basic reproduction ratio. Authors such as You *et al.* (2020), differentiate between a basic reproduction ratio, which they define as the expected number of secondary infections without any infection-control measures, and a 'controlled reproduction ratio', which is the equivalent value when infection-prevention measures are in use. I feel this is an unnecessary complication given the variation in $\mathcal{R}_0$ from measures not under human control, such as seasonal changes in transmission. It also requires us to define a control-free situation for the basic reproduction ratio. For instance, does routine hand washing affect the basic or the controlled reproduction ratio? Throughout this thesis, I assume that $\mathcal{R}_0$ can vary between locations and over time, just as the parameters used in its calculation do.[10] However, I do not follow Alexander *et al.* (2006) in defin-

[10]Reluctance to consider time-varying basic reproduction ratios may come from the subscript 0 in its symbol. This does not refer to the reproduction ratio at time 0 or when 0 per cent of the population has been infected. Rather, in its original use it referred to the zeroth order moment (*i.e.* total mass) of the maternity function (Dietz, 1993).

ing the basic reproduction ratio such that effective vaccination programmes (*i.e.* reductions in the number susceptible) reduce it.

The effective reproduction ratio, $\mathcal{R}_e(t)$, is the expected number of secondary infections per infectious individual, given the population level of immunity at time t. For the models I have shown above, this is simply,

$$\mathcal{R}_e(t) = \frac{s(t)}{n}\,\mathcal{R}_0. \tag{1.29}$$

If $\mathcal{R}_e(t) > 1$ then the outbreak is growing, and if $\mathcal{R}_e(t) < 1$ the outbreak is shrinking. This leads to the 'herd-immunity threshold',

$$\mathcal{R}_e(t) < 1 \;\Leftrightarrow\; \frac{s(t)}{n} < \frac{1}{\mathcal{R}_0} \;\Leftrightarrow\; \frac{r(t)}{n} > 1 - \frac{1}{\mathcal{R}_0}.$$

If $\mathcal{R}_0 < 1$ then $\mathcal{R}_e(t)$ is always less than 1.

Equilibrium points

Equilibrium points are conditions under which the model's conditions will remain constant over time. These are found when each differential equation equals 0. The SIR model (equation 1.18) has two equilibrium points. The disease-free equilibrium is, $s_* = n,\ i_* = r_* = 0$. The endemic equilibrium is,

$$s^* = \frac{[\gamma + \mu]\,n}{\beta}, \quad i^* = \frac{\mu\, n\,[\beta - \gamma - \mu]}{\beta\,[\gamma + \mu]}, \quad r^* = \frac{\gamma\, n\,[\beta - \gamma - \mu]}{\beta\,[\gamma + \mu]}.$$

The endemic proportion susceptible, $s^* = n/\mathcal{R}_0$, is a common result in compartmental models. It gives an effective reproduction ratio of 1 so that, on average, each infectious individual gives rise to one replacement infection, maintaining the equilibrium.

The endemic equilibrium cannot be reached if $\mathcal{R}_0 < 1$ but the disease-free equilibrium is mathematically possible for any set of parameters. We can test the stability of each equilibrium point to determine whether populations would remain near these values if they were perturbed slightly away from the equilibrium. Firstly we calculate the Jacobian matrix (a matrix

of partial derivatives with respect to each compartment) at the equilibrium point, and find the eigenvalues of the Jacobian. These eigenvalues may be real or complex, often appearing as complex conjugate pairs, $\Lambda_{1,2} = a \pm \mathrm{i}b$, where i is the imaginary unit, $\mathrm{i}^2 = -1$. Letting $\boldsymbol{x} = (s, i, r)$, and assuming two eigenvalues, $\Lambda_{1,2} = a \pm \mathrm{i}b$, we can express $\boldsymbol{x}(t)$ as a linear combination of eigenvectors, $\boldsymbol{v_1}$ and $\boldsymbol{v_2}$ (Strogatz, 2018),

$$\begin{aligned}\boldsymbol{x}(t) &= c_1 \, \mathrm{e}^{\Lambda_1 t} \, \boldsymbol{v_1} + c_2 \, \mathrm{e}^{\Lambda_2 t} \, \boldsymbol{v_2} \\ &= c_1 \, \mathrm{e}^{at} \, \mathrm{e}^{\mathrm{i}bt} \, \boldsymbol{v_1} + c_2 \, \mathrm{e}^{at} \, \mathrm{e}^{-\mathrm{i}bt} \, \boldsymbol{v_2}.\end{aligned}$$

Without attempting to solve fully, we see that if we substitute in Euler's formula, $\mathrm{e}^{\mathrm{i}\theta} = \cos\theta + \mathrm{i}\sin\theta$, we have terms of the form, $\mathrm{e}^{at}\cos bt$. From this we infer that if,

1. all eigenvalues are real and negative ($a < 0$, $b = 0$), there are no oscillations and with time the term will tend toward 0 (the equilibrium point is a stable node);[11]
2. if all eigenvalues are real and any are positive ($a > 0$, $b = 0$), there are no oscillations and with time the term will tend toward ∞ (the equilibrium point is an unstable node);
3. if all complex eigenvalues have a negative real parts ($a < 0$, $b \neq 0$), there are sinusoidal oscillations that tend toward 0 with time (the equilibrium point is a stable spiral);
4. if any complex eigenvalues have a positive real part ($a > 0$, $b \neq 0$), there are sinusoidal oscillations that tend toward ∞ with time (the equilibrium point is an unstable spiral); and
5. if all eigenvalues are imaginary ($a = 0$, $b \neq 0$), there are sinusoidal oscillations with a fixed amplitude (the equilibrium point is a centre).

[11] Names of equilibria in this list are from Strogatz, 2018.

Numerical solutions

These models are deterministic. A set of parameters and initial conditions are sufficient to infer the size of each compartment for any time $t \geq t_0$. The canonical approach to find numerical solutions is the Euler method (see Butcher, 2016). For an initial-value problem, $\mathrm{d}y(t)/\mathrm{d}t = f(t, y(t)),\ y(t_0) = y_0$, we can estimate subsequent values at small time-intervals h, by assuming that the gradient of the slope between any time t_x and time $t_x + h$ is constant and equal to the gradient at time t_x. This gives a vector of estimated values such as,

$$\begin{bmatrix} y(t_0 + h) \\ y(t_0 + 2h) \\ y(t_0 + 3h) \\ \vdots \end{bmatrix} \approx \begin{bmatrix} \hat{y}_h \\ \hat{y}_{2h} \\ \hat{y}_{3h} \\ \vdots \end{bmatrix} = \begin{bmatrix} y_0 + h\, f(t_0, y_0) \\ \hat{y}_h + h\, f(t_0 + h, \hat{y}_h) \\ \hat{y}_{2h} + h\, f(t_0 + 2h, \hat{y}_{2h}) \\ \vdots \end{bmatrix}.$$

If $\mathrm{d}y(t)/\mathrm{d}t$ is constant, the Euler method gives the exact solution (but this is an equation for a straight line so numerical methods are probably not needed). For time-varying rates, numerical errors are introduced. Smaller values of h give smaller errors but incur greater computing costs. Runge–Kutta methods calculate intermediate values of $f(t, y(t))$ between time t_x and $t_x + h$. Methods with more intermediate steps (higher-order methods) can provide greater accuracy for lower computational costs (Butcher, 1996, 2016). When these methods are implemented, the step size h can be automatically adjusted to give an optimal balance of speed and accuracy in the estimates (Butcher, 2016; Rackauckas and Nie, 2017). All numerical solutions of ODEs in this thesis are calculated using Runge–Kutta methods in Julia (Bezanson *et al.*, 2017) using `DifferentialEquations.jl` (Rackauckas and Nie, 2017).

Deterministic compartmental models in discrete time

Compartmental models in discrete time are similar to ODE models, but compartment values and time-varying parameters are defined only at discrete times $(t_0, t_1, t_2, \dots)$, where $t_{k+1} = t_k + \Delta t$. We can translate an ODE model to a discrete-time version, by integrating across the interval $[t_k, t_k + \Delta t]$. For example, if members leave a compartment x at a rate α, where $\alpha, x_k, x_{k+1} > 0$,

$$\begin{aligned} \frac{\mathrm{d}x}{\mathrm{d}t} &= -\alpha\, x(t), \\ \Rightarrow \int_{x_k}^{x_{k+1}} \frac{-1}{\alpha\, x(t)}\, \mathrm{d}x &= \int_{t_k}^{t_k+\Delta t} \mathrm{d}t \\ = \frac{\ln x_{k+1} - \ln x_k}{-\alpha} &= \Delta t, \\ \therefore\ x_{k+1} &= x_k \exp(-\alpha\, \Delta t). \end{aligned} \tag{1.30}$$

To simplify notation, constant transitions can be represented by a new constant, such as $\mu' = 1 - \exp(-\mu\, \Delta t)$. A discrete-time version of the SIR model (equation 1.18) is therefore,

$$\left.\begin{aligned} s_{t+1} &= \mu' \left[n - s_t\right] + s_t \exp(-\beta\, i_t\, \Delta t / n), \\ i_{t+1} &= s_t \left[1 - \exp(-\beta\, i_t\, \Delta t / n)\right] + i_t \left[1 - \gamma' - \mu' + \gamma'\mu'\right], \\ r_{t+1} &= n - \left[s_{t+1} + i_{t+1}\right], \end{aligned}\right\} \tag{1.31}$$

where the $\gamma'\mu'$ term in the equation for i_{t+1} prevents double-removal of individuals who would have both recovered and died during the interval $[t, t + \Delta t]$. Similar analysis can be performed with discrete-time models as with ODE models. For example, equilibria are found when $s^* = s_k = s_{k+1}$, $i^* = i_k = i_{k+t}$, and so forth (see Keeling and Rohani, 2008, pp. 46–48).

Renewal equations

The renewal equation is conceptually similar to compartmental models but an individual's infectiousness is not required to be constant over time. Equivalently, the generation interval does not have to follow an exponential or

Erlang distribution. Let $g(x)$ be the distribution of the generation time such that,

$$\int_0^\infty g(x)\,\mathrm{d}x = 1, \quad g(0) = 0, \quad g(x) \geq 0 \text{ for all } x > 0.$$

In a fully-susceptible population, the expected rate of causing secondary infections at a time x after the primary infection is $\mathcal{R}_0\, g(x)$. At a time, t, when $s(t)$ may be less than n, the expected rate of secondary infections from an infectious individual who was infected at a time $x < t$, is $\mathcal{R}_e(t)\, g(t-x)$.

Let $i(t)$ be the number of new infections at time t (note this is an incidence and differs from the prevalence represented by $i(t)$ in compartmental models above), then

$$i(t) = \mathcal{R}_e(t) \int_{-\infty}^{t} i(x)\, g(t-x)\,\mathrm{d}x. \tag{1.32}$$

If $g(x)$ is defined with a domain of natural numbers and $\sum_0^\infty g(x) = 1$, then a discrete-time version of the renewal equation is given by,

$$i_t = \mathcal{R}_e(t) \sum_{x=-\infty}^{t-1} i_x\, g(t-x). \tag{1.33}$$

If the generation interval function is an exponential distribution then the renewal equation gives the same result at the ODE SIR model. Champredon *et al.* (2018) calculated formulae for $g(x)$ to make renewal equations equal any form of the $\mathrm{SE}^m\mathrm{I}^k\mathrm{R}$ model.

Identification of parameters

Structural identifiability

In a general form, let $\hat{\boldsymbol{y}} = \mathcal{M}(\boldsymbol{u_0}, \boldsymbol{\Theta}, t)$ represent an epidemic model, where $\boldsymbol{u_0}$ is a vector of initial conditions, $\boldsymbol{\Theta}$ is a vector of parameters, and t is time. If $\mathcal{M}(\boldsymbol{u_0}, \boldsymbol{\Theta}, t)$=$\mathcal{M}(\boldsymbol{u_0}, \boldsymbol{\Theta}', t)$, and $\boldsymbol{\Theta} \neq \boldsymbol{\Theta}'$, then the parameters are not uniquely described by the output $\hat{\boldsymbol{y}}$. Given a vector of observations, $\boldsymbol{y}$, we might reasonably infer parameters $\boldsymbol{\Theta}$ or $\boldsymbol{\Theta}'$, and would not know

which more accurately described the disease process. In this case we say the parameters are structurally non-identifiable. Castro and de Boer (2020) describe an example of a model in which there are two parameters, the product of which gives a mortality rate,

$$\frac{\mathrm{d}x(t)}{\mathrm{d}t} = -a\,b\,x(t).$$

Even if we were able to infer a total mortality rate, μ, the parameter a could take any non-zero value, with $b = \mu/a$.

To test for identifiability, Castro and de Boer (*ibid.*) propose replacing all unknown parameters (and unobserved compartment sizes) with new parameters of the form $\theta' = u_\theta\,\theta$, then assume that $\mathcal{M}(\boldsymbol{u_0}, \boldsymbol{\Theta}, t)$=$\mathcal{M}(\boldsymbol{u_0}, \boldsymbol{\Theta}', t)$. If this requires each u_θ to take a value of 1 then $\theta' = \theta$. If $u_\theta \neq 1$ then $\theta' \neq \theta$, and the parameter is not uniquely identifiable. Taking the mortality example,

$$-a\,b\,x(t) = -u_a\,u_b\,a\,b\,x(t) \;\Rightarrow u_a\,u_b = 1,$$

u_a and u_b can take values other than 1 so a and b are not identifiable.

Consider a very simple ODE epidemic model with no mortality and no immunity, a susceptible–infectious–susceptible (SIS) model,

$$\left.\begin{aligned} \frac{\mathrm{d}i(t)}{\mathrm{d}t} &= \beta\,s(t)\,\frac{i(t)}{n} - \gamma\,i(t), \\ s_t &= n - i(t). \end{aligned}\right\} \tag{1.34}$$

Castro and de Boer's (*ibid.*) method allows equations to be separated into functionally-independent summands. In this equation the term $\beta\,s(t)\,i(t)/n$, containing the product of two compartments, can be separated from $\gamma\,i(t)$, containing only a single compartment. Assuming prevalence of infection is observed and the total population size is known, so we can fix $u_i = u_n = 1$,

$$\beta\,s(t)\,\frac{i(t)}{n} = u_\beta\,u_s\,\beta\,s(t)\,\frac{i(t)}{n} \qquad \Rightarrow \quad u_\beta\,u_s = 1,$$

$$\gamma\, i(t) = u_\gamma\, \gamma\, i(t) \qquad \Rightarrow \quad u_\gamma = 1,$$
$$n - i(t) = \frac{n - i(t)}{u_s} \qquad \Rightarrow \quad u_s = 1,\ \therefore u_\beta = 1.$$

Each u coefficient must equal unity so the model parameters are identifiable.

Grid-search methods

If parameters are identifiable, the next task is to estimate their values. Let $\boldsymbol{y} = (y_1, y_2, \ldots, y_n)$ be a vector of observed prevalence over time. For the SIS model, we could run the model with a range of possible values for the β and γ parameters, and find the values that produce an output most similar to the dataset, for example,

$$\min_{\beta,\gamma} \sum_{t=1}^{n} \left[y_t - \hat{i}(t) \right]^2.$$

Figure 1.6*A* shows how well different parameters fit to simulated data. The parameters used for the simulation were $\beta = 0.48$ and $\gamma = 0.21$, which are correctly identified.

Bayesian methods and Markov chain Monte Carlo

For models with greater numbers of parameters the grid-search method becomes computationally prohibitive. We may also wish to estimate uncertainty around parameter estimates. Bayesian methods and Markov chain Monte Carlo (MCMC) allow both.

Bayesian inference of model parameters combines the prior plausibility of parameter values with the likelihood of the observed data to generate a posterior probability distribution (McElreath, 2020). Estimates of the prior plausibility of parameter values may be formally elicited from subject experts (Johnson *et al.*, 2010) or calculated from published literature (Tan *et al.*, 2008). Alternatively, a prior distribution may be chosen to indicate no

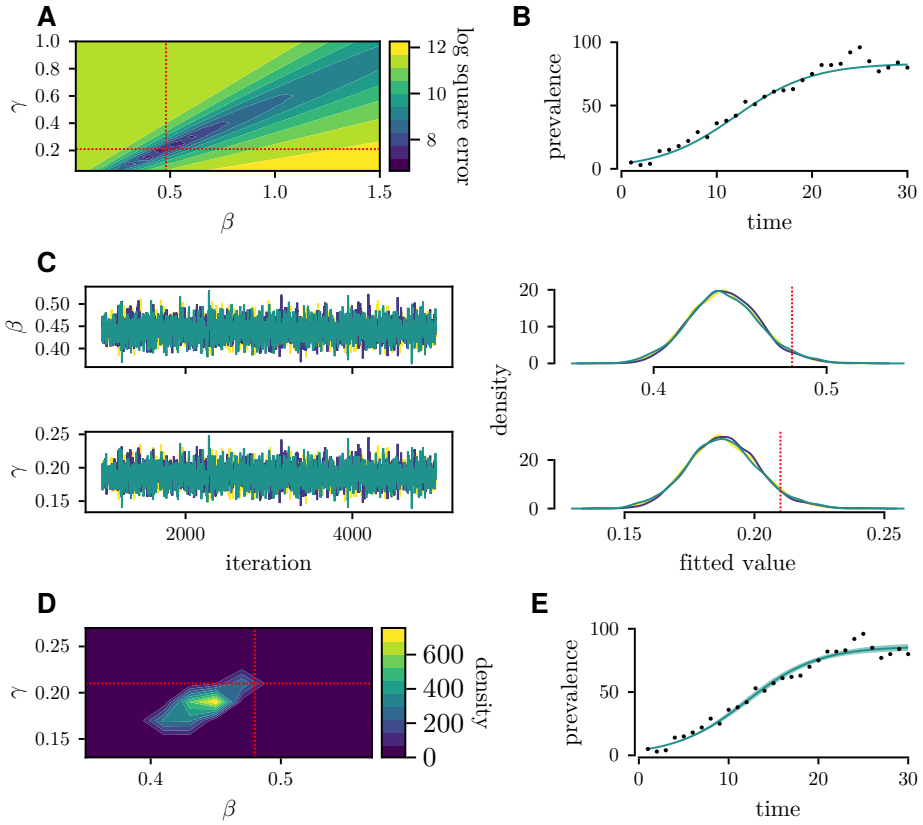

Figure 1.6: Fitting parameters to models. *A* Contour plot of log square error between simulated data and fitted model for a grid-search of β and γ parameters. *B* Simulated data (black points) and model output using β and γ parameters with least square error identified by grid-search. *C* Chains and density plots for β and γ parameters found using no-U-turn sampler. *D* Contour plot of density for parameters found using no-U-turn sampler. *E* Median (solid line) and 90% credible interval of model output using parameters from no-U-turn sampler. Dotted red lines represent 'true' values used to generate the simulated data.

prior knowledge about a parameter. For instance, if a parameter θ represented a proportion, then a prior distribution $\theta \sim \text{Uniform}(0,1)$ would imply that all values between 0 and 1 were equally likely. Between these extremes, weakly-informative prior distributions can be chosen to constrain the model to a reasonable probability space. Prior predictive simulations indicate whether the prior distributions would produce plausible outcomes, where the aim is to exclude parameters that lead to obviously impossible

outcomes but to include a full range of possible values (McElreath, 2020). Prior distributions may also need to be constrained due to structural identifiability of model parameters, where a fixed value of a parameter may be considered an extreme prior distribution in which all probability mass lies on a single parameter value.

For some combinations of data and prior distributions (termed 'conjugate priors'), a precise posterior distribution can be calculated analytically (Sugiyama, 2016). However, these are not available for most models. Rather than attempting to calculate an analytical solution, MCMC methods take series of samples from the posterior distribution (McElreath, 2020). Given sufficient time, the distribution of these samples matches the target posterior probability distribution.

The Metropolis–Hastings algorithm is an MCMC method. Staying with the SIS model example, from a starting point of possible parameters β_1, γ_1, the algorithm chooses nearby parameters, $\beta' = A(\beta_1)$, $\gamma' = B(\gamma_1)$, where A and B are choosing random samples from a probability distribution centred on the previous value of β and γ. If β', γ' fit the data better than β_1, γ_1 then $\beta_2 = \beta'$, $\gamma_2 = \gamma'$. Otherwise, an acceptance probability is calculated based on the relative fit from the two parameter sets. If a random value $X \sim \text{Uniform}(0, 1)$ is less than this probability then $\beta_2 = \beta'$, $\gamma_2 = \gamma'$, otherwise $\beta_2 = \beta_1$, $\gamma_2 = \gamma_1$. The next proposals are based on the values β_2, γ_2, and a chain of accepted parameter values is constructed. Given a sufficiently-long chain, accepted parameter values reproduce the posterior likelihood distribution of the parameters. This is a Bayesian method. The calculation of fit depends on both the distance between the estimates and the data, and the prior distribution for each parameter.

Hamiltonian Monte Carlo and the no-U-turn sampler are similar methods. Instead of a random walk to propose parameter values, the posterior

distribution is modelled as a surface and new values are identified as a function of a modelled particle's kinetic and potential energy moving across this surface.

Although the samples generated by an MCMC algorithm will *eventually* converge on the posterior distribution, there is no guarantee that a finite set of samples has yet converged. Evidence for convergence may be gained by examining trace plots of the chains (McElreath, 2020). Figure 1.6C shows a trace plot for three chains estimating parameters from a simulated dataset. These chains are stable, well mixed, and the separate chains appear to have converged on the same distribution. Diagnostic $\hat{R}$ values quantify convergence of MCMC chains (Vehtari *et al.*, 2021). An $\hat{R}$ value calculated from multiple chains approximately equals 1 when chains have converged and is greater for chains that have not yet converged. A threshold value of less than 1.05 is typically used to indicate that chains converged sufficiently.

MCMC methods can struggle to map a posterior distribution if its gradient suddenly changes or if there are large plateaux in which there is little change in the posterior density. Syed *et al.* (2022) describe a parallel-tempered MCMC method. Parallel chains sample the posterior, the prior distribution, and a series of intermediate distributions. These distributions should be increasingly easy to sample the more they resemble the prior distribution. Samples from simpler distributions are assessed against the next more complicated distribution. The sample from the simpler distribution may be substituted into a chain sampling from a more complicated distribution according to a Metropolis–Hastings acceptance probability. This helps to ensure that the algorithm samples the whole posterior distribution even if it has a very irregular shape. The algorithm is implemented in `Pigeons.jl` (Surjanovic *et al.*, 2023).

I used MCMC methods for each analysis in this thesis. Prior distributions

were selected to reflect subject knowledge (for example plausible basic reproduction ratios and generation intervals) and to produce reasonable prior predictive simulations. Where parameters were structurally unidentifiable (Chapter 3), I fixed values of these parameters at each of a range of values to assess how these affected conclusions drawn from the model. I initially used the no-U-turn sampler but found that this was unable to sample to posterior distributions for the models in Chapters 2 and 3, despite attempts to reparameterize the model and constrain prior distributions. I therefore used the parallel-tempered method in `Pigeons.jl` (version 0.4.5) for each analysis presented in this thesis. For some models, problems with chain convergence persisted. Where trace plots indicated that a chain had not converged with other chains, that chain was discarded. In an extreme case where only one chain produced a stable result, that chain was used alone as an estimate of the posterior distribution.

Stochastic models

In contrast to the deterministic models described so far, stochastic models produce a different output each time they run. In this thesis I use stochastic models to simulate datasets with noise but do not attempt to fit parameters to stochastic models.

Gillespie (1977) described a stochastic compartmental model in continuous time. I used this model to generate the simulated SIS-like data on page 48. To run this algorithm, each transition rate (including those proportional to the force of infection) is described. For the SIS example,

$$\left.\begin{aligned} S \xrightarrow{\alpha_1} I, \quad \alpha_1(t) = \beta\, S(t)\, I(t)/n, \\ I \xrightarrow{\alpha_2} S, \quad \alpha_2(t) = \gamma\, I(t). \end{aligned}\right\} \tag{1.35}$$

The event that occurs next is determined in proportion to their calculated

rates, so

$$\Pr(\text{next event} = x) = \frac{\alpha_x}{\sum_i \alpha_i}.$$

The time interval until the next event is given by,

$$\Delta t = \frac{-\ln X}{\sum_i \alpha_i}, \quad X \sim \text{Uniform}(0, 1).$$

So for the SIS example, if the next event is event 1 ($S \to I$), then $S(t + \Delta t) = S(t) + 1$, $I(t + \Delta t) = I(t) - 1$.

1.5 Ethical considerations

The research in this thesis uses publicly available data that did not allow identification of individuals. Ethics committee approval was not required. See supplementary section A.1 for correspondence with the medical sciences ethics committee confirming this.

1.6 Software and code

All analysis was performed in Julia (version 1.9.3; Bezanson *et al.*, 2017). I used `DrWatson.jl` (version 2.16.0; Datseris *et al.*, 2020) to produce a reproducible environment for each chapter's analysis. Figures from analyses are plotted with `CairoMakie.jl` (version 0.12.9; Danisch and Krumbiegel, 2021). Code for each chapter is available online. Code to reproduce the figures in this introductory chapter is available at `https://github.com/markgpritchard/IntroductoryFigures`.

Chapter Two

Difference-in-differences methods for epidemic diseases

Various infection-prevention measures were applied in different countries and regions during the covid-19 pandemic (Hale *et al.*, 2021). A natural question when seeking to understand effects of these interventions is: *How did outcomes in treated regions differ from those in untreated regions?* It is highly likely that underlying differences (heterogeneity) between treated and untreated regions would have led to differences in outcomes, even under the sharp null hypothesis that the intervention had no effect on any outcomes. It is also likely that some relevant confounding variables are unmeasured.

Bollyky *et al.* (2023) and COVID-19 National Preparedness Collaborators (2022) included vast numbers of covariates, including health-related, social, environmental and economic variables, in regression analyses of differences in SARS-CoV-2 cumulative incidence and mortality between countries and between states in the United States. They included vaccine uptake and use of non-pharmaceutical infection-prevention measures as potential predictors, along with variables without a direct causal path to outcomes, such as proportions voting for each candidate in the 2020 United States presidential election. Within the United States, vaccine uptake and mask use were estimated to have large effects on infections but estimates of effectiveness of

most public-health policy responses showed very large uncertainty intervals (Bollyky *et al.*, 2023). Flaxman *et al.* (2020) and Liu *et al.* (2024) performed regression analyses to estimate effects of non-pharmaceutical interventions on the time-varying (effective) reproduction ratio, as a measure of transmissibility. Flaxman *et al.*'s (2020) analysis considered transmission in eleven European countries until May 2020. They found a substantial effect on transmissibility from large-scale lockdowns but were unable to identify specific effects for most other interventions as the timing of their implementation had been so rapid in most European countries. Liu *et al.* (2024) considered interventions in forty-seven European countries between 2020 and 2022. They divided their analysis into separate periods defined by changes in the most prevalent SARS-CoV-2 variant of concern. For most interventions, estimated effects on transmission differed in both magnitude and direction at different times during the pandemic. They also found temporal clustering of interventions, making estimates of individual interventions' effects difficult to interpret.

We can suppose that some variables affecting outcomes differ between groups. For instance, a densely-populated city might have different levels of disease transmission to a sparsely-populated rural area (Buckee *et al.*, 2021). Other variables will change over time, such as those affected by evolution of a new strain of the virus. If all group-specific variables are constant over time, and all time-varying variables apply to all groups, then in the absence of any intervention, groups' outcomes should vary in parallel. This is the *parallel-trends* or *common-trends assumption* (Roth *et al.*, 2023; Wing *et al.*, 2018). If this holds, we can generate a counterfactual outcome for each treated group that moves in parallel to the untreated groups' observed outcomes (Figure 2.1*A*). The mean difference between this counterfactual and the observed outcomes for treated groups is termed the average treatment

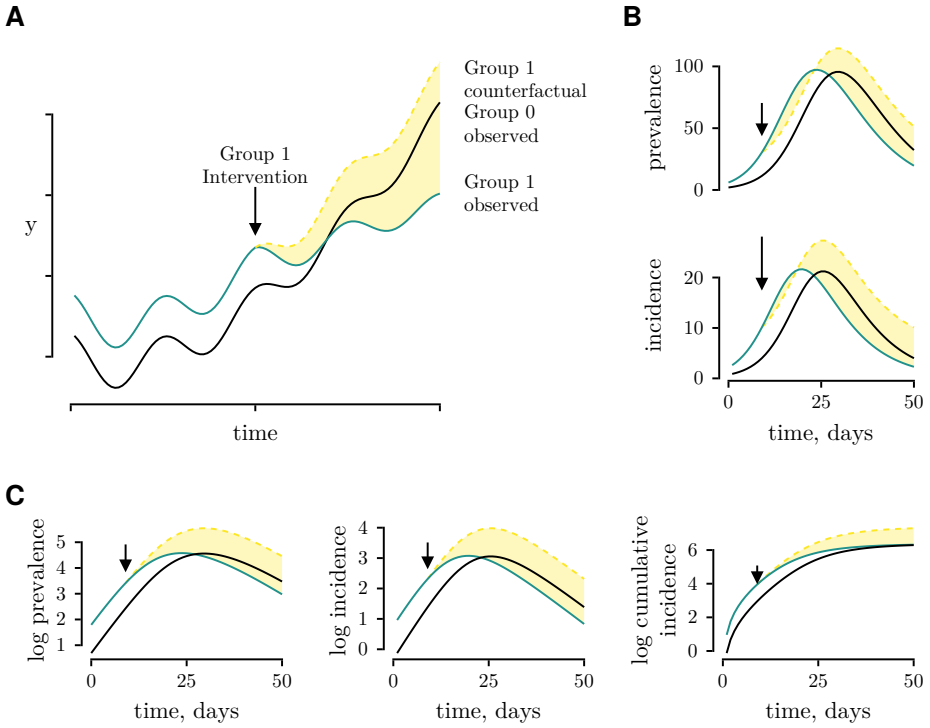

Figure 2.1: Representations of the parallel-trends assumption for two groups. Arrows indicate when Group 1 (green) starts an intervention. Solid lines represent observations, and the dashed yellow line represents a counterfactual outcome for Group 1 parallel to Group 0's (black) observed outcome. The shaded area represents the difference between Group 1's observed and counterfactual outcomes. *A.* Hypothetical representation with an effective intervention. *B.* Simulated data with identical transmission parameters for both groups (no effective intervention; see supplementary text A.2). *C.* Natural logarithms of the simulated data with no effective intervention.

effect in the treated (ATT) (Rothbard *et al.*, 2024).

John Snow's study of the causes of cholera in London in 1855 has been described as the first difference-in-differences analysis (Caniglia and Murray, 2020). Much of the work developing these methods has taken place in the field of econometrics, but differences in differences remain popular in studies of healthcare interventions (Rothbard *et al.*, 2024). Feng and Bilinski (2024) identified twenty-nine studies of interventions against SARS-CoV-2 that used difference-in-differences methods to estimate effects. Most of these

studies assumed parallel trends in numbers or proportions infected or dying from infection. But we should not expect parallel trends in incidence. Figure 2.1*B* shows two simulated outbreaks following SIR dynamics. Both simulations use the same parameters with different initial conditions. No intervention is sumulated but prevalence and incidence are never parallel. A counterfactual outcome for one group (in green) that moves in parallel with the other group initially suggests that an intervention has caused additional infections. After the epidemic has reached its peak, the counterfactual suggests that an intervention averted infections. Logarithms of incidence or prevalence (the second-most popular strategy found by Feng and Bilinski, 2024) initially appear to follow parallel trends, but once the epidemic has reached its peak, the counterfactual again predicts that without an intervention there would have been many more cases (Figure 2.1*C*). Callaway and Li (2023) show that economic effects of an epidemic follow similar non-parallel trends when costs are proportional to prevalence.

I propose that a pathogen's effective reproduction ratio is the target of infection-prevention interventions and should be considered in difference-in-differences analyses of such interventions. Vaccines aim to reduce it by reducing the proportion susceptible to infection. Physical interventions aim to reduce it by reducing the basic reproduction ratio. In this chapter I propose a method of investigating a physical intervention's effect on the basic reproduction ratio. I apply the method to a series of simulated datasets, which are intended to be progressively more like real-world data and less consistent with the parallel-trends assumption. I then apply it to two interventions: recommendations and mandates to use face coverings in the UK in summer 2020; and a community-testing pilot using rapid antigen-test lateral-flow devices in Liverpool City Region in autumn 2020.

2.1 Background and data

Community testing

In early 2020, microbiological diagnosis of SARS-CoV-2 was available only with reverse transcription–polymerase chain reaction testing (UK Health Security Agency, 2022c). People who met symptomatic and epidemiological criteria visited a hospital for testing. During the following months, availability of testing greatly increased, with testing sites opened nationally. By mid-June 2020, 100 000 tests were performed daily across the UK (*ibid.*). The proportion of infections that were recorded is estimated to have been 40–55 per cent by September 2020 (*ibid.*). Tests in the UK were classified into four 'pillars' (Department of Health and Social Care, 2020b):

1. tests of patients in NHS hospitals, and health and care workers;
2. antigen testing of the wider population;
3. serology testing; and
4. testing for the purpose of surveillance.

After the first national restrictions were relaxed in summer 2020, local restrictions were introduced in areas of England with the greatest levels of transmission. From 3 October the six local authorities in the Liverpool City Region (Halton, Knowsley, Liverpool City, Sefton, St Helens and Wirral) and neighbouring Warrington were placed at a higher level of restriction (*The Health Protection (Coronavirus, Restrictions) (Protected Areas and Restriction on Businesses) (Amendment) Regulations* 2020). On 14 October 2020, new regulations were introduced with different levels of restriction for each of three different 'tiers' or 'alert levels' (Brown and Kirk-Wade, 2021). The six local authorities in the Liverpool City region were all in the highest level of restriction (University of Liverpool, 2021b). These restrictions were replaced by national regulations on 5 November (Brown and Kirk-Wade, 2021).

The Liverpool mass testing pilot started on 6 November 2020 (University of Liverpool, 2021b). This pilot allowed asymptomatic people to take a rapid antigen test for SARS-CoV-2. Tests were used to identify infections in vulnerable people such as residents of care homes, to allow key workers to reduce their isolation period if they were recovered, and to allow events to take place that would otherwise have been prohibited under national regulations.

A second set of tiered restrictions was introduced on 2 December, when the Liverpool City Region local authorities were placed in tier 2 (on a scale 1–3; Green *et al.*, 2021). On 3 December, rapid testing was rolled out across the Liverpool City Region (University of Liverpool, 2021a; Zhang *et al.*, 2022). The six Liverpool local authorities moved to tier 3 of restrictions on 31 December (Department of Health and Social Care, 2020a). On 6 January 2021, all parts of England moved to tier 4 (the third national lockdown; Brown and Kirk-Wade, 2021). The neighbouring districts of Warrington, West Lancashire and Wigan did not have mass testing during 2020, and are included as untreated groups in each analysis.

Total and pillar 1 numbers of SARS-CoV-2 diagnoses for each lower-tier local authority are publicly available from the Department of Health and Social Care (2021).

Face coverings

When SARS-CoV-2 first spread to the UK in late January 2020, face coverings were not recommended for the general public due to a lack of evidence of effectiveness (Ruppel Shell, 2020). They were still not recommended in March 2020 when emergency regulations were imposed across the UK. Their use was officially recommended first in Scotland on 28 April, then in England on 12 May, and in Wales on 9 June (Hale *et al.*, 2021). These recommenda-

Table 2.1: Changes in the United Kingdom government's response to covid-19 during 2020 that are included in this analysis

Nation	Mask mandate	Mask recom-mendation	End of stay-at-home restrictions	Reopening businesses*
England	15 June	12 May	15 June	13 May
Northern Ireland	10 July	10 July	31 July	18 May
Scotland	22 June	28 April	NA	29 May
Wales	14 September	9 June	16 August	11 May

* These are dates that any businesses that had been closed were first allowed to reopen, typically outdoors-based businesses such as garden centres.
NA, not applicable (only dates between January and 14 September 2020 are shown)
Data from Hale *et al.*, 2021.

tions became requirements to wear face coverings in at least some settings in Scotland from 22 June, in England from 26 August, and in Wales from 14 September. In Northern Ireland, a requirement for facial coverings was instituted on 10 July, with no prior recommendation for their use having been in place (Hale *et al.*, 2021). Other policies changed at the same time as changes in requirements for face coverings. For instance, recommendations to use face coverings in England accompanied reductions in stay-at-home guidance (*ibid.*). Dates of events included in this analysis are shown in Table 2.1.

The dataset collated by Hale *et al.* (*ibid.*) includes daily incidence of recorded covid-19 diagnoses for each of constituent nation of the UK since January 2020, which I used for the analysis of face coverings.

Evolution of severe acute respiratory syndrome coronavirus 2

Several changes in SARS-CoV-2 were observed during 2020. In April 2020, a point mutation, D614G, in the spike protein was associated with increased infectiousness. This mutation spread across all parts of the UK within a month of emerging, rapidly becoming the dominant variant in all parts of

the country (Korber *et al.*, 2020). In autumn 2020, the first 'variant of concern', Alpha/B.1.1.7, was detected in Kent (Grint *et al.*, 2021). Again, this was associated with greater infectiousness, and the variant rapidly became dominant across the UK (Leung *et al.*, 2021).

Serial interval of severe acute respiratory syndrome coronavirus 2

Cori *et al.* (2013) describe a method for calculating a microparasite's effective reproduction ratio and serial interval. They applied this to a SARS-CoV-2 dataset and published an estimated serial interval distribution for the virus (Figure B.1; Cori *et al.*, 2022).

2.2 Simulated data

I simulated data with an event-driven stochastic SEIR-type model, using Gillespie's (1977) method. A proportion, $\theta \in (0, 1)$ of infectious individuals enter a diagnosed compartment, I'. The model's rates are,

$$\left.\begin{aligned}
S \xrightarrow{\alpha_1} E, &\quad \alpha_1(t,j) = \beta(t,j)\, S(t,j) \left[I(t,j) + I'(t,j) \right] / n(j), \\
E \xrightarrow{\alpha_2} I, &\quad \alpha_2(t,j) = \varsigma\, E(t,j), \\
I \xrightarrow{\alpha_3} I', &\quad \alpha_3(t,j) = \frac{\theta\, \gamma}{1-\theta}\, I(t,j), \\
I \xrightarrow{\alpha_4} R, &\quad \alpha_4(t,j) = \gamma\, I(t,j), \\
I' \xrightarrow{\alpha_5} R, &\quad \alpha_5(t,j) = \gamma\, I'(t,j).
\end{aligned}\right\} \tag{2.1}$$

The number of $I \to I'$ transitions was recorded for each day as the simulated incidence of diagnosis. For all simulations, $\varsigma = 0.5, \gamma = 0.4$. The serial interval distribution was calculated using Champredon *et al.*'s (2018) result for an SEIR model,

$$g(x) = \frac{\varsigma\, \gamma}{\varsigma - \gamma} \left[\exp(-\gamma\, x) - \exp(-\varsigma\, x) \right], \quad x \in \mathbb{N},\ \varsigma,\ \gamma > 0,\ \varsigma \neq \gamma, \tag{2.2}$$

shown in Figure B.2.

For simulations with an effective intervention, the transmission parameter was reduced by 20 per cent. I generated 'true' counterfactual scenarios by repeating simulations without the intervention.

The first simulation represents the 'canonical' difference-in-differences conditions (Roth *et al.*, 2023), with two locations, and a transmission parameter that changes once, on day 50. Subsequent simulations have additional interventions, non-parallel trends and changes in proportions of cases diagnosed. Details of these simulations are given with the results.

2.3 Common trends in a renewal equation model

Renewal equation model

Assume that new infections in group j arise at time t according to a discrete-time renewal equation,

$$i_{jt} = \mathcal{R}_e(j,t) \left[\sum_{x=1}^{t-1} i_{jx}\, g(t-x) \right] \varepsilon_{jt}, \tag{2.3}$$

where i_{jt} is the daily incidence, $\mathcal{R}_e(j,t)$ is the effective reproduction ratio, $g(x)$ is the serial-interval distribution, $\sum_{x=1}^{\infty} g(x) = 1$, and ε_{jt} is a noise term, $\mathbb{E}(\ln \varepsilon_{jt}) = 0$. This assumes that all new infections in a group are caused by infectious individuals within that group.

Assume that the outbreak, or period of study, is sufficiently brief that numbers of births, deaths and migrations do not contribute significantly to disease dynamics, and that each group has a constant population, n_j. Also assume that immunity lasts for at least the remaining duration of the epidemic so no one returns to susceptibility during the period of study. Let $\rho_j(t)$ be the basic reproduction ratio for group j at time t, so

$$\mathcal{R}_e(j,t) = \rho_j(t)\, \frac{s_{jt}}{n_j} = \rho_j(t) \left[\frac{s_{j0} - \sum_{x=0}^{t-1} i_{jx}}{n_j} \right], \tag{2.4}$$

where s_{jt} is the number of susceptible individuals. The value of $\rho_j(t)$ is only evaluated at discrete values of t but is allowed to vary continuously so that a cubic spline can be fitted to it. In this chapter I assume a fully-susceptible population at time t_0, $s_{j0} = n_j$, so

$$\mathcal{R}_e(j,t) = \rho_j(t) \left[1 - \sum_{x=0}^{t-1} \frac{i_{jx}}{n_j}\right]. \tag{2.5}$$

Let $y_{jt} = \theta\, i_{jt}$ be the number of infections recorded in group j at time t, where $\theta \in (0,1]$ is the constant proportion of infections that are diagnosed. Substituting into the renewal equation (2.3),

$$y_{jt} = \mathcal{R}_e(j,t) \left[\sum_{x=1}^{t-1} y_{jx}\, g(t-x)\right] \varepsilon_{jt}, \tag{2.6}$$

and substituting into the effective reproduction ratio (2.5),

$$\mathcal{R}_e(j,t) = \rho_j(t) \left[1 - \sum_{x=0}^{t-1} \frac{y_{jx}}{\theta\, n_j}\right]. \tag{2.7}$$

Causal estimand and the average treatment effect in the treated

Let $A_{jt} = 1$ indicate that the infection-prevention intervention being studied is in use in group j at time t, and $A_{jt} = 0$ otherwise. If the intervention is effective, it is expected to alter the basic reproduction ratio. Let $\rho_j(t)^{A=a}$ represent the potential value of $\rho_j(t)$ if $A_{jt} = a$. Assume infection-prevention interventions have a proportional effect on disease transmission, so the unit-level causal effect of the intervention is the ratio,

$$\tau_{jt} = \frac{\rho_j(t)^{A=1}}{\rho_j(t)^{A=0}}, \tag{2.8}$$

where the 'unit' is the group of individuals.

The parallel-trends assumption allows us to estimate an untreated counterfactual outcome for the groups that were treated. It does not allow estimation of a treated counterfactual for those groups that were untreated. We

therefore estimate the ATT, $\tau_{\text{ATT}} = \mathbb{E}\left(\tau_{jt} \mid A_{jt} = 1\right)$. By the requirement for consistency, $\mathbb{E}(\rho_j(t)^{A=a} \mid A_{jt} = a) = \mathbb{E}(\rho_j(t) \mid A_{jt} = a)$, so

$$\tau_{\text{ATT}} = \mathbb{E}\left(\frac{\rho_j(t)}{\rho_j(t)^{A=0}} \mid A_{jt} = 1\right). \tag{2.9}$$

To estimate a 'difference', groups can be included in the analysis only if they are observed before the intervention starts, so $A_{j0} = 0$ for all groups. If a mandate is reversed, we do not necessarily expect that transmission will return to its pre-transmission level. For example, individuals who started using masks because of a face-covering mandate may continue to use them after the mandate ends (Adjodah *et al.*, 2021). I therefore add a requirement that $A_{j(t+1)} \geq A_{jt}$.

Parameters for the common-trends assumption

For the common-trends assumption to hold, any differences between groups must be constant over time, and any changes with time must apply equally to all groups (Caniglia and Murray, 2020). I assume that such effects are multiplicative. Let ζ_j be a time-invariant parameter to account for location-specific effects on transmission and let $\eta(t)$ be a time-varying parameter to account for population-wide changes in transmission. The expected basic reproduction ratio for each group and time without the intervention is,

$$\mathbb{E}(\rho_j(t)^{A=0}) = \zeta_j\, \eta(t). \tag{2.10}$$

From equations 2.9 and 2.10,

$$\begin{aligned} &\mathbb{E}(\rho_j(t)^{A=1} \mid A_{jt} = 1) = \tau_{\text{ATT}}\, \mathbb{E}(\rho_j(t)^{A=0} \mid A_{jt} = 1), \\ \Rightarrow\ &\mathbb{E}(\rho_j(t)^{A=a'} \mid A_{jt} = a) = \zeta_j\, \eta(t)\ {\tau_{\text{ATT}}}^{a'}, \quad a, a' \in \{0,1\},\ a' \leq a. \end{aligned} \tag{2.11}$$

Identifying parameters

Let $z_{jt} := \ln \mathcal{R}_e(j,t)$. From equation 2.7,

$$z_{jt} = \ln \rho_j(t) + \ln\left(1 - \sum_{x=1}^{t-1} \frac{y_{jx}}{\theta\, n_j}\right). \tag{2.12}$$

From the renewal equation (2.6),

$$z_{jt} = \ln y_{jt} - \ln\left(\sum_{x=1}^{t-1} y_{jx}\, g(t-x)\right) - \ln \varepsilon_{jt}. \tag{2.13}$$

If the generation interval, $g(x)$, is known, z_{jt} can be estimated from incidence data,

$$\hat{z}_{jt} = \ln y_{jt} - \ln\left(\sum_{x=1}^{t-1} y_{jx}\, g(t-x)\right), \tag{2.14}$$

provided $y_{jt} > 0$ and $\sum_{x=1}^{t-1} y_{jx}\, g(t-x) > 0$. When $A_{jt} = a_{jt}$, $\rho_j(t)$ is the potential value $\rho_j(t)^{A=a_{jt}}$. Substituting this into equation 2.12,

$$\mathbb{E}\left(z_{jt} \mid A_{jt} = a_{jt}\right) = \ln \rho_j(t)^{A=a_{jt}} + \ln\left(1 - \sum_{x=0}^{t-1} \frac{y_{jx}}{\theta\, n_j}\right). \tag{2.15}$$

From equation 2.11,

$$\ln \rho_j(t)^{A=a_{jt}} = \ln \zeta_j + \ln \eta(t) + a_{jt}\, \ln \tau_{\text{ATT}}, \tag{2.16}$$

so,

$$\mathbb{E}\left(z_{jt} \mid A_{jt} = a_{jt}\right) = \ln \zeta_j + \ln \eta(t) + a_{jt}\, \ln \tau_{\text{ATT}} + \ln\left(1 - \sum_{x=0}^{t-1} \frac{y_{jx}}{\theta\, n_j}\right). \tag{2.17}$$

Taking the estimated value of $\hat{z}_{jt}$ from equation 2.14, and this equation for $\mathbb{E}\left(z_{jt} \mid A_{jt} = a_{jt}\right)$ (2.17), I estimate values of ζ_j, $\eta(t)$ and τ_{ATT}. Not all values of ζ_j and $\eta(t)$ can be uniquely identified, as their product would be unchanged if all ζ coefficients were multiplied by some non-zero constant and all η coefficients divided by the same constant. I fix one value of $\ln \eta(k) = 0$. I solve for the remaining parameters in a Bayesian regression, defining the model in Turing.jl (version 0.34.1; Ge *et al.*, 2018), and solving

with the parallel-tempered MCMC method (Syed *et al.*, 2022) implemented in Pigeons.jl (version 0.4.5; Surjanovic *et al.*, 2023). I use a hierarchical model to estimate values of ζ and η. For location-specific parameters, I assume the mean has a normal ($\mathcal{N}$) distribution, then that each location's parameter is normally distributed about that mean,

$$\mu_\zeta \sim \mathcal{N}(0,1), \quad \sigma_\zeta^2 \sim \text{Exponential}(1),$$
$$\ln \zeta_j \sim \mathcal{N}(\mu_\zeta, \sigma_\zeta^2), \quad j \in \mathbb{N}_{>0}.$$

I use a similar approach for time-varying parameters. For the initial simulation, I assume that the time-varying parameter moves between discrete values so,

$$\mu_\eta \sim \mathcal{N}(0,1),\ \sigma_\eta^2 \sim \text{Exponential}(1), \quad \ln \eta(t) = \begin{cases} 0, & t < \phi_j, \\ \eta_k \sim \mathcal{N}(\mu_\eta, \sigma_\eta^2), & t \geq \phi_j, \end{cases}$$

where ϕ_j is the time that the intervention is introduced in group j. It is more plausible that the transmission parameter will vary continuously over time. To account for this, I fit a cubic spline (Akima, 1970). The values of $\ln \eta_k$ at each knot are fitted with a hierarchical model, equivalent to that used for ζ, with the second value fixed to 0 to allow unique identifiability.

The prior distribution for θ is assumed to follow a Beta distribution with a mean value equal to the simulated θ when applied to simulations, and equal to 0.4 (matching the 40% estimated detected; UK Health Security Agency, 2022c) when applied to SARS-CoV-2 data. This distribution is then truncated to ensure that,

$$\theta > \max_j \left(\frac{\sum_t y_{jt}}{n_j} \right),$$

preventing θ being less than the proportion of any group's population that has been diagnosed, which would imply $\sum_t i_{jt} > n_j$. The prior distribution for the causal effect is $\ln \tau_{\text{ATT}} \sim \mathcal{N}(0,1)$.

Estimating effect on incidence

Renewal-equation models need to be primed with a history of previous infections. I expect the time-varying parameter to be unreliable for early time points when a large proportion of infections are likely to have been imported rather than arising from infectious individuals within the group. I therefore set the modelled incidence to equal recorded incidence for a period at the start of each analysis (the first ten days for simulated datasets; until 1 April 2020 for analysis of the effects of face coverings; and for eight weeks from the start of June 2020 for analysis of the effect of community testing in Liverpool). For subsequent times I use fitted parameters in the renewal equation to estimate trends in infection numbers.

Running the model with samples from the prior distributions includes most calculated values of the logarithm of the effective reproduction ratio within the central 90% credible interval (CrI) for both the testing (Figure B.3 and B.4) and face coverings (Figure B.5) analyses. Note that the width of the CrI reduces near the knot where the time-varying parameter is fixed at 0, and increases after the intervention as the uncertainty of the τ_{ATT} parameter is added.

2.4 Application to simulated data

'Canonical' difference in difference in transmission (Simulation 1)

In these 'canonical' difference-in-differences conditions, the control group's basic reproduction ratio is 1.5 for the first 50 days, then increases to 1.65. In the intervention group, without an effective intervention the basic reproduction ratio is 1.15 times the other group's. If the intervention is effective then on day 50 the basic reproduction ratio reduces by 20 per cent compared to the untreated potential outcome.

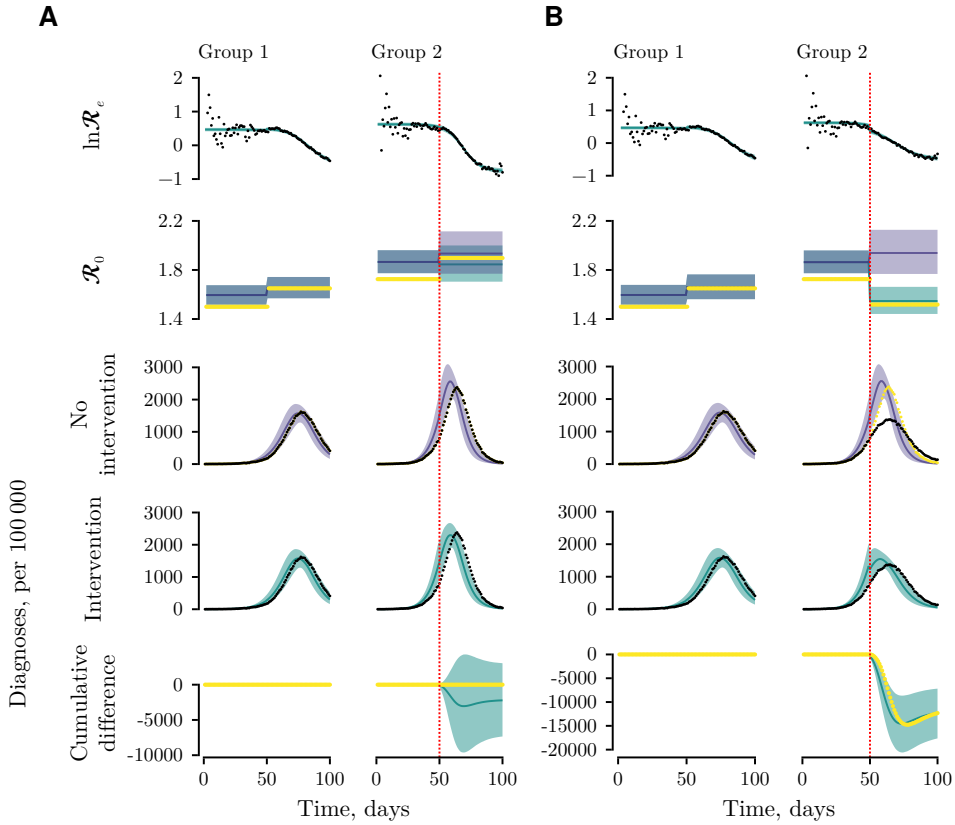

Figure 2.2: Difference-in-differences analysis for simulated groups with two discrete levels of transmission. *A* Simulation with an ineffective intervention; *B* Simulation with an intervention that reduces transmission by 20%. Black dots show simulated data. Yellow dots show values that are known from the simulation but would not be known from data: basic reproduction ratio, simulated counterfactual number of cases, and simulated causal effect on cumulative number of cases. Green lines show results fitted to the data, and purple lines show inferred intervention-free counterfactuals. The bottom panel shows the estimated causal effect as a cumulative difference in numbers diagnosed per 100 000 population. Shaded areas represent the central 90% credible interval. Vertical dotted red lines show times interventions start. $\ln \mathcal{R}_e$ is the natural logarithm of the effective reproduction ratio; $\mathcal{R}_0$ is the basic reproduction ratio.

Fitting with discrete time-varying transmission parameters gives a qualitatively good fit to the logarithm of the effective reproduction ratio (Figure 2.2). For the simulation with no effective intervention, *i.e.* true $\tau_{\text{ATT}} = 1$, fitted parameters give an estimated value of $\hat{\tau}_{\text{ATT}} = 0.95$ (90% CrI 0.86, 1.06). For the simulation with an effective intervention, true $\tau_{\text{ATT}} = 0.8$,

the fitted parameter was $\hat{\tau}_{\text{ATT}} = 0.80$ (90% CrI 0.72, 0.88). Similar results were found with the spline-based method for estimating parameters (Figure B.6). Fitted parameters for the simulation without an effective intervention gave $\hat{\tau}_{\text{ATT}} = 0.97$ (90% CrI 0.87, 1.08), and with an effective intervention, $\hat{\tau}_{\text{ATT}} = 0.82$ (90% CrI 0.75, 0.90).

Known confounding variables (Simulation 2)

A challenge to using difference-in-differences methods during an epidemic is that multiple interventions may be implemented. If other interventions are effective and are introduced in different places at different times, this violates the common trends assumption.

Simulation 2 includes three groups. The intervention of interest is used from day 50 in group 2, and from day 30 in group 3. An additional exposure in group 3 increases transmission by 15% from day 70, representing relaxation of an infection-control measure after the peak of the outbreak. This violates the common-trends assumption.

Fitted parameters map closely to the logarithm of the effective reproduction ratio but do not give a good fit to the basic reproduction ratio (Figure 2.3*A*). This is despite the second intervention coming after most cases have been diagnosed. The estimated effect of the main intervention is $\hat{\tau}_{\text{ATT}} = 0.84$ (90% CrI 0.79, 0.91).

Consider that we know when the additional exposure occurred but do not know the magnitude of its effect. In the same way as for the intervention of interest, we can represent its effectiveness in terms of a change in the basic reproduction ratio. I add an additional parameter to the formulation of $\rho_j(t)$ to account for two possible interventions,

$$\mathbb{E}\left(\rho_j(t)^{A_1=a_1,\, A_2=a_2}\right) = \zeta_j\, \eta(t)\; {\tau_{\text{ATT}\,1}}^{a_1}\; {\tau_{\text{ATT}\,2}}^{a_2}, \quad a_1,\, a_2 \in \{0,\, 1\}. \tag{2.18}$$

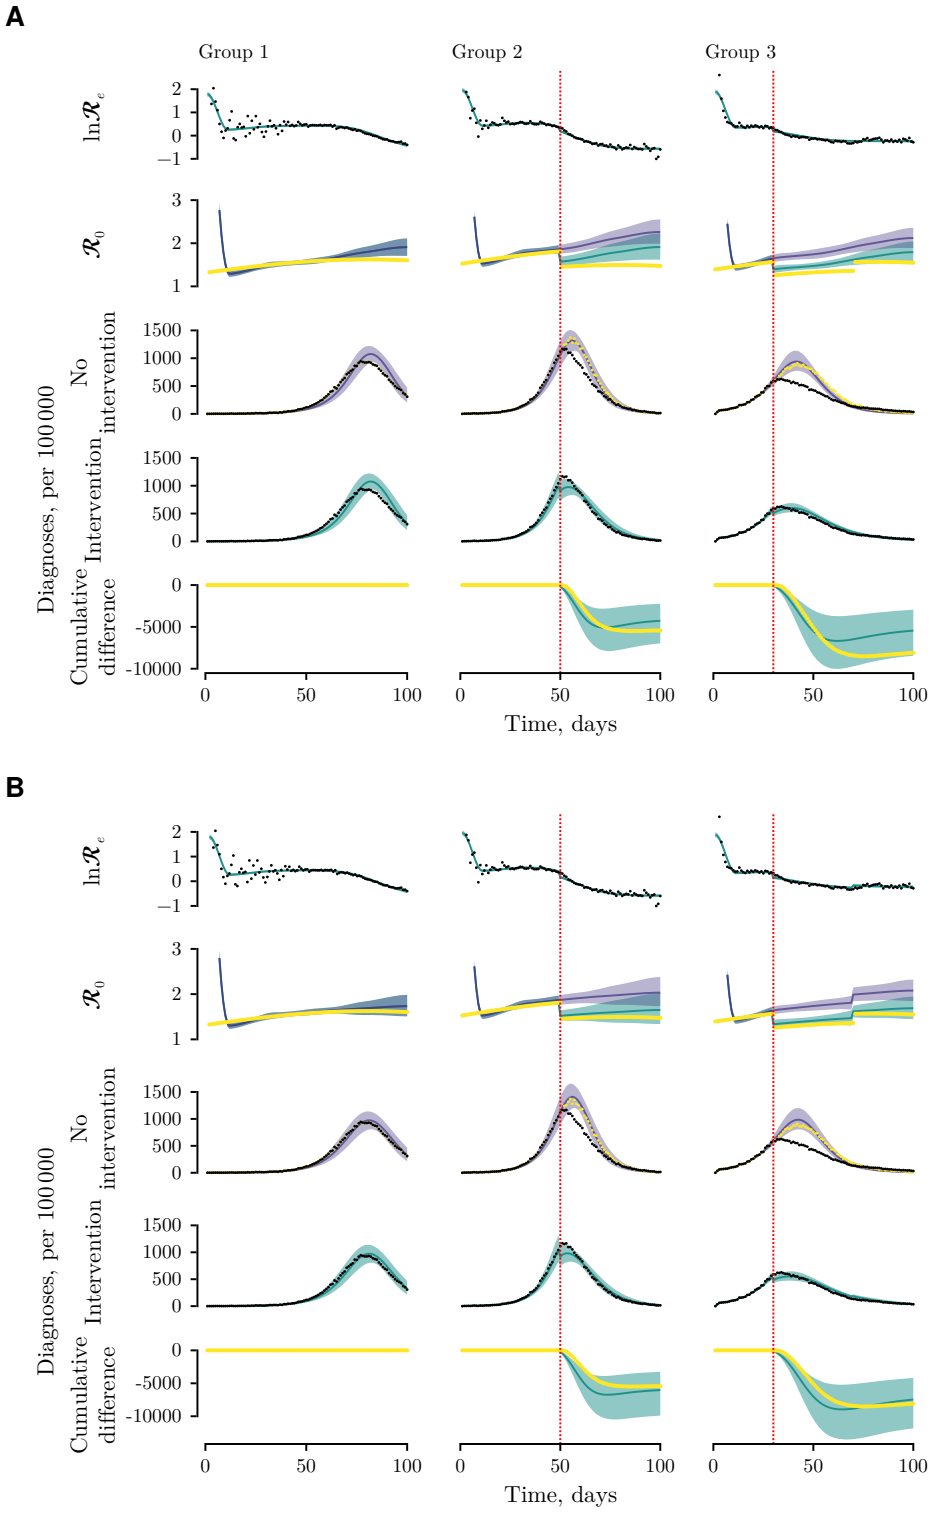

Figure 2.3: Difference-in-differences analysis for a simulation with an additional known intervention (day 70 in group 3). *A* Parameters fit without including details of other intervention; *B* Parameters fit with an additional parameter included for the other intervention. See caption to Figure 2.2 for description of colours.

More generally, we could consider n interventions,

$$\mathbb{E}\Big(\rho_j(t)^{A_1=a_1, A_2=a_2, \ldots, A_n=a_n}\Big) = \zeta_j \eta_t \prod_{x=1}^{n} \tau_{\mathrm{ATT}\,i}{}^{a_x}, \quad a_x \in \{0, 1\} \text{ for all } x \in \{1, 2, \ldots, n\}. \tag{2.19}$$

Substituting this form of $\rho_j(t)$ into Equation 2.17 gives,

$$\mathbb{E}(z_{jt} \mid A_{jt\,1} = a_{jt\,1}, \ldots, A_{jt\,n} = a_{jt\,n}) = \ln \zeta_j + \ln \eta(t) + \sum_{x=1}^{n} \left(a_{jt\,x} \ln \tau_{\mathrm{ATT}x}\right) + \ln\left(1 - \sum_{x=1}^{t-1} \frac{y_{jx}}{\theta\, n_j}\right). \tag{2.20}$$

Applying this formula to the simulation gives a better fit for the basic reproduction ratio and the effect on the number of infections (Figure 2.3*B*), and estimates of effectiveness closer to the simulated values of $\hat{\tau}_{\mathrm{ATT}} = 0.81$ (90% CrI 0.75, 0.88).

Non-parallel transmission trends (simulation 3)

In simulation 3, I add non-parallel trends to the transmission parameter. Such a trend could be due to an unknown time-varying confounding variable, or it may be that location- and time-specific parameters do not have the simple multiplicative relationship described in the model. For this simulation, the basic reproduction number for group 1 is,

$$\rho_1(t)^{A=0} = 1.25 + 0.25 \cos(2\pi(t - 20)/365).$$

For group 2, without an intervention,

$$\rho_2(t)^{A=0} = 0.9 + 0.005t\, \rho_1(t)^{A=0}.$$

Fitting parameters to simulated data with no effective intervention gives a result suggesting that an intervention increased the infectiousness, $\hat{\tau}_{\mathrm{ATT}} = 1.12$ (90% CrI 1.01, 1.25; Figure 2.4*A*). Note that although we can see in the simulation that the fitted basic reproduction ratio and intervention-free

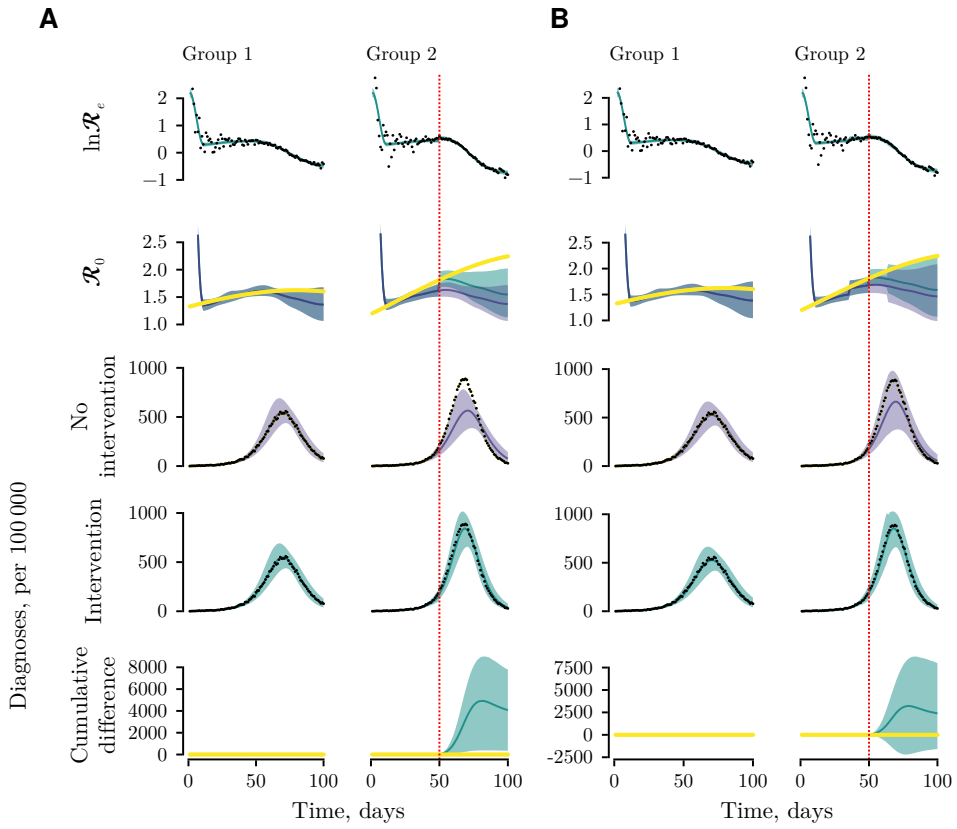

Figure 2.4: Difference-in-differences analysis for a simulation with a non-parallel trend and an ineffective intervention. *A* Parameters fit without including consideration of non-parallel trend; *B* Parameters fit with assumed unobserved changes to infectiousness 14 days before and after the intervention. See caption to Figure 2.2 for description of colours.

counterfactual do not follow the simulated values, the fit to the logarithm of the effective reproduction ratio and to the number of cases with the intervention would not raise any concerns.

The parallel-trends assumption is untestable after treatment (Rambachan and Roth, 2023; Roth *et al.*, 2023). 'Dummy' interventions can be added to the analysis to test for parallel trends prior to the intervention (Callaway and Sant'Anna, 2021; Ryan *et al.*, 2019). If the coefficients for these dummy interventions are large, this suggests that groups are diverging before the intervention (*i.e.* the trends are not parallel). Often a statistical significance test

is performed to compare the value to a null hypothesis that the coefficient is zero (Kahn-Lang and Lang, 2020; Roth, 2022). Demonstration of parallel trends before the intervention is not sufficient, or necessary, to demonstrate that trends would be parallel after the intervention, but it provides some reassurance (Kahn-Lang and Lang, 2020; Sun and Abraham, 2021). Lagged dummy interventions can also be added after the intervention to explore whether groups are diverging further over time (Wing *et al.*, 2018). In this simulation, the parallel-trends assumption is known not to hold before or after the intervention. I added dummy interventions fourteen days before and after the 'true' simulated intervention to explore how these affected the estimate of the intervention's effectiveness. Analysis with the dummy interventions gives a better fit, but still underestimates the basic reproduction ratio (Figure 2.4*B*). The estimated effect of the intervention in this analysis is $\hat{\tau}_{\text{ATT}} = 1.08$ (90% CrI 0.95, 1.23). The dummy interventions on days 36 and 64 have estimated effects $\hat{\tau}_{\text{ATT}\,1} = 1.07$ (90% CrI 0.96, 1.21) and $\hat{\tau}_{\text{ATT}\,2} = 1.00$ (90% CrI 0.86, 1.17).

For the equivalent analysis with an intervention that reduces transmission by 20 per cent (Figures 2.5 and B.7), without dummy interventions, the non-parallel trend effectively cancelled the estimated causal effect of the intervention, $\hat{\tau}_{\text{ATT}} = 1.00$ (90% CrI 0.92, 1.10). With assumed additional interventions, this changed to $\hat{\tau}_{\text{ATT}} = 0.95$ (90% CrI 0.83, 1.08).

Changes in proportion detected (simulation 4)

I have assumed that the proportion of cases diagnosed, θ, is constant over time. In reality, this assumption might be violated. For example, changes in availability of testing or awareness of symptoms could lead to changes in the proportion tested and therefore in the proportion diagnosed. Even if this time-varying proportion were known, incorporating it into the model

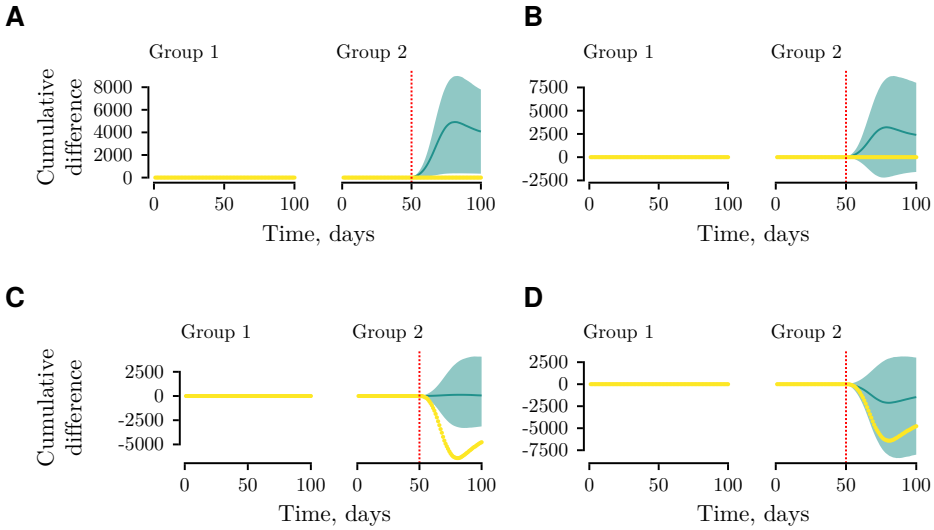

Figure 2.5: Estimated effect on number of infections for a difference-in-differences analysis for simulated data with a non-parallel trend and an effective intervention. *A* Parameters fit without including consideration of non-parallel trend; *B* Parameters fit with assumed unobserved changes to infectiousness 14 days before and after the intervention. See caption to Figure 2.2 for description of colours.

would be cumbersome. With a known $\theta(t)$,

$$\begin{aligned}\mathbb{E}(z_{jt}) &= \ln y_{jt} - \ln \theta(t) - \ln\left(\sum_{x=1}^{t-1} \frac{y_{jx}}{\theta(x)}\, g(t-x)\right) \\ &= \ln \rho_j(t) + \ln\left(1 - \sum_{x=0}^{t-1} \frac{y_{jx}}{\theta(x)\, n_j}\right),\end{aligned}$$

so at each time, estimated numbers infected and immune depend on the vector of all previous values of θ. It would be even more troublesome if a time-varying $\theta(t)$ function needed to be estimated.

I do not incorporate a time-varying value of $\theta(t)$ into the analyses here but I simulate a group with a time-varying proportion diagnosed to explore the effects of this on parameter estimates. In simulation 4, group 1 has a constant proportion diagnosed, $\theta_1(t) = 0.3$. Group 2 has the same proportion diagnosed until time 50, from which point it is increased to 0.36. The simulation's immediate change in diagnosis leads to a discontinuity in numbers diagnosed and in the effective reproduction ratio, which the model

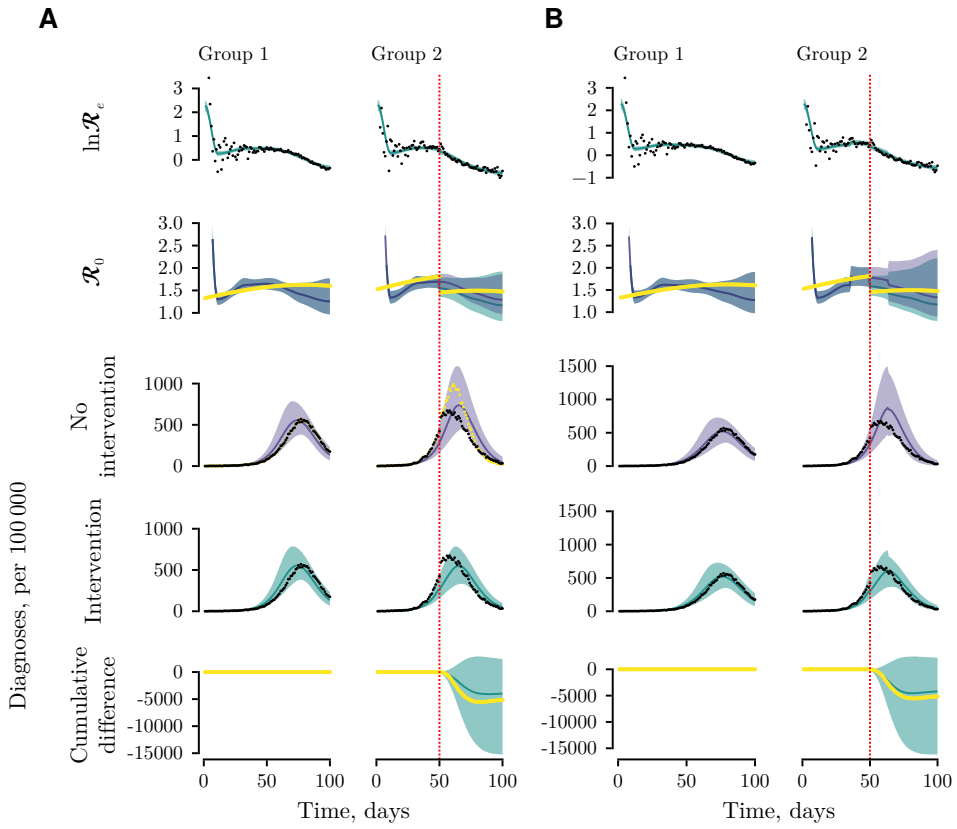

Figure 2.6: Difference-in-differences analysis for a simulation with a time-varying proportion of cases diagnosed. In group 2, there is a 20% increase in the proportion diagnosed at the same time as the intervention. *A* Parameters fit without including consideration of time-varying proportion diagnosed; *B* Parameters fit with assumed unobserved changes to infectiousness 14 days before and after the intervention. See caption to Figure 2.2 for description of colours.

fitting does not follow in simulations with or without an effective intervention (Figures B.8 and 2.6*A*). Fitting to the basic reproduction ratio is poor. The estimated causal effect for the simulation with no effective intervention is close to the 'true' value, $\hat{\tau}_{\text{ATT}} = 1.03$ (90% CrI 0.86, 1.22). For the simulation with a 20 per cent reduction in transmission, the point estimate suggests a smaller reduction than the 'true' value, although the true value is included in the 90% CrI, $\hat{\tau}_{\text{ATT}} = 0.91$ (90% CrI 0.78, 1.08). When parameters for additional hypothetical interventions were added (as in the case

with non-parallel trends), fitted parameters were very slightly closer to the simulated values, $\hat{\tau}_{\text{ATT}} = 0.89$ (90% CrI 0.74, 1.07; Figure 2.6*B*).

2.5 Application to the covid-19 pandemic

Community testing

For the primary analysis of the effect of community testing I used pillar 1 data for residents of each local authority. All chains mixed well (Figure B.9). Fitted parameters gave very high estimates of the basic reproduction ratio (Figure 2.7; local authorities included in the analysis that did not have testing in 2020 are shown in Figure B.10). This led to a large peak in modelled incidence in all districts in August 2020 which far exceeded observed data. The estimated effect of the intervention was $\hat{\tau}_{\text{ATT}} = 1.20$ (90% CrI 1.04, 1.39). This had a minimal effect on the modelled change in incidence as the expected number of susceptible individuals was so low by the time of the intervention.

Chains for analysis using all recorded diagnoses (including any asymptomatic infections that were diagnosed by the testing programme) mixed well (Figure B.11). These gave basic reproduction ratios between 1 and 3 (Figures 2.8 and B.12). The estimated effect of the intervention was $\hat{\tau}_{\text{ATT}} = 1.14$ (90% CrI 1.03, 1.27).

Face coverings

Chains for analysis of masking mandates in the UK mixed well (Figure B.13). Parameters gave a qualitatively good fit for data in England and Scotland, but appeared to underestimate incidence in Northern Ireland and overestimate it in Wales (Figure 2.9) The estimated effect of the intervention was $\hat{\tau}_{\text{ATT}} = 1.23$ (90% CrI 1.08, 1.39).

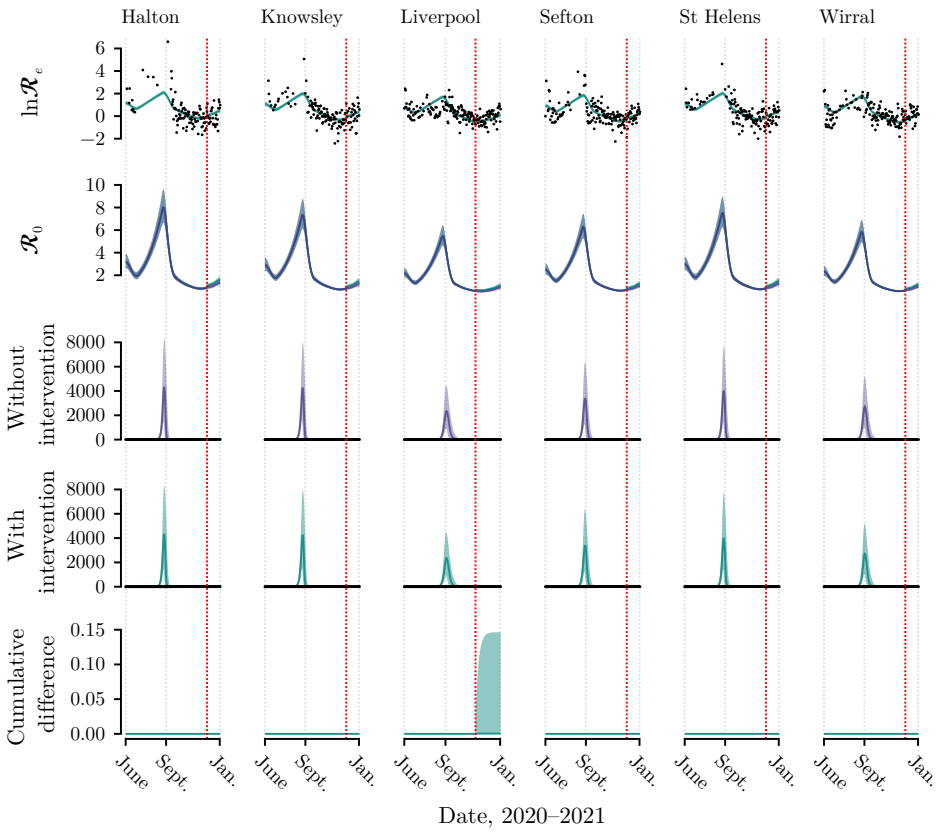

Figure 2.7: Difference-in-differences analysis for the effect of community testing on pillar 1 severe-acute-respiratory-syndrome-coronavirus-2 diagnoses in North West England between June and December 2020. Dotted red lines indicate the start of the testing pilot. See caption to Figure 2.2 for description of colours.

Adding parameters to allow changes in the basic reproduction ratio when stay-at-home rules ended, businesses reopened, and masks were recommended in advance of the mandates had little effect on the estimated effect of the masking mandate, $\hat{\tau}_{\text{ATT}} = 1.18$ (90% CrI 1.04, 1.36; Figure 2.10). The estimated effects of the secondary interventions were:

1. $\hat{\tau}_{\text{ATT}\,1} = 1.11$ (90% CrI 0.98, 1.27) for the effect of ending stay-at-home requirements;
2. $\hat{\tau}_{\text{ATT}\,2} = 1.01$ (90% CrI 0.83, 1.23) for the effect of businesses reopening; and

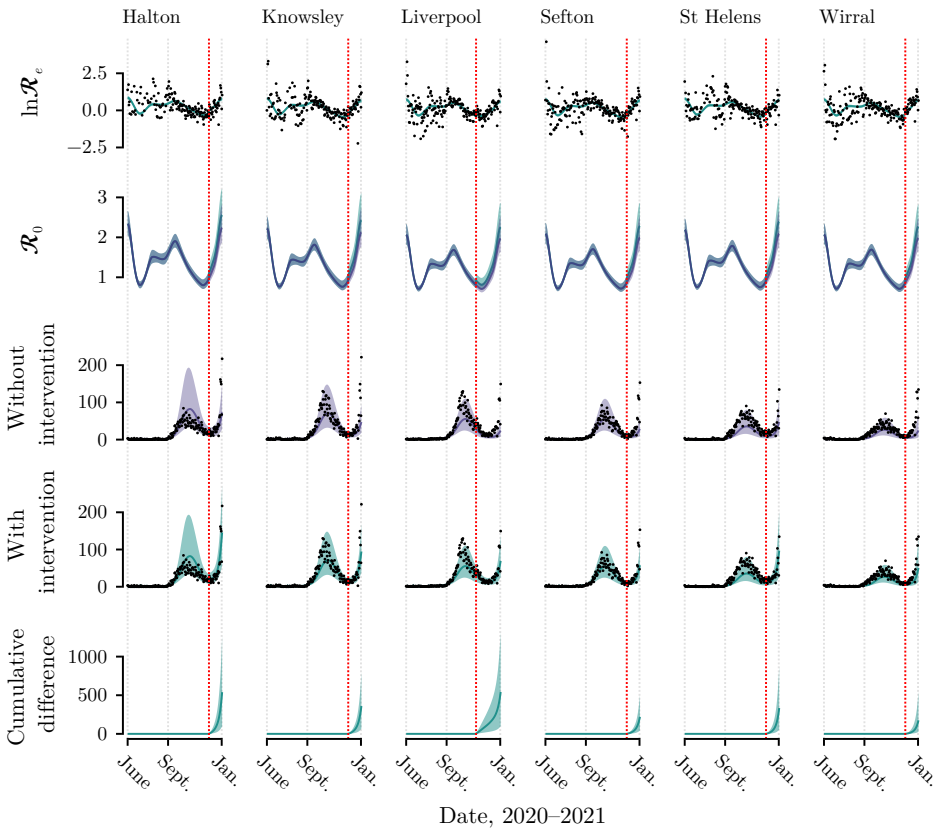

Figure 2.8: Difference-in-differences analysis for the effect of community testing on covid-19 diagnoses in North West England between June and December 2020. Dotted red lines indicate the start of the testing pilot. See caption to Figure 2.2 for description of colours.

3. $\hat{\tau}_{\text{ATT}3} = 1.12$ (90% CrI 0.99, 1.28) for recommendations to use face coverings.

2.6 Discussion

In this chapter I have presented a method combining the common-trends assumption from difference-in-differences methods with a mechanistic model of epidemic disease transmission. Previous difference-in-differences analyses of interventions to reduce infectious disease transmission have used

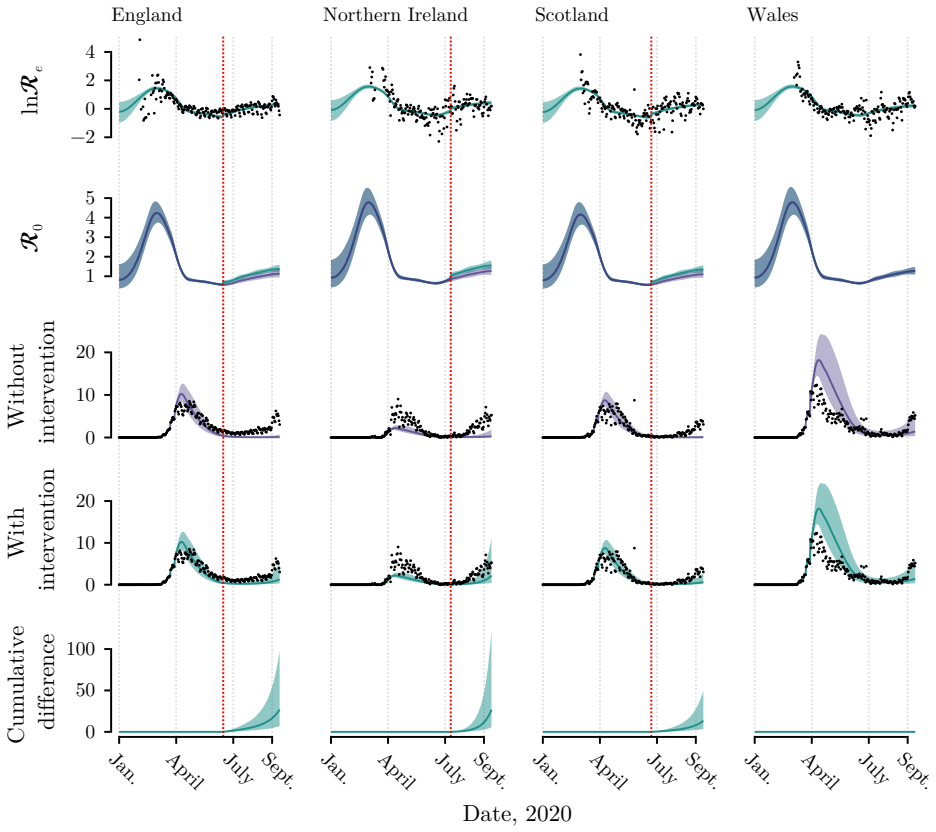

Figure 2.9: Difference-in-differences analysis for the effect of masking mandates on covid-19 diagnoses in the United Kingdom between January and September 2020. Dotted red lines indicate the start of the testing pilot. See caption to Figure 2.2 for description of colours.

incidence or simple transformations such as logarithms as the outcome (see Feng and Bilinski, 2024). These outcomes cannot follow a parallel trend during an outbreak unless groups have identical transmissibility and proportions susceptible and infectious. The use of a mechanistic model does not relax the need to meet other assumptions of difference-in-differences analyses, but through a series of simulations I showed that if certain conditions were met, the method could accurately infer the causal effect of interventions that reduce the basic reproduction ratio. I also showed that violations of the common-trends assumption can lead to inference of causal

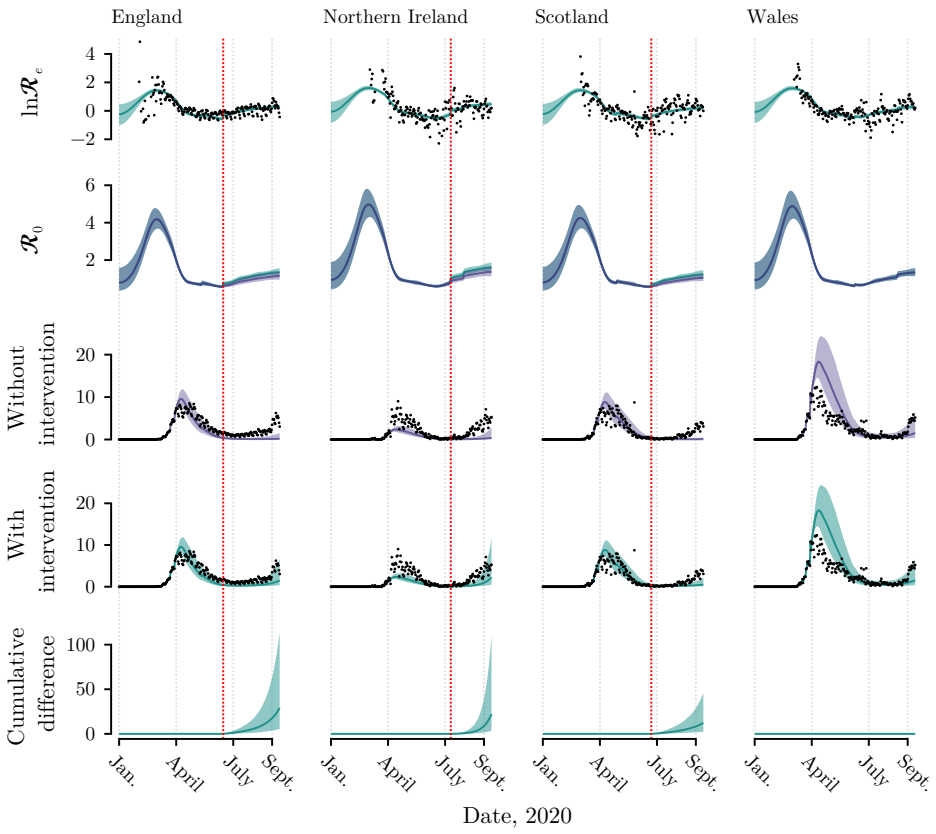

Figure 2.10: Difference-in-differences analysis for the effect of masking mandates on covid-19 diagnoses in the United Kingdom, with parameters fit for additional interventions of the end of stay-at-home rules, businesses reopening, and recommendations to use masks in advance of the mandates. Dotted red lines indicate the start of the testing pilot. See caption to Figure 2.2 for description of colours.

effects when the null hypothesis is true, and apparently null effects when causal effects are present.

Attempts to fit the model to interventions used during the covid-19 pandemic estimated that the intervention increased transmission, with a ten to twenty per cent increase in the basic reproduction ratio each time. The only intervention not associated with an inferred increase in transmission was the reopening of businesses, included as a secondary intervention in the analysis of the effect of face coverings.

For the main analysis of the effect of asymptomatic community testing,

I intended to use pillar 1 data. Reductions in severe disease sufficient to require testing in hospital, and reductions in infection among healthcare workers, are both objectives of public-health policy during an epidemic. That analysis suggested that SARS-CoV-2 would have been almost eliminated after a huge spike in infections in September 2020. Given that these results appeared implausible, I also presented an analysis using all recorded cases. Some infections in this analysis were only recognised because they were diagnosed by the testing programme. This violates the assumption of a constant proportion of cases detected. Even more importantly, in this analysis it is unclear whether recording more infections reflects a good result (better detection) or a poor one (more transmission). My finding of a 14 per cent increase in diagnoses (90% CrI 3%, 27%) is similar to University of Liverpool's (2021b) estimate that the testing programme led to an 18 per cent increase in detection of infections (95% confidence interval (CI) 7%, 29%).

It is possible that the interventions did cause an increase in transmission. A concern in community-level infection control is that adoption of a measure will lead to an inappropriate feeling of security, reductions in adherence to other infection-prevention measures, and a consequent increase in infections (Guglielmi, 2021). Saad-Roy and Traulsen (2023) combined a susceptible–infectious–resistant–susceptible (SIRS) model with a game-theory model of adherence to interventions, and found that new mandates might have no net effect on transmission when individuals down-titrated their use of other interventions to maintain a constant level of risk. With incomplete information, individuals might reduce adherence such that their risk increased.

Alternatively, the analysis may have been incorrect and the interventions did not cause an increase in transmission. For instance, Zhang *et al.* (2022) estimated that it led to a 43 per cent reduction in covid-related hospital

admissions (95% CI 29%, 57%). The analysis relies on many assumptions and we must be cautious in interpreting the output.

The central assumption of the difference-in-differences method is the common-trends assumption. This assumption requires an outcome that has a linear relationship with time-varying parameters. I showed that incidence and prevalence, and natural logarithms of these, do not have this property. The basic reproduction ratio could plausibly have this property, making it an appropriate outcome for difference-in-differences analyses, but we cannot be certain that this relationship is linear. A component of the parallel-trends assumption is the 'common-shocks' assumption (Ryan *et al.*, 2015, 2019). This requires that any events that affect the outcome have the same effect for all groups. This might not be plausible. For instance, relaxation of social-distancing regulations might have a greater effect on transmission in densely-populated urban districts than in rural areas. Similarly, the effect of school closures on transmission could depend on the age structure of each group's population. Despite applying to all groups, these interventions could violate the common-trends assumption. We also cannot be certain that we are correctly inferring the basic reproduction ratio. Erroneous assumptions about the duration of immunity or the proportion of infections reported would lead to errors in the estimate. Such violations would be demonstrated most clearly if the number of infections exceeded population size due to reinfection.

The common-trends assumption is untestable after the intervention is initiated (Rambachan and Roth, 2023; Roth *et al.*, 2023). In each simulated dataset in my analysis, the fitted parameters, which enforced a parallel trend in the basic reproduction ratio, followed the data well before and after interventions, even when the simulation was designed to violate the parallel-trends assumption. However, in the analysis with the pillar 1 data, and

in the plots for Northern Ireland and Wales before masking mandates, the inferred trend had a poor fit to the data.

This analysis assumes that once an intervention is in place, it has a consistent effect over time. The applicability of this assumption is a challenge to all difference-in-differences analyses which, like this one, include multiple time periods (Callaway and Sant'Anna, 2021). As the pandemic progressed, so-called 'pandemic fatigue' was observed, in which adherence to interventions decreased over time (MacIntyre *et al.*, 2021). Also, individual choices to use interventions depend on changes in perceived risk of infection (Saad-Roy and Traulsen, 2023). Therefore, effectiveness of providing or recommending an intervention may be inversely related to prevalence of infection.

The model assumes that all infections arise from infectious individuals within the group. This clearly cannot be the case for the first individual infected. Some number of imported cases is likely inevitable, even when restrictions on movement are in place. A Hawkes process is self exciting (Rizoiu *et al.*, 2018). From an epidemic-modelling perspective, a Hawkes process allows for new infections to arise from existing infectious individuals and also from importation of new cases. However, if we retain a multiplicative relationship between the intervention and the untreated counterfactual, addition of a constant import rate would lead to a non-parallel trend as the relative contribution of imported cases would be greater when transmission was reduced. The rate at which new infections are imported into an area may also depend on prevalence of neighbouring areas. In that case the stability assumption is violated. In the Liverpool City Region, many people in neighbouring districts may visit Liverpool City. Visitors could import cases into the city, or they could be infected during their visit and appear as imported cases for their home districts. In this case, even a Hawkes-process

model with a constant rate of imports would not reproduce true infection dynamics.

A strength of using the difference-in-differences method was that the basic reproduction ratio could be allowed to change constantly over time. Other regression-based estimates of changes in transmissibility assumed piecewise-constant reproduction ratios that changed only when infection-prevention measures were introduced or removed (Flaxman *et al.*, 2020; Liu *et al.*, 2024), or when a new variant of concern reached a threshold prevalence (Liu *et al.*, 2024). Piecewise-constant transmissibility may be violated if changes in behaviour occur between official changes in policy. If these analyses had allowed reproduction ratios to change independently between interventions then causal effects of interventions would not have been identifiable. By constraining groups to follow parallel trends between any interventions, the difference-in-differences analysis explicitly produced an untreated counterfactual scenario for groups that received interventions. This allowed the basic reproduction ratio to change in response to unmeasured influences (such as changes in behaviour or new variants of concern) without impinging on the ability to estimate interventions' causal effects.

An alternative to this analysis would be to discard the difference-in-differences method and simply fit parameters to a mechanistic model. The greatest challenge for that approach with these datasets would be that there do appear to be temporal changes in infectiousness, presumed to be due to changes in restrictions and spread of the Alpha variant. These changes would need to be incorporated into any model. This could be attempted by applying the location- and time-dependent form of the basic reproduction ratio that I have described, but such an approach would then be dependent on the same assumptions as this analysis.

An alternative to differences in differences used in econometric stud-

ies is the 'synthetic control' (Abadie *et al.*, 2010). Weightings are given to untreated groups' outcomes to create time series of untreated potential outcomes for the treated group. Such an analysis still assumes that time-dependent 'shocks' will affect treated and untreated groups equivalently. It also requires that we identify and measure relevant covariates to generate appropriate weights. Synthetic control methods have been used in evaluations of interventions to reduce infectious disease transmission. Notably, University of Liverpool (2021b) and Zhang *et al.* (2022) used synthetic controls to evaluate the effect of community testing in Liverpool. As with the difference-in-differences analyses described in the introduction to this chapter, these analyses use infections and hospital admissions as outcomes. Incorporation of mechanistic understanding of disease transmission into synthetic controls is an attractive future step for assessing effects of infection-prevention interventions. However, adding weights to each untreated group adds further complexity to the analysis. There is currently no consensus about how covariates should be chosen to generate weights, and decisions about variables to use for weighting may allow manipulation of results toward a preferred result (Ferman *et al.*, 2020). I therefore feel that further work would be needed to understand how difference-in-differences methods can apply to infectious-disease transmission before attempting to expand into synthetic control analyses.

2.7 Data and code

Hale *et al.*'s (2021) dataset is available at `https://github.com/OxCGRT/covid-policy-tracker/tree/master/data/United%20Kingdom`, using `OxCGRT_GBR_differentiated_withnotes_2020.csv`, shared under a Creative Commons Attribution 4.0 International Public Licence. Local-

authority-level data are available at `https://www.gov.uk/government/publications/coronavirus-cases-in-england-5-january-2021`, using 'North West epidemiological charts, data set', and `https://www.gov.uk/government/publications/coronavirus-cases-by-local-authority-epidemiological-data-9-june-2021`, using 'North West epidemiological charts by LTLA, data sets', both shared under the Open Government Licence v3.0. Cori *et al.*'s (2022) estimates of the serial interval for SARS-CoV-2 are available at `https://github.com/mrc-ide/EpiEstim/tree/master/data`, using `covid_deaths_2020_uk.rda`, shared under a GNU General Public License, version 2.

Code to reproduce the analysis in this chapter is available at `https://github.com/markgpritchard/RenewalDiffInDiff`.

Chapter Three

Effect of natural immune boosting on dynamics of infections with rapidly waning immunity[1]

Natural immune boosting is a hypothesis first proposed by Hope-Simpson (1965) in the context of varicella zoster virus. This virus causes chickenpox (varicella) and shingles (zoster). In unvaccinated populations, chickenpox is a common childhood illness. It usually causes a mild, self-limiting illness, but can cause more severe disease, especially if contracted as an adult (Gershon *et al.*, 2015). Following clinical recovery from chickenpox, the virus lies dormant in sensory nerve ganglia. Re-infection is rare, but the dormant virus can reactivate to cause shingles, a painful rash that follows the dermatome of the affected nerve root (Hope-Simpson, 1965). Hope-Simpson's (*ibid.*) hypothesis was that re-exposure to the virus during adulthood could re-potentiate the immune response and prevent onset of shingles. This hypothesis contributed to the development of the first shingles vaccine, a live attenuated vaccine that reduced shingles incidence in adults over sixty years

[1]Results in this chapter were partly presented at the 9th International Conference on Infectious Disease Dynamics, M. G. Pritchard, C. A. Donnelly, P. W. Horby, and B. S. Cooper, Natural immune boosting leads to synchrony in immunity and cyclic outbreaks of disease for pathogens with rapidly waning immunity (2023, abstract available at `https://ssrn.com/abstract=4654932`); and in the pre-print, M. G. Pritchard, S. M. Cavany, S. J. Dunachie, G. F. Medley, L. Turtle, C. A. Donnelly, P. W. Horby, and B. S. Cooper, Natural immune boosting can cause synchrony in susceptibility and outbreaks of respiratory infections with rapidly waning immunity, *medRxiv* 2023.11.23.23298952 (2023).

by 51 per cent (Oxman *et al.*, 2005).[2]

Similar immune patterns are seen in *Plasmodium*/malaria epidemiology. In areas with a consistently high burden of malaria, severe disease is most often seen in childhood or in people who have arrived from an area without a high prevalence. People born in high-prevalence areas who have moved away and returned also tend to have a high risk of reinfection (Struik and Riley, 2004). In contrast, in areas where malaria is seasonal, incidence is more evenly distributed across ages (Carneiro *et al.*, 2010).

Immune boosting effects have also been seen for *B. pertussis* (Clark, 2014). Immunity induced by measles and mumps vaccines is also boosted by subclinical exposure (Whittle *et al.*, 1999). Vaccine-induced protection against measles and mumps infections has reduced over time as prevalence of the viruses has decreased, although the risk of severe disease remains much lower than it was in unvaccinated populations (Yang *et al.*, 2020).

De Angelis *et al.* (2021) found that repeated exposure to SARS-CoV-2 might promote immunity and help to reduce the risk of severe disease. Immune boosting could potentially be very important for viral infections of the respiratory tract. These infections cause very brief immunity (Slifka and Ahmed, 1996). Even among people who receive an annual influenza vaccination, immunity can wane during the influenza season leaving people who were vaccinated early in the season vulnerable to infection by the season's end (Domnich *et al.*, 2024; Rambhia and Rambhia, 2019). If natural boosting prolonged immunity, this could have implications for optimal vaccination timing.

In this chapter, I summarize previous mathematical models of natural immune boosting. I then use an ODE model to investigate the effect of boost-

[2]Since September 2023, the UK shingles vaccination programme uses a recombinant varicella zoster glycoprotein E vaccine with efficacy against shingles of 97 per cent (Lal *et al.*, 2015; UK Health Security Agency, 2024b)

ing on disease dynamics for respiratory viruses, using RSV incidence before and during periods when physical distancing measures were in place to reduce the spread of covid-19 as an example.

3.1 Background and data

Respiratory syncytial virus

RSV infection elicits remarkably brief immunity (Slifka and Ahmed, 1996). Reinfection is possible within two months (Hall *et al.*, 1991). Almost all children are infected in the first two years of life (Glezen *et al.*, 1986), then four to ten per cent of adults are reinfected each year (Falsey *et al.*, 2005). It is a seasonal infection, with different regions experiencing peaks twice annually, annually or biennially (Bloom-Feshbach *et al.*, 2013). In the UK, infections typically peak in December each year (UK Health Security Agency, 2021). In 2020 and 2021, non-pharmaceutical interventions to reduce the spread of SARS-CoV-2 disrupted RSV transmission, and the typical annual outbreaks were not seen in 2020 (Friedrich *et al.*, 2021; van Summeren *et al.*, 2021; Yeoh *et al.*, 2021). Unusually early and large outbreaks of RSV occurred in the following two years, attributed to increased susceptibility in the population (Bardsley *et al.*, 2023; Foley *et al.*, 2021; Nygaard *et al.*, 2023).

RSV is not a notifiable pathogen or a cause of a notifiable disease in Scotland (*Public Health etc. (Scotland) Act* 2008). In infants it causes bronchiolitis, which is generally a clinical diagnosis (Dalziel *et al.*, 2022). After infancy, symptoms are non-specific and RSV infections are often undiagnosed (Rozenbaum *et al.*, 2023). In Scotland, weekly numbers of RSV diagnoses confirmed by microbiology laboratories are made publicly available by Public Health Scotland (2023).

Governmental response to the covid-19 pandemic

Hale *et al.* (2021) recorded details of recommendations and mandates to reduce the spread of covid-19. Across a range of policy measures, they allocated zero points when a measure was not in place and greater numbers of points when measures were recommended or mandated. The stringency index is a composite measure, calculated as the points allocated across nine measures compared to the maximum possible points available for those measures, giving a score between 0 (none of these measures recommended) and 100. I used a threshold of 50 to represent the period with a high levels of restrictions. In Scotland, this lasted from 21 March 2020 to 8 August 2021, and corresponded with notable step-changes in increasing and decreasing stringency index (see Figure 1.3, page 17).

3.2 Mathematical modelling of natural immune boosting

The susceptible–infectious–resistant–susceptible compartmental model

Waning immunity can be represented in a compartmental model by allowing movement from the resistant compartment to the susceptible one, as in the SIRS model. An ODE form of this model is,

$$\left.\begin{aligned} \frac{\mathrm{d}s(t)}{\mathrm{d}t} &= \mu\,[n - s(t)] + \omega\, r(t) - \beta\,\frac{s(t)\,i(t)}{n}, \\ \frac{\mathrm{d}i(t)}{\mathrm{d}t} &= \beta\,\frac{s(t)\,i(t)}{n} - [\gamma + \mu]\,i(t), \\ r(t) &= n - [s(t) + i(t)]\,, \end{aligned}\right\} \tag{3.1}$$

where $s(t)$, $i(t)$, $r(t)$ are respectively numbers susceptible, infectious, resistant (immune) to infection at time t, $n := s(t) + i(t) + r(t)$, β is a transmission coefficient, γ is the recovery rate, μ is both the birth and mortality rate, and ω is the rate of immune waning. This gives an exponential distribution to the duration of immunity with a mean duration $1/\omega$.

Models of immune boosting for varicella zoster

Models of immune boosting for varicella zoster include Korostil *et al.*'s (2015) ODE model with a waned compartment from which individuals could be reinfected, immune boosted (returned to full immunity) or could develop shingles. Bakker *et al.* (2022) used a discrete-time model. They assumed that boosting reduced the current risk of shingles, but in their model it did not reduce future risks.

Three models of immune boosting have particularly influenced policy regarding varicella vaccination in the UK. Brisson *et al.* (2010) and van Hoek *et al.* (2011) used ODE models to model immune boosting of varicella zoster. Recovered and vaccinated individuals could have their immunity boosted in proportion to the force of infection. This informed a cost–effectiveness analysis that concluded that benefits of varicella vaccination would be offset by increases in shingles for the first thirty to fifty years after starting a vaccination programme (van Hoek *et al.*, 2012). An updated model (currently unpublished, discussed in Joint Committee on Vaccination and Immunisation, 2023) assumed that boosting provided three years of protection against shingles, in contrast to the twenty years' protection assumed by Brisson *et al.* (2010) and van Hoek *et al.* (2011). This led to a conclusion that varicella vaccination is cost-effective, and the Joint Committee on Vaccination and Immunisation (2023) recommended that varicella vaccination should be added to the childhood vaccination schedule.

Immune boosting for other microparasitic infections

To represent natural immune boosting for other microparasitic infections that do not have a post-recovery dormant phase, we might imagine different forms of boosting. The boosting effect could simply be a result of reinfection that causes no symptoms or is sufficiently mild that people do not seek

medical attention. Modellers might assume that people experiencing such reinfection might be less infectious than those being infected for the first time. Alternatively, we might consider boosting as a process that returns an individual to full immunity without causing any secondary infectiousness.

Some models label a process as 'immune boosting' but allow immune-naïve susceptible individuals to undergo it. Stout *et al.* (2020) describe colonization with *Listeria monocytogenes* as a form of immune boosting. In their model, immune and susceptible individuals undergo such colonization. Alexander *et al.* (2021) describe boosting against dengue virus in terms of immunity generated by a low level of exposure that is insufficient to cause clinical disease. This could be experienced by susceptible individuals who would have been infected if exposed to a larger exposure. Owusu-Ansah *et al.* (2019) included immune boosting to norovirus, moving individuals between partially- and fully-immune compartments, but the mechanism of boosting in their model was unclear.

Mossong *et al.* (1999) produced an ODE model of immune boosting for measles. They discuss the possibility that immune waning without natural immune boosting could lead to a return to susceptibility and reinfection at older ages, but they did not explicitly model immune boosting. Glass and Grenfell (2003) created an ODE model for immunity to measles virus. They represented population-level changes in antibody levels in a separate set of differential equations which gave the first model a time-varying proportion of those vaccinated or recovered who were vulnerable to infection.

Heffernan and Keeling (2009) explicitly included immune boosting and waning in a single set of model equations. Their ODE model of measles immunity had an unlimited number of discrete levels of immunity. People at the lowest level were completely susceptible. Those in the next fifteen levels would become less infectious if infected than those who were completely

susceptible. In the remaining levels of immunity, individuals could not become infectious until they had first waned to a lower level of immunity. Barbarossa and Röst (2015a,b) created a pathogen-agnostic partial differential equations model with immunity on a continuous scale. They also created a delay differential equation model with the duration from recovery to susceptibility affected by boosting (Barbarossa *et al.*, 2017).

Lavine *et al.* (2011) modelled immune boosting in *B. pertussis* by adding a waned compartment to an SIRS model. Immune individuals enter the waned compartment at a rate 2ω. From there they can return to the fully-immune compartment at a rate proportional to the force of infection, or they become fully susceptible at a rate 2ω. Without immune boosting, the expected duration of immunity is $1/\omega$, as in the SIRS model. Individuals being boosted from the waned compartment do not cause secondary infections. Further analysis of the model for *B. pertussis* has been conducted by Dafilis *et al.* (2012, 2014a,b) and Lavine *et al.* (2013). Leung *et al.* (2018) added compartments for vaccinated individuals, with vaccine-induced immunity also susceptible to boosting. Opoku-Sarkodie *et al.* (2022) modified Lavine *et al.*'s (2011) model so that the rate of movement into the waned compartment could differ from the rate of waning into the susceptible compartment.

Wearing and Rohani (2009) describe a model with a second set of susceptible, exposed and infectious compartments representing people who have previously been infected with *B. pertussis* then had their immunity wane. A similar model is described by Blackwood *et al.* (2013). This model allows reinfected individuals to cause fewer secondary infections than those who are infected for the first time. Although not described in terms of immune boosting, a similar model is commonly used to represent repeated RSV infection (Lang, 2022).

Hoyer-Leitzel *et al.* (2023) use a within-host flow–kick model (see Meyer

et al., 2018) in which dynamics generally evolve according to ODEs but are periodically perturbed (kicked) by an added viral load, representing an infection event. These kicks could either boost immunity or cause reinfection, depending on the size of the exposure and the degree of immune waning.

Heffernan and Keeling (2009) noted that in simulations with high levels of vaccination, when an outbreak ended a large proportion of the population was immune. Prevalence was then low while immunity waned. This could then be followed by further outbreaks each time the susceptible fraction became large enough. They found that these cyclical outbreaks could only be maintained by a large number of immunity levels giving a gamma-distributed immune duration. If immunity instead had an exponential distribution then the disease dynamics converged on a stable equilibrium.

Since that result, models without large numbers of immune levels have found equilibria to be unstable when vaccination coverage is high (Lavine *et al.*, 2011, 2013; Leung *et al.*, 2018), and also when the boosting coefficient is large (Barbarossa *et al.*, 2017; Korostil *et al.*, 2015; Lavine *et al.*, 2011, 2013) or life expectancy is high (Barbarossa *et al.*, 2017; Dafilis *et al.*, 2012, 2014b). Dafilis *et al.* (2014b) and Korostil *et al.* (2015) found cyclic epidemics when immune duration was short, and Barbarossa *et al.* (2017) and Leung *et al.* (2018) when immune duration was long.

Except for Hoyer-Leitzel *et al.*'s (2023) within-host model, all models of natural immune boosting have been for microparasites with immune durations typically measured in decades. Here I explore dynamics for microparasites eliciting immunity of short duration, such as viral infections of the respiratory tract.

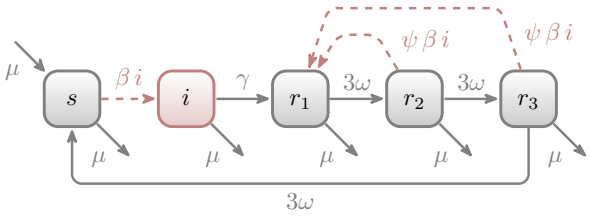

Figure 3.1: Model flow chart. Compartments are, s, susceptible; i, infectious; r_1 to r_3, resistant (immune). Compartments that contribute to the force of infection are shaded in red; events that occur in proportional to the force of infection are represented by dashed red lines. See text for parameter definitions.

3.3 Description of the model

I assume SIRS-like dynamics with two additional compartments for individuals with partially-waned immunity (Figure 3.1). The model has a constant population of 1 so $s(t)$, $i(t)$ and $r(t)$ are respectively proportions susceptible, infectious and resistant. All newborns are susceptible. A series of three immune compartments, $r(t) := r_1(t) + r_2(t) + r_3(t)$, represent waning immunity. The probability of immunity being boosted is proportional to the force of infection. This is represented by a boosting coefficient, labelled ψ, describing the relationship of boosting to the force of infection, such that $\psi = 0$ represents no immune boosting and $\psi > 1$ implies that boosting occurs more readily among those with waning immunity than infection does among those susceptible. The model is described by ODEs,

$$\left.\begin{aligned}
\frac{\mathrm{d}s(t)}{\mathrm{d}t} &= 3\omega\, r_3(t) - \beta\, s(t)\, i(t) + \mu\,[1 - s(t)]\,,\\
\frac{\mathrm{d}i(t)}{\mathrm{d}t} &= \beta\, s(t)\, i(t) - [\gamma + \mu]\, i(t),\\
\frac{\mathrm{d}r_1(t)}{\mathrm{d}t} &= \gamma\, i(t) + \beta\, \psi\, i(t)\,[r_2(t) + r_3(t)] - [3\omega + \mu]\, r_1(t),\\
\frac{\mathrm{d}r_2(t)}{\mathrm{d}t} &= 3\omega\, r_1(t) - [3\omega + \beta\, \psi\, i(t) + \mu]\, r_2(t),\\
r_3(t) &= 1 - [s(t) + i(t) + r_1(t) + r_2(t)]\,.
\end{aligned}\right\} \tag{3.2}$$

All parameters are non-negative.

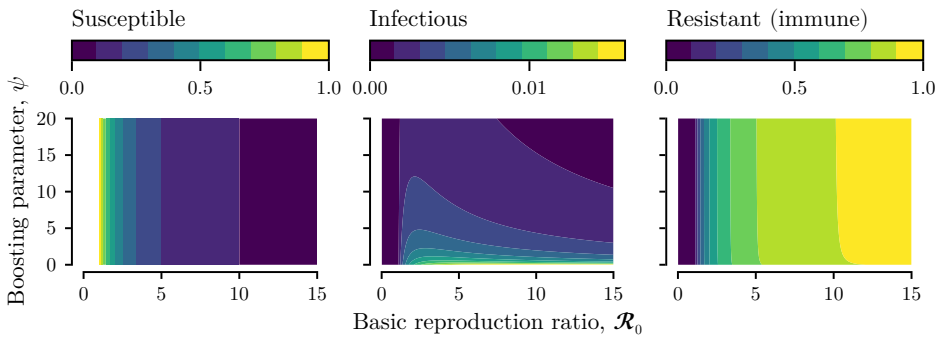

Figure 3.2: Proportions in each compartment at equilibrium. Note that the maximum proportion infectious is much less than proportions susceptible or resistant, and that plot's colour scale is different.

3.4 Analysis

The basic reproduction ratio, $\mathcal{R}_0 = \beta/(\gamma + \mu)$. To find equilibria, set each ODE in 3.2 to equal 0. This is solved by disease-free conditions, $s_* = 1$, $i_* = r_{1*} = r_{2*} = r_{3*} = 0$. Eigenvalues at the disease-free equilibrium are $\Lambda_{\text{DFE}} = \{-\mu,\ \beta - \gamma - \mu, -3\omega - \mu\}$ so the disease-free equilibrium is stable if and only if $\beta < \mu + \gamma$, meaning $\mathcal{R}_0 < 1$.

Endemic equilibrium values are given by,

$$s^* = [\gamma + \mu]\,/\beta, \tag{3.3a}$$

$$i^* = 1 - [s^* + r_1^* + r_2^* + r_3^*]\,, \tag{3.3b}$$

$$r_j^* = \left[\frac{[i^* + \mu/\beta]\,[\gamma + \mu] - \mu}{3\omega}\right]\left[1 + \frac{\beta\,\psi\,i^* + \mu}{3\omega}\right]^{3-j}, \quad j \in \{1,2,3\}. \tag{3.3c}$$

The endemic equilibrium proportion susceptible is $1/\mathcal{R}_0$ (with $\mathcal{R}_0 > 1$), so is independent of waning and boosting parameters (Figure 3.2). Without immune boosting, the model follows SIRS dynamics, with the serial immune compartments giving an Erlang-distributed duration of immunity. In this case, the equilibrium proportion infectious increases monotonically with greater values of the basic reproduction ratio. With immune boosting, $\psi > 0$, there is a turning point beyond which greater infectiousness is associ-

ated with smaller equilibrium proportions infectious and larger proportions immune.

I calculated eigenvalues for endemic equilibria numerically in Julia (version 1.9.3; Bezanson *et al.*, 2017). With no immune boosting, the stationary points are stable. Greater values of the boosting coefficient lead to a Hopf bifurcation beyond which equilibria are unstable. I denote the threshold boosting coefficient value that separates stable and unstable equilibria as ψ^*. For a pathogen eliciting immunity that would last approximately two years without boosting, I found low values of ψ^*, suggesting that boosting at a similar level of exposure as required for infection could lead to instability (Figure 3.3). The magnitude of ψ^* increases rapidly as the unboosted immune duration approaches zero, meaning that as the model approaches SIS dynamics, equilibria are stable even if immune boosting requires much less exposure than infection. For large values of the basic reproduction value and long durations of immunity, equilibria are stable for any value of the boosting coefficient.

3.5 Simulations and seasonal forcing

I performed numeric simulations using Verner's (2010) 9th order Runge–Kutta method in `DifferentialEquations.jl`, version 7.13.0 (Rackauckas and Nie, 2017). I set $\mu = 0.0087$, approximating Scotland's annual birth rate (Scottish Public Health Observatory, 2023), and assumed a mean generation time of 7.5 days. When the boosting coefficient was greater than ψ^*, simulated compartment sizes did not converge on the calculated equilibrium values. Instead they approached a limit cycle, oscillating between a larger and a smaller proportion infectious (Figure 3.4). I estimated discrete Fourier transforms using the fast Fourier transform algorithm in `FFTW.jl`

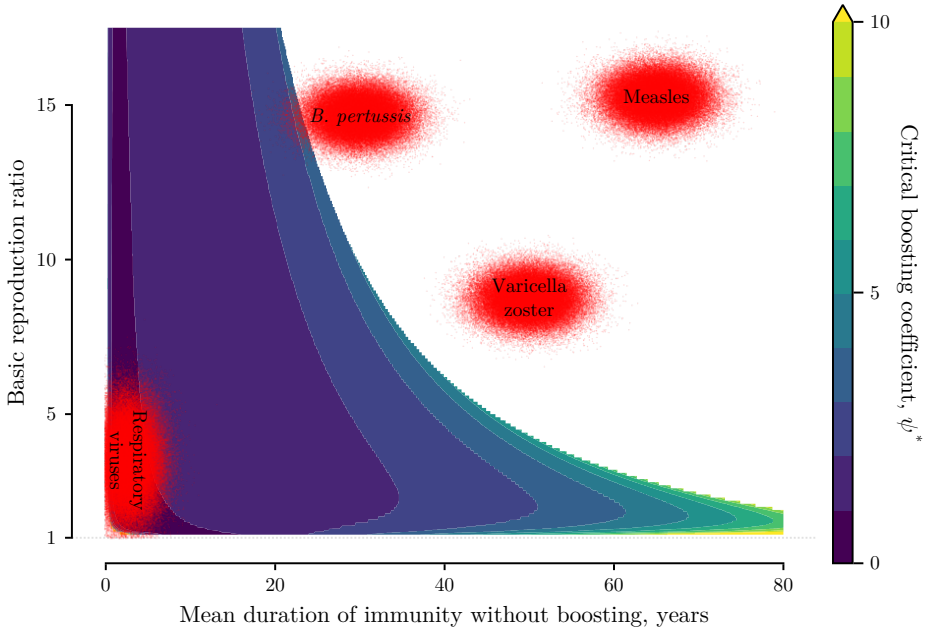

Figure 3.3: Critical values of the boosting coefficient threshold, ψ^*, between stable and unstable equilibria. Values calculated assuming $1/\gamma = 7.5$ days, $\mu = 0.0087$/year. Areas where $\psi^* > 10$ are shown in yellow. Areas where there is no value of ψ^* (equilibria are always stable) are shown in white. Red shaded areas show approximate parameter regions for microparasites that have commonly been included in models of natural immune boosting, and respiratory viruses (based on data from Anderson and May, 1982; Slifka and Ahmed, 1996; Wearing and Rohani, 2009).

(version 1.8.0; Frigo and Johnson, 2005) to examine the relationship between the boosting coefficient and the frequency of simulated outbreaks. The interval between outbreaks increased (reduced frequency) with greater boosting coefficients (Figure 3.5*A*).

To compare cyclical effects of natural immune boosting to any seasonal effects on transmissibility, I allow infectiousness to vary sinusoidally, substituting β in equation 3.2 with,

$$\beta(t) = \beta_0 \left[1 + \beta_1 \cos(2\pi t - \varphi)\right], \tag{3.4}$$

where β_0 is the mean transmission parameter, $\beta_1 \in [0, 1]$ the proportional amplitude of seasonal forcing, $\varphi \in [0, 2\pi)$ a phase offset, and t is time meas-

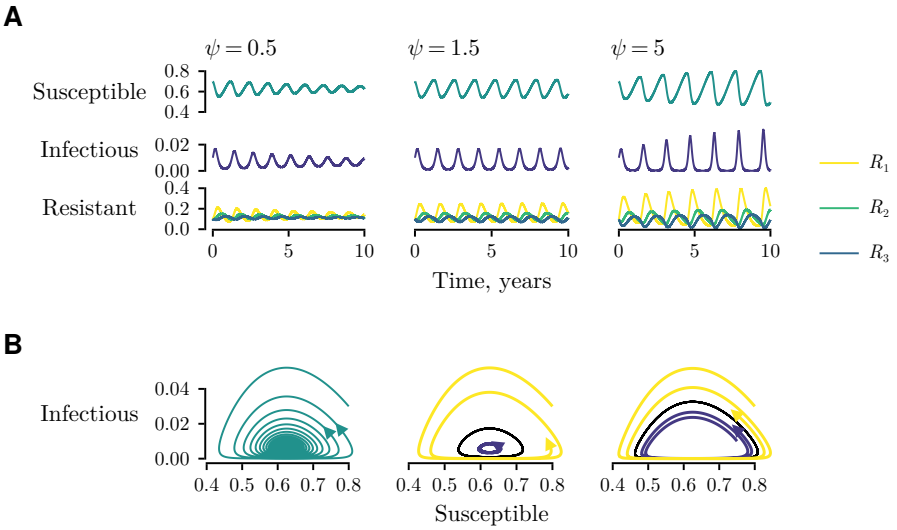

Figure 3.4: Simulations with different boosting coefficients. *A* Simulated proportions in each compartment, in simulations with a basic reproduction ratio, $\mathcal{R}_0 = 1.6$, and a critical boosting coefficient threshold, $\psi^* = 1.18$. Recovered subcompartments R_1 to R_3 are entered in order, with R_1 representing those with the greatest expected time before returning to susceptibility. *B* Infectious–susceptible planes for simulations with the same parameters as above. Black loops indicate limit cycles. In each plot, birth rate is 0.087 per year, mean generation time is 7.5 days, and mean immune duration without boosting is 400 days.

ured in years. I generated qualitatively indistinguishable patterns of incidence in simulations with no seasonal forcing and a large boosting coefficient, and in those with seasonal forcing and no immune boosting (see the period before simulated intervention in Figure 3.6). With intermediate values, if the frequency of seasonal forcing was close to the resonant frequency of the unforced system then stable annual outbreaks were seen (strong signals correlated with annual or biennial outbreaks for a range of boosting coefficients in Figure 3.5*B*). Simulations with more immune boosting required greater mean transmissibility to generate equivalent incidence patterns. Greater immune boosting led to greater proportions in the first resistant compartment during each outbreak. This was followed by a rapid, synchronized return to susceptibility between outbreaks. The proportion susceptible was smaller in

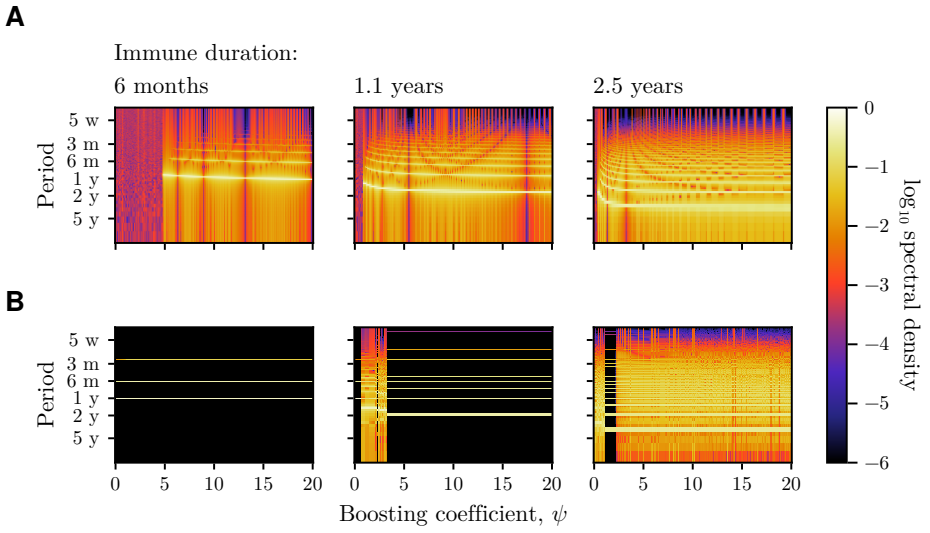

Figure 3.5: Periodicity of the model. *A* Discrete Fourier transform of simulations with constant parameters and *B* with seasonally forced parameters (amplitude 10% of mean transmission parameter). Bright areas show the period of outbreaks. In each plot, $\mathcal{R}_0 = 1.6$, birth rate is 0.087 per year, and mean generation time is 7.5 days.

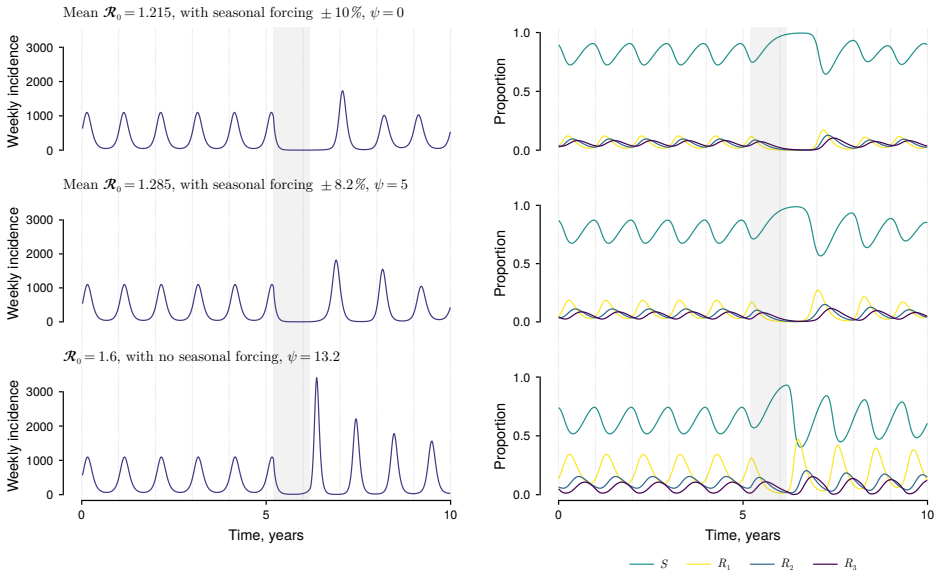

Figure 3.6: Simulated interventions. Left-hand plots show incidence for a series of simulations. Right-hand plots show proportions in the susceptible and resistant compartments for each simulation. The shaded grey area represents a period when the transmission coefficient is reduced for a non-pharmaceutical intervention.

simulations with more immune boosting, which is an expected consequence of a larger basic reproduction ratio.

I simulated adoption of non-pharmaceutical interventions by transiently reducing the transmission parameter by twenty per cent (shaded grey area in Figure 3.6). In all simulations, non-pharmaceutical interventions led to a reduction in incidence and an increase in the proportion susceptible. Once non-pharmaceutical interventions were lifted, simulations with the greatest boosting coefficient (and the greatest basic reproduction ratio) exhibited the largest and soonest rebound outbreaks.

3.6 Fitting model parameters

Using Castro and de Boer's (2020) scaling method I found that model parameters with seasonal forcing were not uniquely identifiable from prevalence data unless the immune waning rate, ω, was known (see supplementary text A.3). I therefore explored a range of scenarios between $\omega = 0.1/\text{year}$ (assuming a mean duration of immunity of ten years without boosting), and $\omega = 6/\text{year}$ (mean immunity of two months). I fit two additional parameters to represent the proportional reduction in transmission when the stringency index exceeded 50 and when it returned to less than 50, and a parameter to represent the proportion of all infections that are reported,

$$\beta(t) = \beta'(t)\,\beta_0\,[1 + \beta_1\,\cos(2\pi t - \varphi)]\,, \tag{3.5}$$

$$\beta'(t) = \begin{cases} 1, & t < 21 \text{ March } 2020, \\ \beta'_1, & 21 \text{ March } 2020 \leq t < 8 \text{ August } 2021, \\ \beta'_2, & 8 \text{ August } 2021 \leq t, \end{cases} \tag{3.6}$$

where β'_1 and β'_2 are coefficients to identify.

I used the parallel-tempered MCMC method (Syed *et al.*, 2022) implemented in `Pigeons.jl`, (version 0.4.1; Surjanovic *et al.*, 2023). I generated five chains of fitted parameters for each value of ω. Chains that failed to generate at least 512 samples or that showed a log likelihood consistently below other chains' results were classified as failing and were discarded. I generated a simulated weekly incidence for each sampled parameter set, and calculated the median and central 90% CrI for the incidence from these simulations.

3.7 Results for respiratory syncytial virus in Scotland

Between 2016 and 2019, RSV diagnoses in Scotland peaked each year between November and January (Figure 3.7*A*). There was a secular trend of increasing numbers of diagnoses each year, from 3100 to almost 5000 (as outbreaks span calendar years, I count each year's cases from 1 April onwards). In the twelve months from April 2020, 61 RSV infections were reported. In 2021, case numbers started to increase in August, leading to an earlier peak in RSV diagnoses but a total number similar to pre-pandemic years (3581 cases). In 2022, reports started in June and the total number of cases reported was 7331. Results stratified by age group (Figure 3.7*B*) showed greater numbers of diagnoses in 2022 for all age groups, suggesting that reinfection was an important contributor to the size of that year's outbreak.

Prior predictive plots showed greater predicted incidence with greater immune waning rates (Figure B.14). Simulations with immune waning rates up to 0.2/year predicted a very low incidence. Chains for fitting parameters with low waning rates did not mix as well as those with greater waning rates (Figures B.15 to B.21). Despite this, each simulation with fitted parameters mostly gave qualitatively good fits to the data (Figure 3.8*A*). In

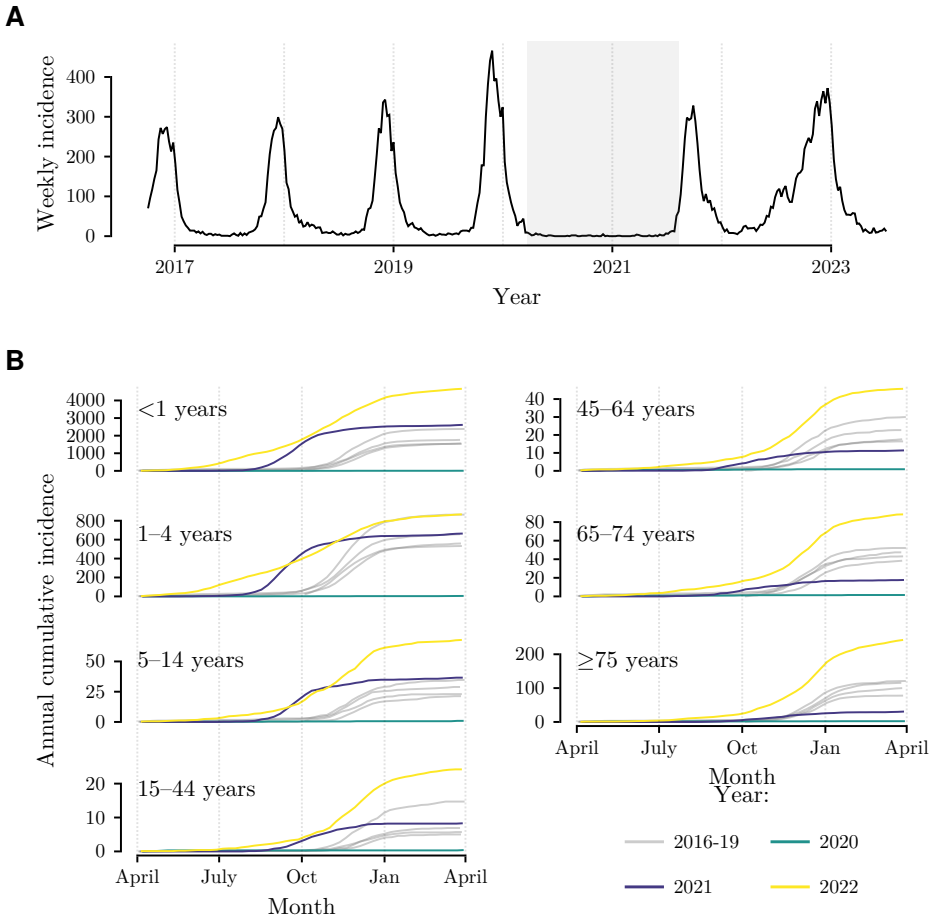

Figure 3.7: Incidence of laboratory-confirmed respiratory syncytial virus infection in Scotland. *A* Weekly incidence. The grey shaded area represents the period of high stringency index (>50) in Scotland. *B* Cumulative incidence (from April each year) stratified by age.

terms of model likelihood, parameters for $\omega = 0.1$/year had the best fit, but these chains were also most likely to fail, apparently because the prior probability of such extreme basic reproduction ratios was very low. These parameters gave a median basic reproduction ratio of 27 (90% CrI 24, 33). Models with more rapid immune waning had progressively lower median basic reproduction ratio estimates. The simulation with $\omega = 1$/year had the greatest proportional magnitude of seasonal forcing in transmissibility, $\beta_1 = 0.178$ (90% CrI 0.175, 0.182), and the lowest immune boosting para-

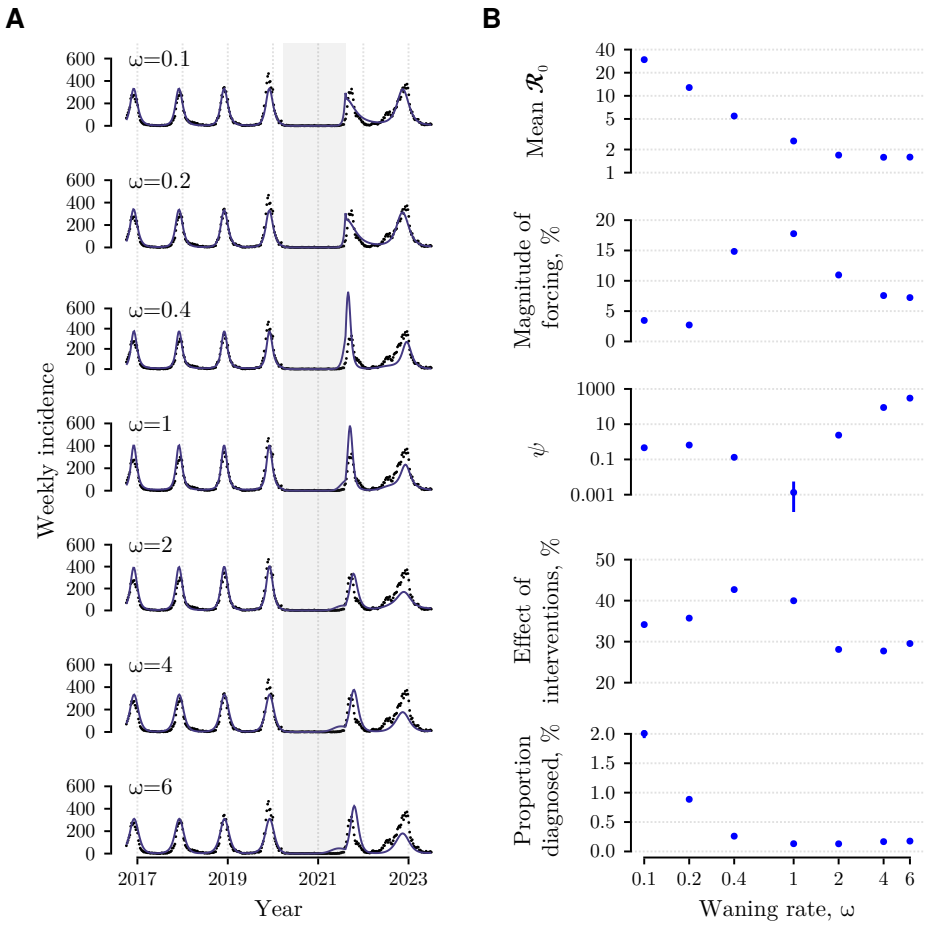

Figure 3.8: Fitted parameters and model outputs for respiratory syncytial virus transmission in Scotland. *A* Purple lines show simulations using fitted parameters for a range of waning rates (ω). The grey shaded area represents the period of high stringency index (>50) in Scotland. Black dots show recorded incidence. *B* Median and central 90% credible intervals of fitted parameters for a range of waning rates. $\mathcal{R}_0$, basic reproduction ratio; ψ, boosting coefficient.

meter, $\psi = 0.0013$ (90% CrI 0.0001, 0.0057). The greatest boosting parameter was for the simulation with $\omega = 10$/year, $\psi = 820$ (90% CrI 790, 850). All simulations with $\omega \geq 2$/year had a fitted boosting parameter greater than 1.

The fitted effect of the greater stringency index varied with rates of immune waning. For a waning rate of $\omega = 0.4$/year, the estimated effect of

non-pharmaceutical interventions was a 43 per cent (90% CrI 42, 44%) reduction in transmission. If $\omega = 2/\text{year}$, the estimated effect was a 26 per cent (90% CrI 26, 29%) reduction.

3.8 Discussion

These results suggest that natural immune boosting to viral respiratory-tract infections with rapidly waning immunity could lead to synchrony in population immunity and susceptibility. With sufficient boosting, this could cause periodic peaks of infection even with no seasonal changes in transmissibility. At lower levels of immune boosting, these inherent oscillations could resonate with, and amplify, effects of seasonal changes.

I demonstrated that observed patterns of RSV infections, before and after restrictions for SARS-CoV-2, could be consistent with a very wide range of boosting coefficients. Modelled scenarios included possibly implausible parameters, such as those requiring a basic reproduction ratio greater than 20 or a boosting coefficient of 800. Nevertheless, within a plausible range of parameters, significant variation in levels of natural immune boosting were possible. The relationship between these parameters showed that mischaracterization of immune boosting could lead to errors in estimates of other parameters. If the effect of immune boosting were assumed to be lower than it is, we could underestimate the pathogen's basic reproduction ratio, and underestimate the magnitude of reduction in transmission required to control spread.

Synchronous and oscillatory dynamics are widely recognised in epidemiology (Earn *et al.*, 1998), and more broadly across biological processes (Strogatz, 2018). Even when disease dynamics tend toward stability, if periodic environmental effects resonate with the system's intrinsic oscillations,

this can lead to large seasonal variation in the incidence of infection (Dushoff *et al.*, 2004). In systems without immune boosting, more frequent outbreaks are predicted when susceptible numbers are rapidly replenished by high birth rates (Earn *et al.*, 2000; Grenfell *et al.*, 2001) or short durations of immunity (Dushoff *et al.*, 2004; Wilson and Worcester, 1945). If these rates are non-constant then outbreaks may synchronize with the susceptible replenishment. For instance, seasonality of animal birth rates contributes to oscillatory dynamics of their diseases (Hosseini *et al.*, 2004). In my model, the effective rate of immune waning is inversely related to the prevalence of infection (boosting prolongs immunity). Immune boosting reduces the return to susceptibility during outbreaks. These outbreaks are followed by more rapid immune waning, and synchrony in the return to susceptibility, priming the population for another large outbreak. Dafilis *et al.* (2012) found that increasing life expectancy in a model of boosting immunity to *B. pertussis* led to oscillatory dynamics. In their model, birth and mortality rates were fixed to be equal so increasing life expectancy equivalently reduced the birth rate. This increased the relative contribution of immune waning to the size of the susceptible compartment, in a similar way to increasing the waning rate (reducing the duration of immunity) does in my model.

Oscillatory behaviour in disease dynamics may be desirable if it is sufficiently extreme to cause local elimination of a microparasite (Earn *et al.*, 1998). However, for viruses such as RSV with rapidly waning immunity and globally high prevalence, reintroduction is expected to occur so readily that local elimination is impossible (Agoti *et al.*, 2015). Rather than benefiting from oscillations, large peaks in numbers infected may overwhelm healthcare systems, and public health authorities may instead need to direct efforts to 'flattening the curve'. This motivated some interventions during the covid-19 pandemic (Kenyon, 2020). The potential for an outbreak

driven by the return to susceptibility suggests a possible strategy for repeat vaccination to be timed with the nadir of the force of infection. Explicit consideration of the balance between the effects of immune waning and seasonal forcing would be needed to determine optimal vaccination strategies (vaccination strategies in the presence of immune boosting are explored in Chapter 4).

Considering RSV specifically, two monoclonal antibodies and two vaccines have recently been licensed in the UK (UK Health Security Agency, 2024a). The global need to reduce the burden of RSV infection is clear. It caused an estimated 100 000 childhood deaths in 2019 (Li *et al.*, 2022) and causes up to a quarter of excess winter deaths among adults (Thompson *et al.*, 2003). Modelling suggests that both the monoclonal antibodies and vaccination would be cost effective (Hodgson *et al.*, 2020, 2022). These models did not include natural immune boosting. If a widespread intervention among infants reduced exposure to the virus among adults, this could potentially cause a paradoxical increase in infection due to reduced immune boosting, similar to that reported for measles and mumps (Yang *et al.*, 2020). RSV tends to be under-diagnosed among adults, including those with severe disease, due to its non-specific symptoms (Zheng *et al.*, 2022). Since September 2024, vaccination against RSV is routinely offered to pregnant women and adults aged 75 to 79 years in the UK (UK Health Security Agency, 2024a). Systematic surveillance for RSV infection among adults admitted to hospital with acute respiratory-tract infections may be needed to identify any unintended consequences of reducing transmission among infants.

This analysis has many simplifying assumptions. I assumed all-or-none immunity, with those whose immunity waned being indistinguishable from those who were immune-naïve. I also assumed all-or-none immune boosting that either fully returned immunity or did not affect it. These do not

reflect realistic patterns of immunity for RSV, influenza (Fonville *et al.*, 2014), or SARS-COV-2 (COVID-19 Forecasting Team, 2023). This model also neglects clinically-important protection against severe disease and death, which persist after mucosal immunity to infection has largely waned (Moore *et al.*, 2023). Even while systemic protection against severe disease persists, mucosal immunity can wane sufficiently for infectious virus to be present in respiratory secretions (Russell and Mestecky, 2022).

I assumed that effects of starting and stopping non-pharmaceutical interventions were immediate, but during the covid-19 pandemic levels of engagement with restrictions varied over time (Smith *et al.*, 2022). My model differs from many RSV-specific models that include distinct sequential susceptible, infectious and immune states to represent accumulating effects of re-exposure (Lang, 2022). I assumed the pathogen was homogeneous, with no variation in strains and no evolution over time. Even for RSV, in which the immunogenic F protein is highly conserved (Ascough *et al.*, 2018), this may not accurately capture all relevant effects, as the dominant RSV group varies between outbreaks (Holmdahl *et al.*, 2024; Waris, 1991).

In models with SIRS dynamics, loss of immunity due to waning is equivalent to loss of immunity due to pathogen evolution (Gomes *et al.*, 2004). Given that exposure to new strains of influenza virus can boost immune response to multiple previous strains (Fonville *et al.*, 2014), it may be plausible that a similar equivalence exists for this model of immune boosting. Alternatively, if a new variant caused infection rather than boosting immunity to previous variants then the model would need to distinguish between immunity lost due to waning *vs.* immunity lost due to virus evolution.

Mathematically, I showed that identification of model parameters required knowledge of the waning rate or observation of each resistant compartment size. These compartments are intended as a discrete abstraction

of continuous changes in immunity. Serial measurements of immune markers such as antibodies and T cells could offer a practical approach toward estimating how immunity changes over time. If we assume a relationship between markers of immunity and the discrete resistant compartments, this could allow estimation of all parameters.

The wide range of plausible boosting coefficient values in our model replicates general uncertainty about dynamics of RSV transmission. White *et al.* (2007) were able to fit dramatically contrasting models to RSV datasets, including one with lifelong partial immunity and $\mathcal{R}_0 = 9.3$, and another with expected immunity lasting less than 7 months and $\mathcal{R}_0 = 1.3$. Baker *et al.*'s (2020, 2022) model of the effects of non-pharmaceutical interventions on transmission produced qualitatively equivalent results to my model, but assumed lifelong immunity and $\mathcal{R}_0 \approx 7$. Each model specification and set of parameters can lead to different predictions for disease dynamics. Underestimating natural immune boosting leads to underestimates of transmissibility and overestimates of the effect on incidence of a reduction in transmission. It is therefore vital to establish more robust understanding of natural immune boosting for respiratory-tract infections to inform successful disease control interventions.

3.9 Data and code

Data used in this analysis are available at `https://www.opendata.nhs.scot/dataset/viral-respiratory-diseases-including-influenza-and-covid-19-data-in-scotland/resource/5032c2b9-a206-4024-9536-5fdf345c8483`, released under the Open Government Licence for public sector information.

Code to reproduce the analysis in this chapter is available at

`https://github.com/markgpritchard/ImmuneBoostingODEs`.

Chapter Four

Should natural immune boosting inform vaccination schedules for healthcare workers?

Healthcare workers are highly exposed to respiratory infections (Mo *et al.*, 2021; Quigley *et al.*, 2021; Xiao *et al.*, 2020). In the first year of the covid-19 pandemic they had the greatest excess mortality of any occupation group in England and Wales (Matz *et al.*, 2022). Healthcare workers' vulnerability differs from that of other groups at high risk of severe disease. Most individuals prioritized for vaccination against SARS-CoV-2 are expected to be vulnerable to severe disease if infected (UK Health Security Agency, 2023). Healthcare workers do not necessarily share this elevated risk (although some healthcare workers will also be in a clinical risk group; Suárez-García *et al.*, 2020). Instead, the risk for healthcare workers relates to their high level of exposure to the virus and high probability of infection.

When the UK's covid-19 vaccination programme started in December 2020, healthcare workers were among the first priority groups to be vaccinated (Hall *et al.*, 2021). For primary vaccination, two doses were given three months apart (UK Health Security Agency, 2023). By September 2021, six months had passed since many healthcare workers had received their second dose. Concerns about waning immunity and spread of the Omicron variant led to roll-out of a third dose from September 2021 (Hall *et al.*, 2024).

In Chapter 3 I used a mathematical model to explore potential effects of immune boosting by re-exposure to a microparasite that induces rapidly-waning immunity. In that model, boosting could lead to a lower prevalence of infection. Immune boosting also led to cycles of infection as numbers in the most-immune compartment increased during outbreaks, followed by an increase in susceptibility as immunity simultaneously waned, fuelling a subsequent outbreak.

The force of infection experienced by healthcare workers need not be proportional to the prevalence of infection among their colleagues. Even if most healthcare workers were immune, they could experience a high force of infection from contact with patients. This could allow a situation in which immune boosting by a consistently high force of infection maintained a low level of susceptibility and a low incidence of infection among healthcare workers.

Immunity following vaccination against influenza can wane within a season (Domnich *et al.*, 2024; Rambhia and Rambhia, 2019). This suggests that people should not be vaccinated too early in the season (unless within-season repeat vaccination is available). However, with sufficient immune boosting, we might prefer to recommend early vaccination that will be maintained by re-exposure to the virus. This leads to two related questions:

1. What is the effect of natural immune boosting on healthcare workers' immunity to SARS-CoV-2? and,
2. If re-exposure to respiratory viruses substantially increases immunity without causing infection, how should this affect vaccination schedules?

In this chapter I explore how natural immune boosting and vaccination might interact in a series of simulations. I use publicly available data from

acute NHS hospitals in the Midlands of England between March 2020 and June 2022 to inform a compartmental model of infections among healthcare workers. I use this to explore counterfactual booster vaccination programmes for healthcare workers starting between July and November 2021.

4.1 Background and data

Vaccination

The first mass vaccination programme against SARS-CoV-2 started in the UK on 8 December 2020 (BBC News, 2020c). Healthcare workers started being vaccinated that very day (Hall *et al.*, 2021). They first became eligible for booster doses from 16 September 2021 (*ibid.*).

Trials of the BNT162b2 (Polack *et al.*, 2020) and mRNA-1273 (Baden *et al.*, 2021) vaccines both found little effect on incidence in the first ten days after vaccination. In their trial of the ChAdOx1 vaccine, Voysey *et al.* (2021) assumed no effect in the first fourteen days after vaccination so excluded events during that period from their analysis. The SIREN study of 20 thousand staff at NHS hospitals in England included questions on vaccination history (Hall *et al.*, 2021). That study estimated that BNT162b2 had an effectiveness among healthcare workers of 70 per cent twenty-one days after a first dose, and of 85 per cent seven days after a second dose (*ibid.*). Effectiveness after a third dose in the booster programme from September 2021 was estimated as 63 per cent against the Delta variant and 35 per cent against the Omicron variant (Hall *et al.*, 2024). Protection against the Omicron variant waned within four months. The ChAdOx1 vaccine was estimated to have a lower effectiveness, but this was received by only eight per cent of SIREN study participants (Hall *et al.*, 2022).

In Wales, uptake of the SARS-CoV-2 vaccine among healthcare workers

was rapid until February 2021, at which time 69 per cent had received at least one dose (Bedston *et al.*, 2022). Coverage then increased more gradually, with 85 per cent receiving at least one dose by June 2021. By April 2022, 68 per cent of frontline healthcare workers across England had received three doses of the SARS-CoV-2 vaccine (UK Health Security Agency, 2022a).

Testing and isolating healthcare workers

A pilot of routine testing of asymptomatic NHS staff was conducted in late 2020 (NHS England, 2020). Throughout 2021, asymptomatic staff were expected to test themselves twice weekly with SARS-CoV-2 rapid antigen tests (NHS England, 2021d). Isolation rules for healthcare workers changed several times during the covid-19 pandemic. At the start of 2021, healthcare workers needed to isolate away from work for ten days if they tested positive for SARS-CoV-2 infection (UK Health Security Agency, 2022b). They also needed to isolate if they were close contacts of someone with SARS-CoV-2. Any contact while wearing appropriate personal protective equipment – which most contacts in healthcare settings were assumed to be – did not lead to isolation. From 16 August 2021, healthcare workers who were contacts of an infected individual did not need to isolate if they were fully vaccinated and had daily negative tests for SARS-CoV-2 antigen (NHS England, 2021b). From 17 January 2022, healthcare workers who had tested positive for SARS-CoV-2 could return to work within ten days if they were consistently testing negative for SARS-CoV-2 antigen (NHS England, 2022b).

Numbers of patients and healthcare workers with covid-19

NHS England (2022a) published daily covid-19 situation data for all NHS hospitals in England from 2 April 2020. Data included total numbers of beds

occupied, numbers of beds occupied by a patient with a positive SARS-CoV-2 test, numbers of staff working in the hospital, and numbers of staff absent due to SARS-CoV-2 infection or exposure. I estimated proportions of infectious patients as the ratio of numbers of beds occupied by patients with confirmed SARS-CoV-2 and total numbers of occupied beds, and proportions of isolating healthcare workers from numbers of staff absent due to infection or exposure to SARS-CoV-2 and total numbers of staff.

4.2 Causal estimands

Effect of immune boosting on healthcare workers' immunity

Assume that the force of infection experienced by healthcare workers is a function of proportions of infectious healthcare workers, patients, and members of the community. This force of infection is expected to cause infections among susceptible healthcare workers. They will go on to become immune. The hypothesis of natural immune boosting suggests that the force of infection can also prolong immunity without causing infection (as in the model in Chapter 3).

Figure 4.1 represents causal relationships of infections, and immune waning and boosting for staff in one hospital. Subscript numbers represent times as generations of the virus's reproduction, so the relationship between numbers of infectious healthcare workers at time k and $k+1$ is explained by transmission, not by individuals remaining infectious over multiple time points. The diagram shows four time points, but healthcare workers' immunity should be assumed to be dependent on many previous generations of infections. The force of infection for healthcare workers and the proportion immune are explicitly represented on the causal diagram. They are shown in circles to signify that they are unobserved. Equivalent intermedi-

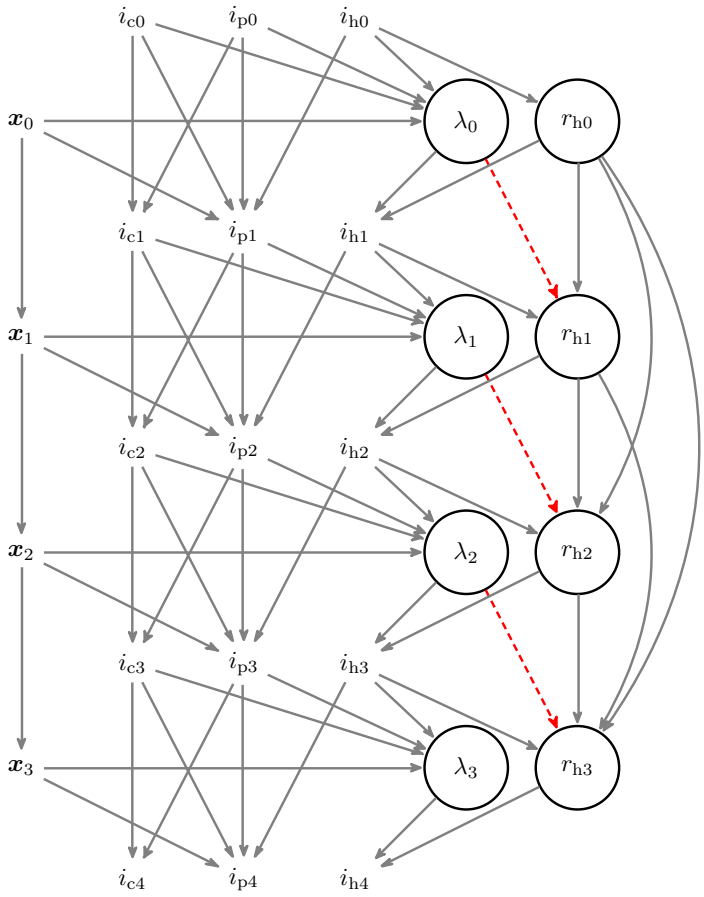

Figure 4.1: Causal diagram showing relationships between numbers of infectious patients (i_p), healthcare workers (i_h), and members of the community (i_c), and immunity among healthcare workers (r_h). Subscript numbers represent times in multiples of the virus' generation interval. The force of infection experienced by healthcare workers is λ, and $\boldsymbol{x}$ is a vector of covariates. Variables in circles are assumed to be unobserved. Dashed red lines represent natural immune boosting.

ate steps would apply for infections in the community and among hospital patients but explicitly listing them would not affect the conclusions drawn about healthcare workers' immunity. Connections between infectious community members and hospital patients also incorporate effects of infectious individuals being admitted or discharged from hospital.

The dashed red line represents natural immune boosting. This is a direct

path from the force of infection to healthcare workers' immunity, without intermediate infections. For the objective of estimating the effect of natural immune boosting, consider causal routes between the number of infected patients at time t_1 and the number of infected healthcare workers at time t_3,

1. $i_{\text{p1}} \to \lambda_1 \to i_{\text{h2}} \to \lambda_2 \to i_{\text{h3}}$,
2. $i_{\text{p1}} \to \lambda_1 \to i_{\text{h2}} \to r_{\text{h2}} \to i_{\text{h3}}$,
3. $i_{\text{p1}} \to i_{\text{p2}} \to \lambda_2 \to i_{\text{h3}}$,
4. $i_{\text{p1}} \to i_{\text{c2}} \to \lambda_2 \to i_{\text{h3}}$, and
5. $i_{\text{p1}} \to \lambda_1 \to r_{\text{h2}} \to i_{\text{h3}}$.

Routes 1 and 2 are mediated by infections among healthcare workers, route 3 by infections among patients, and route 4 by infections in the community. Route 5 operates via natural immune boosting. Without the unobserved nodes these paths are,

1. $i_{\text{p1}} \to i_{\text{h2}} \to i_{\text{h3}}$,
2. $i_{\text{p1}} \to i_{\text{h2}} \to i_{\text{h3}}$ (identical to route 1),
3. $i_{\text{p1}} \to i_{\text{p2}} \to i_{\text{h3}}$,
4. $i_{\text{p1}} \to i_{\text{c2}} \to i_{\text{h3}}$, and
5. $i_{\text{p1}} \to i_{\text{h3}}$.

Route 5 has no observed intermediate steps. To estimate the magnitude of natural immune boosting as a 'direct effect', we would need to adjust for the 'mediators', i_{h2}, i_{p2} and i_{c2}. However, there is mediator–outcome confounding via i_{h1} (see VanderWeele, 2015). Following Greenland *et al.*'s (1999) rules for causal diagrams, if we attempted to close route 1 by controlling for i_{h2}, we open a path,

$$i_{\text{p1}} \to \lambda_1 \to \boxed{i_{\text{h2}}} \leftarrow r_{\text{h1}} \to r_{\text{h2}} \to i_{\text{h3}},$$

as i_{h2} is a collider (the box around i_{h2} indicates that the variable is being controlled for in the analysis).

Without being able to observe the force of infection or the immunity of healthcare workers, it is impossible to account for the mediator–outcome confounding and calculate the direct effect of immune boosting. I therefore use an epidemic model that estimates the force of infection and immunity among healthcare workers. This model has a coefficient for the magnitude of natural immune boosting as a proportion of the force of infection. I use estimates of this coefficient as an estimate of the effect of immune boosting.

Effect of changing booster vaccination dates

In this chapter I use hospital-level data and consider hospitals as units being treated. The intervention is a booster vaccination programme, lasting for 90 days. Let $A = \{-2, -1, 0, 1, 2\}$ represent policies of starting the vaccination programme between two months sooner than the actual start on 16 September 2021, and two months later. Let $\Psi = \psi$ represent the immune boosting coefficient.

Staff at all NHS hospitals in England were eligible for booster vaccination at the same time. I do not have access to daily uptake data per hospital, so methods such as target-trial emulation (Hernán *et al.*, 2008) are not possible. By varying the date of vaccination in an epidemic model I can explore outcomes with counterfactual vaccination dates.

The model assumes that vaccination is effective at moving individuals from susceptibility or waned immunity to the fully-immune state. It is trivial to deduce that vaccination in July reduces the number of infections in July and August, compared to starting vaccination in September. For each scenario, I assume the same rate of vaccination and the same duration of the vaccination programme. I estimate the effect of vaccination on the total

number of infections among healthcare workers between the start of July 2021 and the end of June 2022. Let,

$$Y_j := \sum_{x=1\ \text{July}\ 2021}^{30\ \text{June}\ 2022} i_{\text{h}xj}, \tag{4.1}$$

where Y_j is the total number of infections in hospital j and $i_{\text{h}xj}$ is incidence among healthcare workers in hospital j at time x. Then $Y_j^{A=a,\,\Psi=\psi}$ is the potential value of Y_j if the hospital followed a policy $A = a$, and the immune boosting coefficient $\Psi = \psi$. The potential force of immune boosting is not amenable to manipulation but the value of the boosting coefficient can be varied between simulations. Here, I explicitly state it in the notation for potential outcomes. The causal effect of a vaccination programme $A = a$ is,

$$\left[\tau_{a\,j} \mid \Psi = \psi\right] = Y_j^{A=a,\,\Psi=\psi} - Y_j^{A=0,\,\Psi=\psi}. \tag{4.2}$$

4.3 Description of the model

The model in this chapter is a modified discrete-time version of the previous chapter's model. Added compartments represent an exposed (latent) period and a ten-day isolation period following diagnosis (Figure 4.2). For this model, I focus on the effect of the force of infection on healthcare workers. I do not attempt to model the effect of infectious healthcare workers on infections among patients (Cooper *et al.*, 2023, found that very few SARS-CoV-2 infections among patients were directly caused by infectious healthcare workers).

Susceptible healthcare workers are exposed to a force of infection from infectious healthcare workers, patients, and people in the community. Let n_{c} be the population size of the community, $n_{\text{p}j}$ the number of patients in hospital j, and $n_{\text{h}j}$ the number of healthcare workers at hospital j, each of which is constant over time. Let $i_{\text{c}t}$ be the number of infectious community

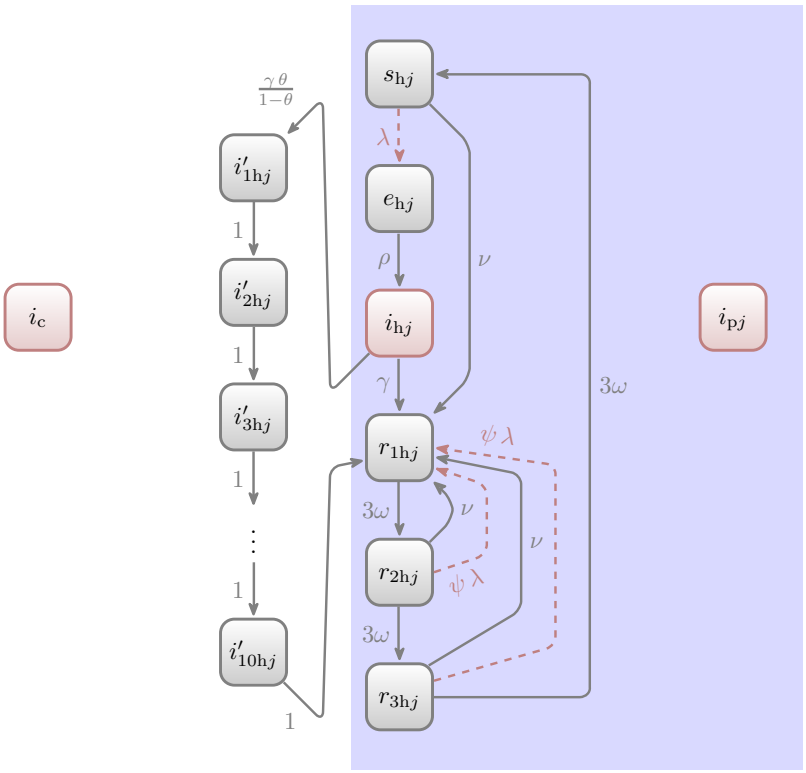

Figure 4.2: Model flow chart. Compartments of healthcare workers are, s, susceptible; e, exposed (latent period); i, infectious, i', diagnosed and isolating; r, resistant (immune). Infectious community members, i_c, and patients, i_p, contribute to the force of infection, but their infection and recovery are not modelled. The purple shaded rectangle contains compartments within the hospital. Events that occur in proportion to the force of infection are represented by dashed red lines. See text for parameter definitions.

members at time t, i_{pjt} the number of infectious patients, and i_{hjt} the number of infectious healthcare workers who are not isolating at home. The force of infection experienced by healthcare workers is,

$$\lambda_{jt} := \beta_{ct} \frac{i_{ct}}{n_c} + \beta_{hj} \frac{i_{hjt}}{n_{hj}} + \beta_{pj} \frac{i_{pjt}}{n_{pj}}, \tag{4.3}$$

where β coefficients represent infectiousness from each infectious source. The community transmission coefficient can vary over time. The healthcare worker and patient transmission coefficients vary between hospitals but are assumed constant over time. The probability of infection for a healthcare

worker during a period $(k, k + \Delta t)$ is,

$$\lambda'_{jt} = 1 - \exp\big(-\lambda_{jt}\,\Delta t\big). \tag{4.4}$$

I run the model with a daily time interval so $\Delta t = 1$.

To avoid exponential terms of constant values, I transform transition probabilities that do not depend on the force of infection. Undiagnosed infectious healthcare workers have a probability, $\theta \in (0, 1)$, of being diagnosed before recovering and a recovery rate of γ. To maintain the mean infectious period for those who are not diagnosed as $1/\gamma$ and the total probability of diagnosis as θ, the daily probability of diagnosis before recovery is,

$$\theta' = \theta\left(1 - \exp\left(\frac{-\gamma}{1-\theta}\right)\right), \tag{4.5}$$

and the daily probability of recovery before diagnosis is,

$$\gamma' = (1-\theta)\left(1 - \exp\left(\frac{-\gamma}{1-\theta}\right)\right). \tag{4.6}$$

When diagnosed, healthcare workers isolate away from the hospital for exactly ten days. This isolation is represented by a series of ten sequential isolating compartments, $i'_{1\mathrm{h}j}$ to $i'_{10\mathrm{h}j}$. While isolating, healthcare workers do not contribute to the force of infection in the hospital but are still counted in the (constant) cohort size, $n_{\mathrm{h}j}$.

Immunity is represented by three sequential compartments, $r_{1\mathrm{h}j}$ to $r_{3\mathrm{h}j}$ (as in Chapter 3). Immunity wanes between compartments at a rate 3ω, $\omega \in (0, 1/3]$, and can be boosted by natural immune boosting or vaccination at a rate,

$$\varphi_{jt} = 1 + (\nu_{jt} - 1)\,\exp\big(-\psi\,\lambda_{jt}\big), \tag{4.7}$$

where ψ is the boosting coefficient, and ν_{jt} is the time-varying vaccination

rate.[1] Disease dynamics for healthcare workers are therefore described by,

$$\left.\begin{aligned}
s_{\mathrm{h}j(t+1)} &= \left(1-\lambda'_{jt}-\nu_{jt}+\nu_{jt}\,\lambda'_{jt}\right)s_{\mathrm{h}jt}+3\omega\left(1-\varphi_{jt}\right)r_{3\mathrm{h}jt},\\
e_{\mathrm{h}j(t+1)} &= \left(1-\varsigma\right)e_{\mathrm{h}jt}+\lambda'_{jt}\,s_{\mathrm{h}jt},\\
i_{\mathrm{h}j(t+1)} &= \left(1-\gamma'-\theta'\right)i_{\mathrm{h}jt}+\varsigma\,e_{\mathrm{h}jt},\\
i'_{1\mathrm{h}j(t+1)} &= \theta'\,i_{\mathrm{h}jt},\\
i'_{k\mathrm{h}j(t+1)} &= i'_{(k-1)\mathrm{h}jt}, \qquad k\in\{2,3,\ldots,10\},\\
r_{1\mathrm{h}j(t+1)} &= \left(1-3\omega+3\omega\,\varphi_{jt}\right)r_{1\mathrm{h}jt}+\nu_{jt}\left(1-\lambda'_{jt}\right)s_{\mathrm{h}jt}\\
&\qquad +\gamma'\,i_{\mathrm{h}jt}+i'_{10\,\mathrm{h}jt}+\varphi_{jt}\left(r_{2\mathrm{h}jt}+r_{3\mathrm{h}jt}\right),\\
r_{2j(t+1)} &= \left(1-3\omega-\varphi_{jt}+3\omega\,\varphi_{jt}\right)r_{2\mathrm{h}jt}+3\omega\left(1-\varphi_{jt}\right)r_{1\mathrm{h}jt},\\
r_{3\mathrm{h}j(t+1)} &= \left(1-3\omega-\varphi_{jt}+3\omega\,\varphi_{jt}\right)r_{3\mathrm{h}jt}+3\omega\left(1-\varphi_{jt}\right)r_{2\mathrm{h}jt},
\end{aligned}\right\} \tag{4.8}$$

where e represents the number exposed who are in the latent (pre-infectious) period, and $1/\varsigma$ is the mean duration of the latent period. Terms such as, $+\nu_{jt}\,\lambda'_{jt}$, avoid double-counting of individuals who experience two events during an interval $[k,\, k+\Delta t]$. The term $+3\omega\,\varphi_{jt}r_{1\mathrm{h}jt}$ in the equation for $r_{1\mathrm{h}j(t+1)}$ accounts for individuals who both waned to $r_{2\mathrm{h}j}$ then were boosted within the time interval $[k,\, k+\Delta t]$.

Vaccination rates

The modelled vaccination rate, ν_{jt}, is intended to move people to immunity ten days later than their vaccination date. It is therefore equal to 0 until ten

[1] Vaccination coverage (the proportion of eligible people who are vaccinated) is sometimes described as a 'vaccination rate'. Here, ν_{jt} is the rate (number per unit time) at which healthcare workers become vaccinated.

days after vaccination became available on 8 December 2020. I use,

$$\nu_{jt} = \begin{cases} 0.035, & \text{18 December 2020} \leq t \leq \text{3 January 2021}, \\ 0.055, & \text{8 January 2021} \leq t < \text{28 January 2021}, \\ 0.025, & \text{28 January 2021} \leq t \leq \text{15 February 2021}, \\ 0.0125, & \xi \leq t < \xi + 90, \\ 0, & \text{otherwise}, \end{cases}$$

where ξ is 10 days after the start date of the booster vaccination programme. These rates give a pattern similar to that seen for the initial vaccine roll-out among SIREN-study participants (Hall *et al.*, 2021) and result in 89 per cent of healthcare workers being vaccinated by 5 February 2021 (immune by 15 February). The rate for the booster vaccination programme results in 64 per cent of healthcare workers receiving a booster vaccine (Figure B.22).

4.4 Modelled outcomes

In each case, $\varsigma = 0.5/\text{day}$ and $\gamma = 0.2/\text{day}$ so the generation time of the virus is 7 days, and $\omega = 0.01/\text{day}$ giving an immune duration of 100 days without boosting. In simulations I set $\theta = 0.5$ so half of all infectious healthcare workers are diagnosed and isolated.

Simulation with sinusoidal changes in force of infection

In the first simulation, I generate a sinusoidal force of infection, $\beta_j = \beta_{\text{p}j} = 0$, $i_{\text{c}t} = n_{\text{c}} = 1$, $\lambda_{jt} = \beta_{\text{c}t} = 0.01 + 0.005\cos(2\pi\,(t-31)/365)$, where t is days since 1 January 2020 and the force of infection peaks on 1 February each year. This force of infection is independent of the number of infected healthcare workers. I produced five simulations with immune boosting coefficients between 0 and 10 (Figure 4.3*A*).

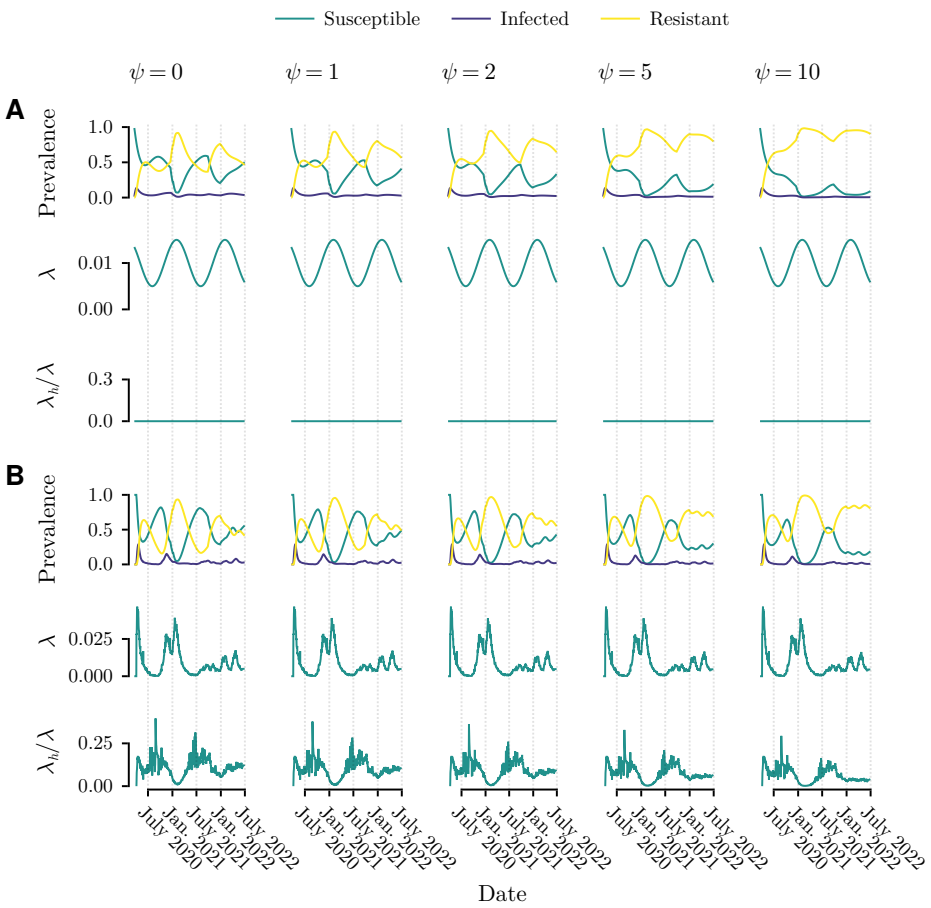

Figure 4.3: Simulated compartments. *A* Simulation with a sinusoidal force of infection that is independent of numbers of infected patients or healthcare workers. *B* Simulation with a force of infection experienced by healthcare workers with parameters for infection from the community, infected patients and other healthcare workers. In each, the top row shows the proportions susceptible, infected (including latent period, undiagnosed and diagnosed) and resistant, the second row shows the per capita force of infection experienced by healthcare workers, and the third row shows the proportion of the force of infection caused by infectious healthcare workers. Columns represent different modelled boosting coefficients.

Simulation with infection from community, patients and healthcare workers

For the second simulation I used recorded prevalence of SARS-COV-2 infection in the community and one hospital's recorded numbers of diagnosed patients. I set $\beta_j = \beta_{pj} = 0.075$, $\beta_{ct} = 0.02$. Some of the force of infection

was generated by infectious healthcare workers, but most was from patients and the community (Figure 4.3*B*).

Parameters fit to data from one hospital

I assumed that all healthcare workers were initially susceptible. I assumed that the force of infection from the community was related to the level of intervention in place to reduce transmission in the community,

$$\beta_{ct} = \max(0,\, \alpha_1 + \alpha_2\, \kappa_t), \tag{4.9}$$

where κ_t is the stringency index (see page 15) on day t. The prior distribution for α_1 was normally distributed, $\alpha_1 \sim \mathcal{N}(0.2,\, 1)$, and truncated, $\alpha_1 \in (-1,\, 10)$, so that the MCMC algorithm would not explore extremely negative values of α_1 that would not affect the transmission parameter. The prior distribution for α_2 was also normally distributed, $\alpha_2 \sim \mathcal{N}(0.1,\, 0.1)$, and truncated, $\alpha_2 \in (0,\, 10)$, to ensure that greater stringency did not cause a greater force of infection. Prior distributions for β_j, $\beta_{\mathrm{p}j}$ and ψ were each Exponential(1), and for θ was Beta$(1,\, 1)$.

Parameters fit to data from multiple hospitals

Parameters were fit for multiple hospitals in the same way as for a single hospital. Additionally, β_j and $\beta_{\mathrm{p}j}$ parameters were selected from a hierarchical model. For each there was assumed to be a shared mean parameter across all hospitals, with additional parameters related to the heated volume of the hospital and proportion of beds in side rooms (variables found to be associated with extent of transmission in hospitals by Cooper *et al.*, 2023). Each hospital's transmission coefficients were sampled from a normal distribution with this calculated mean, truncated to ensure all β coefficients were greater than 0.

Estimating effect of changing vaccination dates

I re-ran simulations with counterfactual vaccination start dates between two months sooner and two months later than the actual start in September 2021. When using fitted parameters, I ran simulations with each sampled parameter set for each counterfactual vaccination start date and report the median and central 90 per cent of outputs.

4.5 Effect of changing vaccination dates in simulations with different levels of natural immune boosting

Simulation with sinusoidal changes in force of infection

Figure 4.4*A* shows cumulative difference in incidence for simulations with different start dates, each compared to starting on 16 September (which I refer to as the 'prime' simulation). Simulations with earlier vaccination initially had fewer infections. Simulations with later vaccination diverged from the prime simulation on 26 September 2021 when that simulation started to show the effect of vaccination.

With no immune boosting, $\psi = 0$, all counterfactual simulations had a greater total number of infections by July 2022 than the prime simulation. With greater levels of immune boosting, the initial effect of earlier vaccination was reduced and the final (by July 2023) effect of early vaccination was reversed, showing a lower total number of infections with early vaccination and a large boosting coefficient. Initial differences caused by later vaccination were also reduced, but by July 2023 each simulation with later vaccination and immune boosting had a greater total number of infections among healthcare workers than the simulation with no immune boosting.

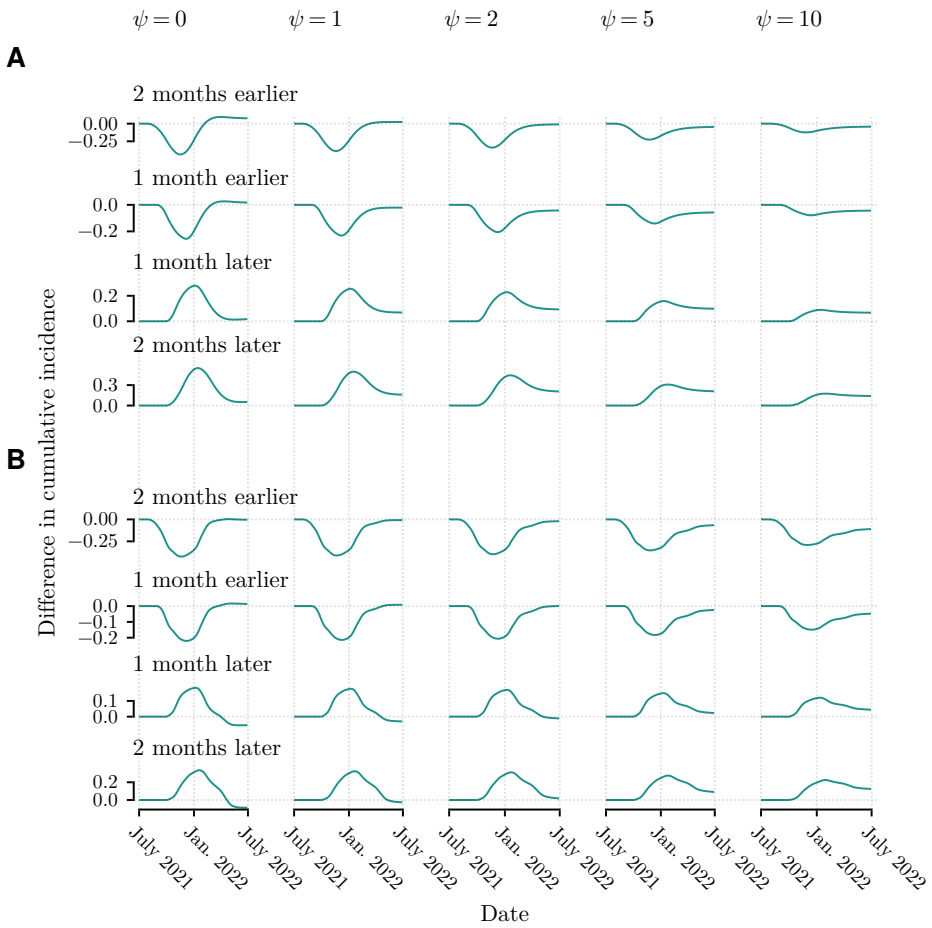

Figure 4.4: Cumulative difference in incidence for simulated compartments, compared to a simulation with vaccination of healthcare workers starting on 16 September 2021. In each plot, the prime simulation and the counterfactual have the same immune boosting coefficient, ψ, as shown at the top of each column. *A* Simulations with a sinusoidal force of infection. *B* Simulations with a force of infection experienced by healthcare workers with parameters for infection from the community, infected patients and other healthcare workers.

Simulation with infection from community, patients and healthcare workers

The simulation with infections caused by community members, patients and healthcare workers initially showed broadly similar results (Figure 4.4*B*) but simulations without immune boosting and a later start to the vaccination programme ended with a net reduction in the total number of infections. Greater levels of immune boosting had the equivalent effect as in the pre-

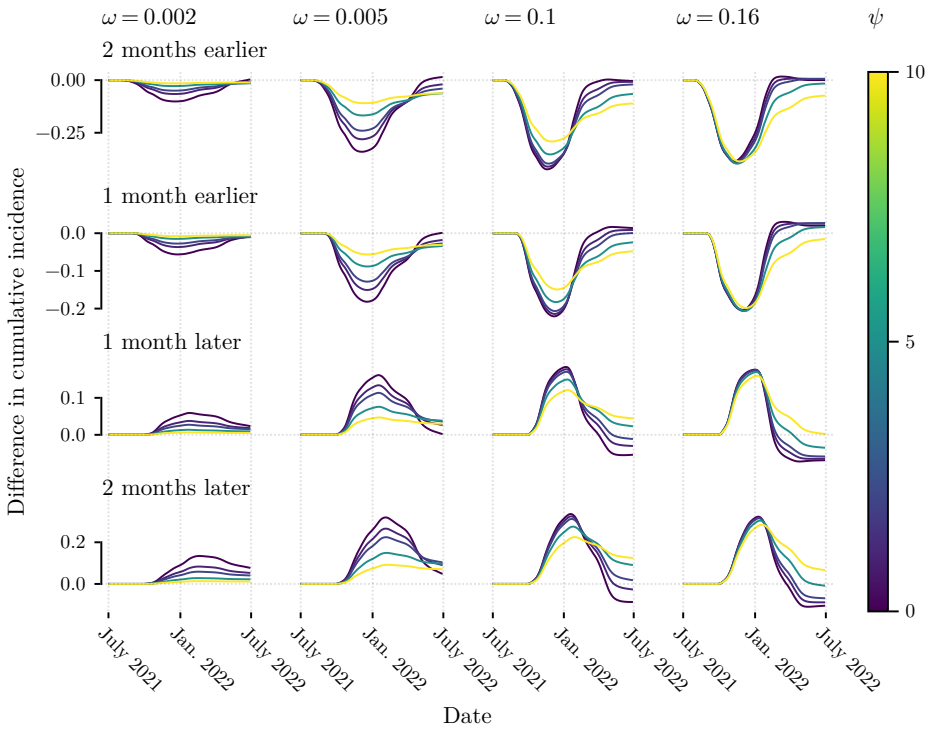

Figure 4.5: Cumulative effect of changing the start date of healthcare worker vaccination programmes for different immune-waning rates, ω. Boosting coefficients, ψ, are represented by different colours.

vious set of simulations. Simulations with earlier vaccination and more boosting had a net reduction in total numbers of infections, and those with later vaccination had a net increase in infections.

Varying the immune boosting coefficient affected the magnitude of the effect of vaccination (Figure 4.5). With lower values of ω (longer-lasting immunity without boosting), there are fewer infections than in simulations with more rapidly waning immunity. The trend that earlier vaccination leads to fewer infections in the presence of natural immune boosting is seen across all simulations. The results for later vaccination are less consistent. For simulations with immunity lasting for 500 days without boosting, greater levels of boosting lead to fewer infections for all vaccination dates,

and later vaccination leads to more infections than earlier vaccination across all values of the immune boosting coefficient.

4.6 Results

Data from one hospital

Not all chains mixed well. The well-mixed chains are shown in Figure B.23. Fitted parameters for the model with one hospital's data were $\beta_{\mathrm{h}j} = 0.013$ (90% CrI 0.0010, 0.059), $\beta_{\mathrm{p}j} = 0.031$ (90% CrI 0.029, 0.033), $\alpha_1 = -0.50$ (90% CrI -0.96, 0.61), $\alpha_2 = 0.040$ (90% CrI 0.022, 0.054), $\theta = 0.97$ (90% CrI 0.91, 0.996). The median boosting coefficient was $\psi = 5.2$ (90% CrI 3.92, 6.80). Modelled numbers of diagnosed healthcare workers are broadly similar to the observed numbers, but the model has a large peak in prevalence in late 2020 followed by a rapid reduction in infections, and it underestimated the size of the peak in early 2022 (Figure 4.6).

The proportion of the force of infection attributable to infectious healthcare workers varied over time. The 90% CrI included the possibility of a very large proportion, but the median estimate was always less than 0.5.

Estimating the effect of changing the start time of vaccination gave a similar result to the simulations. By July 2022, predictions for all vaccination start dates had similar total numbers infected but each was slightly greater than numbers predicted for a vaccination programme that started on 16 September 2021,

$$\tau_{A=-2,\,j=1} = 0.019,\ 90\%\ \mathrm{CrI}\ 0.0074,\ 0.029,$$

$$\tau_{A=-1,\,j=1} = 0.0080,\ 90\%\ \mathrm{CrI}\ 0.0014,\ 0.014,$$

$$\tau_{A=1,\,j=1} = 0.011,\ 90\%\ \mathrm{CrI}\ 0.0019,\ 0.020,$$

$$\tau_{A=2,\,j=1} = 0.042,\ 90\%\ \mathrm{CrI}\ 0.023,\ 0.060.$$

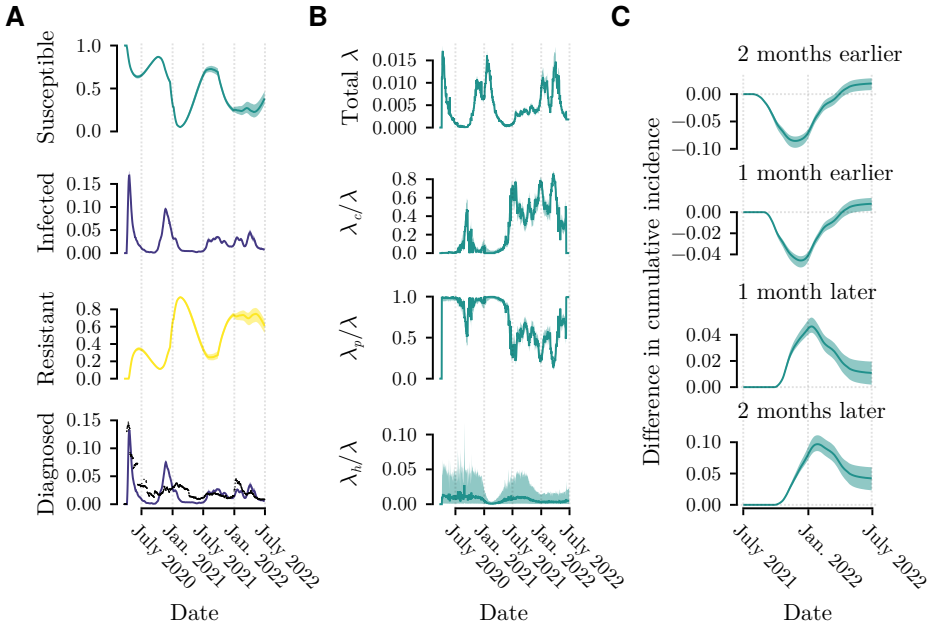

Figure 4.6: Results for healthcare workers with parameters fitted to data from a single hospital. *A* Size of each compartment in simulations using fitted parameters. *B* Total force of infection, λ and fraction of λ caused by community members, λ_c, patients, λ_p, and other healthcare workers, λ_h. *C* Counterfactual scenarios with different start dates for vaccination of healthcare workers. In each plot, the dark line represents the median values and shading represents the central 90% of values.

Data from all hospitals in the Midlands

Chains mixed well (Figures B.24 and B.25). Fitted parameters for most hospitals give a good fit to observed numbers of cases (Figure B.26). The median estimated boosting coefficient, $\psi = 3.5$ (90% CrI 3.3, 3.7), was lower than that estimated for a single hospital, but still greater than 1, suggesting that boosting occurred more readily among immune-waning healthcare workers than infection did among susceptible staff. The trend for most hospitals matched that for the single hospital (Figure 4.7). Almost all hospitals had greater total numbers of infections in counterfactual scenarios than in the model with vaccination starting on 16 September. Point estimates of total effects were that two hospitals would have had fewer infections if vaccination had

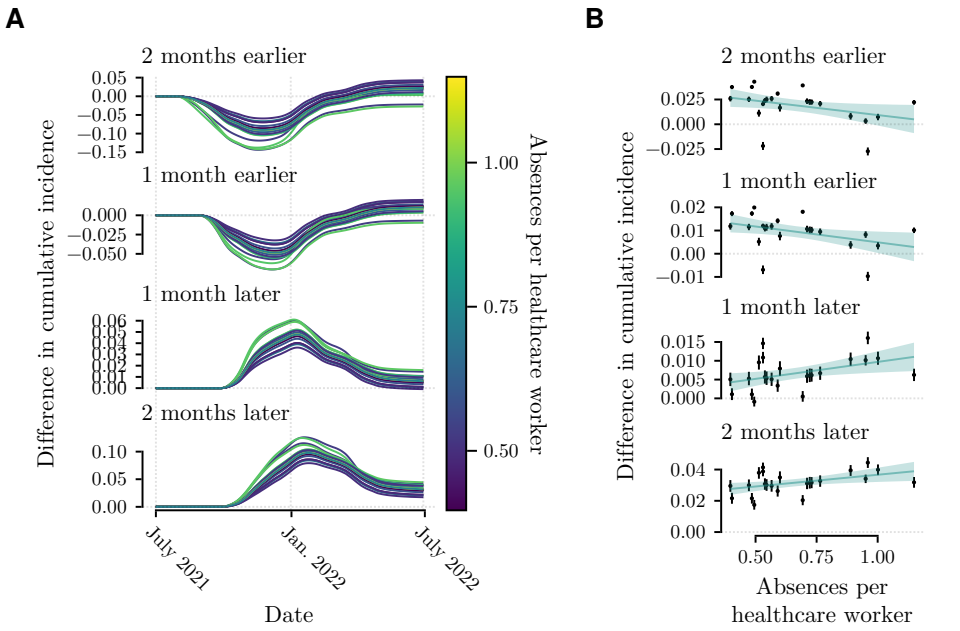

Figure 4.7: Counterfactual scenarios for different vaccination start dates with parameters fitted to data from all hospitals in the Midlands of England. *A* Median estimated cumulative difference in incidence for simulations with different start dates, each compared to starting on 16 September, with colours representing mean number of absences per healthcare worker in the data between 1 July 2021 and 1 July 2022. *B* Median (black points) and central 90% range (black bars) predicted difference in total numbers of infections against observed absences per healthcare worker, with linear trend and 90% credible interval shaded.

been sooner, and one hospital would have had fewer infections if vaccination had started one month later. In general, hospitals with greater numbers of absences among healthcare workers had smaller increases in numbers of infections in scenarios with earlier vaccination, and greater increases in scenarios with later vaccination.

4.7 Discussion

Simulations in this chapter suggest that understanding natural immune boosting could be important for determining optimal schedules for boosting vaccination against respiratory-tract viruses for healthcare workers. For

pathogens with rapidly waning immunity, vaccination *too soon* or *too late* can lead to greater total numbers of infections. A schedule that would lead to greater numbers of infections in scenarios without immune boosting could lead to fewer infections in those with a greater boosting coefficient. Equivalently, a time that was not *too late* without boosting could lead to a net increase in infections in the presence of natural immune boosting.

Application of this model to numbers of infections in acute NHS hospitals in the Midlands of England suggested that starting booster vaccination in July, August, October or November 2021, rather than September, would have led to greater numbers of infections among healthcare workers. Although changing the start date appeared detrimental for almost all hospitals, those with the greatest number of absences among staff (presumed to have the greatest mean force of infection) tended to show smaller increases in infections from vaccinating earlier and larger increases from vaccinating later.

This analysis does not provide guidance on absolute times that are either *too early* or *too late*. Rather, it suggests that *sooner* vaccination will tend to be more beneficial for groups experiencing greater levels of natural immune boosting. In simple simulations, this is produced with greater boosting coefficients. In results fitted to hospitals' data it is seen in the trend that hospitals with more infections have fewer additional infections predicted in counterfactual scenarios with early vaccination and more in counterfactuals with later vaccination.

This analysis suggests that in the presence of immune boosting, key workers – including healthcare workers – whose work cannot be paused or conducted remotely during a pandemic would benefit from earlier booster vaccination. My analysis has not attempted to prioritize people for vaccination in terms of benefits to wider public health. The model explicitly excluded any effects of vaccinating healthcare workers on transmission

to patients. Previous models, which included onward transmission from key workers but not natural immune boosting, concluded that key workers should be prioritized during initial vaccine roll-out against pandemic influenza (Mylius *et al.*, 2008) and SARS-CoV-2 (Lee *et al.*, 2021; Mulberry *et al.*, 2021). Key workers' greater numbers of contacts mean they have a high risk of being infected and, if infected, a high probability of onward transmission. This is especially true when transmission in the community is otherwise suppressed by restrictions on social contacts (Cooper *et al.*, 2023). However, these models were considering prioritization of vaccines as a scarce resource. They guide the decision to prioritize healthcare workers for inclusion in a booster campaign but do not inform timing of that campaign.

Rosenstrom *et al.* (2022) considered timing of booster doses and initial vaccination for children. In the United States, initial SARS-CoV-2 vaccination for children and booster vaccination for adults were both offered from the same month, November 2021 (*ibid.*). Similarly, vaccination of some children in England started in the same month as some booster vaccination for adults, September 2021 (NHS England, 2021c). Rosenstrom *et al.* (2022) found that starting primary vaccination of children sooner would have had a greater effect on mortality from SARS-CoV-2 than changing the start date of the booster campaign. The benefit of vaccinating children was attributed to their large number of contacts and high chance of being in a chain of transmission.

In some simulations, later vaccination led to greater numbers of infections during autumn and winter but a lower total number of infections in the twelve-month period starting in July 2021. This is shown explicitly in plots comparing later vaccination with programmes starting in September, but is also present for vaccination in September compared to programmes starting earlier. This raises a question of how short- and long-term effects should be balanced. Castioni *et al.* (2024) used a model to assess optimal use

of a limited number of vaccine doses to minimize numbers of infections and deaths. They found that rapid vaccination strategies that caused a large decrease in infections could cause a synchronized return to susceptibility and a large second peak in infections. There, vaccination played a similar role in causing cyclic outbreaks as natural immune boosting did in my model in Chapter 3. Castioni *et al.* (2024) found that a slower vaccination programme that lasted for longer had less risk of causing a rebound outbreak. They also found that starting vaccination after an outbreak has reached its peak maximizes the chance of 'flattening' the subsequent outbreak's peak. This fulfils a utilitarian objective of maximizing benefit from available resources (Savulescu *et al.*, 2020). However, the models simplify time-varying conditions. By the following spring there could be changes in use of physical measures to prevent transmission, new variants of the virus that evade existing immunity, or availability of better vaccines and treatments. Accepting short-term harms in the expectation of long-term benefits based on a compartmental model risks net harms from unanticipated violations the model assumptions. In the UK, people at the greatest risk of severe disease have been offered vaccination against SARS-CoV-2 twice per year since spring 2022 (UK Health Security Agency, 2023). In that setting, it is more reasonable to minimize short-term numbers of infections.

I described the question of optimal timings for vaccination in terms of the first booster programme in the UK in 2021. It might be better considered in terms of regular annual vaccination against respiratory viruses. Akhtar *et al.* (2023) report a comparison from an RCT vaccinating people with acute coronary syndrome against influenza. They compared those vaccinated early in the season to those who were vaccinated late in the season. But this was not a randomized comparison. All participants were vaccinated within 72 hours of admission or percutaneous coronary intervention. Observed

differences might be due to confounding variables (such as differences in background risk in those who remained unvaccinated late in the influenza season, compared to those who were eligible for the trial at the start of the season but would soon have gone on to be vaccinated), or might be due to seasonal differences in outcomes following acute myocardial infarction (*e.g.* see Vallabhajosyula *et al.*, 2020). Costantino *et al.* (2019) modelled waning of immunity following influenza vaccination. They concluded that changing vaccination dates in Australia from March to May or June would prevent more infections. They also considered how changing vaccination policy might affect vaccination coverage. If a delay led to a reduction in vaccine uptake this could offset any advantages of optimizing timing.

This analysis provides preliminary evidence on the possible role of natural immune boosting in determining optimal timing for booster vaccination. It suggests the direction that any change to vaccination timing should take. Simulations consistently suggested that immune boosting leads to benefits from vaccinating sooner. Given sufficient evidence of immune boosting, changes in timing of vaccination programmes might be justified. Alternatively, it might simply be an encouragement for those who are most exposed to access the vaccine as soon as it is available. This is an intuitively reasonable proposal even without natural immune boosting, so might be more acceptable than attempts to enforce changes to the vaccination programme for all healthcare workers.

The results also lead to a question: Given a choice between avoiding exposure with high-quality personal protective equipment or maintaining immunity by ensuring ongoing exposure to the virus throughout the winter season, which should a healthcare worker choose? Deliberately increasing exposure is clearly a high-risk strategy and there is currently no evidence to recommend it. If a healthcare worker could always avoid exposure there

would be no need for boosted immunity so it is unlikely that deliberate exposure would ever be advocated. However, intermittent use of high-quality personal protective equipment, alternating periods of exposure and protection, could potentially lead to paradoxical increases in infection that might warrant further study.

This preliminary investigation clearly has limitations. Many important features were assumed. Key among these is that vaccination rates were assumed equal for all hospitals. If differences in numbers of infections were due to differences in uptake of vaccination then fitted transmission coefficients for hospitals with lower uptake could over-estimate the force of infection and lead to incorrect conclusions in counterfactual scenarios. The model also had many simplifying assumptions, such as constant rates of transmission and detection within hospitals. It did not include differences in vaccine efficacy which are observed for different variants of SARS-CoV-2 (Hall *et al.*, 2024).

To gain greater understanding of the relationship between natural immune boosting and effectiveness of booster vaccination, it would be useful to collect individual-level data on vaccination, workplace exposure and infection history. Serial measures of antibodies and T cells would also help to understand trajectories of immunity. I propose that analysis of immune markers would be an important next-step in determining whether assessment of natural immune boosting should be considered in recommendations regarding boosting vaccination.

4.8 Data and code

Numbers of beds occupied by someone with covid-19 and numbers of healthcare workers absent due to covid-19 are available from

https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/, shared under an Open Government Licence, v 3.0. I took numbers from the 'Monthly COVID Publication' files. Details about hospitals' sizes and numbers of side rooms are from Cooper *et al.*'s (2023) repository, https://github.com/BenSCooper/nosocomial_COVID_England/tree/code-for-resubmission, shared under an MIT licence. Hale *et al.*'s (2021) dataset is available at https://github.com/OxCGRT/covid-policy-tracker/tree/master/data/United%20Kingdom, using files OxCGRT_GBR_differentiated_withnotes_2021.csv and OxCGRT_GBR_differentiated_withnotes_2022.csv, shared under a Creative Commons Attribution 4.0 International Public Licence.

Code to reproduce the analysis in this chapter is available at https://github.com/markgpritchard/ImmuneBoostingHealthcare.

Parameter identifiability

Chapter Five

Discussion

> The greater our knowledge increases,
> the greater our ignorance unfolds.
>
> John F. Kennedy
> Address at Rice University (1962)

Perhaps the clearest conclusion from this thesis is that uncertainty about effects of infection-prevention measures persist. These results do not allow us to conclude that any intervention should be used or avoided or delayed. But the results do provide a step along the pathway to improved understanding of how interventions and behaviour and immunity interact.

In Chapter 2 I explored use of difference-in-differences methods to evaluate infection-prevention measures. I modified the method to account for person-to-person transmission dynamics, then applied it to two datasets. The results suggest that widespread community testing and mandatory use of face coverings each increased diagnoses of SARS-CoV-2. I discuss interpretation of this result in more detail below. In Chapter 3 I showed that unless we know the effect of natural immune boosting, magnitudes of interventions' effects are not identifiable. Different assumptions about the rate of immune waning led to very different estimates of model parameters and of the effectiveness of restrictions to reduce the spread of SARS-CoV-2. In Chapter 4 I demonstrated that – subject to a model's assumptions – chan-

ging the date when healthcare workers started to receive booster vaccinations would have had different effects on total numbers infected depending on the strength of natural immune boosting.

5.1 Interpretation

Difference-in-differences analyses

Analyses in Chapter 2 suggested that mass testing for SARS-CoV-2 and mandated mask use increased incidence of diagnosed SARS-CoV-2. These results cannot simply be dismissed but they also must not be accepted without question.

The first consideration for the results of testing in Liverpool is that my preferred analysis would use pillar 1 data, which are results of testing patients in NHS hospitals, and testing health and care workers (Department of Health and Social Care, 2020b). That analysis failed to provide plausible results. It suggested that SARS-CoV-2 would have been practically eliminated due to almost-universal immunity by late 2021, clearly failing to match observed outcomes. An analysis of total numbers of reported cases gave results that appeared more plausible. These data include cases diagnosed by the testing programme. As discussed in that chapter, the inferred effect that testing caused greater numbers of diagnoses might represent detection of a greater proportion of asymptomatic infections so may reflect success of the testing programme.

From a perspective of wanting to minimize the burden of SARS-CoV-2, the worst-case interpretation would be that these results show an increase in incidence, not merely in detection. If we assume that interpretation, how should we interpret the causal results?

One possibility is that the interventions caused harm. This may be plaus-

ible. Early in the pandemic, official advice in the United Kingdom was that masks used incorrectly might do more harm than good (Baynes, 2020), and false reassurance following a negative test result while infectious might lead to more risky behaviour. Jon Deeks expressed concern that widespread testing could lead to unintended harms, saying, 'If people get a false sense of security, they could actually spread the virus around' (Guglielmi, 2021). Batteux *et al.* (2022) estimated that almost thirty per cent of people misunderstood the chance of being infectious following a negative rapid antigen test for SARS-CoV-2, although they found no evidence that this changed adherence to infection-prevention measures. However, if testing and masking were indeed harmful, we would expect to see evidence of that harm in other forms of research (see Lawlor *et al.*, 2016, for discussion of triangulation of causal results from multiple sources).

A second possible explanation is that numbers of infections increased for a reason unrelated to the intervention. Infectious disease outbreaks can cycle between peaks and troughs without any intervention (*e.g.* see Dushoff *et al.*, 2004; Finkenstädt and Grenfell, 2000; and my result in Chapter 3). Such a peak could have coincided with the introduction of infection-prevention measures. Policy makers may have identified areas that were about to experience an increase in infections and, like Rubin's (2004) hypothetical 'perfect doctor', intervened where additional measures were needed most. Such intervention may be an optimal public health approach but it would violate this analysis' common-trends assumption.

Reliance on the common-trends assumption seems to be the fatal weakness of a difference-in-differences analysis. We do not expect populations to respond identically during an epidemic. Inevitably, some groups will be more affected by developments during the epidemic, either by chance or due to background socio-economic and health differences between groups

(Upshaw *et al.*, 2021). Attempts to divert resources to disproportionately affected areas violates the assumptions of the difference-in-differences analysis. Collins and Pinch (1998) describe science as being like the golem of Jewish mythology, a 'bumbling giant' that thoughtlessly follows instructions even when that causes harm. The difference-in-differences method is in danger of being such a golem. It is simple to perform, we are reassured that we do not need to account for confounding variables, and the analysis supplies an answer without complaint. But we do not know if that answer reflects a true causal effect. I have shown that incidence and prevalence data are unlikely to follow parallel trends in an outbreak or epidemic, and such outcomes should be avoided in difference-in-differences analyses of infection-prevention measures. The basic reproduction ratio is a more feasible target for a difference-in-differences analysis but the belief that this measure follows parallel trends remains an untestable assumption.

Natural immune boosting

A common admonition for mathematical modellers is that their models, 'Should be as simple as possible, but no simpler' (Keeling and Rohani, 2008). The art lies in finding this threshold. My model in Chapter 3 adds complexity to an SIRS epidemic model by incorporating natural immune boosting. This makes the model parameters unidentifiable from prevalence data. Outcomes could be consistent with a boosting coefficient of 0, meaning that no natural immune boosting occurs. In that case the added complexity of incorporating possible boosting has not improved the model. With different assumptions about immune waning, the boosting coefficient is estimated to be greater than 800. This value is implausibly extreme but it certainly suggests that boosting should not be ignored.

Adding immune boosting to the model has two notable effects on pre-

dicted disease dynamics. Firstly, it contributes to cyclic changes in prevalence. This could help to explain why respiratory infections have similar seasonal frequencies in locations with very different climates. Secondly, it affects estimates of infection-prevention measures' effectiveness and expectations of how disease dynamics will change when restrictions are lifted. When policy makers are weighing costs and benefits of an intervention, inaccurate predictions could lead to suboptimal infection-prevention strategies.

I further explored potential implications of natural immune boosting in Chapter 4. Immune boosting allows longer durations of immunity in the presence of a greater force of infection. I used simulations to explore how possible levels of waning and boosting could affect optimal timing of vaccination for healthcare workers. In simulations with hypothetical parameters, and in simulations with parameters fitted to data from NHS hospitals, greater levels of immune boosting led to benefits from vaccinating sooner. This could have practical implications for optimal timing of vaccination for a variety of seasonal infections, both for healthcare workers, people vulnerable to severe disease, and members of the wider community.

5.2 Implications for future research

Causal inference from observational data must, to some extent, substitute assumptions for the *magic of randomization* (see Collins *et al.*, 2020). It is therefore worth considering whether RCTs should be performed to answer these causal questions. Indeed, trials have been conducted for some physical interventions. Bundgaard *et al.* (2021) conducted a trial of mask use in Denmark with 6000 participants. Of those allocated to wear a mask, less than half reported adherence to this intervention. This led to large confidence intervals around the estimated effect of the intervention. It may be

that allocation to mask use in an RCT does not lead to a useful reduction in infection, but this does not necessarily translate into equivalent conclusions about the effect of consistently using a mask, or about the effect of regulations mandating mask use.

A trial comparing different booster vaccination schedules for healthcare workers might be able to determine the optimal schedule, provide insight into the extent of natural immune boosting, and test the feasibility of Hudgens and Halloran's (2008) two-stage trial design. The answers to these questions are intellectually interesting. But to justify the expense of a trial, we would need an expectation that the results could change practice. The date on which healthcare workers get vaccinated is likely to be a result of a variety of factors including availability, convenience and personal preference. Pilot work would be needed to determine whether a trial had the potential to change behaviour before its cost could be justified. While RCTs remain the canonical method for answering causal questions, they cannot be our sole means of informing public health decisions.

To conduct causal inference with observational data, a common assumption is that there is no unmeasured confounding. The assumption cannot be tested but for some analyses we might have confidence that unmeasured confounding is unlikely to affect the results. The strength of association between cigarette smoking and lung cancer is so strong that Hill (1965) asserted that any confounding variable that explained the association would need to be so closely linked to cigarette exposure that it would be easy to detect. He wrote: 'In such circumstances I think we are reasonably entitled to reject the vague contention of the armchair critic "you can't prove it, there *may* be such a feature" ' (*ibid.*). Mathur and VanderWeele (2020) quantified the magnitude of unmeasured confounding that would be needed to change the estimate of a causal effect.

In an outbreak or epidemic, there is considerable chance of confounding. When multiple infection-prevention measures are being introduced, it is plausible that other measures will have an effect at least as large as the intervention under study. To gain confidence that there is no unmeasured confounding, we need to have made every effort to identify confounders. Even in a difference-in-differences analysis, which accepts unmeasured confounders provided they do not violate the common-trends assumption, we need to understand the effects of any time-dependent variables. A tremendous effort went into data collection during the covid-19 pandemic. I used Hale *et al.*'s (2021) dataset in each chapter of this thesis, along with routinely collected data that were published by government agencies. ISARIC Clinical Characterization Group *et al.* (2022) collected individual-patient data for more than 700 thousand patients admitted to hospital with covid-19, demonstrating the possibility of collecting data for research purposes during a pandemic. Such data were very useful in describing clinical characteristics of the disease (*e.g.* Docherty *et al.*, 2020; ISARIC Clinical Characterization Group, 2021; Kartsonaki *et al.*, 2023; Swann *et al.*, 2020). Retrospective causal analyses can help to inform data collection for future epidemics to maximize the usefulness of collected data.

A crucial variable for estimates of disease dynamics is the number of susceptible individuals. In each chapter of this thesis I relied on assumptions about susceptibility. The time-series–SIR model can be used to estimate numbers susceptible to childhood infections that lead to lifelong immunity (Finkenstädt and Grenfell, 2000), but for microparasites with rapidly-waning immunity, estimation of the proportion susceptible is more difficult. Even the word 'susceptible' describes a simplified model of reality. People may be susceptible to infection but not to severe disease, or may be susceptible to infection following a large exposure but not following a small one (Alexander

et al., 2021; Halloran, 2019). This suggests a future avenue of research serially measuring markers of immunity. These data could be incorporated into a variety of models, from within-host models of immunity to population-wide models, to explore how immune waning and boosting interact with interventions against respiratory-tract infections.

5.3 Conclusion

Analyses in this thesis highlight potential problems with causal analyses that disregard mechanisms of disease transmission. When an analysis depends on assumptions, we should maximize the chance that such assumptions are accurate. While difference-in-differences analyses combined with mechanistic models may still give inaccurate results, they at least avoid errors caused by assuming implausible linear relationships between interventions and incidence of infections during epidemics. Similarly, assumptions about immunity that may be valid for understanding childhood infections with lifelong immunity may be too much of a simplification for long-term dynamics of respiratory viral infections with rapidly waning immunity. Better understanding of changing immunity over time is necessary to determine optimal approaches to reducing transmission of these respiratory viral infections.

Bibliography

Abadie A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *J Am Stat Assoc* **105**: 493–505. [85]

Adjodah D., K. Dinakar, M. Chinazzi, *et al.* (2021). Association between COVID-19 outcomes and mask mandates, adherence, and attitudes. *PLOS One* **16**: e0252315. [64]

Agoti C. N., J. R. Otieno, M. Ngama, *et al.* (2015). Successive respiratory syncytial virus epidemics in local populations arise from multiple variant introductions, providing insights into virus persistence. *J Virol* **89**: 11630–11642. [106]

Aiello A. E., and E. L. Larson (2002). What is the evidence for a causal link between hygiene and infections? *Lancet Infect Dis* **2**: 103–110. [2]

Akhtar Z., M. Götberg, D. Erlinge, *et al.* (2023). Optimal timing of influenza vaccination among patients with acute myocardial infarction – Findings from the IAMI trial. *Vaccine* **41**: 7159–7165. [134]

Akima H. (1970). A new method of interpolation and smooth curve fitting based on local procedures. *J ACM* **17**: 589–602. [66]

Alexander L. W., R. Ben-Shachar, L. C. Katzelnick, *et al.* (2021). Boosting can explain patterns of fluctuations of ratios of inapparent to symptomatic dengue virus infections. *Proc Natl Acad Sci U S A* **118**: e2013941118. [92, 145]

Alexander M., S. Moghadas, P. Rohani, and A. Summers (2006). Modelling the effect of a booster vaccination on disease epidemiology. *J Math Biol* **52**: 290–306. [41]

Anderson R. M., and R. M. May (1979). Population biology of infectious diseases: Part I. *Nature* **280**: 361–367. [3]

— (1982). Directly transmitted infectious diseases: Control by vaccination. *Science* **215**: 1053–1060. [3, 41, 98]

Anglemyer A., T. H. M. Moore, L. Parker, *et al.* (2020). Digital contact tracing technologies in epidemics: a rapid review. *Cochrane Database Syst Rev*, CD013699. [11]

Arnold A., G. Bickler, and T. S. Harrison (2019). The first 5 years of Part 2A Orders: the use of powers from court applications to protect public health in England 2010–15. *J Public Health* **41**: 27–35. [2]

Ascough S., S. Paterson, and C. Chiu (2018). Induction and subversion of human protective immunity: contrasting influenza and respiratory syncytial virus. *Front Immunol* **9**: 323. [108]

Baden L. R., H. M. El Sahly, B. Essink, *et al.* (2021). Efficacy and safety of the mRNA-1273 SARS-CoV-2 vaccine. *N Engl J Med* **384**: 403–416. [113]

Baker R. E., S. W. Park, W. Yang, G. A. Vecchi, C. J. E. Metcalf, and B. T. Grenfell (2020). The impact of COVID-19 nonpharmaceutical interventions on the future dynamics of endemic infections. *Proc Natl Acad Sci U S A* **117**: 30547–30553. [109]

Baker R. E., C. M. Saad-Roy, S. W. Park, J. Farrar, C. J. E. Metcalf, and B. T. Grenfell (2022). Long-term benefits of nonpharmaceutical interventions for endemic infections are shaped by respiratory pathogen dynamics. *Proc Natl Acad Sci U S A* **119**: e2208895119. [109]

Bakker K. M., M. C. Eisenberg, R. J. Woods, and M. E. Martinez (2022). Identifying optimal vaccination scenarios to reduce varicella zoster virus transmission and reactivation. *BMC Med* **20**: 387. [91]

Baraniuk C. (2021). Covid-19: how the UK vaccine rollout delivered success, so far. *BMJ* **372**: n421. [14]

Barbarossa M. V., M. Polner, and G. Röst (2017). Stability switches induced by immune system boosting in an SIRS model with discrete and distributed delays. *SIAM J Appl Math* **77**: 905–923. [93, 94]

Barbarossa M. V., and G. Röst (2015a). Immuno-epidemiology of a population structured by immune status: a mathematical study of waning immunity and immune system boosting. *J Math Biol* **71**: 1737–1770. [93]

— (2015b). Mathematical models for vaccination, waning immunity and immune system boosting: a general framework. *arXiv*. eprint: 1501.03451. [93]

Bardsley M., R. A. Morbey, H. E. Hughes, *et al.* (2023). Epidemiology of respiratory syncytial virus in children younger than 5 years in England during the COVID-19 pandemic, measured by laboratory, clinical, and syndromic surveillance: a retrospective observational study. *Lancet Infect Dis* **23**: 56–66. [89]

Batteux E., S. Bonfield, L. F. Jones, *et al.* (2022). Impact of residual risk messaging to reduce false reassurance following test-negative results from asymptomatic coronavirus (SARS-CoV-2) testing: an online experimental study of a hypothetical test. *BMJ Open* **12**: e056533. [141]

Baynes C. (2020). *Coronavirus: Face masks could increase risk of infection, medical chief warns*. URL: https://www.independent.co.uk/news/health/coronavirus-news-face-masks-increase-risk-infection-doctor-jenny-harries-a9396811.html (visited on 16 December 2024). [141]

BBC News (2020a). *Coronavirus: First case confirmed in Scotland*. URL: https://www.bbc.co.uk/news/uk-scotland-51698366 (visited on 21 October 2024). [12]

BBC News (2020b). *Coronavirus: First Welsh case among three new UK diagnoses*. URL: https://www.bbc.co.uk/news/uk-51673068 (visited on 21 October 2024). [12]

— (2020c). *Covid-19 vaccine: First person receives Pfizer jab in UK*. URL: https://www.bbc.co.uk/news/uk-55227325 (visited on 5 December 2024). [113]

— (2020d). *Northern Ireland diagnoses first coronavirus case*. URL: https://www.bbc.co.uk/news/uk-51667483 (visited on 21 October 2024). [12]

Bedston S., A. Akbari, C. I. Jarvis, *et al.* (2022). COVID-19 vaccine uptake, effectiveness, and waning in 82,959 health care workers: A national prospective cohort study in Wales. *Vaccine* **40**: 1180–1189. [114]

Begon M., M. Bennett, R. G. Bowers, N. P. French, S. M. Hazel, and J. Turner (2002). A clarification of transmission terms in host-microparasite models: numbers, densities and areas. *Epidemiol Infect* **129**: 147–153. [34]

Bezanson J., A. Edelman, S. Karpinski, and V. B. Shah (2017). Julia: a fresh approach to numerical computing. *SIAM Rev* **59**: 65–98. [44, 53, 97]

Blackwood J. C., D. A. Cummings, H. Broutin, S. Iamsirithaworn, and P. Rohani (2013). Deciphering the impacts of vaccination and immunity on pertussis epidemiology in Thailand. *Proc Natl Acad Sci U S A* **110**: 9595–9600. [93]

Bloom-Feshbach K., W. J. Alonso, V. Charu, *et al.* (2013). Latitudinal variations in seasonal activity of influenza and respiratory syncytial virus (RSV): a global comparative review. *PLOS One* **8**: e54445. [89]

Bollyky T. J., E. Castro, A. Y. Aravkin, *et al.* (2023). Assessing COVID-19 pandemic policies and behaviours and their economic and educational trade-offs across US states from Jan 1, 2020, to July 31, 2022: an observational analysis. *Lancet* **401**: 1341–1360. [54, 55]

Bootsma M. C., and N. M. Ferguson (2007). The effect of public health measures on the 1918 influenza pandemic in U.S. cities. *Proc Natl Acad Sci U S A* **104**: 7588–7593. [2]

Box G. E. P. (1976). Science and statistics. *J Am Stat Assoc* **71**: 791–799. [30, 31]

Bregman D. J., and A. D. Langmuir (1990). Farr's law applied to AIDS projections. *JAMA* **263**: 1522–1525. [32, 33]

Brisson M., G. Melkonyan, M. Drolet, G. De Serres, R. Thibeault, and P. De Wals (2010). Modeling the impact of one-and two-dose varicella vaccination on the epidemiology of varicella and zoster. *Vaccine* **28**: 3385–3397. [91]

Brookmeyer R. (1996). AIDS, epidemics, and statistics. *Biometrics*, 781–796. [32]

Brown J., and E. Kirk-Wade (2021). *Coronavirus: A history of 'Lockdown laws' in England*. London, UK: House of Commons Library. [14, 58, 59]

Brownlee J. (1915). Historical note on Farr's theory of the epidemic. *Br Med J* **ii**: 250–252. [32]

Buckee C., A. Noor, and L. Sattenspiel (2021). Thinking clearly about social aspects of infectious disease transmission. *Nature* **595**: 205–213. [55]

Bundgaard H., J. S. Bundgaard, D. E. T. Raaschou-Pedersen, *et al.* (2021). Effectiveness of adding a mask recommendation to other public health measures to prevent SARS-CoV-2 infection in Danish mask wearers: a randomized controlled trial. *Ann Intern Med* **174**: 335–343. [143]

Burnsa J., A. Movsisyana, J. M. Stratil, *et al.* (2021). International travel-related control measures to contain the COVID-19 pandemic: a rapid review. *Cochrane Database Syst Rev*, CD013717. [11]

Butcher J. (1996). A history of Runge–Kutta methods. *Appl Numer Math* **20**: 247–260. [44]

— (2016). *Numerical methods for ordinary differential equations*. 3rd edition. Chichester, United Kingdom: Wiley. [44]

Cabinet Office (2021). *COVID-19 Response - Spring 2021 (Summary)*. url: `https://www.gov.uk/government/publications/covid-19-response-spring-2021/covid-19-response-spring-2021-summary` (visited on 29 October 2024). [14]

Callaway B., and T. Li (2023). Policy evaluation during a pandemic. *J Econom* **236**: 105454. [57]

Callaway B., and P. H. C. Sant'Anna (2021). Difference-in-differences with multiple time periods. *J Econom* **225**: 200–230. [72, 83]

Caniglia E. C., and E. J. Murray (2020). Difference-in-difference in the time of cholera: a gentle introduction for epidemiologists. *Curr Epidemiol Rep* **7**: 203–211. [56, 64]

Carneiro I., A. Roca-Feltrer, J. T. Griffin, *et al.* (2010). Age-patterns of malaria vary with severity, transmission intensity and seasonality in sub-Saharan Africa: a systematic review and pooled analysis. *PLOS One* **5**: e8988. [88]

Castioni P., S. Gómez, C. Granell, and A. Arenas (2024). Rebound in epidemic control: How misaligned vaccination timing amplifies infection peaks. *NPJ Complex* **1**: 20. [133, 134]

Castro M., and R. J. de Boer (2020). Testing structural identifiability by a simple scaling method. *PLOS Comput Biol* **16**: e1008248. [47, 101, 172]

Chakraborty M., and P. L. Bhukya (2023). Human Bocavirus. In: P. L. Bhukya, S. T. Mhaske, and S. C. Sonkar, editors. *Emerging human viral diseases, Volume I: Respiratory and haemorrhagic fever*. Singapore: Springer, pp. 227–245. [7, 9]

Champredon D., J. Dushoff, and D. J. Earn (2018). Equivalence of the Erlang-distributed SEIR epidemic model and the renewal equation. *SIAM J Appl Math* **78**: 3258–3278. [46, 61]

Chu D. K., E. A. Akl, S. Duda, *et al.* (2020). Physical distancing, face masks, and eye protection to prevent person-to-person transmission of SARS-CoV-2 and COVID-19: a systematic review and meta-analysis. *Lancet* **395**: 1973–1987. [9]

Clark T. A. (2014). Changing pertussis epidemiology: everything old is new again. *J Infect Dis* **209**: 978–981. [88]

Coleman M. P. (1986). A plague epidemic in voluntary quarantine. *Int J Epidemiol* **15**: 379–385. [1]

Collins H., and T. Pinch (1998). *The golem*. 2nd Edition. Cambridge, United Kingdom: Cambridge University Press. [142]

Collins R., L. Bowman, M. Landray, and R. Peto (2020). The magic of randomization versus the myth of real-world evidence. *N Engl J Med* **382**: 674–678. [3, 143]

Cooper B. S., S. Evans, Y. Jafari, *et al.* (2023). The burden and dynamics of hospital-acquired SARS-CoV-2 in England. *Nature* **623**: 132–138. [119, 125, 133, 137]

Cori A., N. M. Ferguson, C. Fraser, and S. Cauchemez (2013). A new framework and software to estimate time-varying reproduction numbers during epidemics. *Am J Epidemiol* **178**: 1505–1512. [61]

Cori A., Z. N. Kamvar, J. Stockwin, *et al.* (2022). *EpiEstim v2.2-4: A tool to estimate time varying instantaneous reproduction number during epidemics*. URL: `https://github.com/mrc-ide/EpiEstim` (visited on 2 October 2024). [61, 86, 174]

Costantino V., M. Trent, and C. R. MacIntyre (2019). Modelling of optimal timing for influenza vaccination as a function of intraseasonal waning of immunity and vaccine coverage. *Vaccine* **37**: 6768–6775. [135]

COVID-19 Forecasting Team (2023). Past SARS-CoV-2 infection protection against re-infection: a systematic review and meta-analysis. *Lancet* **401**: 833–842. [108]

COVID-19 National Preparedness Collaborators (2022). Pandemic preparedness and COVID-19: an exploratory analysis of infection and fatality rates, and contextual factors associated with preparedness in 177 countries, from Jan 1, 2020, to Sept 30, 2021. *Lancet* **399**: 1489–1512. [54]

Cox D. R. (1986). Statistics and causal inference: Comment. *J Am Stat Assoc* **81**: 963–964. [27]

Dafilis M. P., F. Frascoli, J. G. Wood, and J. M. McCaw (2012). The influence of increasing life expectancy on the dynamics of SIRS systems with immune boosting. *ANZIAM J* **54**: 50–63. [93, 94, 106]

Dafilis M. P., F. Frascoli, J. McVernon, J. M. Heffernan, and J. M. McCaw (2014a). Dynamical crises, multistability and the influence of the duration of immunity in a seasonally-forced model of disease transmission. *Theor Biol Med Model* **11**: 1–10. [93]

— (2014b). The dynamical consequences of seasonal forcing, immune boosting and demographic change in a model of disease transmission. *J Theor Biol* **361**: 124–132. [93, 94]

Dalziel S. R., L. Haskell, S. O'Brien, *et al.* (2022). Bronchiolitis. *Lancet* **400**: 392–406. [89]

Dancer S. J. (2016). Infection control: evidence-based common sense. *Infect Dis Health* **21**: 147–153. [2, 4]

Danisch S., and J. Krumbiegel (2021). Makie.jl: flexible high-performance data visualization for Julia. *J Open Source Softw* **6**: 3349. [53]

Darakjy S., J. E. Brady, C. J. DiMaggio, and G. Li (2014). Applying Farr's Law to project the drug overdose mortality epidemic in the United States. *Inj Epidemiol* **1**: 31. [32]

Datseris G., J. Isensee, S. Pech, and T. Gál (2020). DrWatson: the perfect sidekick for your scientific inquiries. *J Open Source Softw* **5**: 2673. [53]

Dattani S. (2023). What were the death tolls from pandemics in history? *Our World in Data*. URL: https://ourworldindata.org/historical-pandemics. [8]

De Angelis M. L., F. Francescangeli, R. Rossi, A. Giuliani, R. De Maria, and A. Zeuner (2021). Repeated exposure to subinfectious doses of SARS-CoV-2 may promote T cell immunity and protection against severe COVID-19. *Viruses* **13**: 961. [88]

Department of Health and Social Care (2020a). *Formal tiering review update: 30 December 2020*. UK Government. URL: https://www.gov.uk/government/news/formal-tiering-review-update-30-december-2020 (visited on 26 November 2024). [59]

— (2020b). *Weekly statistics for NHS Test and Trace (England) and coronavirus testing (UK): 30 July to 5 August*. UK Government. URL: https://www.gov.uk/government/publications/nhs-test-and-trace-england-and-coronavirus-testing-uk-statistics-30-july-to-5-august-2020/weekly-statistics-for-nhs-test-and-trace-england-and-coronavirus-testing-uk-30-july-to-5-august (visited on 26 November 2024). [58, 140]

— (2021). *Coronavirus cases in England: 5 January 2021*. UK Government. URL: https://www.gov.uk/government/publications/coronavirus-cases-in-england-5-january-2021 (visited on 11 July 2024). [59]

Diekmann O., J. Heesterbeek, and M. Roberts (2010). The construction of next-generation matrices for compartmental epidemic models. *J R Soc Interface* **7**: 873–885. [41]

Dietz K. (1993). The estimation of the basic reproduction number for infectious diseases. *Stat Methods Med Res* **2**: 23–41. [40, 41]

Docherty A. B., E. M. Harrison, C. A. Green, *et al.* (2020). Features of 20 133 UK patients in hospital with covid-19 using the ISARIC WHO Clinical Characterisation Protocol: prospective observational cohort study. *BMJ* **369**: m1985. [145]

Domnich A., A. Orsi, A. Signori, *et al.* (2024). Waning intra-season vaccine effectiveness against influenza A (H3N2) underlines the need for more durable protection. *Expert Rev Vaccines* **23**: 380–388. [88, 112]

Dowd B. E. (2011). Separated at birth: statisticians, social scientists, and causality in health services research. *Health Serv Res* **46**: 397–420. [4]

Dushoff J., J. B. Plotkin, S. A. Levin, and D. J. Earn (2004). Dynamical resonance can account for seasonality of influenza epidemics. *Proc Natl Acad Sci U S A* **101**: 16915–16916. [106, 141]

Earn D. J., P. Rohani, and B. T. Grenfell (1998). Persistence, chaos and synchrony in ecology and epidemiology. *Proc R Soc Lond B* **265**: 7–10. [105, 106]

Earn D. J., P. Rohani, B. M. Bolker, and B. T. Grenfell (2000). A simple model for complex dynamical transitions in epidemics. *Science* **287**: 667–670. [106]

Falsey A. R., P. A. Hennessey, M. A. Formica, C. Cox, and E. E. Walsh (2005). Respiratory syncytial virus infection in elderly and high-risk adults. *N Engl J Med* **352**: 1749–1759. [89]

Farmer P. (2020). *Fevers, feuds, and diamonds: Ebola and the ravages of history*. New York: Farrar, Straus and Giroux. [2]

Farr W. (1840). Causes of death in England and Wales: Letter to the Registrar-General from William Farr, Esq. In: *Second annual report of the Registrar-General of births, deaths and marriages in England*. London, United Kingdom: W. Clowes and sons, pp. 69–98. [31, 32]

Feng S., and A. Bilinski (2024). Parallel trends in an unparalleled pandemic: Difference-in-differences for infectious disease policy evaluation. *medRxiv*, 2024.04.08.24305335. [56, 57, 79]

Ferman B., C. Pinto, and V. Possebom (2020). Cherry picking with synthetic controls. *J Policy Anal Manag* **39**: 510–532. [85]

Fine P. E. M. (1979). John Brownlee and the measurement of infectiousness: an historical study in epidemic theory. *J R Stat Soc Ser A Stat Soc* **142**: 347–362. [32]

— (2003). The interval between successive cases of an infectious disease. *Am J Epidemiol* **158**: 1039–1047. [39]

Finkenstädt B. F., and B. T. Grenfell (2000). Time series modelling of childhood diseases: a dynamical systems approach. *J R Stat Soc C Appl Stat* **49**: 187–205. [141, 145]

Flaxman S., S. Mishra, A. Gandy, *et al.* (2020). Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. *Nature* **584**: 257–261. [55, 84]

Foley D. A., D. K. Yeoh, C. A. Minney-Smith, *et al.* (2021). The interseasonal resurgence of respiratory syncytial virus in Australian children following the reduction of coronavirus disease 2019–related public health measures. *Clin Infect Dis* **73**: e2829–e2830. [89]

Fonville J. M., S. Wilks, S. L. James, *et al.* (2014). Antibody landscapes after influenza virus infection or vaccination. *Science* **346**: 996–1000. [108]

Fooks A. R., F. Cliquet, S. Finke, *et al.* (2017). Rabies. *Nat Rev Dis Primers* **3**: 17091. [4]

Freeman M. C., M. E. Stocks, O. Cumming, *et al.* (2014). Hygiene and health: systematic review of handwashing practices worldwide and update of health effects. *Trop Med Int Health* **19**: 906–916. [2]

Friedrich F., R. Ongaratto, M. C. Scotta, *et al.* (2021). Early impact of social distancing in response to coronavirus disease 2019 on hospitaliza-

tions for acute bronchiolitis in infants in Brazil. *Clin Infect Dis* **72**: 2071–2075. [10, 89]

Frigo M., and S. G. Johnson (2005). The design and implementation of FFTW3. *Proc IEEE* **93**: 216–231. [98]

Gail M. H., and R. Brookmeyer (1990). Projecting the incidence of AIDS. *JAMA* **263**: 1538–1539. [32]

Galanti M., R. Birger, M. Ud-Dean, *et al.* (2019). Rates of asymptomatic respiratory virus infection across age groups. *Epidemiol Infect* **147**: e176. [7]

GBD 2013 Mortality and Causes of Death Collaborators (2015). Global, regional, and national age–sex specific all-cause and cause-specific mortality for 240 causes of death, 1990–2013: a systematic analysis for the Global Burden of Disease Study 2013. *Lancet* **385**: 117–171. [1]

GBD 2016 Lower Respiratory Infections Collaborators (2018). Estimates of the global, regional, and national morbidity, mortality, and aetiologies of lower respiratory infections in 195 countries, 1990–2016: a systematic analysis for the Global Burden of Disease Study 2016. *Lancet Infect Dis* **18**: 1191–1210. [8]

GBD 2021 Diseases and Injuries Collaborators (2024). Global incidence, prevalence, years lived with disability (YLDs), disability-adjusted life-years (DALYs), and healthy life expectancy (HALE) for 371 diseases and injuries in 204 countries and territories and 811 subnational locations, 1990–2021: a systematic analysis for the Global Burden of Disease Study 2021. *Lancet* **403**: 2133–2161. [8]

Ge H., K. Xu, and Z. Ghahramani (2018). Turing: a language for flexible probabilistic inference. *Proc Mach Learn Res* **84**: 1682–1690. [65]

Gershon A. A., J. Breuer, J. I. Cohen, *et al.* (2015). Varicella zoster virus infection. *Nat Rev Dis Primers* **1**: 15016. [8, 87]

Gillespie D. T. (1977). Exact stochastic simulation of coupled chemical reactions. *J Phys Chem* **81**: 2340–2361. [52, 61]

Glass K., and B. Grenfell (2003). Antibody dynamics in childhood diseases: waning and boosting of immunity and the impact of vaccination. *J Theor Biol* **221**: 121–131. [92]

Glezen W. P., L. H. Taber, A. L. Frank, and J. A. Kasel (1986). Risk of primary infection and reinfection with respiratory syncytial virus. *Am J Dis Child* **140**: 543–546. [89]

Godfrey K. (1983). *Compartmental models and their application*. London, United Kingdom: Academic Press. [34]

Gomes M. G. M., L. J. White, and G. F. Medley (2004). Infection, reinfection, and vaccination under suboptimal immune protection: epidemiological perspectives. *J Theor Biol* **228**: 539–549. [108]

Green M. A., M. García-Fiñana, B. Barr, *et al.* (2021). Evaluating social and spatial inequalities of large scale rapid lateral flow SARS-CoV-2 antigen testing in COVID-19 management: An observational study of Liverpool,

UK (November 2020 to January 2021). *Lancet Reg Health Eur* **6**: 100107. [59]

Greenland S., and J. M. Robins (1986). Identifiability, exchangeability, and epidemiological confounding. *Int J Epidemiol* **15**: 413–419. [20, 23]

Greenland S., J. Pearl, and J. M. Robins (1999). Causal diagrams for epidemiologic research. *Epidemiology* **10**: 37–48. [21–23, 117]

Grenfell B. T., O. N. Bjørnstad, and J. Kappey (2001). Travelling waves and spatial hierarchies in measles epidemics. *Nature* **414**: 716–723. [106]

Grint D. J., K. Wing, E. Williamson, *et al.* (2021). Case fatality risk of the SARS-CoV-2 variant of concern B.1.1.7 in England, 16 November to 5 February. *Eurosurveillance* **26**: 2100256. [61]

Guglielmi G. (2021). Rapid coronavirus tests: a guide for the perplexed. *Nature* **590**: 202–205. [81, 141]

Guyatt G. H., D. L. Sackett, J. C. Sinclair, *et al.* (1995). Users' guides to the medical literature: IX. A method for grading health care recommendations. *JAMA* **274**: 1800–1804. [3]

Guyatt G. H., A. D. Oxman, R. Kunz, G. E. Vist, Y. Falck-Ytter, and H. J. Schünemann (2008). What is "quality of evidence" and why is it important to clinicians? *BMJ* **336**: 995–998. [3]

Hale T., N. Angrist, R. Goldszmidt, *et al.* (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nat Hum Behav* **5**: 529–538. [2, 12–17, 54, 59, 60, 85, 90, 137, 145]

Hall C. B., E. E. Walsh, C. E. Long, and K. C. Schnabel (1991). Immunity to and frequency of reinfection with respiratory syncytial virus. *J Infect Dis* **163**: 693–698. [89]

Hall V. J., S. Foulkes, A. Saei, *et al.* (2021). COVID-19 vaccine coverage in health-care workers in England and effectiveness of BNT162b2 mRNA vaccine against infection (SIREN): a prospective, multicentre, cohort study. *Lancet* **397**: 1725–1735. [111, 113, 123]

Hall V. J., S. Foulkes, F. Insalata, *et al.* (2022). Protection against SARS-CoV-2 after Covid-19 vaccination and previous infection. *N Engl J Med* **386**: 1207–1220. [113]

Hall V. J., F. Insalata, S. Foulkes, *et al.* (2024). Effectiveness of BNT162b2 mRNA vaccine third doses and previous infection in protecting against SARS-CoV-2 infections during the Delta and Omicron variant waves; the UK SIREN cohort study September 2021 to February 2022. *J Infect* **88**: 30–40. [111, 113, 136]

Halloran M. E. (2019). Analysis of vaccine studies and causal inference. In: L. Held, N. Hens, P. O'Neill, and J. Wallinga, editors. *Handbook of infectious disease data analysis*. New York, New York, United States: Chapman and Hall/CRC, pp. 129–156. [3, 29, 146]

Halloran M. E., and C. J. Struchiner (1995). Causal inference in infectious diseases. *Epidemiology*, 142–151. [29]

Hatoun J., E. T. Correa, S. M. A. Donahue, and L. Vernacchio (2020). Social distancing for COVID-19 and diagnoses of other infectious diseases in children. *Pediatrics* **146**: [10]

Hawker J., N. Begg, R. Reintjes, K. Ekdahl, O. Edeghere, and J. van Steenbergen (2019). *Communicable disease control and health protection handbook*. 4th edition. Hoboken, New Jersey, United States: Wiley–Blackwell. [2]

Heesterbeek J. A. P. (2002). A brief history of R_0 and a recipe for its calculation. *Acta Biotheor* **50**: 189–204. [40]

Heffernan J. M., and M. J. Keeling (2009). Implications of vaccination and waning immunity. *Proc R Soc B* **276**: 2071–2080. [92, 94]

Hernán M. A. (2016). Does water kill? A call for less casual causal inferences. *Ann Epidemiol* **26**: 674–680. [27]

— (2018). The C-word: Scientific euphemisms do not improve causal inference from observational data. *Am J Public Health* **108**: 616–619. [5]

Hernán M. A., and S. Hernández-Díaz (2012). Beyond the intention-to-treat in comparative effectiveness research. *Clin Trials* **9**: 48–55. [27]

Hernán M. A., and J. M. Robins (2020). *Causal inference: What if*. Boca Raton, Florida: CRC Press. [18, 24, 25]

Hernán M. A., and S. L. Taubman (2008). Does obesity shorten life? The importance of well-defined interventions to answer causal questions. *Int J Obes* **32**: S8–S14. [18, 26, 27]

Hernán M. A., A. Alonso, R. Logan, *et al.* (2008). Observational studies analyzed like randomized experiments: an application to postmenopausal hormone therapy and coronary heart disease. *Epidemiology* **19**: 766–779. [118]

Hill A. B. (1965). The environment and disease: Association or causation? *Proc R Soc Med* **58**: 295–300. [144]

HIV-CAUSAL Collaboration (2011). When to initiate combined antiretroviral therapy to reduce mortality and AIDS-defining illness in HIV-infected persons in developed countries: an observational study. *Ann Intern Med* **154**: 509–515. [26]

Hodgson D., R. Pebody, J. Panovska-Griffiths, M. Baguelin, and K. E. Atkins (2020). Evaluating the next generation of RSV intervention strategies: a mathematical modelling study and cost-effectiveness analysis. *BMC Med* **18**: 1–14. [107]

Hodgson D., M. Koltai, F. Krauer, S. Flasche, M. Jit, and K. E. Atkins (2022). Optimal respiratory syncytial virus intervention programmes using nirsevimab in England and Wales. *Vaccine* **40**: 7151–7157. [107]

Hoffmann T. C., C. Erueti, and P. P. Glasziou (2013). Poor description of non-pharmacological interventions: analysis of consecutive sample of randomised trials. *BMJ* **347**: f3755. [4]

Holland P. W. (1986). Statistics and causal inference. *J Am Stat Assoc* **81**: 945–960. [17, 19, 20]

Holmdahl I., S. J. Bents, R. E. Baker, *et al.* (2024). Differential impact of COVID-19 non-pharmaceutical interventions on the epidemiological dynamics of respiratory syncytial virus subtypes A and B. *Sci Rep* **14**: 14527. [108]

Hope-Simpson R. E. (1965). The nature of herpes zoster: a long-term study and a new hypothesis. *Proc R Soc Med* **58**: 9–20. [87]

Hosseini P. R., A. A. Dhondt, and A. Dobson (2004). Seasonality and wildlife disease: how seasonal birth, aggregation and variation in immunity affect the dynamics of *Mycoplasma gallisepticum* in house finches. *Proc R Soc Lond B* **271**: 2569–2577. [106]

Hoyer-Leitzel A., S. Iams, A. Haslam-Hyde, M. L. Zeeman, and N. Fefferman (2023). An immuno-epidemiological model for transient immune protection: A case study for viral respiratory infections. *Infect Dis Model* **8**: 855–864. [93, 94]

Huang C., Y. Wang, X. Li, *et al.* (2020). Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. *Lancet* **395**: 497–506. [12]

Hudgens M. G., and M. E. Halloran (2008). Toward causal inference with interference. *J Am Stat Assoc* **103**: 832–842. [29, 144]

Hwang S. E., J. H. Chang, B. Oh, and J. Heo (2021). Possible aerosol transmission of COVID-19 associated with an outbreak in an apartment in Seoul, South Korea, 2020. *Int J Infect Dis* **104**: 73–76. [10]

ISARIC Clinical Characterization Group (2021). COVID-19 symptoms at hospital admission vary with age and sex: results from the ISARIC prospective multinational observational study. *Infection* **49**: 889–905. [145]

ISARIC Clinical Characterization Group, E. Garcia-Gallo, L. Merson, *et al.* (2022). ISARIC-COVID-19 dataset: a prospective, standardized, global dataset of patients hospitalized with COVID-19. *Sci Data* **9**: 454. [145]

Jefferson T., L. Dooley, E. Ferroni, *et al.* (2023). Physical interventions to interrupt or reduce the spread of respiratory viruses. *Cochrane Database Syst Rev*, CD006207. [4, 10–12]

Johnson S. R., G. A. Tomlinson, G. A. Hawker, J. T. Granton, and B. M. Feldman (2010). Methods to elicit beliefs for Bayesian priors: a systematic review. *J Clin Epidemiol* **63**: 355–369. [48]

Joint Committee on Vaccination and Immunisation (2023). *JCVI statement on a childhood varicella (chickenpox) vaccination programme*. UK Government. URL: `https://www.gov.uk/government/publications/childhood-varicella-vaccination-programme-jcvi-advice-14-november-2023/jcvi-statement-on-a-childhood-varicella-chickenpox-vaccination-programme` (visited on 30 November 2024). [91]

Kahn-Lang A., and K. Lang (2020). The promise and pitfalls of differences-in-differences: Reflections on *16 and Pregnant* and other applications. *J Bus Econ Stat* **38**: 613–620. [73]

Kartsonaki C., J. K. Baillie, N. G. Barrio, *et al.* (2023). Characteristics and outcomes of an international cohort of 600 000 hospitalized patients with COVID-19. *Int J Epidemiol* **52**: 355–376. [145]

Keeling M. J., and P. Rohani (2008). *Modeling infectious diseases in humans and animals*. Princeton, New Jersey, United States: Princeton University Press. [34, 45, 142]

Kenah E., M. Lipsitch, and J. M. Robins (2008). Generation interval contraction and epidemic data analysis. *Math Biosci* **213**: 71–79. [39]

Kendall M., D. Tsallis, C. Wymant, *et al.* (2023). Epidemiological impacts of the NHS COVID-19 app in England and Wales throughout its first year. *Nat Commun* **14**: 858. [11]

Kenyon C. (2020). Flattening-the-curve associated with reduced COVID-19 case fatality rates–an ecological analysis of 65 countries. *J Infect* **81**: e98–e99. [106]

Kermack W. O., and A. G. McKendrick (1927). A contribution to the mathematical theory of epidemics. *Proc R Soc A* **115**: 700–721. [35]

— (1932). Contributions to the mathematical theory of epidemics. II.—The problem of endemicity. *Proc R Soc A* **138**: 55–83. [35]

— (1933). Contributions to the mathematical theory of epidemics. III.—Further studies of the problem of endemicity. *Proc R Soc A* **141**: 94–122. [35]

Kilgore P. E., A. M. Salim, M. J. Zervos, and H.-J. Schmitt (2016). Pertussis: microbiology, disease, treatment, and prevention. *Clin Microbiol Rev* **29**: 449–486. [3]

Korber B., W. M. Fischer, S. Gnanakaran, *et al.* (2020). Tracking changes in SARS-CoV-2 spike: evidence that D614G increases infectivity of the COVID-19 virus. *Cell* **182**: 812–827. [61]

Korostil I. A., J. G. Wood, and D. G. Regan (2015). Periodicity of varicella-zoster virus in the presence of immune boosting and clinical reinfection with varicella. *Theor Biol Med Model* **12**: 6. [91, 94]

Kratzer S., L. M. Pfadenhauer, R. L. Biallas, *et al.* (2022). Unintended consequences of measures implemented in the school setting to contain the COVID-19 pandemic: a scoping review. *Cochrane Database Syst Rev*, CD015397. [11, 12]

Krylova O., and D. J. Earn (2013). Effects of the infectious period distribution on predicted transitions in childhood disease dynamics. *J R Soc Interface* **10**: 20130098. [40]

Kutter J. S., M. I. Spronken, P. L. Fraaij, R. A. M. Fouchier, and S. Herfst (2018). Transmission routes of respiratory viruses among humans. *Curr Opin Virol* **28**: 142–151. [8, 9]

Lal H., A. L. Cunningham, O. Godeaux, *et al.* (2015). Efficacy of an adjuvanted herpes zoster subunit vaccine in older adults. *N Engl J Med* **372**: 2087–2096. [88]

Lang J. C. (2022). Use of mathematical modelling to assess respiratory syncytial virus epidemiology and interventions: a literature review. *J Math Biol* **84**: 26. [93, 108]

Larson E. L., and C. T. Liverman, editors (2011). *Preventing transmission of pandemic influenza and other viral respiratory diseases: Personal protective equipment for healthcare personnel*. Washington, District of Columbia, United States: National Academies Press. [10]

Lavine J. S., A. A. King, and O. N. Bjørnstad (2011). Natural immune boosting in pertussis dynamics and the potential for long-term vaccine failure. *Proc Natl Acad Sci U S A* **108**: 7259–7264. [93, 94]

Lavine J. S., A. A. King, V. Andreasen, and O. N. Bjørnstad (2013). Immune boosting explains regime-shifts in prevaccine-era pertussis dynamics. *PLOS One* **8**: e72086. [93, 94]

Lawlor D. A., K. Tilling, and G. Davey Smith (2016). Triangulation in aetiological epidemiology. *Int J Epidemiol* **45**: 1866–1886. [141]

Lee E. K., Z. L. Li, Y. K. Liu, and J. LeDuc (2021). Strategies for vaccine prioritization and mass dispensing. *Vaccines* **9**: 506. [133]

Leung K., M. H. Shum, G. M. Leung, T. T. Lam, and J. T. Wu (2021). Early transmissibility assessment of the N501Y mutant strains of SARS-CoV-2 in the United Kingdom, October to November 2020. *Eurosurveillance* **26**: 2002106. [61]

Leung N. H. L. (2021). Transmissibility and transmission of respiratory viruses. *Nat Rev Microbiol* **19**: 528–545. [8, 9]

Leung T., P. T. Campbell, B. D. Hughes, F. Frascoli, and J. M. McCaw (2018). Infection-acquired versus vaccine-acquired immunity in an SIRWS model. *Infect Dis Model* **3**: 118–135. [93, 94]

Li Y., X. Wang, D. M. Blau, *et al.* (2022). Global, regional, and national disease burden estimates of acute lower respiratory infections due to respiratory syncytial virus in children younger than 5 years in 2019: a systematic analysis. *Lancet* **399**: 2047–2064. [8, 107]

Liu Y., C. Diamond, S. Abbott, *et al.* (2024). The impact of public health and social measures (PHSMs) on SARS-CoV-2 transmission in the WHO European Region (2020–2022). *Influenza Other Respir Viruses* **18**: e70036. [55, 84]

Livingston E. H. (2020). Study design and statistics. In: AMA Manual of Style Committee, editor. *AMA manual of style: A guide for authors and editors*. 11th Edition. New York, New York, United States: Oxford University Press, pp. 977–1096. [4]

Lloyd A. L. (2001). Destabilization of epidemic models with the inclusion of realistic distributions of infectious periods. *Proc R Soc Lond B* **268**: 985–993. [37]

Lydon E. C., R. Henao, T. W. Burke, *et al.* (2019). Validation of a host response test to distinguish bacterial and viral respiratory infection. *EBioMedicine* **48**: 453–461. [7]

MacIntyre C. R., P.-Y. Nguyen, A. A. Chughtai, *et al.* (2021). Mask use, risk-mitigation behaviours and pandemic fatigue during the COVID-19 pandemic in five cities in Australia, the UK and USA: A cross-sectional survey. *Int J Infect Dis* **106**: 199–207. [83]

Mackenbach J. P. (2021). The rise and fall of diseases: reflections on the history of population health in Europe since ca. 1700. *Eur J Epidemiol* **36**: 1199–1205. [1]

Mathur M. B., and T. J. VanderWeele (2020). Sensitivity analysis for unmeasured confounding in meta-analyses. *J Am Stat Assoc* **115**: 163–172. [144]

Matz M., C. Allemani, M. van Tongeren, *et al.* (2022). Excess mortality among essential workers in England and Wales during the COVID-19 pandemic. *J Epidemiol Community Health* **76**: 660–666. [111]

May R. M., and R. M. Anderson (1979). Population biology of infectious diseases: Part II. *Nature* **280**: 455–461. [3]

McElreath R. (2020). *Statistical rethinking: A Bayesian course with examples in R and Stan*. 2nd Edition. Boca Raton, Florida, United States: CRC Press. [30, 48, 50, 51]

Meyer K., A. Hoyer-Leitzel, S. Iams, *et al.* (2018). Quantifying resilience to recurrent ecosystem disturbances using flow–kick dynamics. *Nat Sustain* **1**: 671–678. [93]

Mo Y., D. W. Eyre, S. F. Lumley, *et al.* (2021). Transmission of community- and hospital-acquired SARS-CoV-2 in hospital settings in the UK: a cohort study. *PLOS Med* **18**: e1003816. [111]

Mo Y., T. M. Pham, C. Lim, *et al.* (2022). The effect of hand hygiene frequency on reducing acute respiratory infections in the community: a meta-analysis. *Epidemiol Infect* **150**: e79. [2, 9]

Moore S. C., B. Kronsteiner, S. Longet, *et al.* (2023). Evolution of long-term vaccine-induced and hybrid immunity in healthcare workers after different COVID-19 vaccine regimens. *Med* **4**: 191–215. [108]

Moriyama M., W. J. Hugentobler, and A. Iwasaki (2020). Seasonality of respiratory viral infections. *Annu Rev Virol* **7**: 83–101. [37]

Moss W. J., and D. E. Griffin (2012). Measles. *Lancet* **379**: 153–164. [8]

Mossong J., D. J. Nokes, W. J. Edmunds, M. J. Cox, S. Ratnam, and C. P. Muller (1999). Modeling the impact of subclinical measles transmission in vaccinated populations with waning immunity. *Am J Epidemiol* **150**: 1238–1249. [92]

Mulberry N., P. Tupper, E. Kirwin, C. McCabe, and C. Colijn (2021). Vaccine rollout strategies: The case for vaccinating essential workers early. *PLOS Glob Public Health* **1**: e0000020. [133]

Mylius S. D., T. J. Hagenaars, A. K. Lugnér, and J. Wallinga (2008). Optimal allocation of pandemic influenza vaccine depends on age, risk and timing. *Vaccine* **26**: 3742–3749. [133]

NHS England (2020). *Novel coronavirus (COVID-19) standard operating procedure*. URL: `https://www.england.nhs.uk/coronavirus/wp-content/`

uploads/sites/52/2020/11/C0873_i_SOP_LFD-rollout-for-asymptomatic-staff-testing_phase-2-trusts-v1.1_16-nov20.pdf (visited on 5 December 2024). [114]

NHS England (2021a). *COVID-19 Vaccination Statistics*. URL: https://www.england.nhs.uk/statistics/wp-content/uploads/sites/2/2021/10/COVID-19-weekly-announced-vaccinations-28-October-2021.pdf (visited on 29 October 2024). [14]

— (2021b). *Lateral flow antigen tests for asymptomatic staff testing: frequently asked questions for – primary care (organisations and staff)*. URL: https://www.england.nhs.uk/coronavirus/documents/lateral-flow-antigen-tests-for-asymptomatic-staff-testing-frequently-asked-questions-for-primary-care-organisations-and-staff/ (visited on 5 December 2024). [114]

— (2021c). *NHS rolls out COVID-19 jab to children aged 12 to 15*. URL: https://www.england.nhs.uk/2021/09/nhs-rolls-out-covid-19-jab-to-children-aged-12-to-15/ (visited on 12 December 2024). [133]

— (2021d). *Use of lateral flow devices for asymptomatic staff testing for SARS CoV-2 in all NHS Staff*. URL: https://www.england.nhs.uk/coronavirus/documents/12184-2/ (visited on 5 December 2024). [114]

— (2022a). *COVID-19 hospital activity*. URL: https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/ (visited on 4 July 2022). [114]

— (2022b). *Updated UK Health Security Agency (UKHSA) guidance –isolation*. URL: https://www.england.nhs.uk/coronavirus/documents/updated-uk-health-security-agency-ukhsa-guidance-isolation/#Self%20isolation%20for%20NHS%20staff (visited on 5 December 2024). [114]

Nishiura H. (2007). Lessons from previous predictions of HIV/AIDS in the United States and Japan: epidemiologic models and policy formulation. *Epidemiol Perspect Innov* 4: 1–7. [32]

Nuffield Council on Bioethics (2007). *Public health: Ethical issues*. London, UK: Nuffield Council on Bioethics. [2]

Nussbaumer-Streit B., V. Mayr, A. I. Dobrescu, *et al.* (2020). Quarantine alone or in combination with other public health measures to control COVID-19: a rapid review. *Cochrane Database Syst Rev*, CD013574. [11]

Nygaard U., U. B. Hartling, J. Nielsen, *et al.* (2023). Hospital admissions and need for mechanical ventilation in children with respiratory syncytial virus before and during the COVID-19 pandemic: a Danish nationwide cohort study. *Lancet Child Adolesc Health* 7: 171–179. [89]

Office for National Statistics (2023a). *Coronavirus (COVID-19) Infection Survey, UK Statistical bulletins*. URL: https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/bulletins/coronaviruscovid19infectionsurveypilot/previousReleases (visited on 29 October 2024). [12, 13]

Office for National Statistics (2023b). *Coronavirus (COVID-19) Infection Survey: methods and further information*. URL: https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/methodologies/covid19infectionsurveypilotmethodsandfurtherinformation (visited on 29 October 2024). [12, 13]

— (2023c). *Coronavirus (COVID-19) Infection Survey: technical data*. URL: https://www.ons.gov.uk/peoplepopulationandcommunity/healthandsocialcare/conditionsanddiseases/datasets/covid19infectionsurveytechnicaldata/2022 (visited on 12 January 2025). [13]

Ogilvie D., M. Egan, V. Hamilton, and M. Petticrew (2005). Systematic reviews of health effects of social interventions: 2. Best available evidence: how low should you go? *J Epidemiol Community Health* **59**: 886–892. [3]

Opoku-Sarkodie R., F. A. Bartha, M. Polner, and G. Röst (2022). Dynamics of an SIRWS model with waning of immunity and varying immune boosting period. *J Biol Dyn* **16**: 596–618. [93]

Owusu-Ansah E. G. J., B. Barnes, R. Abaidoo, *et al.* (2019). Probabilistic modeling for an integrated temporary acquired immunity with norovirus epidemiological data. *Infect Dis Model* **4**: 99–114. [92]

Oxman M. N., M. J. Levin, G. R. Johnson, *et al.* (2005). A vaccine to prevent herpes zoster and postherpetic neuralgia in older adults. *N Engl J Med* **352**: 2271–2284. [88]

Pacheco-Barrios K., A. Cardenas-Rojas, S. Giannoni-Luza, and F. Fregni (2020). COVID-19 pandemic and Farr's law: A global comparison and prediction of outbreak acceleration and deceleration rates. *PLOS One* **15**: e0239175. [32]

Pan K., A. Goel, L. R. Akin, and S. R. Patel (2020). Through plagues and pandemics: the evolution of medical face masks. *R I Med J* **103**: 72–75. [1]

Pearl J. (1995). Causal diagrams for empirical research. *Biometrika* **82**: 669–688. [22]

— (2000). Causal inference without counterfactuals: Comment. *J Am Stat Assoc* **95**: 428–431. [4]

Pearl J., and D. Mackenzie (2018). *The book of why: The new science of cause and effect*. London, United Kingdom: Penguin. [4]

Pearl J., M. Glymour, and N. P. Jewell (2016). *Causal inference in statistics: A primer*. Chichester, United Kingdom: John Wiley & Sons. [23]

Peiris M. (2020). Respiratory tract viruses. In: J. Firth, C. Conlon, and T. Cox, editors. *Oxford textbook of medicine*. 6th edition. Oxford, United Kingdom: Oxford University Press. [7]

Polack F. P., S. J. Thomas, N. Kitchin, *et al.* (2020). Safety and efficacy of the BNT162b2 mRNA Covid-19 vaccine. *N Engl J Med* **383**: 2603–2615. [113]

Public Health etc. (Scotland) Act (2008). URL: https://www.legislation.gov.uk/asp/2008/5/contents (visited on 1 December 2024). [89]

Public Health Scotland (2023). *Respiratory Scotland*. URL: https://www.opendata.nhs.scot/dataset/48a9310e-4db3-44a4-bce8-6b4be9deb88a/resource/37beac86-f8fb-4ab5-9457-2b8ddac9c089/download/respiratory_scot.csv (visited on 18 July 2023). [89]

Quigley A. L., H. Stone, P. Y. Nguyen, A. A. Chughtai, and C. R. MacIntyre (2021). Estimating the burden of COVID-19 on the Australian healthcare workers and health system during the first six months of the pandemic. *Int J Nurs Stud* **114**: 103811. [111]

Rackauckas C., and Q. Nie (2017). DifferentialEquations.jl – A performant and feature-rich ecosystem for solving differential equations in Julia. *J Open Res Softw* **5**: 15. [38, 44, 97]

Rambachan A., and J. Roth (2023). A more credible approach to parallel trends. *Rev Econ Stud* **90**: 2555–2591. [72, 82]

Rambhia K. J., and M. T. Rambhia (2019). Early bird gets the flu: what should be done about waning intraseasonal immunity against seasonal influenza? *Clin Infect Dis* **68**: 1235–1240. [88, 112]

Rizoiu M.-A., S. Mishra, Q. Kong, M. Carman, and L. Xie (2018). SIR-Hawkes: Linking epidemic models and Hawkes processes to model diffusions in finite populations. In: *Proceedings of the 2018 world wide web conference*. Geneva, Switzerland, pp. 419–428. [83]

Robins J. M., M. A. Hernán, and B. Brumback (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology* **11**: 550–560. [25]

Rosenstrom E. T., J. S. Ivy, M. E. Mayorga, and J. L. Swann (2022). Could earlier availability of boosters and pediatric vaccines have reduced impact of COVID-19? In: B. Feng, G. Pedrielli, Y. Peng, *et al.*, editors. *Proceedings of the 2022 winter simulation conference*. Singapore: Institute of Electrical and Electronics Engineers, pp. 1092–1103. [133]

Roth J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *Am Econ Rev Insights* **4**: 305–322. [73]

Roth J., P. H. C. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *J Econom* **235**: 2218–2244. [55, 62, 72, 82]

Rothbard S., J. C. Etheridge, and E. J. Murray (2024). A tutorial on applying the difference-in-differences method to health data. *Curr Epidemiol Rep* **11**: 85–95. [56]

Rozenbaum M. H., J. Judy, D. Tran, K. Yacisin, S. K. Kurosky, and E. Begier (2023). Low levels of RSV testing among adults hospitalized for lower respiratory tract infection in the United States. *Infect Dis Ther* **12**: 677–685. [89]

Rubin D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol* **66**: 688. [3, 21]

— (1978). Bayesian inference for causal effects: The role of randomization. *Ann Stat*, 34–58. [23]

Rubin D. B. (1980). Randomization analysis of experimental data: The Fisher randomization test: Comment. *J Am Stat Assoc* **75**: 591–593. [26]

— (1986). Statistics and causal inference: Comment: Which ifs have causal answers. *J Am Stat Assoc* **81**: 961–962. [26]

— (1990). Comment: Neyman (1923) and causal inference in experiments and observational studies. *Stat Sci* **5**: 472–480. [26]

— (2004). Direct and indirect causal effects via potential outcomes. *Scand J Stat* **31**: 161–170. [141]

— (2005). Causal inference using potential outcomes: Design, modeling, decisions. *J Am Stat Assoc* **100**: 322–331. [17–19, 21]

Ruppel Shell E. (2020). The United Kingdom's mask crusader. *Science* **370**: 276–277. [59]

Russell M. W., and J. Mestecky (2022). Mucosal immunity: The missing link in comprehending SARS-CoV-2 infection and transmission. *Front Immunol* **13**: 957107. [108]

Ryan A. M., J. F. Burgess, and J. B. Dimick (2015). Why we should not be indifferent to specification choices for difference-in-differences. *Health Serv Res* **50**: 1211–1235. [82]

Ryan A. M., E. Kontopantelis, A. Linden, and J. F. Burgess (2019). Now trending: Coping with non-parallel trends in difference-in-differences analysis. *Stat Methods Med Res* **28**: 3697–3711. [72, 82]

Saad-Roy C. M., and A. Traulsen (2023). Dynamics in a behavioral–epidemiological model for individual adherence to a nonpharmaceutical intervention. *Proc Natl Acad Sci U S A* **120**: e2311584120. [27, 81, 83]

Santillana M., A. Tuite, T. Nasserie, *et al.* (2018). Relatedness of the incidence decay with exponential adjustment (IDEA) model, "Farr's law" and SIR compartmental difference equation models. *Infect Dis Model* **3**: 1–12. [32]

Savulescu J., I. Persson, and D. Wilkinson (2020). Utilitarianism and the pandemic. *Bioethics* **34**: 620–632. [134]

Scheibe K. E., and T. R. Sarbin (1965). Towards a theoretical conceptualisation of superstition. *Br J Philos Sci* **16**: 143–158. [15]

Schenzle D. (1984). An age-structured model of pre- and post-vaccination measles transmission. *Math Med Biol* **1**: 169–191. [38]

Schulz K. (2011). "If you'd wiggled A, then B would've changed": Causality and counterfactual conditionals. *Synthese* **179**: 239–251. [17]

Scottish Public Health Observatory (2023). *Pregnancy, births and maternity: Key points for Scotland*. URL: `https://www.scotpho.org.uk/population-dynamics/pregnancy-births-and-maternity/key-points/` (visited on 20 July 2023). [97]

Shaman J., V. E. Pitzer, C. Viboud, B. T. Grenfell, and M. Lipsitch (2010). Absolute humidity and the seasonal onset of influenza in the continental United States. *PLOS Biol* **8**: e1000316. [38]

Shaw-Taylor L. (2020). An introduction to the history of infectious diseases, epidemics and the early phases of the long-run decline in mortality. *Econ Hist Rev* **73**: E1–E19. [1]

Shek L. P.-C., and B.-W. Lee (2003). Epidemiology and seasonality of respiratory tract virus infections in the tropics. *Paediatr Respir Rev* **4**: 105–111. [37]

Showalter D. E., and J. G. Royde-Smith (2024). World War I. In: *Encyclopedia Britannica*. url: `https://www.britannica.com/event/World-War-I` (visited on 16 October 2024). [8]

Slifka M. K., and R. Ahmed (1996). Long-term humoral immunity against viruses: revisiting the issue of plasma cell longevity. *Trends Microbiol* **4**: 394–400. [8, 88, 89, 98]

Smith L. E., H. W. Potts, R. Amlôt, N. T. Fear, S. Michie, and G. J. Rubin (2022). Engagement with protective behaviours in the UK during the COVID-19 pandemic: a series of cross-sectional surveys (the COVID-19 rapid survey of adherence to interventions and responses [CORSAIR] study). *BMC Public Health* **22**: 475. [108]

Steel S. (2012). Causation in English tort law: Still wrong after all these years. *Univ Qld Law J* **31**: 243–264. [17]

Stout A., A. Van Stelten-Carlson, H. Marquis, *et al.* (2020). Public health impact of foodborne exposure to naturally occurring virulence-attenuated *Listeria monocytogenes*: Inference from mouse and mathematical models. *Interface Focus* **10**: 20190046. [92]

Stratford J., E. MacKenzie, and E. Mockford (2020). Balancing speed and safety: the authorisation of Covid-19 vaccines and medicines. *Judic Rev* **25**: 105–117. [3]

Strogatz S. H. (2018). *Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering.* Boca Raton, Florida, United States: CRC Press. [36, 38, 43, 105]

Struik S. S., and E. M. Riley (2004). Does malaria suffer from lack of memory? *Immunol Rev* **201**: 268–290. [88]

Suárez-García I., M. J. Martínez de Aramayona López, A. Sáez Vicente, and P. Lobo Abascal (2020). SARS-CoV-2 infection among healthcare workers in a hospital in Madrid, Spain. *J Hosp Infect* **106**: 357–363. [111]

Sugiyama M. (2016). *Introduction to statistical machine learning*. Waltham, Massachusetts, United States: Morgan Kaufmann. [50]

Sun L., and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *J Econom* **225**: 175–199. [73]

Surjanovic N., M. Biron-Lattes, P. Tiede, S. Syed, T. Campbell, and A. Bouchard-Côté (2023). Pigeons.jl: Distributed sampling from intractable distributions. *arXiv*, 2308.09769. [51, 66, 102]

Svensson Å. (2007). A note on generation times in epidemic models. *Math Biosci* **208**: 300–311. [39]

Swann O. V., K. A. Holden, L. Turtle, *et al.* (2020). Clinical characteristics of children and young people admitted to hospital with covid-19 in United Kingdom: prospective multicentre observational cohort study. *BMJ* **370**: m3249. [145]

Syed S., A. Bouchard-Côté, G. Deligiannidis, and A. Doucet (2022). Non-reversible parallel tempering: a scalable highly parallel MCMC scheme. *J R Stat Soc Series B Stat Methodol* **84**: 321–350. [51, 66, 102]

Tamerius J., M. I. Nelson, S. Z. Zhou, C. Viboud, M. A. Miller, and W. J. Alonso (2011). Global influenza seasonality: reconciling patterns across temperate and tropical regions. *Environ Health Perspect* **119**: 439–445. [37]

Tan S.-B., J. Wee, H.-B. Wong, and D. Machin (2008). Can external and subjective information ever be used to reduce the size of randomised controlled trials? *Contemp Clin Trials* **29**: 211–219. [48]

Tellier R., Y. Li, B. J. Cowling, and J. W. Tang (2019). Recognition of aerosol transmission of infectious agents: a commentary. *BMC Infect Dis* **19**: 101. [9]

The Health Protection (Coronavirus, Restrictions) (Protected Areas and Restriction on Businesses) (Amendment) Regulations (2020). URL: https://www.legislation.gov.uk/uksi/2020/1041/contents/made (visited on 29 November 2024). [58]

Thompson E. (2022). *Escape from Model Land: How mathematical models can lead us astray and what we can do about it*. London, United Kingdom: Basic Books. [30, 31]

Thompson W. W., D. K. Shay, E. Weintraub, *et al.* (2003). Mortality associated with influenza and respiratory syncytial virus in the United States. *JAMA* **289**: 179–186. [8, 107]

Turner R. B. (1997). Epidemiology, pathogenesis, and treatment of the common cold. *Ann Allergy Asthma Immunol* **78**: 531–540. [7]

UK Health Security Agency (2021). *Respiratory syncytial virus (RSV): symptoms, transmission, prevention, treatment*. URL: https://www.gov.uk/government/publications/respiratory-syncytial-virus-rsv-symptoms-transmission-prevention-treatment/respiratory-syncytial-virus-rsv-symptoms-transmission-prevention-treatment (visited on 1 December 2024). [89]

— (2022a). *COVID-19 vaccine uptake in frontline healthcare workers: monthly data, 2021 to 2022*. URL: https://www.gov.uk/government/statistics/covid-19-vaccine-uptake-in-frontline-healthcare-workers-monthly-data-2021-to-2022 (visited on 22 September 2024). [114]

— (2022b). *COVID-19: management of staff and exposed patients or residents in health and social care settings*. URL: https://www.gov.uk/government/publications/covid-19-management-of-exposed-healthcare-workers-and-patients-in-hospital-settings/covid-19-management-of-exposed-healthcare-workers-and-patients-in-hospital-settings (visited on 5 December 2024). [114]

UK Health Security Agency (2022c). *Estimating the proportion of COVID-19 cases detected by testing during the pandemic: September 2020 to March 2022*. URL: https://www.gov.uk/government/publications/estimating-the-proportion-of-covid-19-cases-detected-by-testing (visited on 23 May 2024). [58, 66]

— (2022d). Influenza. In: M. Ramsay, editor. *Immunisation against infectious disease*. UK Health Security Agency. Chap. 19. URL: https://www.gov.uk/government/publications/influenza-the-green-book-chapter-19 (visited on 26 August 2023). [8]

— (2023). COVID-19–SARS-CoV-2. In: M. Ramsay, editor. *Immunisation against infectious disease*. UK Health Security Agency. Chap. 14a. URL: https://www.gov.uk/government/publications/covid-19-the-green-book-chapter-14a (visited on 26 August 2023). [8, 14, 111, 134]

— (2024a). Respiratory syncytial virus. In: M. Ramsay, editor. *Immunisation against infectious disease*. UK Health Security Agency. Chap. 27a. URL: https://www.gov.uk/government/publications/respiratory-syncytial-virus-the-green-book-chapter-27a (visited on 3 December 2024). [107]

— (2024b). Shingles (herpes zoster). In: M. Ramsay, editor. *Immunisation against infectious disease*. UK Health Security Agency. Chap. 28a. URL: https://www.gov.uk/government/publications/shingles-herpes-zoster-the-green-book-chapter-28a (visited on 30 November 2024). [88]

University of Liverpool (2021a). *Covid-SMART asymptomatic testing pilot in Liverpool City Region: quantitative evaluation*. Liverpool, United Kingdom: University of Liverpool. [59]

— (2021b). *Liverpool Covid-SMART Community Testing Pilot*. Liverpool, United Kingdom: University of Liverpool. [58, 59, 81, 85]

Upshaw T. L., C. Brown, R. Smith, M. Perri, C. Ziegler, and A. D. Pinto (2021). Social determinants of COVID-19 incidence and outcomes: A rapid review. *PLOS One* **16**: e0248336. [142]

Vallabhajosyula S., S. H. Patlolla, W. Cheungpasitporn, D. R. Holmes, and B. J. Gersh (2020). Influence of seasons on the management and outcomes acute myocardial infarction: an 18-year US study. *Clin Cardiol* **43**: 1175–1185. [135]

van Hoek A. J., A. Melegaro, E. Zagheni, W. J. Edmunds, and N. Gay (2011). Modelling the impact of a combined varicella and zoster vaccination programme on the epidemiology of varicella zoster virus in England. *Vaccine* **29**: 2411–2420. [91]

van Hoek A. J., A. Melegaro, N. Gay, J. Bilcke, and W. J. Edmunds (2012). The cost-effectiveness of varicella and combined varicella and herpes zoster vaccination programmes in the United Kingdom. *Vaccine* **30**: 1225–1234. [91]

van Summeren J., A. Meijer, G. Aspelund, *et al.* (2021). Low levels of respiratory syncytial virus activity in Europe during the 2020/21 season:

what can we expect in the coming summer and autumn/winter? *Eurosurveillance* **26**: 2100639. [10, 89]

VanderWeele T. J. (2008). The sign of the bias of unmeasured confounding. *Biometrics* **64**: 702–706. [23, 25]

— (2015). *Explanation in causal inference: Methods for mediation and interaction*. Oxford, United Kingdom: Oxford University Press. [24, 28, 117]

VanderWeele T. J., and E. J. Tchetgen Tchetgen (2011). Bounding the infectiousness effect in vaccine trials. *Epidemiology* **22**: 686–693. [29]

Vehtari A., A. Gelman, D. Simpson, B. Carpenter, and P.-C. Bürkner (2021). Rank-normalization, folding, and localization: An improved $\hat{R}$ for assessing convergence of MCMC (with discussion). *Bayesian Anal* **16**: 667–718. [51]

Verner J. H. (2010). Numerically optimal Runge–Kutta pairs with interpolants. *Numer Algorithms* **53**: 383–396. [97]

Viswanathan M., L. Kahwati, B. Jahn, *et al.* (2020). Universal screening for SARS-CoV-2 infection: a rapid review. *Cochrane Database Syst Rev*, CD013718. [10, 11]

Voysey M., S. A. Costa Clemens, S. A. Madhi, *et al.* (2021). Safety and efficacy of the ChAdOx1 nCoV-19 vaccine (AZD1222) against SARS-CoV-2: an interim analysis of four randomised controlled trials in Brazil, South Africa, and the UK. *Lancet* **397**: 99–111. [113]

Wang C. C., K. A. Prather, J. Sznitman, *et al.* (2021). Airborne transmission of respiratory viruses. *Science* **373**: eabd9149. [9, 10]

Ward H., C. Atchison, M. Whitaker, *et al.* (2023). Design and implementation of a national program to monitor the prevalence of SARS-CoV-2 IgG antibodies in England using self-testing: The REACT-2 Study. *Am J Public Health* **113**: 1201–1209. [12]

Waris M. (1991). Pattern of respiratory syncytial virus epidemics in Finland: two-year cycles with alternating prevalence of groups A and B. *J Infect Dis* **163**: 464–469. [108]

Wearing H. J., and P. Rohani (2009). Estimating the duration of pertussis immunity using epidemiological signatures. *PLOS Pathog* **5**: e1000647. [93, 98]

Wearing H. J., P. Rohani, and M. J. Keeling (2005). Appropriate models for the management of infectious diseases. *PLOS Med* **2**: e174. [37]

White L., J. Mandl, M. Gomes, *et al.* (2007). Understanding the transmission dynamics of respiratory syncytial virus using multiple time series and nested models. *Math Biosci* **209**: 222–239. [109]

Whittle H. C., P. Aaby, B. Samb, H. Jensen, J. Bennett, and F. Simondon (1999). Effect of subclinical infection on maintaining immunity against measles in vaccinated children in West Africa. *Lancet* **353**: 98–102. [88]

Wilson E. B., and J. Worcester (1945). Damping of epidemic waves. *Proc Natl Acad Sci U S A* **31**: 294–298. [106]

Wing C., K. Simon, and R. A. Bello-Gomez (2018). Designing difference in difference studies: best practices for public health policy research. *Annu Rev Public Health* **39**: 453–469. [55, 73]

Woodward J. (2003). *Making things happen: A theory of causal explanation*. New York, New York, United States: Oxford University Press. [15, 17]

Wright O. (2021). *Coronavirus: How the UK dealt with its first Covid case*. url: `https://www.bbc.co.uk/news/uk-england-55622386` (visited on 11 June 2024). [12]

Xiao J., M. Fang, Q. Chen, and B. He (2020). SARS, MERS and COVID-19 among healthcare workers: a narrative review. *J Infect Public Health* **13**: 843–848. [111]

Xu J., Y. Cheng, X. Yuan, W. V. Li, and L. Zhang (2020). Trends and prediction in daily incidence of novel coronavirus infection in China, Hubei Province and Wuhan City: an application of Farr's law. *Am J Transl Res* **12**: 1355–1361. [32]

Yang L., B. T. Grenfell, and M. J. Mina (2020). Waning immunity and re-emergence of measles and mumps in the vaccine era. *Curr Opin Virol* **40**: 48–54. [88, 107]

Yeoh D. K., D. A. Foley, C. A. Minney-Smith, *et al.* (2021). Impact of coronavirus disease 2019 public health measures on detections of influenza and respiratory syncytial virus in children during the 2020 Australian winter. *Clin Infect Dis* **72**: 2199–2202. [10, 89]

You C., Y. Deng, W. Hu, *et al.* (2020). Estimation of the time-varying reproduction number of COVID-19 outbreak in China. *Int J Hyg Environ Health* **228**: 113555. [41]

Zhang X., B. Barr, M. Green, *et al.* (2022). Impact of community asymptomatic rapid antigen testing on covid-19 related hospital admissions: synthetic control study. *BMJ* **379**: e071374. [59, 81, 85]

Zheng Z., J. L. Warren, E. D. Shapiro, V. E. Pitzer, and D. M. Weinberger (2022). Estimated incidence of respiratory hospitalizations attributable to RSV infections across age and socioeconomic groups. *Pneumonia* **14**: 6. [107]

Appendix A

Supplementary text

A.1 Correspondence with ethics committee

From: Mark Pritchard < ■■■■■@■■■■■ >
Sent: 30 June 2023 15:57
To: MSD Ethics < ■■■■■@■■■■■ >
Cc: Ben Cooper < ■■■■■@■■■■■ >
Subject: Use of publicly available data from NHS England to inform a mathematical model

Dear Dr Barnby-Porritt,

As part of my DPhil research I am developing a mathematical model that relates numbers of patients and healthcare workers who are testing positive for SARS-CoV-2 in healthcare settings.

There are two NHS datasets that I wish to use to inform my model:
1. Numbers of hospital beds occupied by people with covid-19, available at `https://www.england.nhs.uk/statistics/statistical-work-areas/uec-sitrep/urgent-and-emergency-care-daily-situation-reports-2022-23/`
2. Numbers of healthcare workers absent due to covid-19, available at `https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/`

These datasets are fully available online and they report no individual-level data. It would not be possible to identify any individuals from these datasets.

I therefore feel that formal ethics approval may not be needed. However,

as these data come from NHS England and report on NHS staff and patients, I would be grateful if I could get confirmation of this before proceeding.

Please let me know if you need further information to decide this question.

Best wishes,

Mark

Mark Pritchard

DPhil Student, Pandemic Sciences Institute, University of Oxford
Specialty Registrar, Thames Valley Public Health Training Programme
New Richards Building, Old Road Campus, Oxford OX3 7FZ

From: MSD Ethics < ■■■■■@■■■■■ >
Sent: 03 July 2023 14:06
To: Mark Pritchard < ■■■■■@■■■■■ >
Cc: Ben Cooper < ■■■■■@■■■■■ >
Subject: RE: Use of publicly available data from NHS England to inform a mathematical model

Dear Mark,

Thank you for your email. Because the data is publicly available and does not include any personal information, analysis of this data will not require ethics review. Please do note, for future, that if you needed to gain access to NHS data (i.e. it was not publicly available), then you would need to seek ethics review from an NHS ethics committee.

Best wishes
Helen

A.2 Model used to simulate data for Figure 2.1

The simulated data for panels *B* and *C* in Figure 2.1 were generated using an SIR ODE model (equation 1.18). Parameters used for the simulation are given in Table A.1.

Table A.1: Parameters used to generate simulated data for Figure 2.1

Parameter	Group 0	Group 1
Plotted colour	Green	Black
β	0.4	0.4
γ	0.2	0.2
μ	0	0
$s(0)$	997	989
$i(0)$	2	6
$r(0)$	1	5
Parameters described on page 36.		

A.3 Parameter identifiability for the susceptible–infectious–resistant–susceptible model with three immune compartments

Following the method described by Castro, and de Boer (2020), I re-wrote Equations 3.2 and 3.4 from the main text as functionally-independent summands. I scaled all parameters and compartment sizes by unknown scaling factors, u,

$$s \to u_s\, s,\ i \to u_i\, i,\ r_1 \to u_{r_1}\, r_1, \text{ etc.}$$
$$\beta_0 \to u_{\beta_0}\, \beta_0,\ \beta_1 \to u_{\beta_1}\, \beta_1,\ \gamma \to u_\gamma\, \gamma, \text{ etc.}$$

Assuming that prevalence was known without error, that is assuming that $i(t)$ was observed, set $u_i = 1$. I then equated each equation to its scaled version,

$$\begin{aligned}
f_{s1}: &\quad \mu = \frac{u_\mu}{u_s}\mu,\\
f_{s2}: &\quad \mu\, s + 3\,\omega\, r_3 = u_\mu\, \mu\, s + \frac{3\, u_\omega\, u_{r_3}}{u_s}\,\omega\, r_3,\\
f_{s3}: &\quad \beta_0\, s\, i\,(1 + \beta_1\, \cos(2\pi\, t - \varphi)) = u_{\beta_0}\, \beta_0\, s\, i\,\left(1 + u_{\beta_1}\, \beta_1\, \cos\left(2\pi\, t - u_\varphi\, \varphi\right)\right),\\
f_{i1}: &\quad (\gamma + \mu)\, i = \left(u_\gamma\, \gamma + u_\mu\, \mu\right) i,\\
f_{i2}: &\quad \beta_0\, s\, i\,(1 + \beta_1\, \cos(2\pi\, t - \varphi))\\
&\qquad = u_{\beta_0}\, u_s\, \beta_0\, s\, i\,\left(1 + u_{\beta_1}\, \beta_1\, \cos\left(2\pi\, t - u_\varphi\, \varphi\right)\right),\\
f_{r_1 1}: &\quad \gamma\, i - (3\,\omega + \mu)\, r_1 = \frac{u_\gamma}{u_{r_1}}\,\gamma\, i - \left(3\, u_\omega\, \omega + u_\mu\, \mu\right) r_1,\\
f_{r_1 2}: &\quad \beta_0\, \psi\, i\,(1 + \beta_1\, \cos(2\pi\, t - \varphi))\,(r_2 + r_3)\\
&\qquad = \frac{u_{\beta_0}\, u_\psi}{u_{r_1}}\,\beta_0\, \psi\, i\,\left(1\, + u_{\beta_1}\beta_1\, \cos\left(2\pi\, t - u_\varphi\, \varphi\right)\right)\,\left(u_{r_2}\, r_2 + u_{r_3}\, r_3\right),\\
f_{r_2 1}: &\quad 3\,\omega\, r_1 - (3\,\omega + \mu)\, r_2 = \frac{3\, u_\omega\, u_{r_1}}{u_{r_2}}\,\omega\, r_1 - \left(3\, u_\omega\, \omega + u_\mu\, \mu\right) r_2,
\end{aligned}$$

$$
\begin{aligned}
f_{r_2 2}: \quad & \beta_0\,\psi\,i\,r_2\,(1+\beta_1\,\cos(2\pi\,t\,-\varphi)) \\
& \quad = u_{\beta_0}\,u_\psi\,\beta_0\,\psi\,i\,r_2\,\left(1+u_{\beta_1}\,\beta_1\,\cos\left(2\pi\,t-u_\varphi\,\varphi\right)\right), \\
f_{r_3 1}: \quad & 3\,\omega\,r_2-(3\,\omega+\mu)\,r_3=\frac{3\,u_\omega\,u_{r_2}}{u_{r_3}}\,\omega\,r_2-\left(3\,u_\omega\,\omega+u_\mu\,\mu\right)r_3, \\
f_{r_3 2}: \quad & \beta_0\,\psi\,i\,r_3\,(1+\beta_1\,\cos(2\pi\,t-\varphi))= \\
& \quad u_{\beta_0}\,u_\psi\,\beta_0\,\psi\,i\,r_3\,\left(1+\beta_1\,\cos\left(2\pi\,t-u_\varphi\varphi\right)\right).
\end{aligned}
$$

Rearranging these,

$$
\begin{aligned}
\frac{f_{i2}}{f_{s3}}: \quad & u_s=1, \\
\frac{f_{r_2 2}}{f_{s3}}: \quad & u_\psi=1, \\
u_s=1\to f_{s1}: \quad & u_\mu=1, \\
u_\mu=1\to f_{i1}: \quad & u_\gamma=1, \\
\text{Let } f_{x1}=\frac{\partial f_{s3}}{\partial t}: \quad & \sin(2\pi\,t-\varphi)=u_{\beta_0}\,u_{\beta_1}\,\sin\left(2\pi\,t-u_\varphi\,\varphi\right), \\
\text{Let } f_{x2}=\frac{\partial f_{s3}}{\partial \varphi}: \quad & \sin(2\pi\,t-\varphi)=u_{\beta_0}\,u_{\beta_1}\,u_\varphi\,\sin\left(2\pi t-u_\varphi\varphi\right), \\
\frac{f_{x2}}{f_{x1}}: \quad & u_\varphi=1, \\
u_\varphi=1\to f_{x1}: \quad & u_{\beta_0}\,u_{\beta_1}=1, \\
u_{\beta_1}=\frac{1}{u_{\beta_0}},\ u_\varphi=1\to f_{s3}: \quad & u_{\beta_0}=1, \\
& \Rightarrow u_{\beta_1}=1, \\
u_s=u_\mu=1\to f_{s2}: \quad & u_\omega\,u_{r_3}=1, \\
u_\gamma=u_\mu=1\to f_{r_1 1}: \quad & \gamma\,i-(3\,\omega+\mu)\,r_1=\frac{\gamma\,I}{u_{r_1}}-(3\,u_\omega\,\omega+\mu)\,r_1, \\
u_\psi=1\to\frac{f_{r_1 2}}{f_{s3}}: \quad & u_{r_1}\,(r_2+r_3)=u_{r_2}\,r_2+u_{r_3}\,r_3, \\
u_\mu=1\to f_{r_2 1}: \quad & r_1-r_2=u_\omega\left(\frac{u_{r_1}}{u_{r_2}}r_1-r_2\right), \\
u_\mu=1\to f_{r_3 1}: \quad & r_2-r_3=u_\omega\left(\frac{u_{r_2}}{u_{r_3}}r_2-r_3\right).
\end{aligned}
$$

As these equations allow scaling parameters $u\neq 1$, this model is not identifiable from prevalence data.

Appendix B

Supplementary figures

B.1 For Chapter 2

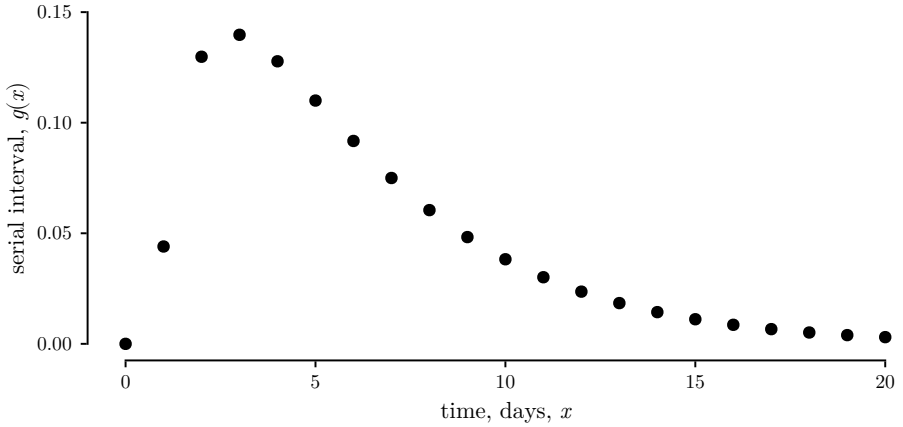

Figure B.1: Serial interval distribution for covid-19. Data from (Cori, *et al.*, 2022). [61]

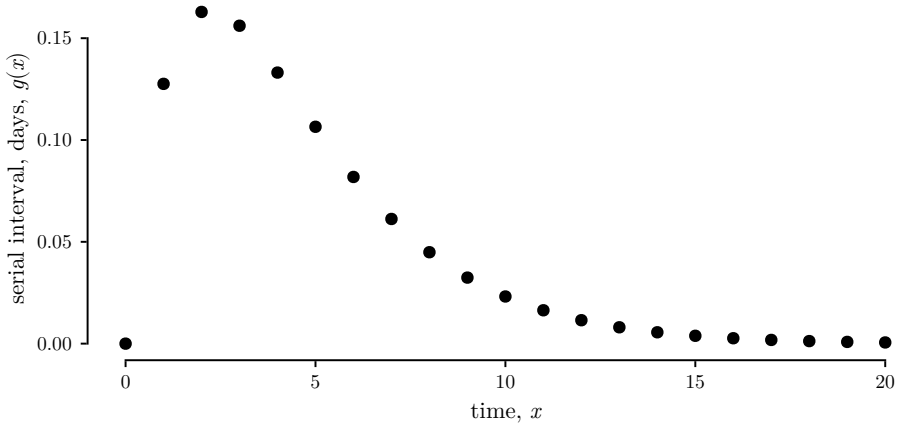

Figure B.2: Serial interval distribution for the simulations used for difference-in-differences analyses. See equation 2.2 (page 61) for details of the calculation; $\varsigma = 0.5/\text{day}$, $\gamma = 0.4/\text{day}$. [61]

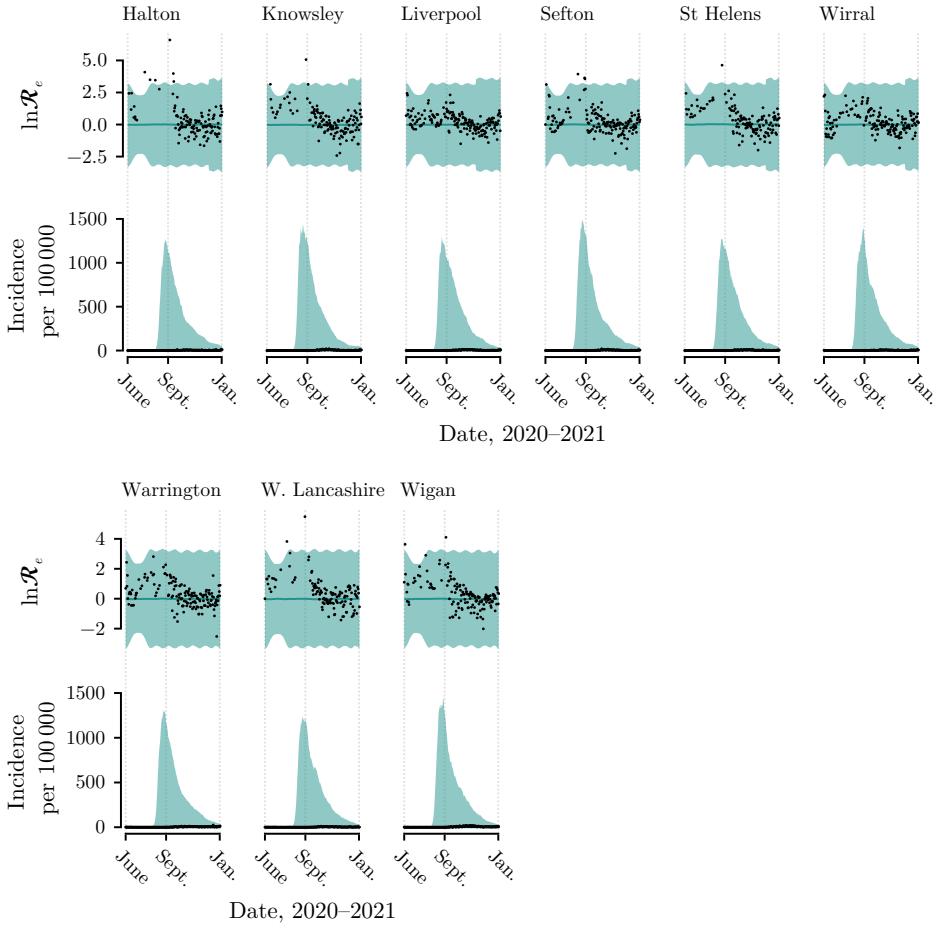

Figure B.3: Prior distribution plot for pillar 1 data on severe acute respiratory syndrome coronavirus 2 diagnoses in the Liverpool City Region when mass testing was introduced. Black dots show incidence and calculated logarithm of the effective reproduction number. Green shading shows the central 90% of outputs from the renewal equation with parameters sampled from the prior distributions. $\ln \mathcal{R}_e$ is the natural logarithm of the effective reproduction ratio. [67]

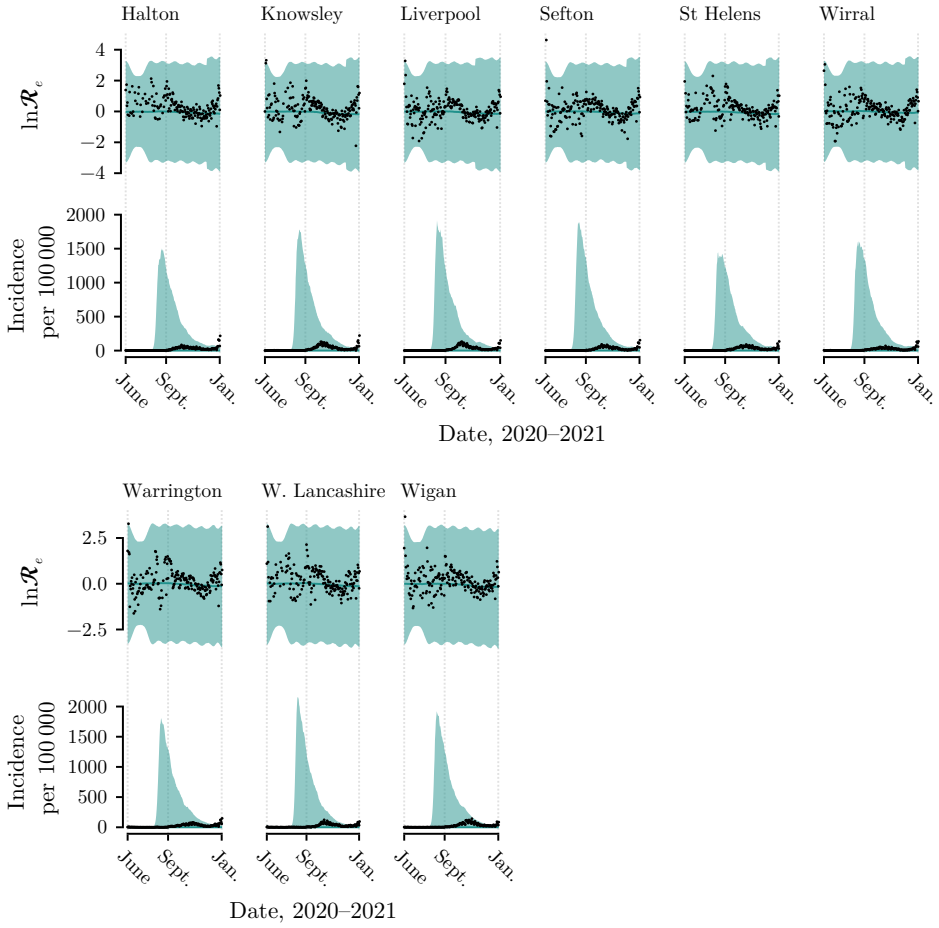

Figure B.4: Prior distribution plot for severe acute respiratory syndrome coronavirus 2 diagnoses in the Liverpool City Region when mass testing was introduced, using all diagnosis data. Black dots show incidence and calculated logarithm of the effective reproduction number. Green shading shows the central 90% of outputs from the renewal equation with parameters sampled from the prior distributions. $\ln \mathcal{R}_e$ is the natural logarithm of the effective reproduction ratio. [67]

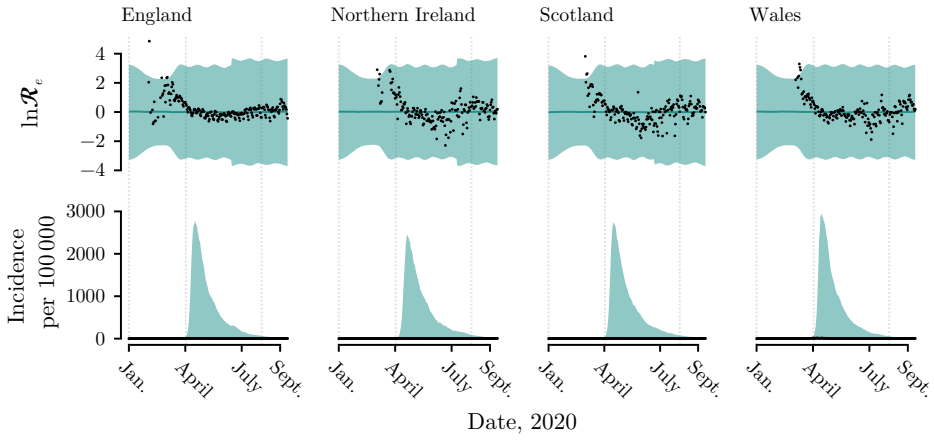

Figure B.5: Prior distribution plot for severe acute respiratory syndrome coronavirus 2 diagnoses in the United Kingdom when mask mandates were introduced. Black dots show incidence and calculated logarithm of the effective reproduction number. Green shading shows the central 90% of outputs from the renewal equation with parameters sampled from the prior distributions. $\ln \mathcal{R}_e$ is the natural logarithm of the effective reproduction ratio. [67]

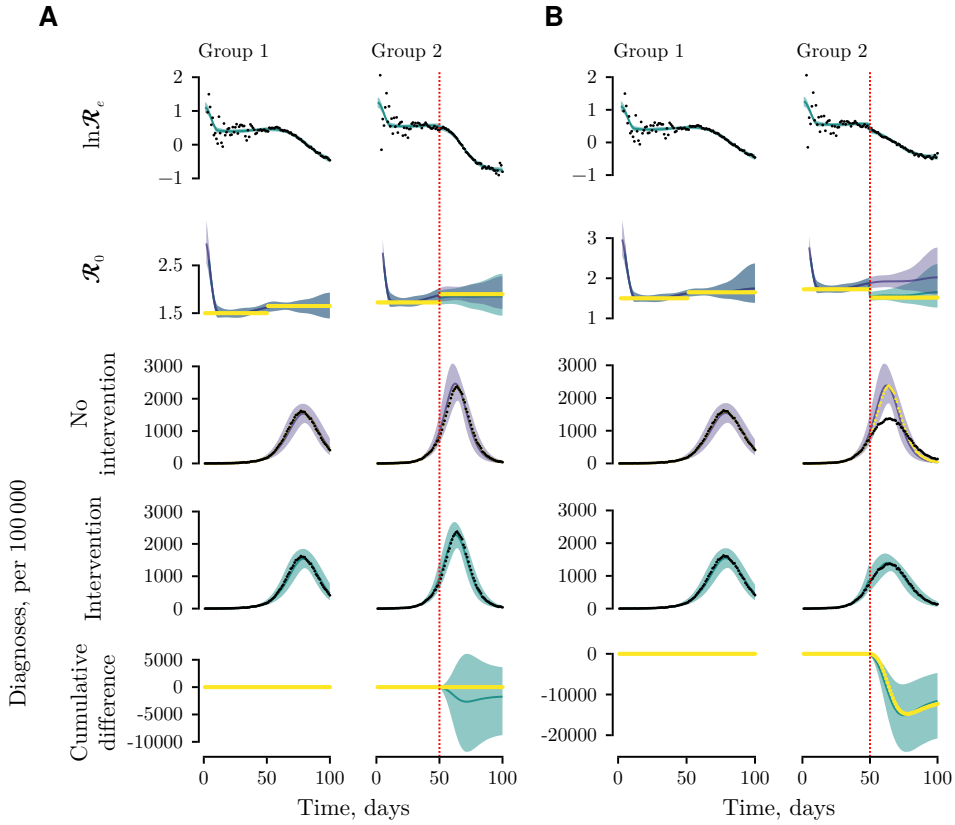

Figure B.6: Difference-in-differences analysis using spline-based time-varying parameters for simulated groups with two discrete levels of transmission. *A* Simulation with an ineffective intervention; *B* Simulation with an intervention that reduces transmission by 20%. Black dots show simulated data. Yellow dots show values that are known from the simulation but would not be known from data: basic reproduction ratio, simulated counterfactual number of cases, and simulative causal effect on cumulative number of cases. Green lines show results fitted to the data, and purple lines show inferred intervention-free counterfactuals. The bottom panel shows the estimated causal effect as a cumulative difference in numbers diagnosed per 100 000 population. Shaded areas represent the central 90% credible interval. Vertical dotted red lines show times interventions start. [69]

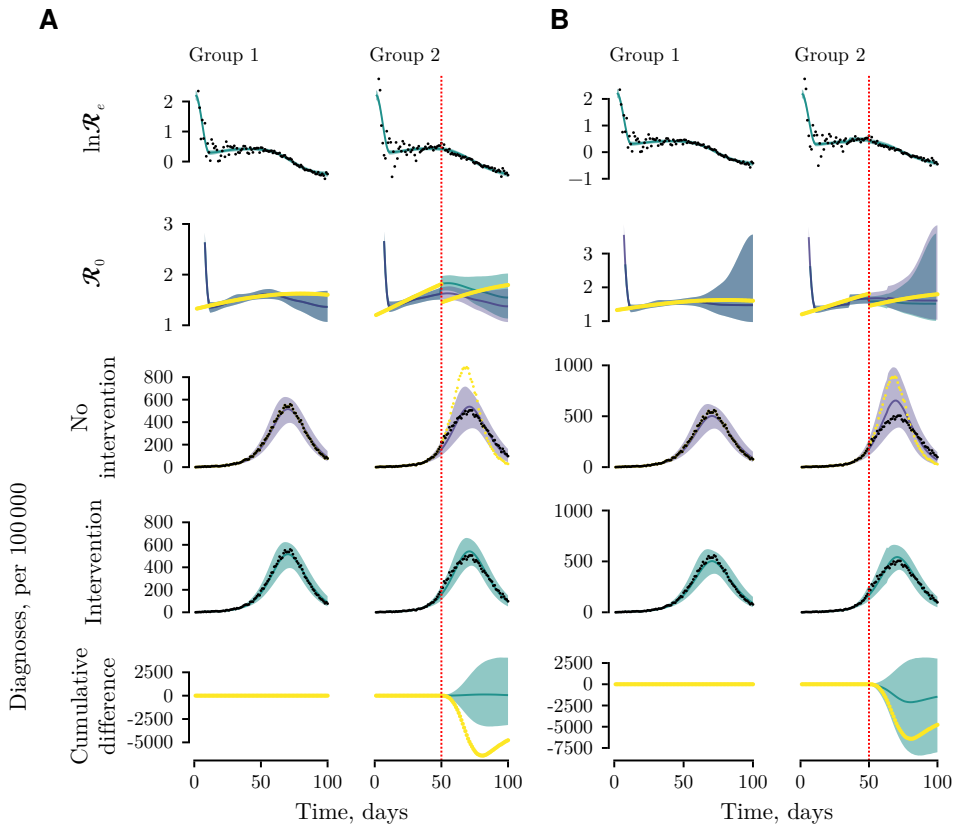

Figure B.7: Difference-in-differences analysis for a simulation with a non-parallel trend and an effective intervention. *A* Parameters fit without including consideration of non-parallel trend; *B* Parameters fit with assumed unobserved changes to infectiousness 14 days before and after the intervention. See caption to Figure B.6 for description of colours. [73]

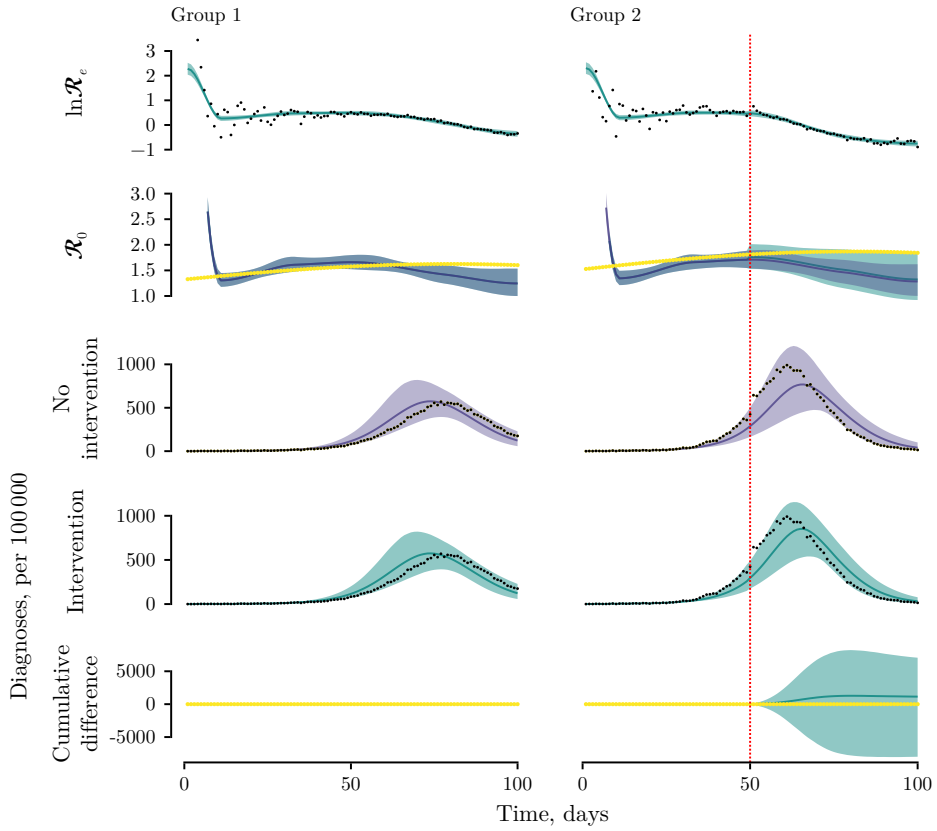

Figure B.8: Difference-in-differences analysis for a simulation with a time-varying proportion of cases diagnosed and no simulated effective intervention. In group 2, there is a 20% increase in the proportion diagnosed at the same time as the intervention. See caption to Figure B.6 for description of colours. [75]

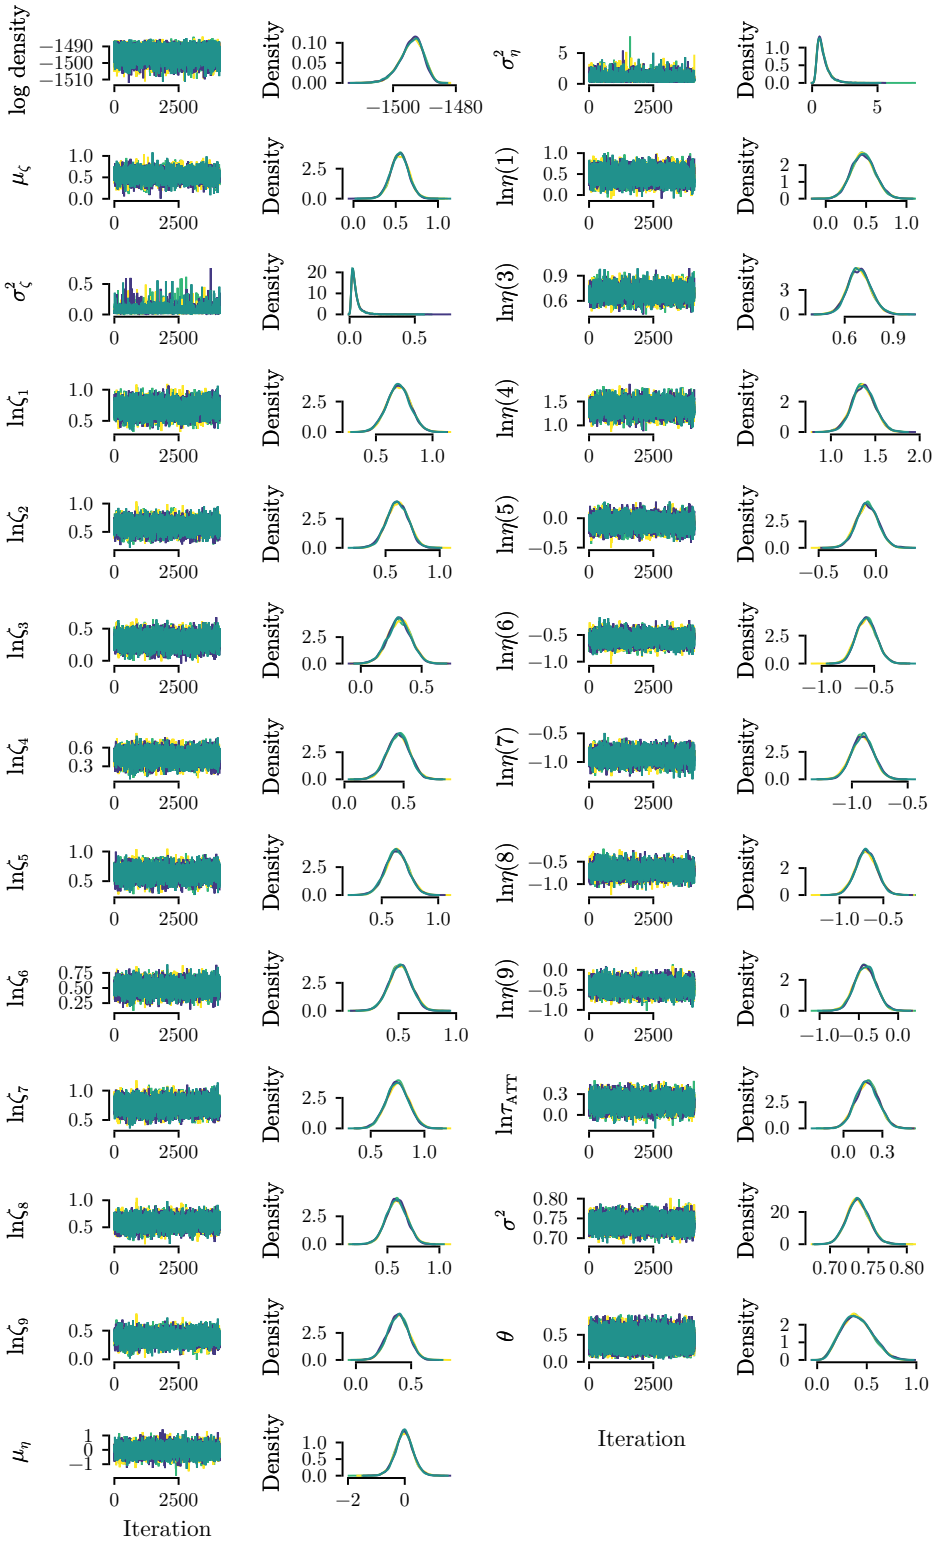

Figure B.9: Chains for analysis of pillar 1 data on severe acute respiratory syndrome coronavirus 2 diagnoses in the Liverpool City Region when mass testing was introduced. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [76]

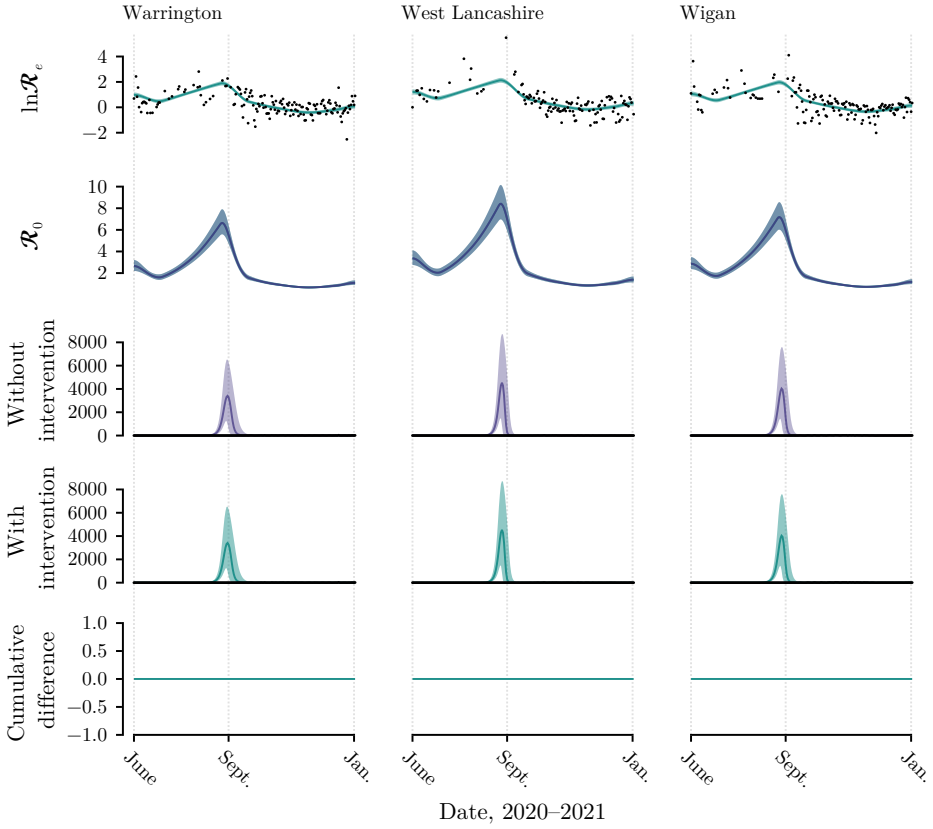

Figure B.10: Plot from difference-in-differences analysis using pillar 1 data for local authorities that did not have mass testing in 2020. See caption to Figure 2.2 for description of colours. [76]

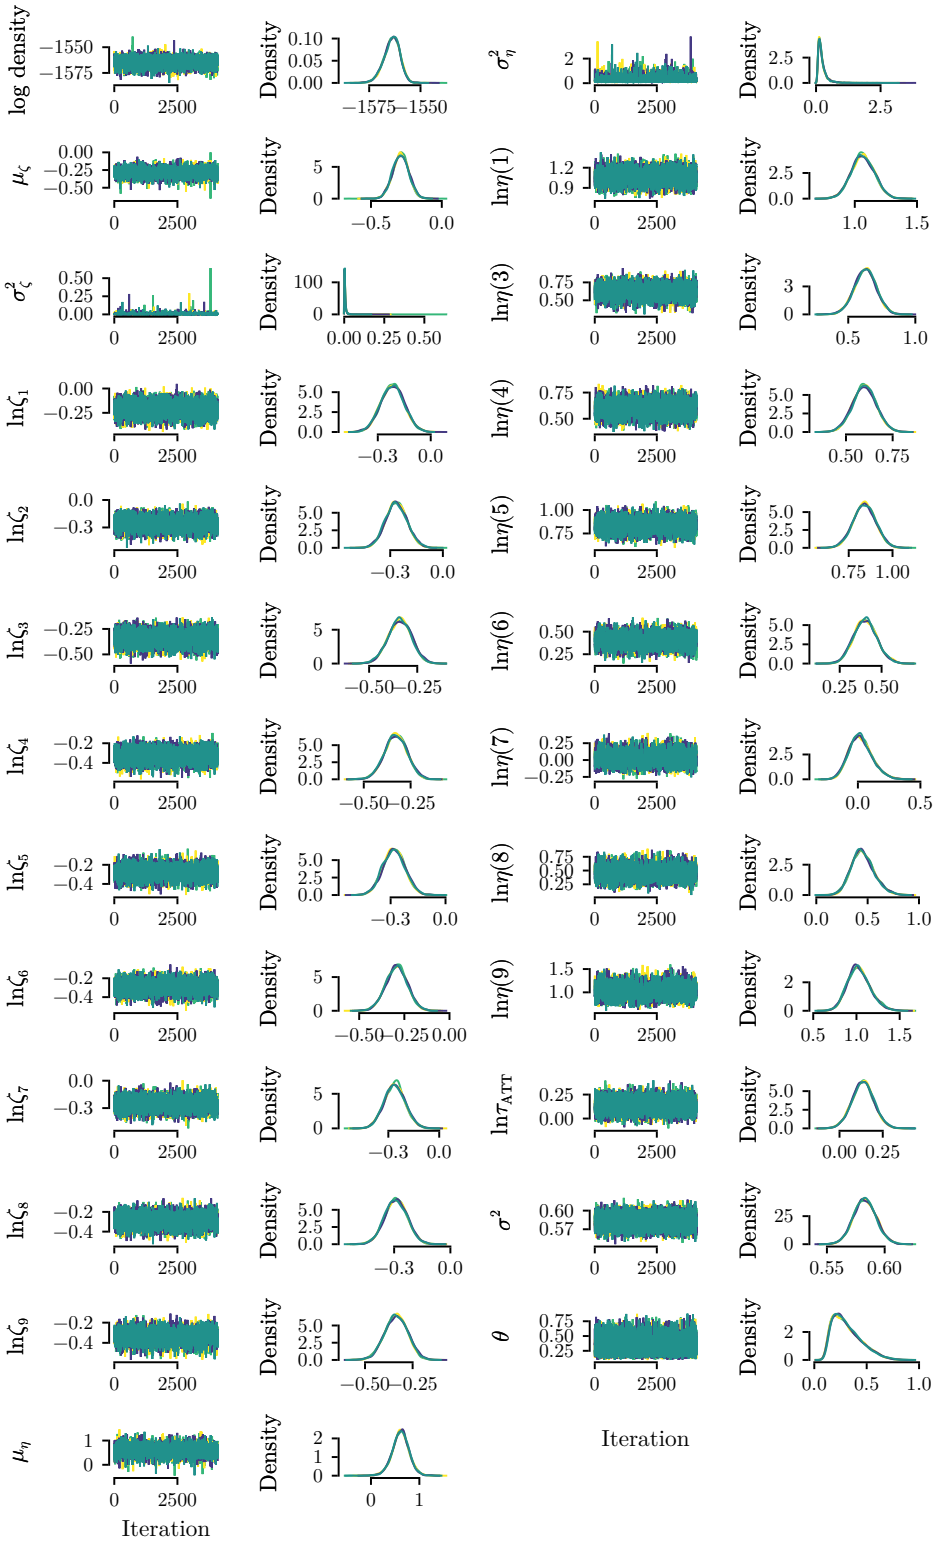

Figure B.11: Chains for analysis of all data on severe acute respiratory syndrome coronavirus 2 diagnoses in the Liverpool City Region when mass testing was introduced. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [76]

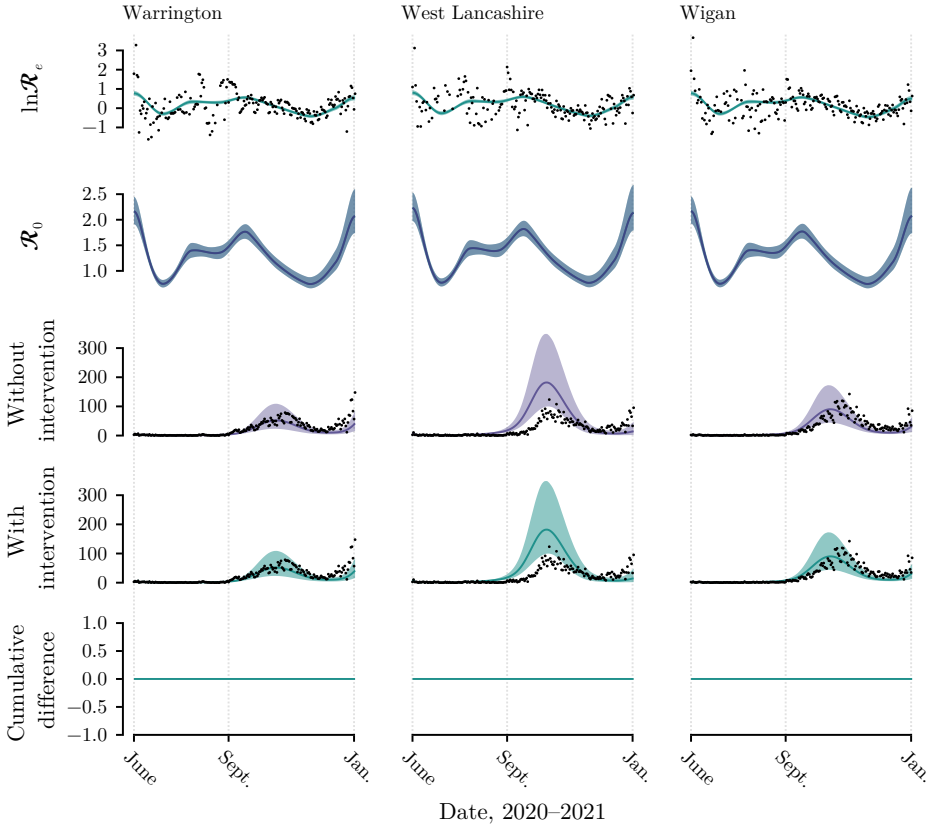

Figure B.12: Plot from difference-in-differences analysis using all data for local authorities that did not have mass testing in 2020. See caption to Figure 2.2 for description of colours. [76]

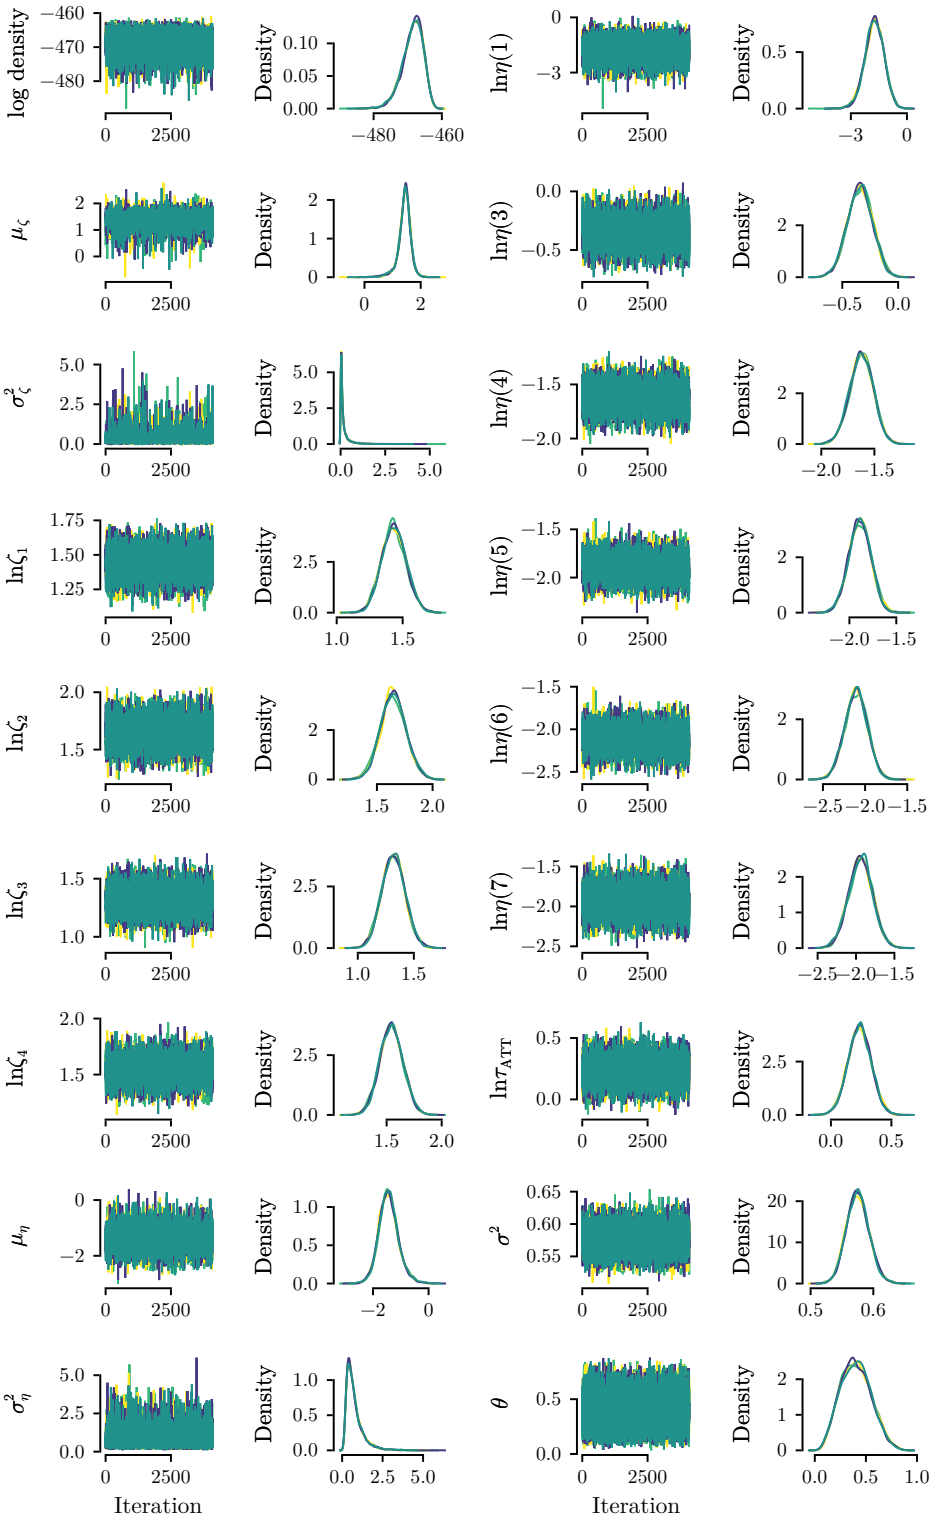

Figure B.13: Chains for analysis of effect of masking mandates in the United Kingdom. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [76]

B.2 For Chapter 3

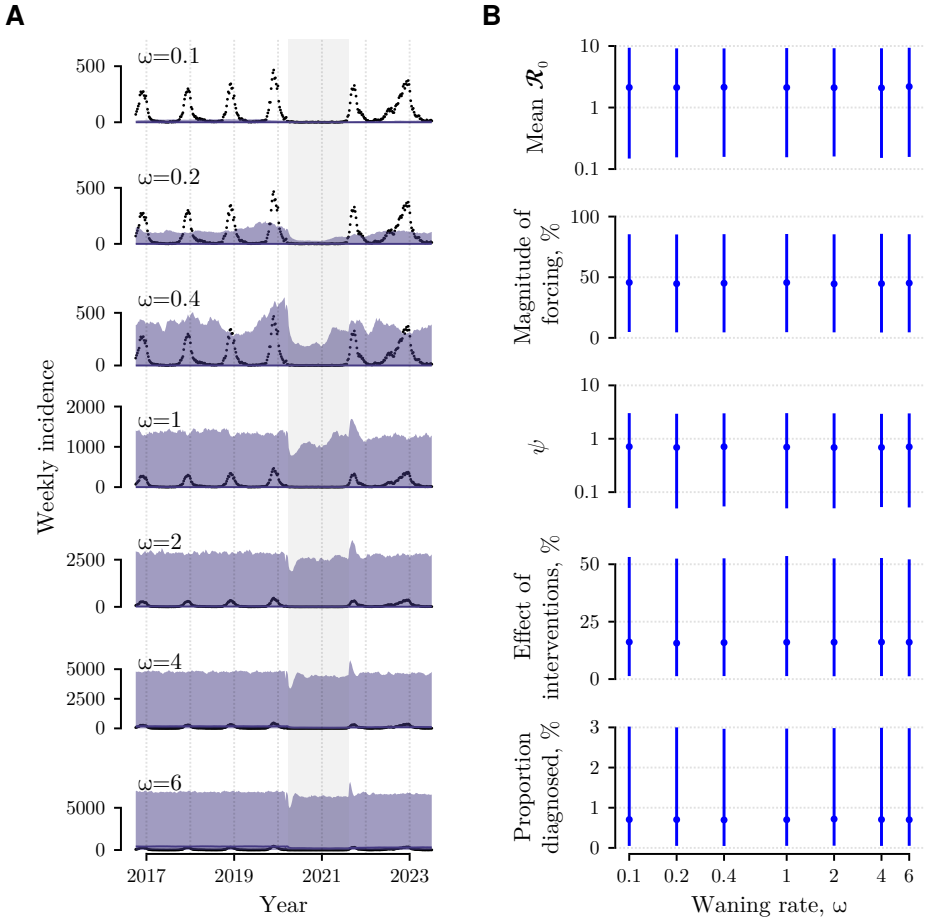

Figure B.14: Prior distribution plot for respiratory syncytial virus incidence in Scotland. *A* Purple shading represents the central 90% of predicted incidence using prior distributions for each parameter. The grey shaded area represents the period of high stringency index (>50) is Scotland. Black dots show recorded incidence. *B* Median and central 90% intervals for prior distributions. $\mathcal{R}_0$, basic reproduction ratio; ψ, boosting coefficient. [102]

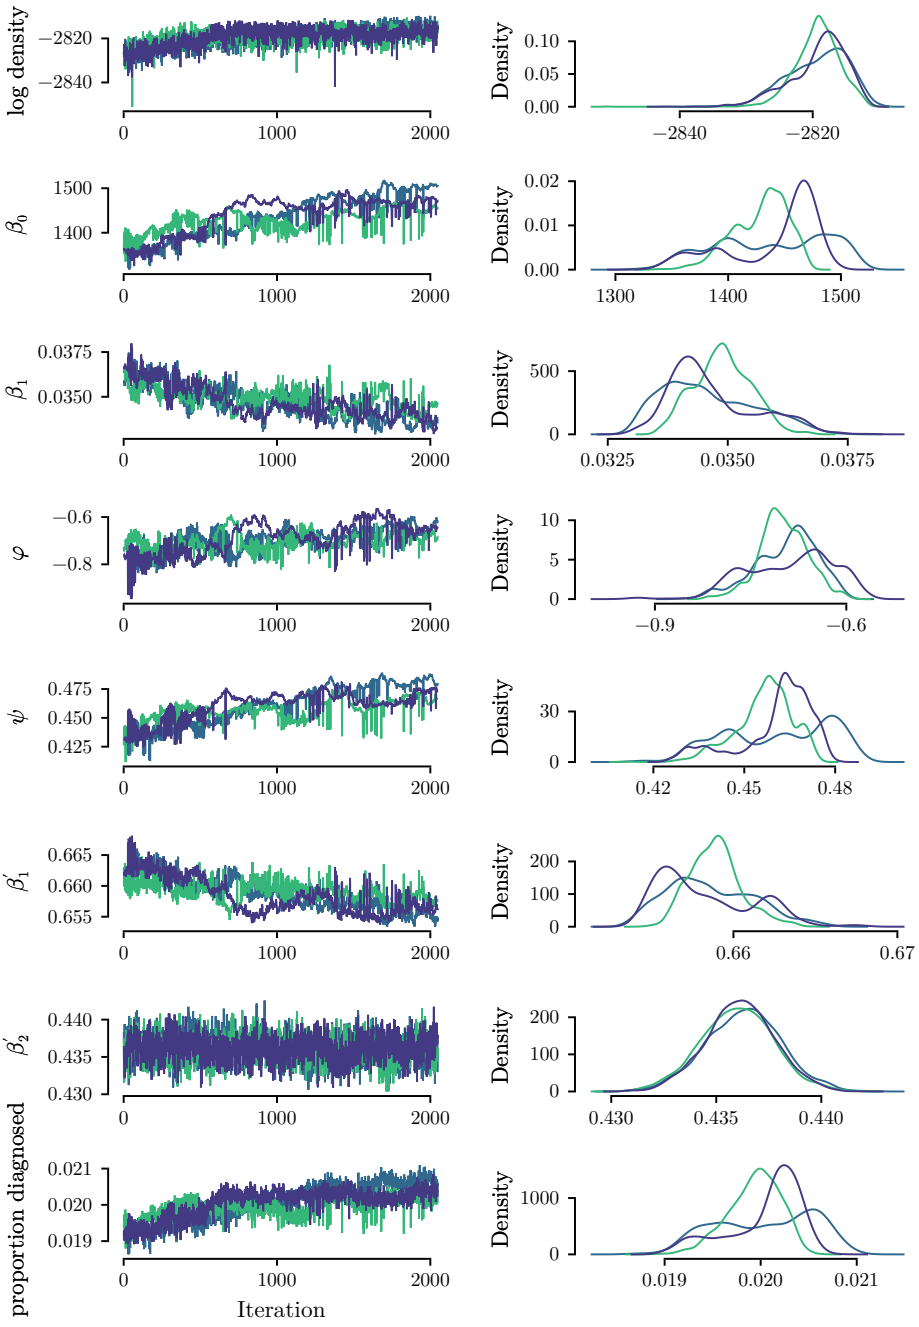

Figure B.15: Chains for analysis of effect of non-pharmaceutical interventions on respiratory syncytial virus incidence in Scotland, assuming immune waning at a rate of 0.1/year. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [102]

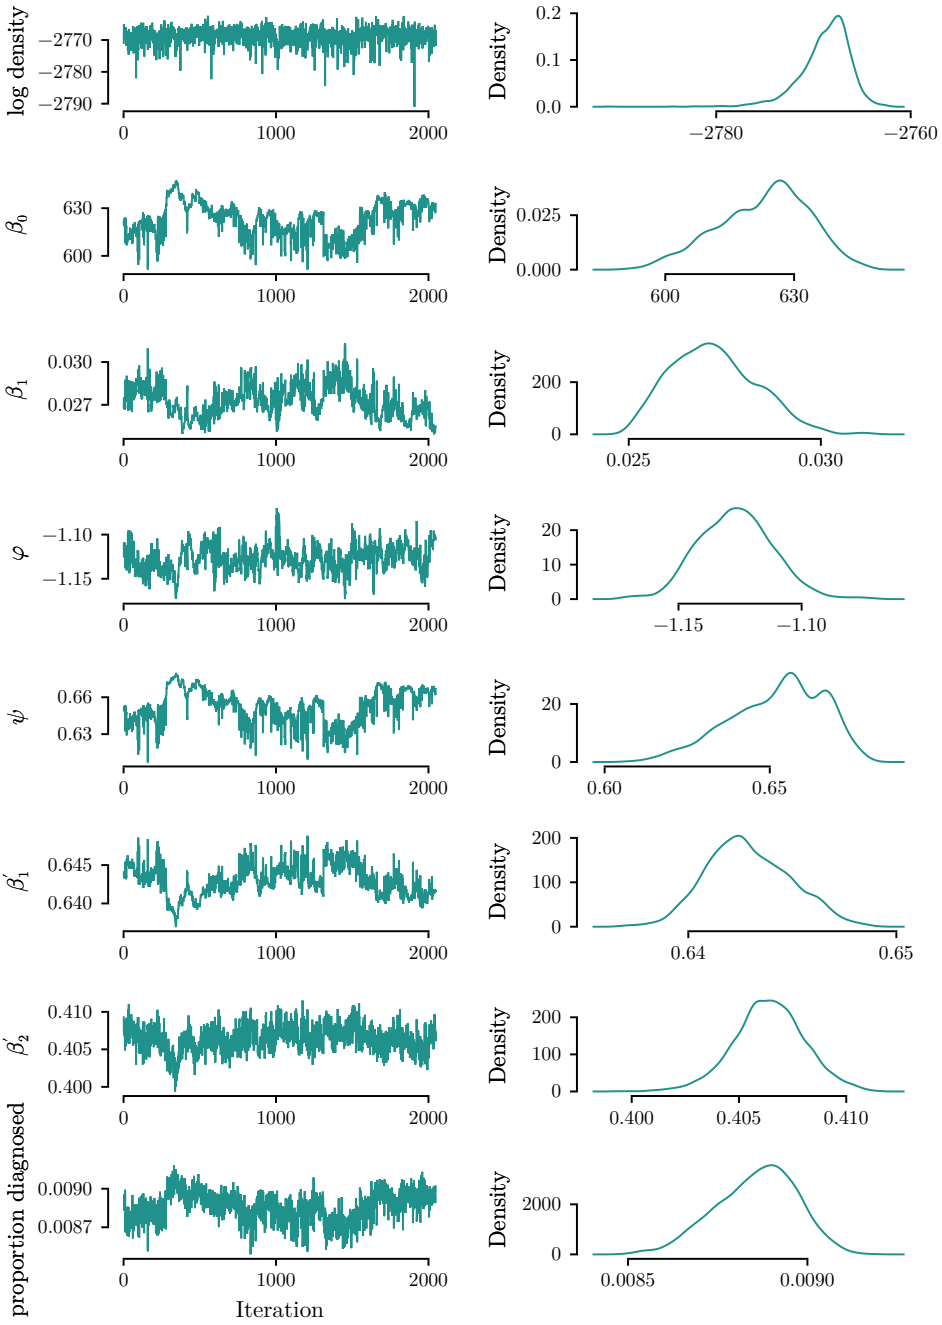

Figure B.16: Chains for analysis of effect of non-pharmaceutical interventions on respiratory syncytial virus incidence in Scotland, assuming immune waning at a rate of 0.2/year. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [102]

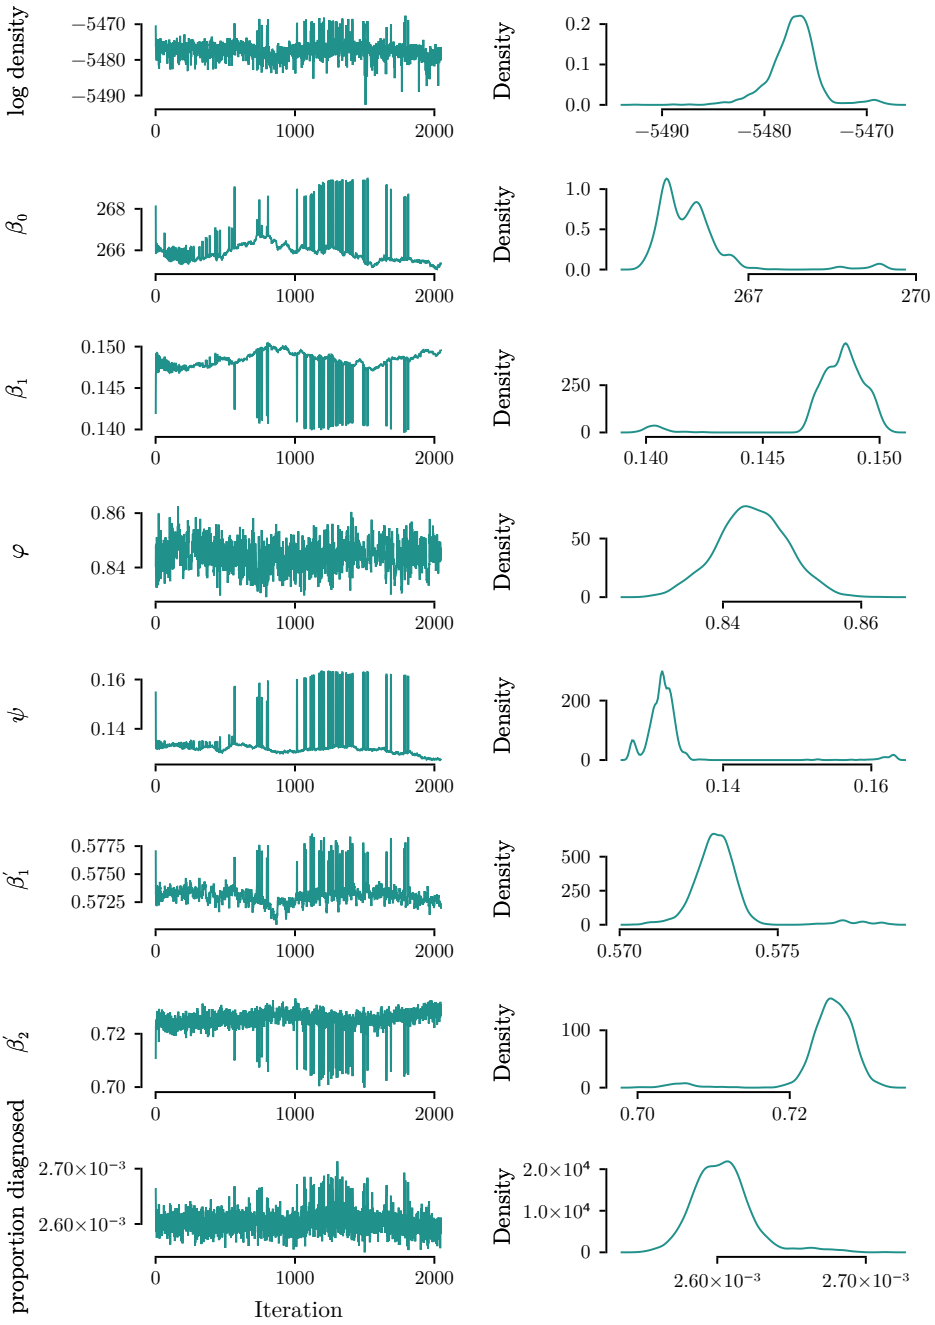

Figure B.17: Chains for analysis of effect of non-pharmaceutical interventions on respiratory syncytial virus incidence in Scotland, assuming immune waning at a rate of 0.4/year. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [102]

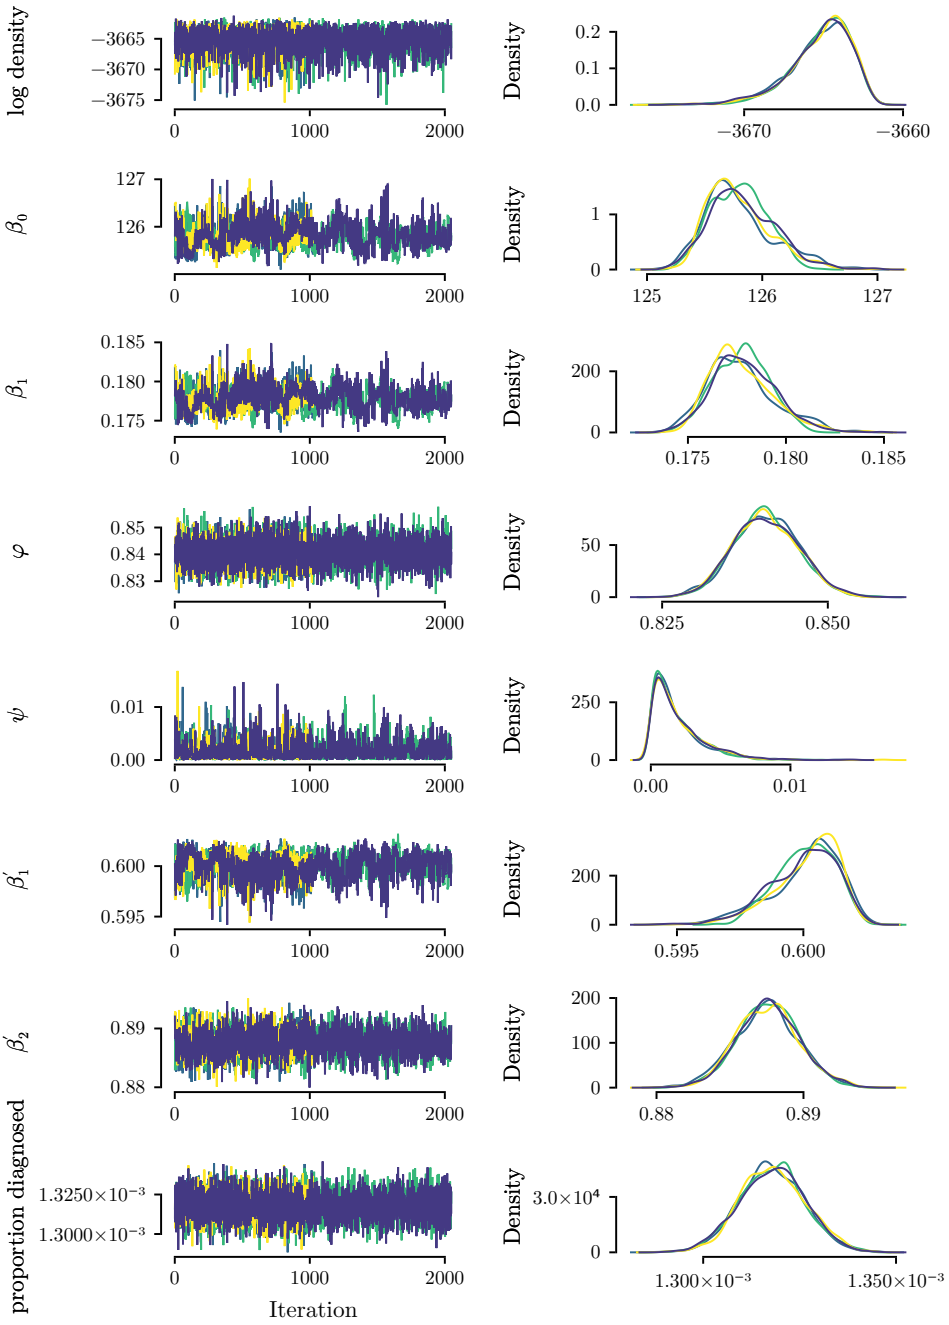

Figure B.18: Chains for analysis of effect of non-pharmaceutical interventions on respiratory syncytial virus incidence in Scotland, assuming immune waning at a rate of 1/year. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [102]

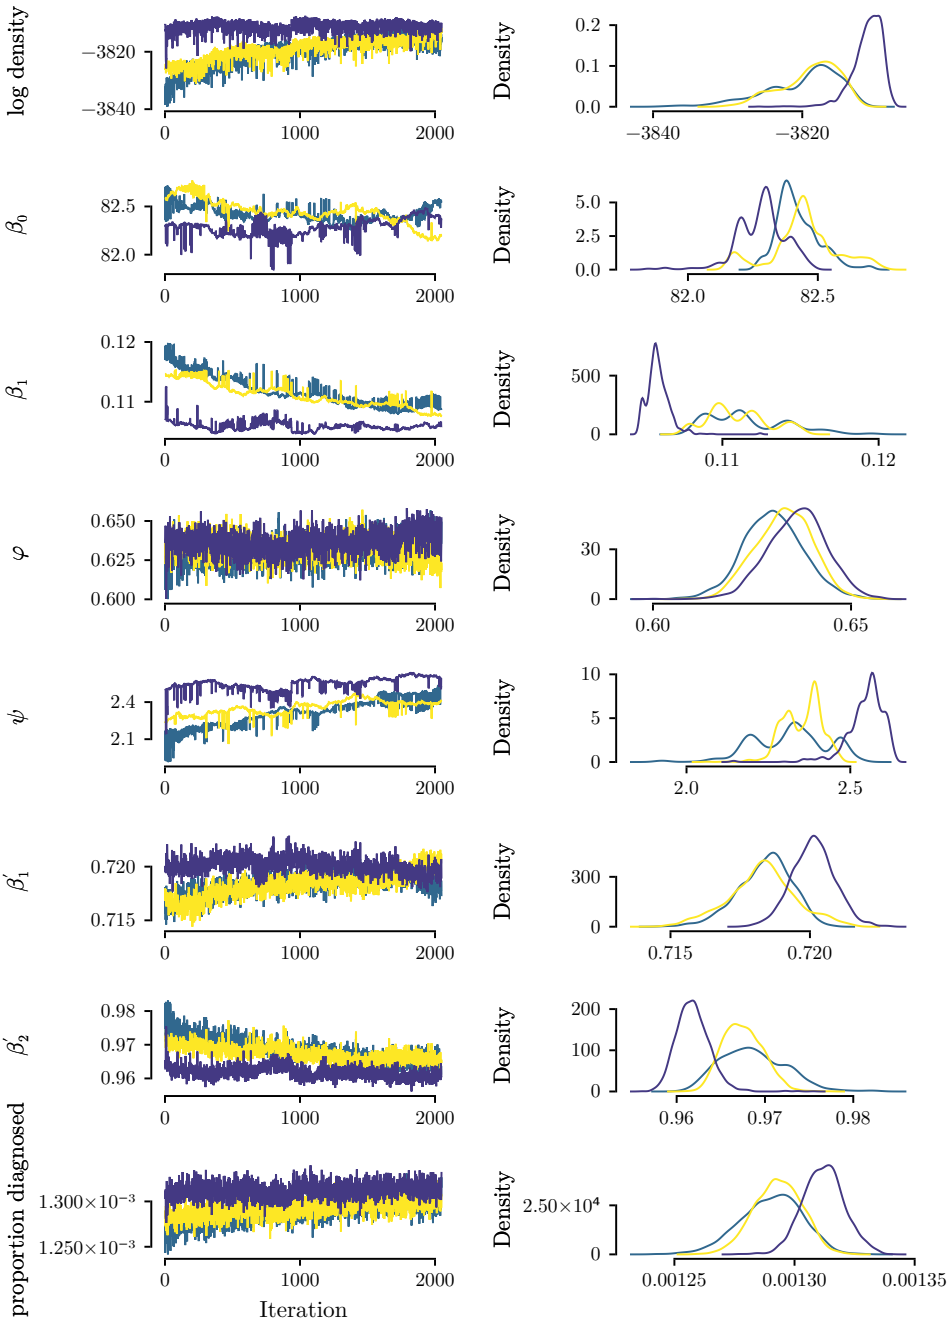

Figure B.19: Chains for analysis of effect of non-pharmaceutical interventions on respiratory syncytial virus incidence in Scotland, assuming immune waning at a rate of 2/year. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [102]

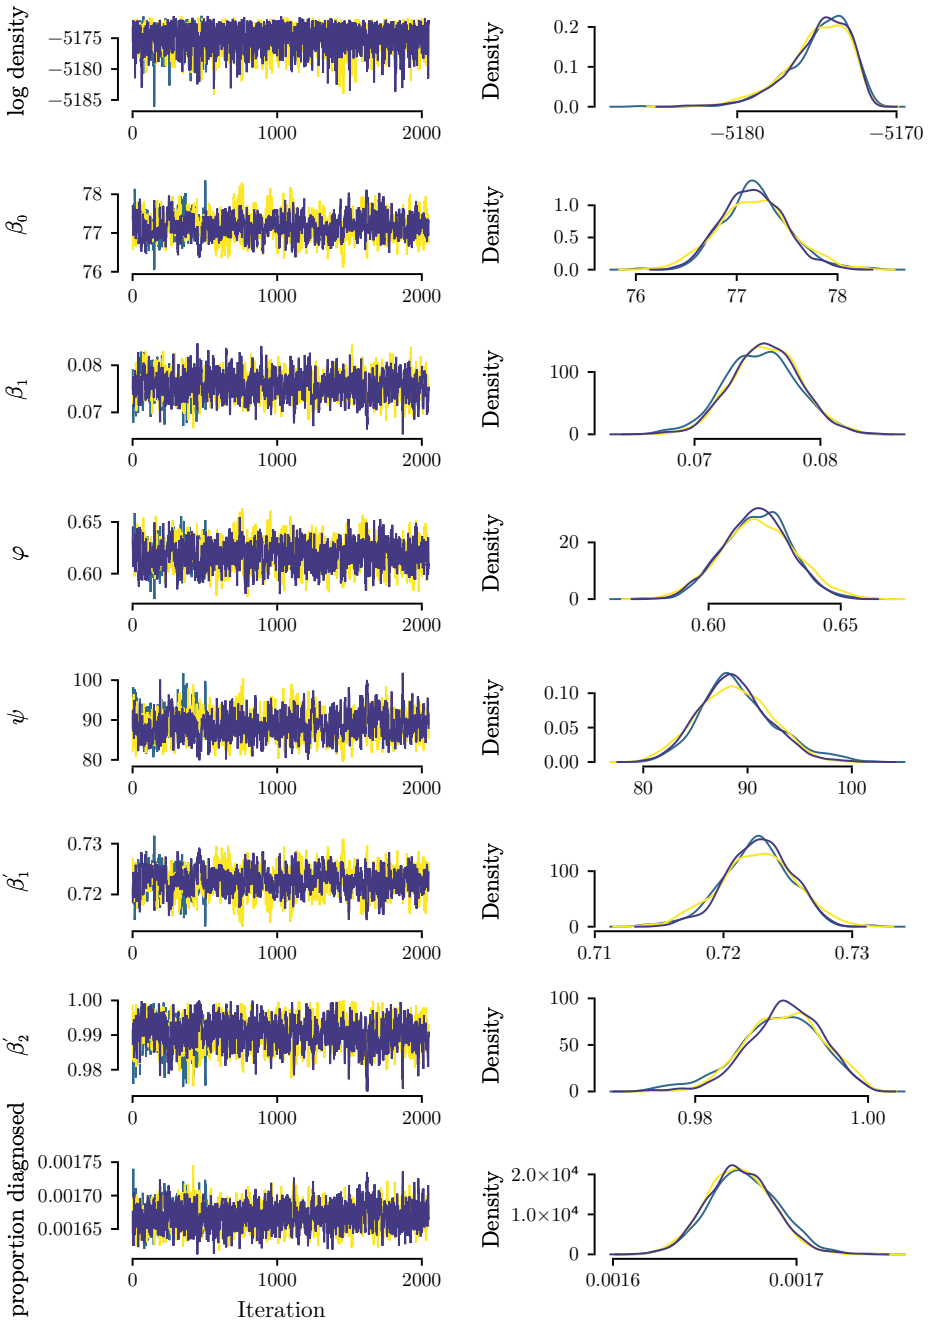

Figure B.20: Chains for analysis of effect of non-pharmaceutical interventions on respiratory syncytial virus incidence in Scotland, assuming immune waning at a rate of 4/year. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [102]

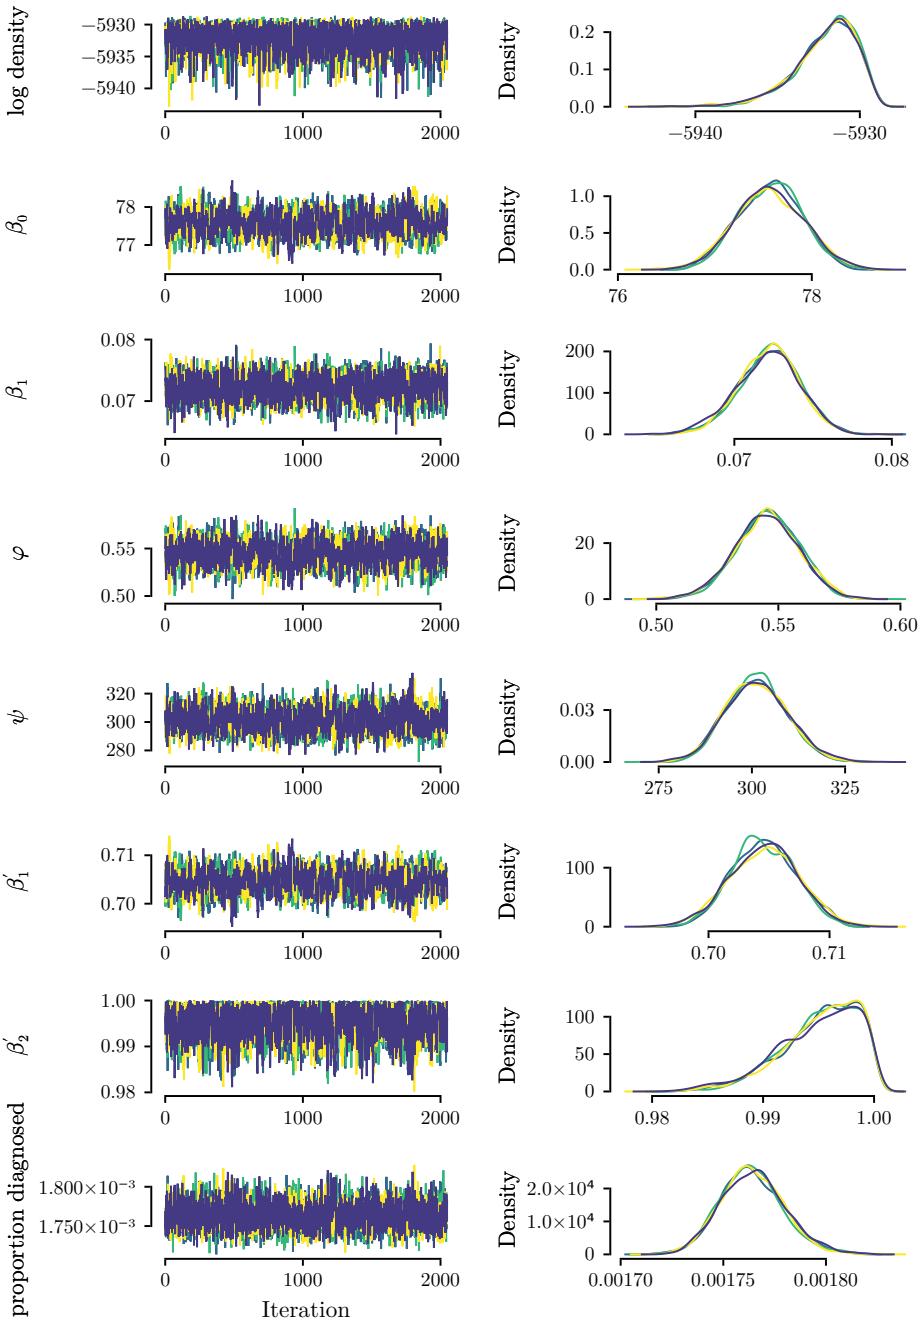

Figure B.21: Chains for analysis of effect of non-pharmaceutical interventions on respiratory syncytial virus incidence in Scotland, assuming immune waning at a rate of 6/year. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [102]

B.3 For Chapter 4

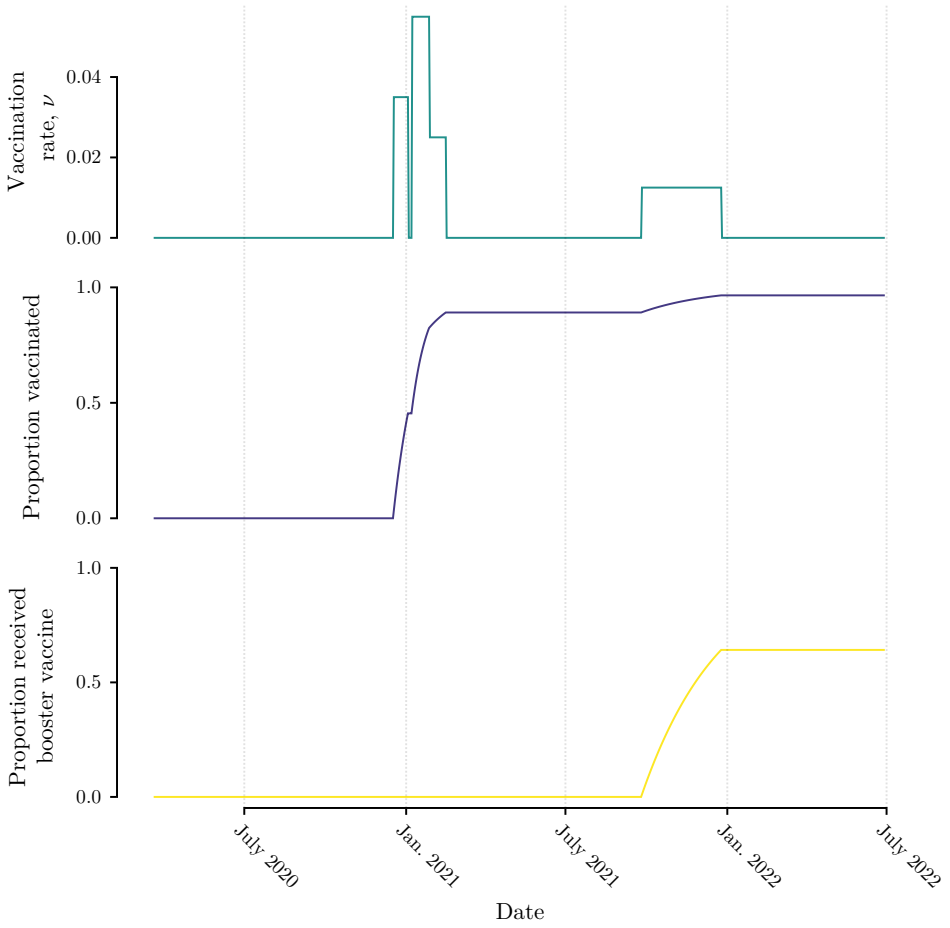

Figure B.22: Simulated proportions of healthcare workers vaccinated over time. [123]

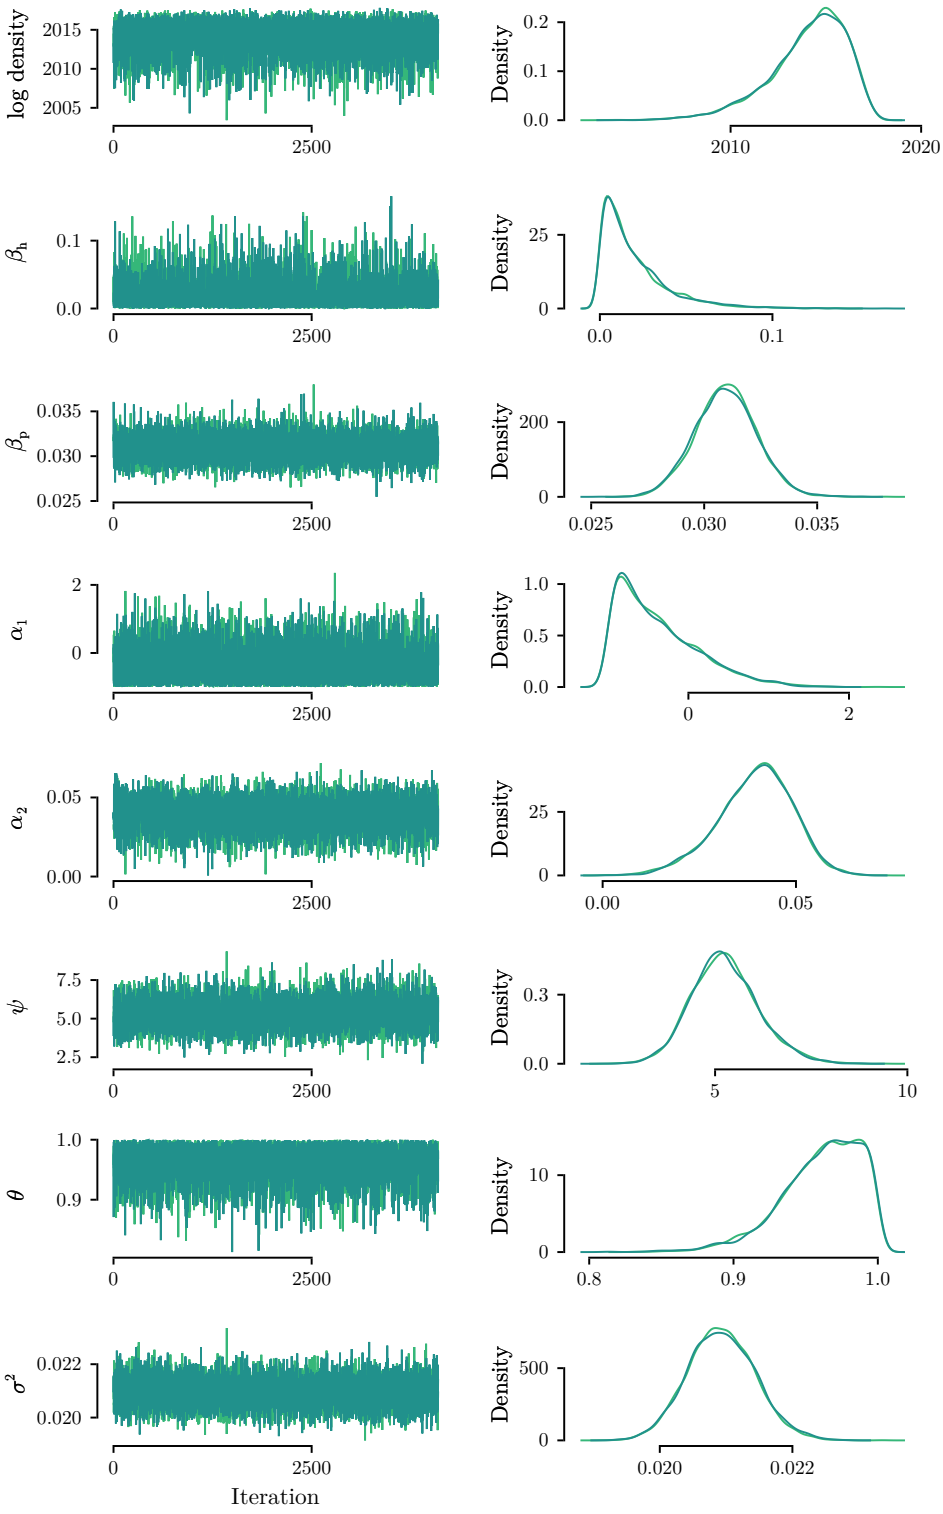

Figure B.23: Chains for analysis of immunity for healthcare workers within a single hospital. Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. [129]

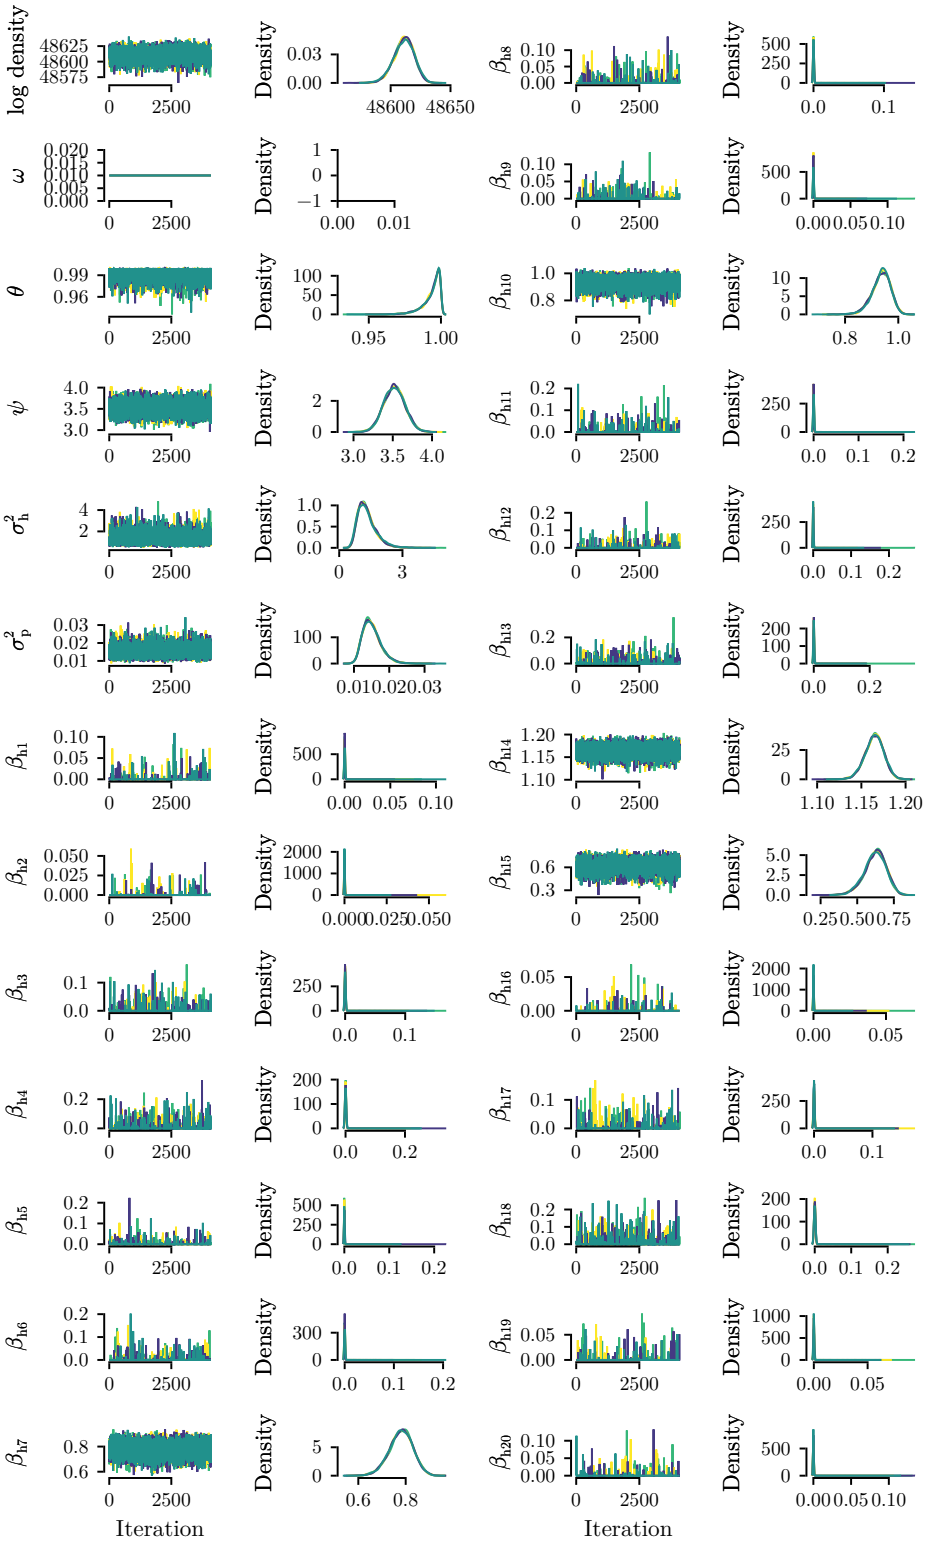

Figure B.24: Chains for analysis of immunity for healthcare workers in hospitals in the Midlands of England (part 1). Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. See Figure B.25 for other parameters. [130]

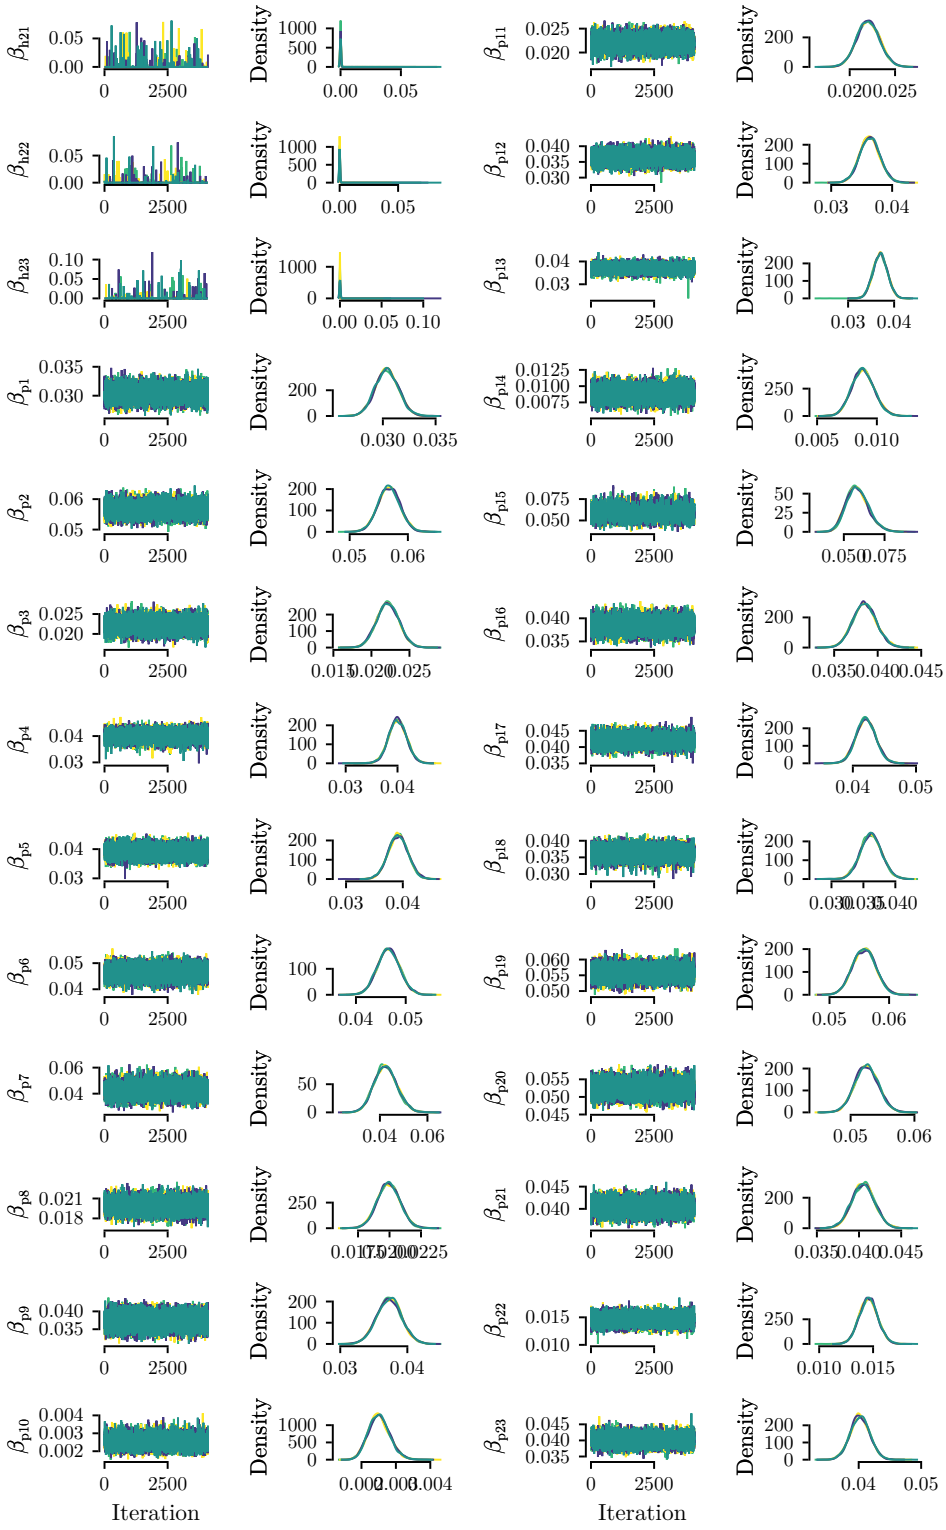

Figure B.25: Chains for analysis of immunity for healthcare workers in hospitals in the Midlands of England (part 2). Each density plot matches the chains to the left. Horizontal axes of density plots not labelled, each matches the vertical axis of the plot of chains to the left. Parameters defined in the text. See Figure B.24 for other parameters. [130]

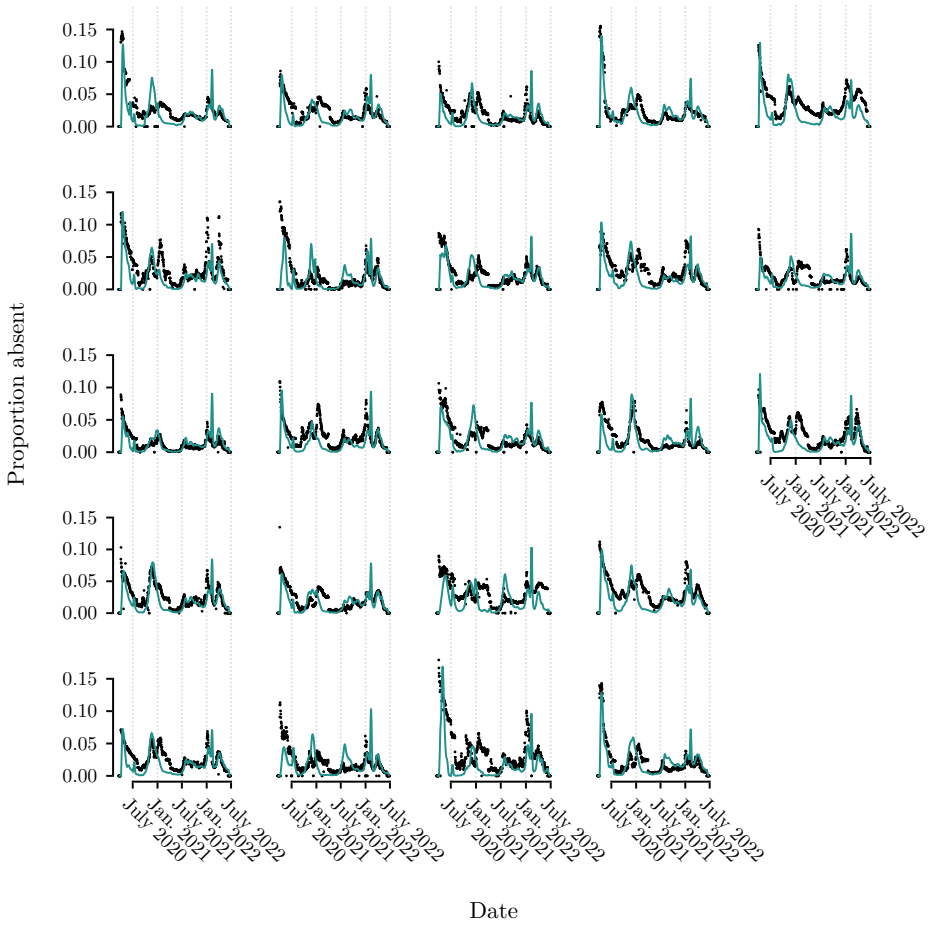

Figure B.26: Proportions of healthcare workers absent from each hospital over time. Each plot represents a hospital in the Midlands of England. Black points represent recorded data; green lines represent median estimates from model with shading representing central 90% credible interval. [130]

Typeset in the Pazo fonts by Diego Puga,
based on the Palatino typeface by Hermann Zapf.
Produced using LaTeX, developed by Leslie Lamport,
based on TeX by Donald Knuth.
Layout uses the Memoir Class by Peter Wilson,
maintained by Lars Madsen.